



---

Junos<sup>®</sup> OS

## Automation Scripting Feature Guide

Release  
16.2



---

Modified: 2016-11-10

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Junos<sup>®</sup> OS Automation Scripting Feature Guide*

16.2

Copyright © 2016, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

#### YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

#### END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	xxv
	Documentation and Release Notes . . . . .	xxv
	Supported Platforms . . . . .	xxv
	Using the Examples in This Manual . . . . .	xxvi
	Merging a Full Example . . . . .	xxvi
	Merging a Snippet . . . . .	xxvii
	Documentation Conventions . . . . .	xxvii
	Documentation Feedback . . . . .	xxix
	Requesting Technical Support . . . . .	xxx
	Self-Help Online Tools and Resources . . . . .	xxx
	Opening a Case with JTAC . . . . .	xxx
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Junos Automation Overview . . . . .</b>	<b>3</b>
	Junos Automation Overview . . . . .	3
	Junos Configuration Automation: Commit Scripts . . . . .	4
	Junos Operations Automation: Op Scripts . . . . .	4
	Junos Event Automation: Event Scripts and Event Policy . . . . .	4
	Event Policy . . . . .	5
	Event Scripts . . . . .	5
	Junos SNMP Automation: SNMP Scripts . . . . .	5
<b>Chapter 2</b>	<b>Junos XML Management Protocol and Junos XML API Overview . . . . .</b>	<b>7</b>
	XML Overview . . . . .	7
	Tag Elements . . . . .	8
	Attributes . . . . .	8
	Namespaces . . . . .	9
	Document Type Definition . . . . .	9
	XML and Junos OS Overview . . . . .	9
	Junos XML Management Protocol and Junos XML API Overview . . . . .	11
	Advantages of Using the Junos XML Management Protocol and Junos XML	
	API . . . . .	12
	Parsing Device Output . . . . .	12
	Displaying Device Output . . . . .	13
<b>Part 2</b>	<b>Automation Scripting Using XSLT</b>	
<b>Chapter 3</b>	<b>XSLT Overview . . . . .</b>	<b>17</b>
	XSLT Overview . . . . .	17
	XSLT Advantages . . . . .	17
	XSLT Engine . . . . .	18

	XSLT Concepts .....	18
	XSLT Namespace .....	19
	XPath Overview .....	19
	Nodes and Axes .....	20
	Path and Predicate Syntax .....	20
	XPath Operators .....	21
	XSLT Templates Overview .....	22
	Unnamed (Match) Templates .....	22
	Named Templates .....	23
	XSLT Parameters Overview .....	24
	Declaring Parameters .....	24
	Passing Parameters .....	25
	Example: Parameters and Match Templates .....	26
	Example: Parameters and Named Templates .....	26
	XSLT Variables Overview .....	27
	XSLT Programming Instructions Overview .....	28
	<xsl:choose> Programming Instruction .....	28
	<xsl:for-each> Programming Instruction .....	29
	<xsl:if> Programming Instruction .....	29
	Sample XSLT Programming Instructions and Pseudocode .....	30
	XSLT Recursion Overview .....	31
	XSLT Context (Dot) Overview .....	31
<b>Chapter 4</b>	<b>Standard XPath and XSLT Functions Used in Automation Scripts .....</b>	<b>33</b>
	concat() .....	33
	contains() .....	34
	count() .....	34
	last() .....	34
	name() .....	35
	not() .....	35
	position() .....	35
	starts-with() .....	36
	string-length() .....	36
	substring-after() .....	37
	substring-before() .....	37
<b>Chapter 5</b>	<b>Standard XSLT Elements and Attributes Used in Automation Scripts ....</b>	<b>39</b>
	xsl:apply-templates .....	40
	xsl:call-template .....	40
	xsl:choose .....	41
	xsl:comment .....	42
	xsl:copy-of .....	42
	xsl:element .....	43
	xsl:for-each .....	43
	xsl:if .....	44
	xsl:import .....	44
	xsl:otherwise .....	45
	xsl:param .....	46
	xsl:stylesheet .....	47
	xsl:template .....	48

	xsl:template match="/" Template .....	49
	xsl:text .....	51
	xsl:value-of .....	51
	xsl:variable .....	52
	xsl:when .....	53
	xsl:with-param .....	53
<b>Part 3</b>	<b>Automation Scripting Using SLAX</b>	
<b>Chapter 6</b>	<b>SLAX Overview .....</b>	<b>57</b>
	SLAX Overview .....	57
	SLAX Advantages .....	57
	How SLAX Works .....	58
	SLAX Syntax Rules Overview .....	59
	Code Blocks .....	59
	Comments .....	59
	Line Termination .....	60
	Strings .....	60
	SLAX Elements and Element Attributes Overview .....	61
	SLAX Elements .....	61
	SLAX Element Attributes .....	61
	SLAX Elements as Function Arguments .....	62
	The Main Template .....	63
	Understanding SLAX Default Namespaces .....	63
	SLAX Errors .....	65
	XPath Expressions Overview for SLAX .....	65
	SLAX Templates Overview .....	66
	Unnamed (Match) Templates .....	67
	Named Templates .....	68
	SLAX Functions Overview .....	70
	SLAX Parameters Overview .....	72
	Declaring Parameters .....	72
	Passing Parameters to Templates .....	73
	Example: Parameters and Match Templates .....	74
	Passing Parameters to Functions .....	75
	SLAX Variables Overview .....	76
	Immutable variables .....	76
	Mutable variables .....	77
	SLAX Statements Overview .....	79
	for-each Statement .....	79
	if, else if, and else Statements .....	80
	match Statement .....	81
	ns Statement .....	81
	version Statement .....	82
	XSLT Elements Without SLAX Equivalents .....	82
	SLAX Operators .....	83
<b>Chapter 7</b>	<b>SLAX Statements .....</b>	<b>87</b>
	append .....	88
	apply-imports .....	89

apply-templates	89
attribute	90
attribute-set	91
call	93
copy-node	94
copy-of	94
decimal-format	95
element	96
else	97
else if	98
expr	99
fallback	100
for	100
for-each	102
function	103
if	104
import	106
key	106
match	108
message	109
mode	109
mvar	111
number	111
output-method	115
param	118
preserve-space	119
priority	119
processing-instruction	120
result	121
set	122
sort	123
strip-space	125
template	125
terminate	126
trace	127
uexpr	128
use-attribute-sets	128
var	129
version	130
while	131
with	131

<b>Chapter 8</b>	<b>Using the libslax Distribution for Automation Scripting . . . . .</b>	<b>135</b>
	libslax Distribution Overview . . . . .	135
	libslax Library and Extension Libraries Overview . . . . .	136
	libslax Library . . . . .	136
	libslax Extension Libraries . . . . .	136
	Downloading and Installing the libslax Distribution . . . . .	137
	libslax Default Extension Libraries: bit, cURL, and xutil . . . . .	138
	libslax bit Extension Library . . . . .	138
	libslax cURL Extension Library . . . . .	140
	Understanding the cURL Extension Library . . . . .	140
	curl:close . . . . .	143
	curl:open . . . . .	143
	curl:perform . . . . .	144
	curl:set . . . . .	145
	curl:single . . . . .	145
	Supported Format Elements . . . . .	145
	cURL Examples . . . . .	146
	libslax xutil Extension Library . . . . .	148
	Understanding the SLAX Processor (slaxproc) . . . . .	148
	slaxproc Overview . . . . .	149
	slaxproc Modes . . . . .	149
	slaxproc Options . . . . .	150
	slaxproc File Argument Handling . . . . .	152
	slaxproc UNIX Scripting Support . . . . .	153
	Using the SLAX Processor (slaxproc) . . . . .	153
	Validating SLAX Script Syntax . . . . .	154
	Converting Scripts Between XSLT and SLAX Formats . . . . .	154
	Running SLAX Scripts . . . . .	156
	Formatting SLAX Scripts . . . . .	157
	SLAX Debugger, Profiler, and callflow . . . . .	158
	SLAX Debugger, Profiler, and callflow Overview . . . . .	158
	Using the SLAX Debugger, Profiler, and callflow . . . . .	160
	Invoking the SLAX Debugger . . . . .	160
	Using the SLAX Debugger (sdb) . . . . .	161
	Using the SLAX Profiler . . . . .	162
	Using callflow . . . . .	163
<b>Part 4</b>	<b>Automation Scripting Using Python</b>	
<b>Chapter 9</b>	<b>Python Overview . . . . .</b>	<b>167</b>
	Understanding Python Automation Scripts for Devices Running Junos OS . . . . .	167
	Requirements for Executing Python Automation Scripts on Devices Running Junos OS . . . . .	169
	Overview of Python Modules Available on Devices Running Junos OS . . . . .	171

<b>Part 5</b>	<b>Automation Script Input</b>	
<b>Chapter 10</b>	<b>Using Global Parameters in Automation Scripts . . . . .</b>	<b>175</b>
	Global Parameters and Variables in Junos OS Automation Scripts . . . . .	175
	Global Parameters Available in SLAX and XSLT Scripts . . . . .	176
	Global Variable Available in Automation Scripts . . . . .	176
<b>Part 6</b>	<b>Extension Functions and Named Templates for Automation Scripting</b>	
<b>Chapter 11</b>	<b>Using Extension Functions for Automation Scripting . . . . .</b>	<b>183</b>
	Understanding Extension Functions in Junos OS Automation Scripts . . . . .	183
	Using Extension Functions in Junos OS Automation Scripts . . . . .	186
	Using Extension Functions in SLAX and XSLT Scripts . . . . .	186
	Using Extension Functions in Python Scripts . . . . .	188
<b>Chapter 12</b>	<b>Extension Functions in the jcs and slax Namespaces . . . . .</b>	<b>189</b>
	base64-decode() Function (SLAX) . . . . .	190
	base64-encode() Function (SLAX) . . . . .	190
	break-lines() Function (SLAX and XSLT) . . . . .	191
	close() Function (SLAX and XSLT) . . . . .	192
	dampen() Function (Python, SLAX, and XSLT) . . . . .	192
	document() Function (SLAX) . . . . .	193
	emit_error() Function (Python) . . . . .	194
	emit_snmp_attributes Function (Python) . . . . .	195
	emit_warning() Function (Python) . . . . .	196
	empty() Function (SLAX and XSLT) . . . . .	197
	evaluate() Function (SLAX) . . . . .	197
	execute() Function (SLAX and XSLT) . . . . .	198
	first-of() Function (SLAX and XSLT) . . . . .	199
	get-command() Function (SLAX) . . . . .	200
	get-hello() Function (SLAX and XSLT) . . . . .	201
	get-input() Function (SLAX and XSLT) and get_input() (Python) . . . . .	202
	get-protocol() Function (SLAX and XSLT) . . . . .	203
	get-secret() Function (SLAX and XSLT) and get_secret() (Python) . . . . .	204
	get_snmp_action() Function (Python) . . . . .	205
	get_snmp_oid() Function (Python) . . . . .	206
	hostname() Function (Python, SLAX, and XSLT) . . . . .	207
	invoke() Function (SLAX and XSLT) . . . . .	208
	open() Function (SLAX and XSLT) . . . . .	209
	output() Function (Python, SLAX, and XSLT) . . . . .	211
	parse-ip() Function (SLAX and XSLT) and parse_ip() (Python) . . . . .	213
	printf() Function (Python, SLAX, and XSLT) . . . . .	214
	progress() Function (Python, SLAX, and XSLT) . . . . .	215
	regex() Function (SLAX and XSLT) . . . . .	216
	sleep() Function (SLAX and XSLT) . . . . .	217
	split() Function (SLAX and XSLT) . . . . .	218
	sysctl() Function (Python, SLAX, and XSLT) . . . . .	220
	syslog() Function (Python, SLAX, and XSLT) . . . . .	220
	trace() Function (Python, SLAX, and XSLT) . . . . .	223



<b>Chapter 13</b>	<b>Using Named Templates for Automation Scripting . . . . .</b>	<b>225</b>
	Understanding Named Templates in Junos OS Automation Scripts . . . . .	225
	Using Named Templates in Junos OS Automation Scripts . . . . .	226
	Using Named Templates in SLAX and XSLT Scripts . . . . .	226
	Using Named Template Functionality in Python Scripts . . . . .	227
<b>Chapter 14</b>	<b>Named Templates in the jcs Namespace . . . . .</b>	<b>229</b>
	jcs:edit-path Template . . . . .	229
	emit-change Template (SLAX and XSLT) and emit_change(Python) . . . . .	230
	jcs:emit-comment Template . . . . .	233
	jcs:grep Template . . . . .	234
	jcs:load-configuration Template . . . . .	235
	jcs:statement Template . . . . .	237
<b>Part 7</b>	<b>Commit Scripts</b>	
<b>Chapter 15</b>	<b>Commit Scripts Overview . . . . .</b>	<b>241</b>
	Commit Script Overview . . . . .	241
	Advantages of Using Commit Scripts . . . . .	242
	How Commit Scripts Work . . . . .	243
	Commit Script Input . . . . .	244
	Commit Script Output . . . . .	245
	Commit Scripts and the Junos OS Commit Model . . . . .	247
	Standard Commit Model . . . . .	247
	Commit Model with Commit Scripts . . . . .	247
<b>Chapter 16</b>	<b>Creating and Executing Commit Scripts . . . . .</b>	<b>251</b>
	Required Boilerplate for Commit Scripts . . . . .	251
	XML Syntax for Common Commit Script Tasks . . . . .	253
	Design Considerations for Commit Scripts . . . . .	254
	Avoiding Potential Conflicts When Using Multiple Commit Scripts . . . . .	256
	Line-by-Line Explanation of Sample Commit Scripts . . . . .	257
	Applying a Change to SONET/SDH Interfaces . . . . .	257
	Applying a Change to ISO-Enabled Interfaces . . . . .	258
	Controlling Execution of Commit Scripts During Commit Operations . . . . .	261
	Enabling Commit Scripts to Execute During Commit Operations . . . . .	262
	Preventing Commit Scripts from Executing During Commit Operations . . . . .	263
	Deactivating Commit Scripts . . . . .	263
	Activating Commit Scripts . . . . .	263
	Configuring Checksum Hashes for a Commit Script . . . . .	264
	Executing Large Commit Scripts . . . . .	265
	Example: Retrieving the Pre-Inheritance Candidate Configuration in a Commit Script . . . . .	265

<b>Chapter 17</b>	<b>Generating a Custom Warning, Error, or System Log Message Using Commit Scripts . . . . .</b>	<b>269</b>
	Overview of Generating Custom Warning, Error, and System Log Messages . . .	269
	Generating a Custom Warning, Error, or System Log Message in Commit Scripts . . . . .	270
	SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages . . . . .	274
	Example: Generating a Custom Warning Message . . . . .	276
	Example: Generating a Custom Error Message . . . . .	279
	Example: Generating a Custom System Log Message . . . . .	283
<b>Chapter 18</b>	<b>Generating Persistent or Transient Configuration Changes Using Commit Scripts . . . . .</b>	<b>287</b>
	Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts . . . . .	287
	Differences Between Persistent and Transient Changes . . . . .	287
	Interaction of Configuration Changes and Configuration Groups . . . . .	290
	Tag Elements and Templates for Generating Changes . . . . .	291
	Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts . . . . .	292
	Generating a Persistent or Transient Configuration Change in Python Commit Scripts . . . . .	297
	SLAX and XSLT Commit Script Tag Elements to Use When Generating Persistent and Transient Configuration Changes . . . . .	299
	Removing a Persistent or Transient Configuration Change Using Commit Scripts . . . . .	300
	Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts . . . . .	301
	Example: Generating a Persistent Change . . . . .	301
	Example: Generating a Transient Change . . . . .	306
<b>Chapter 19</b>	<b>Creating Custom Configuration Syntax with Commit Script Macros . . . . .</b>	<b>313</b>
	Overview of Creating Custom Configuration Syntax with Commit Script Macros . . . . .	313
	Creating Custom Configuration Syntax with Commit Script Macros . . . . .	313
	Creating a Custom Syntax . . . . .	314
	<data> Element . . . . .	315
	Expanding the Custom Syntax . . . . .	316
	Other Ways to Use Macros . . . . .	319
	Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements . . . . .	320
	Example: Creating Custom Configuration Syntax with Commit Script Macros . . .	322
<b>Chapter 20</b>	<b>Commit Script Examples . . . . .</b>	<b>331</b>
	Example: Adding a Final then accept Term to a Firewall . . . . .	331
	Example: Adding T1 Interfaces to a RIP Group . . . . .	336
	Example: Assigning a Classifier . . . . .	339
	Example: Automatically Configuring Logical Interfaces and IP Addresses . . . .	343
	Example: Configuring Administrative Groups for LSPs . . . . .	350
	Example: Configuring a Default Encapsulation Type . . . . .	354

	Example: Configuring Dual Routing Engines . . . . .	357
	Example: Configuring an Interior Gateway Protocol on an Interface . . . . .	362
	Example: Controlling IS-IS and MPLS Interfaces . . . . .	366
	Example: Controlling LDP Configuration . . . . .	370
	Example: Creating a Complex Configuration Based on a Simple Interface Configuration . . . . .	375
	Example: Imposing a Minimum MTU Setting . . . . .	382
	Example: Limiting the Number of ATM Virtual Circuits . . . . .	384
	Example: Limiting the Number of E1 Interfaces . . . . .	388
	Example: Loading a Base Configuration . . . . .	397
	Example: Prepending a Global Policy . . . . .	411
	Example: Preventing Import of the Full Routing Table . . . . .	416
	Example: Requiring Internal Clocking on T1 Interfaces . . . . .	419
	Example: Requiring and Restricting Configuration Statements . . . . .	422
<b>Chapter 21</b>	<b>Junos XML and XSLT Tag Elements Used in Commit Scripts . . . . .</b>	<b>427</b>
	<change> (SLAX and XSLT) . . . . .	427
	<syslog> (Junos XML) . . . . .	427
	<transient-change> (SLAX and XSLT) . . . . .	428
	xnm:error (Junos XML) . . . . .	428
	xnm:warning (Junos XML) . . . . .	430
<b>Chapter 22</b>	<b>Troubleshooting Commit Scripts . . . . .</b>	<b>433</b>
	Displaying Commit Script Output . . . . .	433
	Tracing Commit Script Processing . . . . .	434
	Minimum Configuration for Tracing for Commit Script Operations . . . . .	434
	Example: Minimum Configuration for Enabling Traceoptions for Commit Scripts . . . . .	435
	Configuring Tracing of Commit Scripts . . . . .	436
	Configuring the Commit Script Log Filename . . . . .	437
	Configuring the Number and Size of Commit Script Log Files . . . . .	437
	Configuring Access to Commit Script Log Files . . . . .	437
	Configuring the Commit Script Trace Operations . . . . .	438
	Troubleshooting Commit Scripts . . . . .	438
<b>Part 8</b>	<b>Op Scripts</b>	
<b>Chapter 23</b>	<b>Op Scripts Overview . . . . .</b>	<b>443</b>
	Op Script Overview . . . . .	443
	How Op Scripts Work . . . . .	444
<b>Chapter 24</b>	<b>Creating and Executing Op Scripts . . . . .</b>	<b>445</b>
	Required Boilerplate for Op Scripts . . . . .	445
	Mapping Operational Mode Commands and Output Fields to Junos XML Notation . . . . .	447
	Using RPCs and Operational Mode Commands in Op Scripts . . . . .	448
	Using RPCs in Op Scripts . . . . .	448
	Displaying the RPC Tags for a Command . . . . .	450

	Using Operational Mode Commands in Op Scripts . . . . .	451
	Declaring and Using Command-Line Arguments in Op Scripts . . . . .	452
	Declaring Op Script Command-Line Arguments . . . . .	453
	Using Command-Line Arguments in Op Scripts . . . . .	454
	Example: Declaring Arguments in XSLT Scripts . . . . .	455
	Example: Declaring and Using Arguments in Python Scripts . . . . .	456
	Configuring Help Text for Op Scripts . . . . .	457
	Examples: Configuring Help Text for Op Scripts . . . . .	457
	Defining Commands to Allow in an Op Script . . . . .	458
	Enabling an Op Script and Defining a Script Alias . . . . .	459
	Configuring Checksum Hashes for an Op Script . . . . .	460
	Executing an Op Script . . . . .	461
	Executing an Op Script by Issuing the op Command . . . . .	462
	Executing an Op Script at Login . . . . .	462
	Executing an Op Script from a Remote Site . . . . .	462
	Disabling an Op Script . . . . .	464
<b>Chapter 25</b>	<b>Op Script Examples . . . . .</b>	<b>467</b>
	Example: Changing the Configuration Using SLAX and XSLT Op Scripts . . . . .	467
	Example: Changing the Configuration Using Python Op Scripts . . . . .	472
	Example: Customizing Output of the show interfaces terse Command Using an Op Script . . . . .	477
	Example: Displaying DNS Hostname Information Using an Op Script . . . . .	487
	Example: Finding LSPs to Multiple Destinations Using an Op Script . . . . .	491
	Example: Restarting an FPC Using an Op Script . . . . .	495
	Example: Exporting Files Using an Op Script . . . . .	498
	Example: Importing Files Using an Op Script . . . . .	505
	Example: Searching Files Using an Op Script . . . . .	510
<b>Chapter 26</b>	<b>Provisioning Services Using Service Template Automation . . . . .</b>	<b>515</b>
	Service Template Automation Overview . . . . .	515
	Example: Configuring Service Template Automation . . . . .	516
<b>Chapter 27</b>	<b>Troubleshooting Op Scripts . . . . .</b>	<b>529</b>
	Tracing Op Script Processing . . . . .	529
	Minimum Configuration for Enabling Traceoptions for Op Scripts . . . . .	529
	Example: Minimum Configuration for Enabling Traceoptions for Op Scripts . . . . .	530
	Configuring Tracing of Local Op Scripts . . . . .	530
	Configuring the Op Script Log Filename . . . . .	531
	Configuring the Number and Size of Op Script Log Files . . . . .	531
	Configuring Access to Op Script Log Files . . . . .	532
	Configuring the Op Script Trace Operations . . . . .	532
<b>Part 9</b>	<b>Event Policies and Event Scripts</b>	
<b>Chapter 28</b>	<b>Event Policy Overview . . . . .</b>	<b>537</b>
	Event Notifications and Policies Overview . . . . .	537
	How Event Policies Work . . . . .	538

<b>Chapter 29</b>	<b>Triggering an Event Policy . . . . .</b>	<b>541</b>
	Using Correlated Events to Trigger an Event Policy . . . . .	541
	Representing the Correlating Event in an Event Policy . . . . .	544
	Example: Representing the Correlating Event in an Event Policy . . . . .	545
	Example: Correlating Events Based on Receipt of Other Events Within a Specified Time Interval . . . . .	546
	Example: Correlating Events Based on Event Attributes . . . . .	546
	Triggering an Event Policy Based on Event Count . . . . .	547
	Example: Triggering a Policy Based on Event Count . . . . .	547
	Using Regular Expressions to Refine the Set of Events That Trigger a Policy . . .	549
	Example: Controlling Event Policy Using a Regular Expression . . . . .	550
	Generating Internal Events to Trigger Event Policies . . . . .	551
	Example: Generating Internal Events to Trigger Event Policies . . . . .	552
	Using Nonstandard System Log Messages to Trigger Event Policies . . . . .	552
	Example: Using Nonstandard System Log Messages to Trigger an Event Policy . . . . .	553
<b>Chapter 30</b>	<b>Configuring Event Policy Actions . . . . .</b>	<b>555</b>
	Configuring an Event Policy to Execute Operational Mode Commands . . . . .	555
	Configuring an Event Policy to Change the Configuration . . . . .	558
	Example: Changing the Configuration Using an Event Policy . . . . .	559
	Example: Changing the Interface Configuration in Response to an Event . . . .	566
	Executing Event Scripts in an Event Policy . . . . .	573
	Configuring an Event Policy to Pass Arguments to an Event Script . . . . .	576
	Configuring Event Script Arguments in an Event Policy . . . . .	576
	Using Arguments in an Event Script . . . . .	578
	Configuring Event Policies to Ignore an Event . . . . .	578
	Example: Ignoring Events Based on Receipt of Other Events . . . . .	579
	Overview of Using Event Policies to Raise SNMP Traps . . . . .	580
	Example: Raising an SNMP Trap in Response to an Event . . . . .	581
	Understanding the Event System Log Priority in an Event Policy . . . . .	583
	Example: Configuring the Event System Log Priority in an Event Policy . . . . .	584
	Example: Limiting Event Script Output Based on a Specific Event Type . . . . .	588
<b>Chapter 31</b>	<b>Configuring Event Policy File Archiving . . . . .</b>	<b>591</b>
	Event Policy File Archiving Overview . . . . .	591
	Example: Defining Destinations for File Archiving by Event Policies . . . . .	592
	Example: Configuring an Event Policy to Upload Files . . . . .	595
	Configuring the Delay Before Files Are Uploaded by an Event Policy . . . . .	602
	Example: Configuring the Delay Before Files Are Uploaded by an Event Policy . .	603
	Configuring an Event Policy to Retry the File Upload Action . . . . .	605
	Example: Configuring an Event Policy to Retry the File Upload Action . . . . .	606
<b>Chapter 32</b>	<b>Configuring Event Policy Privileges . . . . .</b>	<b>609</b>
	Changing the User Privilege Level for an Event Policy Action . . . . .	609
	Example: Associating an Optional User with an Event Policy Action . . . . .	610
<b>Chapter 33</b>	<b>Event Scripts Overview . . . . .</b>	<b>613</b>
	Event Scripts Overview . . . . .	613

<b>Chapter 34</b>	<b>Creating and Executing Event Scripts . . . . .</b>	<b>615</b>
	Required Boilerplate for Event Scripts . . . . .	615
	Using Event and Remote Execution Details in Event Scripts . . . . .	617
	Using RPCs and Operational Mode Commands in Event Scripts . . . . .	619
	Executing RPCs on a Local Device . . . . .	620
	Executing RPCs on a Remote Device . . . . .	620
	Displaying the RPC Tags for a Command . . . . .	623
	Using Operational Mode Commands in Event Scripts . . . . .	624
	Enabling and Executing Event Scripts . . . . .	626
	Replacing an Event Script . . . . .	628
	Configuring Checksum Hashes for an Event Script . . . . .	628
<b>Chapter 35</b>	<b>Troubleshooting Event Policies and Event Scripts . . . . .</b>	<b>631</b>
	Tracing Event Policy Processing . . . . .	631
	Configuring the Event Policy Log Filename . . . . .	632
	Configuring the Number and Size of Event Policy Log Files . . . . .	632
	Configuring Access to the Log File . . . . .	632
	Configuring a Regular Expression for Lines to Be Logged . . . . .	633
	Configuring the Trace Operations . . . . .	633
	Tracing Event Script Processing . . . . .	634
	Minimum Configuration for Enabling Traceoptions for Event Scripts . . . . .	634
	Example: Minimum Configuration for Enabling Traceoptions for Event Scripts . . . . .	635
	Configuring Tracing of Event Scripts . . . . .	635
	Configuring the Event Script Log Filename . . . . .	636
	Configuring the Number and Size of Event Script Log Files . . . . .	636
	Configuring Access to Event Script Log Files . . . . .	637
	Configuring the Event Script Trace Operations . . . . .	637
<b>Part 10</b>	<b>SNMP Scripts</b>	
<b>Chapter 36</b>	<b>SNMP Scripts Overview . . . . .</b>	<b>641</b>
	SNMP Scripts Overview . . . . .	641
<b>Chapter 37</b>	<b>Creating and Executing SNMP Scripts . . . . .</b>	<b>643</b>
	Required Boilerplate for SNMP Scripts . . . . .	643
	Enabling an SNMP Script . . . . .	646
	Configuring Checksum Hashes for an SNMP Script . . . . .	647
<b>Chapter 38</b>	<b>SNMP Script Example . . . . .</b>	<b>649</b>
	Example: Processing Unsupported OIDs with an SNMP Script . . . . .	649
<b>Chapter 39</b>	<b>Troubleshooting SNMP Scripts . . . . .</b>	<b>655</b>
	Tracing SNMP Script Processing . . . . .	655
	Minimum Configuration for Enabling Traceoptions for SNMP Scripts . . . . .	655
	Example: Minimum Configuration for Enabling Traceoptions for SNMP Scripts . . . . .	656
	Configuring Tracing of SNMP Scripts . . . . .	656
	Configuring the SNMP Script Log Filename . . . . .	657
	Configuring the Number and Size of SNMP Script Log Files . . . . .	657

	Configuring Access to SNMP Script Log Files . . . . .	657
	Configuring the SNMP Script Trace Operations . . . . .	658
<b>Part 11</b>	<b>Managing Automation Scripts</b>	
<b>Chapter 40</b>	<b>Storing and Enabling Scripts . . . . .</b>	<b>661</b>
	Storing and Enabling Scripts . . . . .	661
	Storing Scripts in Flash Memory . . . . .	664
	Storing and Using Imported Scripts and Script Functionality . . . . .	665
<b>Chapter 41</b>	<b>Configuring a Remote Source for Scripts . . . . .</b>	<b>667</b>
	Overview of Updating Scripts from a Remote Source . . . . .	667
	Configuring and Using a Master Source Location for a Script . . . . .	669
	Configuring the Master Source for a Script . . . . .	669
	Updating a Script from the Master Source . . . . .	669
	Example: Configuring and Refreshing from the Master Source for a Script . . . .	671
	Using an Alternate Source Location for a Script . . . . .	673
	Refreshing a Script from an Alternate Location . . . . .	673
	Example: Refreshing a Script from an Alternate Source . . . . .	675
<b>Chapter 42</b>	<b>Configuring the Session Protocol for Scripts . . . . .</b>	<b>679</b>
	Understanding the Session Protocol in Automation Scripts . . . . .	679
	Example: Specifying the Session Protocol for a Connection within Scripts . . .	681
<b>Chapter 43</b>	<b>Controlling Execution of Scripts . . . . .</b>	<b>693</b>
	Understanding Limits on Executed Event Policies and Memory Allocation for Scripts . . . . .	693
	Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts . . . . .	695
	Dampening Script Execution . . . . .	697
<b>Chapter 44</b>	<b>Synchronizing Scripts Between Routing Engines . . . . .</b>	<b>701</b>
	Understanding Script Synchronization Between Routing Engines . . . . .	701
	Synchronizing Scripts Between Routing Engines . . . . .	702
	Configuring Script Synchronization Between Routing Engines for Commit Synchronize Operations . . . . .	703
	Synchronizing Scripts Between Routing Engines on a Per-Commit Basis . .	704
	Synchronizing Scripts Between Routing Engines from Operational Mode . .	704
	Synchronizing a Script Between Routing Engines After a Refresh . . . . .	705
	Example: Synchronizing Scripts Between Routing Engines . . . . .	706
<b>Chapter 45</b>	<b>Converting Scripts Between SLAX and XSLT . . . . .</b>	<b>711</b>
	Converting Scripts Between SLAX and XSLT . . . . .	711
	Converting a Script from SLAX to XSLT . . . . .	711
	Converting a Script from XSLT to SLAX . . . . .	712
<b>Part 12</b>	<b>Configuration Statements</b>	
<b>Chapter 46</b>	<b>Configuration Statements Common to All Scripts . . . . .</b>	<b>717</b>
	language (Scripts) . . . . .	718
	load-scripts-from-flash (Scripts) . . . . .	719

	max-datasize .....	720
	synchronize (Scripts) .....	722
<b>Chapter 47</b>	<b>Commit Script Configuration Statements .....</b>	<b>723</b>
	allow-transients .....	724
	apply-macro .....	725
	checksum .....	726
	commit .....	727
	cpu-factor (Dampen Options) .....	728
	dampen .....	729
	dampen-options .....	730
	direct-access .....	730
	file (Commit Scripts) .....	731
	line-interval (Dampen Options) .....	732
	max-datasize .....	733
	optional .....	734
	refresh (Commit Scripts) .....	735
	refresh-from (Commit Scripts) .....	736
	scripts .....	737
	source (Commit Scripts) .....	739
	time-interval (Dampen Options) .....	740
	traceoptions (Commit, Op, and SNMP Scripts) .....	741
<b>Chapter 48</b>	<b>Op Script Configuration Statements .....</b>	<b>743</b>
	allow-commands (Op Scripts) .....	744
	allow-url-for-python (Op Scripts) .....	745
	arguments (Op Scripts) .....	746
	checksum .....	747
	command .....	748
	cpu-factor (Dampen Options) .....	749
	dampen .....	750
	dampen-options .....	751
	description .....	751
	file (Op Scripts) .....	752
	line-interval (Dampen Options) .....	753
	max-datasize .....	754
	op .....	756
	no-allow-url (Op Scripts) .....	757
	refresh (Op Scripts) .....	758
	refresh-from (Op Scripts) .....	759
	scripts .....	760
	source (Op Scripts) .....	762
	time-interval (Dampen Options) .....	763
	traceoptions (Commit, Op, and SNMP Scripts) .....	764
<b>Chapter 49</b>	<b>Event Policy Configuration Statements .....</b>	<b>767</b>
	archive-sites (Event Policy) .....	769
	arguments (Event Policy) .....	770
	attributes-match .....	771
	change-configuration (Event Policy) .....	772



commands (Event Policy Change Configuration) . . . . .	773
commands (Event Policy Execute Commands) . . . . .	774
commit-options . . . . .	775
destination (Event Policy) . . . . .	776
destinations (Event Policy) . . . . .	777
equals (Event Policy) . . . . .	778
event-options . . . . .	779
event-script (Event Policy) . . . . .	781
events (Event Policy Trigger Events) . . . . .	782
events (Event Policy Correlated Events) . . . . .	782
execute-commands . . . . .	783
facility . . . . .	784
generate-event . . . . .	784
ignore . . . . .	785
matches (Event Policy) . . . . .	785
max-policies . . . . .	786
no-drift (Generate Events) . . . . .	786
not (Event Policy) . . . . .	787
output-filename (Event Policy) . . . . .	788
output-format (Event Policy) . . . . .	789
policy (Event Policy) . . . . .	790
priority-override (Event Policy) . . . . .	792
raise-trap (Event Policy) . . . . .	792
retry (Event Policy) . . . . .	793
retry-count (Event Policy) . . . . .	794
severity (Event Policy) . . . . .	795
starts-with (Event Policy) . . . . .	796
then (Event Policy) . . . . .	797
time-interval (Generate Events) . . . . .	798
time-of-day . . . . .	799
traceoptions (Event Options) . . . . .	800
transfer-delay (Event Policy) . . . . .	802
trigger . . . . .	803
upload . . . . .	804
user-name (Event Policy) . . . . .	805
within . . . . .	806
<b>Chapter 50 Event Script Configuration Statements . . . . .</b>	<b>807</b>
checksum . . . . .	808
cpu-factor (Dampen Options) . . . . .	809
dampen . . . . .	810
dampen-options . . . . .	811
event-script (Event Options) . . . . .	812
file (Event Scripts) . . . . .	813
line-interval (Dampen Options) . . . . .	814
max-datasize . . . . .	815
optional (Event Scripts) . . . . .	816
python-script-user (Event and SNMP Scripts) . . . . .	817
refresh (Event Scripts) . . . . .	818

	refresh-from (Event Scripts) .....	819
	remote-execution .....	820
	source (Event Scripts) .....	821
	time-interval (Dampen Options) .....	822
	traceoptions (Event Scripts) .....	823
<b>Chapter 51</b>	<b>SNMP Script Configuration Statements .....</b>	<b>825</b>
	checksum .....	826
	file (SNMP Scripts) .....	827
	max-datasize .....	828
	oid .....	829
	python-script-user (Event and SNMP Scripts) .....	830
	refresh (SNMP Scripts) .....	831
	refresh-from (SNMP Scripts) .....	832
	snmp .....	833
	source (SNMP Scripts) .....	834
	traceoptions (Commit, Op, and SNMP Scripts) .....	835
<b>Chapter 52</b>	<b>Translation Script Configuration Statements .....</b>	<b>837</b>
	max-datasize .....	838
	translation .....	839
<b>Part 13</b>	<b>Operational Commands</b>	
<b>Chapter 53</b>	<b>Operational Commands Common to All Scripts .....</b>	<b>843</b>
	file checksum md5 .....	844
	file checksum sha-256 .....	845
	file checksum sha1 .....	846
	request system scripts convert .....	847
	request system scripts refresh-from .....	849
	request system scripts synchronize .....	851
<b>Chapter 54</b>	<b>Op Script Operational Commands .....</b>	<b>853</b>
	op .....	854
	op invoke-debugger cli .....	856
<b>Chapter 55</b>	<b>Event Script Operational Commands .....</b>	<b>859</b>
	request system scripts event-scripts .....	860

# List of Figures

<b>Part 2</b>	<b>Automation Scripting Using XSLT</b>	
<b>Chapter 3</b>	<b>XSLT Overview</b> .....	<b>17</b>
	Figure 1: Flow of XSLT Commit Script Through the XSLT Engine .....	18
<b>Chapter 5</b>	<b>Standard XSLT Elements and Attributes Used in Automation Scripts</b> ....	<b>39</b>
	Figure 2: Commit Script Input and Output .....	49
<b>Part 3</b>	<b>Automation Scripting Using SLAX</b>	
<b>Chapter 6</b>	<b>SLAX Overview</b> .....	<b>57</b>
	Figure 3: SLAX Script Input and Output .....	58
<b>Part 7</b>	<b>Commit Scripts</b>	
<b>Chapter 15</b>	<b>Commit Scripts Overview</b> .....	<b>241</b>
	Figure 4: Standard Commit Model .....	247
	Figure 5: Commit Model with Commit Scripts Added .....	248
<b>Chapter 16</b>	<b>Creating and Executing Commit Scripts</b> .....	<b>251</b>
	Figure 6: Configuration Evaluation by Multiple Commit Scripts .....	256
<b>Chapter 19</b>	<b>Creating Custom Configuration Syntax with Commit Script Macros</b> .....	<b>313</b>
	Figure 7: Macro Input and Output .....	314
<b>Part 9</b>	<b>Event Policies and Event Scripts</b>	
<b>Chapter 28</b>	<b>Event Policy Overview</b> .....	<b>537</b>
	Figure 8: Interaction of eventd Process with Other Junos OS Processes .....	537



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xxv</b>
	Table 1: Notice Icons . . . . .	xxviii
	Table 2: Text and Syntax Conventions . . . . .	xxviii
<b>Part 2</b>	<b>Automation Scripting Using XSLT</b>	
<b>Chapter 3</b>	<b>XSLT Overview</b> . . . . .	<b>17</b>
	Table 3: XSLT Concepts . . . . .	18
	Table 4: Examples and Pseudocode for XSLT Variable Declaration . . . . .	28
	Table 5: Examples and Pseudocode for XSLT Programming Instructions . . . . .	30
<b>Part 3</b>	<b>Automation Scripting Using SLAX</b>	
<b>Chapter 6</b>	<b>SLAX Overview</b> . . . . .	<b>57</b>
	Table 6: Prefix Set with libslax . . . . .	64
	Table 7: Error Elements . . . . .	65
	Table 8: SLAX Operators . . . . .	83
<b>Chapter 7</b>	<b>SLAX Statements</b> . . . . .	<b>87</b>
	Table 9: Numbering Styles for SLAX Statement number, format Option . . . . .	113
<b>Chapter 8</b>	<b>Using the libslax Distribution for Automation Scripting</b> . . . . .	<b>135</b>
	Table 10: Functions in the bit Extension Library . . . . .	138
	Table 11: Functions in the cURL Extension Library . . . . .	141
	Table 12: Web Services Elements in the cURL Extension Library . . . . .	141
	Table 13: E-mail Elements in the cURL Extension Library . . . . .	142
	Table 14: curl:perform Reply Elements . . . . .	144
	Table 15: curl:perform <header> Elements . . . . .	144
	Table 16: Functions in the xutil Extension Library . . . . .	148
	Table 17: Slaxproc Modes . . . . .	149
	Table 18: slaxproc Common Options and File Options . . . . .	150
	Table 19: SLAX Debugger Commands . . . . .	158
	Table 20: Profile Command Options . . . . .	162
<b>Part 4</b>	<b>Automation Scripting Using Python</b>	
<b>Chapter 9</b>	<b>Python Overview</b> . . . . .	<b>167</b>
	Table 21: Python Automation Script Input . . . . .	168
	Table 22: Python Automation Script Requirements . . . . .	170
	Table 23: Python Modules Available on Devices Running Junos OS . . . . .	171

<b>Part 5</b>	<b>Automation Script Input</b>	
<b>Chapter 10</b>	<b>Using Global Parameters in Automation Scripts . . . . .</b>	<b>175</b>
	Table 24: Predefined Parameters Available in SLAX and XSLT Scripts . . . . .	176
	Table 25: Global Variable \$junos-context Available to SLAX and XSLT Scripts . . . . .	178
<b>Part 6</b>	<b>Extension Functions and Named Templates for Automation Scripting</b>	
<b>Chapter 11</b>	<b>Using Extension Functions for Automation Scripting . . . . .</b>	<b>183</b>
	Table 26: Extension Functions for Use in Automation Scripts . . . . .	184
<b>Chapter 12</b>	<b>Extension Functions in the jcs and slax Namespaces . . . . .</b>	<b>189</b>
	Table 27: Options for slax:document Function . . . . .	194
	Table 28: Facility Strings . . . . .	221
	Table 29: Severity Strings . . . . .	222
<b>Chapter 13</b>	<b>Using Named Templates for Automation Scripting . . . . .</b>	<b>225</b>
	Table 30: Junos OS Named Templates . . . . .	225
<b>Part 7</b>	<b>Commit Scripts</b>	
<b>Chapter 15</b>	<b>Commit Scripts Overview . . . . .</b>	<b>241</b>
	Table 31: Commit Scripts Actions and Output . . . . .	245
<b>Chapter 16</b>	<b>Creating and Executing Commit Scripts . . . . .</b>	<b>251</b>
	Table 32: XML Syntax for Common Commit Script Tasks . . . . .	253
<b>Chapter 17</b>	<b>Generating a Custom Warning, Error, or System Log Message Using Commit Scripts . . . . .</b>	<b>269</b>
	Table 33: Tags and Attributes for Creating Custom Warning, Error, and System Log Messages . . . . .	274
<b>Chapter 18</b>	<b>Generating Persistent or Transient Configuration Changes Using Commit Scripts . . . . .</b>	<b>287</b>
	Table 34: Differences Between Persistent and Transient Changes . . . . .	289
	Table 35: Tags and Attributes for Creating Configuration Changes in SLAX and XSLT Commit Scripts . . . . .	300
<b>Chapter 19</b>	<b>Creating Custom Configuration Syntax with Commit Script Macros . . . . .</b>	<b>313</b>
	Table 36: Sample Macro and Junos OS CLI Expansion . . . . .	322
<b>Chapter 22</b>	<b>Troubleshooting Commit Scripts . . . . .</b>	<b>433</b>
	Table 37: Commit Script Configuration and Operational Mode Commands . . . . .	433
	Table 38: Commit Script Tracing Operational Mode Commands . . . . .	435
	Table 39: Commit Script Tracing Flags . . . . .	438
	Table 40: Troubleshooting Commit Scripts . . . . .	438
<b>Part 8</b>	<b>Op Scripts</b>	
<b>Chapter 27</b>	<b>Troubleshooting Op Scripts . . . . .</b>	<b>529</b>
	Table 41: Op Script Tracing Operational Mode Commands . . . . .	530

	Table 42: Op Script Tracing Flags . . . . .	532
<b>Part 9</b>	<b>Event Policies and Event Scripts</b>	
<b>Chapter 29</b>	<b>Triggering an Event Policy . . . . .</b>	<b>541</b>
	Table 43: Event Count Triggers Policy . . . . .	548
	Table 44: Regular Expression Operators for the matches Statement . . . . .	549
	Table 45: Event ID by System Log Message Origin . . . . .	553
<b>Chapter 35</b>	<b>Troubleshooting Event Policies and Event Scripts . . . . .</b>	<b>631</b>
	Table 46: Event Policy Tracing Flags . . . . .	633
	Table 47: Event Script Tracing Operational Mode Commands . . . . .	635
	Table 48: Event Script Tracing Flags . . . . .	637
<b>Part 10</b>	<b>SNMP Scripts</b>	
<b>Chapter 36</b>	<b>SNMP Scripts Overview . . . . .</b>	<b>641</b>
	Table 49: SNMP Script Tasks . . . . .	641
<b>Chapter 39</b>	<b>Troubleshooting SNMP Scripts . . . . .</b>	<b>655</b>
	Table 50: SNMP Script Tracing Operational Mode Commands . . . . .	656
	Table 51: SNMP Script Tracing Flags . . . . .	658
<b>Part 12</b>	<b>Configuration Statements</b>	
<b>Chapter 49</b>	<b>Event Policy Configuration Statements . . . . .</b>	<b>767</b>
	Table 52: System Log Message Severity Levels . . . . .	795





# About the Documentation

- [Documentation and Release Notes on page xxv](#)
- [Supported Platforms on page xxv](#)
- [Using the Examples in This Manual on page xxvi](#)
- [Documentation Conventions on page xxvii](#)
- [Documentation Feedback on page xxix](#)
- [Requesting Technical Support on page xxx](#)

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

## Supported Platforms

---

For the features described in this document, the following platforms are supported:

- [ACX Series](#)
- [EX Series](#)
- [M Series](#)
- [MX Series](#)
- [QFX Series](#)
- [SRX Series](#)
- [T Series](#)
- [PTX Series](#)

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

### Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

## Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {  
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]  
user@host# edit system scripts  
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]  
user@host# load merge relative /var/tmp/ex-script-snippet.conf  
load complete
```

For more information about the **load** command, see [CLI Explorer](#).

---

## Documentation Conventions

[Table 1 on page xxviii](#) defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xxviii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop <i>address</i>;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.

- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

---

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.





## PART 1

# Overview

- [Junos Automation Overview on page 3](#)
- [Junos XML Management Protocol and Junos XML API Overview on page 7](#)



## CHAPTER 1

# Junos Automation Overview

- [Junos Automation Overview on page 3](#)

## Junos Automation Overview

---

Junos automation consists of a suite of tools used to automate operational and configuration tasks on network devices running the Junos<sup>®</sup> operating system (Junos OS). The Junos automation toolkit is part of the standard Junos OS available on all switches, routers, and security devices running Junos OS. Junos automation tools, which leverage the native XML capabilities of Junos OS, include commit scripts, operation (op) scripts, event policies and event scripts, SNMP scripts, and macros.

Junos automation simplifies complex configurations and reduces potential configuration errors. It saves time by automating operational and configuration tasks. It also speeds troubleshooting and maximizes network uptime by warning of potential problems and automatically responding to system events.

Junos automation can capture the knowledge and expertise of experienced network operators and administrators and allow a business to leverage this combined expertise across the organization.

Junos automation scripts can be written in several different scripting languages: Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python. XSLT is a standard for processing Extensible Markup Language (XML) data and is designed to convert one XML document into another. SLAX is an alternative to XSLT. It has a simple syntax that follows the style of C and PERL, but retains the same semantics as XSLT. Programmers who are familiar with C often find it easier to learn and use SLAX. You can easily convert SLAX scripts into XSLT and convert XSLT scripts into SLAX. Python is a commonly-used, open-source programming language with extensive standard and community libraries. On devices running Junos OS, Python automation scripts can leverage [Junos PyEZ](#) APIs to simplify many operational and configuration tasks.

The following sections describe the different types of functionality for Junos automation:

- [Junos Configuration Automation: Commit Scripts on page 4](#)
- [Junos Operations Automation: Op Scripts on page 4](#)

- [Junos Event Automation: Event Scripts and Event Policy on page 4](#)
- [Junos SNMP Automation: SNMP Scripts on page 5](#)

## Junos Configuration Automation: Commit Scripts

Junos configuration automation uses commit scripts to automate the commit process. Junos OS commit scripts enforce custom configuration rules. When a candidate configuration is committed, it is inspected by each active commit script. If a configuration violates your custom rules, the script can instruct Junos OS to take appropriate action. A commit script can perform the following actions:

- Generate and display custom warning messages to the user
- Generate and log custom system log (syslog) messages
- Change the configuration to conform to the custom configuration rules
- Generate a commit error and halt the commit operation

Commit scripts, when used in conjunction with macros, allow you to simplify the Junos configuration and, at the same time, extend it with your own custom configuration syntax.

## Junos Operations Automation: Op Scripts

Junos operations automation uses op scripts to automate operational tasks and network troubleshooting. Junos OS op scripts can be executed manually in the CLI or upon user login, or they can be called from another script. Op scripts can process user arguments and can be constructed to:

- Create custom operational mode commands
- Execute a series of operational mode commands
- Customize the output of operational mode commands
- Shorten troubleshooting time by gathering operational information and iteratively narrowing down the cause of a network problem
- Perform controlled configuration changes
- Monitor the overall status of a device by creating a general operation script that periodically checks network warning parameters, such as high CPU usage.

## Junos Event Automation: Event Scripts and Event Policy

Junos event automation uses event policy and event scripts to instruct Junos OS to perform actions in response to system events.

### Event Policy

---

An event policy is an if-then-else construct that defines actions to be executed by the software on receipt of an event such as a system log message or SNMP trap. Event policies can be executed in response to a single system event or to correlated system events. For each policy, you can configure multiple actions including:

- Ignore the event
- Upload a file to a specified destination
- Execute Junos OS operational mode commands
- Execute Junos OS event scripts
- Modify the configuration

### Event Scripts

---

Junos OS event scripts are triggered automatically by defined event policies in response to a system event and can instruct Junos OS to take immediate action. An event script automates network troubleshooting and network management by:

- Automatically diagnosing and fixing problems in the network
- Monitoring the overall status of a device
- Running automatically as part of an event policy that detects periodic error conditions
- Changing the configuration in response to a problem

## Junos SNMP Automation: SNMP Scripts

Junos OS SNMP scripts provide the flexibility to support custom MIBs until they are implemented in Junos OS. SNMP scripts are triggered automatically when the SNMP manager requests information from the SNMP agent for an object identifier (OID) that is mapped to an SNMP script for an unsupported OID. The script acts like an SNMP subagent, and the system sends the return value from the script to the network management system (NMS).

You can map an SNMP script to one or more OIDs using the `oid` statement at the `[edit system scripts snmp file script-name]` hierarchy level.

#### Related Documentation

- [Commit Script Overview on page 241](#)
- [Event Notifications and Policies Overview on page 537](#)
- [Event Scripts Overview on page 613](#)
- [Op Script Overview on page 443](#)
- [SNMP Scripts Overview on page 641](#)



## CHAPTER 2

# Junos XML Management Protocol and Junos XML API Overview

- XML Overview on page 7
- XML and Junos OS Overview on page 9
- Junos XML Management Protocol and Junos XML API Overview on page 11
- Advantages of Using the Junos XML Management Protocol and Junos XML API on page 12

## XML Overview

---

Extensible Markup Language (XML) is a language for defining a set of markers, called *tags*, that are applied to a data set or document to describe the function of individual elements and codify the hierarchical relationships between them. Tags look much like Hypertext Markup Language (HTML) tags, but XML is actually a metalanguage used to define tags that best suit the kind of data being marked.

For more details about XML, see *A Technical Introduction to XML* at <http://www.xml.com/pub/a/98/10/guide0.html> and the additional reference material at the <http://www.xml.com> site.

The official XML specification from the World Wide Web Consortium (W3C), *Extensible Markup Language (XML) 1.0*, is available at <http://www.w3.org/TR/REC-xml>.

The following sections discuss general aspects of XML:

- Tag Elements on page 8
- Attributes on page 8
- Namespaces on page 9
- Document Type Definition on page 9

## Tag Elements

XML has three types of tags: opening tags, closing tags, and empty tags. XML tag names are enclosed in angle brackets and are case sensitive. Items in an XML-compliant document or data set are always enclosed in paired opening and closing tags, and the tags must be properly nested. That is, you must close the tags in the same order in which you opened them. XML is stricter in this respect than HTML, which sometimes uses only opening tags. The following examples show paired opening and closing tags enclosing a value. The closing tags are indicated by the forward slash at the start of the tag name.

```
<interface-state>enabled</interface-state>
<input-bytes>25378</input-bytes>
```

The term *tag element* refers to a three-part set: opening tag, contents, and closing tag. The content can be an alphanumeric character string as in the preceding examples, or can itself be a *container* tag element, which contains other tag elements. For simplicity, the term *tag* is often used interchangeably with *tag element* or *element*.

If a tag element is *empty*—has no contents—it can be represented either as paired opening and closing tags with nothing between them, or as a single tag with a forward slash after the tag name. For example, the notation **<snmp-trap-flag/>** is equivalent to **<snmp-trap-flag></snmp-trap-flag>**.

As the preceding examples show, angle brackets enclose the name of the tag element. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

Junos XML and Junos XML protocol tag elements obey the XML convention that the tag element name indicates the kind of information enclosed by the tags. For example, the name of the Junos XML **<interface-state>** tag element indicates that it contains a description of the current status of an interface on the device, whereas the name of the **<input-bytes>** tag element indicates that its contents specify the number of bytes received.

When discussing tag elements in text, this documentation conventionally uses just the opening tag to represent the complete tag element (opening tag, contents, and closing tag). For example, the documentation refers to the **<input-bytes>** tag to indicate the entire **<input-bytes>number-of-bytes</input-bytes>** tag element.

## Attributes

XML elements can contain associated properties in the form of *attributes*, which specify additional information about an element. Attributes appear in the opening tag of an element and consist of an attribute name and value pair. The attribute syntax consists of the attribute name followed by an equals sign and then the attribute value enclosed in quotation marks. An XML element can have multiple attributes. Multiple attributes are separated by spaces and can appear in any order.

In the following example, the **configuration** element has two attributes, **junos:changed-seconds** and **junos:changed-localtime**.



```
<configuration junos:changed-seconds="1279908006"
junos:changed-localtime="2010-07-23 11:00:06 PDT">
```

The value of the **junos:changed-seconds** attribute is "1279908006", and the value of the **junos:changed-localtime** attribute is "2010-07-23 11:00:06 PDT".

## Namespaces

*Namespaces* allow an XML document to contain the same tag, attribute, or function names for different purposes and avoid name conflicts. For example, many namespaces may define a **print** function, and each may exhibit a different functionality. To use the functionality defined in one specific namespace, you must associate that function with the namespace that defines the desired functionality.

To refer to a tag, attribute, or function from a defined namespace, you must first provide the namespace Uniform Resource Identifier (URI) in your style sheet declaration. You then qualify a tag, attribute, or function from the namespace with the URI. Since a URI is often lengthy, generally a shorter prefix is mapped to the URI.

In the following example the **jcs** prefix is mapped to the namespace identified by the URI <http://xml.juniper.net/junos/commit-scripts/1.0>, which defines extension functions used in commit, op, event, and SNMP scripts. The **jcs** prefix is then prepended to the **output** function, which is defined in that namespace.

```
<?xml version="1.0"?>
<xsl:stylesheet version="1.0"
xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
...
<xsl:value-of select="jcs:output('The VPN is up.')"/>
</xsl:stylesheet>
```

During processing, the prefix is expanded into the URI reference. Although there may be multiple namespaces that define an **output** element or function, the use of **jcs:output** explicitly defines which **output** function is used. You can choose any prefix to refer to the contents in a namespace, but there must be an existing declaration in the XML document that binds the prefix to the associated URI.

## Document Type Definition

An XML-tagged document or data set is *structured*, because a set of rules specifies the ordering and interrelationships of the items in it. The rules define the contexts in which each tagged item can—and in some cases must—occur. A file called a *document type definition*, or *DTD*, lists every tag element that can appear in the document or data set, defines the parent-child relationships between the tags, and specifies other tag characteristics. The same DTD can apply to many XML documents or data sets.

- Related Documentation**
- [Junos XML Management Protocol and Junos XML API Overview on page 11](#)
  - [XML and Junos OS Overview on page 9](#)

## XML and Junos OS Overview

Extensible Markup Language (XML) is a standard for representing and communicating information. It is a metalanguage for defining customized tags that are applied to a data

set or document to describe the function of individual elements and codify the hierarchical relationships between them. Junos OS natively supports XML for the operation and configuration of devices running Junos OS.

The Junos OS command-line interface (CLI) and the Junos OS infrastructure communicate using XML. When you issue an operational mode command in the CLI, the CLI converts the command into XML format for processing. After processing, Junos OS returns the output in the form of an XML document, which the CLI converts back into a readable format for display. Remote client applications also use XML-based data encoding for operational and configuration requests on devices running Junos OS.

The Junos XML API is an XML representation of Junos OS configuration statements and operational mode commands. It defines an XML equivalent for all statements in the Junos OS configuration hierarchy and many of the commands that you issue in CLI operational mode. Each operational mode command with a Junos XML counterpart maps to a request tag element and, if necessary, a response tag element.

To display operational mode command output as Junos XML tag elements instead of as the default formatted ASCII, issue the command, and pipe the output to the **display xml** command. Infrastructure tag elements in the response belong to the Junos XML management protocol. The tag elements that describe Junos OS configuration or operational data belong to the Junos XML API, which defines the Junos OS content that can be retrieved and manipulated by both the Junos XML management protocol and the NETCONF XML management protocol operations. The following example compares the text and XML output for the **show chassis alarms** operational mode command:

```
user@host> show chassis alarms
No alarms currently active

user@host> show chassis alarms | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.4R1/junos">
  <alarm-information xmlns="http://xml.juniper.net/junos/10.4R1/junos-alarm">
    <alarm-summary>
      <no-active-alarms/>
    </alarm-summary>
  </alarm-information>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```

To display the Junos XML API representation of any operational mode command, issue the command, and pipe the output to the **display xml rpc** command. The following example shows the Junos XML API request tag for the **show chassis alarms** command.

```
user@host> show chassis alarms | display xml rpc
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.4R1/junos">
  <rpc>
    <get-alarm-information>
    </get-alarm-information>
  </rpc>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```

As shown in the previous example, the **| display xml rpc** option displays the Junos XML API request tag that is sent to Junos OS for processing whenever the command is issued. In contrast, the **| display xml** option displays the actual output of the processed command in XML format.

When you issue the **show chassis alarms** operational mode command, the CLI converts the command into the Junos XML API **<get-alarm-information>** request tag and sends the XML request to the Junos OS infrastructure for processing. Junos OS processes the request and returns the **<alarm-information>** response tag element to the CLI. The CLI then converts the XML output into the “No alarms currently active” message that is displayed to the user.

Junos OS automation scripts use XML to communicate with the host device. Junos OS provides XML-formatted input to a script. The script processes the input source tree and then returns XML-formatted output to Junos OS. The script type determines the XML input document that is sent to the script as well as the output document that is returned to Junos OS for processing. Commit script input consists of an XML representation of the post-inheritance candidate configuration file. Event scripts receive an XML document containing the description of the triggering event. All script input documents contain information pertaining to the Junos OS environment, and some scripts receive additional script-specific input that depends on the script type.

- Related Documentation**
- *Junos XML API Configuration Developer Reference*
  - *Junos XML API Operational Developer Reference*

---

## Junos XML Management Protocol and Junos XML API Overview

The Junos XML Management Protocol is an XML-based protocol that client applications use to request information from and manage the configuration on routing, switching, and security devices running Junos OS. It uses an XML-based data encoding for the configuration data and remote procedure calls. The Junos XML protocol defines basic operations that are equivalent to configuration mode commands in the Junos OS command-line interface (CLI). Applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as **show**, **set**, and **commit** to perform those operations.

The Junos XML API is an XML representation of Junos configuration statements and operational mode commands. Junos XML configuration tag elements are the content to which the Junos XML protocol operations apply. Junos XML operational tag elements are equivalent in function to operational mode commands in the CLI, which administrators use to retrieve status information for a device.

Client applications request information and change the configuration on a switch, router, or security device by encoding the request with tag elements from the Junos XML management protocol and Junos XML API and sending it to the Junos XML protocol server on the device. The Junos XML protocol server is integrated into Junos OS and does not appear as a separate entry in process listings. The Junos XML protocol server directs the request to the appropriate software modules within the device, encodes the response

in Junos XML protocol and Junos XML API tag elements, and returns the result to the client application. For example, to request information about the status of a device's interfaces, a client application sends the Junos XML API `<get-interface-information>` request. The Junos XML protocol server gathers the information from the interface process and returns it in the Junos XML API `<interface-information>` response tag element.

You can use the Junos XML management protocol and Junos XML API to configure devices running Junos OS or request information about the device configuration or operation. You can write client applications to interact with the Junos XML protocol server, and you can also use the Junos XML protocol to build custom end-user interfaces for configuration and information retrieval and display, such as a Web browser-based interface.

- Related Documentation**
- [Advantages of Using the Junos XML Management Protocol and Junos XML API on page 12](#)
  - [XML and Junos OS Overview on page 9](#)
  - [XML Overview on page 7](#)

---

## Advantages of Using the Junos XML Management Protocol and Junos XML API

The Junos XML management protocol and Junos XML API fully document all options for every supported Junos OS operational request, all statements in the Junos OS configuration hierarchy, and basic operations that are equivalent to configuration mode commands. The tag names clearly indicate the function of an element in an operational or configuration request or a configuration statement.

The combination of meaningful tag names and the structural rules in a DTD makes it easy to understand the content and structure of an XML-tagged data set or document. Junos XML and Junos XML protocol tag elements make it straightforward for client applications that request information from a device to parse the output and find specific information.

### Parsing Device Output

The following example illustrates how the Junos XML API makes it easier to parse device output and extract the needed information. The example compares formatted ASCII and XML-tagged versions of output from a device running Junos OS.

The formatted ASCII follows:

```
Physical interface: fxp0, Enabled, Physical link is Up
Interface index: 4, SNMP ifIndex: 3
```

The corresponding XML-tagged version is:

```
<interface>
  <name>fxp0</name>
  <admin-status>enabled</admin-status>
  <operational-status>up</operational-status>
  <index>4</index>
  <snmp-index>3</snmp-index>
</interface>
```

When a client application needs to extract a specific value from formatted ASCII output, it must rely on the value's location, expressed either absolutely or with respect to labels or values in adjacent fields. Suppose that the client application wants to extract the interface index. It can use a regular-expression matching utility to locate specific strings, but one difficulty is that the number of digits in the interface index is not necessarily predictable. The client application cannot simply read a certain number of characters after the **Interface index:** label, but must instead extract everything between the label and the subsequent label **SNMP ifIndex:** and also account for the included comma.

A problem arises if the format or ordering of text output changes in a later version of the Junos OS. For example, if a **Logical index:** field is added following the interface index number, the new formatted ASCII might appear as follows:

```
Physical interface: fxp0, Enabled, Physical link is Up
Interface index: 4, Logical index: 12, SNMP ifIndex: 3
```

An application that extracts the interface index number delimited by the **Interface index:** and **SNMP ifIndex:** labels now obtains an incorrect result. The application must be updated manually to search for the **Logical index:** label as the new delimiter.

In contrast, the structured nature of XML-tagged output enables a client application to retrieve the interface index by extracting everything within the opening `<index>` tag and closing `</index>` tag. The application does not have to rely on an element's position in the output string, so the Junos XML protocol server can emit the child tag elements in any order within the `<interface>` tag element. Adding a new `<logical-index>` tag element in a future release does not affect an application's ability to locate the `<index>` tag element and extract its contents.

## Displaying Device Output

XML-tagged output is also easier to transform into different display formats than formatted ASCII output. For instance, you might want to display different amounts of detail about a given device component at different times. When a device returns formatted ASCII output, you have to write special routines and data structures in your display program to extract and show the appropriate information for a given detail level. In contrast, the inherent structure of XML output is an ideal basis for a display program's own structures. It is also easy to use the same extraction routine for several levels of detail, simply ignoring the tag elements you do not need when creating a less detailed display.

### Related Documentation

- [Junos XML Management Protocol and Junos XML API Overview on page 11](#)
- [XML Overview on page 7](#)



## PART 2

# Automation Scripting Using XSLT

- [XSLT Overview on page 17](#)
- [Standard XPath and XSLT Functions Used in Automation Scripts on page 33](#)
- [Standard XSLT Elements and Attributes Used in Automation Scripts on page 39](#)





## CHAPTER 3

# XSLT Overview

- [XSLT Overview on page 17](#)
- [XSLT Namespace on page 19](#)
- [XPath Overview on page 19](#)
- [XSLT Templates Overview on page 22](#)
- [XSLT Parameters Overview on page 24](#)
- [XSLT Variables Overview on page 27](#)
- [XSLT Programming Instructions Overview on page 28](#)
- [XSLT Recursion Overview on page 31](#)
- [XSLT Context \(Dot\) Overview on page 31](#)

## XSLT Overview

---

Commit scripts, op scripts, event scripts, and SNMP scripts can be written in Extensible Stylesheet Language Transformations (XSLT), which is a standard for processing Extensible Markup Language (XML) data. XSLT is developed by the World Wide Web Consortium (W3C) and is accessible at <http://www.w3c.org/TR/xslt>.

- [XSLT Advantages on page 17](#)
- [XSLT Engine on page 18](#)
- [XSLT Concepts on page 18](#)

## XSLT Advantages

XSLT is a natural match for Junos OS, with its native XML capabilities. XSLT performs XML-to-XML transformations, turning one XML hierarchy into another. It offers a great degree of freedom and power in the way in which it transforms the input XML, allowing everything from making minor changes to the existing hierarchy (such as additions or deletions) to building a completely new document hierarchy.

Because XSLT was created to allow generic XML-to-XML transformations, it is a natural choice for both inspecting configuration syntax (which Junos OS can easily express in XML) and for generating errors and warnings (which Junos OS communicates internally as XML). XSLT includes powerful mechanisms for finding configuration statements that match specific criteria. XSLT can then generate the appropriate XML result tree from

these configuration statements to instruct the Junos OS user-interface (UI) components to perform the desired behavior.


Although XSLT provides a powerful scripting ability, its focus is specific and limited. It does not make Junos OS vulnerable to arbitrary or malicious programmers. XSLT restricts programmers from performing haphazard operations, such as opening random Transmission Control Protocol (TCP) ports, forking numerous processes, or sending e-mail. The only action available in XSLT is to generate XML, and the XML is interpreted by the UI according to fixed semantics. An XSLT script can output only XML data, which is directly processed by the UI infrastructure to allow only the specific abilities listed above—generating error, warning, and system log messages, and persistent and transient configuration changes. This means that the impact of commit scripts, op scripts, event scripts, and SNMP scripts on the device is well-defined and can be viewed inside the command-line interface (CLI), using commands added for that purpose.

## XSLT Engine

XSLT is a language for transforming one XML document into another XML document. The basic model is that an XSLT engine (or processor) reads a script (or style sheet) and an XML document. The XSLT engine uses the instructions in the script to process the XML document by traversing the document's hierarchy. The script indicates what portion of the tree should be traversed, how it should be inspected, and what XML should be generated at each point. For commit scripts, op scripts, event scripts, and SNMP scripts, the XSLT engine is a function of the Junos OS management process (mgd).

[Figure 1 on page 18](#) shows the relationship between an XSLT commit script and the XSLT engine.

**Figure 1: Flow of XSLT Commit Script Through the XSLT Engine**



## XSLT Concepts

XSLT has seven basic concepts. These are summarized in [Table 3 on page 18](#).

**Table 3: XSLT Concepts**

XSLT Concepts	Description
XPath	Expression syntax for specifying a node in the input document
Templates	Mechanism for mapping input hierarchies to instructions that handle them
Parameters	Mechanism for passing arguments to templates

Table 3: XSLT Concepts (*continued*)

XSLT Concepts	Description
Variables	Mechanism for defining read-only references to nodes
Programming instructions	Mechanism for defining logic in XSLT
Recursion	Mechanism by which templates call themselves to facilitate looping
Context (Dot)	Node currently being inspected in the input document

#### Related Documentation

- [XPath Overview on page 19](#)
- [XSLT Context \(Dot\) Overview on page 31](#)
- [XSLT Parameters Overview on page 24](#)
- [XSLT Programming Instructions Overview on page 28](#)
- [XSLT Recursion Overview on page 31](#)
- [XSLT Templates Overview on page 22](#)
- [XSLT Variables Overview on page 27](#)

## XSLT Namespace

The XSLT namespace has the Uniform Resource Identifier (URI) <http://www.w3.org/1999/XSL/Transform>. The namespace must be included in the style sheet declaration of a script in order for the XSLT processor to recognize and use XSLT elements and attributes. The following example declares the XSLT namespace and associates the **xsl** prefix with the URI.

```
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
  <xsl:template match="route">
    ...
  </xsl:template>
</xsl:stylesheet>
```

Once the XSLT namespace is declared in a script, you use elements and attributes from the namespace by adding the associated prefix, which in this case is **xsl**, to the tag or attribute name. In the preceding example, the XSLT processor knows to treat **xsl:template** as an XSLT instruction. During processing, the **xsl** prefix is expanded into the URI reference, and the functionality of the **template** element is defined by the XSLT namespace. For more information about namespaces, see [“XML Overview” on page 7](#).

## XPath Overview

XSLT uses the XML Path Language (XPath) standard to specify and locate elements in the input document's XML hierarchy. XPath's powerful expression syntax enables you to define complex criteria for selecting portions of the XML input document.

## Nodes and Axes

XPath views every piece of the document hierarchy as a *node*. For commit scripts, op scripts, event scripts, and SNMP scripts, the important types of nodes are *element nodes*, *text nodes*, and *attribute nodes*. Consider the following XML tags:

```
<system>
  <host-name>my-router</host-name>
  <accounting inactive="inactive">
</system>
```

These XML tag elements show examples of the following types of XPath nodes:

- **<host-name>my-router</host-name>**—Element node
- **my-router**—Text node
- **inactive="inactive"**—Attribute node

Nodes are viewed as being arranged in certain *axes*. The *ancestor axis* points from a node up through its series of parent nodes. The *child axis* points through the list of an element node's direct child nodes. The *attribute axis* points through the list of an element node's set of attributes. The *following-sibling axis* points through the nodes that follow a node but are under the same parent. The *descendant axis* contains all the descendents of a node. There are numerous other axes that are not listed here.

Each XPath expression is evaluated from a particular node, which is referred to as the *context node* (or simply *context*). The context node is the node at which the XSLT processor is currently looking. XSLT changes the context as the document's hierarchy is traversed, and XPath expressions are evaluated from that particular context node.



**NOTE:** In Junos OS commit scripts, the context node concept corresponds to Junos OS hierarchy levels. For example, the `/configuration/system/domain-name` XPath expression sets the context node to the [edit system domain-name] hierarchy level.

We recommend including the `<xsl:template match="configuration">` template in all commit scripts. This element allows you to exclude the `/configuration/` root element from all XPath expressions in programming instructions (such as `<xsl:for-each>` or `<xsl:if>`) in the script, thus allowing you to begin XPath expressions at a Junos hierarchy level (for example, `system/domain-name`). For more information, see [“Required Boilerplate for Commit Scripts” on page 251](#).

---

## Path and Predicate Syntax

An XPath expression contains two types of syntax, a path syntax and a predicate syntax. Path syntax specifies which nodes to inspect in terms of their path locations on one of the axes in the document's hierarchy from the current context node. Several examples of path syntax follow:

- **accounting-options**—Selects an element node named **accounting-options** that is a child of the current context.
- **server/name**—Selects an element node named **name** that is a child of an element named **server** that is a child of the current context.
- **/configuration/system/domain-name**—Selects an element node named **domain-name** that is the child of an element named **system** that is the child of the root element of the document (**configuration**).
- **parent::system/host-name**—Selects an element node named **host-name** that is the child of an element named **system** that is the parent of the current context node. The **parent::** axis can be abbreviated as two periods (**..**).

The predicate syntax allows you to perform tests at each node selected by the path syntax. Only nodes that pass the test are included in the result set. A predicate appears inside square brackets (**[ ]**) after a path node. Following are several examples of predicate syntax:

- **server[name = '10.1.1.1']**—Selects an element named **server** that is a child of the current context and has a child element named **name** whose value is 10.1.1.1.
- **\*[@inactive]**—Selects any node (**\*** matches any node) that is a child of the current context and that has an attribute (**@** selects nodes from the **attribute** axis) named **inactive**.
- **route[starts-with(next-hop, '10.10.')]** —Selects an element named **route** that is a child of the current context and that has a child element named **next-hop** whose value starts with the string 10.10..

The **starts-with** function is one of many functions that are built into XPath. XPath also supports relational tests, equality tests, and many more features not listed here.

## XPath Operators

XPath supports standard logical operators, such as **AND** and **|** (or); comparison operators, such as **=**, **!=**, **<**, and **>**; and numerical operators, such as **+**, **-**, and **\***.

In XSLT, you always have to represent the less-than (**<**) operator as **&lt;**; and the less-than-or-equal-to (**<=**) operator as **&lt;=** because XSLT scripts are XML documents, and less-than signs are represented this way in XML.

For more information about XPath functions and operators, consult a comprehensive XPath reference guide. XPath is fully described in the W3C specification at <http://w3c.org/TR/xpath>.

## XSLT Templates Overview

---

An XSLT script consists of one or more sets of rules called *templates*. Each template is a segment of code that contains rules to apply when a specified node is matched. You use the `<xsl:template>` element to build templates.

There are two types of templates, named and unnamed (or match), and they are described in the following sections.

- [Unnamed \(Match\) Templates on page 22](#)
- [Named Templates on page 23](#)

### Unnamed (Match) Templates

Unnamed templates, also known as match templates, include a **match** attribute that contains an XPath expression to specify the criteria for nodes upon which the template should be invoked. In the following example, the template applies to the element named **route** that is a child of the current context and that has a child element named **next-hop** whose value starts with the string **10.10..**

```
<xsl:template match="route[starts-with(next-hop, '10.10.')] ">
  <!-- ... body of the template ... -->
</xsl:template>
```

By default, when XSLT processes a document, it recursively traverses the entire document hierarchy, inspecting each node, looking for a template that matches the current node. When a matching template is found, the contents of that template are evaluated.

The `<xsl:apply-templates>` element can be used inside an unnamed template to limit and control XSLT's default, hierarchical traversal of nodes. If the `<xsl:apply-templates>` element has a **select** attribute, only nodes matching the XPath expression defined by the attribute are traversed. Otherwise all children of the context node are traversed. If the **select** attribute is included, but does not match any nodes, nothing is traversed and nothing happens.

In the following example, the template rule matches the `<route>` element in the XML hierarchy. All the nodes containing a **changed** attribute are processed. All `<route>` elements containing a **changed** attribute are replaced with a `<new>` element.

```
<xsl:template match="route">
  <new>
    <xsl:apply-templates select="*[@changed]" />
  </new>
</xsl:template>
```

Using unnamed templates allows the script to ignore the location of a tag in the XML hierarchy. For example, if you want to convert all `<author>` tags into `<div class="author">` tags, using templates enables you to write a single rule that converts all `<author>` tags, regardless of their location in the input XML document.

For more information about how unnamed templates are used in scripts, see [xsl:template match="/" Template](#).

## Named Templates

Named templates operate like functions in traditional programming languages, although with a verbose syntax. When the complexity of a script increases or a code segment appears in multiple places, you can modularize the code and create named templates. Like functions, named templates accept arguments and run only when explicitly called.

You create a named template by using the `<xsl:template>` element and defining the **name** attribute, which is similar to a function name in traditional programming languages. Use the `<xsl:param>` tag and its **name** attribute to define parameters for the named template, and optionally include the **select** attribute to declare default values for each parameter. The **select** attribute can contain XPath expressions. If the **select** attribute is not defined, the parameter defaults to an empty string.

The following example creates a template named **my-template** and defines three parameters, one of which defaults to the string **false**, and one of which defaults to the contents of the element node named **name** that is a child of the current context node. If the script calls the template and does not pass in a parameter, the default value is used.

```
<xsl:template name="my-template">
  <xsl:param name="a"/>
  <xsl:param name="b" select="'false'"/>
  <xsl:param name="c" select="name"/>
  <!-- ... body of the template ... -->
</xsl:template>
```

To invoke a named template in a script, use the `<xsl:call-template>` element. The **name** attribute is required and defines the name of the template being called. When processed, the `<xsl:call-template>` element is replaced by the contents of the `<xsl:template>` element it names.

When you invoke a named template, you can pass arguments into the template by including the `<xsl:with-param>` child element and specifying the **name** attribute. The value of the `<xsl:with-param>` **name** attribute must match a parameter defined in the actual template; otherwise the parameter is ignored. Optionally, you can set a value for each parameter with either the **select** attribute or the content of the `<xsl:with-param>` element. If you do not define a value for the parameter in the calling environment, the script passes in the current value of the parameter if it was previously initialized, or it generates an error if the parameter was never declared. For more information about passing parameters, see [“XSLT Parameters Overview” on page 24](#).

In the following example, the template **my-template** is called with the parameter **c** containing the contents of the element node named **other-name** that is a child of the current context node.

```
<xsl:call-template name="my-template">
  <xsl:with-param name="c" select="other-name"/>
</xsl:call-template>
```

For an example showing how to use named templates in a commit script, see [“Example: Requiring and Restricting Configuration Statements” on page 422](#).

### Related Documentation

- [XSLT Parameters Overview on page 24](#)

- [xsl:apply-templates on page 40](#)
- [xsl:call-template on page 40](#)
- [xsl:param on page 46](#)
- [xsl:template on page 48](#)
- [xsl:template match="/" Template on page 49](#)
- [xsl:with-param on page 53](#)

---

## XSLT Parameters Overview

Parameters can be passed to either named or unnamed templates. Inside the template, parameters must be declared and can then be referenced by prefixing their name with the dollar sign (\$).

### Declaring Parameters

The scope of a parameter can be global or local. A parameter whose value is set by Junos OS at script initialization must be defined as a global parameter. Global parameter declarations are placed just after the style sheet declarations. A script can assign a default value to the global parameter, which is used in the event that Junos OS does not give a value to the parameter.

```
<?xml version="1.0" standalone="yes"?>
  <xsl stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
    xmlns:junos="http://xml.juniper.net/junos/*/junos"
    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"
    xmlns:ext="http://xmlsoft.org/XSLT/namespace" version="1.0">

  <!-- global parameter -->
  <xsl:param name="interface1"/>
```

Local parameters must be declared at the beginning of a block and their scope is limited to the block in which they are declared. Inside a template, you declare parameters using the **<xsl:param>** tag and **name** attribute. Optionally, declare default values for each parameter by including the **select** attribute, which can contain XPath expressions. If a template is invoked without the parameter, the default expression is evaluated, and the results are assigned to the parameter. If you do not define a default value in the template, the parameter defaults to an empty string.

The following named template **print-host-name** declares the parameter **message** and defines a default value:

```
<xsl:template name="print-host-name">
  <xsl:param name="message"
    select="concat('host-name: ', system/host-name)"/>
  <xsl:value-of select="$message"/>
</xsl:template>
```

The template accesses the value of the **message** parameter by prefixing the parameter name with the dollar sign (\$).



## Passing Parameters

When you invoke a template, you pass arguments into the template using the `<xsl:with-param>` element and `name` attribute. The value of the `<xsl:with-param>` `name` attribute must match the name of a parameter defined in the actual template; otherwise the parameter is ignored. Optionally, for each parameter you pass to a template, you can define a value using either the `select` attribute or the contents of the `<xsl:with-param>` element.

The parameter value that gets used in a template depends on how the template is called. The following three examples, which call the `print-host-name` template, illustrate the possible calling environments.

If you call a template but do not include the `<xsl:with-param>` element for a specific parameter, the default expression defined in the template is evaluated, and the results are assigned to the parameter. If there is no default value for that parameter in the template, the parameter defaults to an empty string. The following example calls the named template `print-host-name` but does not include any parameters in the call. In this case, the named template will use the default value for the `message` parameter that was defined in the `print-host-name` template, or an empty string if no default exists.

```
<xsl:template match="configuration">
  <xsl:call-template name="print-host-name"/>
</xsl:template>
```

If you call a template and include a parameter, but do not define a value for the parameter in the calling environment, the script passes in the current value of the parameter if it was previously initialized, or it generates an error if the parameter was never declared. The following example calls the named template `print-host-name` and passes in the `message` parameter, but does not include a value. If `message` is declared and initialized in the script, and the scope is visible to the block, the current value of `message` is used. If `message` is declared in the script but not initialized, the value of `message` will be an empty string. If `message` has not been declared, the script produces an error.

```
<xsl:template match="configuration">
  <xsl:call-template name="print-host-name">
    <xsl:with-param name="message"/>
  </xsl:call-template>
</xsl:template>
```

If you call a template, include the parameter, and define a value for the parameter, the template uses the provided value. The following example calls the named template `print-host-name` with the `message` parameter and a defined value, so the template uses the new value.

```
<xsl:template match="configuration">
  <xsl:call-template name="print-host-name">
    <xsl:with-param name="message"
      select=concat('Host-name passed in: ', system/host-name)"/>
  </xsl:call-template>
</xsl:template>
```

## Example: Parameters and Match Templates

The following template matches on `/`, the root of the XML document. It then generates an element named `<outside>`, which is added to the output document, and instructs the Junos OS management process (mgd) to recursively apply templates to the **configuration/system** subtree. The parameter **host** is used in the processing of any matching nodes. The value of the **host** parameter is the value of the **host-name** statement at the `[edit system]` level of the configuration hierarchy.

```
<xsl:template match="/">
  <outside>
    <xsl:apply-templates select="configuration/system">
      <xsl:with-param name="host"
select="configuration/system/host-name"/>
    </xsl:apply-templates>
  </outside>
</xsl:template>
```

The following template matches the `<system>` element, which is the top of the subtree selected in the previous example. The **host** parameter is declared with no default value. An `<inside>` element is generated, which contains the value of the **host** parameter that was defined in the `<xsl:with-param>` tag in the previous example.

```
<xsl:template match="system">
  <xsl:param name="host"/>
  <inside>
    <xsl:value-of select="$host"/>
  </inside>
</xsl:template>
```

## Example: Parameters and Named Templates

The following named template **report-changed** declares two parameters: **dot**, which defaults to the current node, and **changed**, which defaults to the **changed** attribute of the node **dot**.

```
<xsl:template name="report-changed">
  <xsl:param name="dot" select="."/>
  <xsl:param name="changed" select="$dot/@changed"/>
  <!-- ... -->
</xsl:template>
```

The next stanza calls the **report-changed** template and defines a source for the **changed** attribute different from the default source defined in the **report-changed** template. When the **report-changed** template is invoked, it will use the newly defined source for the **changed** attribute in place of the default source.

```
<xsl:template match="system">
  <xsl:call-template name="report-changed">
    <xsl:with-param name="changed" select="../@changed"/>
  </xsl:call-template>
</xsl:template>
```

Likewise, the template call can include the **dot** parameter and define a source other than the default current node, as shown here:

```
<xsl:template match="system">
  <xsl:call-template name="report-changed">
    <xsl:with-param name="dot" select="../.."/>
  </xsl:call-template>
```

```
</xsl:template>
```

- Related Documentation**
- [XSLT Templates Overview on page 22](#)
  - [xsl:param on page 46](#)
  - [xsl:with-param on page 53](#)

## XSLT Variables Overview

In XSLT scripts, you declare variables using the `<xsl:variable>` element. The **name** attribute specifies the name of the variable, which is case-sensitive. Once you declare a variable, you can reference it within an XPath expression using the variable name prefixed with a dollar sign (\$).

Variables are immutable; you can set the value of a variable only when you declare the variable, after which point, the value is fixed. You initialize a variable by including the **select** attribute and an expression in the `<xsl:variable>` tag. The following example declares and initializes the variable **location**. The **location** variable is then used to initialize the **message** variable.

```
<xsl:variable name="location" select="$dot/@location"/>
<xsl:variable name="message" select="concat('We are in ', $location, ' now.')" />
```

You can define both local and global variables. Variables are global if they are children of the `<xsl:stylesheet>` element. Otherwise, they are local. The value of a global variable is accessible anywhere in the style sheet. The scope of a local variable is limited to the template or code block in which it is defined.

XSLT variables can store any values that you can calculate or statically define. This includes data structures, XML hierarchies, and combinations of text and parameters. For example, you could assign the XML output of an operational mode command to a variable and then access the hierarchy within the variable.

The following template declares the **message** variable. The **message** variable includes both text and parameter values. The template generates a system log message by referring to the value of the message variable.

```
<xsl:template name="emit-syslog">
  <xsl:param name="user"/>
  <xsl:param name="date"/>
  <xsl:param name="device"/>
  <xsl:variable name="message">
    <xsl:text>Device </xsl:text>
    <xsl:value-of select="$device"/>
    <xsl:text> was changed on </xsl:text>
    <xsl:value-of select="$date"/>
    <xsl:text> by user '</xsl:text>
    <xsl:value-of select="$user"/>
    <xsl:text>.' </xsl:text>
  </xsl:variable>
  <syslog>
    <message>
      <xsl:value-of select="$message"/>
    </message>
  </syslog>
</xsl:template>
```

The resulting system log message is as follows:

Device *device-name* was changed on *date* by user '*user*.'

[Table 4 on page 28](#) provides examples of XSLT variable declarations along with pseudocode explanations.

**Table 4: Examples and Pseudocode for XSLT Variable Declaration**

Variable Declaration	Pseudocode Explanation
<code>&lt;xsl:variable name="mpls" select="protocols/mpls"/&gt;</code>	Assigns the <b>[edit protocols mpls]</b> hierarchy level to the variable named <b>mpls</b> .
<code>&lt;xsl:variable name="color" select="data[name = 'color']/value"/&gt;</code>	Assigns the value of the <b>color</b> macro parameter to a variable named <b>color</b> . The <b>&lt;data&gt;</b> element in the XPath expression is useful in commit script macros. For more information, see <a href="#">“Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements” on page 320</a> .

Related Documentation

- [xsl:variable on page 52](#)

## XSLT Programming Instructions Overview

XSLT has a number of traditional programming instructions. Their form tends to be verbose, because their syntax is built from XML elements.

The XSLT programming instructions most commonly used in commit, op, event, and SNMP scripts, which provide flow control within a script, are described in the following sections:

- [<xsl:choose> Programming Instruction on page 28](#)
- [<xsl:for-each> Programming Instruction on page 29](#)
- [<xsl:if> Programming Instruction on page 29](#)
- [Sample XSLT Programming Instructions and Pseudocode on page 30](#)

### <xsl:choose> Programming Instruction

The **<xsl:choose>** instruction is a conditional construct that causes different instructions to be processed in different circumstances. It is similar to a switch statement in traditional programming languages. The **<xsl:choose>** instruction contains one or more **<xsl:when>** elements, each of which tests an XPath expression. If the test evaluates to true, the XSLT processor executes the instructions in the **<xsl:when>** element. After the XSLT processor finds an XPath expression in an **<xsl:when>** element that evaluates to true, the XSLT processor ignores all subsequent **<xsl:when>** elements contained in the **<xsl:choose>** instruction, even if their XPath expressions evaluate to true. In other words, the XSLT

processor processes only the instructions contained in the first `<xsl:when>` element whose **test** attribute evaluates to true. If none of the `<xsl:when>` elements' **test** attributes evaluate to true, the content of the optional `<xsl:otherwise>` element, if one is present, is processed.

The `<xsl:choose>` instruction is similar to a switch statement in other programming languages. The `<xsl:when>` element is the “case” of the switch statement, and you can add any number of `<xsl:when>` elements. The `<xsl:otherwise>` element is the “default” of the switch statement.

```
<xsl:choose>
  <xsl:when test="xpath-expression">
    ...
  </xsl:when>
  <xsl:when test="another-xpath-expression">
    ...
  </xsl:when>
  <xsl:otherwise>
    ...
  </xsl:otherwise>
</xsl:choose>
```

### `<xsl:for-each>` Programming Instruction

The `<xsl:for-each>` element tells the XSLT processor to gather together a set of nodes and process them one by one. The nodes are selected by the XPath expression specified by the **select** attribute. Each of the nodes is then processed according to the instructions held in the `<xsl:for-each>` construct.

```
<xsl:for-each select="xpath-expression">
  ...
</xsl:for-each>
```

Code inside the `<xsl:for-each>` instruction is evaluated recursively for each node that matches the XPath expression. That is, the current context is moved to each node selected by the `<xsl:for-each>` clause, and processing is relative to that current context.

In the following example, the `<xsl:for-each>` construct recursively processes each node in the [system syslog file] hierarchy. It updates the current context to each matching node and prints the value of the **name** element, if one exists, that is a child of the current context.

```
<xsl:for-each select="system/syslog/file">
  <xsl:value-of select="name"/>
</xsl:for-each>
```

### `<xsl:if>` Programming Instruction

An `<xsl:if>` programming instruction is a conditional construct that causes instructions to be processed if the XPath expression held in the **test** attribute evaluates to **true**.

```
<xsl:if test="xpath-expression">
  ...executed if test expression evaluates to true
</xsl:if>
```

There is no corresponding else clause.

## Sample XSLT Programming Instructions and Pseudocode

Table 5 on page 30 presents examples that use several XSLT programming instructions along with pseudocode explanations.

Table 5: Examples and Pseudocode for XSLT Programming Instructions

Programming Instruction	Pseudocode Explanation
<pre> &lt;xsl:choose&gt;   &lt;xsl:when test="system/host-name"&gt;     &lt;change&gt;       &lt;system&gt;         &lt;host-name&gt;M320&lt;/host-name&gt;       &lt;/system&gt;     &lt;/change&gt;   &lt;/xsl:when&gt;   &lt;xsl:otherwise&gt;     &lt;xnm:error&gt;       &lt;message&gt;         Missing [edit system host-name]         M320.       &lt;/message&gt;     &lt;/xnm:error&gt;   &lt;/xsl:otherwise&gt; &lt;/xsl:choose&gt; </pre>	<p>When the <b>host-name</b> statement is included at the [edit system] hierarchy level, change the hostname to <b>M320</b>.</p> <p>Otherwise, issue the warning message: <b>Missing [edit system host-name] M320</b>.</p>
<pre> &lt;xsl:for-each select="interfaces/ interface[starts-with(name, 'ge-')]/unit"&gt; </pre>	<p>For each Gigabit Ethernet interface configured at the [edit interfaces <i>ge-fpc/pic/port</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
<pre> &lt;xsl:for-each select="data[not(value)]/name"&gt; </pre>	<p>Select any macro parameter that does not contain a parameter value.</p> <p>In other words, match all <b>apply-macro</b> statements of the following form:</p> <pre>   apply-macro <i>apply-macro-name</i> {     <i>parameter-name</i>;   } </pre> <p>And ignore all <b>apply-macro</b> statements of the form:</p> <pre>   apply-macro <i>apply-macro-name</i> {     <i>parameter-name</i> <i>parameter-value</i>;   } </pre>
<pre> &lt;xsl:if test="not(system/host-name)"&gt; </pre>	<p>If the <b>host-name</b> statement is not included at the [edit system] hierarchy level.</p>
<pre> &lt;xsl:if test="apply-macro[name = 'no-igp']" </pre>	<p>If the <b>apply-macro</b> statement named <b>no-igp</b> is included at the current hierarchy level.</p>
<pre> &lt;xsl:if test="not(..//apply-macro[name = 'no-ldp'])" </pre>	<p>If the <b>apply-macro</b> statement with the name <b>no-ldp</b> is not included two hierarchy levels above the current hierarchy level.</p>

Related • [xsl:choose on page 41](#)  
Documentation

- [xsl:for-each on page 43](#)
- [xsl:if on page 44](#)
- [xsl:otherwise on page 45](#)
- [xsl:when on page 53](#)

## XSLT Recursion Overview

XSLT depends on recursion as a looping mechanism. Recursion occurs when a section of code calls itself, either directly or indirectly. Both named and unnamed templates can use recursion, and different templates can use mutual recursion, one calling another that in turn calls the first.

To avoid infinite recursion and excessive consumption of system resources, the Junos OS management process (mgd) limits the maximum recursion to 5000 levels. If this limit is reached, the script fails.

In the following example, an unnamed template matches on a `<count>` element. It then calls the `<count-to-max>` template, passing the value of the `count` element as `max`. The `<count-to-max>` template starts by declaring both the `max` and `cur` parameters and setting the default value of each to 1 (one). Although the optional default value for `max` is one, the template will use the value passed in from the `count` template. Then the current value of `cur` is emitted in an `<out>` element. Finally, if `cur` is less than `max`, the `<count-to-max>` template recursively invokes itself, passing `cur + 1` as `cur`. This recursive pass then outputs the next number and repeats the recursion until `cur` equals `max`.

```
<xsl:template match="count">
  <xsl:call-template name="count-to-max">
    <xsl:with-param name="max" select="."/>
  </xsl:call-template>
</xsl:template>

<xsl:template name="count-to-max">
  <xsl:param name="cur" select="'1'"/>
  <xsl:param name="max" select="'1'"/>

  <out><xsl:value-of select="$cur"/></out>

  <xsl:if test="$cur < $max">
    <xsl:call-template name="count-to-max">
      <xsl:with-param name="cur" select="$cur + 1"/>
      <xsl:with-param name="max" select="$max"/>
    </xsl:call-template>
  </xsl:if>
</xsl:template>
```

Given a `max` value of 10, the values contained in the `<out>` tag are 1, 2, 3, 4, 5, 6, 7, 8, 9, and 10.

## XSLT Context (Dot) Overview

The current context node changes as an `<xsl:apply-templates>` instruction traverses the document hierarchy and as an `<xsl:for-each>` instruction examines each node that

matches an XPath expression. All relative node references are relative to the current context node. This node is abbreviated “.” (read: dot) and can be referred to in XPath expressions, allowing explicit references to the current node.

The following example contains four uses for “.”. The **system** node is saved in the **system** variable for use inside the **<xsl:for-each>** instruction, where the value of “.” will have changed. The **for-each select** expression uses “.” to mean the value of the **name** element. The “.” is then used to pull the value of the **name** element into the **<tag>** element. The **<xsl:if>** test then uses “.” to reference the value of the current context node.

```
<xsl:template match="system">
  <xsl:variable name="system" select="."/>
  <xsl:for-each select="name-server/name[starts-with(., '10.')] ">
    <tag><xsl:value-of select="."/></tag>
    <xsl:if test=".= '10.1.1.1'">
      <match>
        <xsl:value-of select="$system/host-name"/>
      </match>
    </xsl:if>
  </xsl:for-each>
</xsl:template>
```



## CHAPTER 4

# Standard XPath and XSLT Functions Used in Automation Scripts

- [concat\(\)](#) on page 33
- [contains\(\)](#) on page 34
- [count\(\)](#) on page 34
- [last\(\)](#) on page 34
- [name\(\)](#) on page 35
- [not\(\)](#) on page 35
- [position\(\)](#) on page 35
- [starts-with\(\)](#) on page 36
- [string-length\(\)](#) on page 36
- [substring-after\(\)](#) on page 37
- [substring-before\(\)](#) on page 37

### concat()

---

<b>Syntax</b>	<code>string concat(string, string+)</code>
<b>Description</b>	Return the concatenation of the arguments.
<b>Usage Examples</b>	See “Example: Limiting the Number of E1 Interfaces” on page 388, “Example: Controlling IS-IS and MPLS Interfaces” on page 366, “Example: Adding T1 Interfaces to a RIP Group” on page 336, “Example: Configuring Administrative Groups for LSPs” on page 350, and “Example: Configuring Dual Routing Engines” on page 357.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">contains()</a> on page 34</li><li>• <a href="#">starts-with()</a> on page 36</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li><li>• <a href="#">substring-before()</a> on page 37</li></ul>

## contains()

---

<b>Syntax</b>	<i>boolean</i> contains( <i>string</i> , <i>string</i> )
<b>Description</b>	Return TRUE if the first string argument contains the second string argument, otherwise return FALSE.
<b>Usage Examples</b>	See “Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">concat()</a> on page 33</li><li>• <a href="#">starts-with()</a> on page 36</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li><li>• <a href="#">substring-before()</a> on page 37</li></ul>

## count()

---

<b>Syntax</b>	<i>number</i> count( <i>node-set</i> )
<b>Description</b>	Return the number of nodes in the argument node-set.
<b>Usage Examples</b>	See “Example: Limiting the Number of EI Interfaces” on page 388.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">last()</a> on page 34</li><li>• <a href="#">name()</a> on page 35</li><li>• <a href="#">not()</a> on page 35</li><li>• <a href="#">position()</a> on page 35</li></ul>

## last()

---

<b>Syntax</b>	<i>number</i> last()
<b>Description</b>	Return the index of the last node in the list that is currently being evaluated, which is equal to the number of items in the processed node list.
<b>Usage Examples</b>	See “Example: Limiting the Number of EI Interfaces” on page 388.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">count()</a> on page 34</li><li>• <a href="#">name()</a> on page 35</li><li>• <a href="#">not()</a> on page 35</li><li>• <a href="#">position()</a> on page 35</li></ul>

## name()

<b>Syntax</b>	<i>string</i> <code>name(&lt;node-set&gt;)</code>
<b>Description</b>	Return the full name of the first node in the node set, including the prefix for its namespace declared in the source document. If no argument is passed, the function returns the full name of the context node.
<b>Usage Examples</b>	See <a href="#">emit-change Template (SLAX and XSLT)</a> and <a href="#">emit_change(Python)</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">count()</a> on page 34</li> <li>• <a href="#">last()</a> on page 34</li> <li>• <a href="#">not()</a> on page 35</li> <li>• <a href="#">position()</a> on page 35</li> </ul>

## not()

<b>Syntax</b>	<i>boolean</i> <code>not(boolean)</code>
<b>Description</b>	Return TRUE if the argument is FALSE, and FALSE if the argument is TRUE.
<b>Usage Examples</b>	See “Example: Requiring and Restricting Configuration Statements” on page 422, “Example: Controlling IS-IS and MPLS Interfaces” on page 366, “Example: Configuring a Default Encapsulation Type” on page 354, “Example: Controlling LDP Configuration” on page 370, “Example: Adding a Final then accept Term to a Firewall” on page 331, “Example: Configuring Administrative Groups for LSPs” on page 350, “Example: Configuring Dual Routing Engines” on page 357, and “Example: Preventing Import of the Full Routing Table” on page 416.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">count()</a> on page 34</li> <li>• <a href="#">last()</a> on page 34</li> <li>• <a href="#">name()</a> on page 35</li> <li>• <a href="#">position()</a> on page 35</li> </ul>

## position()

<b>Syntax</b>	<i>number</i> <code>position()</code>
<b>Description</b>	Return the position of the context node among the list of nodes that are currently being evaluated.
<b>Usage Examples</b>	See “Example: Adding a Final then accept Term to a Firewall” on page 331 and “Example: Prepending a Global Policy” on page 411.

- |                       |   |
|-----------------------|---|
| Related Documentation | <ul style="list-style-type: none"><li>• <a href="#">count()</a> on page 34</li><li>• <a href="#">last()</a> on page 34</li><li>• <a href="#">name()</a> on page 35</li><li>• <a href="#">not()</a> on page 35</li></ul> |
|-----------------------|---|

---

## starts-with()

---

- |                       |   |
|-----------------------|---|
| Syntax                | <i>boolean</i> starts-with( <i>string</i> , <i>string</i> )   |
| Description           | Return TRUE if the first string argument starts with the second string argument, otherwise return FALSE.  |
| Usage Examples        | See “Example: Imposing a Minimum MTU Setting” on page 382, “Example: Limiting the Number of E1 Interfaces” on page 388, “Example: Limiting the Number of ATM Virtual Circuits” on page 384, “Example: Adding T1 Interfaces to a RIP Group” on page 336, “Example: Configuring a Default Encapsulation Type” on page 354, and “Example: Configuring Dual Routing Engines” on page 357.                                 |
| Related Documentation | <ul style="list-style-type: none"><li>• <a href="#">concat()</a> on page 33</li><li>• <a href="#">contains()</a> on page 34</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li><li>• <a href="#">substring-before()</a> on page 37</li></ul> |

---

## string-length()

---

- |                       |   |
|-----------------------|---|
| Syntax                | <i>number</i> string-length(< <i>string</i> >)  |
| Description           | Return the number of characters in the string. If the argument is omitted, it returns the string value of the context node.   |
| Usage Examples        | See “Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343.   |
| Related Documentation | <ul style="list-style-type: none"><li>• <a href="#">concat()</a> on page 33</li><li>• <a href="#">contains()</a> on page 34</li><li>• <a href="#">starts-with()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li><li>• <a href="#">substring-before()</a> on page 37</li></ul> |

## substring-after()

---

<b>Syntax</b>	<i>string</i> substring-after( <i>string</i> , <i>string</i> )
<b>Description</b>	Return the portion of the first string argument that follows the occurrence of the second argument substring within the first. If the second string is not contained in the first string, or if the second string is empty, the function returns an empty string.
<b>Usage Examples</b>	See “ <a href="#">Example: Limiting the Number of EI Interfaces</a> ” on page 388 and “ <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses</a> ” on page 343.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">concat()</a> on page 33</li><li>• <a href="#">contains()</a> on page 34</li><li>• <a href="#">starts-with()</a> on page 36</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-before()</a> on page 37</li></ul>

## substring-before()

---

<b>Syntax</b>	<i>string</i> substring-before( <i>string</i> , <i>string</i> )
<b>Description</b>	Return the portion of the first string argument that precedes the occurrence of the second argument substring within the first. If the second string is not contained in the first string, or if the second string is empty, the function returns an empty string.
<b>Usage Examples</b>	See “ <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses</a> ” on page 343.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">concat()</a> on page 33</li><li>• <a href="#">contains()</a> on page 34</li><li>• <a href="#">starts-with()</a> on page 36</li><li>• <a href="#">string-length()</a> on page 36</li><li>• <a href="#">substring-after()</a> on page 37</li></ul>



## CHAPTER 5

# Standard XSLT Elements and Attributes Used in Automation Scripts

- [xsl:apply-templates on page 40](#)
- [xsl:call-template on page 40](#)
- [xsl:choose on page 41](#)
- [xsl:comment on page 42](#)
- [xsl:copy-of on page 42](#)
- [xsl:element on page 43](#)
- [xsl:for-each on page 43](#)
- [xsl:if on page 44](#)
- [xsl:import on page 44](#)
- [xsl:otherwise on page 45](#)
- [xsl:param on page 46](#)
- [xsl:stylesheet on page 47](#)
- [xsl:template on page 48](#)
- [xsl:template match="/" Template on page 49](#)
- [xsl:text on page 51](#)
- [xsl:value-of on page 51](#)
- [xsl:variable on page 52](#)
- [xsl:when on page 53](#)
- [xsl:with-param on page 53](#)

## xsl:apply-templates

---

<b>Syntax</b>	<pre>&lt;xsl:apply-templates select="node-set-expression"&gt;   &lt;xsl:with-param name="qualified-name" select="expression"&gt;     ...   &lt;/xsl:with-param&gt; &lt;/xsl:apply-templates&gt;</pre>
<b>Description</b>	Apply one or more templates, according to the value of the <b>select</b> attribute. If the <b>select</b> attribute is not included, the script recursively processes all child nodes of the current node. If the <b>select</b> attribute is present, the processor only applies templates to the child elements that match the expression of the <b>select</b> attribute, which must evaluate to a node-set. The <b>&lt;xsl:template&gt;</b> instruction dictates which elements are transformed according to which template. The templates that are applied are passed the parameters specified by the <b>&lt;xsl:with-param&gt;</b> elements within the <b>&lt;xsl:apply-templates&gt;</b> instruction.
<b>Attributes</b>	<b>select</b> —(Optional) Selects the nodes to which the processor applies templates. By default, the processor applies templates to the child nodes of the current node.
<b>Usage Examples</b>	See <a href="#">“Example: Adding a Final then accept Term to a Firewall” on page 331</a> and <a href="#">“Example: Preventing Import of the Full Routing Table” on page 416</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">XSLT Templates Overview on page 22</a></li><li>• <a href="#">xsl:call-template on page 40</a></li><li>• <a href="#">xsl:param on page 46</a></li><li>• <a href="#">xsl:template on page 48</a></li><li>• <a href="#">xsl:variable on page 52</a></li><li>• <a href="#">xsl:with-param on page 53</a></li></ul>

## xsl:call-template

---

<b>Syntax</b>	<pre>&lt;xsl:call-template name="qualified-name"&gt;   &lt;xsl:with-param name="qualified-name" select="expression"&gt;     ...   &lt;/xsl:with-param&gt; &lt;/xsl:call-template&gt;</pre>
<b>Description</b>	Call a named template. The <b>&lt;xsl:with-param&gt;</b> elements within the <b>&lt;xsl:call-template&gt;</b> instruction define the parameters that are passed to the template.
<b>Attributes</b>	<b>name</b> —Specifies the name of the template to call.
<b>Usage Examples</b>	See <a href="#">“Example: Requiring and Restricting Configuration Statements” on page 422</a> , <a href="#">“Example: Imposing a Minimum MTU Setting” on page 382</a> , and <a href="#">“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343</a> .



- Related Documentation**
- [XSLT Templates Overview on page 22](#)
  - [xsl:apply-templates on page 40](#)
  - [xsl:param on page 46](#)
  - [xsl:template on page 48](#)
  - [xsl:variable on page 52](#)
  - [xsl:with-param on page 53](#)

---

## xsl:choose

**Syntax**

```
<xsl:choose>
  <xsl:when test="boolean-expression">
    ...
  </xsl:when>
  <xsl:otherwise>
    ...
  </xsl:otherwise>
</xsl:choose>
```

**Description** Evaluate multiple conditional tests, and execute instructions for the first test that evaluates to TRUE or execute an optional default set of instructions if all tests evaluate to FALSE. The **<xsl:choose>** instruction contains one or more **<xsl:when>** elements, each of which tests a Boolean expression. If the test evaluates to TRUE, the XSLT processor executes the instructions in the **<xsl:when>** element, and ignores all subsequent **<xsl:when>** elements. The XSLT processor processes only the instructions contained in the first **<xsl:when>** element whose **test** attribute evaluates to TRUE. If none of the **<xsl:when>** elements' **test** attributes evaluate to TRUE, the content of the optional **<xsl:otherwise>** element, if one is present, is processed.

**Usage Examples** See “Example: Configuring Dual Routing Engines” on page 357, “Example: Preventing Import of the Full Routing Table” on page 416, and “Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343.

- Related Documentation**
- [XSLT Programming Instructions Overview on page 28](#)
  - [xsl:for-each on page 43](#)
  - [xsl:if on page 44](#)
  - [xsl:otherwise on page 45](#)
  - [xsl:when on page 53](#)

## xsl:comment

---

<b>Syntax</b>	<pre>&lt;xsl:comment&gt; ... &lt;/xsl:comment&gt;</pre>
<b>Description</b>	<p>Generate a comment node within the final document. The content within the <b>&lt;xsl:comment&gt;</b> element determines the value of the comment. The content must not contain two hyphens next to each other (- -); this sequence is not allowed in comments.</p> <p>XSLT files can contain ordinary comments delimited by <b>&lt;!--</b> and <b>--&gt;</b> such as <b>&lt;!-- ... Insert your comment here ... --&gt;</b>, but these are ignored by the processor. To generate a comment within the final document, use an <b>&lt;xsl:comment&gt;</b> element.</p>
<b>Usage Examples</b>	See <a href="#">“Example: Adding a Final then accept Term to a Firewall” on page 331</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">xsl:import on page 44</a></li><li>• <a href="#">xsl:stylesheet on page 47</a></li></ul>

## xsl:copy-of

---

<b>Syntax</b>	<pre>&lt;xsl:copy-of select="expression"/&gt;</pre>
<b>Description</b>	Create a copy of what is selected by the expression defined in the <b>select</b> attribute. Namespace nodes, child nodes, and attributes of the current node are automatically copied as well.
<b>Attributes</b>	<b>select</b> —XPath expression specifying which nodes to copy.
<b>Usage Examples</b>	See <a href="#">“Example: Requiring and Restricting Configuration Statements” on page 422</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">xsl:element on page 43</a></li><li>• <a href="#">xsl:text on page 51</a></li><li>• <a href="#">xsl:value-of on page 51</a></li></ul>

## xsl:element

<b>Syntax</b>	<code>&lt;xsl:element name="expression"/&gt;</code>
<b>Description</b>	Create an element node in the output document.
<b>Attributes</b>	name—Specifies the name of the element to be created. The value of the <b>name</b> attribute can be set to an expression that is extracted from the input XML document and evaluated at run time. To do this, enclose an XML element in curly brackets, as in <code>&lt;xsl:element name="{Sisis-level-1}"</code> .
<b>Usage Examples</b>	See “ <a href="#">Example: Creating a Complex Configuration Based on a Simple Interface Configuration</a> ” on page 375.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">xsl:copy-of on page 42</a></li> <li>• <a href="#">xsl:text on page 51</a></li> <li>• <a href="#">xsl:value-of on page 51</a></li> </ul>

## xsl:for-each

<b>Syntax</b>	<pre>&lt;xsl:for-each select="node-set-expression"&gt;     ... &lt;/xsl:for-each&gt;</pre>
<b>Description</b>	Include a looping mechanism that repeats XSL processing for each XML element in the specified node-set. The element nodes are selected by the XPath expression defined by the <b>select</b> attribute. Each of the nodes is then processed according to the instructions contained in the <code>&lt;xsl:for-each&gt;</code> element.
<b>Attributes</b>	select—Specifies an XPath expression that selects the nodes to be processed.
<b>Usage Examples</b>	See “ <a href="#">Example: Requiring and Restricting Configuration Statements</a> ” on page 422, “ <a href="#">Example: Imposing a Minimum MTU Setting</a> ” on page 382, “ <a href="#">Example: Limiting the Number of E1 Interfaces</a> ” on page 388, “ <a href="#">Example: Adding T1 Interfaces to a RIP Group</a> ” on page 336, “ <a href="#">Example: Configuring Administrative Groups for LSPs</a> ” on page 350, and “ <a href="#">Example: Configuring Dual Routing Engines</a> ” on page 357.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">XSLT Programming Instructions Overview on page 28</a></li> <li>• <a href="#">XPath Overview on page 19</a></li> <li>• <a href="#">xsl:choose on page 41</a></li> <li>• <a href="#">xsl:if on page 44</a></li> <li>• <a href="#">xsl:otherwise on page 45</a></li> <li>• <a href="#">xsl:when on page 53</a></li> </ul>

## xsl:if

---

<b>Syntax</b>	<pre>&lt;xsl:if test="expression"&gt;     ... &lt;/xsl:if&gt;</pre>
<b>Description</b>	Include a conditional construct that causes instructions to be processed if the expression held in the <b>test</b> attribute evaluates to <b>TRUE</b> .
<b>Attributes</b>	<b>test</b> —Specifies the expression to evaluate.
<b>Usage Examples</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Adding T1 Interfaces to a RIP Group on page 336</a></li><li>• <a href="#">Example: Configuring Dual Routing Engines on page 357</a></li><li>• <a href="#">Example: Limiting the Number of E1 Interfaces on page 388</a></li><li>• <a href="#">Example: Requiring and Restricting Configuration Statements on page 422</a></li></ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">XSLT Programming Instructions Overview on page 28</a></li><li>• <a href="#">xsl:choose on page 41</a></li><li>• <a href="#">xsl:for-each on page 43</a></li><li>• <a href="#">xsl:otherwise on page 45</a></li><li>• <a href="#">xsl:when on page 53</a></li></ul>

## xsl:import

---

<b>Syntax</b>	<pre>&lt;xsl:import href="../import/junos.xsl"/&gt;</pre>
<b>Description</b>	<p>Import rules from an external style sheet. Provides access to all the declarations and templates within the imported style sheet, and allows you to override them with your own if needed. Any <b>&lt;xsl:import&gt;</b> elements must be the first elements within the style sheet, the first children of the <b>&lt;xsl:stylesheet&gt;</b> document element. The path can be any URI. The <b>../import/junos.xsl</b> path shown in the syntax is standard for all commit scripts, op scripts, and event scripts.</p> <p>Imported rules are overwritten by any subsequent matching rules within the importing style sheet. If more than one style sheet is imported, the style sheets imported last override each previous import where the rules match.</p>
<b>Attributes</b>	<b>href</b> —Specifies the location of the imported style sheet.
<b>Usage Examples</b>	See <a href="#">“Example: Adding a Final then accept Term to a Firewall” on page 331</a> , <a href="#">“Example: Configuring a Default Encapsulation Type” on page 354</a> , <a href="#">“Example: Configuring Dual Routing Engines” on page 357</a> , <a href="#">“Example: Controlling IS-IS and MPLS Interfaces” on page 366</a> , <a href="#">“Example: Prepending a Global Policy” on page 411</a> , and <a href="#">“Example: Preventing Import of the Full Routing Table” on page 416</a> .

- Related Documentation**
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)
  - [Required Boilerplate for Commit Scripts on page 251](#)
  - [Required Boilerplate for Event Scripts on page 615](#)
  - [Required Boilerplate for Op Scripts on page 445](#)
  - [xsl:stylesheet on page 47](#)

---

## xsl:otherwise

---

<b>Syntax</b>	<pre>&lt;xsl:otherwise&gt;     ... &lt;/xsl:otherwise&gt;</pre>
<b>Description</b>	Within an <b>&lt;xsl:choose&gt;</b> instruction, include a default set of instructions that are processed if none of the expressions defined in the <b>test</b> attributes of the <b>&lt;xsl:when&gt;</b> elements evaluate to <b>TRUE</b> .
<b>Usage Examples</b>	See “ <a href="#">Example: Configuring Dual Routing Engines</a> ” on page 357 and “ <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses</a> ” on page 343.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">XSLT Programming Instructions Overview on page 28</a></li><li>• <a href="#">xsl:choose on page 41</a></li><li>• <a href="#">xsl:for-each on page 43</a></li><li>• <a href="#">xsl:if on page 44</a></li><li>• <a href="#">xsl:when on page 53</a></li></ul>

## xsl:param

---

<b>Syntax</b>	<pre>&lt;xsl:param name="<i>qualified-name</i>" select="<i>expression</i>"&gt;   ... &lt;/xsl:param&gt;</pre>
<b>Description</b>	Declare a parameter for a template or for the style sheet as a whole. A template parameter must be declared within the template element. A global parameter, the scope of which is the entire style sheet, must be declared at the top level of the style sheet.
<b>Attributes</b>	<p><b>name</b>—Defines the name of the parameter.</p> <p><b>select</b>—(Optional) XPath expression defining the default value for the parameter, which is used if the person or client application that executes the script does not explicitly provide a value. The <b>select</b> attribute or the content of the <b>&lt;xsl:param&gt;</b> element can define the default value. Do not specify both a <b>select</b> attribute and content; we recommend using the <b>select</b> attribute so as not to create a result tree fragment.</p>
<b>Usage Examples</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Imposing a Minimum MTU Setting on page 382</a></li><li>• <a href="#">Example: Limiting the Number of ATM Virtual Circuits on page 384</a></li><li>• <a href="#">Example: Limiting the Number of E1 Interfaces on page 388</a></li><li>• <a href="#">Example: Preventing Import of the Full Routing Table on page 416</a></li><li>• <a href="#">Example: Requiring and Restricting Configuration Statements on page 422</a></li></ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">XSLT Parameters Overview on page 24</a></li><li>• <a href="#">XSLT Templates Overview on page 22</a></li><li>• <a href="#">xsl:apply-templates on page 40</a></li><li>• <a href="#">xsl:call-template on page 40</a></li><li>• <a href="#">xsl:template on page 48</a></li><li>• <a href="#">xsl:variable on page 52</a></li><li>• <a href="#">xsl:with-param on page 53</a></li></ul>

## xsl:stylesheet

---

<b>Syntax</b>	<pre>&lt;xsl:stylesheet version="1.0" xmlns:ext="URI"&gt;   &lt;xsl:import href="../import/junos.xsl"/&gt;   ... &lt;/xsl:stylesheet&gt;</pre>
<b>Description</b>	<p>Include the document element for the style sheet. This element defines the root element of the style sheet, which contains all the top-level elements such as global variable and parameter declarations, import elements, and templates. Optionally, namespace mappings, which include an extension prefix and Uniform Resource Identifier (URI), can be included as attributes in the opening <b>&lt;xsl:stylesheet&gt;</b> tag.</p> <p>Any <b>&lt;xsl:import&gt;</b> elements must be the first elements within the style sheet, the first children of the <b>&lt;xsl:stylesheet&gt;</b> document element. The path can be any Uniform Resource Identifier (URI). The <b>../import/junos.xsl</b> path shown in the syntax is standard for all commit scripts, op scripts, and event scripts.</p>
<b>Attributes</b>	<p><b>version</b>—Specifies the version of XSLT that is being used. Junos OS supports XSLT version 1.0.</p> <p><b>xmlns:ext="URI"</b>—(Optional) Maps a namespace prefix to the URI for extension elements.</p>
<b>Usage Examples</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Adding a Final then accept Term to a Firewall on page 331</a></li><li>• <a href="#">Example: Configuring Administrative Groups for LSPs on page 350</a></li><li>• <a href="#">Example: Configuring a Default Encapsulation Type on page 354</a></li><li>• <a href="#">Example: Customizing Output of the show interfaces terse Command Using an Op Script on page 477</a></li></ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Required Boilerplate for Commit Scripts on page 251</a></li><li>• <a href="#">Required Boilerplate for Event Scripts on page 615</a></li><li>• <a href="#">Required Boilerplate for Op Scripts on page 445</a></li><li>• <a href="#">XSLT Namespace on page 19</a></li><li>• <a href="#">xsl:import on page 44</a></li></ul>

## xsl:template

---

**Syntax**

```
<xsl:template match="pattern" mode="qualified-name" name="qualified-name"
priority="integer">
  <xsl:param name="qualified-name" select="expression">
    ...
  </xsl:param>
  ...
</xsl:template>
```

**Description** Declare a template that contains rules to apply when a specified node is matched. The **match** attribute associates the template with an XML element. The **match** attribute can also be used to define a template for a whole branch of an XML document. For example, **match="/"** matches the root element of the document. Although the **match** and **name** attributes are optional, one of the two attributes must be included in the template definition.

When templates are applied to a node set using the **<xsl:apply-templates>** instruction, they might be applied in a particular mode; the **mode** attribute in the **<xsl:template>** instruction indicates the mode in which a template needs to be applied for the template to be used. If templates are applied in the specified mode, the **match** attribute is used to determine whether the template can be used with the particular node. If more than one template matches a node in the specified mode, the priority attribute determines which template is used. The highest priority wins. If no priority is specified explicitly, the priority of a template is determined by the **match** attribute.

You can pass template parameters using the **<xsl:with-param>** element. To receive a parameter, the template must contain an **<xsl:param>** element that declares a parameter of that name. These parameters are listed before the body of the template, which is used to process the node and create a result.

**Attributes** **match**—(Optional) XPath expression specifying the nodes to which to apply the template. If this attribute is omitted, the **name** attribute must be included.

**mode**—(Optional) Indicate the mode in which a template needs to be applied for the template to be used.

**name**—(Optional) Specify a name for the template. Named templates can be explicitly called with the **<xsl:call-template>** element. If the **name** attribute is omitted, the **match** attribute must be included.

**priority**—(Optional) Specify a numeric priority for the template.

- Usage Examples**
- [Example: Adding a Final then accept Term to a Firewall on page 331](#)
  - [Example: Adding T1 Interfaces to a RIP Group on page 336](#)
  - [Example: Automatically Configuring Logical Interfaces and IP Addresses on page 343](#)
  - [Example: Customizing Output of the show interfaces terse Command Using an Op Script on page 477](#)
  - [Example: Requiring and Restricting Configuration Statements on page 422](#)



- Related Documentation**
- [XSLT Templates Overview on page 22](#)
  - [XSLT Parameters Overview on page 24](#)
  - [xsl:apply-templates on page 40](#)
  - [xsl:call-template on page 40](#)
  - [xsl:param on page 46](#)
  - [xsl:variable on page 52](#)
  - [xsl:with-param on page 53](#)


## xsl:template match="/" Template

**Syntax**      `<xsl:template match="/">`

**Description**    The `<xsl:template match="/">` template is an unnamed template in the `junos.xsl` file that allows you to use shortened XPath expressions in commit scripts. You must import the `junos.xsl` file to use this template. However, because this template is not in the `jcs` namespace, you do not need to map to the `jcs` namespace in your style sheet declaration in order to use this template.

Junos OS provides XML-formatted input to a script. Commit script input consists of an XML representation of the post-inheritance candidate configuration file. When you execute a script, the Junos OS management process (mgd) generates an XML-formatted output document as the product of its evaluation of the input document, as shown in [Figure 2 on page 49](#).

**Figure 2: Commit Script Input and Output**



Generally, an XSLT engine uses recursion to evaluate the entire input document. However, the `<xsl:apply-templates>` instruction allows you to limit the scope of the evaluation so that the management process (the Junos OS's XSLT engine) must evaluate only a subset of the input document.

The `<xsl:template match="/">` template is an unnamed template that uses the `<xsl:apply-templates>` instruction to specify the contents of the input document's `<configuration>` element as the only node to be evaluated in the generation of the output document.

The `<xsl:template match="/">` template contains the following tags:

```

1  <xsl:template match="/">
2    <commit-script-results>
3      <xsl:apply-templates select="commit-script-input/configuration"/>
4    </commit-script-results>
  
```

```
5    </xsl:template>
```

Line 1 matches the root node of the input document. When the management process sees the root node of the input document, this template is applied.

```
1    <xsl:template match="/">
```

Line 2 designates the root, top-level tag of the output document. Thus, Line 2 specifies that the evaluation of the input document results in an output document whose top-level tag is **<commit-script-results>**.

```
2        <commit-script-results>
```

Line 3 limits the scope of the evaluation of the input document to the contents of the **<configuration>** element, which is a child of the **<commit-script-input>** element.

```
3        <xsl:apply-templates select="commit-script-input/configuration"/>
```

Lines 4 and 5 are closing tags.

You do not need to explicitly include the **<xsl:template match="/">** template in your scripts because this template is included in the import file **junos.xsl**.

When the **<xsl:template match="/">** template executes the **<xsl:apply-templates>** instruction, the script jumps to a template that matches the **<configuration>** tag. This template, **<xsl:template match="configuration">**, is part of the commit script boilerplate that you must include in all of your commit scripts:

```
<xsl:template match="configuration">
    <!-- ... insert your code here ... -->
</xsl:template>
```

Thus, the import file **junos.xsl** contains a template that points to a template explicitly referenced in your script.

## Usage Examples

The following example contains the **<xsl:if>** programming instruction and the **<xnm:warning>** element. The logical result of both templates is:

```
<commit-script-results>    <!-- from template in junos.xsl import file
-->
    <xsl:if test="not(system/host-name)"> <!-- from "configuration" template
-->
        <xnm:warning xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm">
            <edit-path>[edit system]</edit-path>
            <statement>host-name</statement>
            <message>Missing a hostname for this device.</message>
        </xnm:warning>
    </xsl:if>    <!-- end of "configuration" template -->
</commit-script-results>    <!-- end of template in junos.xsl import file -->
```

When you import the **junos.xsl** file and explicitly include the **<xsl:template match="configuration">** tag in your commit script, the context (**dot**) moves to the **<configuration>** node. This allows you to write all XPath expressions relative to that point. This technique allows you to simplify the XPath expressions you use in your commit scripts. For example, instead of writing this, which matches the device with hostname **atlanta**:

```
<xsl:if test="starts-with(commit-script-input/configuration/system/host-name,
'atlanta')">
```

You can write this:

```
<xsl:if test="starts-with(system/host-name, 'atlanta')">
```

**Related Documentation**

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)
- [apply-templates on page 89](#)
- [xsl:apply-templates on page 40](#)
- [xsl:template on page 48](#)

---

## xsl:text

<b>Syntax</b>	<pre>&lt;xsl:text&gt; ... &lt;/xsl:text&gt;</pre>
<b>Description</b>	Insert literal text in the output.
<b>Usage Examples</b>	See “Example: Requiring and Restricting Configuration Statements” on page 422, “Example: Imposing a Minimum MTU Setting” on page 382, “Example: Limiting the Number of E1 Interfaces” on page 388, “Example: Controlling IS-IS and MPLS Interfaces” on page 366, and “Example: Adding a Final then accept Term to a Firewall” on page 331.

---

## xsl:value-of

<b>Syntax</b>	<pre>&lt;xsl:value-of select="expression"/&gt;</pre>
<b>Description</b>	Extract the value of an XML element and insert it into the output. The <b>select</b> attribute specifies the XPath expression that is evaluated. In the XPath expression, use <b>@</b> to access attributes of elements. Use <b>.</b> to access the contents of the element itself. If the result is a node set, the <b>&lt;xsl:value-of&gt;</b> instruction adds the string value of the first node in that node set; none of the structure of the node is preserved. To preserve the structure of the node, you must use the <b>&lt;xsl:copy-of&gt;</b> instruction instead.
<b>Attributes</b>	<b>select</b> —XPath expression specifying the node or attribute to evaluate.
<b>Usage Examples</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses on page 343</a></li><li>• <a href="#">Example: Configuring Administrative Groups for LSPs on page 350</a></li><li>• <a href="#">Example: Controlling IS-IS and MPLS Interfaces on page 366</a></li><li>• <a href="#">Example: Imposing a Minimum MTU Setting on page 382</a></li><li>• <a href="#">Example: Limiting the Number of E1 Interfaces on page 388</a></li></ul>

- Related Documentation
- [xsl:copy-of on page 42](#)

---

## xsl:variable

---

Syntax	<pre>&lt;xsl:variable name="<i>qualified-name</i>" select="<i>expression</i>"&gt;     ... &lt;/xsl:variable&gt;</pre>
Description	Declare a local or global variable. If the <b>&lt;xsl:variable&gt;</b> instruction appears at the top level of the style sheet as a child of the <b>&lt;xsl:stylesheet&gt;</b> document element, it is a global variable with a scope that includes the entire style sheet. Otherwise, it is a local variable with a scope of its following siblings and their descendants.
Attributes	<p><b>name</b>—Specifies the name of the variable. After declaration, the variable can be referred to within XPath expressions using this name, prefixed with the <b>\$</b> character.</p> <p><b>select</b>—(Optional) Determines the value of the variable. The value of the variable is determined either by the <b>select</b> attribute or by the contents of the <b>&lt;xsl:variable&gt;</b> element. Do not specify both a <b>select</b> attribute and some content; we recommend using the <b>select</b> attribute so as not to create a result tree fragment.</p>
Usage Examples	See “ <a href="#">Example: Limiting the Number of E1 Interfaces</a> ” on page 388, “ <a href="#">Example: Limiting the Number of ATM Virtual Circuits</a> ” on page 384, “ <a href="#">Example: Configuring Administrative Groups for LSPs</a> ” on page 350, and “ <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses</a> ” on page 343.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">XSLT Variables Overview on page 27</a></li><li>• <a href="#">xsl:apply-templates on page 40</a></li><li>• <a href="#">xsl:call-template on page 40</a></li><li>• <a href="#">xsl:param on page 46</a></li><li>• <a href="#">xsl:template on page 48</a></li><li>• <a href="#">xsl:with-param on page 53</a></li></ul>

## xsl:when

<b>Syntax</b>	<pre>&lt;xsl:when test="boolean-expression"&gt;   ... &lt;/xsl:when&gt;</pre>
<b>Description</b>	<p>Within an <b>&lt;xsl:choose&gt;</b> instruction, specify a set of processing instructions that are executed when the expression specified in the <b>test</b> attribute evaluates to <b>TRUE</b>. The XSLT processor processes only the instructions contained in the first <b>&lt;xsl:when&gt;</b> element whose <b>test</b> attribute evaluates to <b>TRUE</b>. If none of the <b>&lt;xsl:when&gt;</b> elements' <b>test</b> attributes evaluate to <b>TRUE</b>, the content of the <b>&lt;xsl:otherwise&gt;</b> element, if there is one, is processed.</p>
<b>Attributes</b>	<p><b>test</b>—Specifies a Boolean expression.</p>
<b>Usage Examples</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses on page 343</a></li> <li>• <a href="#">Example: Configuring Dual Routing Engines on page 357</a></li> <li>• <a href="#">Example: Preventing Import of the Full Routing Table on page 416</a></li> </ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">XSLT Programming Instructions Overview on page 28</a></li> <li>• <a href="#">xsl:choose on page 41</a></li> <li>• <a href="#">xsl:for-each on page 43</a></li> <li>• <a href="#">xsl:if on page 44</a></li> <li>• <a href="#">xsl:otherwise on page 45</a></li> </ul>

## xsl:with-param

<b>Syntax</b>	<pre>&lt;xsl:with-param name="qualified-name" select="expression"&gt;   ... &lt;/xsl:with-param&gt;</pre>
<b>Description</b>	<p>Specify a parameter to pass into a template. This element can be used when applying templates with the <b>&lt;xsl:apply-templates&gt;</b> instruction or when calling templates with the <b>&lt;xsl:call-template&gt;</b> instruction.</p>
<b>Attributes</b>	<p><b>name</b>—Specifies the name of the parameter.</p> <p><b>select</b>—(Optional) XPath expression specifying the value of the parameter. The value of the parameter is determined either by the <b>select</b> attribute or by the contents of the <b>&lt;xsl:with-param&gt;</b> element. Do not specify both a <b>select</b> attribute and content. We recommend using the <b>select</b> attribute to set the parameter so as to prevent the parameter from being passed a result tree fragment as its value.</p>
<b>Usage Examples</b>	<p>See “<a href="#">Example: Configuring Dual Routing Engines</a>” on page 357, “<a href="#">Example: Preventing Import of the Full Routing Table</a>” on page 416, and “<a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses</a>” on page 343.</p>

- Related Documentation**
- [XSLT Templates Overview on page 22](#)
  - [xsl:apply-templates on page 40](#)
  - [xsl:call-template on page 40](#)
  - [xsl:param on page 46](#)
  - [xsl:template on page 48](#)
  - [xsl:variable on page 52](#)

## PART 3

# Automation Scripting Using SLAX

- [SLAX Overview on page 57](#)
- [SLAX Statements on page 87](#)
- [Using the libslax Distribution for Automation Scripting on page 135](#)





## CHAPTER 6

# SLAX Overview

- [SLAX Overview on page 57](#)
- [SLAX Syntax Rules Overview on page 59](#)
- [SLAX Elements and Element Attributes Overview on page 61](#)
- [SLAX Elements as Function Arguments on page 62](#)
- [Understanding SLAX Default Namespaces on page 63](#)
- [SLAX Errors on page 65](#)
- [XPath Expressions Overview for SLAX on page 65](#)
- [SLAX Templates Overview on page 66](#)
- [SLAX Functions Overview on page 70](#)
- [SLAX Parameters Overview on page 72](#)
- [SLAX Variables Overview on page 76](#)
- [SLAX Statements Overview on page 79](#)
- [XSLT Elements Without SLAX Equivalents on page 82](#)
- [SLAX Operators on page 83](#)

## SLAX Overview

---

Stylesheet Language Alternative Syntax (SLAX) is a language for writing Junos OS commit scripts, op scripts, event scripts, and SNMP scripts. It is an alternative to Extensible Stylesheet Language Transformations (XSLT). SLAX has a distinct syntax similar to that of C and Perl, but the same semantics as XSLT.

- [SLAX Advantages on page 57](#)
- [How SLAX Works on page 58](#)

## SLAX Advantages

XSLT is a powerful and effective tool for handling Extensible Markup Language (XML) that works well for machine-to-machine communication, but its XML-based syntax is inconvenient for the development of complex programs.

SLAX has a simple syntax that follows the style of C and PERL. It provides a practical and succinct way to code, thus enabling you to create readable, maintainable commit,

op, event, and SNMP scripts. SLAX removes XPath expressions and programming instructions from XML elements. XML angle brackets and quotation marks are replaced by parentheses and curly brackets ( { } ), which are the familiar delimiters of C and PERL.

The benefits of SLAX are particularly strong for programmers who are not already accustomed to XSLT, because SLAX enables them to concentrate on the new programming topics introduced by XSLT, rather than concentrating on learning a new syntax. For example, SLAX enables you to:

- Use **if**, **else if**, and **else** statements instead of `<xsl:choose>` and `<xsl:if>` elements
- Put test expressions in parentheses ( )
- Use the double equal sign (==) to test equality instead of the single equal sign (=)
- Use curly braces to show containment instead of closing tags
- Perform concatenation using the underscore ( \_ ) operator, as in PERL, version 6
- Write text strings using simple quotation marks ( " " ) instead of the `<xsl:text>` element
- Define named templates with a syntax resembling a function definition
- Invoke named templates with a syntax resembling a function call
- Simplify namespace declarations
- Reduce the clutter in your scripts
- Write more readable scripts


## How SLAX Works

SLAX functions as a preprocessor for XSLT. Junos OS internally translates SLAX programming instructions (such as **if** and **else** statements) into the equivalent XSLT instructions (such as `<xsl:if>` and `<xsl:choose>` elements). After this translation, the XSLT transformation engine—which, for Junos OS, is the Junos OS management (mgd) process—is invoked.

SLAX does not affect the expressiveness of XSLT; it only makes XSLT easier to use. The underlying SLAX constructs are completely native to XSLT. SLAX adds nothing to the XSLT engine. The SLAX parser parses an input document and builds an XML tree identical to the one produced when the XML parser reads an XSLT document.

Figure 3 on page 58 shows the flow of SLAX script input and output.

**Figure 3: SLAX Script Input and Output**



- Related Documentation**
- [Converting Scripts Between SLAX and XSLT on page 711](#)
  - [SLAX Elements and Element Attributes Overview on page 61](#)
  - [SLAX Statements Overview on page 79](#)
  - [SLAX Syntax Rules Overview on page 59](#)
  - [SLAX Templates Overview on page 66](#)
  - [SLAX Variables Overview on page 76](#)
  - [XPath Expressions Overview for SLAX on page 65](#)
  - [XSLT Overview on page 17](#)

## SLAX Syntax Rules Overview

SLAX syntax rules are similar to those of traditional programming languages like C and PERL. The following sections discuss general aspects of SLAX syntax rules:

- [Code Blocks on page 59](#)
- [Comments on page 59](#)
- [Line Termination on page 60](#)
- [Strings on page 60](#)

### Code Blocks

SLAX delimits blocks of code with curly braces. Code blocks, which may define the boundaries of an element, a hierarchy, or a segment of code, can be at the same level as or nested within other code blocks. Declarations defined within a particular code block have a scope that is limited to that block.

The following example shows two blocks of code. Curly braces define the bounds of the **match** / block. The second block, containing the **<op-script-results>** element, is nested within the first.

```
match / {
  <op-script-results> {
    <output> "Script summary:";
  }
}
```

### Comments

In SLAX, you can add comments anywhere in a script. Commenting a script increases readability for all users, including the author, who may need to return to a script long after it was originally written. It is recommended that you add comments throughout a script as you write it.

In SLAX, you insert comments in the traditional C style, beginning with **/\*** and ending with **\*/**. For example:

```
/* This is a comment. */
```

Multi-line comments follow the same format. In the following example, the additional `/*` characters are added to the beginning of the lines for readability, but they are not required.

```
/* Script Title
 * Author: Jane Doe
 * Last modified: 01/01/10
 * Summary of modifications: ...
 */
```

The XSLT equivalent is:

```
<!-- Script Title
Author: Jane Doe
Last modified: 01/01/10
Summary of modifications: ...
-->
```

The following example inserts a comment into the script to remind the programmer that the output is sent to the console.

```
match / {
  <op-script-results> {
    /* Output script summary to the console */
    <output> "Script summary: ...";
  }
}
```

## Line Termination

As with many traditional programming languages, SLAX statements are terminated with a semicolon.

In the following example, the namespace declarations, import statement, and output element are all terminated with a semicolon. Lines that begin or end a block are not terminated with a semicolon.

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match / {
  <op-script-results> {
    <output> "Script summary:";
    /* ... */
  }
}
```

## Strings

Strings are sequences of text characters. SLAX strings can be enclosed in either single quotes or double quotes. However, you must close the string with the same type of quote used to open the string. Strings can be concatenated together using the SLAX concatenation operation, which is the underscore (`_`).

For example:

```
match / {
```

```

    <op-script-results> {
      /* Output script summary to the console */
      <output> "Script" _ "summary: ...";
    }

```

#### Related Documentation

- [SLAX Elements and Element Attributes Overview on page 61](#)
- [SLAX Overview on page 57](#)
- [SLAX Statements Overview on page 79](#)
- [SLAX Templates Overview on page 66](#)
- [SLAX Variables Overview on page 76](#)

## SLAX Elements and Element Attributes Overview

### SLAX Elements

SLAX elements are written with only the opening tag. The contents of the tag appear immediately following the opening tag. The contents can be either a simple expression or a more complex expression placed inside curly braces. For example:

```

<top> {
  <one>;
  <two> {
    <three>;
    <four>;
    <five> {
      <six>;
    }
  }
}

```

The XSLT equivalent is:

```

<top>
  <one/>
  <two>
    <three/>
    <four/>
    <five>
      <six/>
    </five>
  </two>
</top>

```

Using these nesting techniques and removing the closing tag reduces clutter and increases code clarity.

### SLAX Element Attributes

SLAX element attributes follow the style of XML. Attributes are included in the opening tag and consist of an attribute name and value pair. The attribute syntax consists of the attribute name followed by an equals sign and then the attribute value enclosed in quotation marks. Multiple attributes are separated by spaces.

```

<element attr1="one" attr2="two">;

```

Where XSLT allows attribute value templates using curly braces, SLAX uses the normal expression syntax. Attribute values can include any XPath syntax, including quoted strings, parameters, variables, numbers, and the SLAX concatenation operator, which is an underscore (\_). In the following example, the SLAX element **location** has two attributes, **state** and **zip**:

```
<location state=$location/state zip=$location/zip5 _ "-" _ $location/zip4>;
```

The XSLT equivalent is:

```
<location state="{ $location/state }"
  zip="{concat($location/zip5, '-', $location/zip4) }"/>
```

In SLAX, curly braces placed inside quote strings are not interpreted as attribute value templates. Instead, they are interpreted as plain-text curly braces.

An escape sequence causes a character to be treated as plain text and not as a special operator. For example, in HTML, an ampersand (&) followed by **lt** causes the less-than symbol (<) to be printed.

In XSLT, the double curly braces ({ and }) are escape sequences that cause opening and closing curly braces to be treated as plain text. When a SLAX script is converted to XSLT, the curly braces inside quote strings are converted to double curly braces:

```
<avt sign="{here} ">;
```

The XSLT equivalent is:

```
<avt sign="{{here} }"/>
```

**Related Documentation** • [XML Overview on page 7](#)

---

## SLAX Elements as Function Arguments

Starting with version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, you can use SLAX elements directly as arguments for both functions and templates. Function arguments can either be a single element or a block of SLAX code, placed inside braces as shown in the following example:

```
var $a = my:function(<elt>, <max> 15);

var $b = my:test ({
  <min> 5;
  <max> 15;
  if ($step) {
    <step> $step;
  }
});

var $c = my:write(<content> {
  <document> "total.txt";
  <size> $file/size;
  if (node[@type == "full"]) {
    <full>;
  }
});
```

For templates, you still need to include the argument name, but the value can be inline. For example:

```
call my:valid($name = <name> $input, $object = {
  <min> $min;
  <max> $max;
});
```

## The Main Template

Starting with version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, the **main** template is introduced in SLAX. The **main** template enables you to process the input XML document and provide the top-level element for the output hierarchy. The **main** template is equivalent to using **match /** but slightly more convenient.

You can use the **main** statement to match the top of the input data hierarchy and create the top-level tag of the output hierarchy. You can use the statement in two forms: with or without the output tag. When you omit the output element, **main** is just followed by a block of statements within a set of braces, as shown in the following example:

```
main {
  <top> {
    <answer> 42;
  }
}
```

The **main** template can also be used with a top-level output element following the **main** token, as shown in the following example:

```
main <top> {
  <answer> 42;
}
```

Both of the preceding examples are equivalent to the following XSLT version:

```
<xsl:template match="/">
  <top>
    <answer>42</answer>
  </top>
</xsl:template>
```

### Related Documentation

- [SLAX Elements and Element Attributes Overview on page 61](#)

## Understanding SLAX Default Namespaces

Version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, introduces default namespaces for prefixes. [Table 6 on page 64](#) lists the prefixes that have default namespaces installed with the libslax software distribution.

When a prefix is used without the corresponding **ns** statement in scope, SLAX refers to the set of default namespaces. If the prefix has a default namespace, that namespace is automatically mapped to the prefix.

Table 6: Prefix Set with libslax

Prefix	Source	Default URI
bit	libslax	xml.libslax.org/bit
curl	libslax	xml.libslax.org/curl
exsl	exslt	http://exslt.org/common
crypto	exslt	http://exslt.org/crypto
math	exslt	http://exslt.org/math
set	exslt	http://exslt.org/sets
func	exslt	http://exslt.org/functions
str	exslt	http://exslt.org/strings
date	exslt	http://exslt.org/dates-and-times
dyn	exslt	http://exslt.org/dynamic
saxon	libxslt	http://icl.com/saxon
os	libslax	http://xml.libslax.org/os
xutil	libslax	http://xml.libslax.org/xutil

When using the **slaxproc** tool with the **--format** or **--slax-to-xslt** command-line options, the namespace is properly displayed, as shown in the following example:

```
% cat /tmp/foo.slax
version 1.1;
match / {
  <top> {
    expr date:time();
  }
}

% slaxproc --format /tmp/foo.slax
version 1.1;
ns date extension = "http://exslt.org/dates-and-times";
match / {
  <top> date:time();
}
```

**Related Documentation**

- [SLAX Syntax Rules Overview on page 59](#)



## SLAX Errors

Starting with version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, you can use the **<errors>** element to control the handling of HTML and XML parsing errors. The default behavior is to display errors to the standard error output stream (stderr). By supplying different values for the **<errors>** element, you can also choose to ignore, log, or record the error. [Table 7 on page 65](#) lists the values that you can use for the **errors** element.

**Table 7: Error Elements**

Value	Special Behavior
default	Errors are displayed on stderr
ignore	Errors are discarded
log	Errors are logged (via <code>slaxLog()</code> )
record	Errors are recorded

When the value of **<errors>** is set to "record", all errors appear in a string under the **<errors>** element in the XML node (as returned, for example by **curl:perform**). If no errors are generated, the **<errors>** element is not present, which can be used as a test for errors.

```
var $opt = {
  <url> $url;
  <format> "html";
  <errors> "record";
}
var $res = curl:single($opts);
if ($res/errors) {
  terminate "failure: " _ $res/errors;
}
```

**Related Documentation** • [SLAX Elements and Element Attributes Overview on page 61](#)

## XPath Expressions Overview for SLAX

XPath expressions can appear either as the contents of an XML element or as the contents of an **expr** (expression) statement. In either case, the value is translated to either an **<xsl:text>** element, which outputs literal text, or to an **<xsl:value-of>** element, which extracts data from an XML structure.

You encode strings using quotation marks (single or double). The concatenation operator is the underscore (`_`), as in PERL 6.

In this example, the contents of the `<three>` and `<four>` elements are identical, and the content of the `<five>` element differs only in the use of the XPath `concat()` function. The resulting output is the same in all three cases.

```
<top> {
  <one> "test";
  <two> "The answer is " _ results/answer _ ".";
  <three> results/count _ " attempts made by " _ results/user;
  <four> {
    expr results/count _ " attempts made by " _ results/user;
  }
  <five> {
    expr results/count;
    expr " attempts made by ";
    expr results/user;
  }
  <six> results/message;
}
```

The XSLT equivalent is:

```
<top>
  <one><xsl:text>test</xsl:text></one>
  <two>
    <xsl:value-of select='concat("The answer is ", results/answer, ".")' />
  </two>
  <three>
    <xsl:value-of select='concat(results/count, " attempts made by ",
results/user)' />
  </three>
  <four>
    <xsl:value-of select='concat(results/count, " attempts made by ",
results/user)' />
  </four>
  <five>
    <xsl:value-of select="results/count" />
    <xsl:text> attempts made by </xsl:text>
    <xsl:value-of select="results/user" />
  </five>
  <six><xsl:value-of select='results/message' /></six>
</top>
```

#### Related Documentation

- [concat\(\) on page 33](#)
- [SLAX Elements and Element Attributes Overview on page 61](#)
- [SLAX Syntax Rules Overview on page 59](#)
- [XPath Overview on page 19](#)
- [xsl:text on page 51](#)
- [xsl:value-of on page 51](#)

---

## SLAX Templates Overview

A SLAX script consists of one or more sets of rules called *templates*. Each template is a segment of code that contains rules to apply when a specified node is matched.



**NOTE:** Version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, supports SLAX elements as arguments to both templates and functions.

There are two types of templates, named and unnamed (or match), described in the following sections.

- [Unnamed \(Match\) Templates on page 67](#)
- [Named Templates on page 68](#)

## Unnamed (Match) Templates

Unnamed templates, also known as match templates, contain a **match** statement with an XPath expression to specify the criteria for nodes upon which the template should be invoked. In the following commit script sample, the template matches the top-level element in the configuration hierarchy:

```
match configuration {
    /* ...body of the template goes here */
}
```

By default, the processor recursively traverses the entire document hierarchy, inspecting each node and looking for a template that matches the current node. When a matching template is found, the contents of that template are evaluated.

The **apply-templates** statement can be used inside an unnamed template to limit and control the default, hierarchical traversal of nodes. This statement accepts an optional XPath expression, which is equivalent to the **select** attribute in an **<xsl:apply-templates>** element. If an optional XPath expression is included, only nodes matching the XPath expression are traversed. Otherwise, all children of the context node are traversed. If the XPath expression is included but does not match any nodes, nothing is traversed and nothing happens.

In the following example, the template rule matches the **<route>** element in the XML hierarchy. All the nodes containing a **changed** attribute are processed. All **route** elements containing a **changed** attribute are replaced with a **new** element.

```
match route {
    <new> {
        apply-templates *[@changed];
    }
}
```

The XSLT equivalent:

```
<xsl:template match="route">
    <new>
        <xsl:apply-templates select="*[@changed]" />
    </new>
</xsl:template>
```

Using unnamed templates allows the script to ignore the location of a tag in the XML hierarchy. For example, if you want to convert all **<author>** tags into **<div class="author">**

tags, using templates enables you to write a single rule that converts all **<author>** tags, regardless of their location in the input XML document.

## Named Templates

Named templates operate like functions in traditional programming languages. When the complexity of a script increases or a code segment appears in multiple places, you can modularize the code and create named templates. Like functions, named templates accept arguments and run only when explicitly called.

In SLAX, the named template definition consists of the **template** keyword, the template name, a set of parameters, and a braces-delimited block of code. Parameter declarations can be inline and consist of the parameter name, and, optionally, a default value. Alternatively, you can declare parameters inside the template block using the **param** statement. If a default value is not defined, the parameter defaults to an empty string.

The following example creates a template named **my-template** and defines three parameters, one of which defaults to the string **false**, and one of which defaults to the contents of the element node named **name** that is a child of the current context node. If the script calls the template and does not pass in a parameter, the default value is used.

```
template my-template ($a, $b = "false", $c = name) {  
    /* ... body of the template ... */  
}
```

An alternate method is to declare the parameters within the template using the **param** statement. The following code is identical to the previous example:

```
template my-template {  
    param $a;  
    param $b = "false";  
    param $c = name;  
    /* ... body of the template ... */  
}
```

In SLAX, you invoke named templates using the **call** statement, which consists of the **call** keyword and template name, followed by a set of parameter bindings. These bindings are a comma-separated list of parameter names that are passed into the template from the calling environment. Parameter assignments are made by name and not by position in the list. Alternatively, you can declare parameters inside the **call** block using the **with** statement. Parameters passed into a template must match a parameter defined in the actual template; otherwise the parameter is ignored. Optionally, you can set a value for each parameter. If you do not define a value for the parameter in the calling environment, the script passes in the current value of the parameter if it was previously initialized, or it generates an error if the parameter was never declared. For more information about passing parameters, see [“SLAX Parameters Overview” on page 72](#).

In the following example, the template **my-template** is called with the parameter **c** containing the contents of the element node named **other-name** that is a child of the current context node:

```
call my-template {  
    with $c = other-name;  
}
```

In the following example, the **name-servers-template** declares two parameters: **name-servers** and **size**. The **size** parameter is given a default value of zero. The match template, which declares and initializes **name-servers**, calls the **name-servers-template** three times.

The first call to the template does not include any parameters. Thus **name-servers** will default to an empty string, and **size** will default to a value of zero as defined in the template. The second call includes the **name-servers** and **size** parameters, but only supplies a value for the **size** parameter. Thus **name-servers** has the value defined by its initialization in the script, and **size** is equal to the number of **name-servers** elements in the configuration hierarchy. The last call is identical to the second call, but it supplies the parameters using the **with** statement syntax.

```
match configuration {
  param $name-servers = name-servers/name;
  call name-servers-template();
  call name-servers-template($name-servers, $size = count($name-servers));
  call name-servers-template() {
    with $name-servers;
    with $size = count($name-servers);
  }
}

template name-servers-template($name-servers, $size = 0) {
  <output> "template called with size " _ $size;
}
```

The XSLT equivalent is:

```
<xsl:template match="configuration">
  <xsl:variable name="name-servers" select="name-servers/name"/>
  <xsl:call-template name="name-servers-template"/>
  <xsl:call-template name="name-servers-template">
    <xsl:with-param name="name-servers" select="$name-servers"/>
    <xsl:with-param name="size" select="count($name-servers)"/>
  </xsl:call-template>
  <xsl:call-template name="name-servers-template">
    <xsl:with-param name="name-servers" select="$name-servers"/>
    <xsl:with-param name="size" select="count($name-servers)"/>
  </xsl:call-template>
</xsl:template>

<xsl:template name="name-servers-template">
  <xsl:param name="name-servers"/>
  <xsl:param name="size" select="0"/>
  <output>
    <xsl:value-of select="concat('template called with size ', $size)"/>
  </output>
</xsl:template>
```

- Related Documentation**
- [SLAX Parameters Overview on page 72](#)
  - [XSLT Templates Overview on page 22](#)
  - [apply-templates on page 89](#)
  - [call on page 93](#)
  - [match on page 108](#)
  - [param on page 118](#)

- [with on page 131](#)

## SLAX Functions Overview

---

Version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases, supports functions. When the complexity of a script increases or a code segment appears in multiple places, you can modularize the code and create functions. Functions accept arguments and run only when explicitly called.



**NOTE:** Version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, supports SLAX elements as arguments to both templates and functions.

Functions have several advantages over templates, including the following:

- Arguments are passed by position rather than name.
- Return values can be objects as opposed to result tree fragments.
- Functions can be used in expressions.
- Functions can be resolved dynamically (using EXSLT `dyn:evaluate()`).

In SLAX, you define a function definition as a top-level statement in the script. The function definition consists of the **function** keyword, the function name, a set of arguments, and a braces-delimited block of code. The function name must be a qualified name. The argument list is a comma-separated list of parameter names, which are positionally assigned based on the function call. Trailing arguments can have default values. Alternatively, you can define function parameters inside the function block using the **param** statement. The syntax is:

```
function function-name (argument-list) {  
    ...  
    result return-value;  
}  
  
function function-name () {  
    param param-name1;  
    param param-name2;  
    param param-name3 = default-value;  
    ...  
    result return-value;  
}
```

The return value can be a scalar value, an XML element or XPath expression, or a set of instructions that emit the value to be returned.

If there are fewer arguments in the function invocation than in the definition, the default values are used for any trailing arguments. If there are more arguments in the function invocation than in the definition, the function call generates an error.

The following example defines the function **size**, which has three parameters: **width**, **height**, and **scale**. The default value for **scale** is 1. If the function call argument list does not include the **scale** argument, the calculation uses the default value of 1 for that argument. The function's return value is the product of the **width**, **height**, and **scale** variables enclosed in a **<size>** element.

In the main match template, the function call uses width and height data selected from each **graphic/dimension** element in the source XML file. The script evaluates the function, and the **copy-of** statement emits the return value to the result tree as the contents of the **<out>** element.

```
version 1.1;
ns my = "http://www.example.com/myfunctions";

function my:size ($width, $height, $scale = 1) {
  result <size> {
    expr $width * $height * $scale;
  }
}

match / {
  for-each (graphic/dimension) {
    <out> {
      copy-of my:size((width/.), (height/.));
    }
  }
}
```

The following function definition uses **param** statements to define the parameters instead of a comma-separated list. The behavior of the function is identical to that in the previous example.

```
version 1.1;
ns my = "http://www.example.com/myfunctions";

function my:size () {
  param $width;
  param $height;
  param $scale = 1;
  result <size> {
    expr $width * $height * $scale;
  }
}

match / {
  for-each (graphic/dimension) {
    <out> {
      copy-of my:size((width/.), (height/.));
    }
  }
}
```

- Related Documentation
- [function on page 103](#)
  - [param on page 118](#)
  - [result on page 121](#)

## SLAX Parameters Overview

---

Parameters may be passed to named or unnamed templates or to functions. After declaring a parameter, you can reference it by prefixing the parameter name with the dollar sign (\$).

- [Declaring Parameters on page 72](#)
- [Passing Parameters to Templates on page 73](#)
- [Example: Parameters and Match Templates on page 74](#)
- [Passing Parameters to Functions on page 75](#)

### Declaring Parameters

In SLAX scripts, you declare parameters using the **param** statement. Optionally, you can define an initial value for each parameter in the declaration. For example:

```
param $dot = .;
```

The scope of a parameter can be local or global. Local parameters must be declared at the beginning of a block, and their scope is limited to the block in which they are declared. A parameter whose value is set by Junos OS at script initialization must be defined as a global parameter. Global parameter declarations are placed just after the style sheet declarations. A script can assign a default value to the global parameter, which is used in the event that Junos OS does not give a value to the parameter.

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

ns ext = "http://xmlsoft.org/XSLT/namespace";

/* global parameter */
param $interface1 = "fxp0";
```

In a template, you declare parameters either in a parameter list or by using the **param** statement in the template block. Optionally, declare default values for each template parameter. If a template is invoked without the parameter, the default expression is evaluated, and the results are assigned to the parameter. If you do not define a default value in the template, the parameter defaults to an empty string.

The following named template **print-host-name** declares the parameter **message** and defines a default value:

```
template print-host-name ($message = "host name: " _ system/host-name) {
  <xnm:warning> {
    <message> $message;
  }
}
```

An alternative, but equivalent, declaration is:

```
template print-host-name () {
  param $message = "host name: " _ system/host-name;
  <xnm:warning> {
    <message> $message;
  }
}
```



```
    }
}
```

The template declares **message** and accesses its value by prefixing the parameter name with the dollar sign (\$). In XSLT, the parameter name is prefixed by the dollar sign when you access it but not when you declare it.

In a function, you declare parameters either in a parameter list or by using the **param** statement in the function block. Optionally, you can declare default values for trailing parameters. If you invoke a function without that trailing parameter, the default expression is evaluated, and the results are assigned to the parameter. If you do not define a default value, the parameter defaults to an empty string.

The following example defines a function named **size**, which has three parameters: **width**, **height**, and **scale**. The default value for **scale** is 1. If the function call argument list does not include the **scale** argument, the calculation uses the default value of 1 for that argument. The return value for the function is the product of the **width**, **height**, and **scale** variables enclosed in a **<size>** element.

```
function my:size ($width, $height, $scale = 1) {
    result <size> {
        expr $width * $height * $scale;
    }
}
```

An alternative, but equivalent declaration, which uses the **param** statement, is:

```
function my:size () {
    param $width;
    param $height;
    param $scale = 1;
    result <size> {
        expr $width * $height * $scale;
    }
}
```

## Passing Parameters to Templates

When you invoke a template, you pass arguments into the template either in an argument list or by using the **with** statement. The name of the parameter supplied in the calling environment must match the name of a parameter defined in the actual template. Otherwise, the parameter is ignored. Optionally, for each parameter you pass to a template, you can define a value using an equal sign (=) and a value expression. In the following example, the two calls to the named template **print-host-name** are identical:

```
match configuration {
    call print-host-name($message = "passing in host name: " _
system/host-name);
}
match configuration {
    call print-host-name() {
        with $message = "passing in host name: " _ system/host-name;
    }
}
```

The parameter value that gets used in a template depends on how the template is called. The following three examples, which call the **print-host-name** template, illustrate the possible calling environments.

If you call a template but do not include a specific parameter, the default expression defined in the template is evaluated, and the results are assigned to the parameter. If there is no default value for that parameter in the template, the parameter defaults to an empty string. The following example calls the named template **print-host-name** but does not include any parameters in the call. In this case, the named template will use the default value for the **message** parameter that was defined in the **print-host-name** template, or an empty string if no default exists.

```
match configuration {
    call print-host-name();
}
```

If you call a template and include a parameter, but do not define a value for the parameter in the calling environment, the script passes in the current value of the parameter if it was previously initialized, or it generates an error if the parameter was never declared. The following example calls the named template **print-host-name** and passes in the **message** parameter but does not include a value. If **message** is declared and initialized in the script, and the scope is visible to the block, the current value of **message** is used. If **message** is declared in the script but not initialized, the value of **message** will be an empty string. If **message** has not been declared, the script produces an error.

```
match configuration {
    call print-host-name($message);
    /* If $message was initialized previously, the current value is used;
     * If $message was declared but not initialized, an empty string is
    used;
     * If $message was never declared, the call generates an error. */
}
```

If you call a template, include the parameter, and define a value for the parameter, the template uses the provided value. The following example calls the named template **print-host-name** with the **message** parameter and a defined value, so the template uses the new value:

```
match configuration {
    call print-host-name($message = "passing in host name: " _
system/host-name);
}
```

## Example: Parameters and Match Templates

The following example matches the top level **configuration** hierarchy element and then instructs the Junos OS management process (mgd) to recursively apply templates to the **system/host-name** subtree. The parameters **message** and **domain** are used in the processing of any matching nodes.

```
match configuration {
    var $domain = domain-name;
    apply-templates system/host-name {
        with $message = "Invalid host-name";
        with $domain;
    }
}
```

```

}

match host-name {
  param $message = "Error";
  param $domain;
  <hello> $message _ ":: " _ . _ " (" _ $domain _ ")";
}

```

The XSLT equivalent is:

```

<xsl:template match="configuration">
  <xsl:apply-templates select="system/host-name">
    <xsl:with-param name="message" select="'Invalid host-name'"/>
    <xsl:with-param name="domain" select="$domain"/>
  </xsl:apply-templates>
</xsl:template>

<xsl:template match="host-name">
  <xsl:param name="message" select="'Error'"/>
  <xsl:param name="domain"/>
  <hello>
    <xsl:value-of select="concat($message, ':: ', '(', $domain, ')')"/>
  </hello>
</xsl:template>

```

## Passing Parameters to Functions

Version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases, supports functions. Although you can use the **param** statement to define function parameters, you cannot use the **with** statement to pass parameter values into the function from the calling environment. When you call a function, you pass arguments into the function in a comma-separated list. Function arguments are passed to the function by position rather than by name as in a template.

A function declaration can define default values for trailing arguments. If there are fewer arguments in the function invocation than in the definition, the default values are used for any trailing arguments. If there are more arguments in the function invocation than in the definition, the function call generates an error.

In the following match template, the function call uses width and height data selected from each **graphic/dimension** element in the source XML file. The script evaluates the function, and the **copy-of** statement emits the return value to the result tree as the contents of the **<out>** element. The function call includes arguments for **width** and **height**, but not for **scale**. The default value of 1 is used for **scale** within the function block.

```

version 1.1;
ns my = "http://www.example.com/myfunctions";

function my:size () {
  param $width;
  param $height;
  param $scale = 1;
  result <size> {
    expr $width * $height * $scale;
  }
}

match / {
  for-each (graphic/dimension) {

```

```
        <out> {  
            copy-of my:size((width/.), (height/.));  
        }  
    }  
}
```

- Related Documentation**
- [SLAX Functions Overview on page 70](#)
  - [SLAX Templates Overview on page 66](#)
  - [function on page 103](#)
  - [param on page 118](#)
  - [template on page 125](#)
  - [with on page 131](#)

---

## SLAX Variables Overview

SLAX variables can store any values that you can calculate or statically define. This includes data structures, XML hierarchies, and combinations of text and parameters. For example, you could assign the XML output of an operational mode command to a variable and then access the hierarchy within the variable.

You can define both local and global variables. Variables are global if they are defined outside of any template. Otherwise, they are local. The value of a global variable is accessible anywhere in the script. The scope of a local variable is limited to the template or code block in which it is defined.

Version 1.0 of the SLAX language supports immutable variables, which are declared using the **var** statement. Version 1.1 of the SLAX language, which is supported starting with Junos OS Release 12.2, introduces mutable variables, which are declared using the **mvar** statement. Mutable and immutable variables are discussed in the following sections:

- [Immutable variables on page 76](#)
- [Mutable variables on page 77](#)

### Immutable variables

In version 1.0 of the SLAX language, you declare variables using the **var** statement. Variables declared using the **var** statement are immutable. You can set the value of an immutable variable only when you declare it, after which point the value is fixed.

In the declaration, the variable name is prefixed with the dollar sign (\$), which is unlike the XSLT declaration, where the dollar sign does not prefix the value of the **name** attribute of the **<xsl:variable>** element. Once you declare a variable, you can reference it within an XPath expression using the variable name prefixed with a dollar sign (\$). You initialize a variable by following the variable name with an equal sign (=) and an expression.

The following example declares and initializes the variable **location**, which is then used to initialize the variable **message**:

```
var $location = $dot/@location;
```

```
var $message = "We are in " _ $location _ " now.";
```

The XSLT equivalent is:

```
<xsl:variable name="location" select="$dot/@location"/>
<xsl:variable name="message" select="concat('We are in ', $location, ' now.')" />
```

Variables declared using the **var** statement are immutable. As such, you can never change the value of the variable after it is initialized in the declaration. Although you cannot directly update the value of the variable, you can mimic the effect by recursively calling a function and passing in the value of the variable as a parameter. For example:

```
var $count = 1;
match / {
  call update-count($myparam = $count);
}
template update-count($myparam) {
  expr $count _ ", " $myparam _ "\n";
  if ($myparam != 4) {
    call update-count($myparam = $myparam + 1)
  }
}
```

Executing the op script in the CLI produces the following output in the log file. Although the **count** variable must remain fixed, **myparam** is updated with each call to the template.

```
1, 1
1, 2
1, 3
1, 4
1, 5
```

## Mutable variables

Version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases, introduces mutable variables. Unlike variables declared using the **var** statement, the value of a mutable variable can be modified by a script. You can set the initial value of a mutable variable at the time you declare it or at any point in the script.

You declare mutable variables using the **mvar** statement. In the declaration, the variable name is prefixed with the dollar sign (\$). Once you declare a mutable variable, you can reference it within an XPath expression using the variable name prefixed with a dollar sign (\$). You initialize the variable by following the variable name with an equal sign (=) and an expression.

The following example declares and initializes the mutable variable **location**, which is then used to initialize the mutable variable **message**:

```
mvar $location = $dot/@location;
mvar $message = "We are in " _ $location _ " now.";
```

Mutable variables can be initialized or updated after they are declared. To initialize or update the value of a mutable variable, use the **set** statement. The following example declares the variable, **block**, and initializes it with the element **<block>**:

```
mvar $block;
set $block = <block> "start here";
```

For mutable variables that represent a node set, use the **append** statement to append a new node to the existing node set. The following example creates the mutable variable **\$mylist**, which is initialized with one **<item>** element. For each grocery item in the **\$list** variable, the script appends an **<item>** element to the **\$mylist** node set with the name and size of the item.

```
version 1.1;

var $list := {
  <list> {
    <grocery> {
      <name> "milk";
      <type> "nonfat";
      <brand> "any";
      <size> "gallon";
    }
    <grocery> {
      <name> "orange juice";
      <type> "no pulp";
      <brand> "any";
      <size> "half gallon";
    }
    <drugstore>{
      <name> "aspirin";
      <brand> "any";
      <size> "50 tablets";
    }
  }
}

match / {
  mvar $mylist;
  set $mylist = <item> {
    <name> "coffee";
    <size> "1 lb";
  }
  for $item ($list/list/grocery) {
    append $mylist += <item> {
      <name> $item/name;
      <size> $item/size;
    }
  }
  <grocery-short-list> {
    copy-of $mylist;
  }
}
```

The output from the script is:

```
<?xml version="1.0"?>
<grocery-short-list>
  <item>
    <name>coffee</name>
    <size>1 lb</size>
  </item>
  <item>
    <name>milk</name>
    <size>gallon</size>
  </item>
  <item>
```

```

        <name>orange juice</name>
        <size>half gallon</size>
    </item>
</grocery-short-list>

```

#### Related Documentation

- [XSLT Variables Overview on page 27](#)
- [SLAX Parameters Overview on page 72](#)
- [append on page 88](#)
- [mvar on page 111](#)
- [set on page 122](#)
- [var on page 129](#)

## SLAX Statements Overview

This section lists some commonly used SLAX statements, with brief examples and XSLT equivalents.

- [for-each Statement on page 79](#)
- [if, else if, and else Statements on page 80](#)
- [match Statement on page 81](#)
- [ns Statement on page 81](#)
- [version Statement on page 82](#)

### for-each Statement

The SLAX **for-each** statement functions like the `<xsl:for-each>` element. The statement consists of the **for-each** keyword, a parentheses-delimited expression, and a curly braces-delimited block. The **for-each** statement tells the processor to gather together a set of nodes and process them one by one. The nodes are selected by the specified XPath expression. Each of the nodes is then processed according to the instructions held in the **for-each** code block.

```

for-each (xpath-expression) {
    ...
}

```

Code inside the **for-each** instruction is evaluated recursively for each node that matches the XPath expression. That is, the current context is moved to each node selected by the **for-each** clause, and processing is relative to that current context.

In the following example, the **inventory** variable stores the inventory hierarchy. The **for-each** statement recursively processes each **chassis-sub-module** node that is a child of **chassis-module** that is a child of the **chassis** node. For each **chassis-sub-module** element that contains a **part-number** with a value equal to the specified part number, a **message** element is created that includes the name of the chassis module and the name and description of the chassis sub module.

```

for-each ($inventory/chassis/chassis-module/
    chassis-sub-module[part-number == '750-000610']) {
    <message> "Down rev PIC in " _ ../name _ ", " _ name _ ": " _ description;
}

```

```
}
```

The XSLT equivalent is:

```
<xsl:for-each select="$inventory/chassis/chassis-module/  
                chassis-sub-module[part-number = '750-000610']">  
  <message>  
    <xsl:value-of select="concat('Down rev PIC in ', ../name, ', ', name,  
    ': ',  
                description)"/>  
  </message>  
</xsl:for-each>
```

## if, else if, and else Statements

SLAX supports **if**, **else if**, and **else** statements. The **if** statement is a conditional construct that causes instructions to be processed if the specified XPath expression evaluates to true. The **if** construct may have one or more associated **else if** clauses, each of which tests an XPath expression. If the expression in the **if** statement evaluates to false, the processor checks each **else if** expression. If a statement evaluates to true, the script executes the instructions in the associated block and ignores all subsequent **else if** and **else** statements. The optional **else** clause is the default code that is executed in the event that all associated **if** and **else-if** expressions evaluate to false. If all of the **if** and **else if** statements evaluate to false, and the **else** statement is not present, no action is taken.

The expressions that appear in parentheses are extended XPath expressions, which support the double equal sign (==) in place of XPath's single equal sign (=).

```
if (expression) {  
  /* If block Statement */  
}  
else if (expression) {  
  /* else if block statement */  
}  
else {  
  /* else block statement */  
}
```

During script processing, an **if** statement that does not have an associated **else if** or **else** statement is transformed into an **<xsl:if>** element. If either the **else if** or **else** clauses are present, the **if** statement and associated **else if** and **else** blocks are transformed into an **<xsl:choose>** element.

```
if (starts-with(name, "fe-")) {  
  if (mtu < 1500) {  
    /* Select Fast Ethernet interfaces with low MTUs */  
  }  
}  
else {  
  if (mtu > 8096) {  
    /* Select non-Fast Ethernet interfaces with high MTUs */  
  }  
}
```

The XSLT equivalent is:

```
<xsl:choose>  
  <xsl:when test="starts-with(name, 'fe-')">  
    <xsl:if test="mtu < 1500">  
      <!-- Select Fast Ethernet interfaces with low MTUs -->  
    </xsl:if>
```



```

</xsl:when>
<xsl:otherwise>
  <xsl:if test="mtu > 8096">
    <!-- Select non-Fast Ethernet interfaces with high MTUs -->
    </xsl:if>
  </xsl:otherwise>
</xsl:choose>

```

## match Statement

You specify basic match templates using the **match** statement, followed by an expression specifying when the template should be allowed and a block of statements enclosed in a set of braces.

```

match configuration {
  <xnm:error> {
    <message> "...";
  }
}

```

The XSLT equivalent is:

```

<xsl:template match="configuration">
  <xnm:error>
    <message> ...</message>
  </xnm:error>
</xsl:template>

```

For more information about constructing match templates, see [“SLAX Templates Overview” on page 66](#).

## ns Statement

You specify namespace definitions using the SLAX **ns** statement. This consists of the **ns** keyword, a prefix string, an equal sign, and a namespace Uniform Resource Identifier (URI). To define the default namespace, use only the **ns** keyword and a namespace URI.

```

ns junos = "http://www.juniper.net/junos/";

```

The **ns** statement can appear after the **version** statement at the beginning of the style sheet or at the beginning of any block.

```

ns a = "http://example.com/1";
ns "http://example.com/global";
ns b = "http://example.com/2";
match / {
  ns c = "http://example.com/3";
  <top> {
    ns a = "http://example.com/4";
    apply-templates commit-script-input/configuration;
  }
}

```

When it appears at the beginning of the style sheet, the **ns** statement can include either the **exclude** or **extension** keyword. The keyword instructs the parser to add the namespace prefix to the **exclude-result-prefixes** or **extension-element-prefixes** attribute.

```

ns exclude foo = "http://example.com/foo";
ns extension jcs = "http://xml.juniper.net/jcs";

```

The XSLT equivalent is:

```
<xsl:stylesheet xmlns:foo="http://example.com/foo"
  xmlns:jcs="http://xml.juniper.net/jcs"
  exclude-result-prefixes="foo"
  extension-element-prefixes="jcs">
  <!-- ... -->
</xsl:stylesheet>
```

## version Statement

All SLAX style sheets must begin with a **version** statement, which specifies the version number for the SLAX language. Supported versions include 1.0 and 1.1. SLAX version 1.0 uses XML version 1.0 and XSLT version 1.1.

```
version 1.0;
```

The XSLT equivalent is:

```
<xsl:stylesheet version="1.0">
```

### Related Documentation

- [else on page 97](#)
- [for-each on page 102](#)
- [if on page 104](#)
- [match on page 108](#)
- [version on page 130](#)

---

## XSLT Elements Without SLAX Equivalents

Some XSLT elements are not directly translated into SLAX statements. Some examples of XSLT elements for which there are no SLAX equivalents in SLAX version 1.0 are **<xsl:fallback>**, **<xsl:output>**, and **<xsl:sort>**.

You can encode these elements directly as normal SLAX elements in the XSLT namespace. For example, you can include the **<xsl:output>** and **<xsl:sort>** elements in a SLAX script, as shown here:

```
<xsl:output method="xml" indent="yes" media-type="image/svg">
match * {
  for-each (configuration/interfaces/unit) {
    <xsl:sort order="ascending">
  }
}
```

When you include XSLT namespace elements in a SLAX script, do not include closing tags. For empty tags, do not include a forward slash (/) after the tag name. The examples shown in this section demonstrate the correct syntax.

The following XSLT snippet contains a combination of elements, some of which have SLAX counterparts and some of which do not:

```
<xsl:loop select="title">
  <xsl:fallback>
    <xsl:for-each select="title">
      <xsl:value-of select="."/>
    </xsl:for-each>
  </xsl:fallback>
</xsl:loop>
```

```

        </xsl:for-each>
    </xsl:fallback>
</xsl:loop>

```

The SLAX conversion uses the XSLT namespace for XSLT elements that do not have SLAX counterparts:

```

<xsl:loop select = "title"> {
  <xsl:fallback> {
    for-each (title) {
      expr .;
    }
  }
}

```

## SLAX Operators

SLAX provides a variety of operators, which add great versatility to the SLAX scripting language. [Table 8 on page 83](#) summarizes the available operators and provides an example and an explanation of each.

**Table 8: SLAX Operators**

Name	Operator	Example / Explanation
Addition	+	<pre>var \$result = 1 + 1;</pre> <p>Return the sum of the operands. This example assigns a value of 2 to the <b>\$result</b> variable.</p>
And	&&	<pre>(\$byte-count &gt; 500000) &amp;&amp; (\$byte-count &lt; 1000000)</pre> <p>Evaluate two expressions and return one boolean result. If either of the two expressions evaluates to <b>false</b>, then the combined expression evaluates to <b>false</b>.</p>
Assignment	=	<pre>var \$mtu = 1500; mvar \$mtu2 = 48; set \$mtu2 = 1500;</pre> <p>Assign a value to a variable or parameter or assign a namespace to a prefix. The example assigns a value of 1500 to both the <b>\$mtu</b> variable and the <b>\$mtu2</b> mutable variable. <b>\$mtu2</b> was originally initialized with a value of 48.</p>
Conditional	?:	<pre>var \$result = (\$a &lt; 10) ? \$b : \$c;</pre> <p>Provide conditional assignment based on the boolean value of the evaluated condition. If the conditional expression evaluates to <b>true</b>, the entire expression assumes the value of the operand to the left of the colon. If the conditional expression evaluates to <b>false</b>, the entire expression assumes the value of the operand to the right of the colon. This operator was introduced in version 1.1 of the SLAX language, which is supported starting with Junos OS Release 12.2.</p> <p>In the example, if the value stored in the variable <b>\$a</b> is less than 10, <b>\$result</b> is assigned the value stored in <b>\$b</b>. Otherwise, <b>\$result</b> is assigned the value stored in <b>\$c</b>.</p>
Division	div	<pre>&lt;output&gt;\$bit-count div 8;</pre> <p>Return the result of dividing the left operand by the right operand. If the remainder of the division is nonzero, the result is expressed in decimal floating-point notation. The example divides the <b>\$bit-count</b> variable by eight, returning the byte count (requires that <b>\$bit-count</b> has been initialized).</p>

Table 8: SLAX Operators (*continued*)

Name	Operator	Example / Explanation
Equality	==	<p>\$mtu == 1500</p> <p>Return <b>true</b> if the values of the left and right operands are equal; otherwise, the expression returns <b>false</b>. In the example, if <b>\$mtu</b> equals 1500, then the expression resolves to <b>true</b>; otherwise, it returns <b>false</b> (requires that <b>\$mtu</b> has been initialized).</p>
Greater than	>	<p>\$hop-count &gt; 0</p> <p>Return <b>true</b> if the value of the left operand is greater than the value of the right operand; otherwise, the expression returns <b>false</b>. In this example, if <b>\$hop-count</b> is greater than zero, the expression returns <b>true</b> (requires that <b>\$hop-count</b> has been initialized).</p>
Greater than or equal to	>=	<p>\$hop-count &gt;= 1</p> <p>Return <b>true</b> if the value of the left operand is either greater than or equal to the value of the right operand; otherwise, the expression returns <b>false</b>. In this example, if <b>\$hop-count</b> is 1 or greater, the expression returns <b>true</b> (requires that <b>\$hop-count</b> has been initialized).</p>
Inequality	!=	<p>\$mtu != 1500</p> <p>Return <b>true</b> if the values of the left and right operands are not equal; otherwise, the expression returns <b>false</b>. In the example, if <b>\$mtu</b> does not equal 1500, then the expression resolves to <b>true</b>; otherwise, the expression returns <b>false</b> (requires that <b>\$mtu</b> has been initialized).</p>
Iteration	...	<pre>for \$i (1 ... 10) {   &lt;player number=\$i&gt;; }</pre> <p>Iterate through a range of integer values with a start value equal to the left operand and an end value equal to the right operand. If the left operand is greater than the right, the numbers are generated in decreasing order. The operator translates into an XPath function that generates the sequence as a node set. This operator was introduced in version 1.1 of the SLAX language, which is supported starting with Junos OS Release 12.2.</p>
Less than	<	<p>\$hop-count &lt; 15</p> <p>Return <b>true</b> if the value of the left operand is less than the value of the right operand; otherwise, the expression returns <b>false</b>. In this example, if <b>\$hop-count</b> is less than 15, the expression returns <b>true</b> (requires that <b>\$hop-count</b> has been initialized).</p>
Less than or equal to	<=	<p>\$hop-count &lt;= 14</p> <p>Return <b>true</b> if the value of the left operand is either less than or equal to the value of the right operand; otherwise, the expression returns <b>false</b>. In this example, if <b>\$hop-count</b> is 14 or less, the expression returns <b>true</b> (requires that <b>\$hop-count</b> has been initialized).</p>
Modulo	mod	<p>&lt;output&gt; 10 mod 3;</p> <p>Return the division remainder of two numbers. In this example, the expression outputs a value of 1.</p>
Multiplication	*	<p>&lt;output&gt; 5 * 10;</p> <p>Return the product of the operands. In this example, the expression outputs a value of 50.</p>

Table 8: SLAX Operators (*continued*)

Name	Operator	Example / Explanation
Node Set, append to	<code>+=</code>	<pre>mvar \$block = &lt;block&gt; "start here"; append \$block += &lt;block&gt; "next block";</pre> <p>Append a value to a node set contained in a mutable variable, which is defined with the <b>mvar</b> statement. This operator was introduced in version 1.1 of the SLAX language, which is supported starting with Junos OS Release 12.2.</p>
Node Set Conversion	<code>:=</code>	<pre>var \$new-node-set := \$rtf-variable;</pre> <p>Convert a result tree fragment into a node set. A result tree fragment contains an unparsed XML data structure. It is not possible to retrieve any of the embedded XML information from this data type. A script can convert the result tree fragment into a node set and then search the node set for the appropriate information and extract it. This operator is supported in Junos OS Release 9.2 and later releases.</p>
Or	<code>  </code>	<pre>(\$mtu-size != 1500)    (\$mtu-size &gt; 2000)</pre> <p>Evaluate two expressions and return one boolean result. If either of the two expressions evaluates to true, then the combined expression evaluates to true.</p>
Parentheses	<code>()</code>	<pre>var \$result = ( \$byte-count * 8 ) + 150;</pre> <p>Create complex expressions. Parentheses function the same way as in a mathematical expression, where the expression within the parentheses is evaluated first. Parentheses can be nested; the innermost set of parentheses is evaluated first, then the next set, and so on.</p>
String concatenation	<code>_</code> (underscore)	<pre>var \$combined-string = \$host-name _ " is located at " _ \$location;</pre> <p>Concatenate multiple strings (note that strings cannot be combined using the <code>+</code> operator in SLAX). In the example, if <b>\$host-name</b> is "r1" and <b>\$location</b> is "HQ", then the value of <b>\$combined-string</b> is "r1 is located at HQ".</p>
Subtraction	<code>-</code>	<pre>var \$result = 64 - 14;</pre> <p>Return the difference between the left operand and the right operand. This example assigns a value of 50 to the <b>\$result</b> variable.</p>
Unary Minus	<code>-</code>	<pre>mvar \$number = 5. set \$number = - \$number;</pre> <p>Negate the value of the operand, changing a positive value to a negative value or a negative value to a positive value. The example negates the value stored in <b>\$number</b> and reassigns the new value of -5 to the variable.</p>
Union	<code> </code>	<pre>var \$all-interface-nodes = \$fe-interface-nodes   \$ge-interface-nodes;</pre> <p>Create a union of two node sets. All the nodes from one set combine with the nodes in the second set. This is useful when a script needs to perform a similar operation over XML nodes that are pulled from multiple sources.</p>

**Related Documentation**

- [SLAX Elements and Element Attributes Overview on page 61](#)
- [SLAX Overview on page 57](#)

- [SLAX Statements Overview on page 79](#)
- [SLAX Syntax Rules Overview on page 59](#)
- [SLAX Variables Overview on page 76](#)

## CHAPTER 7

# SLAX Statements

- [append on page 88](#)
- [apply-imports on page 89](#)
- [apply-templates on page 89](#)
- [attribute on page 90](#)
- [attribute-set on page 91](#)
- [call on page 93](#)
- [copy-node on page 94](#)
- [copy-of on page 94](#)
- [decimal-format on page 95](#)
- [element on page 96](#)
- [else on page 97](#)
- [else if on page 98](#)
- [expr on page 99](#)
- [fallback on page 100](#)
- [for on page 100](#)
- [for-each on page 102](#)
- [function on page 103](#)
- [if on page 104](#)
- [import on page 106](#)
- [key on page 106](#)
- [match on page 108](#)
- [message on page 109](#)
- [mode on page 109](#)
- [mvar on page 111](#)
- [number on page 111](#)
- [output-method on page 115](#)
- [param on page 118](#)
- [preserve-space on page 119](#)

- [priority on page 119](#)
- [processing-instruction on page 120](#)
- [result on page 121](#)
- [set on page 122](#)
- [sort on page 123](#)
- [strip-space on page 125](#)
- [template on page 125](#)
- [terminate on page 126](#)
- [trace on page 127](#)
- [uexpr on page 128](#)
- [use-attribute-sets on page 128](#)
- [var on page 129](#)
- [version on page 130](#)
- [while on page 131](#)
- [with on page 131](#)

---

## append

---

<b>Syntax</b>	<code>append <i>name</i> += <i>value</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Append a value to the node set contained in a mutable variable. The variable must be defined using the <b>mvar</b> statement.
<b>Attributes</b>	<p><b><i>name</i></b>—Name of the mutable variable.</p> <p><b><i>value</i></b>—Value to append to the node set.</p>
<b>SLAX Example</b>	<p>The following snippet appends the <code>&lt;item&gt;</code> element and <code>&lt;name&gt;</code> and <code>&lt;size&gt;</code> child elements to the node set contained in the mutable variable <b>block</b>:</p> <pre>mvar \$block; set \$block = &lt;block&gt; "item list";  for \$item (list) {   append \$block += &lt;item&gt; {     &lt;name&gt; \$item/name;     &lt;size&gt; \$item/size;   } }</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SLAX Variables Overview on page 76</a></li><li>• <a href="#">mvar on page 111</a></li><li>• <a href="#">set on page 122</a></li></ul>



## apply-imports

<b>Syntax</b>	<code>apply-imports;</code>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Apply a template rule from an external file or style sheet. By default, template rules in the main script have precedence over equivalent imported template rules. Use this statement to process the context node using the imported match template rule from the external source.</p> <p>The <b>apply-imports</b> statement mimics the <code>&lt;xsl:apply-imports&gt;</code> element, allowing the script to invoke any imported templates.</p>
<b>SLAX Example</b>	<p>In the example, the main script imports the file <b>route-rules.slax</b>. The <b>apply-imports</b> statement invokes the imported template rule for <b>&lt;route&gt;</b> elements.</p> <pre> version 1.1; import "route-rules.slax";  match route {   &lt;routes&gt; {     apply-imports;   } } </pre> <p>The imported file contains a template rule for <b>&lt;route&gt;</b> elements.</p> <pre> /* route-rules.slax */ version 1.1;  match route {   &lt;new&gt; {     apply-templates *[@changed];   } } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">import on page 106</a></li> </ul>

## apply-templates

<b>Syntax</b>	<code>apply-templates <i>expression</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Apply one or more templates, according to the value of the node-set expression. If a node-set expression is not specified, the script recursively processes all child nodes of the current node. If a node-set expression is specified, the processor only applies templates to the child elements that match the node-set expression. The <b>template</b> statement dictates which elements are transformed according to which template. The templates that are applied are passed the parameters specified by the <b>with</b> statement within the <b>apply-templates</b> statement block.</p>

<b>Attributes</b>	<b><i>expression</i></b> —(Optional) Selects the nodes to which the processor applies templates. By default, the processor applies templates to the child nodes of the current node.
<b>SLAX Example</b>	<pre>match configuration {     apply-templates system/host-name; }</pre>
<b>XSLT Equivalent</b>	<pre>&lt;xsl:template match="configuration"&gt;     &lt;xsl:apply-templates select="system/host-name"/&gt; &lt;/xsl:template&gt;</pre>
<b>Usage Examples</b>	See <a href="#">“Example: Adding a Final then accept Term to a Firewall” on page 331</a> and <a href="#">“Example: Preventing Import of the Full Routing Table” on page 416</a> .
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SLAX Templates Overview on page 66</a></li><li>• <a href="#">call on page 93</a></li><li>• <a href="#">match on page 108</a></li><li>• <a href="#">mode on page 109</a></li><li>• <a href="#">priority on page 119</a></li><li>• <a href="#">template on page 125</a></li><li>• <a href="#">with on page 131</a></li></ul>

---

## attribute

---

<b>Syntax</b>	<pre>attribute <i>attribute-name</i> {     <i>attribute-value</i>; }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Create an attribute with the given name. The attribute value is defined by a block of statements, which must be placed inside a set of braces.
<b>Attributes</b>	<p><b><i>attribute-name</i></b>—Name of the attribute, which can be an XPath expression or a string. Enclose string arguments in quotes.</p> <p><b><i>attribute-value</i></b>—A block of statements enclosed in curly braces that defines the attribute value.</p>
<b>SLAX Example</b>	<p>In the following example, the <b>&lt;book&gt;</b> element is output to the result tree with an attribute named <b>format</b>, which has the value "PDF":</p> <pre>&lt;book&gt; {     attribute "format" {         expr "PDF";     } }</pre>

In the following example, the value of the `<name>` node (rather than the literal string "name") is used to create an XML attribute with a value of "from-" concatenated with the contents of the address node. Node values are selected from the current context.

```
<source> {
  attribute name {
    expr "from-" _ address;
  }
}
```

- Related Documentation**
- [SLAX Elements and Element Attributes Overview on page 61](#)
  - [attribute-set on page 91](#)
  - [element on page 96](#)
  - [use-attribute-sets on page 128](#)

## attribute-set

**Syntax**

```
attribute-set attribute-set-name {;
  attribute attribute-name1 { attribute-value1; }
  attribute attribute-name2 { attribute-value2; }
  use-attribute-sets attribute-set-name2;
  ...
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Define a collection of attributes that can be used repeatedly. The **attribute-set** statement must be defined as a top-level statement in the script. The attribute set name is a string argument. The attribute set contents define the attributes to include in the collection. The contents can include individual **attribute** statements, which define attributes as a name and value pair, and they can include **use-attribute-sets** statements, which add the attributes from a previously defined attribute set to the current set.

To apply the attributes in an attribute set to a specific element, include the **use-attribute-sets** statement under that element and reference the attribute set name.

**Attributes** *attribute-set-name*—Name of the attribute set, which must be a string. To add the attribute set to an element, reference this name in the **use-attribute-sets** statement.

*attribute-name*—Name of the individual attribute to add to the set.

*attribute-value*—A block of statements enclosed in curly braces that defines the attribute value.

**SLAX Example** The following example creates two attribute sets: **table-attributes** and **table-attributes-ext**. The **table-attributes-ext** set includes all of the attributes that are already defined in the **table-attributes** set through use of the **use-attribute-sets** statement. In the main script body, the **table-attributes-ext** attribute set is applied to the **<table>** element. The **<table>** element includes the four attributes: **order**, **cellpadding**, **cellspacing**, and **border**.

```
version 1.1;

var $cellpadding = "0";
var $cellspacing = "10";

attribute-set table-attributes {
  attribute "order" { expr "0"; }
  attribute "cellpadding" { expr $cellpadding; }
  attribute "cellspacing" { expr $cellspacing; }
}
attribute-set table-attributes-ext {
  use-attribute-sets table-attributes;
  attribute "border" { expr "0"; }
}

match / {
  ...
  <table> {
    use-attribute-sets table-attributes-ext;
  }
}
```

**XSLT Equivalent**

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform" version="1.0">

  <xsl:variable name="cellpadding" select="0"/>
  <xsl:variable name="cellspacing" select="10"/>
  <xsl:attribute-set name="table-attributes">
    <xsl:attribute name="order">
      <xsl:text>0</xsl:text>
    </xsl:attribute>
    <xsl:attribute name="cellpadding">
      <xsl:value-of select="$cellpadding"/>
    </xsl:attribute>
    <xsl:attribute name="cellspacing">
      <xsl:value-of select="$cellspacing"/>
    </xsl:attribute>
  </xsl:attribute-set>
  <xsl:attribute-set name="table-attributes-ext"
    use-attribute-sets="table-attributes">
    <xsl:attribute name="border">
      <xsl:text>0</xsl:text>
    </xsl:attribute>
  </xsl:attribute-set>
  <xsl:template match="/">
    <table use-attribute-sets="table-attributes-ext"/>
  </xsl:template>
</xsl:stylesheet>
```

- Related Documentation**
- [SLAX Elements and Element Attributes Overview on page 61](#)
  - [attribute on page 90](#)
  - [element on page 96](#)

- [use-attribute-sets on page 128](#)

## call

**Syntax**      `call template-name (parameter-name = value) {`  
                   `/* code */`  
                   `}`

**Release Information**    Statement introduced in version 1.0 of the SLAX language.

**Description**          Call a named template. You can pass parameters into the template by including a comma-separated list of parameters, with the parameter name and an optional equal sign (=) and value expression. If a value is not specified, the current value of the parameter is passed to the template.

You can declare additional parameters inside the code block using the **with** statement.

**Attributes**          *template-name*—Specifies the name of the template to call.

**SLAX Example**

```
match configuration {
    var $name-servers = name-servers/name;
    call temp();
    call temp($name-servers, $size = count($name-servers));
    call temp() {
        with $name-servers;
        with $size = count($name-servers);
    }

    template temp($name-servers, $size = 0) {
        <output> "template called with size " _ $size;
    }
}
```

**XSLT Equivalent**

```
<xsl:template match="configuration">
    <xsl:variable name="name-servers" select="name-servers/name"/>
    <xsl:call-template name="temp"/>
    <xsl:call-template name="temp">
        <xsl:with-param name="name-servers" select="$name-servers"/>
        <xsl:with-param name="size" select="count($name-servers)"/>
    </xsl:call-template>
    <xsl:call-template name="temp">
        <xsl:with-param name="name-servers" select="$name-servers"/>
        <xsl:with-param name="size" select="count($name-servers)"/>
    </xsl:call-template>
</xsl:template>

<xsl:template name="temp">
    <xsl:param name="name-servers"/>
    <xsl:param name="size" select="0"/>
    <output>
        <xsl:value-of select="concat('template called with size ', $size)"/>
    </output>
</xsl:template>
```

**Usage Examples** See [“Example: Requiring and Restricting Configuration Statements” on page 422](#), [“Example: Imposing a Minimum MTU Setting” on page 382](#), and [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

**Related Documentation**

- [SLAX Templates Overview on page 66](#)
- [apply-templates on page 89](#)
- [match on page 108](#)
- [mode on page 109](#)
- [priority on page 119](#)
- [template on page 125](#)
- [with on page 131](#)

---

## copy-node

---

<b>Syntax</b>	<pre>copy-node;  copy-node {     /* body */ }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Copy the current node including namespace nodes to the result tree, but do not copy any attribute or child nodes. The optional body is a block of statements that emit additional nodes inside that copy.
<b>SLAX Example</b>	<pre>copy-node {     &lt;that&gt; "one"; }</pre>
<b>XSLT Equivalent</b>	<pre>&lt;xsl:copy&gt;     &lt;that&gt;         &lt;xsl:value-of select="one"/&gt;     &lt;/that&gt; &lt;/xsl:copy&gt;</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">copy-of on page 94</a></li><li>• <a href="#">xsl:copy-of on page 42</a></li></ul>

---

## copy-of

---

<b>Syntax</b>	<pre>copy-of <i>expression</i>;</pre>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	Copy the specified node including namespace nodes, child nodes, and attributes of that node. The argument is an XPath expression that specifies which nodes to copy.

Attributes	<i>expression</i> —XPath expression that specifies which nodes to copy.
SLAX Example	copy-of configuration/protocols/bgp;
XSLT Equivalent	<xsl:copy-of select="configuration/protocols/bgp"/>
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">copy-node on page 94</a></li> <li>• <a href="#">xsl:copy-of on page 42</a></li> </ul>

## decimal-format

Syntax	<pre>decimal-format <i>format-name</i> {     decimal-separator <i>character</i>;     digit <i>character</i> ;     grouping-separator <i>character</i>;     infinity <i>string</i>;     minus-sign <i>character</i>;     nan <i>string</i>;     pattern-separator <i>character</i>;     percent <i>character</i>;     per-mille <i>character</i>;     zero-digit <i>character</i>; }</pre>
Release Information	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
Description	Define formatting parameters for use by the <b>format-number()</b> XPath function. The <b>decimal-format</b> statement must be defined as a top-level statement in the script.
Attributes	<p><b>decimal-format <i>format-name</i></b>—Decimal-format identifier, which is passed as the third argument to the <b>format-number()</b> XPath function.</p> <p><b>decimal-separator <i>character</i></b>—Character used as the decimal sign. The default is the period (.).</p> <p><b>digit <i>character</i></b>—Character used to represent a digit in a pattern. The default is the number sign (#).</p> <p><b>grouping-separator <i>character</i></b>—Character used as the digit group separator or the thousands separator. The default is the comma (,).</p> <p><b>infinity <i>string</i></b>—String used to represent infinity. The default is "Infinity".</p> <p><b>minus-sign <i>character</i></b>—Character used as the minus sign. The default is the hyphen (-).</p> <p><b>nan <i>string</i></b>—String used to represent NaN. The default is "NaN".</p> <p><b>pattern-separator <i>character</i></b>—Character used to separate patterns. The first pattern is used for positive numbers, and the second pattern is used for negative numbers. The default is the semicolon (;).</p>

**percent *character***—Character used as the percent sign. The default is the percent character (%).

**per-mille *character***—Character used as a per mille sign. The default is the Unicode per mille sign (\x2030 or ‰).

**zero-digit *character***—Character used as zero. The default is the number zero (0).

**SLAX Example** The following code snippet lists the defaults for the decimal-format parameters, and uses the defined decimal format in the **format-number** XPath function:

```
version 1.1;

decimal-format default-format {
    decimal-separator "." ;
    digit "#" ;
    grouping-separator "," ;
    infinity "Infinity" ;
    minus-sign "-" ;
    nan "NaN";
    pattern-separator ";" ;
    percent "%";
    per-mille "\x2030";
    zero-digit "0" ;
}

match / {
    ...
    var $number = -14560302.5;
    expr format-number($number, "###,###.00", "default-format");
}

/* output is -14,560,302.50 */
```

**XSLT Equivalent**

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform" version="1.0">

    <xsl:decimal-format name="default-format" decimal-separator="." digit="#"
        grouping-separator="," infinity="Infinity" minus-sign="-" NaN="NaN"
        pattern-separator=";" percent="%" per-mille="\x2030" zero-digit="0"/>

    <xsl:template match="/">
        <xsl:variable name="number" select="-14560302.5"/>
        <xsl:value-of select="format-number($number, '###,###.00',
            'default-format')"/>
    </xsl:template>
</xsl:stylesheet>
```

**Related Documentation** • [output-method on page 115](#)

---

## element

<b>Syntax</b>	<pre>element <i>name</i> {     /* element contents */ }</pre>
---------------	---



<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Create an element node with the given name. The element name can be an XPath expression or a string. The element contents must be placed inside curly braces.
<b>Attributes</b>	<b>name</b> —Name of the element, which can be an XPath expression or a string. Enclose string arguments in quotes.
<b>SLAX Example</b>	The following sample code uses the value of the <b>name</b> node (rather than the literal string "name") to create an XML element, whose contents are an empty element with a name of "from-" concatenated with the value of the address node. Node values are selected from the current context.
	<pre> for-each (list/item) {   element name {     element "from-" _ address;   } } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SLAX Elements and Element Attributes Overview on page 61</a></li> </ul>

## else

<b>Syntax</b>	<pre> if (expression) {     /* code */ } else {     /* code */ } </pre>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	Include a default set of instructions that are processed if the preceding <b>if</b> and <b>else if</b> statements evaluate to <b>FALSE</b> .
<b>SLAX Example</b>	<pre> if (starts-with(name, "fe-")) {   if (mtu &lt; 1500) {     /* Select the Fast Ethernet interfaces with low MTUs */   } } else {   if (mtu &gt; 8096) {     /* Select the non-Fast Ethernet interfaces with high MTUs */   } } </pre>
<b>XSLT Equivalent</b>	<pre> &lt;xsl:choose&gt;   &lt;xsl:when select="starts-with(name, 'fe-')"&gt;     &lt;xsl:if test="mtu &lt; 1500"&gt;       &lt;!-- Select with Fast Ethernet interfaces with low MTUs --&gt;     &lt;/xsl:if&gt;   &lt;/xsl:when&gt;   &lt;xsl:otherwise&gt; </pre>

```
<xsl:if test="mtu > 8096">
  <!-- Select the non-Fast Ethernet interfaces with high MTUs -->
</xsl:if>
</xsl:otherwise>
</xsl:choose>
```

**Usage Examples** See [“Example: Configuring Dual Routing Engines” on page 357](#) and [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

**Related Documentation**

- [SLAX Statements Overview on page 79](#)
- [else if on page 98](#)
- [for-each on page 102](#)
- [if on page 104](#)

---

## else if

---

**Syntax**

```
if (expression) {
    /* code */
}
else if (expression) {
    /* code */
}
```

**Release Information** Statement introduced in version 1.0 of the SLAX language.

**Description** Include instructions that are processed if the expression defined in the preceding **if** statement evaluates to **FALSE** and the expression defined in the **else if** statement evaluates to **TRUE**. Multiple **else if** statements can be included, but the processor only executes the instructions contained in the first **else if** statement whose expression evaluates to **TRUE**. All subsequent **else if** statements are ignored.

**SLAX Example**

```
var $description2 = {
  if (description) {
    expr description;
  }
  else if (../description) {
    expr ../description;
  }
  else {
    expr "no description found";
  }
}
```

**XSLT Equivalent**

```
<xsl:variable name="description2">
  <xsl:choose>
    <xsl:when test="description">
      <xsl:value-of select="description"/>
    </xsl:when>
    <xsl:when test="../description">
      <xsl:value-of select="../description"/>
    </xsl:when>
    <xsl:otherwise>unknown</xsl:otherwise>
  </xsl:choose>
</xsl:variable>
```

```
</xsl:choose>
</xsl:variable>
```

**Usage Examples** See [“Example: Configuring Dual Routing Engines” on page 357](#) and [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

**Related Documentation**

- [SLAX Statements Overview on page 79](#)
- [else on page 97](#)
- [for-each on page 102](#)
- [if on page 104](#)

## expr

<b>Syntax</b>	<code>expr <i>expression</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Generate the string value of an XPath expression and add it to the result tree. The XPath expression might consist of a function call, a location path, a literal number, or a string. SLAX-specific operators are permitted. This statement cannot be used at the top-level of a script. It can only appear within a code block. By default, characters such as "&lt;", "&gt;", and "&amp;" are escaped into proper XML as "&amp;lt;", "&amp;gt;", and "&amp;amp;", respectively.</p> <p>The <b>expr</b> statement is most commonly used to invoke functions that return no results, for conditional variable assignment, and to return text content from a template.</p>
<b>Attributes</b>	<b>expression</b> —XPath expression to evaluate. The resulting string is added to the result tree.
<b>SLAX Example</b>	<pre>expr "Test: "; expr substring-before(name, "."); expr status; expr jcs:output("Test");</pre>
<b>XSLT Equivalent</b>	<pre>&lt;xsl:text&gt;Test: &lt;/xsl:text&gt; &lt;xsl:value-of select="substring-before(name, '.')" /&gt; &lt;xsl:value-of select="status" /&gt; &lt;xsl:value-of select="jcs:output('Test')" /&gt;</pre>
<b>Usage Examples</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Adding a Final then accept Term to a Firewall on page 331</a></li> <li>• <a href="#">Example: Automatically Configuring Logical Interfaces and IP Addresses on page 343</a></li> </ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">XPath Expressions Overview for SLAX on page 65</a></li> <li>• <a href="#">message on page 109</a></li> <li>• <a href="#">terminate on page 126</a></li> <li>• <a href="#">trace on page 127</a></li> <li>• <a href="#">uexpr on page 128</a></li> </ul>

## fallback

---

<b>Syntax</b>	<pre>fallback {     /* body to execute if extension function        or element is unavailable */ }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Specify statements to use when an extension function or element is not available in the current implementation. The <b>fallback</b> statement is enclosed within another instruction element to indicate what fallback code should be run if the script processor does not recognize the enclosing instruction element. The script executes the body of the fallback statement to handle this error condition.</p> <p>A script might utilize this statement when it is run in environments that support different extension elements.</p>
<b>SLAX Example</b>	<p>The following example script declares the namespace binding <b>test</b> with a URI of "test". The code attempts to reference the nonexistent extension element <b>&lt;test:fake&gt;</b>, which is not supported, and the code instead executes the fallback instructions.</p> <pre>version 1.1; ns junos = "http://xml.juniper.net/junos/*/junos"; ns xnm = "http://xml.juniper.net/xnm/1.1/xnm"; ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0"; ns test extension = "test";  match / {     &lt;op-script-results&gt; {         /* Fake extension element */         &lt;test:fake&gt; {             expr slax:output( "&lt;test:fake&gt; exists!" );             fallback {                 expr slax:output( "&lt;test:fake&gt; does not exist." );             }         }     } }</pre>

## for

---

<b>Syntax</b>	<pre>for <i>name</i> (<i>expression</i>) {     /* code */ }  for <i>name</i> (<i>min</i> ... <i>max</i>) {     /* code */ }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Iterate through an integer set or a node set without changing the context, and execute a block of statements using each member of the integer or node set as the value of the given variable.

If the argument is an XPath expression, the variable is assigned each member of the node set selected by the expression in sequence. If the argument is an integer set, the iteration operator (...) generates a sequence of nodes with the value of each integer between the left and right operands. If the left operand is greater than the right operand, the numbers are generated in decreasing order. The variable takes on the value of each integer in sequence. For each iteration, the contents are then evaluated, processed according to the instructions contained in the **for** code block.

**Attributes** *expression*—XPath expression that selects the nodes to be processed.

*max*—Integer or variable that defines the end value of the integer sequence. If the end value is less than the start value, the numbers are generated in decreasing order.

*min*—Integer or variable that defines the starting value of the integer sequence. If the start value is greater than the end value, the numbers are generated in decreasing order.

*name*—Identifier of the **for** loop variable, which takes on the values of each member of the integer or node set. This variable can be referenced within the **for** loop code block.

**SLAX Example** In the following example, the **for** loop iterates over the **interfaces** node. The XPath expression selects each **name** node that is a child of the **interface** node and that has a value beginning with the 'ge-' designator. The selection is assigned to the **\$name** variable, which is used within that iteration of the **for** loop code block. The **for** loop outputs a **<name>** element for each selection. The content of each **<name>** element is the interface name currently stored in the **\$name** variable for that iteration. The end result is a list of all Gigabit Ethernet interfaces on the device.

```
for $name (interfaces/interface[starts-with(name, 'ge-')]) {
  <name> {
    expr $name;
  }
}
```

In the following example, the **for** loop iterates over the integers 1 through 3, and the variable **\$int** assumes each integer value. For each iteration, the code block generates an **<item>** element, which contains the attribute **item-number** with a value equal to the current integer value of **\$int**.

```
for $int (1 ... 3) {
  <item> {
    attribute "item-number" {
      expr $int;
    }
  }
}

/* Output: <item item-number="1"/><item item-number="2"/><item item-number="3"/>
*/
```

- Related Documentation**
- [SLAX Statements Overview on page 79](#)
  - [XPath Overview on page 19](#)
  - [for-each on page 102](#)

---

## for-each

---

**Syntax**

```
for-each (expression) {  
    /* code */  
}  
  
/* Syntax added in version 1.1 of the SLAX language.*/  
for-each (min ... max) {  
    /* code */  
}
```

**Release Information** Statement introduced in version 1.0 of the SLAX language.  
Support for iteration operator (...) added in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Include a looping mechanism that repeats script processing for each XML element in the specified node set or each value in the integer set.

If the argument is an XPath expression, the element nodes are selected by the value of the XPath expression. If the argument is an integer set, the iteration operator (...) generates a sequence of nodes with the value of each integer between the left and right operands. If the left operand is greater than the right operand, the numbers are generated in decreasing order. For each iteration, the contents are then evaluated, processed according to the instructions contained in the **for-each** code block.

**Attributes** **for-each *expression***—XPath expression that selects the nodes to be processed.

***max***—Integer or variable that defines the end value of the integer sequence. If the end value is less than the start value, the numbers are generated in decreasing order.

***min***—Integer or variable that defines the starting value of the integer sequence. If the start value is greater than the end value, the numbers are generated in decreasing order.

**SLAX Example** The following code iterates over each **chassis-sub-module** element that has a part-number child element equal to 750-000610. For each match, the script outputs a **<message>** element with the name of the module and the name and description of the submodule.

```
for-each ($inventory/chassis/chassis-module/  
    chassis-sub-module[part-number == '750-000610']) {  
    <message> "Down rev PIC in " _ ../name _ ", " _ name _ ": " _ description;  
}
```

The following code iterates over the integers 1 through 3. For each iteration, the code block generates an **<item>** element, which contains the attribute **item-number** with a value equal to the current integer value of the set.

```
for-each (1 ... 3) {  
    <item> {
```

```

        attribute "item-number" {
            expr .;
        }
    }

    /* Output: <item item-number="1"/><item item-number="2"/><item item-number="3"/>
    */

```

- Usage Examples**
- [Example: Adding T1 Interfaces to a RIP Group on page 336](#)
  - [Example: Configuring Administrative Groups for LSPs on page 350](#)
  - [Example: Configuring Dual Routing Engines on page 357](#)
  - [Example: Imposing a Minimum MTU Setting on page 382](#)
  - [Example: Limiting the Number of E1 Interfaces on page 388](#)
  - [Example: Requiring and Restricting Configuration Statements on page 422](#)

- Related Documentation**
- [SLAX Statements Overview on page 79](#)
  - [XPath Overview on page 19](#)
  - [for on page 100](#)
  - [xsl:for-each on page 43](#)

## function

**Syntax**

```

function function-name (argument-list) {
    ...
    result return-value;
}

function function-name () {
    param param-name1;
    param param-name2;
    param param-name3 = default-value;
    ...
    result return-value;
}

```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Define an extension function that can be used in XPath expressions. The **function** statement must be defined as a top-level statement in the script using a qualified name for the function identifier. The argument list is a comma-separated list of parameter names, which are positionally assigned based on the function call. Trailing arguments can have default values. Alternatively, you can define function parameters inside the function block using the **param** statement. The function body is a set of statements, which should include a **result** statement that defines the return value for the function.

If there are fewer arguments in the function invocation than in the definition, the default values are used for any trailing arguments. If there are more arguments in the function invocation than in the definition, the function call generates an error.

**Attributes**    *function-name*—Specifies the name of the function as a qualified name.

*argument-list*—Comma-separated list of parameter names, which are positionally assigned based on the function call. Trailing arguments can have default values.

*return-value*—XML element or XPath expression, scalar value, or a set of instructions providing the return value of the function.

**SLAX Example**    The following example defines the function **size**, which has three parameters: **width**, **height**, and **scale**. The default value for **scale** is 1. If the function call argument list does not include the **scale** argument, the calculation uses the default value of 1 for that argument. The function's return value is the product of the **width**, **height**, and **scale** variables enclosed in a **<size>** element.

In the main match template, the function call uses width and height data selected from each **graphic/dimension** element in the source XML file. The script evaluates the function, and the **copy-of** statement emits the return value to the result tree as the contents of the **<out>** element.

```
version 1.1;
ns my = "http://www.example.com/myfunctions";

function my:size ($width, $height, $scale = 1) {
    result <size> {
        expr $width * $height * $scale;
    }
}

match / {
    for-each (graphic/dimension) {
        <out> {
            copy-of my:size((width/.), (height/.));
        }
    }
}
```

**Related Documentation**

- [SLAX Functions Overview on page 70](#)
- [param on page 118](#)
- [result on page 121](#)

---

## if

**Syntax**

```
if (expression) {
    /* code */
}
else if (expression) {
    /* code */
}
else {
    /* code */
}
```



}

**Release Information** Statement introduced in version 1.0 of the SLAX language.

**Description** Include a conditional construct that causes instructions to be processed if the Boolean expression evaluates to **TRUE**.

Optionally, you can include multiple **else if** statements following an **if** statement to perform additional conditional tests if the expression in the **if** statement evaluates to **FALSE**. Multiple **else if** statements can be included, but the processor only executes the instructions contained in the first **else if** statement whose expression evaluates to **TRUE**; all subsequent **else if** statements are ignored. The optional **else** statement includes a default set of instructions that are processed if the expressions defined in all associated **if** and **else if** statements evaluate to **FALSE**.

**Attributes** *expression*—Specifies the expression to evaluate.

**SLAX Example**

```
var $description2 = {
  if (description) {
    expr description;
  }
  else if (../description) {
    expr ../description;
  }
  else {
    expr "no description found";
  }
}
```

**XSLT Equivalent**

```
<xsl:variable name="description2">
  <xsl:choose>
    <xsl:when test="description">
      <xsl:value-of select="description"/>
    </xsl:when>
    <xsl:when test="../description">
      <xsl:value-of select="../description"/>
    </xsl:when>
    <xsl:otherwise>unknown</xsl:otherwise>
  </xsl:choose>
</xsl:variable>
```

**Usage Examples** See [“Example: Configuring Dual Routing Engines” on page 357](#), [“Example: Preventing Import of the Full Routing Table” on page 416](#), and [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

**Related Documentation**

- [SLAX Statements Overview on page 79](#)
- [else on page 97](#)
- [else if on page 98](#)
- [for-each on page 102](#)

## import

---

<b>Syntax</b>	<code>import href;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Import rules from an external file or style sheet, which provide access to all the declarations and templates within the imported item. Any <b>import</b> statements must be the first elements in the script or style sheet. The path can be any URI. The path <code>../import/junos.xsl</code> is standard for all commit scripts, op scripts, and event scripts.</p> <p>Imported rules are overwritten by any subsequent matching rules within the importing script. If more than one file or style sheet is imported, the items imported last override each previous import where the rules match.</p>
<b>Attributes</b>	<b>href</b> —Specifies the location of the imported file or style sheet.
<b>SLAX Example</b>	<p>In the example, the main script imports the file <b>route-rules.slax</b>, which contains a template rule for &lt;route&gt; elements.</p> <pre>version 1.1; import "route-rules.slax";  match route {   &lt;routes&gt; {     apply-imports;   } }</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Named Templates in Junos OS Automation Scripts on page 225</a></li><li>• <a href="#">Global Parameters and Variables in Junos OS Automation Scripts on page 175</a></li><li>• <a href="#">Required Boilerplate for Commit Scripts on page 251</a></li><li>• <a href="#">Required Boilerplate for Event Scripts on page 615</a></li><li>• <a href="#">Required Boilerplate for Op Scripts on page 445</a></li><li>• <a href="#">apply-imports on page 89</a></li></ul>

## key

---

<b>Syntax</b>	<pre>key name {   match pattern;   value expression; }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Define a key for use with the <b>key()</b> XPath function. Keys are an alternative to IDs and are used to index the nodes within an XML document. The key must be defined as a top-level statement in the script. A <b>key</b> definition consists of the key identifier, the nodes to index,

and the value that is paired with the key name to reference the matching nodes. The **key()** function is then used to locate the appropriate nodes.

The **key()** function works with the XML document of the current node and uses the specified **key** definition to retrieve nodes that are referenced by a particular name and value. The function arguments are the key name and the desired key's value. The return value is a node set that includes all nodes referenced by that key name and value. If the desired key value is provided as a node set, rather than a string, the returned node set is a union of all the referenced nodes for the key values expressed by the nodes within the node set.

For example, if you define the key:

```
key func {
  match prototype;
  value @name;
}
```

the following code would select **<prototype>** elements that have a **name** attribute with a value of "trace", and then output the value of the child element **<return-type>**:

```
for-each ( key("func", "trace") ) {
  <out> return-type/.;
}
```

**Attributes** **key name**—Key identifier, which uniquely identifies the key within the script and is passed as the first argument to the **key()** function.

**match pattern**—XPath expression that selects the set of nodes to index.

**value expression**—XPath expression that defines the value of the key.

**SLAX Example** The following op script creates two **key** definitions, **protocol** and **next-hop**, which are used to retrieve and display all static routes and all routes with a next hop of ge-0/0/0.0 on a device. The script invokes the Junos XML API **get-route-information** command to obtain the route information for the device. The **for-each( \$results )** statement changes the current node to the **\$results** XML document. The subsequent **for-each** loops use the keys to retrieve all nodes that are indexed according to the key names and values.

The **for-each( key( "protocol", "Static" )** statement uses the **protocol** key definition, which matches on **route-table/rt** elements, to retrieve the desired nodes. The **rt-entry/protocol-name** key value matches the **<protocol-name>** child elements that have the value "Static". The code block executes using **<rt>** as the context node. For each match, the script outputs the value of the **<rt-destination>** element.

The **for-each( key( "next-hop", "ge-0/0/0.0" )** statement uses the "next-hop" key definition, which matches on **route-table/rt** elements, to retrieve the desired nodes. The **rt-entry/nh/via** key value matches the **<via>** child elements that have the value "ge-0/0/0.0". The code block executes using **<rt>** as the context node. For each match, the script outputs the value of the **<rt-destination>** element.

```
version 1.1;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
```

```

key protocol {
  match route-table/rt;
  value rt-entry/protocol-name;
}
key next-hop {
  match route-table/rt;
  value rt-entry/nh/via;
}

match / {
  <op-script-results> {
    var $results = jcs:invoke("get-route-information");

    for-each( $results ) {
      /* Display all static routes */
      <output> "Static routes: ";
      for-each( key( "protocol", "Static" ) ) {
        <output> rt-destination;
      }
      /* Display all routes with next-hop of ge-0/0/0.0 */
      <output> "Next-hop ge-0/0/0.0: ";
      for-each( key( "next-hop", "ge-0/0/0.0" ) ) {
        <output> rt-destination;
      }
    }
  }
}

```

- Related Documentation**
- [XPath Overview on page 19](#)
  - [XPath Expressions Overview for SLAX on page 65](#)

## match

<b>Syntax</b>	<pre> match <i>expression</i> {   <i>statements</i>; } </pre>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	Declare a template that contains rules to apply when a specified node is matched. The <b>match</b> statement associates the template with an XML element. The <b>match</b> statement can also be used to define a template for a whole branch of the XML document. For example, <b>match /</b> matches the root element of the document.
<b>Attributes</b>	<b>expression</b> —XPath expression specifying the nodes to which to apply the template.
<b>SLAX Example</b>	<pre> match host-name {   &lt;hello&gt; .; } </pre>
<b>XSLT Equivalent</b>	<pre> &lt;xsl:template match="host-name"&gt;   &lt;hello&gt;     &lt;xsl:value-of select="."/&gt;   &lt;/hello&gt; &lt;/xsl:template&gt; </pre>

**Usage Examples** [“Example: Adding a Final then accept Term to a Firewall” on page 331](#), [“Example: Configuring Dual Routing Engines” on page 357](#), and [“Example: Preventing Import of the Full Routing Table” on page 416](#).

**Related Documentation**

- [apply-templates on page 89](#)
- [call on page 93](#)
- [mode on page 109](#)
- [priority on page 119](#)
- [template on page 125](#)
- [with on page 131](#)

## message

<b>Syntax</b>	<pre>message <i>expression</i>;  message {     /* body */ }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Generate an error message that is immediately displayed to the user, typically on the standard error file descriptor. This is different from most script output, which is displayed only after the script generates the final result tree.</p> <p>Junos OS op scripts, event scripts, and commit scripts prepend "error:" to the displayed message when generating text output. When generating XML output, the scripts place the output inside a <code>&lt;message&gt;</code> element, which is enclosed in an <code>&lt;xmn:error&gt;</code> element.</p> <p>If the <b>message</b> statement is used in a commit script, the script will generate two errors and terminate the commit process. If the <b>message</b> statement is used in an event script, the script writes the message to the output file, if one is configured.</p>
<b>Attributes</b>	<b>message <i>expression</i></b> —XPath expression or string emitted as output.
<b>SLAX Example</b>	<pre>if (not(valid)) {     message "The " _ name() _ " node is not valid"; }</pre>

## mode

<b>Syntax</b>	<code>mode <i>qualified-name</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Indicate the mode in which a template needs to be applied for the template to be used. If templates are applied in the specified mode, the <b>match</b> statement is used to determine whether the template can be used with the particular node. If more than one template</p>

matches a node in the specified mode, the priority statement determines which template is used. The highest priority wins. If no priority is specified explicitly, the priority of a template is determined by the **match** statement.

This statement is comparable to the **mode** attribute of the `<xsl:template>` element. You can include this statement inside a SLAX **match** or **apply-templates** statement.

**SLAX Example**

```
match * {
  mode "one";
  <one> .;
}

match * {
  mode "two";
  <two> string-length(.);
}

match / {
  apply-templates version {
    mode "one";
  }
  apply-templates version {
    mode "two";
  }
}
```

**XSLT Equivalent**

```
<xsl:template match="*" mode="one">
  <one>
    <xsl:value-of select="."/>
  </one>
</xsl:template>

<xsl:template match="*" mode="two">
  <two>
    <xsl:value-of select="string-length(.)"/>
  </two>
</xsl:template>


<xsl:template match="/">
  <xsl:apply-templates select="version" mode="one"/>
  <xsl:apply-templates select="version" mode="two"/>
</xsl:template>
```

**Usage Examples**    See [“Example: Adding a Final then accept Term to a Firewall” on page 331](#).

**Related  
Documentation**

- [apply-templates on page 89](#)
- [call on page 93](#)
- [match on page 108](#)
- [priority on page 119](#)
- [template on page 125](#)
- [with on page 131](#)
- [xsl:template on page 48](#)

## mvar

<b>Syntax</b>	<code>mvar \$name[=<i>initial-value</i>];</code>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Declare a mutable variable in a SLAX script. You can initialize a mutable variable when you declare it by following the variable name with an equal sign (=) and a value.</p> <p>Mutable variables differ from variables declared using the <b>var</b> statement in that you can change the value of a mutable variable after it is declared. To initialize or set the value of a mutable variable after you declare it, use the <b>set</b> statement. To append a value to the node set contained in a mutable variable, use the <b>append</b> statement.</p>
	<p> <b>NOTE:</b> Mutable variables use non-standard SLAX specific extension elements, which can affect the portability of a script.</p>
<b>Attributes</b>	<p><b>name</b>—Mutable variable identifier. After declaration, you can reference the variable within expressions by using the identifier prefixed with the dollar sign (\$) character.</p> <p><b>initial-value</b>—Initial value assigned to the mutable variable.</p>
<b>SLAX Example</b>	<p>The following example creates the mutable variable <b>block</b>, and initializes it. The <b>set</b> statement assigns a new value to the <b>block</b> variable, overwriting the initial value set in the declaration. In the <b>for</b> loop, the code iterates over each item in the specified list and appends an <b>&lt;item&gt;</b> element with two child elements, <b>&lt;name&gt;</b> and <b>&lt;size&gt;</b>, to the node set stored in the <b>block</b> variable.</p> <pre> mvar \$block= &lt;block&gt; "start here"; set \$block = &lt;block&gt; "item list";  for \$item (list) {     append \$block += &lt;item&gt; {         &lt;name&gt; \$item/name;         &lt;size&gt; \$item/size;     } } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SLAX Variables Overview on page 76</a></li> <li>• <a href="#">append on page 88</a></li> <li>• <a href="#">set on page 122</a></li> <li>• <a href="#">var on page 129</a></li> </ul>

## number

**Syntax**      `number expression {  
                  format numbering-style;`

```
        grouping-separator character;  
        grouping-size number;  
    }  
  
    number {  
        count nodes;  
        format numbering-style;  
        from nodes;  
        grouping-separator character;  
        grouping-size number;  
        level "single" | "multiple" | "any";  
    }
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Generate a formatted number string, which is output to the result tree. When used with an argument, the statement formats the number given by that XPath expression. When used without an argument, the statement uses the **count**, **from**, and **level** options to generate the number based on the position of one or more nodes within the current XML document. In both cases, optional statements specify the formatting for that number. If needed, you can also redirect the formatted number string to a variable or output method instead of the result tree.

**Attributes** **number *expression***—XPath expression providing the number to format.

**count *nodes***—XPath expression specifying which nodes should be counted. If **count** is omitted, it defaults to nodes with the same name as the current node.

**format *numbering-style***—A string, variable, or XPath expression that defines the number formatting.

The **format** option can include the following:

- **start string**—Any non-alphanumeric characters that precede the first number token in the format string. The start string is prepended to the formatted number string.
- **number token**—One or more number tokens that indicate what numbering format to use for the included numbers. The formatted number string only includes more than one number if the **level** option is set to "multiple". [Table 9 on page 113](#) lists format values and corresponding styles. The default value is "1", which uses a decimal format style. When using decimal format, you can specify the minimum length of the formatted number string by preceding the "1" with one or more zeros.
- **token separator**—Non-alphanumeric characters that separate number tokens in the format string. These characters are included in the formatted number string between the computed numbers.
- **end string**—Any non-alphanumeric characters that follow the last number token in the format string. The end string is appended to the formatted number string.



Table 9: Numbering Styles for SLAX Statement `number, format` Option

Format Value	Style	Example
1	Decimal format	1 2 3 ...10 11 ...
01	Decimal format with a minimum output string length of 2	01 02 03 ... 10 11 ...
001	Decimal format with a minimum output string length of 3	001 002 003 ... 010 011 012 ... 100, 101
a	Lowercase alphabetic numbering	a b c ... z ... aa ab ... az ... ba bb ...
A	Uppercase alphabetic numbering	A B C ... Z ... AA AB ... BA BB ...
i	Lowercase Roman numbering	i ii iii iv v ...
I	Uppercase Roman numbering	I II III IV V ...

**from *nodes***—XPath expression specifying from which element to start the count. When **level** is set to **single** or **multiple**, this option constrains the counting to only node descendants of the nearest ancestor that matches the expression. When **level** is set to **any**, this option constrains the counting to only nodes that follow the nearest ancestor or preceding node of the current node that matches the expression.

**grouping-separator *character***—Character used to delimit groups of digits for numbers expressed in decimal format. For example, decimal notation uses a comma as the delimiter between digit groupings.

**grouping-size *number***—Defines the number of digits in a group for numbers expressed in decimal format. Setting this option causes the formatted number to be split into multiple groups according to the grouping size, with the grouping separator delimiting the groups. For example, decimal notation often uses a grouping size of 3.

**level**—Specifies what type of counting to perform. Accepted values are **single**, **multiple**, and **any**. The default is **single**. Specifying **single** starts the counting from the first ancestor node, specifying **multiple** starts the counting from any ancestor node, and specifying **any** starts the counting from any node.

- **single**—Perform only one count. The current node, if it matches the count expression, or the nearest ancestor that matches the count expression, is counted. The position of the node in document order, relative to its siblings that also match the count parameter, is used as the number to be formatted.
- **multiple**—Separately count all nodes that match the count expression and are either the current node or an ancestor of the current node. The position of each node in document order, relative to its siblings that also match the count parameter, is used as one of the numbers to be formatted.

- **any**—Perform only one count. The current node, if it matches the count pattern, or its nearest ancestor or preceding node that matches the count pattern, is counted. The position of the node in document order, relative to all other matching nodes that are ancestors or precede the node, is used as the number to be formatted.



**NOTE:** Currently libxslt (1.1.26) does not support the “language” and “letter-value” options for the `<xsl:number>` element. While SLAX provides a means of encoding these XSLT constructs, they are not usable under Junos OS.

**SLAX Example** The following sample code iterates from 1 through 5. For each integer, the **number** statement outputs the equivalent uppercase Roman numeral value.

```
for $i (1 ... 5) {
  number $i {
    format "I ";
  }
}
```

I II III IV V

The following sample code provides the string “1234567890” to the **number** statement, which formats the output in decimal format with a group size of 3 and a comma as a group delimiter.

```
number "1234567890" {
  grouping-size 3;
  grouping-separator ",";
  format "1";
}
```

1,234,567,890

The following sample code counts all the **name** elements in the configuration hierarchy stored in the variable **\$data**. The **count** option combined with the **level "multiple"** option tracks the count for any **name** elements under the **interface**, **unit**, and **address** elements.

The **format** option (1.A.a) includes a start string, which is an open parenthesis, and an end string, which is a close parenthesis and a space character. The number tokens are “1”, “A”, and “a”, which define the formatting of the numbers as decimal format, uppercase alphabetic numbering, and lowercase alphabetic numbering, respectively. The token separator is a period, which is also included in the output.

```
var $data := {
  <interfaces> {
    <interface> {
      <name> "ge-0/0/0";
      <unit> {
        <name> "0";
      }
      <unit> {
        <name> "1";
      }
    }
  }
}
```

```

    }
    <interface> {
      <name> "ge-0/1/0";
      <unit> {
        <name> "10";
        <family> {
          <inet>;
        }
      }
    }
    <interface> {
      <name> "ge-2/0/2";
      <unit> {
        <name> "0";
        <family> {
          <inet> {
            <address> {
              <name> "10.1.1.1/24";
            }
          }
        }
      }
    }
  }
}

for-each ($data//name) {
  number {
    level "multiple";
    count interface|unit|address;
    format "(1.A.a) ";
  }
  expr . _ "\n";
}

```

For the generated numbers displayed in the result tree, the decimal number in parentheses is associated with a particular interface. For each interface, the uppercase letter is associated with each logical unit name, and any lowercase letter is associated with the address **name** element for that logical unit, which is the IP address.

```

(1) ge-0/0/0
(1.A) 0
(1.B) 1
(2) ge-0/1/0
(2.A) 10
(3) ge-2/0/2
(3.A) 0
(3.A.a) 10.1.1.1/24

```

- Related Documentation
- [decimal-format on page 95](#)
  - [output-method on page 115](#)

## output-method

**Syntax**

```

output-method output-format {
  cdata-section-elements name-list;
  doctype-public string;
  doctype-system string;
}

```

```
    encoding string;  
    indent "yes" | "no";  
    media-type string;  
    omit-xml-declaration "yes" | "no";  
    standalone "yes" | "no";  
    version string;  
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Define the style used for result tree output. The **output-method** statement must be defined as a top-level statement in the script. Output formats include HTML, text, or XML. The default is XML, unless the first child element of the root node is **<html>** and there are no preceding text nodes, in which case the default output format is HTML.

**Attributes** **output-format**—Specify the format of the output. Acceptable values are “html”, “text”, “xml”, or a qualified name. The default is XML, unless the first child element of the root node is **<html>** and there are no preceding text nodes, in which case the default output format is HTML. Specifying a format of XML adds the XML declaration (**<?xml ?>**) to the result tree file.

**cdata-section-elements *name-list***—Specify a space-delimited list of the names of output elements whose text contents should be output to the result tree using CDATA sections. A CDATA section starts with **<![CDATA[** and ends with **]]>**, and the contents of the section are interpreted by an XML parser as character data only, rather than markup.

**doctype-public *string***—Add the DOCTYPE declaration to the result tree, and specify the value of the **PUBLIC** attribute, which tells the parser where to locate the Document Type Definition (DTD) file.

**doctype-system *string***—Add the DOCTYPE declaration to the result tree, and specify the value of the **SYSTEM** attribute, which tells the parser where to locate the DTD file on the system.

**encoding *string***—Explicitly add the pseudo-attribute **encoding** to the XML declaration in the output, and specify the character encoding used to encode the document, for example UTF-8, UTF-16, or ISO-8859-1.

**indent "yes" | "no"**—Specify whether to indent the result tree output according to the hierarchical structure. Acceptable values are “yes” and “no”. The default is no indentation.

**media-type *string***—Define the MIME content type of the output. The default is “text/xml”.

**omit-xml-declaration "yes" | "no"**—Specify whether to include or omit the XML declaration (**<?xml ?>**) in the output. The default is “no”.

**standalone "yes" | "no"**—Explicitly add the pseudo-attribute **standalone** with the given string value to the XML declaration (**<?xml ?>**) in the output. Acceptable values are “yes” and “no”. The **standalone** attribute is only relevant if the document uses a

DTD. If the **standalone** option is not included in the **output-method** statement, there is no explicit declaration in the result tree, which is identical to **standalone="no"**.

**version string**—For HTML and XML formats, set the W3C version for the output format. The pseudo-attribute **version** is included in the XML declaration (`<?xml ?>`) with the given version number.

**SLAX Example** The following example uses the output method XML, which creates an XML declaration in the result tree output and adds the pseudo-attributes **version**, **encoding**, and **standalone** to the declaration. The DOCTYPE declaration has the root element `<html>` and provides values for both the **PUBLIC** and the **SYSTEM** attributes.

```
version 1.1;

output-method xml {
  doctype-public "-//W3C//DTD XHTML 1.0 Transitional//EN";
  doctype-system "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd";

  encoding "utf-8";
  indent "yes";
  omit-xml-declaration "no";
  standalone "no";
  version "1.0";
}

match / {
  <html> {
    <script type="text/javascript" src="/assets/js/api.js">;
    /* ... */
  }
}
```

The script produces the following output:

```
<?xml version="1.0" encoding="utf-8" standalone="no"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
  <script type="text/javascript" src="/assets/js/api.js"></script>
  ...
</html>
```

The following example is similar to the previous example except that the script does not specify an output format. Since the first child element of the root node is `<html>`, the output format defaults to HTML.

```
version 1.1;

output-method {
  doctype-public "-//W3C//DTD XHTML 1.0 Transitional//EN";
  doctype-system "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd";

  encoding "utf-8";
  indent "yes";
  omit-xml-declaration "no";
  standalone "no";
  version "1.0";
}
```

```

match / {
  <html> {
    <script type="text/javascript" src="/assets/js/api.js">;
    /* ... */
  }
}

```

The default output format is HTML. The XML declaration is omitted from the output.

```

<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html><script type="text/javascript" src="/assets/js/api.js"></script></html>

```

## param

<b>Syntax</b>	<code>param \$name=value;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Declare a parameter for a template or for the style sheet as a whole. Template parameters declared with the <b>param</b> statement must be placed inside the template code block. A global parameter, the scope of which is the entire style sheet, must be declared at the top level of the style sheet. You can include an initial value by following the parameter name with an equal sign (=) and a value expression. A parameter whose value is set by Junos OS at script initialization must be defined as a global parameter.</p> <p>In SLAX, parameter and variable names are declared and accessed using the dollar sign (\$). This is unlike the <b>name</b> attribute of <code>&lt;xsl:variable&gt;</code> and <code>&lt;xsl:parameter&gt;</code> elements, which do not include the dollar sign in the declaration.</p>
<b>Attributes</b>	<p><b>name</b>—Defines the name of the parameter.</p> <p><b>value</b>—Defines the default value for the parameter, which is used if the person or client application that executes the script does not explicitly provide a value.</p>
<b>SLAX Example</b>	<pre> param \$vrf; param \$dot = .; </pre>
<b>XSLT Equivalent</b>	<pre> &lt;xsl:param name="vrf"/&gt; &lt;xsl:param name="dot" select="."/&gt; </pre>
<b>Usage Examples</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Imposing a Minimum MTU Setting on page 382</a></li> <li>• <a href="#">Example: Limiting the Number of ATM Virtual Circuits on page 384</a></li> <li>• <a href="#">Example: Limiting the Number of E1 Interfaces on page 388</a></li> <li>• <a href="#">Example: Preventing Import of the Full Routing Table on page 416</a></li> <li>• <a href="#">Example: Requiring and Restricting Configuration Statements on page 422</a></li> </ul>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SLAX Parameters Overview on page 72</a></li> <li>• <a href="#">SLAX Templates Overview on page 66</a></li> <li>• <a href="#">template on page 125</a></li> </ul>

- [var on page 129](#)
- [with on page 131](#)

## preserve-space

<b>Syntax</b>	<code>preserve-space <i>element-list</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Preserve whitespace-only child text nodes for the source tree element nodes listed, but not for the child text nodes of the element node children. To preserve whitespace-only child text nodes of the element node children, specify the child nodes as separate entries in the preserve-space element list. Specifying an asterisk preserves whitespace-only child elements for all elements, which is the default behavior. A text node is considered whitespace-only if it includes only spaces, tabs, newlines, and carriage returns.</p> <p>The <b>preserve-space</b> statement is only needed if the <b>strip-space</b> statement has been used with an asterisk, indicating that whitespace-only child text nodes should be removed from all element nodes. In this case, use the <b>preserve-space</b> statement to indicate specific element nodes that should not have their whitespace-only child text nodes stripped.</p> <p>This statement must be defined as a top-level statement in the script.</p>
<b>Attributes</b>	<b><i>element-list</i></b> —Space-separated list of element names for which to preserve whitespace-only child text nodes.
<b>SLAX Example</b>	<p>The following example removes all whitespace-only text nodes from the source tree except for child elements of <code>&lt;user-context&gt;</code>:</p> <pre> version 1.1;  preserve-space user-context; strip-space *;  match / { ... } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">strip-space on page 125</a></li> </ul>

## priority

<b>Syntax</b>	<code>priority <i>number</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>If more than one template matches a node in the specified mode, this statement determines which template is used. The highest priority wins. If no priority is specified explicitly, the priority of a template is determined by the <b>match</b> statement.</p>

This statement is comparable to the **priority** attribute of the `<xsl:template>` element. You can include this statement inside a SLAX **match** statement.

SLAX Example	<pre>match * {     priority 10;     &lt;output&gt; .; }</pre>
XSLT Equivalent	<pre>&lt;xsl:template match="*" priority="10"&gt;     &lt;output&gt;         &lt;xsl:value-of select="."/&gt;     &lt;/output&gt; &lt;/xsl:template&gt;</pre>
Usage Examples	None of the examples in this manual use this statement.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">apply-templates on page 89</a></li><li>• <a href="#">call on page 93</a></li><li>• <a href="#">match on page 108</a></li><li>• <a href="#">mode on page 109</a></li><li>• <a href="#">template on page 125</a></li><li>• <a href="#">with on page 131</a></li><li>• <a href="#">xsl:template on page 48</a></li></ul>

---

## processing-instruction

---

Syntax	<pre>processing-instruction <i>instruction-name</i>;  processing-instruction <i>instruction-name</i> {     <i>instruction-value</i>; }</pre>
Release Information	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
Description	<p>Add an XML processing instruction to the result tree. A processing instruction is a mechanism to convey application-specific information inside an XML document. The application can detect processing instructions and change their behavior accordingly. The instruction name is mandatory and becomes the target of the processing instruction. It can be a hard-coded string, a variable, or an XPath expression. The optional body generates the processing instruction's content, which consists of one or more name-value pairs. The generated instruction is enclosed within the tags <code>&lt;?&gt;</code> and <code>?&gt;</code>.</p> <p>Junos OS SLAX scripts generally do not require the <b>processing-instruction</b> statement, because the result tree is processed directly by Junos OS. However, you might add a processing instruction to an XML document that is written to disk through the <code>&lt;xsl:document&gt;</code> instruction element or one of its related extension elements.</p>



**Attributes** *instruction-name*—Identifier for the processing instruction, which can be a string, a variable, or an XPath expression.

*instruction-value*—Instruction content, which consists of name-value pairs.

**SLAX Example** The following code creates the processing instruction `xml-stylesheet`. The instruction content contains two name-value pairs: `type` and `href`.

```
processing-instruction "xml-stylesheet" {
  expr 'type="text/css" ';
  expr 'href="style.css"';
}
```

The corresponding output in the result tree is:

```
<?xml-stylesheet type="text/css" href="style.css"?>
```

The following example writes an XML document to the file `/var/tmp/output.xml` using the `<xsl:document>` instruction element. The script adds a processing instruction named `instruction` to the document.

```
version 1.1;

match / {
  <op-script-results> {
    <xsl:document href="/var/tmp/output.xml" indent="yes" method="xml">
  {
    <document-element> {
      <element>;
      processing-instruction "instruction" {
        expr 'name="testing"';
      }
      <element>;
    }
  }
}
```

The script generates the file `/var/tmp/output.xml`, which contains the processing instruction enclosed within `<?>` and `?>` tags.

```
<?xml version="1.0"?>
<document-element>
  <element/>
  <?instruction name="testing"?>
  <element/>
</document-element>
```

## result

**Syntax**

```
result expression;

result {
  /* body */
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Define the return value for a function. The value can be a simple scalar value, an XML element or XPath expression, or a set of instructions that emit the value to be returned.

**Attributes** **result *expression***—XPath expression defining the return value of the function.

**SLAX Example** The following example defines three extension functions, **my:size()**, **my:box-parts()**, and **my:ark()**. The **my:ark()** function returns a node set containing an **<ark>** element that encloses the node set returned by the **my:box-parts()** function. The **my:box-parts()** function returns a node set containing a **<box>** element enclosing three **<part>** child elements. The content of each **<part>** element is the value returned by the **my:size()** function. The return value of the **my:size()** function is the product of the three parameters **width**, **height**, and **scale**.

```
version 1.1;
ns my exclude = "http://www.example.com/myfunctions";

function my:size ($x, $y, $scale = 1) {
    result $x * $y * $scale;
}
function my:box-parts ($width, $height, $depth, $scale = 1) {
    result <box> {
        <part count=2> my:size($width, $depth);
        <part count=2> my:size($width, $height);
        <part count=2> my:size($depth, $height);
    }
}
function my:ark () {
    result {
        <ark> {
            copy-of my:box-parts(2.5, 1.5, 1.5);
        }
    }
}

match / {
    var $res = my:ark();
    copy-of $res;
}
```

**Related Documentation**

- [SLAX Functions Overview on page 70](#)
- [copy-of on page 94](#)
- [function on page 103](#)

---

## set

**Syntax**      **set** *\$name* = *value*;

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Assign a value to a mutable variable. The variable must be defined using the **mvar** statement.

**Attributes** *name*—Name of the mutable variable.

*value*—Value to assign to the mutable variable.

**SLAX Example** The following example creates the mutable variable, **block**. The **set** statement assigns an initial value to the **block** variable. In the **for** loop, the code iterates over each item in the specified list and appends an **<item>** element with two child elements, **<name>** and **<size>**, to the node set stored in the **block** variable.

```
mvar $block;
set $block = <block> "item list";

for $item (list) {
  append $block += <item> {
    <name> $item/name;
    <size> $item/size;
  }
}
```

**Related Documentation**

- [SLAX Variables Overview on page 76](#)
- [append on page 88](#)
- [mvar on page 111](#)
- [var on page 129](#)

## sort

**Syntax**

```
sort expression {;
  case-order "upper-first" | "lower-first";
  data-type "text" | "number" | type-name;
  order "ascending" | "descending";
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Control the order in which the **for-each** and **apply-templates** statements iterate through the current node list. By default, the **for-each** and **apply-templates** statements consider nodes in document order, but the **sort** statement defines the order prior to iterating through the node list. Insert the **sort** statement immediately after the **for-each** or **apply-templates** statement. The **sort** statement is only processed when the loop is first initiated.

The **sort** statement has an optional XPath expression and three optional parameters: **case-order**, **data-type**, and **order**. The XPath expression determines each node's comparison string used for sorting. The script evaluates the expression with the node as its context, and then translates the result into the comparison string for that node. If you do not specify an XPath expression, the default value is ".", which causes the string content of each node in the list to be compared. SLAX-specific operators such as **==** and **\_** cannot be used within the expression string. If the **sort** statement does not include any optional parameters, the list is sorted based on the string value of each node.

The **sort** statement does not permanently sort the underlying XML data structure, only the order of the current node list being used by the **for-each** or **apply-templates** statement. Multiple **sort** statements can be assigned to a single **for-each** or **apply-templates** statement. They are applied, in order, until a difference is found.

**Attributes** **expression**—XPath expression that determines each node's comparison string used for sorting. The default value is ".".

**case-order**—Specify whether to sort lowercase first or uppercase first. Acceptable values are "lower-first" or "upper-first". The default is "upper-first".

**data-type**—Specify the element type, which determines whether a numerical, lexical, or other sort is performed. Acceptable values are "number" and "text". The default is "text".

Setting **data-type** to "text" compares the strings based on their character values (that is ASCII code), so "0" is less than "9", which is less than "A", which is less than "Z", which is less than "a", which is less than "z". Setting **data-type** to "number" converts the strings to numbers and compares them numerically. With ascending text sorting, "100" would come before "11" because "0" has a lower ASCII code than "1", but with ascending number sorting, 11 would come before 100 because 11 is a smaller number than 100.

**order**—Specify whether to sort in ascending or descending order. Acceptable values are "descending" or "ascending". The default is "ascending".

**SLAX Example** The following example SLAX script executes the Junos XML API **get-interface-information** command and parses the resulting output. The **for-each** loop prints the name of each physical interface on the device sorted in ascending order.

```
version 1.1;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

match / {
  <op-script-results> {
    var $results = jcs:invoke("get-interface-information");
    for-each ($results/physical-interface/name) {
      sort . {
        data-type "text";
        order "ascending";
      }
      <interface-name> .;
    }
  }
}
```

**Related Documentation**

- [apply-templates on page 89](#)
- [for-each on page 102](#)

## strip-space

<b>Syntax</b>	<code>strip-space <i>element-list</i>;</code>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Remove whitespace-only child text nodes from the source tree element nodes listed, but not from the child text nodes of the element node children. To perform whitespace stripping on the child text nodes of the element node children, specify the child nodes as separate entries in the strip-space element list. Specifying an asterisk removes whitespace-only child elements from all elements. A text node is considered whitespace-only if it includes only spaces, tabs, newlines, and carriage returns.</p> <p>This statement must be defined as a top-level statement in the script. The default is to preserve all whitespace-only elements.</p>
<b>Attributes</b>	<i>element-list</i> —List of element names separated by spaces.
<b>SLAX Example</b>	<p>The following example removes all whitespace-only text nodes from the source tree except for child elements of <code>&lt;user-context&gt;</code>:</p> <pre> version 1.1;  preserve-space user-context; strip-space *;  match / { ... } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">preserve-space on page 119</a></li> </ul>

## template

<b>Syntax</b>	<pre> template <i>qualified-name</i> (<i>parameter-name</i> = <i>value</i>) {     /* code */ } </pre>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	<p>Declare a named template. You can include a comma-separated list of parameter declarations, with the parameter name and an optional equal sign (=) and value expression. You can declare additional parameters inside the code block using the <b>param</b> statement. You can invoke the template using the <b>call</b> statement.</p>
<b>SLAX Example</b>	<pre> match configuration {     var \$name-servers = name-servers/name;     call temp();     call temp(\$name-servers, \$size = count(\$name-servers));     call temp() {         with \$name-servers;         with \$size = count(\$name-servers);     } } </pre>

```
    }  
  
    template temp($name-servers, $size = 0) {  
        <output> "template called with size " _ $size;  
    }  
}
```

**XSLT Equivalent**

```
<xsl:template match="configuration">  
  <xsl:variable name="name-servers" select="name-servers/name"/>  
  <xsl:call-template name="temp"/>  
  <xsl:call-template name="temp">  
    <xsl:with-param name="name-servers" select="$name-servers"/>  
    <xsl:with-param name="size" select="count($name-servers)"/>  
  </xsl:call-template>  
  <xsl:call-template name="temp">  
    <xsl:with-param name="name-servers" select="$name-servers"/>  
    <xsl:with-param name="size" select="count($name-servers)"/>  
  </xsl:call-template>  
</xsl:template>  
  
<xsl:template name="temp">  
  <xsl:param name="name-servers"/>  
  <xsl:param name="size" select="0"/>  
  <output>  
    <xsl:value-of select="concat('template called with size ', $size)"/>  
  </output>  
</xsl:template>
```

**Usage Examples** See [“Example: Adding a Final then accept Term to a Firewall” on page 331](#) and [“Example: Adding T1 Interfaces to a RIP Group” on page 336](#).

**Related Documentation**

- [SLAX Parameters Overview on page 72](#)
- [SLAX Templates Overview on page 66](#)
- [apply-templates on page 89](#)
- [call on page 93](#)
- [match on page 108](#)
- [mode on page 109](#)
- [priority on page 119](#)
- [with on page 131](#)

---

## terminate

---

**Syntax**

```
terminate expression;  
  
terminate {  
    /* body */  
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Generate an error message that is immediately displayed to the user, and exit the script.

Junos OS op scripts, event scripts, and commit scripts prepend "error:" to the displayed message when generating text output. When generating XML output, the scripts place the output inside a **<message>** element, which is enclosed in an **<xmn:error>** element.

If the **terminate** statement is used in a commit script, the script will generate two errors and terminate the script and the commit process. If the **terminate** statement is used in an event script, the script writes the message to the output file, if one is configured, and terminates the script.

**Attributes** *expression*—XPath expression or string emitted as output.

**SLAX Example**

```
if (not(valid)) {
    terminate "The " _ name() _ " node is not valid. Exiting script.";
}
```

**Related Documentation**

- [message on page 109](#)

## trace

**Syntax**

```
trace expression;

trace {
    /* body */
}
```

**Release Information** Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Write a message to the trace file, if tracing is enabled. If tracing is not enabled, trace output is not generated. The **trace** message can be an XPath expression or string, or it can be generated by the contents of a **trace** statement block.

Enabling of tracing is typically a feature of the environment in which a SLAX script is called. When executing a script using the **slaxproc** command, include the **--trace** or **-t** option to enable tracing. For more information about slaxproc, see "[Understanding the SLAX Processor \(slaxproc\)](#)" on page 148.

**Attributes** *trace expression*—XPath expression or string written to the trace file.

**SLAX Example** The following examples demonstrate the **trace** statement syntax. The first example writes a concatenated string to the trace file. The second example uses a code block to output a **<max>** element and a **<min>** element and the values of the **max** and **min** variables. The third example uses a conditional statement to specify when to output trace data. If the expression evaluates to **true**, the code block writes the string and the **<options>** element hierarchy to the trace file.

```
trace "max " _ $max _ "; min " _ $min;

trace {
    <max> $max;
    <min> $min;
}
```

```
trace {
  if ($my-trace-flag) {
    expr "max " _ $max _ "; min " _ $min;
    copy-of options;
  }
}
```

**Related  
Documentation**

- [message on page 109](#)
- [terminate on page 126](#)
- [Understanding the SLAX Processor \(slaxproc\) on page 148](#)

---

## uexpr

---

**Syntax**      `uexpr expression;`

**Release Information**    Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description**            Generate the string value of an XPath expression and add it to the result tree, but do not escape special characters. The **uexpr** statement behaves identically to the **expr** statement, except that the contents are not escaped. By default, characters such as "<", ">", and "&" are escaped into proper XML as "&lt;", "&gt;", and "&amp;", respectively, but **uexpr** does not execute this escaping mechanism.

**Attributes**            *expression*—XPath expression to add to the result tree.

**SLAX Example**           The following statement outputs the string to the result tree exactly as it appears in the statement. If **expr** is used in place of **uexpr**, the script would output the string "&lt;:&amp;&gt;".

```
uexpr "<: -&gt;";
```

**Related  
Documentation**

- [expr on page 99](#)
- [message on page 109](#)
- [terminate on page 126](#)
- [trace on page 127](#)

---

## use-attribute-sets

---

**Syntax**      `use-attribute-sets attribute-set-name;`

**Release Information**    Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description**            Add the attributes in the attribute set to the current element. The **use-attribute-sets** statement can be used under the **attribute-set**, **copy-node**, and **element** statements, as well as under a normal element.



<b>Attributes</b>	<b><i>attribute-set-name</i></b> —Name of the attribute set, which is defined using an <b>attribute-set</b> statement.
<b>SLAX Example</b>	<p>The following example creates two attribute sets: <b>table-attributes</b> and <b>table-attributes-ext</b>. The <b>table-attributes-ext</b> set includes all of the attributes that are already defined in the <b>table-attributes</b> set through use of the <b>use-attributes-sets</b> statement. In the main script body, the <b>table-attributes-ext</b> attribute set is applied to the <b>&lt;table&gt;</b> element. The <b>&lt;table&gt;</b> element includes the four attributes: <b>order</b>, <b>cellpadding</b>, <b>cellspacing</b>, and <b>border</b>.</p> <pre> version 1.1;  var \$cellpadding = "0"; var \$cellspacing = "10";  attribute-set table-attributes {   attribute "order" { expr "0"; }   attribute "cellpadding" { expr \$cellpadding; }   attribute "cellspacing" { expr \$cellspacing; } } attribute-set table-attributes-ext {   use-attribute-sets table-attributes;   attribute "border" { expr "0"; } }  match / {   ...   &lt;table&gt; {     use-attribute-sets table-attributes-ext;   } } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SLAX Elements and Element Attributes Overview on page 61</a></li> <li>• <a href="#">attribute on page 90</a></li> <li>• <a href="#">attribute-set on page 91</a></li> <li>• <a href="#">element on page 96</a></li> </ul>

## var

<b>Syntax</b>	<b>var</b> <i>\$name=</i> <i>value</i> ;
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.
<b>Description</b>	Declare a local or global variable. A variable is global if it is defined outside of any template. Otherwise, it is local. The value of a global variable is accessible anywhere in the style sheet. The scope of a local variable is limited to the template or code block in which it is defined. Variables declared in this manner are immutable. You initialize a variable by following the variable name with an equal sign (=) and an expression.
<b>Attributes</b>	<b><i>name</i></b> —Specifies the name of the variable. After declaration, the variable can be referred to within expressions using this name, including the \$ character.

**value**—Defines the default value for the variable, which is used if the person or client application that executes the script does not explicitly provide a value.

SLAX Example	<pre>var \$vrf; var \$location = \$dot/@location; var \$message = "We are in "_ \$location _" now.";</pre>
XSLT Equivalent	<pre>&lt;xsl:variable name="vrf"/&gt; &lt;xsl:variable name="location" select="\$dot/location"/&gt; &lt;xsl:variable name="message" select="concat('We are in ', \$location, now.)"/&gt;</pre>
Usage Examples	See <a href="#">“Example: Limiting the Number of E1 Interfaces” on page 388</a> , <a href="#">“Example: Limiting the Number of ATM Virtual Circuits” on page 384</a> , <a href="#">“Example: Configuring Administrative Groups for LSPs” on page 350</a> , and <a href="#">“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343</a> .
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">SLAX Variables Overview on page 76</a></li><li>• <a href="#">mvar on page 111</a></li><li>• <a href="#">param on page 118</a></li></ul>

---

## version

---

Syntax	<code>version 1.0;</code>
Release Information	Statement introduced in version 1.0 of the SLAX language.
Description	<p>Specify the version of SLAX that is being used. All SLAX style sheets must begin with a <b>version</b> statement.</p> <p>Version 1.0 uses XML version 1.0 and XSLT version 1.1.</p> <p>In addition, the <code>xsl</code> namespace is implicitly defined as follows:</p> <pre>xmlns:xsl="http://www.w3.org/1999/XSL/Transform"</pre>
Attributes	<b>version-number</b> —Specifies the version of SLAX.
SLAX Example	<code>version 1.0;</code>
XSLT Equivalent	<code>&lt;xsl:stylesheet version="1.0"&gt;</code>
Usage Examples	<ul style="list-style-type: none"><li>• <a href="#">Example: Adding a Final then accept Term to a Firewall on page 331</a></li><li>• <a href="#">Example: Changing the Configuration Using SLAX and XSLT Op Scripts on page 467</a></li><li>• <a href="#">Example: Assigning a Classifier on page 339</a></li><li>• <a href="#">Example: Imposing a Minimum MTU Setting on page 382</a></li><li>• <a href="#">Example: Restarting an FPC Using an Op Script on page 495</a></li></ul>

- Related Documentation**
- [Required Boilerplate for Commit Scripts on page 251](#)
  - [Required Boilerplate for Event Scripts on page 615](#)
  - [Required Boilerplate for Op Scripts on page 445](#)
  - [SLAX Syntax Rules Overview on page 59](#)

## while

<b>Syntax</b>	<pre>while (<i>expression</i>) {     /* body */ }</pre>
<b>Release Information</b>	Statement introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Repeatedly execute a block of statements until the specified condition evaluates to <b>false</b> . The condition is an XPath expression that is converted to a boolean type. If the expression evaluates to <b>true</b> , the contents of the while loop are executed. The loop continues to execute until the expression evaluates to <b>false</b> . During execution, the context is not changed. In the expression, you should use a mutable variable, which is declared using the <b>mvar</b> statement, to avoid creating an infinite loop.
<b>Attributes</b>	<b>expression</b> —XPath expression, which is cast to boolean type and used as the condition for the while loop. The code block contents are executed as long as the condition evaluates to <b>true</b> .
<b>SLAX Example</b>	<p>In the example, the while loop parses through the item list until the desired value is found. When that value is detected, <b>\$seen</b> is set to true, and the while loop exits.</p> <pre>mvar \$seen = false(); mvar \$count = 1;  while (not(\$seen)) {     if (item[\$count]/value) {         set \$seen = true();     }     set \$count = \$count + 1; }</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SLAX Variables Overview on page 76</a></li> <li>• <a href="#">XPath Overview on page 19</a></li> <li>• <a href="#">mvar on page 111</a></li> </ul>

## with

<b>Syntax</b>	<code>with \$name = value;</code>
<b>Release Information</b>	Statement introduced in version 1.0 of the SLAX language.

**Description** Specify a parameter to pass into a template. You can use this statement when you apply templates with the **apply-templates** statement or invoke templates with the **call** statement.

Optionally, you can specify a value for the parameter by including an equal sign (=) and a value expression. If no value is specified, the current value of the parameter is passed to the template.

**Attributes** *name*—Name of the variable or parameter for which the value is being passed.

*value*—Value of the parameter being passed to the template.

**SLAX Example**

```
match configuration {
    var $domain = domain-name;
    apply-templates system/host-name {
        with $message = "Invalid host-name";
        with $domain;
    }
}

match host-name {
    param $message = "Error";
    param $domain;
    <hello> $message _ ":: " _ . _ " (" _ $domain _ ")";
}
```

**XSLT Equivalent**

```
<xsl:template match="configuration">
    <xsl:apply-templates select="system/host-name">
        <xsl:with-param name="message" select="'Invalid host-name'"/>
        <xsl:with-param name="domain" select="$domain"/>
    </xsl:apply-templates>
</xsl:template>

<xsl:template match="host-name">
    <xsl:param name="message" select="'Error'"/>
    <xsl:param name="domain"/>
    <hello>
        <xsl:value-of select="concat($message, ':: ', ' (' , $domain, ')')"/>
    </hello>
</xsl:template>
```

**Usage Examples** See [“Example: Configuring Dual Routing Engines” on page 357](#), [“Example: Preventing Import of the Full Routing Table” on page 416](#), and [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

**Related Documentation**

- [SLAX Parameters Overview on page 72](#)
- [SLAX Templates Overview on page 66](#)
- [apply-templates on page 89](#)
- [call on page 93](#)
- [match on page 108](#)
- [mode on page 109](#)

- [priority on page 119](#)
- [template on page 125](#)



## CHAPTER 8

# Using the libslax Distribution for Automation Scripting

- [libslax Distribution Overview on page 135](#)
- [libslax Library and Extension Libraries Overview on page 136](#)
- [Downloading and Installing the libslax Distribution on page 137](#)
- [libslax Default Extension Libraries: bit, cURL, and xutil on page 138](#)
- [Understanding the SLAX Processor \(slaxproc\) on page 148](#)
- [Using the SLAX Processor \(slaxproc\) on page 153](#)
- [SLAX Debugger, Profiler, and callflow on page 158](#)

### libslax Distribution Overview

---

SLAX is an alternative syntax for XSLT and is tailored for readability and familiarity, following the style of C and Perl. In the SLAX language, programming constructs and XPath expressions are moved from the XML elements and attributes used in XSLT to first class language constructs. SLAX was originally developed as part of Junos OS for the purpose of on-box scripting to allow users to customize and enhance the command-line interface (CLI).

libslax is an open-source implementation of the SLAX language using the "New BSD License". libslax is written in C and is built on top of the libxml2, libxslt, and libexslt libraries. The libslax distribution contains the libslax library, which incorporates a SLAX writer and SLAX parser, a debugger, a profiler, and the SLAX processor (slaxproc). The SLAX processor is a command-line tool that can validate SLAX script syntax, convert between SLAX and XSLT formats, and format, debug, or run SLAX scripts.

The libslax tools are included as part of the standard Junos OS. However, you can download and install the libslax distribution on a computer with a UNIX-like operating system to develop SLAX scripts outside of Junos OS. Current releases, source code, additional documentation, and support materials for libslax are available at:

<https://github.com/Juniper/libslax>

The SLAX community and support site is available at:

<https://github.com/Juniper/libslax/wiki>

**Related Documentation**

- [Downloading and Installing the libslax Distribution on page 137](#)
- [SLAX Overview on page 57](#)
- [Understanding the SLAX Processor \(slaxproc\) on page 148](#)
- [libslax Library and Extension Libraries Overview on page 136](#)
- [Using the SLAX Processor \(slaxproc\) on page 153](#)

---

## libslax Library and Extension Libraries Overview

- [libslax Library on page 136](#)
- [libslax Extension Libraries on page 136](#)

### libslax Library

libslax is an open-source implementation of the SLAX language using the "New BSD License". libslax is written in C and is built on top of the libxml2, libxslt, and libexslt libraries.

The core of the libslax distribution is the libslax library, which incorporates a SLAX parser to read SLAX files, a SLAX writer to write SLAX files, a debugger, a profiler, and the SLAX processor (slaxproc) command-line tool. The parser turns a SLAX source file into an XSLT tree (xmlDocPtr) using the `xsltSetLoaderFunc()` hook. The writer turns an XSLT tree (xmlDocPtr) into a file containing SLAX statements.

### libslax Extension Libraries

libslax supports a way to dynamically load extension libraries. The libslax distribution includes the xutil, bit, and cURL extension libraries. The source files for the default extension libraries are stored in the **libslax-release/extensions** directory of the distribution. You can supply additional extension libraries beyond the default extension libraries supported by the libslax distribution. Extension library locations can be specified statically at build time or dynamically at run time.

By default, libslax installs architecture-independent files, including extension library files, in the **/usr/local** directories. Specifically, libslax installs the extension libraries in the **/usr/local/lib/slax/extensions** directory. If you do not specify a different installation directory for the extension libraries at build time, the SLAX processor checks this directory for extension libraries when executing a script.

There are several ways to specify extension library locations at build time. During installation, to specify a directory prefix other than **/usr/local** for all installation files, including the libraries, execute the **./configure** command and include the **--prefix=prefix** option specifying the location to install the files. The default extension libraries are installed in the **prefix/lib/slax/extensions** directory, and the SLAX processor checks this directory for extension libraries when executing a script. To install just the extension library files in a different, user-defined location, execute the **./configure** command and include the **--with-extensions-dir=dir** option specifying the location where the extension libraries live. The SLAX processor will then automatically check the specified directory for extension libraries when executing a script. For more information about installing libslax, see ["Downloading and Installing the libslax Distribution" on page 137](#).



There are several ways to specify extension library locations dynamically after installation is complete. You can define or update the `SLAXEXTPATH` environment variable to include the directory locations of additional extension libraries. The variable value is a colon-separated list of directories. The SLAX processor automatically checks these directories for extension libraries when executing a script. Alternatively, you can specify the extension library location when you execute a script by using the `slaxproc` command with the `--lib` or `-L` option.

To summarize, extension library locations are supplied to the SLAX processor in one of the following ways:

- By default, in the `/usr/local/lib/slax/extensions` directory.
- In `lib/slax/extensions/` under the directory specified by the `./configure --prefix` option given at build time.
- In the user-defined directory specified by the `./configure --with-extension-dir` option given at build time.
- In a directory included in the colon-separated list of the `SLAXEXTPATH` environment variable.
- In a directory provided using the `--lib` or `-L` argument to the `slaxproc` command.

**Related  
Documentation**

- [libslax bit Extension Library on page 138](#)
- [libslax cURL Extension Library on page 140](#)
- [Downloading and Installing the libslax Distribution on page 137](#)
- [libslax Distribution Overview on page 135](#)

---

## Downloading and Installing the libslax Distribution

The libslax distribution contains the libslax library, which incorporates a SLAX writer and SLAX parser, a debugger, a profiler, and the SLAX processor (`slaxproc`). The SLAX processor is a command-line tool that can validate SLAX script syntax, convert between SLAX and XSLT formats, and format, debug, or run SLAX scripts.

The libslax tools are included as part of the standard Junos OS. However, you can download and install the libslax distribution on a computer with a UNIX-like operating system to develop SLAX scripts outside of Junos OS.

To download and install the libslax distribution, follow the instructions on the libslax wiki.

<https://github.com/Juniper/libslax/wiki/Building>

Review the **INSTALL** file that comes with the distribution for additional information.

**Related  
Documentation**

- [libslax Distribution Overview on page 135](#)
- [libslax Library and Extension Libraries Overview on page 136](#)
- [Understanding the SLAX Processor \(`slaxproc`\) on page 148](#)

## libslax Default Extension Libraries: bit, cURL, and xutil

- [libslax bit Extension Library on page 138](#)
- [libslax cURL Extension Library on page 140](#)
- [libslax xutil Extension Library on page 148](#)

### libslax bit Extension Library

The bit extension library contains functions that create and manipulate bit strings. The functions support 64-bit integer arguments. To incorporate functions from the bit extension library into SLAX scripts, include the namespace statement for that library in the script.

```
ns bit extension = "http://xml.libslax.org/bit";
```

Call the bit extension functions using the **bit** prefix and the function name. For example:

```
version 1.1 ;
ns bit extension = "http://xml.libslax.org/bit";

var $a = 63;
var $b = { expr "10111"; }

match / {
  <out> {
    <bit-and> {
      <a1> bit:and("101100", "100101");
      <a2> bit:and($a, $b);
      <a3> bit:and($a, number($b));
    }
    <bit-or> {
      <a1> bit:or("101100", "100101");
      <a2> bit:or($a, $b);
      <a3> bit:or($a, number($b));
    }
    <bit-mask> {
      <a1> bit:mask(0);
      <a2> bit:mask(8, 32);
    }
    <ops> {
      <a1> bit:to-int("10101");
    }
  }
}
```

[Table 10 on page 138](#) lists the functions available in the bit extension library and which are supported in SLAX 1.1 scripts.

**Table 10: Functions in the bit Extension Library**

Function and Arguments	Description	Example
bit:and(b1, b2)	Return the logical AND of two bit strings.	bit:and("101100", "100101") return value: "100100"

Table 10: Functions in the bit Extension Library (*continued*)

Function and Arguments	Description	Example
<code>bit:clear(b1,bitnum)</code>	Set the specified bit in the bit string to zero and return the new bit string. Bits are numbered starting from zero. If the integer argument is greater than the bit string length, the bit string is extended.	<code>bit:clear("1111", 0)</code> return value: "1110"  <code>bit:clear("1111", 6)</code> return value: "001111"
<code>bit:compare(value1, value2)</code>	Compare two values and return an integer less than, equal to, or greater than zero, if the first argument is found to be less than, equal to, or greater than the second argument, respectively.	<code>bit:compare("10000", 16)</code> return value: 0  <code>bit:compare("1111", "10000")</code> return value: 1
<code>bit:from-hex(string, len?)</code>	Return the value of the hex argument as a bit string. The optional second argument pads the bit string with leading 0's until it is the specified length.	<code>bit:from-hex("0x45", 8)</code> return value: "01000101"
<code>bit:from-int(integer, len?)</code>	Return the value of the integer argument as a bit string. The optional second argument pads the bit string with leading 0's until it is the specified length.	<code>bit:from-int(65,8)</code> return value: "01000001"
<code>bit:mask(count, len?)</code>	Return a bit string with count low-order bits set to one. The optional second argument pads the bit string with leading 0's until it is the specified length.	<code>bit:mask(4, 8)</code> return value: "00001111"
<code>bit:nand(b1, b2)</code>	Return the logical NAND of two bit strings.	<code>bit:nand("101100", "100101")</code> return value: "010010"
<code>bit:nor(b1, b2)</code>	Return the logical NOR of two bit strings.	<code>bit:nor("101100", "100101")</code> return value: "011011"
<code>bit:not(b1)</code>	Return the inversion (NOT) of a bit string.	<code>bit:not("101100")</code> return value: "010011"
<code>bit:or(b1, b2)</code>	Return the logical OR of two bit strings.	<code>bit:or("101100", "100101")</code> return value: "101101"
<code>bit:set(b1, bitnum)</code>	Set the specified bit in the bit string and return the new bit string. Bits are numbered starting from zero. If the integer argument is greater than the bit string length, the bit string is extended.	<code>bit:set("1001", 2)</code> return value: "1101"  <code>bit:set("1001", 6)</code> return value: "1001001"

Table 10: Functions in the bit Extension Library (*continued*)

Function and Arguments	Description	Example
bit:to-int(b1)	Return the value of the bit string argument as an integer.	bit:to-int("101100") return value: 44
bit:to-hex(b1)	Return the value of the bit string argument as a string representation of the hex value.	bit:to-hex("101100") return value: "0x2c"
bit:xor(b1, b2)	Return the logical XOR of two bit strings.	bit:xor("101100", "100101") return value: "001001"
bit:xnor(b1, b2)	Return the logical XNOR of two bit strings.	bit:xnor("101100", "100101") return value: "110110"

## libslax cURL Extension Library

- [Understanding the cURL Extension Library on page 140](#)
- [curl:close on page 143](#)
- [curl:open on page 143](#)
- [curl:perform on page 144](#)
- [curl:set on page 145](#)
- [curl:single on page 145](#)
- [Supported Format Elements on page 145](#)
- [cURL Examples on page 146](#)

### Understanding the cURL Extension Library

cURL is a command-line tool that uses the libcurl library and supports data transfers using a number of protocols, including FTP, FTPS, HTTP, HTTPS, SCP, and SMTP. For more information about cURL, see the cURL website at <http://curl.haxx.se/>.

To incorporate functions from the cURL extension library into SLAX scripts, include the namespace statement for that library in the script.

```
ns curl extension = "http://xml.libslax.org/curl";
```

Call the cURL extension functions using the **curl** prefix and the function name. cURL operations are directed using a set of elements passed to the cURL extension functions.

[Table 11 on page 141](#) lists the supported operations in the cURL extension library and includes a description of each function. [Table 12 on page 141](#) and [Table 13 on page 142](#) list the supported elements and include the syntax and a description of each element. [Table 12 on page 141](#) lists elements used for web services operations, and [Table 13 on page 142](#) lists the elements used for e-mail operations.

Table 11: Functions in the cURL Extension Library

Function	Description
curl:close	Close an open connection. Further operations cannot be performed over the connection. See <a href="#">"curl:close" on page 143</a> .
curl:open	Open a connection to a remote server, allowing multiple operations over a single connection. See <a href="#">"curl:open" on page 143</a> .
curl:perform	Perform simple transfers using a persistent connection handle provided by curl:open. See <a href="#">"curl:perform" on page 144</a> .
curl:set	Record a set of parameters that persists for the lifespan of a connection. See <a href="#">"curl:set" on page 145</a> .
curl:single	Perform transfer operations without using a persistent connection. See <a href="#">"curl:single" on page 145</a> .

Table 12: Web Services Elements in the cURL Extension Library

Element	Description	Syntax
<content-type>	Provide the MIME type for the transfer payload.	<content-type> "mime/type";
<fail-on-error>	Indicate that the transfer should fail if any errors, including insignificant ones, are detected.	<fail-on-error>;
<format>	Specify the expected format of returned results, allowing the cURL extension to automatically make the content available in the native format. Formats include "html", "name", "text", "url-encoded", and "xml".	<format> "xml";
<header>	Provide additional header fields for the request.	<header name="name"> "value";
<insecure>	Indicate a willingness to tolerate insecure communications operations. Specifically, allow SSL Certs without checking the common name.	<insecure>;
<method>	Set the method used to transfer data. This controls the HTTP request type, as well as triggering other transfer mechanisms.  Acceptable method names include "get", "post", "delete", "head", "email", "put", and "upload". The "get" method is the default.	<method> "get";

Table 12: Web Services Elements in the cURL Extension Library (*continued*)

Element	Description	Syntax
<param>	Provide additional parameter values for the request. These parameters are typically encoded into the URL.	<param name="x"> "y";
<password>	Set the user's password for the transfer.	<password> "password";
<secure>	Request the use of the secure version of a protocol, including HTTPS and FTPS.	<secure>;
<upload>	Indicate this is a file upload request.	<upload>;
<url>	Set the base URL for the request.	<url> "target-url";
<username>	Set the username to use for the transfer.	<username> "username";
<verbose>	Request detailed debug information about the operations and communication of the cURL transfer.	<verbose>;

Table 13: E-mail Elements in the cURL Extension Library

Element	Description	Syntax
<cc>	Set the "Cc" address for e-mail (SMTP) requests. For multiple addresses, use multiple <cc> elements.	<cc> "cc-user@email.example.com";
<contents>	Specify the contents to be transferred.	<contents> "multi-\nline\ncontents\n";
<from>	Set the "From" address for e-mail (SMTP) requests.	<from> "source-user@email.example.com";
<header>	Provide additional header fields for the request.	<header name="name"> "value";
<local>	Set the local hostname for e-mail (SMTP) requests.	<local> "local host name";
<server>	Set the outgoing SMTP server name. Currently, MX records are not processed.	<server> "email-server.example.com";
<subject>	Set the "Subject" field for e-mail (SMTP) requests.	<subject> "email subject string";

Table 13: E-mail Elements in the cURL Extension Library (*continued*)

Element	Description	Syntax
<to>	Set the "To" address for e-mail (SMTP) requests. For multiple addresses, use multiple <to> elements.	<to> "to-user@email.examplecom";
<verbose>	Request detailed debug information about the operations and communication of the cURL transfer.	<verbose>;

The libcurl elements closely mimic the options used by the native C libcurl API in libcurl's `curl_easy_setopt()` function. Once the options are set, a call to `curl_easy_perform()` performs the requested transfer. For more information about the `curl_easy_setopt()` function, see [http://curl.haxx.se/libcurl/c/curl\\_easy\\_setopt.html](http://curl.haxx.se/libcurl/c/curl_easy_setopt.html).

In the SLAX cURL extension library, the libcurl API options are represented as individual elements. For example, the <url> element is mapped to the CURLOPT\_URL option, the <method> element is mapped to the CURLOPT\_CUSTOMREQUEST option, and so forth.

These elements can be used in three ways:

- The `curl:single()` extension function supports using a set of options in a single transfer operation with no persistent connection handle.
- The `curl:perform()` extension function supports using a set of options with a persistent connection handle. The handle is returned from the `curl:open()` extension function and can be closed with the `curl:close()` extension function.
- The `curl:set()` extension function records a set of options for a connection handle and keeps those options active for the lifetime of the connection. For example, if the script needs to transfer a number of files, it can record the <username> and <password> options and avoid repeating them in every `curl:perform()` call.

### curl:close

The `curl:close()` extension function closes an open connection. Further operations cannot be performed over the connection once it is closed.

The syntax is:

```
node-set[empty] curl:close(node-set[connection]);
```

The argument is the connection handle to close.

### curl:open

The `curl:open()` extension function opens a connection to a remote server, allowing multiple operations over a single connection.

The syntax is:

```
node-set[connection] curl:open();
```

The returned object is a connection handle that can be passed to **curl:perform()** or **curl:close()**.

### curl:perform

The **curl:perform()** extension function performs simple transfers using a persistent connection handle provided by **curl:open()**.

The syntax is:

```
node-set[object] curl:perform(node-set[connection], node-set[options])
```

The arguments are the connection handle and a set of option elements. Supported cURL extension library elements are defined in [Table 12 on page 141](#) and [Table 13 on page 142](#).

The returned object is an XML hierarchy containing the results of the transfer. [Table 14 on page 144](#) lists the possible elements in the reply, and [Table 15 on page 144](#) lists the possible elements contained within the **<header>** element.

**Table 14: curl:perform Reply Elements**

Element	Contents
<curl-success>	Empty element which Indicates success
<data>	Parsed data
<error>	Error message text, if any
<header>	Parsed header fields
<raw-data>	Raw data from the reply
<raw-headers>	Raw header fields from the reply
<url>	Requested URL

**Table 15: curl:perform <header> Elements**

Element	Contents
<code>	HTTP reply code
<field>	HTTP reply field (with @name and value)
<message>	HTTP reply message
<version>	HTTP reply version string



The following example shows the `<header>` element with header fields parsed into `<field>` elements:

```
<header>
  <version>HTTP/1.1</version>
  <code>404</code>
  <message>Not Found</message>
  <field name="Content-Type">text/html</field>
  <field name="Content-Length">345</field>
  <field name="Date">Mon, 08 Aug 2011 03:40:21 GMT</field>
  <field name="Server">lighttpd/1.4.28 juisebox</field>
</header>
```

### curl:set

The `curl:set()` extension function records a set of parameters that persist for the lifespan of a connection.

The syntax is:

```
node-set[empty] curl:set(node-set[handle], node-set[options]);
```

The arguments are the connection handle and a set of option elements. Supported cURL extension library elements are defined in [Table 12 on page 141](#) and [Table 13 on page 142](#).

### curl:single

The `curl:single()` extension function performs transfer operations without using a persistent connection.

The syntax is:

```
node-set[result] curl:single(node-set[options]);
```

The returned object is identical in structure to the one returned by `curl:perform()`. Refer to ["curl:perform" on page 144](#) for additional information.

### Supported Format Elements

Starting with version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases, the **format** element supports two new formats in addition to "html", "text", and "xml": "name" and "url-encoded". The "name" encoding is used for name=value pairs that are separated by newlines, whereas the "url-encoded" encoding is used when the name=value pairs are separated by an ampersand (&).

The parsed data is returned in the `<data>` element, using `<name>` elements. In the following example, **format** is set to "url-encoded":

```
<results>
  <url>https://api.example.com/request_token</url>
  <curl-success/>
  <raw-headers>HTTP/1.1 200 OK&#xD;
  Server: XXXX&#xD;
  Date: Tue, 18 Jun 2013 18:56:31 GMT&#xD;
  Content-Type: application/x-www-form-urlencoded&#xD;
  Transfer-Encoding: chunked&#xD;
  Connection: keep-alive&#xD;
  x-server-response-time: 69&#xD;
  x-example-request-id: 123456&#xD;
```

```

pragma: no-cache&#xD;
cache-control: no-cache&#xD;
x-http-protocol: None&#xD;
x-frame-options: SAMEORIGIN&#xD;
X-RequestId: 12345&#xD;
&#xD;
</raw-headers>
<headers>
  <version>HTTP/1.1</version>
  <code>200</code>
  <message>OK</message>
  <header name="Server">XXXXX</header>
  <header name="Date">Tue, 18 Jun 2013 18:56:31 GMT</header>
  <header name="Content-Type"
    >application/x-www-form-urlencoded</header>
  <header name="Transfer-Encoding">chunked</header>
  <header name="Connection">keep-alive</header>
  <header name="x-server-response-time">69</header>
  <header name="x-example-request-id">123456</header>
  <header name="pragma">no-cache</header>
  <header name="cache-control">no-cache</header>
  <header name="x-http-protocol">None</header>
  <header name="x-frame-options">SAMEORIGIN</header>
  <header name="X-RequestId">12345</header>
</headers>
<raw-data>oauth_token_secret=s&amp;oauth_token=t</raw-data>
<data format="url-encoded">
  <name name="oauth_token_secret">s</name>
  <name name="oauth_token">t</name>
</data>
</results>

```

### cURL Examples

The following examples show SLAX version 1.1 scripts that use the cURL extension library functions to perform operations.

The following SLAX script performs a simple GET operation to retrieve a web page. The script specifies the header field for the HTTP header and a parameter that is incorporated into the requested URL.

```

version 1.1;

ns curl extension = "http://xml.libslax.org/curl";

param $url = "http://www.juniper.net";

match / {
  <op-script-results> {
    var $options = {
      <header name="client"> "slaxproc";
      <param name="smokey"> "bandit";
    }
    var $results = curl:single($url, $options);
    message "completed: " _ $results/headers/message;
    <curl> {
      copy-of $results;
    }
  }
}

```

The following SLAX script takes a username and password and uses the Google login services to translate them into an "Authorization" string:

```
version 1.1;

ns curl extension = "http://xml.libslax.org/curl";

param $url = "https://www.google.com/accounts/ClientLogin";
param $username;
param $password;

var $auth-params := {
  <url> $url;
  <method> "post";
  <insecure>;
  <param name="Email"> $username;
  <param name="Passwd"> $password;
  <param name="accountType"> "GOOGLE";
  <param name="service"> "wise";
  <param name="source"> "test-app";
}

match / {
  var $curl = curl:open();
  var $auth-cred = curl:perform($curl, $auth-params);

  <options> {
    for-each(slx:break-lines( $auth-cred/raw-data )) {
      if(starts-with(., "Auth")) {
        <header name="GData-Version"> "3.0";
        <header name="Authorization"> "GoogleLogin " _ .;
      }
    }
  }
  expr curl:close($curl);
}
```

The following SLAX script sends an e-mail by way of a server, which is provided as a parameter:

```
version 1.1;

ns curl extension = "http://xml.libslax.org/curl";

param $server;

match / {
  <out> {
    var $info = {
      <method> "email";
      <server> $server;
      <from> "muffin@example.com";
      <to> "phil@example.net";
      <subject> "Testing...";
      <contents> "Hello,
This is an email.
Thanks,
Phil
";
    }
    var $res = curl:single($info);
  }
}
```

```

        <res> {
            copy-of $res;
        }
    }
}

```

## libslax xutil Extension Library

The xutil extension library contains functions that convert between strings and XML node sets. To incorporate functions from the xutil extension library into SLAX scripts, include the namespace statement for that library in the script.

```
ns xutil extension = "http://xml.libslax.org/xutil";
```

Call the xutil extension functions using the **xutil** prefix and the function name.

[Table 16 on page 148](#) lists the functions available in the xutil extension library and which are supported in SLAX 1.1 scripts.

**Table 16: Functions in the xutil Extension Library**

Function and Arguments	Description
<b>xutil:max-call-depth(number)</b>	Change the limit on the depth of recursive calls. The default limit is 3,000.  The <b>xutil:max-call-depth()</b> function was added to version 1.1 of the SLAX language in Junos OS Release 14.1.
<b>xutil:string-to-xml(string+)</b>	Return a node set of the concatenated string arguments.
<b>xutil:xml-to-string(node-set+)</b>	Return a string representation of the XML hierarchies provided as arguments.

### Related Documentation

- [libslax Library and Extension Libraries Overview on page 136](#)
- [libslax Distribution Overview on page 135](#)
- [Downloading and Installing the libslax Distribution on page 137](#)
- [Understanding the SLAX Processor \(slaxproc\) on page 148](#)
- [Using the SLAX Processor \(slaxproc\) on page 153](#)
- [SLAX Debugger, Profiler, and callflow on page 158](#)

## Understanding the SLAX Processor (slaxproc)

- [slaxproc Overview on page 149](#)
- [slaxproc Modes on page 149](#)
- [slaxproc Options on page 150](#)
- [slaxproc File Argument Handling on page 152](#)
- [slaxproc UNIX Scripting Support on page 153](#)

## slaxproc Overview

The libslax distribution contains the libslax library, which incorporates a SLAX writer and SLAX parser, a debugger, a profiler, and the SLAX processor (slaxproc). The SLAX processor is a command-line tool that can validate SLAX script syntax, convert between SLAX and XSLT formats, and format, debug, or run SLAX scripts.

The SLAX processor is invoked on the command line using the **slaxproc** command. The **slaxproc** command accepts command-line arguments that specify the mode of the processor, any behavioral options, and required input and output files.

The syntax for the **slaxproc** command is:

```
slaxproc [mode] [options] [script] [files]
```

The slaxproc mode defines what function the processor performs. slaxproc options include file options and common options. File options are used to specify the script file, input file, output file, and trace file. Common options include additional functionality provided by the SLAX processor such as verbose debugging output.

You can access the slaxproc help by issuing the **slaxproc** command with the **--help** or **-h** option.

```
$ slaxproc -h
```

## slaxproc Modes

The slaxproc mode defines what function the processor performs. The default mode is **--run** or **-r**. If you do not explicitly specify a mode, the SLAX processor executes a script. [Table 17 on page 149](#) outlines the slaxproc modes and their functions.

**Table 17: Slaxproc Modes**

Mode	Description
<b>--check</b> <b>-c</b>	Perform a syntax and content check on a SLAX script, reporting any errors. This mode is useful for off-box syntax checks before installing or uploading scripts to a device running Junos OS.
<b>--format</b> <b>-F</b>	Format a SLAX script, correcting indentation and spacing to the preferred style.
<b>--json-to-xml</b>	<p>Transform JSON input into XML using the conventions defined in <a href="http://juniper.github.io/libslax/slax-manual.html#json-elements">http://juniper.github.io/libslax/slax-manual.html#json-elements</a>.</p> <p>The <b>--json-to-xml</b> slaxproc mode is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--run</b> <b>-r</b>	Run a SLAX script. This is the default mode. The script name, input filename, and output filename can be provided using command-line options, positional arguments, or a mix of both. Input defaults to standard input, and output defaults to standard output.

Table 17: Slaxproc Modes (*continued*)

Mode	Description
<b>--show-select</b>	<p>Show an XPath selection from the input document. This mode is used to extract out selections from a script for external consumption. This enables the consumer to avoid a SLAX parser but still have visibility into the contents of the script.</p> <p>The <b>--show-select</b> slaxproc mode is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--show-variable</b>	<p>Show the contents of a global variable. This mode is used to extract static variable contents for external consumption. This enables the consumer of the data to avoid a SLAX parser but still have access to the static contents of global variables, such as the <b>\$arguments</b> variable.</p> <p>The <b>--show-variable</b> slaxproc mode is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--slax-to-xslt</b> <b>-x</b>	<p>Convert a SLAX script into XSLT format. The script filename and output filename are provided using command-line options, positional arguments, or a mix of both.</p>
<b>--xml-to-json</b>	<p>Convert XML into JSON using the conventions defined in <a href="http://juniper.github.io/libslax/slax-manual.html#json-elements">http://juniper.github.io/libslax/slax-manual.html#json-elements</a>.</p> <p>The <b>--xml-to-json</b> slaxproc mode is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--xpath <i>xpath</i></b> <b>-X <i>xpath</i></b>	<p>Select data matching an XPath expression from the input document.</p> <p>The <b>--xpath</b> and <b>-X</b> slaxproc modes are available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--xslt-to-slax</b> <b>-s</b>	<p>Convert an XSLT script into SLAX format. The script filename and output filename are provided using command-line options, positional arguments, or a mix of both.</p>

## slaxproc Options

The slaxproc options include file options and common options. File options are used specify the script file, input file, output file, and trace file. Common options include additional functionality and options provided by the SLAX processor such as verbose debugging output. [Table 18 on page 150](#) lists the slaxproc common options and file options.

Table 18: slaxproc Common Options and File Options

Option	Description
<b>--debug</b> <b>-d</b>	Enable the SLAX/XSLT debugger.
<b>--empty</b> <b>-E</b>	Provide an empty document as the input data set. This is useful for scripts that do not expect or need meaningful input.

Table 18: slaxproc Common Options and File Options (*continued*)

Option	Description
<b>--expression &lt;expr&gt;</b>	<p>Convert a SLAX expression to an XPATH expression, or vice versa, depending on the presence of <b>--slax-to-xslt</b> or <b>--xslt-to-slax</b>.</p> <p>The <b>--expression</b> slaxproc option is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--exslt</b> <b>-e</b>	Enables the EXSLT library, which provides a set of standard extension functions. See <a href="http://www.exslt.org">http://www.exslt.org</a> for more information.
<b>--help</b> <b>-h</b>	Displays the help message and exit.
<b>--html</b> <b>-H</b>	Parse input data using the HTML parser, which differs from XML.
<b>--include &lt;dir&gt;</b> <b>-I &lt;dir&gt;</b>	Add a directory to the list of directories searched when using include and import files. Alternatively, you can define the SLAXPATH environment variable to specify a colon-delimited list of directories to search.
<b>--indent</b> <b>-g</b>	Indent output. This option is identical to the behavior triggered by <b>output-method { indent 'true'; }</b> .
<b>--input &lt;file&gt;</b> <b>-i &lt;file&gt;</b>	Read input from the specified file.
<b>--json-tagging</b>	<p>Tag JSON elements with the 'json' attribute as the input is converted to XML. This enables you to use the <b>--format</b> mode to transform the data back into JSON format.</p> <p>The <b>--json-tagging</b> slaxproc option is available starting in version 1.2 of the SLAX language, which is supported in Junos OS Release 14.2 and later releases.</p>
<b>--lib &lt;dir&gt;</b> <b>-L &lt;dir&gt;</b>	Adds a directory to the list of directories searched when using extension libraries. Alternatively, you can define the SLAXEXTPATH environment variable to specify a colon-delimited list of extension library locations to search.
<b>--name &lt;file&gt;</b> <b>-n &lt;file&gt;</b>	Reads the SLAX script from the specified file.
<b>--no-randomize</b>	Do not initialize the random number generator. This is useful if you want the script to return identical data for a series of invocations. This option is typically only used during testing.
<b>--no-tty</b>	Do not use tty for the SLAX debugger and other tty-related input needs.

Table 18: slaxproc Common Options and File Options (*continued*)

Option	Description
<b>--output &lt;file&gt;</b> <b>-o &lt;file&gt;</b>	Write output to the specified file.
<b>--param &lt;name&gt; &lt;value&gt;</b> <b>-a &lt;name&gt; &lt;value&gt;</b>	Pass a parameter to the script using the name and value pair provided. Note that all parameters are string parameters, so normal quoting rules apply.
<b>--partial</b> <b>-p</b>	Allow the input data to contain a partial SLAX script, which can be used with the <b>--slax-to-xslt</b> or <b>-x</b> mode to perform partial transformations.
<b>--trace &lt;file&gt;</b> <b>-t &lt;file&gt;</b>	Writes trace data to the specified file.
<b>--verbose</b> <b>-v</b>	Add verbose internal debugging output to the trace data output, including calls to the <b>slaxLog()</b> function.
<b>--version</b> <b>-V</b>	Show version information and exit.
<b>--write-version &lt;version&gt;</b> <b>-w &lt;version&gt;</b>	Write the specified version number to the output file when converting a script using the <b>--xslt-to-slax</b> or <b>-s</b> mode. This can be used to limit the conversion to avoid using SLAX 1.1 features. Acceptable values are 1.0 and 1.1. If this option is not specified, the SLAX script version defaults to 1.1.

## slaxproc File Argument Handling

For all modes except **check**, you have the option to reference file arguments positionally or use the file options to specify input and output files. If you use the file options, the files can be referenced in any order on the command line, and the file options can be interspersed among other command-line options.

If no input file is required, use the **-E** option to indicate an empty input document. Additionally, if the input or output option argument has the value "-", the standard input or standard output file is used. When using standard input, press Ctrl+d to signal the end-of-file.

To reference files positionally on the command line, specify the script file first if it is required for that mode, then specify the input file, and lastly specify the output file. Referencing the files positionally allows slaxproc to be plug compatible with xsltproc.

```
$ slaxproc script.slax input.xml output.xml
```

To reference files using explicit file option values, include **--name** or **-n**, **--input** or **-i**, and **--output** or **-o**, to specify the SLAX script file, and the input and output files, respectively.

```
$ slaxproc -i input.xml -n script.slax -o output.xml
```



If a file option is not provided, the filename is parsed positionally. In the following command, the input and output filenames are specified using the file options, but the script filename is referenced positionally:

```
$ slaxproc -i input.xml -o output.xml -g -v script.slax
```

To execute a script that requires no input file, include the **-E** option to indicate an empty input document.

```
$ slaxproc -E script.slax output.xml
```

## slaxproc UNIX Scripting Support

SLAX supports the shebang construct (**#!**), allowing the first line of a script to begin with the characters **"#"** and **"!"** followed by a path to the executable that runs the script and a set of command-line arguments. For example:

```
#!/usr/bin/slaxproc -n
```

or

```
#!/opt/local/bin/slaxproc -n
```

The operating system adds the name of the scripts and any command-line arguments to the command line that follows the **"#!"**. Adding the **-n** option allows additional arguments to be passed in on the command line. Flexible argument parsing allows aliases. For example, if the first line of the script is:

```
#!/usr/bin/slaxproc -E -n
```

additional arguments can be provided:

```
$ that-script -g output.xml
```

and the resulting command becomes:

```
/usr/bin/slaxproc -E -n /path/to/that-script -g output.xml
```

If the input or output argument has the value **"-"**, the standard input or standard output file is used. This allows slaxproc to be used as a traditional UNIX filter.

### Related Documentation

- [Using the SLAX Processor \(slaxproc\) on page 153](#)
- [libslax Distribution Overview on page 135](#)
- [libslax Library and Extension Libraries Overview on page 136](#)
- 

## Using the SLAX Processor (slaxproc)

The SLAX processor (slaxproc) is a command-line tool that can validate SLAX script syntax, convert between SLAX and XSLT formats, and format or run SLAX scripts. The slaxproc mode options define what function the processor performs. The following sections outline each of the modes:

- [Validating SLAX Script Syntax on page 154](#)
- [Converting Scripts Between XSLT and SLAX Formats on page 154](#)

- [Running SLAX Scripts on page 156](#)
- [Formatting SLAX Scripts on page 157](#)

## Validating SLAX Script Syntax

The SLAX processor provides an option to check the syntax of a SLAX script.

- To check the syntax of a SLAX script, issue the **slaxproc** command with the **--check** or **-c** mode option and the script filename.

```
$ slaxproc --check script1.slax
```

OR

```
$ slaxproc -c script1.slax
```

If the script syntax is correct, the SLAX processor issues a "script check succeeds" message. Otherwise, the processor issues a list of error messages detected during script parsing. Fix any indicated errors, and repeat the check.

## Converting Scripts Between XSLT and SLAX Formats

The SLAX processor supports converting scripts between SLAX and XSLT formats. When you convert a script, you have the option to reference the file arguments positionally or use the command-line file options, **--input** or **-i** and **--output** or **-o**, to specify the original input script and the converted output script, respectively. If you use the command-line file options, the files can be referenced in any order on the command line, and the file options can be interspersed among other command-line options.

If you do not provide an argument specifying an input file or an output file, the standard input or standard output file is used. When using standard input, press Ctrl+d to signal the end-of-file.

To convert a SLAX script to XSLT, issue the **slaxproc** command with the **--slax-to-xslt** or **-x** mode option. To reference the files positionally, specify the input SLAX file as the first argument and the desired output path and filename of the converted XSLT script as the second argument. To reference the files using command-line file options, include the file options in any order. For example:

```
$ slaxproc --slax-to-xslt test/script2.slax test/script2.xsl
```

OR

```
$ slaxproc -x -i test/script2.slax -o test/script2.xsl
```

To convert an XSLT script to SLAX, issue the **slaxproc** command with the **--xslt-to-slax** or **-s** mode option. To reference the files positionally, specify the input XSLT file as the first argument and the desired output path and filename of the converted SLAX script as the second argument. To reference the files using command-line file options, include the file options in any order.

Optionally, when converting a script from XSLT to SLAX, include the **--write-version** or **-w** option to specify the SLAX version of the converted script. Acceptable values are 1.0 and 1.1. The default is version 1.1. Use the **-p** option for partial input when you do not require the SLAX script boilerplate in the output.

The following example converts the XSLT script **script1.xsl** to the SLAX script **script1.slax**. The SLAX script will include the statement "version 1.0;" as the first line of the script.

```
$ slaxproc --xslt-to-slax -w 1.0 test/script1.xsl test/script1.slax
```

OR

```
$ slaxproc -s -w 1.0 -i test/script1.xsl -o test/script1.slax
```

The **slaxproc --xslt-to-slax** mode with the **-p** option is useful for quickly converting Junos OS hierarchies from XML format into SLAX. The following example provides the Junos OS **[edit policy-options]** hierarchy in XML format as input to the SLAX processor. The **-p** option indicates partial script input as opposed to a full script.

```
$ slaxproc -s -p
<policy-options>
  <policy-statement>
    <name>export-policy</name>
    <term>
      <name>term1</name>
      <from>
        <route-filter>
          <address>10.0.4.4/30</address>
          <prefix-length-range>/30-/30</prefix-length-range>
        </route-filter>
      </from>
      <then>
        <accept/>
      </then>
    </term>
  </policy-statement>
</policy-options>
[Ctrl+d]
```

The SLAX processor returns the SLAX formatting for the hierarchy.

```
<policy-options> {
  <policy-statement> {
    <name> "export-policy";
    <term> {
      <name> "term1";
      <from> {
        <route-filter> {
          <address> "10.0.4.4/30";
          <prefix-length-range> "/30-/30";
        }
      }
    }
    <then> {
      <accept>;
    }
  }
}
```

```
    }  
  }  
}
```

## Running SLAX Scripts

The SLAX processor supports executing SLAX scripts from the command line. This is the default **slaxproc** mode. To explicitly use this mode, issue the **slaxproc** command with the **--run** or **-r** command-line mode option.

When you execute a script, you have the option to reference the file arguments positionally or use the command-line file options, **--name** or **-n**, **--input** or **-i**, and **--output** or **-o**, to specify the SLAX script file, and the input and output files, respectively. If you use the command-line file options, the files can be referenced in any order on the command line, and the file options can be interspersed among other command-line options.

If no input file is required, use the **-E** option to indicate an empty input document. Additionally, if the input or output argument has the value "-", the standard input or standard output file is used. When using standard input, press Ctrl+d to signal the end-of-file.

The syntax for executing a script is:

```
$ slaxproc script input-file output-file
```

or

```
$ slaxproc (--name | -n) script (--input | -i) input-file (--output | -o)  
output-file
```

To execute a script using the **slaxproc** command-line tool:

1. Create a script using your favorite editor.
2. (Optional) Check the script syntax by invoking the processor with the **--check** or **-c** mode option, and fix any indicated errors.

```
$ slaxproc -c test/script1.slax
```

3. Execute the script and provide the required input and output files as well as any desired **slaxproc** options.

You can reference files positionally or use the command-line file options.

- To execute a script named **script1.slax** using **input.xml** as the input document and **output.xml** as the output document, issue either of the following commands. The two commands are identical in execution.

```
$ slaxproc script1.slax input.xml output.xml
```

```
$ slaxproc -n script1.slax -i input.xml -o output.xml
```

- To execute a script that requires no input file, include the **-E** option to indicate an empty input document. For example:

```
$ slaxproc -E script1.slax output.xml
```

```
$ slaxproc -n script1.slax -o output.xml -E
```

- To execute a script and use standard input as the input document, issue the **slaxproc** command with no input file argument. At the prompt, enter the input and press Ctrl+d to signal the end-of-file. For example:

```
$ slaxproc -n script1.slax -o output.xml
<user input>
[Ctrl+d]
```

## Formatting SLAX Scripts

The SLAX processor provides the option to format a script to correct the indentation and spacing to the preferred style. When you format a script, you have the option to reference the file arguments positionally or use the command-line file options, **--input** or **-i** and **--output** or **-o**, to specify the unformatted input file and the formatted output file, respectively. If you use the command-line file options, the files can be referenced in any order on the command line.

To format a SLAX script, issue the **slaxproc** command with the **--format** or **-F** mode option. To reference the files positionally, specify the unformatted SLAX script as the first argument and the desired output path and filename of the formatted SLAX script as the second argument. To reference the files using command-line file options, include the file options in any order. For example:

```
$ slaxproc --format script1.slax script1-format.slax
```

OR

```
$ slaxproc -F -i script1.slax -o script1-format.slax
```

Given the following unformatted SLAX script as input:

```
version 1.1;

decimal-format default-format {
decimal-separator "." ;
digit "#" ;
grouping-separator "," ;
infinity "Infinity" ;
minus-sign "-" ;
nan "NaN";
pattern-separator ";" ;
percent "%";
per-mille "\x2030";
zero-digit "0" ;
}

match / {
var $number = -14560302.5;
expr format-number($number, "###,###.00", "default-format");
}
```

the SLAX processor outputs the following formatted SLAX script:

```
version 1.1;

decimal-format default-format {
    decimal-separator ".";
    digit "#";
    grouping-separator ",";
}
```

```

infinity "Infinity";
minus-sign "-";
pattern-separator ";";
percent "%";
per-mille " 30";
zero-digit "0";
nan "NaN";
}

match / {
    var $number = -14560302.5;

    expr format-number($number, "###,###.00", "default-format");
}

```

- Related Documentation**
- [libslax Distribution Overview on page 135](#)
  - [libslax Library and Extension Libraries Overview on page 136](#)
  - [Understanding the SLAX Processor \(slaxproc\) on page 148](#)

## SLAX Debugger, Profiler, and callflow

- [SLAX Debugger, Profiler, and callflow Overview on page 158](#)
- [Using the SLAX Debugger, Profiler, and callflow on page 160](#)

### SLAX Debugger, Profiler, and callflow Overview

The Junos OS command-line interface (CLI) and the libslax distribution include the SLAX debugger (sdb), which is used to trace the execution of SLAX scripts. The SLAX debugger enables you to step through script execution, pause script execution at defined breakpoints, and review the value of script variables at any point.

The SLAX debugger operation and command syntax resemble that of the GNU Project Debugger (GDB). Many of the sdb commands follow their GDB counterparts, to the extent possible. [Table 19 on page 158](#) lists the SLAX debugger commands and a brief description of each command.

The SLAX debugger includes a profiler that can report information about the activity and performance of a script. The profiler, which is automatically enabled when you start the debugger, tracks script execution until the script terminates. At any point, profiling information can be displayed or cleared, and the profiler can be temporarily disabled or enabled. The SLAX debugger **callflow** command enables printing of informational data when you enter or exit levels of the script.

**Table 19: SLAX Debugger Commands**

Command	Description
<code>break [loc]</code>	Add a breakpoint to the script at the current line of execution. Optionally specify <i>[file:]line</i> or a template name to create a breakpoint at that position.

Table 19: SLAX Debugger Commands (*continued*)

Command	Description
callflow [on   off]	Enable or disable callflow tracing. You can explicitly specify the <b>on</b> or <b>off</b> value. Omitting the value toggles callflow on and off.
continue [ <i>loc</i> ]	Continue running the script until it reaches the next breakpoint. If there are no defined breakpoints, the script runs in its entirety. Optionally, specify [ <i>file:</i> ] <i>line</i> or a template name. When you include the optional argument, script execution continues until it reaches either a breakpoint or the specified line number or template name, whichever comes first.
delete [ <i>num</i> ]	Delete one or all breakpoints. Breakpoints are numbered sequentially as they are created. Omit the optional argument to delete all breakpoints. Include the breakpoint number as an argument to delete only the specified breakpoint. View currently active breakpoints with the <b>info</b> command.
finish	Finish executing the current template.
help	Display the help message.
info [breakpoints   profile   profile brief]	Display information about the current script. The default command lists all breakpoints in the script. Optionally specify the <b>profile</b> or <b>profile brief</b> arguments to display profiling information.
list [ <i>loc</i> ]	List the contents of the current script. Optionally specify [ <i>file:</i> ] <i>line</i> or a template name from which point the debugger lists partial script contents. The output includes the filename, line number, and code.
next	Execute the next instruction, stepping over any function or template calls.
over	Execute the next instruction, stepping over any function or template calls or instruction hierarchies.
print < <i>xpath</i> >	Print the value of the XPath expression.
profile [clear   on   off   report   report brief]	Enable or disable the profiler. The profiler is enabled by default. Include the <b>clear</b> option to clear profiling information. Include the <b>report</b> or <b>report brief</b> option to display profiling information for the current script.
quit	Exit debugging mode.
reload	Reload the script.
run	Restart script execution from the beginning of the script.
step	Execute the next instruction, stepping into any function or template calls or instruction hierarchies.
where	Show the backtrace of template calls.

## Using the SLAX Debugger, Profiler, and callflow

- [Invoking the SLAX Debugger on page 160](#)
- [Using the SLAX Debugger \(sdb\) on page 161](#)
- [Using the SLAX Profiler on page 162](#)
- [Using callflow on page 163](#)

---

### Invoking the SLAX Debugger

Both the Junos OS CLI and the SLAX processor in the libslax distribution include the SLAX debugger (sdb), which is used to trace the execution of SLAX scripts.

When you invoke the SLAX debugger, the command-line prompt changes to (sdb) to indicate that you are in debugging mode. For example:

```
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```

When using the SLAX debugger from the Junos OS CLI, you can only use the debugger with op scripts that are enabled in the configuration. To invoke the SLAX debugger from the CLI on a device running Junos OS, issue the **op invoke-debugger cli** operational mode command, include the op script name, and optionally include any necessary script arguments.

```
user@host>op invoke-debugger cli script <argument-name argument-value>
```

The following example invokes the SLAX debugger for the op script **ge-interfaces.slax**, which has two parameters, **interface** and **protocol**. Values are supplied for both arguments.

```
user@host> op invoke-debugger cli ge-interfaces interface ge-0/2/0.0 protocol inet
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```

To invoke the SLAX debugger when using the SLAX processor, issue the **slaxproc** command with the **--debug** or **-d** option. Specify the script file and any input or output files. If no input file is required, use the **-E** option to indicate an empty input document. If the **-i** or **--input** argument has the value "-", or if you do not include the input option or an input file, standard input is used. When using standard input, press Ctrl+d to signal the end-of-file. The general syntax is:

```
$ slaxproc --debug [options] [script] [files]
```

The following example invokes the SLAX debugger for the script **script1.slax** with an empty input document and an output file **script1-output.xml**

```
$ slaxproc --debug -n script1.slax -o script1-output.xml -E
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```



### Using the SLAX Debugger (sdb)

To view the SLAX debugger help message, issue the **help** command at the (sdb) prompt. To display the help message for a single command, issue **help *command***, where *command* is the sdb command for which you want more information. For example:

```
(sdb) help break
break [loc]      Add a breakpoint at [file:]line or template
```

The process for debugging a script varies depending on the script. A generic outline is presented here:

1. Enter debugging mode.
2. Insert breakpoints in the script using the **break** command.

During execution, the debugger pauses at defined breakpoints.

The breakpoint location can be the name of a template or a line number in the current script, or the filename and a line number separated by a colon. If you do not include an argument, a breakpoint is created at the current line of execution. Breakpoints are numbered sequentially as you create them. To view a list of breakpoints, issue the **info breakpoints** command. To delete a breakpoint, issue the **delete *num*** command, and specify the breakpoint number. To delete all breakpoints, issue the **delete** command with no argument.

The following example creates three breakpoints, the first at line 7, the second at line 25, and the third at the template named "three":

```
(sdb) break 7
Breakpoint 1 at file script1.slax, line 7
(sdb) break 25
Breakpoint 2 at file script1.slax, line 25
(sdb) break three
Breakpoint 3 at file script1.slax, line 51
(sdb) info breakpoints
List of breakpoints:
#1 [global] at script1.slax:7
#2 template two at script1.slax:25
#3 template three at script1.slax:51
```

3. Increment script execution by issuing the **continue**, **finish**, **next**, **over**, and **step** commands at the debugger prompt.

For example:

```
(sdb) next
Reached breakpoint 1, at script1.slax:7
script1.slax:3: var $byte = "10011001";
```

4. Review the value of variables as the program executes to ensure that they have the expected value.

```
print xpath-expression
```

5. To reload the script contents at any point and restart script execution from the beginning, issue the **reload** command.

```
(sdb) reload
The script being debugged has been started already.
```

```

Reload and restart it from the beginning? (y or n) y
Reloading script...
Reloading complete.

```

## Using the SLAX Profiler

The SLAX debugger includes a profiler that can report information about the activity and performance of a script. The profiler, which is automatically enabled when you start the debugger, tracks script execution until the script terminates. At any point, profiling information can be displayed or cleared, and the profiler can be temporarily disabled or enabled.

To access the profiler, issue the **profile** command at the SLAX debugger prompt, (sdb), and include any options. The profile command syntax is:

```
(sdb) profile [options]
```

[Table 20 on page 162](#) lists the profile command options. Issuing the **profile** command with no additional options toggles the profiler on and off.

```

(sdb) profile
Disabling profiler
(sdb)

```

You can access the profiler help by issuing the **help profile** command at the (sdb) prompt.

**Table 20: Profile Command Options**

Option	Description
clear	Clear profiling information
off	Disable profiling
on	Enable profiling
report [brief]	Report profiling information

To enable the profiler and print a report:

1. Enter debugging mode. The profiler is enabled by default.
2. Step through script execution, or execute a script in its entirety.

```

(sdb) run
<?xml version="1.0"?>
<message>Down rev PIC in Fruvenator, Fru-Master 3000</message>
Script exited normally.

```

3. At any point during script execution, display profiling information.

The **brief** option instructs sdb to avoid showing lines that were not hit, since there is no valid information. If you omit the **brief** option, dashes are displayed.

```
(sdb) profile report brief
```

The following sample output shows a profile report with and without the **brief** option. The source code data in the example is truncated for display purposes.

```
(sdb) profile report
Line Hits User U/Hit System S/Hit Source
1 - - - - - - version 1.0;
2 - - - - - -
3 2 4 2.00 8 4.00 match / {
4 1 25 25.00 13 13.00 var ....
5 - - - - - -
6 - - - - - - for-each....
7 1 45 45.00 10 10.00 ..
8 1 12 12.00 5 5.00 <message>
9 1 45 45.00 15 15.00 ....
10 - - - - - - }
11 - - - - - - }
Total 6 131 51 Total

(sdb) pro rep b
Line Hits User U/Hit System S/Hit Source
3 2 4 2.00 8 4.00 match / {
4 1 25 25.00 13 13.00 var ....
7 1 45 45.00 10 10.00 ....
8 1 12 12.00 5 5.00 <message>
9 1 45 45.00 15 15.00 ....
Total 6 131 51 Total
```

The profile report includes the following information:

- **Line**—Line number in the source file.
- **Hits**—Number of times this line was executed.
- **User**—Number of microseconds of "user" time spent processing this line.
- **U/Hit**—Average number of microseconds of "user" time per hit.
- **System**—Number of microseconds of "system" time spent processing this line.
- **S/Hit**—Average number of microseconds of "system" time per hit.
- **Source**—Source code line.

This information not only shows how much time is spent during code execution, but can also show which lines are being executed, which can help debug scripts where the execution does not match expectations.

### Using callflow

The SLAX debugger **callflow** command enables printing of informational data when you enter or exit levels of the script.

To enable callflow and view callflow data for a script:

1. Enter debugging mode.
2. Issue the **callflow** command at the SLAX debugger prompt, (sdb).

```
(sdb) callflow
Enabling callflow
```

3. Step through script execution, or execute a script in its entirety.

Callflow prints information as it enters and exits different levels of the script. Each output line references the instruction, filename, and line number of the frame.

```
(sdb) run
callflow: 0: enter <xsl:template> in match / at script3.slax:5
callflow: 1: enter <xsl:template> in template one at script3.slax:14
callflow: 2: enter <xsl:template> in template two at script3.slax:20
callflow: 3: enter <xsl:call-template> at script3.slax:22
...
<?xml version="1.0"?>
<message>Down rev PIC in Fruvenator, Fru-Master 3000</message>
Script exited normally.
```

**Related  
Documentation**

- [op invoke-debugger cli on page 856](#)
- [Understanding the SLAX Processor \(slaxproc\) on page 148](#)
- [Using the SLAX Processor \(slaxproc\) on page 153](#)
- [libslax Distribution Overview on page 135](#)

## PART 4

# Automation Scripting Using Python

- [Python Overview on page 167](#)



## CHAPTER 9

# Python Overview

- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
- [Requirements for Executing Python Automation Scripts on Devices Running Junos OS on page 169](#)
- [Overview of Python Modules Available on Devices Running Junos OS on page 171](#)

### Understanding Python Automation Scripts for Devices Running Junos OS

Starting in Junos OS Release 16.1, as an alternative to SLAX and XSLT, you can use Python to author Junos OS commit, event, op, and SNMP scripts on certain devices that include the Python extensions package in the software image. Supported Python versions include Python 2.7.x. In order to execute Python scripts on a device, you must configure the **language python** statement at the **[edit system scripts]** hierarchy level, and you must enable individual scripts by configuring the script filename under the hierarchy level appropriate to the script type. For detailed information about the requirements for executing Python automation scripts on devices running Junos OS, see [“Requirements for Executing Python Automation Scripts on Devices Running Junos OS” on page 169](#).

Python is a practical, feature-rich language with extensive library support that enables you to create customized and readable commit, event, op, and SNMP scripts on devices running Junos OS. Python automation scripts support [Junos PyEZ](#) APIs and constructs, which make it easy to execute RPCs, and perform operational and configuration tasks on devices running Junos OS. In addition, Python automation scripts can leverage all of the features and constructs in the Python language, which provides increased flexibility over SLAX and XSLT. Python scripts can also leverage the `lxml` library, which simplifies XPath handling, and the `Jinja2` library, which is a template engine that enables you to generate content from predefined templates. For more information about the modules that are available to Python scripts on devices running Junos OS, see [“Overview of Python Modules Available on Devices Running Junos OS” on page 171](#).

Python enables you to perform operations that are difficult or impossible to perform in SLAX and XSLT. SLAX and XSLT are designed for processing XML data, but Python is suited for processing any format. Python supports an extensive list of data types, including lists, dictionaries, sets, and tuples, and all Python variables are mutable, unlike most SLAX and XSLT variables. Loop control is limited in SLAX and XSLT scripts, but Python scripts can utilize statements like **break** and **continue** to precisely control loop behavior. Python also has extensive support for file operations, and you can invoke the standard Python debugger on the command line to debug Python op scripts.

Devices that support Python automation scripts include the **junos** module, which contains the **Junos** class and script inputs, and the **jcs** module, which provides access to a subset of the extension functions and templates that are available to SLAX and XSLT scripts.

Junos OS automatically provides input to automation scripts when they are executed. [Table 21 on page 168](#) outlines the objects defined in the **junos** module that are accessible to the different Python automation scripts.

**Table 21: Python Automation Script Input**

Python Object	Type	Script Type	Description	XSLT/SLAX Equivalent
<b>Junos_Context</b>	Dictionary	Commit	Device-specific information about the script execution environment such as the device hostname, the script type, and the user executing the script.	commit-script-input/junos-context
		Event		event-script-input/junos-context
		Op		op-script-input/junos-context
<b>Junos_Configuration</b>	lxml.etree_Element	Commit	Post-inheritance candidate configuration in Junos XML format.	commit-script-input/configuration
<b>Junos_Trigger_Event</b>	lxml.etree_Element	Event	Details of the event that triggered the policy.	event-script-input/trigger-event
<b>Junos_Received_Events</b>	lxml.etree_Element	Event	Details of correlated events that occurred before the triggering event.	event-script-input/received-events
<b>Junos_Remote_Execution_Details</b>	Generator function that produces a sequence of remote devices.	Event	Connection details for a local or remote host configured for an event script at the <b>[edit event-options event-script file filename remote-execution]</b> hierarchy level.	event-script-input/remote-execution-details

For more information about the Python programming language, see <https://www.python.org/>.

For more information about Junos PyEZ, see [Junos PyEZ](#).



- Related Documentation**
- [Requirements for Executing Python Automation Scripts on Devices Running Junos OS on page 169](#)
  - [Overview of Python Modules Available on Devices Running Junos OS on page 171](#)
  - [language on page 718](#)

## Requirements for Executing Python Automation Scripts on Devices Running Junos OS

You can use Python to author Junos OS commit, event, op, and SNMP automation scripts. To prevent the execution of unauthorized Python code on devices running Junos OS, an unsigned Python automation script must meet certain requirements before you can execute the script on a device.

By default, you cannot execute unsigned Python scripts on devices running Junos OS. To enable the execution of unsigned Python automation scripts that meet the requirements outlined in this topic, you must configure the **language python** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
language python;
```

As with SLAX and XSLT automation scripts, Python automation scripts must be stored in the appropriate directory on the device, and you must enable individual scripts by configuring the script filename under the hierarchy level appropriate to the script type in the configuration. For information about storing and enabling automation scripts, see [“Storing and Enabling Scripts” on page 661](#).

In Junos OS Release 16.1R2 and earlier releases, unsigned Python commit, event, op, and SNMP scripts must be owned by the root user, and Junos OS executes the scripts using the access privileges of the \*nix user and group **nobody**, which is the generic, unprivileged system account. Starting in Junos OS Release 16.1R3, unsigned Python automation scripts must be owned by either the root user or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file. Furthermore, Python automation scripts can be executed with the access privileges of authorized users.



**NOTE:** To enable a user who does not belong to the file's user or group class to execute an unsigned Python automation script, the script's file permissions must include read permission for others.

Starting in Junos OS Release 16.1R3, interactive Python scripts, such as commit and op scripts, run with the access privileges of the user who executes the command or operation that invokes the script. Non-interactive Python scripts, such as event and SNMP scripts, by default, still execute under the privileges of the user and group **nobody**. To execute the scripts using the access privileges of a specific user, you must configure the **python-script-user *username*** statement at the **[edit event-options event-script file *filename*]** hierarchy level for event scripts, or the **[edit system scripts snmp file *filename*]** hierarchy level for SNMP scripts, and specify a user configured at the **[edit system login]** hierarchy level.

For example:

```
[edit event-options event-script]
file filename {
  python-script-user username;
}

[edit system scripts snmp]
file filename {
  python-script-user username;
}
```



**NOTE:** You cannot configure Python event and SNMP scripts to execute with root access privileges.

Table 22 on page 170 outlines the requirements for executing unsigned Python automation scripts in the different Veriexec-enabled versions of Junos OS.

**Table 22: Python Automation Script Requirements**

Requirement/Restriction	Junos OS Release 16.1R2 or Earlier Release	Junos OS Release 16.1R3 or Later Release
File owner	Root user	Root user or a user in the Junos OS <b>super-user</b> login class.
File write permissions	Any	File owner only
<b>language python</b> statement must be configured at the <b>[edit system scripts]</b> hierarchy level	Yes	Yes
Script must be enabled in the configuration under the hierarchy appropriate to that script type	Yes	Yes
Access privileges	All Python automation scripts execute with the access privileges of the user and group <b>nobody</b>	Python commit and op scripts execute with the access privileges of the user who invokes the script.  Python event and SNMP scripts execute with the access privileges of the user configured in the <b>python-script-user</b> statement, or if the <b>python-script-user</b> statement is not configured, as user and group <b>nobody</b> .

## Release History Table

Release	Description
16.1R3	In Junos OS Release 16.1R2 and earlier releases, unsigned Python commit, event, op, and SNMP scripts must be owned by the root user, and Junos OS executes the scripts using the access privileges of the *nix user and group <b>nobody</b> , which is the generic, unprivileged system account. Starting in Junos OS Release 16.1R3, unsigned Python automation scripts must be owned by either the root user or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Furthermore, Python automation scripts can be executed with the access privileges of authorized users.

## Related Documentation

- [Storing and Enabling Scripts on page 661](#)
- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
- [python-script-user on page 817](#)

## Overview of Python Modules Available on Devices Running Junos OS

The Python interpreter is included as part of the Junos operating system (Junos OS). Starting in Junos OS Release 16.1, devices running Junos OS include several open-source Python modules that can be used in Python applications including commit, event, op, and SNMP automation scripts, Juniper Extension Toolkit (JET) applications, and YANG action scripts. Usage of a specific module depends on the permissions of the user who is executing it.

[Table 23 on page 171](#) outlines the Python modules that are available on devices running Junos OS and provides a brief description of each module.

**Table 23: Python Modules Available on Devices Running Junos OS**

Module	Description
<b>Crypto</b>	Collection of Python modules implementing cryptographic algorithms and protocols. The modules provide various secure hash and encryption functions.
<b>ecdsa</b>	Provides an implementation of ECDSA cryptography, which can be used to create key pairs, sign messages, and verify the signatures.
<b>jinja2</b>	Fast, secure, designer-friendly templating language for Python.  For more information about Jinja2, see <a href="http://jinja.pocoo.org/docs/dev/">http://jinja.pocoo.org/docs/dev/</a> .
<b>jnpr.junos (Junos PyEZ)</b>	Microframework for Python that enables you to automate devices running Junos OS. Junos PyEZ is designed to provide the capabilities that a user would have on the Junos OS command-line interface (CLI) in an environment built for automation tasks.  For more information, see <a href="#">Junos PyEZ</a> .

**Table 23: Python Modules Available on Devices Running Junos OS (continued)**

Module	Description
<b>lxml</b>	XML processing library that combines the speed and XML feature completeness of the C libraries libxml2 and libxslt with the simplicity of a native Python API.  For more information about the <b>lxml</b> Python module, see <a href="http://lxml.de/tutorial.html">http://lxml.de/tutorial.html</a> .
<b>markupsafe</b>	Provides the ability to escape and format an XML, HTML, or XHTML markup safe string.
<b>ncclient</b>	Facilitates client scripting and application development through the NETCONF protocol.  For more information about <b>ncclient</b> , including documentation for some of the external APIs, see <a href="http://ncclient.grnet.gr/">http://ncclient.grnet.gr/</a> .
<b>netaddr</b>	Network address manipulation library that enables processing of Layer 2 and Layer 3 network addresses.
<b>paho</b>	Serves as a client class that enables applications to connect to an MQTT broker to publish messages, and to subscribe to topics and receive published messages.
<b>paramiko</b>	SSH2 protocol library that provides the ability to make SSH2 protocol-based connections. This module supports all major ciphers and hash methods for both client and server modes.
<b>pyang</b>	Extensible YANG validator and converter that enables the processing, validation, and conversion of YANG modules.
<b>pyyaml</b>	Python module used to serialize and deserialize data in YAML format.
<b>thrift</b>	Provides Python bindings for the Apache Thrift framework.
<b>scp</b>	Implementation of the SCP protocol for Paramiko that uses Paramiko transport to send and receive files via the SCP protocol.

**Related Documentation**

- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)

## PART 5

# Automation Script Input

- [Using Global Parameters in Automation Scripts on page 175](#)



## CHAPTER 10

# Using Global Parameters in Automation Scripts

- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)

## Global Parameters and Variables in Junos OS Automation Scripts

---

Junos OS automatically provides input to automation scripts when they are executed. The script can reference this input, which includes device-specific information about the script execution environment such as the device hostname, the script type, and the user executing the script. This information is useful for creating scripts that respond to a variety of complex scenarios.

SLAX and XSLT scripts that import the **junos.xml** file can reference this information using the **\$junos-context** global variable, which is a node-set. The **junos.xml** import file also declares several predefined global parameters that enable the scripts to more easily reference a subset of this information. Python scripts can reference this information through the **junos.Junos\_Context** dictionary, which must be imported into the script.

To use the pre-defined parameters or global variable in SLAX and XSLT scripts, you must import the **junos.xml** file by including the **<xsl:import>** tag in the style sheet declaration of an XSLT script or by including the **import** statement in a SLAX script and specifying the **junos.xml** file location as shown in the following sample code:

**XSLT Syntax**

```
<?xml version="1.0"?>
<xsl:stylesheet version="1.0">
  <xsl:import href="../import/junos.xml"/>
  ...
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
import "../import/junos.xml";
```

**Python** To reference the information in Python scripts, import the **Junos\_Context** dictionary.

```
from junos import Junos_Context
```

The script input is described in detail in the following sections:

- [Global Parameters Available in SLAX and XSLT Scripts on page 176](#)
- [Global Variable Available in Automation Scripts on page 176](#)

## Global Parameters Available in SLAX and XSLT Scripts

Several predefined global parameters are available for use in SLAX and XSLT automation scripts that import the **junos.xml** file. The parameters provide information about the Junos OS environment. [Table 24 on page 176](#) describes the built-in arguments.

**Table 24: Predefined Parameters Available in SLAX and XSLT Scripts**

Name	Description	Example
<code>\$hostname</code>	Hostname of the local device	Tokyo
<code>\$localtime</code>	Local time when the script is executed	Fri Dec 10 11:42:21 2010
<code>\$localtime-iso</code>	Local time, in ISO format, when the script is executed	2010-12-10 11:42:21 PST
<code>\$product</code>	Model of the local device	m10i
<code>\$script</code>	Filename of the executing script	<b>test.slax</b>
<code>\$user</code>	Local name of the user executing the script	root

The predefined global parameters are declared in the **junos.xml** file. You do not need to declare these parameters in a script in order to use them. Access the value of the global parameters in a script by prefixing the parameter name with the dollar sign (\$), as shown in the following example:

SLAX syntax:

```
if ($user != "root") {
    var $script-message = $user _ " does not have permission to execute " _
    $script;
    expr jcs:output($script-message);
}
```

XSLT syntax:

```
<xsl:if test="$user != 'root'">
    <xsl:variable name="script-message"
        select="concat($user, ' does not have permission to execute ', $script)"/>
    <xsl:value-of select="jcs:output($script-message)"/>
</xsl:if>
```

## Global Variable Available in Automation Scripts

Commit, event, and op scripts can access specific environment information that is provided to the script upon execution. To access this information, Python scripts must import and reference the **junos.Junos\_Context** dictionary, and SLAX and XSLT scripts that import the **junos.xml** file can reference the **\$junos-context** global variable. **\$junos-context** and **Junos\_Context** contain identical information but in a format that is suitable for the respective script language.



The `$junos-context` variable is a node-set that contains the `<junos-context>` node and the following hierarchy, which is common to and embedded in the source tree of all scripts:

```
<junos-context>
  <chassis></chassis>
  <hostname></hostname>
  <localtime></localtime>
  <localtime-iso></localtime-iso>
  <pid></pid>
  <product></product>
  <re-master/>
  <routing-engine-name></routing-engine-name>
  <script-type></script-type>
  <tty></tty>
  <user-context>
    <class-name></class-name>
    <login-name></login-name>
    <uid></uid>
    <user></user>
  </user-context>
</junos-context>
```

Additionally, script-specific information is available depending on the type of script executed. For op scripts, the `<op-context>` element is also included in the source tree provided to an op script:

```
<junos-context>
  <op-context>
    <via-url/>
  </op-context>
</junos-context>
```

For commit scripts, the `<commit-context>` element is also included in the source tree provided to a commit script:

```
<junos-context>
  <commit-context>
    <commit-comment>"This is a test commit"</commit-comment>
    <commit-boot/>
    <commit-check/>
    <commit-sync/>
    <commit-confirm/>
    <database-path/>
  </commit-context>
</junos-context>
```

[Table 25 on page 178](#) identifies each node of the `$junos-context` variable node-set, provides a brief description of the node, and gives examples of values for any elements that are not input to a script as an empty tag.

Table 25: Global Variable \$junos-context Available to SLAX and XSLT Scripts

Parent Node	Node	Description	Example Content
<junos-context>	<chassis>	Specifies whether the script is executed on a component of a routing matrix, the Root System Domain (RSD), or a Protected System Domain (PSD)	scc, lcc (TX Matrix) psd, rsd (JCS) others
	<hostname>	Hostname of the local device	Tokyo
	<localtime>	Local time when the script is executed	Fri Dec 10 11:42:21 2010
	<localtime-iso>	Local time, in ISO format, when the script is executed	2010-12-10 11:42:21 PST
	<pid>	cscript process ID	5257
	<product>	Model of the local device	m10i
	<re-master/>	Empty element included if the script is executed on the master Routing Engine	
	<routing-engine-name>	Routing Engine on which the script is executed	re0
	<script-type>	Type of script being executed	op
	<tty>	TTY of the user's session	/dev/tty1
<junos-context> <user-context>	<class-name>	Login class of the user executing the script	superuser
	<login-name>	Login name of the user executing the script. For AAA access, this is the RADIUS/TACACS username.	jsmith
	<uid>	User ID number of the user executing the script as defined in the device configuration	2999
	<user>	Local name of the user executing the script. Junos OS uses the local name for authentication. It might differ from the <b>login-name</b> used for AAA authentication.	root
<junos-context> <op-context>  (op scripts only)	<via-url>	Empty element included if the remote op script is executed using the <b>op url</b> command	

Table 25: Global Variable \$junos-context Available to SLAX and XSLT Scripts (*continued*)

Parent Node	Node	Description	Example Content
<junos-context> <commit-context>  (commit scripts only)	<commit-boot/>	Empty element included when the commit occurs at boot time	
	<commit-check/>	Empty element included when a <b>commit check</b> is performed	
	<commit-comment>	User comment regarding the commit	Commit to fix forwarding issue
	<commit-confirm/>	Empty element included when a <b>commit confirmed</b> is performed	
	<commit-sync/>	Empty element included when a <b>commit synchronize</b> is performed	
	<database-path/>	Element specifying the location of the session's pre-inheritance candidate configuration. For normal configuration sessions, the value of the element is the location of the normal candidate database. For private configuration sessions, the value of the element is the location of the private candidate database. When the <b>&lt;get-configuration&gt; database-path</b> attribute is set to this value, the commit script retrieves the corresponding pre-inheritance candidate configuration.	

The **\$junos-context** variable is a node-set. Therefore, you can access the child elements throughout a script by including the proper XPath expression. The following SLAX commit script writes a message to the system log file if the commit is performed during initial boot-up. The message is given a facility value of **daemon** and a severity value of **info**. For more information, see [syslog\(\)](#).

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

import "../import/junos.xsl";

match configuration {
  if ($junos-context/commit-context/commit-boot) {
    expr jcs:syslog("daemon.info", "This is boot-time commit");
  }
  else {
    /* Do this ... */
  }
}

```

Python scripts must import the **Junos\_Context** dictionary from the **junos** module to access the environment information provided to scripts. The names of the keys in the **Junos\_Context** dictionary are identical to the names of the **\$junos-context** nodes outlined in [Table 25 on page 178](#). Nodes with child elements that are nested under the **junos-context**

node such as **user-context**, **op-context**, and **commit-context** map to items in **Junos\_Context**, where the key is the node name and the value is a dictionary of the node's child elements. For example:

```
'user-context': {'login-name': 'bsmith', 'user': 'bsmith', 'class-name':  
'j-superuser', 'uid': '9999'}
```

The following example output displays the **Junos\_Context** dictionary for an op script that was executed locally. Note that the op script input contains the **op-context** key, which in this scenario is empty.

```
{'product': 'm7i', 'user-context': {'login-name': 'bsmith', 'user': 'bsmith',  
'class-name': 'j-superuser', 'uid': '9999'}, 'routing-engine-name': 're0',  
'script-type': 'op', 're-master': None, 'hostname': 'R1', 'pid': '7136', 'tty':  
'/dev/tty1', 'chassis': 'others', 'op-context': ' ', 'localtime': 'Thu Jan  
22 11:45:47 2015', 'localtime-iso': '2015-01-22 11:45:47 PST'}
```

The following example output displays the **Junos\_Context** dictionary for a commit script that was executed during a **commit check** operation. Note that the commit script input contains the **commit-context** key.

```
{'product': 'm7i', 'user-context': {'login-name': 'bsmith', 'user': 'bsmith',  
'class-name': 'j-superuser', 'uid': '9999'}, 'routing-engine-name': 're0',  
'script-type': 'commit', 're-master': None, 'hostname': 'R1', 'pid': '11201',  
'tty': '/dev/tty1', 'commit-context': {'database-path':  
'/var/run/db/juniper.db', 'commit-check': None}, 'chassis': 'others',  
'localtime': 'Thu Jan 22 16:23:55 2015', 'localtime-iso': '2015-01-22 16:23:55  
PST'}
```

To access individual values in the dictionary, specify the key name. For example:

```
Junos_Context['hostname']  
Junos_Context['user-context']['class-name']
```

#### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [SLAX Parameters Overview on page 72](#)
- [SLAX Variables Overview on page 76](#)
- [XSLT Parameters Overview on page 24](#)
- [XSLT Variables Overview on page 27](#)

## PART 6

# Extension Functions and Named Templates for Automation Scripting

- [Using Extension Functions for Automation Scripting on page 183](#)
- [Extension Functions in the jcs and slax Namespaces on page 189](#)
- [Using Named Templates for Automation Scripting on page 225](#)
- [Named Templates in the jcs Namespace on page 229](#)



# Using Extension Functions for Automation Scripting

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Using Extension Functions in Junos OS Automation Scripts on page 186](#)

## Understanding Extension Functions in Junos OS Automation Scripts

---

Junos OS provides automation scripting tools, including extension functions and named templates, that can be used in commit, op, event, and SNMP scripts to more easily accomplish scripting tasks on devices running Junos OS. The libraries provide logic, data manipulation, input and output, and utility functions and enable you to perform operations that are difficult or impossible to perform in XPath. [Table 26 on page 184](#) summarizes the Juniper extension functions.

To use the extension functions, Python automation scripts must import the **jcs** module, and SLAX and XSLT scripts must declare the appropriate namespace URI in the style sheet declaration. A function in the **jcs** namespace is defined in the namespace URI <http://xml.juniper.net/junos/commit-scripts/1.0>, and a function in the **slax** namespace is defined in the namespace URI <http://xml.libslax.org/slax>.

Functions using the **slax** namespace are supported starting in Junos OS Release 12.2. Scripts using Junos OS-independent extension functions that existed in earlier releases in the **jcs** namespace can use either the **jcs** or the **slax** namespace starting in Junos OS Release 12.2. However, to use any of these functions in earlier Junos OS releases, scripts must use the **jcs** namespace URI.

Functions introduced in version 1.0 of the SLAX language can be used in SLAX scripts that include either a "version 1.0" statement or a "version 1.1" statement, but functions introduced in version 1.1 of the SLAX language can only be used in SLAX scripts that include the "version 1.1" statement. Scripts written in version 1.1 of the SLAX language are supported starting in Junos OS Release 12.2.

Python automation scripts only support a subset of the extension functions. For many of the unsupported functions, you can use standard Python libraries to achieve the same functionality.

Table 26: Extension Functions for Use in Automation Scripts

Function	Name-spaces	SLAX Version	Type	Description	Support in Python Scripts
<code>base64-decode()</code>	slax	1.1	Data manipulation	Decode BASE64 encoded data and return a string.	–
<code>base64-encode()</code>	slax	1.1	Data manipulation	Encode a string of data in the BASE64 encoding format.	–
<code>break-lines()</code>	jcs, slax	1.0	Data manipulation	Break a simple element into multiple elements, delimited by newlines.	–
<code>close()</code>	jcs	1.0	Utility	Close a previously opened connection handle.	–
<code>dampen()</code>	jcs, slax	1.0	Utility	Prevent the same operation from being repeatedly executed within a script.	Y
<code>document()</code>	slax	1.1	Input/output control	Read data from a file or URL and return a string.	–
<code>emit_change()</code>	–	–	Utility	Generate a persistent or transient change to the configuration in a commit script.	Y
<code>emit_error()</code>	–	–	Input/output control	Generate an error message on the console.	Y
<code>emit_snmp_attributes()</code>	–	–	Input/output control	Return information for the requested MIB object from an SNMP script.	Y
<code>emit_warning()</code>	–	–	Input/output control	Generate a warning message on the console.	Y
<code>empty()</code>	jcs, slax	1.0	Logic	Evaluate a node set or string argument to determine if it is an empty value.	–
<code>evaluate()</code>	slax	1.1	Input/output control	Evaluate a SLAX expression and return the result.	–
<code>execute()</code>	jcs	1.0	Utility	Execute a remote procedure call (RPC) within the context of a specified connection handle.	–
<code>first-of()</code>	jcs, slax	1.0	Logic	Return the first nonempty (non-null) item in a list. If all objects in the list are empty, the default expression is returned.	–
<code>get-command()</code>	jcs, slax	1.1	Input/output control	Prompt the user for command input and return the input as a string.	–



Table 26: Extension Functions for Use in Automation Scripts (*continued*)

Function	Name-spaces	SLAX Version	Type	Description	Support in Python Scripts
<code>get-hello()</code>	jcs	1.0	Utility	Return the session ID and the capabilities of the NETCONF server during a NETCONF session.	–
<code>get-input()</code> (XSLT/SLAX) <code>get_input()</code> (Python)	jcs, slax	1.0	Input/output control	Invoke a CLI prompt and wait for user input. If the script is run non-interactively, the function returns an empty value. This function cannot be used with event scripts.	Y
<code>get-protocol()</code>	jcs	1.0	Utility	Return the session protocol associated with the connection handle.	–
<code>get-secret()</code> (XSLT/SLAX) <code>get_secret()</code> (Python)	jcs, slax	1.0	Input/output control	Invoke a CLI prompt and wait for user input. The input is not echoed back to the user.	Y
<code>get_snmp_action()</code>	–	–	Input/output control	Retrieve the action value passed to the SNMP script.	Y
<code>get_snmp_oid()</code>	–	–	Input/output control	Retrieve the OID value passed to the SNMP script.	Y
<code>hostname()</code>	jcs	1.0	Utility	Return the fully qualified domain name associated with a given IPv4 or IPv6 address, provided the DNS server is configured on the device.	Y
<code>invoke()</code>	jcs	1.0	Utility	Invoke an RPC on a local device running Junos OS.	–
<code>open()</code>	jcs	1.0	Utility	Return a connection handle that can be used to execute RPCs.	–
<code>output()</code>	jcs, slax	1.0	Input/output control	Generate unformatted output text that is immediately sent to the CLI session.	Y
<code>parse-ip()</code>	jcs	1.0	Data manipulation	Parse an IPv4 or IPv6 address and return the host IP address, protocol family, prefix length, network address, and network mask.	Y
<code>printf()</code>	jcs, slax	1.0	Input/output control	Generate formatted output text. Most standard <b>printf</b> formats are supported, in addition to some Junos OS-specific formats. The function returns a formatted string but does not print it on call.	Y

Table 26: Extension Functions for Use in Automation Scripts (*continued*)

Function	Name-spaces	SLAX Version	Type	Description	Support in Python Scripts
<code>progress()</code>	jcs, slax	1.0	Input/output control	Issue a progress message containing the single argument immediately to the CLI session provided that the <b>detail</b> flag was specified when the script was invoked.	Y
<code>regex()</code>	jcs, slax	1.0	Data manipulation	Evaluate a regular expression against a given string argument and return any matches.	–
<code>sleep()</code>	jcs, slax	1.0	Utility	Cause the script to sleep for a specified time.	–
<code>split()</code>	jcs, slax	1.0	Data manipulation	Split a string into an array of substrings delimited by a regular expression pattern.	–
<code>sysctl()</code>	jcs, slax	1.0	Utility	Return the value of the given <b>sysctl</b> value as a string or an integer.	Y
<code>syslog()</code>	jcs, slax	1.0	Input/output control	Log messages with the specified priority to the system log file.	Y
<code>trace()</code>	jcs, slax	1.0	Input/output control	Issue a trace message, which is sent to the trace file.	Y

#### Related Documentation

- [Using Extension Functions in Junos OS Automation Scripts on page 186](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)
- [SLAX Variables Overview on page 76](#)
- [XSLT Variables Overview on page 27](#)

## Using Extension Functions in Junos OS Automation Scripts

Junos OS provides extension functions that can be used in commit, op, event, and SNMP scripts to more easily accomplish scripting tasks on devices running Junos OS. The following sections outline how to import and use the extension functions for different script languages:

- [Using Extension Functions in SLAX and XSLT Scripts on page 186](#)
- [Using Extension Functions in Python Scripts on page 188](#)

### Using Extension Functions in SLAX and XSLT Scripts

To use the extension functions in SLAX and XSLT scripts, the script must declare the appropriate namespace Uniform Resource Identifier (URI) in the style sheet declaration.

Junos OS extension functions, which have functionality that is specific to devices running Junos OS, are defined in the namespace with the associated URI `http://xml.juniper.net/junos/commit-scripts/1.0`. SLAX extension functions are defined in the namespace with the associated URI `http://xml.libslax.org/slax`.

SLAX and XSLT scripts generally map the **jcs** or **slax** prefix to its respective URI to avoid name conflicts with standard XSLT functions and user-defined templates. The scripts then qualify the extension functions with the appropriate prefix, which is expanded into its associated URI reference during processing.

For example, the following SLAX namespace statement maps the **jcs** prefix to the namespace URI that defines the Junos OS extension functions used in automation scripts:

```
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
```

The following SLAX namespace statement maps the **slax** prefix to the namespace URI that defines SLAX extension functions:

```
ns slax = "http://xml.libslax.org/slax";
```

To call an extension function in a SLAX or XSLT script, include any required variable declarations, call the function using **jcs:function-name()** or **slax:function-name()** as appropriate, and pass along any required or optional arguments. Arguments must be passed into the function in the precise order specified by the function definition. This is different from a template, where the parameters are assigned by name and can appear in any order. The return value of an extension function must always either be assigned to a variable or designated as output.

The following example maps the **jcs** prefix to the namespace identified by the URI `http://xml.juniper.net/junos/commit-scripts/1.0`. The script then calls the **jcs:invoke()** function with one argument.

<b>XSLT Syntax</b>	<pre>&lt;?xml version="1.0"?&gt; &lt;xsl:stylesheet version="1.0"   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"&gt;   ...   &lt;xsl:variable name="result"     select="jcs:invoke('get-software-information')"/&gt;   ... &lt;/xsl:stylesheet&gt;</pre>
--------------------	---

<b>SLAX Syntax</b>	<pre>version 1.0; ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0"; ... var \$result = jcs:invoke('get-software-information'); ...</pre>
--------------------	--

The following example maps the **slax** prefix to the namespace identified by the URI `http://xml.libslax.org/slax`. The script then calls the **slax:get-input()** function with one string argument. The **version** statement specifies version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

<b>SLAX Syntax</b>	<pre>version 1.1; ns slax = "http://xml.libslax.org/slax"; ... var \$input = slax:get-input(\$prompt); ...</pre>
--------------------	--

## Using Extension Functions in Python Scripts

Python automation scripts that import the **jcs** module can use a Python version of supported Junos OS and SLAX extension functions. To determine which extension functions are supported in Python scripts, see [“Understanding Extension Functions in Junos OS Automation Scripts” on page 183](#).

To call the equivalent extension function in a Python script, include the **import jcs** statement and any required variable declarations, call the function using **jcs.function\_name()**, and pass along any required or optional arguments. Note that in Python scripts, the extension function names must use underscores instead of hyphens. For example:

<b>Python Syntax</b>	<pre>import jcs  if __name__ == '__main__':     name = jcs.get_input("Enter name: ")     jcs.output(name)</pre>
----------------------	---

- |                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">Global Parameters and Variables in Junos OS Automation Scripts on page 175</a></li><li>• <a href="#">Using Named Templates in Junos OS Automation Scripts on page 226</a></li></ul> |
|------------------------------|--|

## CHAPTER 12

# Extension Functions in the jcs and slax Namespaces

- `base64-decode()` Function (SLAX) on page 190
- `base64-encode()` Function (SLAX) on page 190
- `break-lines()` Function (SLAX and XSLT) on page 191
- `close()` Function (SLAX and XSLT) on page 192
- `dampen()` Function (Python, SLAX, and XSLT) on page 192
- `document()` Function (SLAX) on page 193
- `emit_error()` Function (Python) on page 194
- `emit_snmp_attributes` Function (Python) on page 195
- `emit_warning()` Function (Python) on page 196
- `empty()` Function (SLAX and XSLT) on page 197
- `evaluate()` Function (SLAX) on page 197
- `execute()` Function (SLAX and XSLT) on page 198
- `first-of()` Function (SLAX and XSLT) on page 199
- `get-command()` Function (SLAX) on page 200
- `get-hello()` Function (SLAX and XSLT) on page 201
- `get-input()` Function (SLAX and XSLT) and `get_input()` (Python) on page 202
- `get-protocol()` Function (SLAX and XSLT) on page 203
- `get-secret()` Function (SLAX and XSLT) and `get_secret()` (Python) on page 204
- `get_snmp_action()` Function (Python) on page 205
- `get_snmp_oid()` Function (Python) on page 206
- `hostname()` Function (Python, SLAX, and XSLT) on page 207
- `invoke()` Function (SLAX and XSLT) on page 208
- `open()` Function (SLAX and XSLT) on page 209
- `output()` Function (Python, SLAX, and XSLT) on page 211
- `parse-ip()` Function (SLAX and XSLT) and `parse_ip()` (Python) on page 213
- `printf()` Function (Python, SLAX, and XSLT) on page 214

- [progress\(\) Function \(Python, SLAX, and XSLT\) on page 215](#)
- [regex\(\) Function \(SLAX and XSLT\) on page 216](#)
- [sleep\(\) Function \(SLAX and XSLT\) on page 217](#)
- [split\(\) Function \(SLAX and XSLT\) on page 218](#)
- [sysctl\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
- [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
- [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

---

## base64-decode() Function (SLAX)

---

<b>Namespaces</b>	<code>http://xml.libslax.org/slax</code>
<b>SLAX Syntax</b>	<code>string slax:base64-decode(string, &lt;control-string&gt;)</code>
<b>Release Information</b>	Function introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	<p>Decode BASE64 encoded data. BASE64 is a means of encoding arbitrary data into a radix-64 format that is more easily transmitted, typically using STMP or HTTP.</p> <p>Include the optional control string argument to replace any non-XML control characters in the decoded string with the specified string. If the argument is an empty string, non-XML characters are removed. The decoded data is returned to the caller.</p>
<b>Parameters</b>	<p><i>control-string</i>—(Optional) String to replace non-XML control characters in the decoded string. Use an empty string argument to remove the non-XML characters.</p> <p><i>string</i>—BASE64 encoded data.</p>
<b>Return Value</b>	<p><i>string</i>—Decoded data.</p> <p>—</p>
<b>Usage Examples</b>	<pre>var \$real-data = slax:base64-decode(\$encoded-data, "@");</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">base64-encode() Function (SLAX) on page 190</a></li></ul>

---

## base64-encode() Function (SLAX)

---

<b>Namespaces</b>	<code>http://xml.libslax.org/slax</code>
<b>SLAX Syntax</b>	<code>string slax:base64-encode(string)</code>
<b>Release Information</b>	Function introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

<b>Description</b>	Encode a string of data in the BASE64 encoding format. BASE64 is a means of encoding arbitrary data into a radix-64 format that is more easily transmitted, typically using STMP or HTTP.
<b>Parameters</b>	<i>string</i> —Input data string.
<b>Return Value</b>	<i>string</i> —Encoded data.
<b>Usage Examples</b>	<pre>var \$encoded-data = slax:base64-encode(\$real-data);</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">base64-decode() Function (SLAX) on page 190</a></li></ul>

---

### break-lines() Function (SLAX and XSLT)

---

<b>Namespaces</b>	<pre>http://xml.juniper.net/junos/commit-scripts/1.0 http://xml.libslax.org/slax</pre>
<b>SLAX Syntax</b>	<pre>var \$lines = prefix:break-lines(<i>expression</i>);</pre>
<b>XSLT Syntax</b>	<pre>&lt;xsl:variable name="lines" select="prefix:break-lines(<i>expression</i>)"/&gt;</pre>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6 Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.
<b>Description</b>	Break a simple element into multiple elements, delimited by newlines. This is especially useful for large output elements such as those returned by the <b>show pfe</b> command.  The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.
<b>Parameters</b>	<i>expression</i> —Original output.
<b>Return Value</b>	<i>\$lines</i> —Output broken up into lines.
<b>Usage Examples</b>	<pre>var \$lines = jcs:break-lines(\$output); for-each (\$lines) {     ... }</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">parse-ip() Function (SLAX and XSLT) and parse_ip() (Python) on page 213</a></li><li>• <a href="#">regex() Function (SLAX and XSLT) on page 216</a></li><li>• <a href="#">split() Function (SLAX and XSLT) on page 218</a></li></ul>

## close() Function (SLAX and XSLT)

---

Namespaces	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code>
SLAX Syntax	<code>var \$results = jcs:close(<i>connection</i>);</code>
XSLT Syntax	<code>&lt;xsl:variable name="results" select="jcs:close(<i>connection</i>)"/&gt;</code>
Release Information	Function introduced in Junos OS Release 9.3.
Description	Close a previously opened connection handle.
Parameters	<b><i>connection</i></b> —Connection handle generated by a call to the <b>jcs:open()</b> function.
Usage Examples	<p>The following example closes the connection handle <b>\$connection</b>, which was originally generated by a call to the <b>jcs:open()</b> function:</p> <pre>var \$connection = jcs:open(); ... var \$result = jcs:close(\$connection);</pre>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">execute() Function (SLAX and XSLT) on page 198</a></li><li>• <a href="#">open() Function (SLAX and XSLT) on page 209</a></li></ul>

## dampen() Function (Python, SLAX, and XSLT)

---

Namespaces	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
Python Syntax	<code>result = jcs.dampen(<i>tag-string</i>, <i>max</i>, <i>interval</i>)</code>
SLAX Syntax	<code>var \$result = <i>prefix</i>:dampen(<i>tag-string</i>, <i>max</i>, <i>interval</i>);</code>
XSLT Syntax	<code>&lt;xsl:variable name="result" select="<i>prefix</i>:dampen(<i>tag-string</i>, <i>max</i>, <i>interval</i>)"/&gt;</code>
Release Information	<p>Function introduced in Junos OS Release 9.4.</p> <p>Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2.</p> <p>Support for Python added in Junos OS Release 16.1.</p>
Description	<p>Prevent the same operation from being repeatedly executed within a script. The <b>dampen()</b> function returns <b>false</b> if the number of calls to the <b>jcs:dampen()</b> function exceeds a <b><i>max</i></b> number of calls in the time interval <b><i>interval</i></b>. Otherwise the function returns <b>true</b>. The function parameters include an arbitrary string that is used to distinguish different calls to the <b>jcs:dampen()</b> function. This tag is stored in the <code>/var/run</code> directory on the device.</p> <p>The <b><i>prefix</i></b> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.</p>



**Parameters** *interval*—Time interval, in minutes.

*max*—Maximum number of calls to the `jcs:dampen()` function with a given tag allowed before the function returns **false**. This limit is based on the number of calls within a specified time interval.

*tag-string*—Arbitrary string used to distinguish different calls to the `jcs:dampen()` function.

**Return Value** *result*—Boolean value based on the number of calls to `jcs:dampen()` with a given tag and within a specified time. If the number of calls for a given tag exceeds *max*, the return value is **false**. If the number of calls is less than *max*, the return value is **true**.

**Usage Examples** In the following example, if the `jcs:dampen()` function with the tag 'mytag1' is called less than three times in a 10-minute interval, the function returns **true**. If the function is called more than three times within 10 minutes, the function returns **false**.

```
if (jcs:dampen('mytag1', 3, 10)) {
    /* Code for situations when jcs:dampen() with */
    /* the tag 'mytag1' is called less than three times */
    /* within 10 minutes */
} else {
    /* Code for situations when jcs:dampen() with */
    /* the tag 'mytag1' exceeds the three call maximum */
    /* limit within 10 minutes */
}
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [hostname\(\) Function \(Python, SLAX, and XSLT\) on page 207](#)
  - [sleep\(\) Function \(SLAX and XSLT\) on page 217](#)
  - [sysctl\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)

## document() Function (SLAX)

**Namespaces** `http://xml.libslax.org/slax`

**SLAX Syntax** `string slax:document(url, <options>)`

**Release Information** Function introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.

**Description** Read data from a file or URL. The data can be encoded in any character set and can be BASE64 encoded. The default character set is "utf-8". Optional arguments specify the character encoding scheme and the encoding format, and define the replacement string for non-XML control characters. [Table 27 on page 194](#) lists the available options.



**NOTE:** Calling the `document()` function to read the same file multiple times within a script does not reflect any changes made to that file in the interim, because the file is cached when first read, and the original contents are returned every time the `document()` function retrieves the same file.

**Parameters** **options**—(Optional) Specify the character encoding scheme and format of the data, and define the replacement string for non-XML control characters.

**url**—File or URL from which to read data.

**Table 27: Options for slax:document Function**

Option	Description
<b>&lt;encoding&gt; string</b>	Character encoding scheme. For example "ascii" or "utf-8".
<b>&lt;format&gt; string</b>	"base64" for BASE64-encoded data.
<b>&lt;non-xml&gt; string</b>	String used to replace non-XML control characters. If the value is an empty string, non-XML characters are removed.

**Return Value** **string**—String representing the data.

**Usage Examples**

```
var $data = slax:document($url);

var $options := {
  <encoding> "ascii";
  <format> "base64";
  <non-xml> "#";
}
var $data2 = slax:document($url, $options);
```

**Related Documentation**

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)

## emit\_error() Function (Python)

**Python Syntax** `jcs.emit_error(string)`

**Release Information** Function introduced in Junos OS Release 16.1.

**Description** Generate an error message on the console from a Python automation script. The error message is preceded by the text **error:** .

When used in commit scripts, this function also causes the commit operation to fail.



**NOTE:** If the `jcs.emit_error` string argument includes a null character sequence (`\0`), the script generates an Invalid number of arguments error. Additionally, if the string argument consists exclusively of an empty string or escape sequences (`\n\t\b\v\f\r`) and spaces, Junos OS does not emit any message.

**Parameters** **string**—String describing the error.

**Usage Examples**

```
import jcs
```

```
def main():
    ...
    jcs.emit_error("Error message from a Python automation script")

if __name__ == '__main__':
    main()
```

- Related Documentation**
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)
  - [Example: Generating a Custom Error Message on page 279](#)
  - [emit\\_warning\(\) Function \(Python\) on page 196](#)

## emit\_snmp\_attributes Function (Python)

**Python Syntax** `jcs.emit_snmp_attributes(snmp_oid, snmp_oid_type, snmp_oid_value)`

**Release Information** Function introduced in Junos OS Release 16.1.

**Description** Return the attributes for the requested MIB object from a Python SNMP script.

SNMP scripts provide the flexibility to support custom MIBs. SNMP scripts are triggered automatically when the SNMP manager requests information for an unsupported object identifier (OID) that is mapped to an SNMP script. When the script is invoked, the `jcs.emit_snmp_attributes()` function emits the data for the requested object. The script acts like an SNMP subagent, and the system sends the return value from the script to the network management system (NMS).

**Parameters** *snmp\_oid*—String containing the OID about which the script is returning information.

*snmp\_oid\_type*—String that specifies the type of the data being returned for the requested MIB object. Acceptable values are **Counter32**, **Counter64**, **Integer32**, **Unsigned32**, and **Octet String**.

*snmp\_oid\_value*—Data to return for the requested MIB object.

**Usage Examples** The following example Python SNMP script processes two unsupported object identifiers (OIDs) that are mapped to the script in the Junos OS configuration. The `emit_snmp_attributes()` function returns the attributes for the requested OID.

```
import jcs

def main():

    snmp_action = jcs.get_snmp_action()
    snmp_oid = jcs.get_snmp_oid()

    jcs.syslog("8", "snmp_action = ", snmp_action, " snmp_oid = ", snmp_oid)

    if snmp_action == 'get':
        if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "10")
        elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "211")
```

```

elif snmp_action == 'get-next':
    if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
        jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.1",
                                "Integer32", "211")
    elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
        jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.2",
                                "Integer32", "429")

if __name__ == '__main__':
    main()

```

- Related Documentation**
- [get\\_snmp\\_action\(\) Function \(Python\) on page 205](#)
  - [get\\_snmp\\_oid\(\) Function \(Python\) on page 206](#)
  - [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)

## emit\_warning() Function (Python)

**Python Syntax** `jcs.emit_warning(string)`

**Release Information** Function introduced in Junos OS Release 16.1.

**Description** Generate a warning message on the console from a Python automation script. The warning message is preceded by the text **warning:**.



**NOTE:** If the `jcs.emit_warning` string argument includes a null character sequence (`\0`), the script generates an Invalid number of arguments error. Additionally, if the string argument consists exclusively of an empty string or escape sequences (`\n \t \b \v \f \r`) and spaces, Junos OS does not emit any message.

**Parameters** *string*—String describing the warning.

**Usage Examples**

```

import jcs

def main():
    ...
    jcs.emit_warning("Warning message from a Python automation script")

if __name__ == '__main__':
    main()

```

- Related Documentation**
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)
  - [Example: Generating a Custom Warning Message on page 276](#)
  - [emit\\_error\(\) Function \(Python\) on page 194](#)

## empty() Function (SLAX and XSLT)

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
<b>SLAX Syntax</b>	<code>var \$result = prefix:empty(<i>node-set</i>   <i>string</i>);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="result" select="prefix:empty(<i>node-set</i>   <i>string</i>)"/&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6 Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2.
<b>Description</b>	Test for the presence of a value and return <b>true</b> if the node set or string argument evaluates to an empty value.  The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.
<b>Parameters</b>	<b>(<i>node-set</i>   <i>string</i>)</b> —Argument to test for the presence of a value.
<b>Return Value</b>	<b>result</b> —Boolean value, which is <b>true</b> if the argument is empty.
<b>Usage Examples</b>	In the following example, if <b>\$set</b> is empty, the script executes the enclosed code block: <pre> if ( jcs:empty(\$set) ) {     /* Code to handle true value (\$set is empty) */ } </pre> <p>The following example tests whether the <b>description</b> node for interface fe-0/0/0/0 is empty. If the description is missing, a <b>&lt;message&gt;</b> tag is output.</p> <pre> if (jcs:empty(interfaces/interface[name="fe-0/0/0"]/description)) {     &lt;message&gt; "interface " _ name _ " is missing description"; } </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> <li>• <a href="#">first-of() Function (SLAX and XSLT) on page 199</a></li> </ul>

## evaluate() Function (SLAX)

<b>Namespaces</b>	<code>http://xml.libslax.org/slax</code>
<b>SLAX Syntax</b>	<code>object slax:evaluate(<i>expression</i>);</code>
<b>Release Information</b>	Function introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.
<b>Description</b>	Evaluate a SLAX expression and return the results of the expression. This supports expressions using the extended syntax provided by SLAX in addition to what is allowed in XPath.

<b>Parameters</b>	<i>expression</i> —SLAX expression to evaluate.
<b>Return Value</b>	<i>object</i> —Result of the expression.
<b>Usage Examples</b>	<pre>var \$result = slax:evaluate("expr[name == '&amp;']");</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> </ul>

## execute() Function (SLAX and XSLT)

<b>Namespaces</b>	http://xml.juniper.net/junos/commit-scripts/1.0
<b>SLAX Syntax</b>	<pre>var \$result = jcs:execute(connection, rpc);</pre>
<b>XSLT Syntax</b>	<pre>&lt;xsl:variable name="result" select="jcs:execute(connection, rpc)"/&gt;</pre>
<b>Release Information</b>	Function introduced in Junos OS Release 9.3.
<b>Description</b>	Execute a remote procedure call (RPC) within the context of a specified connection handle in SLAX and XSLT automation scripts. Any number of RPCs may be executed within the context of the connection handle until it is closed with the <b>jcs:close()</b> function.



**NOTE:** Python automation scripts can use [Junos PyEZ APIs](#) to execute RPCs on a local or remote device.

<b>Parameters</b>	<p><i>connection</i>—Connection handle generated by a call to the <b>jcs:open()</b> function.</p> <p><i>rpc</i>—Remote procedure call (RPC) to execute.</p>
<b>Return Value</b>	<i>result</i> —Results of the executed RPC, which include the contents of the <b>&lt;rpc-reply&gt;</b> element, but not the <b>&lt;rpc-reply&gt;</b> tag itself. This <b>\$result</b> variable is the same as that produced by the <b>jcs:invoke()</b> function. By default, the results are in XML format equivalent to the output produced with the <b>  display xml</b> option in the CLI.
<b>Usage Examples</b>	<p>In the following example, the <b>\$rpc</b> variable is declared and initialized with the Junos XML <b>&lt;get-interface-information&gt;</b> element. A call to the <b>jcs:open()</b> function generates a connection handle to the remote device at IP address 198.51.100.1. The user's login and password are provided as arguments to <b>jcs:open()</b> to provide access to the remote device. The code calls <b>jcs:execute()</b> and passes in the connection handle and RPC as arguments. Junos OS on the remote device processes the RPC and returns the results, which are stored in the <b>\$results</b> variable.</p> <pre>var \$rpc = &lt;get-interface-information&gt;; var \$connection = jcs:open('198.51.100.1', 'bsmith', 'test123'); var \$results = jcs:execute(\$connection, \$rpc); expr \$results;</pre>

The equivalent XSLT code is:

```
<xsl:variable name="connection" select="jcs:open('198.51.100.1', 'bsmith',
'test123')"/>
<xsl:variable name="rpc">
  <get-interface-information/>
</xsl:variable>
<xsl:variable name="results" select="jcs:execute($connection, $rpc)"/>
<xsl:value-of select="$results"/>
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [close\(\) Function \(SLAX and XSLT\) on page 192](#)
  - [invoke\(\) Function \(SLAX and XSLT\) on page 208](#)
  - [open\(\) Function \(SLAX and XSLT\) on page 209](#)

## first-of() Function (SLAX and XSLT)

<b>Namespaces</b>	http://xml.juniper.net/junos/commit-scripts/1.0 http://xml.libslax.org/slax
<b>SLAX Syntax</b>	var \$result = <i>prefix</i> :first-of( <i>object</i> , " <i>expression</i> ");
<b>XSLT Syntax</b>	<xsl:variable name="result" select=" <i>prefix</i> :first-of( <i>object</i> , ' <i>expression</i> ')"/>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6. Support for the slax namespace http://xml.libslax.org/slax added in Junos OS Release 12.2.
<b>Description</b>	Return the first nonempty (non-null) item in a list. If all objects in the list are empty, the default expression is returned. This function provides the same functionality as an <b>if / else-if / else</b> construct but in a much more concise format.  The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.
<b>Parameters</b>	<b><i>expression</i></b> —Default value returned if all objects in the list are empty.  <b><i>object</i></b> —List of objects.
<b>Return Value</b>	<b>result</b> —First nonempty (non-null) item in the object list. If all objects in the list are empty, the default expression is returned.
<b>Usage Examples</b>	In the following example, if the value of <b>\$a</b> is empty, <b>\$b</b> is checked. If the value of <b>\$b</b> is empty, <b>\$c</b> is checked. If the value of <b>\$c</b> is empty, <b>\$d</b> is checked. If the value of <b>\$d</b> is empty, the string "none" is returned.  jcs:first-of(\$a, \$b, \$c, \$d, "none")

In the following example, for each physical interface, the script checks for a description of each logical interface. If a logical interface description does not exist, the function returns the description of the (parent) physical interface. If the parent physical interface description does not exist, the function returns a message that no description was found.

```
var $rpc = <get-interface-information>;
var $results = jcs:invoke($rpc);
for-each ($results/physical-interface/logical-interface) {
    var $description = jcs:first-of(description, ../description, "no description
found");
}
```

The equivalent XSLT code is:

```
<xsl:variable name="rpc">
    <get-interface-information/>
</xsl:variable>
<xsl:variable name="results" select="jcs:invoke($rpc)"/>
<xsl:for-each select="$results/physical-interface/logical-interface">
    <xsl:variable name="description"
        select="jcs:first-of(description, ../description, 'no description
found')"/>
</xsl:for-each>
```

The code for the **description** variable declaration in the previous examples would be equivalent to the following more verbose **if / else-if / else** construct:

```
var $description = {
    if (description) {
        expr description;
    }
    else if (../description) {
        expr ../description;
    }
    else {
        expr "no description found";
    }
}
```

See also “[Example: Displaying DNS Hostname Information Using an Op Script](#)” on page 487.

- Related Documentation**
- [Example: Displaying DNS Hostname Information Using an Op Script](#) on page 487
  - [Understanding Extension Functions in Junos OS Automation Scripts](#) on page 183
  - [empty\(\) Function \(SLAX and XSLT\)](#) on page 197

---

## get-command() Function (SLAX)

---

Namespaces	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
SLAX Syntax	<i>string</i> = <i>prefix</i> :get-command( <i>string</i> );
Release Information	Function introduced in version 1.1 of the SLAX language, which is supported in Junos OS Release 12.2 and later releases.



<b>Description</b>	<p>Prompt the user for input and return the input as a string. If the readline (or libedit) library was found at install time, the return string is entered in the readline history, and will be available using the readline history keystrokes (Ctrl+P and Ctrl+N).</p> <p>The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.</p>
<b>Parameters</b>	<i>string</i> —Prompt text.
<b>Return Value</b>	<i>string</i> —Command text entered by the user.
<b>Usage Examples</b>	<pre>var \$response = slax:get-command("# ");</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li></ul>

---

## get-hello() Function (SLAX and XSLT)

---

<b>Namespaces</b>	http://xml.juniper.net/junos/commit-scripts/1.0
<b>SLAX Syntax</b>	<pre>var \$capabilities = jcs:get-hello(<i>connection</i>);</pre>
<b>XSLT Syntax</b>	<pre>&lt;xsl:variable name="capabilities" select="jcs:get-hello(<i>connection</i>)"/&gt;</pre>
<b>Release Information</b>	Function introduced in Junos OS Release 11.4.
<b>Description</b>	<p>Return the session ID and the capabilities of the NETCONF server during a NETCONF session.</p> <p>During session establishment, the NETCONF server and client application each emit a <b>&lt;hello&gt;</b> element to specify which operations, or <i>capabilities</i>, they support from among those defined in the NETCONF specification or published as proprietary extensions. The <b>&lt;hello&gt;</b> element encloses the <b>&lt;capabilities&gt;</b> element and the <b>&lt;session-id&gt;</b> element, which specifies the session ID for this NETCONF session.</p> <p>Within the <b>&lt;capabilities&gt;</b> element, a <b>&lt;capability&gt;</b> element specifies each supported function. Each capability defined in the NETCONF specification is represented by a uniform resource name (URN). Capabilities defined by individual vendors are represented by uniform resource identifiers (URIs), which can be URNs or URLs.</p>
<b>Parameters</b>	<i>connection</i> —Connection handle generated by a call to the <b>jcs:open()</b> function.
<b>Return Value</b>	<b>capabilities</b> —XML node set that specifies which operations, or <i>capabilities</i> , the NETCONF server supports. The node set also includes the session ID.

**Usage Examples** In the following code snippet, the user, bsmith, establishes a NETCONF session on the default port with the remote device, fivestar, which is running Junos OS. Since the code does not specify a value for the password, the user is prompted for a password during script execution. Once authentication is established, the code calls the `jcs:get-hello()` function and stores the return value in the variable `$hello`, which is then printed to the CLI.

```
var $netconf := {  
    <method> "netconf";  
    <username> "bsmith";  
}  
var $connection = jcs:open("fivestar", $netconf);  
var $hello = jcs:get-hello($connection);  
expr jcs:output($hello);  
expr jcs:close($connection);
```

The CLI displays the following output:

```
bsmith@fivestar's password:
```

```
urn:ietf:params:xml:ns:netconf:base:1.0  
urn:ietf:params:xml:ns:netconf:capability:candidate:1.0  
urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0  
urn:ietf:params:xml:ns:netconf:capability:validate:1.0  
urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file  
http://xml.juniper.net/netconf/junos/1.0  
http://xml.juniper.net/dmi/system/1.0
```

```
20847
```

**Related Documentation**

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding the Session Protocol in Automation Scripts on page 679](#)
- [Example: Specifying the Session Protocol for a Connection within Scripts on page 681](#)
- [get-protocol\(\) Function \(SLAX and XSLT\) on page 203](#)
- [open\(\) Function \(SLAX and XSLT\) on page 209](#)

---

## get-input() Function (SLAX and XSLT) and get\_input() (Python)

---

Namespaces	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
Python Syntax	<code>input = jcs.get_input('string')</code>
SLAX Syntax	<code>var \$input = <i>prefix</i>:get-input(<i>string</i>);</code>
XSLT Syntax	<code>&lt;xsl:variable name="input" select="<i>prefix</i>:get-input(<i>string</i>)"/&gt;</code>
Release Information	Function introduced in Junos OS Release 9.4. Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2. Support for Python added in Junos OS Release 16.1.

**Description** Invoke a CLI prompt and wait for user input. The user input is defined as a string for subsequent use. If the script is run non-interactively, the function returns an empty value. This function cannot be used with event scripts.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.



**NOTE:** If the `jcs.get_input` string argument includes a null character sequence (`\0`), Python automation scripts generate an **Invalid number of arguments** error.

**Parameters** *string*—CLI prompt text.

**Return Value** *input*—Text typed by the user and stored as a string. The return value will be empty if the script is run non-interactively.

**Usage Examples** In the following SLAX example, the user is prompted to enter a login name. The user's input is stored in the variable `$username`.

```
var $username = jcs:get-input("Enter login id: ");
```

In Python:

```
username = jcs.get_input("Enter login id: ")
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [get-secret\(\) Function \(SLAX and XSLT\) and get\\_secret\(\) \(Python\) on page 204](#)
  - [output\(\) Function \(Python, SLAX, and XSLT\) on page 211](#)
  - [printf\(\) Function \(Python, SLAX, and XSLT\) on page 214](#)
  - [progress\(\) Function \(Python, SLAX, and XSLT\) on page 215](#)
  - [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
  - [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

## get-protocol() Function (SLAX and XSLT)

**Namespaces** `http://xml.juniper.net/junos/commit-scripts/1.0`

**SLAX Syntax** `var $protocol = jcs:get-protocol(connection);`

**XSLT Syntax** `<xsl:variable name="protocol" select="jcs:get-protocol(connection)"/>`

**Release Information** Function introduced in Junos OS Release 11.4.

**Description** Return the session protocol associated with the connection handle. The protocol values are `junoscript`, `netconf`, and `junos-netconf`.

<b>Parameters</b>	<b>connection</b> —Connection handle generated by a call to the <b>jcs:open()</b> function.
<b>Return Value</b>	<b>protocol</b> —Session protocol associated with the connection handle. The values are <b>junoscript</b> , <b>netconf</b> , and <b>junos-netconf</b> .
<b>Usage Examples</b>	<p>In the following code snippet, the user, bsmith, establishes a NETCONF session on the default port with the remote device, fivestar. Since the code does not specify a value for the password, the user is prompted for a password during script execution. Once authentication is established, the code calls the <b>jcs:get-protocol()</b> function and stores the return value in the variable <b>\$protocol</b>, which is then printed to the CLI.</p> <pre>var \$netconf := {     &lt;method&gt; "netconf";     &lt;username&gt; "bsmith"; } var \$connection = jcs:open("fivestar", \$netconf); var \$protocol = jcs:get-protocol(\$connection); expr jcs:output(\$protocol); expr jcs:close(\$connection);</pre> <p>The CLI displays the following output:</p> <pre>bsmith@fivestar's password:  netconf</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">Understanding the Session Protocol in Automation Scripts on page 679</a></li><li>• <a href="#">Example: Specifying the Session Protocol for a Connection within Scripts on page 681</a></li><li>• <a href="#">get-hello() Function (SLAX and XSLT) on page 201</a></li><li>• <a href="#">open() Function (SLAX and XSLT) on page 209</a></li></ul>

---

## get-secret() Function (SLAX and XSLT) and get\_secret() (Python)

---

<b>Namespaces</b>	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
<b>Python Syntax</b>	<code>input = jcs.get_secret('string')</code>
<b>SLAX Syntax</b>	<code>var \$input = <i>prefix</i>:get-secret(<i>string</i>);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="input" select="<i>prefix</i>:get-secret(<i>string</i>)"/&gt;</code>
<b>Release Information</b>	<p>Function introduced in Junos OS Release 9.5R2.</p> <p>Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.</p> <p>Support for Python added in Junos OS Release 16.1.</p>
<b>Description</b>	Invoke a CLI prompt and wait for user input. Unlike the <b>jcs:get-input()</b> function, the input is not echoed back to the user, which makes the function useful for obtaining passwords.

The user input is defined as a string for subsequent use. This function cannot be used with event scripts.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

<b>Parameters</b>	<i>string</i> —CLI prompt text.
<b>Return Value</b>	<i>input</i> —Text typed by the user and stored as a string.
<b>Usage Examples</b>	<p>The following SLAX example shows how to prompt for a password that is not echoed back to the user:</p> <pre>var \$password = jcs:get-secret("Enter password: ");</pre> <p>In Python:</p> <pre>password = jcs.get_secret("Enter password")</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">get-input() Function (SLAX and XSLT) and get_input() (Python) on page 202</a></li><li>• <a href="#">output() Function (Python, SLAX, and XSLT) on page 211</a></li><li>• <a href="#">printf() Function (Python, SLAX, and XSLT) on page 214</a></li><li>• <a href="#">progress() Function (Python, SLAX, and XSLT) on page 215</a></li><li>• <a href="#">syslog() Function (Python, SLAX, and XSLT) on page 220</a></li><li>• <a href="#">trace() Function (Python, SLAX, and XSLT) on page 223</a></li></ul>

---

## get\_snmp\_action() Function (Python)

---

<b>Python Syntax</b>	<code>snmp_action = jcs.get_snmp_action()</code>
<b>Release Information</b>	Function introduced in Junos OS Release 16.1.
<b>Description</b>	Retrieve the SNMP action passed to a Python SNMP script.
<b>Return Value</b>	<i>snmp_action</i> —SNMP action value passed to the script.
<b>Usage Examples</b>	<p>The following example Python SNMP script processes two unsupported object identifiers (OIDs) that are mapped to the script in the Junos OS configuration. The <code>get_snmp_action()</code> function returns the value of the action argument passed into the script.</p> <pre>import jcs  def main():      snmp_action = jcs.get_snmp_action()     snmp_oid = jcs.get_snmp_oid()      jcs.syslog("8", "snmp_action = ", snmp_action, " snmp_oid = ", snmp_oid)</pre>

```
if snmp_action == 'get':
    if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
        jcs.emit_snmp_attributes(snmp_oid, "Integer32", "10")
    elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
        jcs.emit_snmp_attributes(snmp_oid, "Integer32", "211")

elif snmp_action == 'get-next':
    if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
        jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.1",
                                "Integer32", "211")
    elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
        jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.2",
                                "Integer32", "429")

if __name__ == '__main__':
    main()
```

- Related Documentation**
- [emit\\_snmp\\_attributes Function \(Python\) on page 195](#)
  - [get\\_snmp\\_oid\(\) Function \(Python\) on page 206](#)
  - [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)

---

## get\_snmp\_oid() Function (Python)

---

**Python Syntax**    `snmp_oid = jcs.get_snmp_oid()`

**Release Information**    Function introduced in Junos OS Release 16.1.

**Description**    Retrieve the SNMP object identifier (OID) passed to a Python SNMP script.

**Return Value**    `snmp_oid`—SNMP OID value to process.

**Usage Examples**    The following example Python SNMP script processes two unsupported object identifiers (OIDs) that are mapped to the script in the Junos OS configuration. The `get_snmp_oid()` function returns the value of the OID argument passed into the script.

```
import jcs

def main():

    snmp_action = jcs.get_snmp_action()
    snmp_oid = jcs.get_snmp_oid()

    jcs.syslog("8", "snmp_action = ", snmp_action, " snmp_oid = ", snmp_oid)

    if snmp_action == 'get':
        if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "10")
        elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "211")

    elif snmp_action == 'get-next':
        if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
            jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.1",
                                    "Integer32", "211")
        elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
            jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.2",
```

```

    "Integer32", "429")

    if __name__ == '__main__':
        main()

```

- Related Documentation**
- [emit\\_snmp\\_attributes Function \(Python\) on page 195](#)
  - [get\\_snmp\\_action\(\) Function \(Python\) on page 205](#)
  - [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)

## hostname() Function (Python, SLAX, and XSLT)

<b>Namespaces</b>	http://xml.juniper.net/junos/commit-scripts/1.0
<b>Python Syntax</b>	<code>name = jcs.hostname(address)</code>
<b>SLAX Syntax</b>	<code>var \$name = jcs:hostname(address);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="name" select="jcs:hostname(address)"/&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 9.3. Support for Python added in Junos OS Release 16.1.
<b>Description</b>	Return the fully qualified domain name associated with a given IPv4 or IPv6 address. The DNS server must be configured on the device in order to resolve the domain name.
<b>Parameters</b>	<b>address</b> —String containing an IPv4 or IPv6 address.
<b>Return Value</b>	<b>name</b> —Hostname associated with the IP address.
<b>Usage Examples</b>	<p>The following SLAX example initializes the variable <b>address</b> with the IP address 198.51.100.1. The <b>\$address</b> variable is passed as the argument to the <b>jcs:hostname()</b> function. If the DNS server is configured on the device, the function will resolve the IP address and return the fully qualified domain name, which is stored in the variable <b>host</b>.</p> <pre> var \$address = "198.51.100.1"; var \$host = jcs:hostname(\$address); </pre> <p>In XSLT:</p> <pre> &lt;xsl:variable name="address" select="198.51.100.1"&gt; &lt;xsl:variable name="host" select="jcs:hostname(\$address)"/&gt; </pre> <p>In Python:</p> <pre> host = jcs.hostname("198.51.100.1") </pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Finding LSPs to Multiple Destinations Using an Op Script on page 491</a></li> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> <li>• <a href="#">dampen() Function (Python, SLAX, and XSLT) on page 192</a></li> <li>• <a href="#">parse-ip() Function (SLAX and XSLT) and parse_ip() (Python) on page 213</a></li> </ul>

- [sleep\(\) Function \(SLAX and XSLT\) on page 217](#)
- [sysctl\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)

## invoke() Function (SLAX and XSLT)

Namespaces	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code>
SLAX Syntax	<code>var \$result = jcs:invoke(<i>rpc</i>);</code>
XSLT Syntax	<code>&lt;xsl:variable name="result" select="jcs:invoke(<i>rpc</i>)"/&gt;</code>

**Release Information** Function introduced in Junos OS Release 7.6.

**Description** Invoke a remote procedure call (RPC) on the local device. The function is called with one argument, either a string containing a Junos XML API RPC, or a tree containing an RPC. The result contains the contents of the `<rpc-reply>` element, not including the `<rpc-reply>` tag. An RPC allows you to perform functions equivalent to Junos OS operational mode commands.



**NOTE:** Python automation scripts can use [Junos PyEZ](#) APIs to execute RPCs on a local or remote device.

**Parameters** *rpc*—String containing a Junos XML API RPC or a tree containing an RPC.

**Return Value** *result*—Results of the executed RPC, which include the contents of the `<rpc-reply>` element, but not the `<rpc-reply>` tag itself. By default, the results are in XML format equivalent to the output produced with the `| display xml` option in the CLI.

**Usage Examples** The following example tests to see if the **interface** argument is included on the command line when the script is executed. If the argument is provided, the output of the **show interfaces terse** operational mode command is narrowed to include only information about the specified interface.

```
<xsl:param name="interface"/>
<xsl:variable name="rpc">
  <get-interface-information>
    <terse/>
    <xsl:if test="$interface">
      <interface-name>
        <xsl:value-of select="$interface"/>
      </interface-name>
    </xsl:if>
  </get-interface-information>
</xsl:variable>
<xsl:variable name="out" select="jcs:invoke($rpc)"/>
```

In this example, the `jcs:invoke()` function calls the Junos XML API RPC **get-software-information**, and stores the unmodified output in the variable **sw**:

```
<xsl:variable name="sw" select="jcs:invoke('get-software-information')"/>
```



- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [execute\(\) Function \(SLAX and XSLT\) on page 198](#)
  - [hostname\(\) Function \(Python, SLAX, and XSLT\) on page 207](#)
  - [sleep\(\) Function \(SLAX and XSLT\) on page 217](#)

## open() Function (SLAX and XSLT)

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code>
<b>SLAX Syntax</b>	<pre>var \$connection = jcs:open(); var \$connection = jcs:open(<i>remote-hostname</i>, &lt;username&gt;, &lt;passphrase&gt;,     &lt;routing-instance-name&gt;); var \$connection = jcs:open(<i>remote-hostname</i>, &lt;session-options&gt;);</pre>
<b>XSLT Syntax</b>	<pre>&lt;xsl:variable name="connection" select="jcs:open()"/&gt; &lt;xsl:variable name="connection" select="jcs:open(<i>remote-hostname</i>, &lt;username&gt;,     &lt;passphrase&gt;, &lt;routing-instance-name&gt;)/"&gt; &lt;xsl:variable name="connection" select="jcs:open(<i>remote-hostname</i>, &lt;session-options&gt;)/"&gt;</pre>
<b>Release Information</b>	<p>Function introduced in Junos OS Release 9.3.</p> <p>Support for NETCONF sessions added in Junos OS Release 11.4R1.</p> <p>Support for routing instances added in Junos OS Release 12.2R1.</p>
<b>Description</b>	<p>Return a connection handle that can be used to execute remote procedure calls (RPCs) using the <b>jcs:execute()</b> extension function. To execute an RPC on a remote device, an SSH session must be established. In order for the script to establish the connection, you must either configure the SSH host key information for the remote device on the local device where the script will be executed, or the SSH host key information for the remote device must exist in the known hosts file of the user executing the script.</p> <p>To redirect the SSH connection to originate from within a specific routing instance, include the routing instance name in the connection parameters. The routing instance must be configured at the <b>[edit routing-instances]</b> hierarchy level, and the remote device must be reachable either using the routing table for that routing instance or from one of the interfaces configured under that routing instance.</p> <p>Starting in Junos OS Release 11.4, the new parameter, <b>session-options</b>, supports the option to create a session either with the Junos XML protocol server on devices running Junos OS or with the NETCONF server on devices where NETCONF service over SSH is enabled. Previously, the function supported only sessions with the Junos XML protocol server on devices running Junos OS.</p>



**NOTE:** Python automation scripts can use [Junos PyEZ APIs](#) to establish a session and execute RPCs on a local or remote device.

The connection handle is closed with the **jcs:close()** function.

**Parameters** *passphrase*—(Optional) User's login passphrase. If you do not specify a passphrase and it is required for authentication, you should be prompted for one during script execution by the device to which you are connecting.

*remote-hostname*—Domain name or IP address of the remote router, switch, or security device. If you are opening a local connection, do not pass this value. If you specify a session type, this parameter is required.

*routing-instance-name*—(Optional) Routing instance from within which the SSH connection originates.

*session-options*—(Optional) XML node set that specifies the session protocol and connection parameters. The structure of the node set is:

```
var $session-options := {  
  <instance> "routing-instance-name";  
  <method> ("junoscript" | "netconf" | "junos-netconf");  
  <passphrase> "passphrase";  
  <password> "password";  
  <port> "port-number";  
  <routing-instance> "routing-instance-name";  
  <username> "username";  
}
```

- **<instance>**—(Optional) Routing instance from within which the SSH connection originates. This element is identical to **<routing-instance>**.
- **<method>**—(Optional) Session protocol. The protocol is one of three values: **junoscript**, **netconf**, or **junos-netconf**. If you do not specify a protocol, a **junoscript** session is created by default. A **<method>** value of **junoscript** establishes a session with the Junos XML protocol server on a device running Junos OS. A **<method>** value of **netconf** establishes a session with a NETCONF server over an SSHv2 connection. A **<method>** value of **junos-netconf** establishes a session with a NETCONF server over an SSHv2 connection on a device running Junos OS.
- **<passphrase> or <password>**—(Optional) User's login passphrase. If you do not specify a passphrase and it is required for authentication, you should be prompted for one during script execution by the device to which you are connecting.
- **<port>**—(Optional) Server port number for **netconf** and **junos-netconf** sessions. For NETCONF sessions, **jcs:open()** connects to the NETCONF server at the default port 830. If you specify a value for **<port>**, **jcs:open()** connects to the given port instead. Specifying a port number has no impact on **junoscript** sessions, which are always established over SSH port 22.
- **<routing-instance>**—(Optional) Routing instance from within which the SSH connection originates. This element is identical to **<instance>**.
- **<username>**—(Optional) User's login name. If you do not specify a username and it is required for the connection, the script uses the local name of the user executing the script.

*username*—(Optional) User's login name. If you do not specify a username and it is required for the connection, the script uses the local name of the user executing the script.

<b>Return Value</b>	<b>connection</b> —Connection handle to the remote host.
<b>Usage Examples</b>	<p>The following example shows how to connect to a local device:</p> <pre>var \$connection = jcs:open();</pre> <p>The following example shows how to connect to a remote device:</p> <pre>var \$connection = jcs:open(remote-hostname);</pre> <p>The following example shows how the user, bsmith, with the passphrase “test123” obtains a connection handle to the remote device, fivestar:</p> <pre>var \$connection = jcs:open("fivestar", "bsmith", "test123");</pre> <p>The following example shows how the user, bsmith, with the passphrase “test123” creates a <b>junos-netconf</b> session with a device running Junos OS:</p> <pre>var \$options := {   &lt;method&gt; "junos-netconf";   &lt;username&gt; "bsmith";   &lt;passphrase&gt; "test123"; } var \$connection = jcs:open("fivestar", \$options);</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> <li>• <a href="#">Understanding the Session Protocol in Automation Scripts on page 679</a></li> <li>• <a href="#">Example: Specifying the Session Protocol for a Connection within Scripts on page 681</a></li> <li>• <a href="#">close() Function (SLAX and XSLT) on page 192</a></li> <li>• <a href="#">execute() Function (SLAX and XSLT) on page 198</a></li> <li>• <a href="#">get-hello() Function (SLAX and XSLT) on page 201</a></li> <li>• <a href="#">get-protocol() Function (SLAX and XSLT) on page 203</a></li> </ul>

## output() Function (Python, SLAX, and XSLT)

<b>Namespaces</b>	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
<b>Python Syntax</b>	<code>jcs.output(string, &lt;string&gt;)</code>
<b>SLAX Syntax</b>	<code>expr prefix:output(string, &lt;string&gt;);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:value-of select="prefix:output(string, &lt;string&gt;)" /&gt;</code>
<b>Release Information</b>	<p>Function introduced in Junos OS Release 7.6.</p> <p>Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.</p> <p>Support for Python added in Junos OS Release 16.1.</p>
<b>Description</b>	Display one or more lines of output text, either on the CLI (when used in op scripts), or to the output file (when used in event scripts). The function can be called with either a

single string argument or with multiple string arguments. Multiple arguments are concatenated into a single string. A newline terminates the output text.

The **output()** function is not supported in commit scripts. SLAX and XSLT commit scripts use the `<xnm:warning>` and `<xnm:error>` result tree elements to display text on the CLI, and Python commit scripts use the **emit\_warning()** and **emit\_error()** functions.

The **output()** function displays its text immediately rather than waiting until the conclusion of the script, which differs from the SLAX and XSLT `<output>` element. This makes it suitable for scripts where user interaction is required or when status messages should be displayed in the midst of script processing. In SLAX and XSLT scripts, **jcs:output()** returns an empty node set, which can be ignored. Therefore, the **jcs:output()** function is normally called with the **expr** statement, rather than assigning its result to a variable.

The following escape characters are supported in the output text:

- `\\` –Backslash (as of Junos OS Release 10.2)
- `\r` –Carriage Return
- `\"` –Double quote (as of Junos OS Release 10.1R2)
- `\n` –Newline
- `\'` –Single quote
- `\t` –Tab

Starting in Junos OS Release 10.2, the maximum length for output text is 10 KB, and longer strings are truncated to the supported length.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

**Parameters**    *string*—Text that is output immediately to the CLI session.

**Usage Examples**    SLAX syntax:

```
expr jcs:output('The VPN is up.');
```

XSLT syntax:

```
<xsl:value-of select="jcs:output('The VPN is up.')" />
```

Python syntax:

```
jcs.output('The VPN is up.')
```

**Related Documentation**

- [Understanding Extension Functions in Junos OS Automation Scripts](#) on page 183
- [get-input\(\) Function \(SLAX and XSLT\) and get\\_input\(\) \(Python\)](#) on page 202
- [get-secret\(\) Function \(SLAX and XSLT\) and get\\_secret\(\) \(Python\)](#) on page 204
- [printf\(\) Function \(Python, SLAX, and XSLT\)](#) on page 214
- [progress\(\) Function \(Python, SLAX, and XSLT\)](#) on page 215

- [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
- [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

## parse-ip() Function (SLAX and XSLT) and parse\_ip() (Python)

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code>
<b>Python Syntax</b>	<code>result = jcs.parse_ip("ipaddress/(prefix-length   netmask)")</code>
<b>SLAX Syntax</b>	<code>var \$result = jcs:parse-ip("ipaddress/(prefix-length   netmask)");</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="result" select="jcs:parse-ip('ipaddress/(prefix-length   netmask)')"/&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 9.0. Support for Python added in Junos OS Release 16.1.
<b>Description</b>	Parse an IPv4 or IPv6 address.
<b>Parameters</b>	<p><i>ipaddress</i>—IPv4 or IPv6 address.</p> <p><i>netmask</i>—Netmask defining the network portion of the address.</p> <p><i>prefix-length</i>—Prefix length defining the number of bits used in the network portion of the address.</p>
<b>Return Value</b>	<p><b>result</b>—An array containing the following information. In SLAX and XSLT scripts, the array index starts at 1. In Python scripts, the array index begins at 0.</p> <ul style="list-style-type: none"> <li>• Host IP address (or <b>NULL</b> in the case of an error)</li> <li>• Protocol family (inet for IPv4 or inet6 for IPv6)</li> <li>• Prefix length</li> <li>• Network address</li> <li>• Network mask in dotted decimal notation for IPv4 addresses (left blank for IPv6 addresses)</li> </ul>
<b>Usage Examples</b>	<p>In the following two SLAX examples, an IPv4 address and an IPv6 address are parsed and the resulting output is detailed:</p> <pre>var \$addr = jcs:parse-ip("10.1.2.10/255.255.255.0");</pre> <ul style="list-style-type: none"> <li>• <b>\$addr[1]</b> contains the host address <b>10.1.2.10</b>.</li> <li>• <b>\$addr[2]</b> contains the protocol family <b>inet</b>.</li> <li>• <b>\$addr[3]</b> contains the prefix length <b>24</b>.</li> <li>• <b>\$addr[4]</b> contains the network address <b>10.1.2.0</b>.</li> <li>• <b>\$addr[5]</b> contains the netmask for IPv4 <b>255.255.255.0</b>.</li> </ul> <pre>var \$addr = jcs:parse-ip("2001:DB8::c50:8a:800:200C:417A/32");</pre>

- `$addr[1]` contains the host address `2001:db8:0:c50:8a:800:200c:417a`.
- `$addr[2]` contains the protocol family `inet6`.
- `$addr[3]` contains the prefix length `32`.
- `$addr[4]` contains the network address `2001:db8::`.
- `$addr[5]` is blank for IPv6 (`""`).

The following Python statement parses an IPv4 address. The values in the `addr` array are shown. Note that the array index begins at 0.

```
addr = jcs.parse_ip('10.1.2.10/255.255.255.0')
```

- `$addr[0]` contains the host address `10.1.2.10`.
- `$addr[1]` contains the protocol family `inet`.
- `$addr[2]` contains the prefix length `24`.
- `$addr[3]` contains the network address `10.1.2.0`.
- `$addr[4]` contains the netmask for IPv4 `255.255.255.0`.

#### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [break-lines\(\) Function \(SLAX and XSLT\) on page 191](#)
- [hostname\(\) Function \(Python, SLAX, and XSLT\) on page 207](#)
- [regex\(\) Function \(SLAX and XSLT\) on page 216](#)
- [split\(\) Function \(SLAX and XSLT\) on page 218](#)

---

## printf() Function (Python, SLAX, and XSLT)

Namespaces	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
Python Syntax	<code>jcs.printf(<i>expression</i>)</code>
SLAX Syntax	<code>expr <i>prefix</i>:printf(<i>expression</i>);</code>
XSLT Syntax	<code>&lt;xsl:value-of select="<i>prefix</i>:printf(<i>expression</i>)"/&gt;</code>
Release Information	Function introduced in Junos OS Release 7.6. Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2. Support for Python added in Junos OS Release 16.1.
Description	Generate formatted output text. Most standard <b>printf</b> formats are supported, in addition to some Junos OS–specific formats. The function returns a formatted string but does not print it on call. To use the following Junos OS modifiers, place the modifier between the percent sign (%) and the conversion specifier.

- **j1**—Operator that emits the field only if it changed from the last time the function was called. This assumes that the expression's format string is unchanged.
- **jc**—Operator that capitalizes the first letter of the associated output string.
- **jt{TAG}**—Operator that emits the tag if the associated argument is not empty.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

**Parameters** *expression*—Format string containing an arbitrary number of format specifiers and associated arguments to output.

**Usage Examples** In the following example, the **j1** operator suppresses printing the interface identifier so-0/0/0 in the second line of output, because the identifier argument has not changed from the first printing. The **jc** operator capitalizes the output strings **up** and **down**. The **jt{--}** operator does not print the **{--}** tag in the first line of output, because the associated output argument is an empty string. However, the tag is printed in the second line because the associated output is the non-empty string **test**.

```
<xsl:value-of select="jcs:printf('%-24j1s %-5jcs %-5jcs %s%jt{ -- }s\n',
    'so-0/0/0', 'up', 'down', '10.1.2.3', '')"/>
<xsl:value-of select="jcs:printf('%-24j1s %-5jcs %-5jcs %s%jt{ -- }s\n',
    'so-0/0/0', 'down', 'down', '10.1.2.3', 'test')"/>
```

produces the following output:

```
so-0/0/0      Up      Down  10.1.2.3
              Down    Down  10.1.2.3 -- test
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [get-input\(\) Function \(SLAX and XSLT\) and get\\_input\(\) \(Python\) on page 202](#)
  - [get-secret\(\) Function \(SLAX and XSLT\) and get\\_secret\(\) \(Python\) on page 204](#)
  - [output\(\) Function \(Python, SLAX, and XSLT\) on page 211](#)
  - [progress\(\) Function \(Python, SLAX, and XSLT\) on page 215](#)
  - [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
  - [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

## progress() Function (Python, SLAX, and XSLT)

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
<b>Python Syntax</b>	<code>jcs.progress(string)</code>
<b>SLAX Syntax</b>	<code>expr prefix:progress(string);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:value-of select="prefix:progress(string)"/&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6.

Support for the slax namespace <http://xml.libslax.org/slax> added in Junos OS Release 12.2.

Support for Python added in Junos OS Release 16.1.

**Description** Issue a progress message containing the single argument immediately to the CLI session provided that the **detail** flag was specified when the script was invoked.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

**Parameters** *string*—Text that is output to CLI session.

**Usage Examples** SLAX syntax:

```
expr jcs:progress('Working...');
```

XSLT syntax:

```
<xsl:value-of select="jcs:progress('Working...')"/>
```

Python syntax:

```
jcs.progress('Working...')
```

The script must be invoked with the **detail** flag in order for the progress message to appear in the CLI session.

```
user@host> op script1.slax detail
2010-10-01 16:27:54 PDT: running op script 'script1.slax'
2010-10-01 16:27:54 PDT: opening op script '/var/db/scripts/op/script1.slax'
2010-10-01 16:27:54 PDT: reading op script 'script1.slax'
2010-10-01 16:27:54 PDT: Working...
2010-10-01 16:28:14 PDT: inspecting op output 'script1.slax'
2010-10-01 16:28:14 PDT: finished op script 'script1.slax'
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [get-input\(\) Function \(SLAX and XSLT\) and get\\_input\(\) \(Python\) on page 202](#)
  - [get-secret\(\) Function \(SLAX and XSLT\) and get\\_secret\(\) \(Python\) on page 204](#)
  - [output\(\) Function \(Python, SLAX, and XSLT\) on page 211](#)
  - [printf\(\) Function \(Python, SLAX, and XSLT\) on page 214](#)
  - [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
  - [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

---

## regex() Function (SLAX and XSLT)

---

**Namespaces** <http://xml.juniper.net/junos/commit-scripts/1.0>  
<http://xml.libslax.org/slax>

**SLAX Syntax** `var $result = prefix:regex(pattern, string);`

**XSLT Syntax** `<xsl:variable name="result" select="prefix:regex(pattern, string)"/>`



<b>Release Information</b>	Function introduced in Junos OS Release 7.6 Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.
<b>Description</b>	Evaluate a regular expression against a given string argument and return any matches. This function requires two arguments: the regular expression and the string to which the regular expression is compared.  The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.
<b>Parameters</b>	<i>pattern</i> —Regular expression that is evaluated against the string argument.  <i>string</i> —String within which to search for matches of the specified regular expression.
<b>Return Value</b>	<i>result</i> —Array of strings that match the given regex pattern within the string argument.
<b>Usage Examples</b>	In the following example, the regex pattern consists of four distinct groups. The first group consists of the entire expression. The three subsequent groups are each of the parentheses–enclosed expressions within the main expression. The results for each <code>jcs:regex()</code> function call contain an array of the matches of the regex pattern to each of the specified strings. <pre> var \$pattern = "([0-9]+)(:*)([a-z]*)"; var \$a = jcs:regex(\$pattern, "123:xyz"); var \$b = jcs:regex(\$pattern, "r2d2"); var \$c = jcs:regex(\$pattern, "test999!!!");  \$a[1] == "123:xyz"    # string that matches the full reg expression \$a[2] == "123"        # ([0-9]+) \$a[3] == ":"          # (:*) \$a[4] == "xyz"        # ([a-z]*) \$b[1] == "2d"         # string that matches the full reg expression \$b[2] == "2"          # ([0-9]+) \$b[3] == ""           # (:*) [empty match] \$b[4] == "d"          # ([a-z]*) \$c[1] == "999"        # string that matches the full reg expression \$c[2] == "999"        # ([0-9]+) \$c[3] == ""           # (:*) [empty match] \$c[4] == ""           # ([a-z]*) [empty match]</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> <li>• <a href="#">break-lines() Function (SLAX and XSLT) on page 191</a></li> <li>• <a href="#">parse-ip() Function (SLAX and XSLT) and parse_ip() (Python) on page 213</a></li> <li>• <a href="#">split() Function (SLAX and XSLT) on page 218</a></li> </ul>

## sleep() Function (SLAX and XSLT)

<b>Namespaces</b>	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
<b>SLAX Syntax</b>	<code>expr prefix:sleep(<i>seconds</i>, &lt;<i>milliseconds</i>&gt;);</code>

<b>XSLT Syntax</b>	<code>&lt;xsl:value-of select="prefix:sleep(seconds, &lt;milliseconds&gt;)" /&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6. Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.
<b>Description</b>	<p>Cause the script to pause for a specified number of seconds and (optionally) milliseconds. You can use this function to help determine how a device component works over time. To do this, write a script that issues a command, calls the <code>jcs:sleep()</code> function, and then reissues the same command.</p> <p>The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.</p>
<b>Parameters</b>	<p><i>milliseconds</i>—(Optional) Number of milliseconds the script should sleep.</p> <p><i>seconds</i>—Number of seconds the script should sleep.</p>
<b>Usage Examples</b>	<p>In the following example, <code>jcs:sleep(1)</code> causes the script to sleep for 1 second, and <code>jcs:sleep(0, 10)</code> causes the script to sleep for 10 milliseconds:</p> <p>SLAX syntax:</p> <pre>expr jcs:sleep(1); expr jcs:sleep(0, 10);</pre> <p>XSLT syntax:</p> <pre>&lt;xsl:value-of select="jcs:sleep(1)" /&gt; &lt;xsl:value-of select="jcs:sleep(0, 10)" /&gt;</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">dampen() Function (Python, SLAX, and XSLT) on page 192</a></li><li>• <a href="#">hostname() Function (Python, SLAX, and XSLT) on page 207</a></li><li>• <a href="#">sysctl() Function (Python, SLAX, and XSLT) on page 220</a></li></ul>

---

## split() Function (SLAX and XSLT)

---

<b>Namespaces</b>	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
<b>SLAX Syntax</b>	<code>var \$substrings = prefix:split(expression, string, &lt;limit&gt;);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="substrings" select="prefix:split(expression, string, &lt;limit&gt;)" /&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 8.4 Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.
<b>Description</b>	Split a string into an array of substrings delimited by a regular expression pattern. If the optional integer argument <i>limit</i> is specified, the function splits the entire string into <i>limit</i>

number of substrings. If there are more than *limit* number of matches, the substrings include the first *limit*-1 matches as well as the remaining portion of the original string for the last match.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

**Parameters** *expression*—Regular expression pattern used as the delimiter.

*limit*—(Optional) Number of substrings into which to break the original string.

*string*—Original string.

**Return Value** *\$substrings*—Array of *limit* number of substrings. If *limit* is not specified, the result array size is equal to the number of substrings extracted from the original string as determined by the specified delimiter.

**Usage Examples** In the following example, the original string is "123:abc:456:xyz:789". The `jcs:split()` function breaks this string into substrings that are delimited by the regular expression pattern, which in this case is a colon(:). The optional parameter *limit* is not specified, so the function returns an array containing all the substrings that are bounded by the delimiter(:).

```
var $pattern = "(:)";
var $substrings = jcs:split($pattern, "123:abc:456:xyz:789");
```

returns:

```
$substrings[1] == "123"
$substrings[2] == "abc"
$substrings[3] == "456"
$substrings[4] == "xyz"
$substrings[5] == "789"
```

The following example uses the same original string and regular expression as the previous example, but in this case, the optional parameter *limit* is included. Specifying *limit*=2 causes the function to return an array containing only two substrings. The substrings include the first match, which is "123" (the same first match as in the previous example), and a second match, which is the remaining portion of the original string after the first occurrence of the delimiter.

```
var $pattern = "(:)";
var $substrings = jcs:split($pattern, "123:abc:456:xyz:789", 2);
```

returns:

```
$substrings[1] == "123"
$substrings[2] == "abc:456:xyz:789"
```

- Related Documentation**
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
  - [break-lines\(\) Function \(SLAX and XSLT\) on page 191](#)
  - [parse-ip\(\) Function \(SLAX and XSLT\) and parse\\_ip\(\) \(Python\) on page 213](#)
  - [regex\(\) Function \(SLAX and XSLT\) on page 216](#)

## sysctl() Function (Python, SLAX, and XSLT)

---

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
<b>Python Syntax</b>	<code>value = jcs.sysctl(sysctl-value, "(i   s)")</code>
<b>SLAX Syntax</b>	<code>var \$value = prefix:sysctl(sysctl-value, "(i   s)");</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="value" select="prefix:sysctl(sysctl-value, '(i   s)')"/&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6 Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2. Support for Python added in Junos OS Release 16.1.
<b>Description</b>	Return the given <b>sysctl</b> value as a string or an integer. Use the "i" argument to specify an integer. Use the "s" argument to specify a string.  The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.
<b>Parameters</b>	<b>sysctl-value</b> —Name of the <b>sysctl</b> value to return.
<b>Return Value</b>	<b>value</b> —Returned string or integer value.
<b>Usage Examples</b>	<code>var \$value = jcs:sysctl("kern.hostname", "s");</code>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li><li>• <a href="#">dampen() Function (Python, SLAX, and XSLT) on page 192</a></li><li>• <a href="#">hostname() Function (Python, SLAX, and XSLT) on page 207</a></li><li>• <a href="#">sleep() Function (SLAX and XSLT) on page 217</a></li></ul>

## syslog() Function (Python, SLAX, and XSLT)

---

<b>Namespaces</b>	<code>http://xml.juniper.net/junos/commit-scripts/1.0</code> <code>http://xml.libslax.org/slax</code>
<b>Python Syntax</b>	<code>jcs.syslog(priority, message, &lt;message2&gt;)</code>
<b>SLAX Syntax</b>	<code>expr prefix:syslog(priority, message, &lt;message2&gt;);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:value-of select="prefix:syslog(priority, message, &lt;message2&gt;)" /&gt;</code>
<b>Release Information</b>	Function introduced in Junos OS Release 7.6 Support for the slax namespace <code>http://xml.libslax.org/slax</code> added in Junos OS Release 12.2. Support for Python added in Junos OS Release 16.1.

**Description** Log messages with the specified priority to the system log file.

The priority can be expressed as a **facility.severity** string or as a calculated integer. The **message** argument is a string that is written to the system log file. Optionally, additional strings can be included in the argument list. The **message** argument is concatenated with any additional arguments, and the concatenated string is written to the system log file. The syslog file is specified at the **[edit system syslog]** hierarchy level of the configuration.

The *prefix* associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.

**Parameters** **message**—String that is output to the system log file.

**message2**—(Optional) Any additional number of strings passed as arguments to the function. These are concatenated with the **message** argument and output to the system log file.

**priority**—Priority given to the syslog message.

The priority can be specified as a **facility.severity** string, or it can be expressed as an integer calculated from the corresponding numeric values of the facility and severity strings. In Python scripts, the priority integer value must be passed in as a string.

Table 28 on page 221 and Table 29 on page 222 show the available facility and severity strings and their corresponding numeric values. The integer value of the **priority** parameter is calculated by multiplying the numeric value of the facility string by 8 and adding the numeric value of the severity string. For example, if the **facility.severity** string pair is "pfe.alert", the priority value is 161 ((20 x 8) + 1).

**Table 28: Facility Strings**

Facility String	Description	Numeric Value
auth	Authorization system	4
change	Configuration change log	22
conflict	Configuration conflict log	21
daemon	Various system processes	3
external	Local external applications	18
firewall	Firewall filtering system	19
ftp	FTP processes	11
interact	Commands executed by the UI	23
pfe	Packet Forwarding Engine	20
user	User processes	1

Table 29: Severity Strings

Severity String	Description	Numeric Value
alert	Conditions that should be corrected immediately	1
crit	Critical conditions	2
debug	Debug messages	7
emerg or panic	Panic conditions	0
err or error	Error conditions	3
info	Informational messages	6
notice	Conditions that should be specially handled	5
warn or warning	Warning messages	4

**Usage Examples** The following three SLAX examples log **pfe** messages with an **alert** priority. The string **"mymessage"** is output to the system log file. All three examples are equivalent.

```
expr jcs:syslog("pfe.alert", "mymessage");

expr jcs:syslog(161, "mymessage");

var $message = "mymessage";
expr jcs:syslog("pfe.alert", $message);
```

The following example logs **pfe** messages with an **alert** priority similar to the previous example. In this example, however, there are additional string arguments. For this case, the concatenated string **"mymessage mymessage2"** is output to the system log file.

```
expr jcs:syslog("pfe.alert", "mymessage ", "mymessage2");
```

Similarly, in Python:

```
jcs.syslog("pfe.alert", "message")
jcs.syslog("161", "message")
```

#### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [get-input\(\) Function \(SLAX and XSLT\) and get\\_input\(\) \(Python\) on page 202](#)
- [get-secret\(\) Function \(SLAX and XSLT\) and get\\_secret\(\) \(Python\) on page 204](#)
- [output\(\) Function \(Python, SLAX, and XSLT\) on page 211](#)
- [printf\(\) Function \(Python, SLAX, and XSLT\) on page 214](#)
- [progress\(\) Function \(Python, SLAX, and XSLT\) on page 215](#)
- [trace\(\) Function \(Python, SLAX, and XSLT\) on page 223](#)

## trace() Function (Python, SLAX, and XSLT)

<b>Namespaces</b>	<a href="http://xml.juniper.net/junos/commit-scripts/1.0">http://xml.juniper.net/junos/commit-scripts/1.0</a> <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a>
<b>Python Syntax</b>	<code>jcs.trace(<i>string</i>)</code>
<b>SLAX Syntax</b>	<code>expr <i>prefix</i>:trace(<i>string</i>);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:value-of select="<i>prefix</i>:trace(<i>string</i>)"/&gt;</code>
<b>Release Information</b>	<p>Function introduced in Junos OS Release 7.6.</p> <p>Support for the slax namespace <a href="http://xml.libslax.org/slax">http://xml.libslax.org/slax</a> added in Junos OS Release 12.2.</p> <p>Support for Python added in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Issue a trace message, which is sent to the trace file.</p> <p>You must configure <b>traceoptions</b> under the respective script type in the configuration hierarchy in order to output a message to the trace file using the <b>trace()</b> function. The output goes to the configured trace file. If <b>traceoptions</b> is enabled, but no trace file is explicitly configured, the output goes to the default trace file for that script type.</p> <p>The <i>prefix</i> associated with the namespace URI should be defined in the prefix-to-namespace mapping in the style sheet.</p>
<b>Parameters</b>	<i>string</i> —String that is output to the trace file.
<b>Usage Examples</b>	<p>SLAX syntax:</p> <pre>expr jcs:trace('test');</pre> <p>XSLT syntax:</p> <pre>&lt;xsl:value-of select="jcs:trace('test')"/&gt;</pre> <p>Python syntax:</p> <pre>jcs.trace('test')</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li> <li>• <a href="#">get-input() Function (SLAX and XSLT) and get_input() (Python) on page 202</a></li> <li>• <a href="#">get-secret() Function (SLAX and XSLT) and get_secret() (Python) on page 204</a></li> <li>• <a href="#">output() Function (Python, SLAX, and XSLT) on page 211</a></li> <li>• <a href="#">printf() Function (Python, SLAX, and XSLT) on page 214</a></li> <li>• <a href="#">progress() Function (Python, SLAX, and XSLT) on page 215</a></li> <li>• <a href="#">syslog() Function (Python, SLAX, and XSLT) on page 220</a></li> </ul>





# Using Named Templates for Automation Scripting

- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Using Named Templates in Junos OS Automation Scripts on page 226](#)

## Understanding Named Templates in Junos OS Automation Scripts

Junos OS provides several named templates that can be used in commit, op, event, and SNMP scripts to more easily accomplish scripting tasks on devices running Junos OS. The named templates reside in import files, which are included with the standard Junos OS installation available on switches, routers, and security devices running Junos OS. [Table 30 on page 225](#) summarizes the Junos OS named templates.

For information about using named templates in automation scripts, see “[Using Named Templates in Junos OS Automation Scripts](#)” on page 226.

**Table 30: Junos OS Named Templates**

Template	Description	Supported in Python Scripts
<code>jcs:edit-path</code>	Generate an <code>&lt;edit-path&gt;</code> element suitable for inclusion in an <code>&lt;xnm:error&gt;</code> or <code>&lt;xnm:warning&gt;</code> element.	—
<code>jcs:emit-change</code> (XSLT/SLAX) <code>jcs.emit_change</code> (Python)	Generate a persistent or transient change to the configuration.	Y
<code>jcs:emit-comment</code>	Emit a simple comment that indicates a change was made by a commit script.	—
<code>jcs:grep</code>	Search a file for all instances matching a specified regular expression and write the matching strings and corresponding lines to the result tree.	—
<code>jcs:load-configuration</code>	Make structured changes to the Junos OS configuration using an op script.	—

Table 30: Junos OS Named Templates (*continued*)

Template	Description	Supported in Python Scripts
<code>jcs:statement</code>	Generate a <code>&lt;statement&gt;</code> element suitable for inclusion in an <code>&lt;xnm:error&gt;</code> or <code>&lt;xnm:warning&gt;</code> element.	–

#### Related Documentation

- [Using Named Templates in Junos OS Automation Scripts on page 226](#)
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)

## Using Named Templates in Junos OS Automation Scripts

Junos OS provides several named templates that can be used in commit, op, event, and SNMP scripts to more easily accomplish scripting tasks on devices running Junos OS. The following sections outline how to import and use the templates for different script languages:

- [Using Named Templates in SLAX and XSLT Scripts on page 226](#)
- [Using Named Template Functionality in Python Scripts on page 227](#)

### Using Named Templates in SLAX and XSLT Scripts

To use the named templates within SLAX and XSLT scripts, the script must import the `junos.xml` file and also declare the appropriate namespace Uniform Resource Identifier (URI) in the style sheet declaration. The Junos OS named templates are defined in the namespace with the associated URI `http://xml.juniper.net/junos/commit-scripts/1.0`.

Import the `junos.xml` file into the script by including the `<xsl:import/>` tag element in XSLT scripts or the `import` statement in SLAX scripts and specifying the `junos.xml` file location.

SLAX and XSLT scripts generally map the `jcs` prefix to the URI to avoid name conflicts with standard XSLT or user-defined templates. The scripts then qualify the named templates with the appropriate prefix, which is expanded into its associated URI reference during processing. Map the `jcs` prefix to the URI by including the `xmlns:jcs` attribute in the opening `<xsl:stylesheet>` tag element for XSLT scripts or by including the `ns jcs` statement in SLAX scripts.

To call a named template in a script, include the `<xsl:call-template name="template-name">` element in XSLT scripts or the `call` statement in SLAX scripts and pass along any required or optional parameters. Template parameters are assigned by name and can appear in any order. This differs from functions where the arguments must be passed into the function in the precise order specified by the function definition.

The following example imports the **junos.xml** file into a script and maps the **jcs** prefix to the namespace identified by the URI <http://xml.juniper.net/junos/commit-scripts/1.0>. The script demonstrates a call to the **jcs:edit-path** template.

<b>XSLT Syntax</b>	<pre> &lt;?xml version="1.0"?&gt; &lt;xsl:stylesheet version="1.0"     xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"&gt;   &lt;xsl:import href="../import/junos.xml"/&gt;   ...   &lt;xsl:for-each select="interfaces/interface[starts-with(name, 'so-')] "&gt;     &lt;xnm:warning&gt;       &lt;xsl:call-template name="jcs:edit-path"/&gt;       &lt;message&gt;interface configured&lt;/message&gt;     &lt;/xnm:warning&gt;   &lt;/xsl:for-each&gt;   ... &lt;/xsl:stylesheet&gt; </pre>
<b>SLAX Syntax</b>	<pre> version 1.0; ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0"; import "../import/junos.xml"; ... for-each ( interfaces/interface[starts-with(name,'so-') ] ) {   &lt;xnm:warning&gt; {     call jcs:edit-path();     &lt;message&gt; "interface configured";   } } ... </pre>

## Using Named Template Functionality in Python Scripts

Python scripts that import the **jcs** module can execute functions that are the Python equivalent of the named templates used in SLAX and XSLT scripts. The functions provide essentially the same functionality in the Python script as the corresponding named template does in a SLAX or XSLT script. To determine which named templates have functionality that is supported in Python scripts, see [“Understanding Named Templates in Junos OS Automation Scripts” on page 225](#).

To call the function equivalent to a named template within a Python script, include any required variable declarations, call the function, and pass along any required or optional arguments. Note that in Python scripts, the function names must use underscores instead of hyphens. For example:

<b>Python Syntax</b>	<pre> import jcs  if __name__ == '__main__':     script = "system-check.py"     change_xml = """&lt;system&gt;&lt;scripts&gt;&lt;op&gt;         &lt;file&gt;&lt;name&gt;{0}&lt;/name&gt;&lt;/file&gt;&lt;/op&gt;         &lt;/scripts&gt;&lt;/system&gt;""".format(script)     jcs.emit_change(change_xml, "change", "xml") </pre>
----------------------	--

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Named Templates in Junos OS Automation Scripts on page 225</a></li> <li>• <a href="#">template on page 125</a></li> </ul>
------------------------------	--

- [xsl:template on page 48](#)
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Using Extension Functions in Junos OS Automation Scripts on page 186](#)

# Named Templates in the jcs Namespace

- [jcs:edit-path Template on page 229](#)
- [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)
- [jcs:emit-comment Template on page 233](#)
- [jcs:grep Template on page 234](#)
- [jcs:load-configuration Template on page 235](#)
- [jcs:statement Template on page 237](#)

## jcs:edit-path Template

---

<b>XSLT Syntax</b>	<pre>&lt;xsl:call-template name="jcs:edit-path"&gt;   &lt;xsl:with-param name="dot" select="expression"/&gt; &lt;/xsl:call-template&gt;</pre>
<b>SLAX Syntax</b>	<pre>call jcs:edit-path(\$dot=expression);</pre>
<b>Description</b>	<p>Generate an <b>&lt;edit-path&gt;</b> element suitable for inclusion in an <b>&lt;xnm:error&gt;</b> or <b>&lt;xnm:warning&gt;</b> element. This template converts a location in the configuration hierarchy into the standard text representation that you would see in the Junos OS configuration mode banner. By default, the location of the configuration error is passed into the <b>jcs:edit-path</b> template as the value of <b>dot</b>. This location defaults to ".", the current position in the XML hierarchy. You can alter the default by including a valid XPath expression for the <b>dot</b> parameter when you call the template.</p>
<b>Parameters</b>	<p><b>dot</b>—XPath expression specifying the hierarchy level. The default location is the position in the XML hierarchy that the script is currently evaluating. You can alter the default when you call the template by including a valid XPath expression either for the <b>dot</b> parameter in SLAX scripts or for the <b>select</b> attribute of the <b>dot</b> parameter in XSLT scripts.</p>
<b>Usage Examples</b>	<p>The following example demonstrates how to call the <b>jcs:edit-path</b> template in a commit script and set the context to the <b>[edit chassis]</b> hierarchy level:</p> <pre>&lt;xsl:if test="not(chassis/source-route)"&gt;   &lt;xnm:warning&gt;     &lt;xsl:call-template name="jcs:edit-path"&gt;       &lt;xsl:with-param name="dot" select="chassis"/&gt;     &lt;/xsl:call-template&gt;     &lt;message&gt;IP source-route processing is not enabled.&lt;/message&gt;</pre>

```

    </xnm:warning>
  </xsl:if>

```

When you commit a configuration that does not enable IP source routing, the code generates an `<xnm:warning>` element, which results in the following command-line interface (CLI) output:

```

user@host# commit
[edit chassis] # The hierarchy level is generated by the jcs:edit-path template.
warning: IP source-route processing is not enabled.
commit complete

```

- Related Documentation**
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [call on page 93](#)
  - [xnm:error \(Junos XML\) on page 428](#)
  - [xnm:warning \(Junos XML\) on page 430](#)
  - [xsl:call-template on page 40](#)

## emit-change Template (SLAX and XSLT) and emit\_change(Python)

<b>Python Syntax</b>	<code>jcs.emit_change(<i>content</i>, <i>tag</i>, <i>format</i>)</code>
<b>SLAX Syntax</b>	<pre> call jcs:emit-change(\$dot=<i>expression</i>, \$name = name(\$dot), \$tag = "(change   transient-change)" {   with \$content = {     ...   }   with \$message = {     expr "<i>message</i>";   } } </pre>
<b>XSLT Syntax</b>	<pre> &lt;xsl:call-template name="jcs:emit-change"&gt;   &lt;xsl:with-param name="content"&gt;     ...   &lt;/xsl:with-param&gt;   &lt;xsl:with-param name="dot" select="<i>expression</i>"&gt;   &lt;xsl:with-param name="message"&gt;     &lt;xsl:text&gt;<i>message</i>&lt;/xsl:text&gt;   &lt;/xsl:with-param&gt;   &lt;xsl:with-param name="name" select="name(\$dot)"&gt;   &lt;xsl:with-param name="tag" select="'(change   transient-change)'"&gt; &lt;/xsl:call-template&gt; </pre>
<b>Description</b>	Generate a persistent or transient change to the configuration.
<b>Parameters</b>	<p>This template includes the following optional parameters:</p> <p><b>content</b>—Content of the persistent or transient change. In SLAX and XSLT scripts, this is relative to <b>dot</b>. Python scripts must include the full configuration path representing all levels of the configuration hierarchy formatted as an XML string.</p>

**dot**—XPath expression specifying the hierarchy level at which the change will be made. The default location is the position in the XML hierarchy that the script is currently evaluating. You can alter the default when you call the template by including a valid XPath expression either for the **dot** parameter in SLAX scripts or for the **select** attribute of the **dot** parameter in XSLT scripts.

**format**—Format of the configuration data loaded through a Python commit script. The only supported format is **xml**.

**message**—Warning message displayed in the CLI notifying the user that the configuration has been changed. The message parameter automatically includes the edit path, which defaults to the current location in the XML hierarchy. To change the default edit path, specify a valid XPath expression either for the **dot** parameter in SLAX scripts or for the **select** attribute of the **dot** parameter in XSLT scripts.

**name**—Allows you to refer to the current element or attribute. The **name()** XPath function returns the name of an element or attribute. The **name** parameter defaults to the value **name(\$dot)**, which is the name of the element in **dot** (which in turn defaults to “.”, which is the current element).

**tag**—Type of change to generate. Specify **'change'** to generate a persistent change, or specify **'transient-change'** to generate a transient change. If you omit this parameter in SLAX and XSLT scripts, the **jcs:emit-change** template defaults to generating a persistent change.

**Usage Examples** The following example demonstrates how to call the **jcs:emit-change** template in an XSLT commit script:

```
<xsl:template match="configuration">
  <xsl:for-each select="interfaces/interface/unit[family/iso]">
    <xsl:if test="not(family/mp1s)">
      <xsl:call-template name="jcs:emit-change">
        <xsl:with-param name="message">
          <xsl:text>Adding 'family mp1s' to ISO-enabled
interface</xsl:text>
        </xsl:with-param>
        <xsl:with-param name="content">
          <family>
            <mp1s/>
          </family>
        </xsl:with-param>
      </xsl:call-template>
    </xsl:if>
  </xsl:for-each>
</xsl:template>
```

When you commit a configuration that includes one or more interfaces that have IS-IS enabled but do not have the **family mp1s** statement included at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level, the **jcs:emit-change** template adds the **family mp1s** statement to the configuration and generates the following CLI output:

```
[edit]
user@host# commit
```

```
[edit interfaces interface so-1/2/3 unit 0]
  warning: Adding 'family mpls' to ISO-enabled interface
[edit interfaces interface so-1/2/3 unit 0]
  warning: Adding ISO-enabled interface so-1/2/3.0 to [protocols mpls]
[edit interfaces interface so-1/3/2 unit 0]
  warning: Adding 'family mpls' to ISO-enabled interface
[edit interfaces interface so-1/3/2 unit 0]
  warning: Adding ISO-enabled interface so-1/3/2.0 to [protocols mpls]
commit complete
```

The **content** parameter of the **jcs:emit-change** template provides a simpler method for specifying a change to the configuration. For example, consider the following code:

```
<xsl:with-param name="content">
  <family>
    <mpls/>
  </family>
</xsl:with-param>
```

In SLAX and XSLT scripts, the **jcs:emit-change** template converts the **content** parameter into a **<change>** request. The **<change>** request inserts the provided partial configuration content into the complete hierarchy of the current context node. Thus, the **jcs:emit-change** template changes the hierarchy information in the **content** parameter into the following code:

```
<change>
  <interfaces>
    <interface>
      <name><xsl:value-of select="name"/></name>
      <unit>
        <name><xsl:value-of select="unit/name"/></name>
        <family>
          <mpls/>
        </family>
      </unit>
    </interface>
  </interfaces>
</change>
```

If a transient change is required, the **tag** parameter can be passed in as **'transient-change'**, as shown here:

```
<xsl:with-param name="tag" select="'transient-change'"/>
```

The extra quotation marks are required to allow XSLT to distinguish between the string **"transient-change"** and the contents of a node named **"transient-change"**. If the change is relative to a node other than the context node, the parameter **dot** can be set to that node, as shown in the following example, where context is set to the **[edit chassis]** hierarchy level:

```
<xsl:for-each select="interfaces/interface/unit">
  ...
  <xsl:call-template name="jcs:emit-change">
    <xsl:with-param name="dot" select="chassis"/>
  ...
</xsl:for-each>
```

The following Python commit script generates a persistent change to the configuration:

```
import jcs

if __name__ == '__main__':
```



```
script = "system-check.py"
change_xml = """<system><scripts><op>
    <file><name>{0}</name></file></op>
    </scripts></system>""".format(script)
jcs.emit_change(change_xml, "change", "xml")
```

- Related Documentation**
- [Example: Imposing a Minimum MTU Setting on page 382](#)
  - [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [call on page 93](#)
  - [xsl:call-template on page 40](#)

## jcs:emit-comment Template

**XSLT Syntax**

```
<junos:comment>
  <xsl:text>...</xsl:text>
</junos:comment>
```

**Description** Generate a comment in the configuration from within a commit script. The template contains a **<junos:comment>** element. You never call the **jcs:emit-comment** template directly. Rather, you include its **<junos:comment>** element and the child element **<xsl:text>** inside a call to the **jcs:emit-change** template, a **<change>** element, or a **<transient-change>** element.

**Usage Examples** The following example demonstrates how to call this template in a commit script:

```
<xsl:call-template name="jcs:emit-change">
  <xsl:with-param name="content">
    <term>
      <name>very-last</name>
      <junos:comment>
        <xsl:text>This term was added by a commit script</xsl:text>
      </junos:comment>
      <then>
        <accept/>
      </then>
    </term>
  </xsl:with-param>
</xsl:call-template>
```

When you issue the **show firewall** configuration mode command, the following output appears:

```
[edit]
user@host# show firewall
family inet {
  term very-last {
    /* This term was added by a commit script */
    then accept;
  }
}
```

- Related Documentation**
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)

## jcs:grep Template

---

**XSLT Syntax**

```
<xsl:call-template name="jcs:grep">
  <xsl:with-param name="filename" select="filename" />
  <xsl:with-param name="pattern" select="pattern" />
</xsl:call-template>
```

**SLAX Syntax**

```
call jcs:grep($filename=filename, $pattern=pattern);
```

**Description**      Search the given input file for all instances matching the specified regular expression and write the matching strings and corresponding lines to the result tree. The pattern is matched to each line of the file. The template does not support matching a pattern spanning multiple lines.

If the regular expression contains a syntax error, the template generates an error for every line of the file. For each match, the template adds a **<match>** element, which contains **<input>** and **<output>** child tags, to the result tree. The template writes the matching string to the **<output>** element and writes the corresponding matching line to the **<input>** element.

```
<match> {
  <input>
  <output>
}
```

Starting in Junos OS Release 11.1, if an absolute path is not specified for the input file, the default path is relative to the user's home directory for op scripts, and it is relative to the **/var/tmp/** directory for commit scripts and for event scripts that are enabled at the **[edit event-options event-script]** hierarchy level. For event scripts that are enabled at the **[edit system scripts]** hierarchy level, the default path is relative to the top-level directory, **/**.

**Parameters**      **filename**—Absolute or relative path and filename of the file to search.

Starting in Junos OS Release 11.1, if you do not specify an absolute path, the path is relative to the user's home directory for op scripts, and it is relative to the **/var/tmp/** directory for commit scripts and for event scripts that are enabled at the **[edit event-options event-script]** hierarchy level. For event scripts that are enabled at the **[edit system scripts]** hierarchy level, the default path is relative to the top-level directory, **/**.

**pattern**—Regular expression.

**Usage Examples**      ["Example: Searching Files Using an Op Script" on page 510](#)

**Related Documentation**

- [Example: Searching Files Using an Op Script on page 510](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [SLAX Operators on page 83](#)
- [regex\(\) Function \(SLAX and XSLT\) on page 216](#)

## jcs:load-configuration Template

<b>XSLT Syntax</b>	<pre> &lt;xsl:call-template name="jcs:load-configuration"&gt;   &lt;xsl:with-param name="action" select="(merge   override   replace)"/&gt;   &lt;xsl:with-param name="commit-options" select="node-set"/&gt;   &lt;xsl:with-param name="configuration" select="configuration-data"/&gt;   &lt;xsl:with-param name="connection" select="connection-handle"/&gt;   &lt;xsl:with-param name="rollback" select="number"/&gt; &lt;/xsl:call-template&gt; </pre>
<b>SLAX Syntax</b>	<pre> call jcs:load-configuration(\$action="(merge   override   replace)", \$commit-options=node-set, \$configuration=configuration-data, \$connection=connection-handle, \$rollback=number); </pre>
<b>Description</b>	<p>Make structured changes to the Junos OS configuration using an op script. When called, the template locks the configuration database, loads the configuration changes, commits the configuration, and then unlocks the configuration database.</p> <p>The <b>jcs:load-configuration</b> template makes changes to the configuration in <b>configure exclusive</b> mode. In this mode, Junos OS locks the candidate <i>global</i> configuration for as long as the script accesses the shared database and makes changes to the configuration without interference from other users.</p> <p>If another user is currently editing the configuration in <b>configure exclusive</b> mode or if the database is already locked when the template is called, the call fails. In addition, if there are existing, uncommitted changes to the configuration when the template is called, the commit will fail. If the template call is successful but the commit fails, Junos OS discards the uncommitted changes and rolls back the configuration.</p>
<b>Parameters</b>	<p><b>action</b>—Specifies how to load the configuration changes with respect to the candidate configuration. The following options are supported:</p> <ul style="list-style-type: none"> <li>• <b>merge</b>—Combine the candidate configuration and the incoming configuration changes. If the candidate configuration and the incoming configuration contain conflicting statements, the incoming statements override those in the candidate configuration.</li> <li>• <b>override</b>—Replace the entire candidate configuration.</li> <li>• <b>replace</b>—Replace existing statements in the candidate configuration with the tags of the same name that are marked with <b>replace</b>: in the incoming configuration. If there is no existing statement of the same name in the candidate configuration, the statement is added to the candidate configuration.</li> </ul> <p><b>commit-options</b>—Node set defining options that customize the commit command. The default value is null. Supported commit options are:</p> <ul style="list-style-type: none"> <li>• <b>check</b>—Check the correctness of the candidate configuration syntax, but do not commit the changes.</li> <li>• <b>force-synchronize</b>—Force the commit on the other Routing Engine (ignore any warnings).</li> </ul>

- **log**—Write the specified message to the commit log. This is identical to the CLI configuration mode command **commit comment**.
- **synchronize**—Synchronize the commit on both Routing Engines.

**configuration**—XML configuration changes. The configuration changes are incorporated into the candidate configuration as specified by the action parameter. Starting in Junos OS Release 12.2, you can also supply a NULL configuration. If the configuration data value is NULL, the template performs a simple commit of the candidate configuration.

**connection**—Connection handle generated by a call to the **jcs:open()** function.

**rollback**—Return to a previously committed configuration. Specify the rollback number of the configuration, and the configuration you specify is loaded from the associated file. The software saves the last 50 committed configurations. The rollback parameter is available starting in Junos OS Release 12.2.

**Usage Examples** The following example calls the **jcs:load-configuration** template to modify the configuration to disable an interface. The interface name is supplied by the user and stored in the variable **interface-name**. All of the values required for the **jcs:load-configuration** template are defined as variables, which are then passed into the template as arguments.

In this example, the configuration data that includes the changes to the configuration are stored in the variable **disable**. This is the value used for the **configuration** parameter of the **jcs:load-configuration** template. The **load-action** variable is initialized to **merge**, which merges the configuration changes in the **disable** variable with the candidate configuration. This is the equivalent of the CLI configuration mode command **load merge**.

The **options** variable uses the **:=** operator to create a node-set, which is passed to the template as the value of the **commit-options** parameter. This example uses the **synchronize** commit option. If the commit succeeds, it will commit the configuration changes on both Routing Engines. The **log** tag is also included to add the description of the commit to the commit log file for future reference.

The call to the **jcs:open()** function opens a connection with the Junos OS management process (mgd) and returns a connection handle that is stored in the **conn** variable. All of the defined variables are passed as arguments to the **jcs:load-configuration** template at the time that it is called.

SLAX syntax:

```
var $disable = {
  <configuration> {
    <interfaces> {
      <interface> {
        <name> $interface-name;
        <disable>;
      }
    }
  }
}
```

```

var $load-action = "merge";
var $options := {
  <commit-options> {
    <synchronize>;
    <log> "disabling interface on both routing engines";
  }
}
var $conn = jcs:open();

var $disable-results := {
  call jcs:load-configuration($action=$load-action, $commit-options=$options,

    $configuration = $disable, $connection = $conn);
}
if ($disable-results//xnm:error) {
  for-each ($disable-results//xnm:error) {
    <output> message;
  }
}
var $close-results = jcs:close($conn);

```

The `:=` operator copies the results of the `jcs:load-configuration` template call to a temporary variable and runs the `node-set` function on that variable. The `:=` operator ensures that the `disable-results` variable is a node-set rather than a result tree fragment so that the script can access the contents. The `if` code block is included to output any error messages that may indicate a problem in committing the configuration. The `jcs:close` function closes the connection.

In XSLT, the code corresponding to the SLAX call to `jcs:load-configuration` template is:

```

<xsl:variable name="disable-results-temp">
  <xsl:call-template name="jcs:load-configuration">
    <xsl:with-param name="action" select="$load-action"/>
    <xsl:with-param name="commit-options" select="$options"/>
    <xsl:with-param name="configuration" select="$disable"/>
    <xsl:with-param name="connection" select="$conn"/>
  </xsl:call-template>
</xsl:variable>

<xsl:variable xmlns ext="http://xmlsoft.org/XSLT/namespace" \
  name="disable-results" select="ext:node-set($disable-results-temp)"/>

```

- Related Documentation**
- [Example: Changing the Configuration Using SLAX and XSLT Op Scripts on page 467](#)
  - [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [SLAX Operators on page 83](#)
  - [close\(\) Function \(SLAX and XSLT\) on page 192](#)
  - [open\(\) Function \(SLAX and XSLT\) on page 209](#)

## jcs:statement Template

XSLT Syntax	<pre> &lt;xsl:call-template name="jcs:statement"&gt;   &lt;xsl:with-param name="dot" select="expression"/&gt; &lt;/xsl:call-template&gt; </pre>
SLAX Syntax	<pre> call jcs:statement(\$dot=expression); </pre>

- Description** Generate a `<statement>` element suitable for inclusion in an `<xnm:error>` or `<xnm:warning>` element. This location defaults to “.”, the current position in the XML hierarchy. If the error is not at the current position in the XML hierarchy, you can alter the default when you call the template by including a valid XPath expression either for the **dot** parameter in SLAX scripts or for the **select** attribute of the **dot** parameter in XSLT scripts.
- Parameters** **dot**—XPath expression specifying the hierarchy level. The default location is the position in the XML hierarchy that the script is currently evaluating. You can alter the default when you call the template by including a valid XPath expression either for the **dot** parameter in SLAX scripts or for the **select** attribute of the **dot** parameter in XSLT scripts.
- Usage Examples** The following example demonstrates how to call the **jcs:statement** template in a commit script:

```
<xnm:error>
  <xsl:call-template name="jcs:edit-path"/>
  <xsl:call-template name="jcs:statement">
    <xsl:with-param name="dot" select="mtu"/>
  </xsl:call-template>
  <message>
    <xsl:text>SONET interfaces must have a minimum MTU of </xsl:text>
    <xsl:value-of select="$min-mtu"/>
    <xsl:text>.</xsl:text>
  </message>
</xnm:error>
```

When you commit a configuration that includes a SONET/SDH interface with a maximum transmission unit (MTU) setting less than a specified minimum, the `<xnm:error>` element results in the following CLI output:

```
[edit]
user@host# commit
[edit interfaces interface so-1/2/3]
'mtu 576;' # mtu statement generated by the jcs:statement template
SONET interfaces must have a minimum MTU of 2048.
error: 1 error reported by commit scripts
error: commit script failure
```

The test of the MTU setting is not performed in the `<xnm:error>` element. For the full example, see [“Example: Imposing a Minimum MTU Setting” on page 382](#).

- Related Documentation**
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
  - [call on page 93](#)
  - [xnm:error \(Junos XML\) on page 428](#)
  - [xnm:warning \(Junos XML\) on page 430](#)
  - [xsl:call-template on page 40](#)

## PART 7

# Commit Scripts

- [Commit Scripts Overview on page 241](#)
- [Creating and Executing Commit Scripts on page 251](#)
- [Generating a Custom Warning, Error, or System Log Message Using Commit Scripts on page 269](#)
- [Generating Persistent or Transient Configuration Changes Using Commit Scripts on page 287](#)
- [Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Commit Script Examples on page 331](#)
- [Junos XML and XSLT Tag Elements Used in Commit Scripts on page 427](#)
- [Troubleshooting Commit Scripts on page 433](#)





## CHAPTER 15

# Commit Scripts Overview

- [Commit Script Overview on page 241](#)
- [Advantages of Using Commit Scripts on page 242](#)
- [How Commit Scripts Work on page 243](#)

### Commit Script Overview

---

Junos OS commit scripts enable users to customize the validation process of their configurations in accordance with their own practices and policies and enforce custom configuration rules during the commit process. When a candidate configuration is committed, it is inspected by each active commit script. If a configuration violates your custom rules, the script can instruct Junos OS to take appropriate action. A commit script can:

- Generate and display custom warning messages to the user
- Generate and log custom system log (syslog) messages
- Change the configuration to conform to the custom business rules
- Generate a commit error and halt the commit operation

Commit scripts are based on the Junos XML management protocol and the Junos XML API. The Junos XML management protocol is an XML based RPC mechanism, and the Junos XML API is an XML representation of Junos OS configuration statements and operational mode commands.

Commit scripts can be written in Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX). The Junos XML API defines an XML equivalent for all statements in the Junos configuration hierarchy. Commit scripts use XML Path Language (XPath) to locate the configuration objects to be inspected and automation script constructs to specify the actions to perform on the configuration objects. The actions can change the configuration or generate messages about it.

Additionally, you can create *macros*, which allow you to create custom configuration syntax that simplifies the task of configuring a device running Junos OS. By itself, your custom syntax has no operational impact on the device. A corresponding commit script macro uses your custom syntax as input data for generating standard Junos OS configuration statements that execute your intended operational impact.

To view the device's current configuration in Extensible Markup Language (XML) in the command-line interface (CLI), issue the **show configuration | display xml** command in operational mode. To view your configuration in commit-script-style XML, issue the **show configuration | display commit-scripts view** command. Commit-script-style XML view displays the configuration in the format that would be input to a commit script.

**Related  
Documentation**

- [Junos XML Management Protocol and Junos XML API Overview on page 11](#)
- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
- [SLAX Overview on page 57](#)
- [XSLT Overview on page 17](#)
- [XPath Overview on page 19](#)

---

## Advantages of Using Commit Scripts

Reducing human error in a network configuration can significantly improve network uptime. Commit scripts enable you to control operational practices and enforce operational policy, thereby decreasing the possibility of human error. Restricting device configurations in accordance with custom design rules can vastly improve network reliability.

Consider the following examples of actions you can perform with commit scripts:

- Basic sanity test—Ensure that the **[edit interfaces]** and **[edit protocols]** hierarchies have not been accidentally deleted.
- Consistency check—Ensure that every T1 interface configured at the **[edit interfaces]** hierarchy level is also configured at the **[edit protocols rip]** hierarchy level.
- Dual Routing Engine configuration test—Ensure that the **re0** and **re1** configuration groups are set up correctly. When you use configuration groups, the inherited values can be overridden in the target configuration. A commit script can determine if an individual target configuration element is blocking proper inheritance of the configuration group settings.
- Interface density—Ensure that a channelized interface does not have too many channels configured.
- Link scaling—Ensure that SONET/SDH interfaces never have a maximum transmission unit (MTU) size less than 4 kilobytes (KB).
- Import policy check—Ensure that an interior gateway protocol (IGP) does not use an import policy that imports the full routing table.
- Cross-protocol checks—Ensure that all LDP-enabled interfaces are configured for an IGP, or ensure that all IGP-enabled interfaces are configured for LDP.
- IGP design check—Ensure that Level 1 IS-IS routers are never enabled.

When a candidate configuration does not adhere to your design rules, a commit script can instruct Junos OS to generate custom warnings, system log messages, or error messages that block the commit operation from succeeding. In addition, the commit

script can change the configuration in accordance with your rules and then proceed with the commit operation.

Consider a network design that requires every interface on which the International Organization for Standardization (ISO) family of protocols is enabled to also have MPLS enabled. At commit time, a commit script inspects the configuration and issues an error if this requirement is not met. This error causes the commit operation to fail and forces the user to update the configuration to comply.

Instead of an error, the commit script can issue a warning about the configuration problem and then automatically correct it by changing the configuration to enable MPLS on all interfaces. A system log message can also be generated, indicating that corrective action was taken.

Another option is to define a macro that enables ISO protocols and MPLS when the macro is applied to an interface. Configuring this macro simplifies the configuration task while ensuring that both protocols are configured together.

Finally, you can have the commit script correct the configuration using a *transient change*. In our example, a transient change allows MPLS to always be enabled on ISO-enabled interfaces without having the configuration statements appear in the candidate configuration.



**NOTE:** Transient changes cause a change to be generated in the *checkout configuration* but not in the candidate configuration. The checkout configuration is the configuration database that is checked for standard Junos OS syntax just before a configuration becomes active. This means transient changes are not saved in the configuration if the associated commit script is deleted or deactivated. The `show configuration | display commit-scripts` command displays all the statements that are in the configuration, including statements that were generated by transient changes. For more information, see [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts” on page 287](#).

**Related Documentation**

- [Commit Script Overview on page 241](#)

## How Commit Scripts Work

Commit scripts contain instructions that enforce custom configuration rules and are invoked during the commit process before the standard Junos OS validity checks are performed. You enable commit scripts by listing the names of one or more commit script files at the `[edit system scripts commit]` hierarchy level. These files must be added to the appropriate commit script directory on the device.

When you perform a commit operation, Junos OS executes each script in turn, passing the information in the post-inheritance candidate configuration to the scripts. The script inspects the configuration, performs the necessary tests and validations, and generates a set of instructions for performing certain actions. After all commit scripts have been

executed, Junos OS then processes all of the scripts' instructions. If the commit process is not halted by a commit script, then Junos OS applies all the commit script changes and performs its final inspection of the checkout configuration.

Commit script actions can include generating error, warning, and system log messages. If errors are generated, the commit operation fails and the candidate configuration remains unchanged. This is the same behavior that occurs with standard commit errors. Commit scripts can also generate changes to the system configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

The following sections discuss several important concepts related to the commit script input and output:

- [Commit Script Input on page 244](#)
- [Commit Script Output on page 245](#)
- [Commit Scripts and the Junos OS Commit Model on page 247](#)

## Commit Script Input

The input for a commit script is the post-inheritance candidate configuration in Junos XML API format. The term *post-inheritance* means that all configuration group values have been inherited by their targets in the candidate configuration and that the inactive portions of the configuration have been removed. For more information about configuration groups, see the *CLI User Guide*.

When you issue the **commit** command, Junos OS automatically generates the candidate configuration in XML format and reads it into the management (mgd) process, at which time the input is evaluated by any commit scripts.

To display the XML format of the post-inheritance candidate configuration in the CLI, issue the **show | display commit-scripts view** command.

[edit]

```
user@host# show | display commit-scripts view
```

To display all configuration groups data, including script-generated changes to the groups, issue the **show groups | display commit-scripts** command.

[edit]

```
user@host# show groups | display commit-scripts
```

## Commit Script Output

During the commit process, enabled commit scripts are executed sequentially, and the commit script output, or instruction set, is provided to Junos OS. After all commit scripts have been executed, Junos OS then processes all of the scripts' instructions.

Commit script actions can include generating warning, error, and system log messages, and making persistent and transient changes to the configuration. [Table 31 on page 245](#) briefly outlines the various elements, templates, and functions that commit scripts can use to instruct Junos OS to perform various actions during the commit process. In some cases, there are multiple ways to perform the same action. Because SLAX and XSLT scripts return a result tree, output elements like `<syslog><message>` that are present in SLAX and XSLT scripts are added directly into the result tree.

**Table 31: Commit Scripts Actions and Output**

Commit Script Output	SLAX / XSLT	Python
Generate a warning message to the committing user.	<code>&lt;xnm:warning&gt;</code>	<code>jcs.emit_warning()</code>
Generate an error message and cause the commit operation to fail.	<code>&lt;xnm:error&gt;</code>	<code>jcs.emit_error()</code>
Generate a system log message.	<code>jcs.syslog()</code>  <code>&lt;syslog&gt;&lt;message&gt;</code>	<code>jcs.syslog()</code>
Generate a persistent change to the configuration.	<code>&lt;change&gt;</code>	<code>emit_change(content, 'change', format)</code>
Generate a transient change to the configuration.	<code>&lt;transient-change&gt;</code>	<code>emit_change(content, 'transient-change', format)</code>
Generate a persistent change relative to the current context node as defined by an XPath expression.	<b>XSLT</b> <code>&lt;xsl:call-template name="jcs:emit-change"&gt;</code> <code>  &lt;xsl:with-param name="content"&gt;</code>  <b>SLAX</b> <code>call jcs:emit-change() {</code> <code>  with \$content = {    }</code> <code>}</code>	—

Table 31: Commit Scripts Actions and Output (*continued*)

Commit Script Output	SLAX / XSLT	Python
Generate a transient change relative to the current context node as defined by an XPath expression.	<b>XSLT</b> <pre>&lt;xsl:call-template name="jcs:emit-change"&gt;   &lt;xsl:with-param name="tag" select="'transient-change'"/&gt;   &lt;xsl:with-param name="content"&gt;</pre> <b>SLAX</b> <pre>call jcs:emit-change() {   with \$tag = "transient-change";   with \$content = {   } }</pre>	—
Generate a warning message in conjunction with a configuration change. You can use this set of tags to generate a notification that the configuration has been changed.	<b>XSLT</b> <pre>&lt;xsl:call-template name="jcs:emit-change"&gt;   &lt;xsl:with-param name="message"&gt;     &lt;xsl:text&gt;</pre> <b>SLAX</b> <pre>call jcs:emit-change() {   with \$message = {     expr "message";   } }</pre>	jcs.emit_warning()

Junos OS processes this output and performs the appropriate actions. Errors and warnings are passed back to the Junos OS CLI or to a Junos XML protocol client application. The presence of an error automatically causes the commit operation to fail. Persistent and transient changes are loaded into the appropriate configuration database.

To test the output of error, warning, and system log messages from commit scripts, issue the **commit check | display xml** command.

```
[edit]
user@host# commit check | display xml
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command.

```
[edit]
user@host# commit check | display detail
```



**NOTE:** System log messages do not appear in the trace output, so you cannot use the commit check operation to test script-generated system log messages. Furthermore, system log messages are written to the system log during a commit operation, but not during a commit check operation.

#### Related Documentation

- *Example: Protecting the Junos OS Configuration from Modification or Deletion.*

- [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)

## Commit Scripts and the Junos OS Commit Model

Junos OS uses a commit model to update the device's configuration. This model allows you to make a series of changes to a candidate configuration without affecting the operation of the device. When the changes are complete, you can commit the configuration. The commit operation saves the candidate configuration changes into the current configuration.

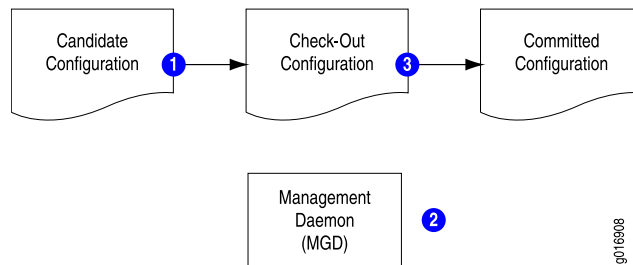
When you commit a set of changes in the candidate configuration, two methods are used to forward these changes to the current configuration:

- Standard commit model—Used when no commit scripts are active on the device.
- Commit script model—Incorporates commit scripts into the commit model.

### Standard Commit Model

In the standard commit model, the management (mgd) process validates the candidate configuration based on standard Junos OS validation rules. If the configuration file is valid, it becomes the current active configuration. [Figure 4 on page 247](#) and the accompanying discussion explain how the standard commit model works.

**Figure 4: Standard Commit Model**



In the standard commit model, the software performs the following steps:

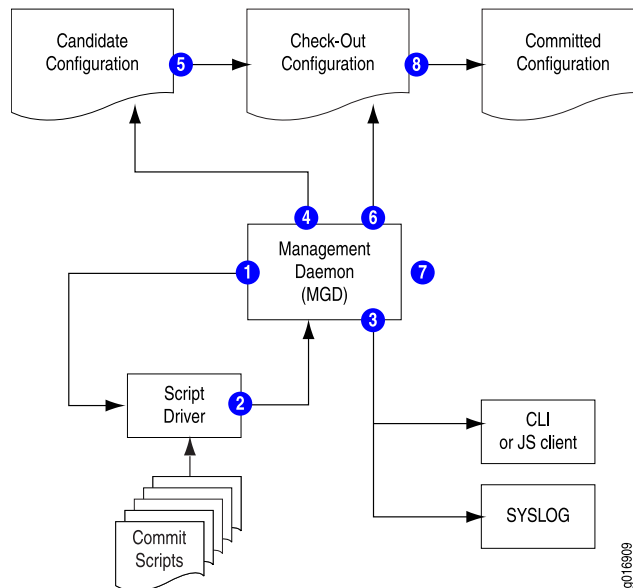
1. When the candidate configuration is committed, it is copied to become the checkout configuration.
2. The mgd process validates the checkout configuration.
3. If no error occurs, the checkout configuration is copied as the current active configuration.

### Commit Model with Commit Scripts

When commit scripts are added to the standard commit model, the process becomes more complex. The mgd process first passes an XML-formatted checkout configuration to a script driver, which handles the verification of the checkout configuration by the commit scripts. When verification is complete, the script driver returns an *action file* to the mgd process. The mgd process follows the instructions in the action file to update the candidate and checkout configurations, issue messages to the CLI or client application,

and write information to the system log as required. After processing the action file, the mgd process performs the standard Junos OS validation. [Figure 5 on page 248](#) and the accompanying discussion explain this process.

**Figure 5: Commit Model with Commit Scripts Added**



In the commit script model, Junos OS performs the following steps:

1. When the candidate configuration is committed, the mgd process sends the XML-formatted candidate configuration to the script driver.
2. Each enabled commit script is invoked against the candidate configuration, and each script can generate a set of actions for the mgd process to perform. The actions are collected in an action file.
3. The mgd process performs the following actions for commit script error, warning, and system log messages in the action file:
  - error—The mgd process halts the commit process (that is, the commit operation fails), returns an error message to the CLI or Junos XML protocol client, and takes no further action.
  - warning—The mgd process forwards the message to the CLI or the Junos XML protocol client.
  - system log message—The mgd process forwards the message to the system log process.
4. If the action file includes any persistent changes, the mgd process loads the requested changes into the candidate configuration.
5. The candidate configuration is copied to become the checkout configuration.
6. If the action file includes any transient changes, the mgd process loads the requested changes into the checkout configuration.



7. The mgd process validates the checkout configuration.
8. If there are no validation errors, the checkout configuration is copied to become the current active configuration.



**NOTE:** Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

Changes that are made to the candidate configuration during the commit operation are not evaluated by the custom rules during that commit operation. However, persistent changes are maintained in the candidate configuration and are evaluated by the custom rules during subsequent commit operations. For more information about how commit scripts change the candidate configuration, see [“Avoiding Potential Conflicts When Using Multiple Commit Scripts” on page 256](#).

Transient changes are never evaluated by the custom rules in commit scripts, because they are made to the checkout configuration only after the commit scripts have evaluated the candidate configuration and the candidate is copied to become the checkout configuration. To remove a transient change from the configuration, remove, disable, or deactivate the commit script (as discussed in [“Controlling Execution of Commit Scripts During Commit Operations” on page 261](#)), or comment out the code that generates the transient change.

For more information about differences between persistent and transient changes, see [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts” on page 287](#).

**Related  
Documentation**

- [Controlling Execution of Commit Scripts During Commit Operations on page 261](#)
- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 256](#)



# Creating and Executing Commit Scripts

- [Required Boilerplate for Commit Scripts on page 251](#)
- [XML Syntax for Common Commit Script Tasks on page 253](#)
- [Design Considerations for Commit Scripts on page 254](#)
- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 256](#)
- [Line-by-Line Explanation of Sample Commit Scripts on page 257](#)
- [Controlling Execution of Commit Scripts During Commit Operations on page 261](#)
- [Configuring Checksum Hashes for a Commit Script on page 264](#)
- [Executing Large Commit Scripts on page 265](#)
- [Example: Retrieving the Pre-Inheritance Candidate Configuration in a Commit Script on page 265](#)

## Required Boilerplate for Commit Scripts

---

Junos OS commit scripts can be written in Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python. Commit scripts must include the necessary boilerplate required for that script language for both basic script functionality as well as any optional functionality used within the script such as the Junos OS extension functions and named templates. This topic provides standard boilerplate that can be used in XSLT, SLAX, and Python scripts.

SLAX and XSLT commit scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

### XSLT Boilerplate for Commit Scripts

The XSLT commit script boilerplate is as follows:

```
1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7      <xsl:import href="../import/junos.xsl"/>
```

```

8      <xsl:template match="configuration">
          <!-- ... insert your code here ... -->
9      </xsl:template>
10     </xsl:stylesheet>

```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI). This PI specifies that the code is written in XML using version 1.0. The XML PI, if present, must be the first noncomment token in the script file.

```
1    <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2    <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in commit scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```

3        xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4        xmlns:junos="http://xml.juniper.net/junos/*/junos"
5        xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6        xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">

```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of the Junos OS. The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in [“Understanding Named Templates in Junos OS Automation Scripts” on page 225](#).

```
7        <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `<configuration>` element, which is the node selected by the `<xsl:template match="/">` template, contained in the `junos.xsl` import file. The `<xsl:template match="configuration">` element allows you to exclude the `/configuration/` root element from all XPath expressions in the script and begin XPath expressions with the top Junos OS hierarchy level. For more information, see [“XPath Overview” on page 19](#).

```
8        <xsl:template match="configuration">
```

Add your code between Lines 8 and 9.

Line 9 closes the template.

```
9        </xsl:template>
```

Line 10 closes the style sheet and the commit script.

```
10     </xsl:stylesheet>
```

### SLAX Boilerplate for Commit Scripts

The corresponding SLAX commit script boilerplate is as follows:

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    /*

```

```

        * insert your code here
    */
}

```

### Python Boilerplate for Commit Scripts

Python commit scripts do not have a required boilerplate, but they must import any objects that are used in the script. Python automation scripts import the **Junos\_Context** dictionary when the script must access information about the script execution environment. The **Junos\_Configuration** object is required when the commit script must read and manipulate the post-inheritance candidate configuration. Python automation scripts that include the **import jcs** statement can use Junos OS extension functions and Junos OS named template functionality in the script.

```

from junos import Junos_Context
from junos import Junos_Configuration
import jcs

if __name__ == '__main__':

```

Python automation scripts do not need to include an interpreter directive line (**#!/usr/bin/env python**) at the start of the script. However, the program will still execute correctly if one is present.

### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)

## XML Syntax for Common Commit Script Tasks

Junos OS commit scripts can generate persistent or transient changes to the configuration during the commit process. A commit script can perform common configuration tasks by adding the appropriate attribute to a specific XML tag in the configuration data. [Table 32 on page 253](#) summarizes the tasks and the syntax for each task.

**Table 32: XML Syntax for Common Commit Script Tasks**

Action	Syntax	Example
Add a data element	normal XML	<pre>&lt;address&gt;   &lt;name&gt;192.168.1.1&lt;/name&gt; &lt;/address&gt;</pre>
Remove the inactive tag from a statement	active="active"	<pre>&lt;address active="active"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>
Delete a data element	delete="delete"	<pre>&lt;address delete="delete"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>
Add the inactive tag to a statement	inactive="inactive"	<pre>&lt;address inactive="inactive"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>

Table 32: XML Syntax for Common Commit Script Tasks (*continued*)

Action	Syntax	Example
Insert a new ordered data element	<code>insert="(before   after)" name="reference-value"</code>	<pre>&lt;address insert="before" name="192.168.1.5/30"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>
Add the protect tag to a statement or node to prevent configuration changes to that element	<code>protect="protect"</code>	<pre>&lt;address protect="protect"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>
Rename a statement	<code>rename="rename" name="new-name"</code>	<pre>&lt;address rename="rename" name="192.168.1.1/30"&gt;   &lt;name&gt;192.168.1.5/30&lt;/name&gt; &lt;/address&gt;</pre>
Replace a node or statement in the hierarchy	<code>replace="replace"</code>	<pre>&lt;system&gt;   &lt;services     replace="replace"&gt;     [...]   &lt;/services&gt; &lt;/system&gt;</pre>
Unprotect a statement or node in the hierarchy	<code>unprotect="unprotect"</code>	<pre>&lt;address unprotect="unprotect"&gt;   &lt;name&gt;192.168.1.1/30&lt;/name&gt; &lt;/address&gt;</pre>
Annotate a configuration statement with a comment	<code>&lt;junos:comment&gt;</code>	<pre>&lt;system&gt;   &lt;junos:comment&gt;     /* added by username   */   &lt;/junos:comment&gt;   &lt;services&gt;     [...]   &lt;/services&gt; &lt;/system&gt;</pre>

## Design Considerations for Commit Scripts

After you have some experience looking at Junos OS configuration data in XML, creating commit scripts is fairly straightforward. This section provides some advice and common patterns for developing commit scripts using XSLT.

XSLT is an interpreted language, making performance an important consideration. For best performance, minimize node traversals and testing performed on each node. When possible, use the **select** attribute on a recursive `<xsl:apply-templates>` invocation to limit the portion of the document hierarchy being visited.

For example, the following **select** attribute limits the nodes to be evaluated by specifying SONET/SDH interfaces that have the **inet** (IPv4) protocol family enabled:

```
<xsl:apply-templates select="interfaces/interface[starts-with(name, 'so-')
and unit/family/inet]"/>
```

The following example contains two `<xsl:apply-templates>` instructions that limit the scope of the script to the `import` statements configured at the `[edit protocols ospf]` and `[edit protocols isis]` hierarchy levels:

```
<xsl:template match="configuration">
  <xsl:apply-templates select="protocols/ospf/import"/>
  <xsl:apply-templates select="protocols/isis/import"/>
  <!-- ... body of template ... -->
</xsl:template>
```

In an interpreted language, doing anything more than once can affect performance. If the script needs to reference a node or node set repeatedly, make a variable that holds the node set, and then make multiple references to the variable. For example, the following variable declaration creates a variable called `mpls` that resolves to the `[edit protocols mpls]` hierarchy level. This allows the script to traverse the `/protocols/` hierarchy searching for the `mpls/` node only once.

```
<xsl:variable name="mpls" select="/protocols/mpls"/>
<xsl:choose>
  <xsl:when test="$mpls/path-mtu/allow-fragmentation">
    <!-- ... -->
  </xsl:when>
  <xsl:when test="$mpls/hop-limit > 40">
    <!-- ... -->
  </xsl:when>
</xsl:choose>
```

Variables are also important when using `<xsl:for-each>` instructions, because the current context node examines each node selected by the `<xsl:for-each>` instruction. For example, the following script uses multiple variables to store and refer to values as the `<xsl:for-each>` instruction evaluates the E1 interfaces that are configured on all channelized STM1 (cstm1-) interfaces:

```
<xsl:param name="limit" select="16"/>
<xsl:template match="configuration">
  <xsl:variable name="interfaces" select="interfaces"/>
  <xsl:for-each select="$interfaces/interface[starts-with(name, 'cstm1-')]">
    <xsl:variable name="triple" select="substring-after(name, 'cstm1-')"/>
    <xsl:variable name="e1name" select="concat('e1-', $triple)"/>
    <xsl:variable name="count"
      select="count($interfaces/interface[starts-with(name, $e1name)])/>
    <xsl:if test="$count > $limit">
      <xnm:error>
        <edit-path>[edit interfaces]</edit-path>
        <statement><xsl:value-of select="name"/></statement>
        <message>
          <xsl:text>E1 interface limit exceeded on CSTM1 IQ PIC.
          </xsl:text>
          <xsl:value-of select="$count"/>
          <xsl:text> E1 interfaces are configured, but only
          </xsl:text>
          <xsl:value-of select="$limit"/>
          <xsl:text> are allowed.</xsl:text>
        </message>
      </xnm:error>
    </xsl:if>
  </xsl:for-each>
</xsl:template>
```

If you channelize a `cstm1-0/1/0` interface into 17 E1 interfaces, the script causes the following error message to appear when you issue the `commit` command. (For more

information about this example, see [“Example: Limiting the Number of E1 Interfaces” on page 388.](#))

[edit]

```
user@host# commit
```

```
[edit interfaces]
```

```
'cstm1-0/1/0'
```

```
E1 interface limit exceeded on CSTM1 IQ PIC.
```

```
17 E1 interfaces are configured, but only 16 are allowed.
```

```
error: 1 error reported by commit scripts
```


```
error: commit script failure
```

Related •  
Documentation

## Avoiding Potential Conflicts When Using Multiple Commit Scripts

When you use multiple commit scripts, each script evaluates the original candidate configuration file. Changes made by one script are not evaluated by the other scripts. This means that conflicts between scripts might not be resolved when the scripts are first applied to the configuration. The commit scripts are executed in the order they are listed at the `[edit system scripts commit]` hierarchy level, as illustrated in [Figure 6 on page 256.](#)

**Figure 6: Configuration Evaluation by Multiple Commit Scripts**



As an example of a conflict between commit scripts, suppose that commit script **A.xsl** is created to ensure that the device uses the domain name server with IP address 192.168.0.255. Later, the DNS server's address is changed to 192.168.255.255 and a second script, **B.xsl**, is added to check that the device uses the DNS server with that address. However, script **A.xsl** is not removed or disabled.

Because each commit script evaluates the original candidate configuration, the final result of executing both scripts **A.xsl** and **B.xsl** depends on which DNS server address is configured in the original candidate configuration. If the now outdated address of 192.168.0.255 is configured, script **B.xsl** changes it to 192.168.255.255. However, if the



correct address of 192.168.255.255 is configured, script **A.xsl** changes it to the incorrect value 192.168.0.255.

As another example of a potential conflict between commit scripts, suppose that a commit script protects a hierarchy using the **protect** attribute. If a second commit script attempts to modify or delete the hierarchy or the statements within the hierarchy, Junos OS issues a warning during the commit process and prevents the configuration change.

Exercise care to ensure that you do not introduce conflicts between scripts like those described in the examples. As a method of checking for conflicts with persistent changes, you can issue two separate **commit** commands.

**Related Documentation**

- [How Commit Scripts Work on page 243](#)

## Line-by-Line Explanation of Sample Commit Scripts

- [Applying a Change to SONET/SDH Interfaces on page 257](#)
- [Applying a Change to ISO-Enabled Interfaces on page 258](#)

### Applying a Change to SONET/SDH Interfaces

The following XSLT commit script applies a transient change to each interface whose name begins with **so-**, setting the encapsulation to **ppp**. For information about transient changes, see “[Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts](#)” on page 287. For a SLAX version of this example, see “[Example: Generating a Transient Change](#)” on page 306.

```

1  <?xml version="1.0"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7      <xsl:import href="../import/junos.xsl"/>

8      <xsl:template match="configuration">
9          <xsl:for-each select="interfaces/interface[starts-with(name, 'so-')
10              \
11                  and unit/family/inet]">
12              <transient-change>
13                  <interfaces>
14                      <interface>
15                          <name><xsl:value-of select="name"/></name>
16                          <encapsulation>ppp</encapsulation>
17                      </interface>
18                  </interfaces>
19              </transient-change>
20          </xsl:for-each>
21      </xsl:template>
22  </xsl:stylesheet>

```

Lines 1 through 8 are boilerplate as described in “[Required Boilerplate for Commit Scripts](#)” on page 251 and are omitted here for brevity.

Line 9 is an `<xsl:for-each>` programming instruction that examines each interface node whose names starts with 'so-' and that has **family inet** enabled on any logical unit. (It appears here on two lines only for brevity.)

```

9      <xsl:for-each select="interfaces/interface[starts-with(name, 'so-')
\
      and unit/family/inet]">

```

Line 10 is the open tag for a transient change. The possible contents of the `<transient-change>` element are the same as the contents of the `<configuration>` tag element in the Junos XML protocol operation `<load-configuration>`.

```

10      <transient-change>

```

Lines 11 through 16 represent the content of the transient change. The encapsulation is set to **ppp**.

```

11      <interfaces>
12      <interface>
13      <name><xsl:value-of select="name"/></name>
14      <encapsulation>ppp</encapsulation>
15      </interface>
16      </interfaces>

```

Lines 17 through 19 close all open tags in this template.

```

17      </transient-change>
18      </xsl:for-each>
19      </xsl:template>

```

Line 20 closes the style sheet and the commit script.

```

20 </xsl:stylesheet>

```

## Applying a Change to ISO-Enabled Interfaces

The following sample XSLT script ensures that interfaces that are enabled for an International Organization for Standardization (ISO) protocol also have MPLS enabled and are included at the **[edit protocols mpls interface]** hierarchy level. For a SLAX version of this example, see ["Example: Controlling IS-IS and MPLS Interfaces" on page 366](#).

```

1  <?xml version="1.0"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7      <xsl:import href="../import/junos.xsl"/>
8
9      <xsl:template match="configuration">
10         <xsl:variable name="mpls" select="protocols/mpls"/>
11         <xsl:for-each select="interfaces/interface/unit[family/iso]">
12             <xsl:variable name="ifname" select="concat(..name, '.',
13 name)"/>
14             <xsl:if test="not(family/mpls)">
15                 <xsl:call-template name="jcs:emit-change">
16                     <xsl:with-param name="message">
17                         <xsl:text>
18                             Adding 'family mpls' to ISO-enabled interface
19                         </xsl:text>
20                     </xsl:with-param>
21                     <xsl:with-param name="content">

```

```

20             <family>
21                 <mpls/>
22             </family>
23         </xsl:with-param>
24     </xsl:call-template>
25 </xsl:if>
26 <xsl:if test="$mpls and not($mpls/interface[name = $ifname])">
27     <xsl:call-template name="jcs:emit-change">
28         <xsl:with-param name="message">
29             <xsl:text>Adding ISO-enabled interface </xsl:text>
30             <xsl:value-of select="$ifname"/>
31             <xsl:text> to [protocols mpls]</xsl:text>
32         </xsl:with-param>
33         <xsl:with-param name="dot" select="$mpls"/>
34         <xsl:with-param name="content">
35             <interface>
36                 <name>
37                     <xsl:value-of select="$ifname"/>
38                 </name>
39             </interface>
40         </xsl:with-param>
41     </xsl:call-template>
42 </xsl:if>
43 </xsl:for-each>
44 </xsl:template>
45 </xsl:stylesheet>

```

Lines 1 through 8 are boilerplate as described in [“Required Boilerplate for Commit Scripts” on page 251](#) and are omitted here for brevity.

Line 9 saves a reference to the **[edit protocols mpls]** hierarchy level so that it can be referenced in the following **for-each** loop.

```

9         <xsl:variable name="mpls" select="protocols/mpls"/>

```

Line 10 examines each interface unit (logical interface) on which ISO is enabled. The **select** stops at the **unit**, but the predicate limits the selection to only those units that contain an **<iso>** element nested under a **<family>** element.

```

10     <xsl:for-each select="interfaces/interface/unit[family/iso]">

```

Line 11 builds the interface name in a variable. First, the **name** attribute of the variable declaration is set to **ifname**. In Junos OS, an interface name is the concatenation of the device name, a period, and the unit number. At this point in the script, the context node is the unit number, because Line 10 changes the context to **interfaces/interface/unit**. The **../name** refers to the **<name>** element of the parent node of the context node, which is the device name (*type-fpc/pic/port*). The **"name"** token in the XPath expression refers to the **<name>** element of the context node, which is the unit number (*unit-number*). After the concatenation is performed, the XPath expression in Line 11 resolves to *type-fpc/pic/port.unit-number*. As the **<xsl:for-each>** instruction in Line 10 traverses the hierarchy and locates ISO-enabled interfaces, the interface names are recursively stored in the **ifname** variable.

```

11         <xsl:variable name="ifname" select="concat(../name, '.', name)"/>

```

Line 12 evaluates as true for each ISO-enabled interface that does not have MPLS enabled.

```

12     <xsl:if test="not(family/mpls)">

```

Line 13 calls the **jcs:emit-change** template, which is a helper or convenience template in the **junos.xsl** file. This template is discussed in [emit-change Template \(SLAX and XSLT\)](#) and [emit\\_change\(Python\)](#).

```
13          <xsl:call-template name="jcs:emit-change">
```

Lines 14 through 18 use the **message** parameter from the **jcs:emit-change** template. The **message** parameter is a shortcut you can use instead of explicitly including the **<warning>**, **<edit-path>**, and **<statement>** elements.

```
14          <xsl:with-param name="message">
15              <xsl:text>
16                  Adding 'family mpls' to ISO-enabled interface
17              </xsl:text>
18          </xsl:with-param>
```

Lines 19 through 23 use the **content** parameter from the **jcs:emit-change** template. The **content** parameter specifies the change to make, relative to the current context node.

```
19          <xsl:with-param name="content">
20              <family>
21                  <mpls/>
22              </family>
23          </xsl:with-param>
```

Lines 24 and 25 close the tags opened in Lines 13 and 12, respectively.

```
24          </xsl:call-template>
25      </xsl:if>
```

Line 26 tests whether MPLS is already enabled and if this interface is not configured at the **[edit protocols mpls interface]** hierarchy level.

```
26          <xsl:if test="$mpls and not($mpls/interface[name = $ifname])">
```

Lines 27 through 41 contain another invocation of the **jcs:emit-change** template. In this invocation, the interface is added at the **[edit protocols mpls interface]** hierarchy level.

```
27          <xsl:call-template name="jcs:emit-change">
28              <xsl:with-param name="message">
29                  <xsl:text>Adding ISO-enabled interface </xsl:text>
30                  <xsl:value-of select="$ifname"/>
31                  <xsl:text> to [edit protocols mpls]</xsl:text>
32              </xsl:with-param>
33              <xsl:with-param name="dot" select="$mpls"/>
34              <xsl:with-param name="content">
35                  <interface>
36                      <name>
37                          <xsl:value-of select="$ifname"/>
38                      </name>
39                  </interface>
40              </xsl:with-param>
41          </xsl:call-template>
```

Lines 42 through 45 close all open elements.

```
42          </xsl:if>
43      </xsl:for-each>
44  </xsl:template>
45 </xsl:stylesheet>
```

- Related Documentation**
- [Example: Generating a Transient Change on page 306](#)

## Controlling Execution of Commit Scripts During Commit Operations

Commit scripts are stored on a device's hard disk in the `/var/db/scripts/commit` directory or on the flash drive in the `/config/scripts/commit` directory. Only users in the Junos OS superuser login class can access and edit files in these directories. For information about setting the storage location for scripts, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

A commit script is not actually executed during commit operations unless its filename is included at the `[edit system scripts commit file]` hierarchy level. When you configure the script filename, you must include the appropriate filename extension for SLAX (`.slax`) and Python (`.py`) scripts. XSLT scripts do not require a filename extension, but we strongly recommend that you append the `.xsl` extension for clarity. To prevent execution of a commit script, delete the commit script's filename at that hierarchy level.

By default, the commit operation fails unless all scripts included at the `[edit system scripts commit file]` hierarchy level actually exist in the commit script directory. To enable the commit operation to succeed even if a script is missing, include the **optional** statement at the `[edit system scripts commit file filename]` hierarchy level. For example, you might want to mark a script as optional if you anticipate the need to quickly remove it from operation by deleting it from the commit script directory, but do not want to remove the commit script filename at the `[edit system scripts commit file]` hierarchy level. To enable use of the script again later, you simply replace the file in the commit script directory.



**CAUTION:** When you include the **optional** statement at the `[edit system scripts commit file filename]` hierarchy level, no error message is generated during the commit operation if the file does not exist. As a result, you might not be aware that a script is not executed as you expect.

You can also deactivate and reactivate commit scripts by issuing the **deactivate** and **activate** configuration mode commands. When a commit script is deactivated, the script is marked as inactive in the configuration and does not execute during the commit operation. When a commit script is reactivated, the script is again executed during the commit operation.

To determine which commit scripts are currently enabled on the device, use the **show** command to display the files configured at the `[edit system scripts commit]` hierarchy level. To ensure that the enabled files are on the device, list the contents of the `/var/run/scripts/commit/` directory by using the **file list /var/run/scripts/commit** operational mode command.

See the following sections:

- [Enabling Commit Scripts to Execute During Commit Operations on page 262](#)
- [Preventing Commit Scripts from Executing During Commit Operations on page 263](#)

- [Deactivating Commit Scripts on page 263](#)
- [Activating Commit Scripts on page 263](#)

## Enabling Commit Scripts to Execute During Commit Operations

To configure a commit script to execute during a commit operation:

1. Ensure that the commit script is located in the correct directory:

- `/var/db/scripts/commit` directory on the hard disk
- `/config/scripts/commit` directory on the flash drive

For more information about script storage location, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

2. Enable the commit script by including the `file filename` statement at the `[edit system scripts commit]` hierarchy level. Only users who belong to the Junos OS **super-user** login class can enable commit scripts.

```
[edit system scripts commit]
user@host# set file filename <optional>
```

- **filename**—Name of the commit script.
  - **optional**—Enable the commit operation to succeed when the script file does not exist in the script directory. If this statement is omitted, the commit operation fails if the script does not exist.
3. For unsigned Python scripts, ensure that the following requirements are met:
- File owner is either root or a user in the Junos OS **super-user** login class.
  - Only the file owner has write permission for the file.
  - The `language python` statement is configured at the `[edit system scripts]` hierarchy level.

```
[edit system scripts]
user@host# set language python
```



**NOTE:** In Junos OS Release 16.1R2 and earlier releases, unsigned Python scripts must be owned by the root user. Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS **super-user** login class.

---

4. Commit the configuration.

```
[edit]
user@host# commit
```

## Preventing Commit Scripts from Executing During Commit Operations

To prevent a commit script from executing during a commit operation:

1. Delete the commit script filename at the **[edit system scripts commit]** hierarchy level.

```
[edit system scripts commit]
user@host# delete file filename
```

*filename*—Name of the commit script.

2. Remove the commit script from the commit script directory. Although removing the commit script from the commit script directory is not required, it is always a good policy to delete unused files from the system.
3. Commit your changes.

```
[edit]
user@host# commit
```

## Deactivating Commit Scripts

To deactivate a commit script so that it does not execute during the commit operation:

1. Issue the **deactivate** command.

```
[edit]
user@host# deactivate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@host# commit
```

A deactivated commit script is marked as **inactive:** and ignored during a commit operation.

In this example, the script **mycommit.slax** is deactivated:

```
[edit]
user@host# deactivate system scripts commit file mycommit.slax
[edit]
user@host# show system scripts commit
inactive: file mycommit.slax
```

## Activating Commit Scripts

To activate a commit script that was previously deactivated:

1. Issue the **activate** command.

```
[edit]
user@host# activate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@host# commit
```

## Release History Table

Release	Description
16.1R3	In Junos OS Release 16.1R2 and earlier releases, unsigned Python scripts must be owned by the root user. Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class.

## Configuring Checksum Hashes for a Commit Script

You can configure one or more checksum hashes that can be used to verify the integrity of a commit script before the script runs on the switch, router, or security device.

To configure a checksum hash:

1. Create the script.
2. Place the script in the `/var/db/scripts/commit` directory on the device.
3. Run the script through one or more hash functions to calculate hash values.

Junos OS supports MD5, SHA-1, and SHA-256 hash functions.

```
user@host> file checksum md5 /var/db/scripts/commit/script1.slax
MD5 (/var/db/scripts/commit/script1.slax) = 3af7884eb56e2d4489c2e49b26a39a97
```

```
user@host> file checksum sha1 /var/db/scripts/commit/script1.slax
SHA1 (/var/db/scripts/commit/script1.slax) =
00dc690fb08fb049577d012486c9a6dad34212c0
```

```
user@host> file checksum sha-256 /var/db/scripts/commit/script1.slax
SHA256 (/var/db/scripts/commit/script1.slax) =
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

4. Configure the script.

```
[edit system scripts commit]
user@host# set file script1.slax checksum md5 3af7884eb56e2d4489c2e49b26a39a97
```

```
[edit system scripts commit]
user@host# set file script1.slax checksum
sha-1 00dc690fb08fb049577d012486c9a6dad34212c0
```

```
[edit system scripts commit]
user@host# set file script1.slax checksum
sha-256 150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

During the execution of the script, Junos OS recalculates the checksum value using the configured hash and verifies that the calculated value matches the configured value. If the values differ, the execution of the script fails. When you configure multiple checksum values with different hash algorithms, all the configured values must match the calculated values; otherwise, the script execution fails. The commit operation also fails.

### Related Documentation

- [Configuring Checksum Hashes for an Event Script on page 628](#)
- [Configuring Checksum Hashes for an Op Script on page 460](#)



- [Configuring Checksum Hashes for an SNMP Script on page 647](#)

## Executing Large Commit Scripts

---

When you use large commit scripts, the standard commit model can have trouble reading these scripts. When this occurs, you can include the **direct-access** statement at the **[edit system scripts commit]** hierarchy level. When the **direct-access** statement is included, the script driver retrieves the candidate configuration directly from the configuration database. Once the candidate configuration is retrieved, the script driver processes this configuration file against the commit scripts and returns any generated actions to the management (mgd) process.

Directly accessing the configuration data and processing non-XML converted data are processor-intensive compared to the standard commit model. You should only use this feature to handle large files, because system performance is affected.

To set the script driver to directly access the candidate configuration, include the **direct-access** statement at the **[edit system scripts commit]** hierarchy level.

```
[edit system scripts commit]  
direct-access;
```

## Example: Retrieving the Pre-Inheritance Candidate Configuration in a Commit Script

---

This example shows how to construct a commit script to retrieve the pre-inheritance candidate configuration for either a normal or private configuration session.

- [Requirements on page 265](#)
- [Overview on page 265](#)
- [Configuration on page 266](#)

### Requirements

- Routing, switching, or security device running Junos OS Release 12.2 or later.

### Overview

In Junos OS, when a candidate configuration is committed, it is inspected by each active commit script. The normal input for a commit script is the post-inheritance candidate configuration, in which all configuration group values have been inherited by their targets and the inactive portions of the configuration have been removed.

At times, a commit script requires access to the pre-inheritance candidate configuration rather than the post-inheritance configuration it receives by default. Since normal configuration sessions use the regular candidate database, and private configuration sessions use a dedicated, private candidate database, the candidate database location depends on the session type.

Within a commit script, invoking the **<get-configuration>** remote procedure call (RPC) with the **database="candidate"** attribute retrieves the normal pre-inheritance candidate

configuration. The `<get-configuration>` RPC also has a `database-path` attribute, which is used to specify the location of the pre-inheritance configuration database for either a normal or private configuration session. This attribute is an alternative to the `database` attribute and indicates which database file to load. Commit scripts can invoke the `<get-configuration>` RPC with the `database-path` attribute to retrieve the pre-inheritance candidate configuration specific to that session.

The global variable, `$junos-context` contains the `commit-context/database-path` element, which stores the location of the session's pre-inheritance candidate configuration. In a normal configuration session, the `<database-path>` element contains the location of the normal candidate database:

```
<junos-context>
  <commit-context>
    <database-path>/var/run/db/juniper.db</database-path>
  </commit-context>
</junos-context>
```

In a private configuration session, the `<database-path>` element contains the location of the session-specific, private candidate database. For example:

```
<junos-context>
  <commit-context>
    <commit-private/>
    <database-path>/var/run/db/private/juniper-1396.db</database-path>
  </commit-context>
</junos-context>
```

To construct a commit script that retrieves the pre-inheritance candidate configuration specific to that session, include the `<get-configuration>` RPC in the commit script, and set the `<database-path>` attribute to `$junos-context/commit-context/database-path`. For normal configuration sessions, the commit-script retrieves the normal pre-inheritance candidate configuration, and for private configuration sessions, the commit-script retrieves the private, pre-inheritance candidate configuration.



**NOTE:** If a commit script includes both the `database` and the `database-path` attributes in the `<get-configuration>` tag, the `database` attribute takes precedence.

## Configuration

### Configuring the Commit Script

#### Step-by-Step Procedure

To construct a commit script that retrieves the pre-inheritance candidate configuration specific to that session:

1. In a text editor, add the commit script boilerplate to a file.

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

import "../import/junos.xsl";
```

```

match configuration {
}

```

2. Create a variable that stores the `<get-configuration>` RPC with the `database-path` attribute set to `$junos-context/commit-context/database-path`.

```

var $rpc = <get-configuration
database-path=$junos-context/commit-context/database-path>;

```

3. Add a statement that invokes the `<get-configuration>` RPC and stores the resulting configuration in a variable.

```

var $config = jcs:invoke( $rpc );

```

4. Refer to the desired hierarchy levels and statements in the pre-inheritance candidate configuration using normal XPath constructs, for example:

```

var $hostname = $config/system/host-name;

```

5. Include any statements required to enforce your custom configuration rules during the commit process.
6. Copy the script to the `/var/run/scripts/commit` directory on the device.
7. In configuration mode, configure the `file` statement to enable the commit script.

```

[edit system scripts commit]
user@R1# set file script-name.slax

```

8. Issue the `commit` command to commit the configuration.

```

[edit]
user@R1# commit

```

The commit script is executed during the commit operation.

## Results

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

import "../import/junos.xsl";

match configuration {

    var $rpc =
<get-configuration database-path=$junos-context/commit-context/database-path>;

    var $config = jcs:invoke( $rpc );
    ...
<!-- commit script rules -->
}

```

**Related Documentation** • [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)



## CHAPTER 17

# Generating a Custom Warning, Error, or System Log Message Using Commit Scripts

- [Overview of Generating Custom Warning, Error, and System Log Messages on page 269](#)
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)
- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages on page 274](#)
- [Example: Generating a Custom Warning Message on page 276](#)
- [Example: Generating a Custom Error Message on page 279](#)
- [Example: Generating a Custom System Log Message on page 283](#)

### Overview of Generating Custom Warning, Error, and System Log Messages

You can use a commit script to specify configuration rules that you always want to enforce. If a rule is broken, the commit script can emit a warning, error, or system log message.

In the Junos OS command-line interface (CLI), warning messages are emitted during commit operations to alert you that the configuration is not complete or contains a syntax error. If a custom configuration rule is broken, a custom warning message notifies you about the problem. The commit script causes the warning message to be passed back to the Junos OS CLI or to a Junos XML protocol client application. Unlike error messages, warning messages do not cause the commit operation to fail, so they are used for configuration problems that do not affect network traffic. A warning is best used as a response to configuration settings that do not adhere to recommended practices. An example of this type of configuration setting might be assignment of the same user ID to different users.

Alternatively, you can generate a custom warning message for a serious configuration problem, and specify an automatic configuration change that rectifies the problem. For more information about the use of warning messages in conjunction with automatic configuration changes, see [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts” on page 287](#).

Unlike warning messages, a custom error message causes the commit operation to fail and notifies the user about the configuration problem. The commit script causes the error message to be passed back to the Junos OS CLI or to a Junos XML protocol client application. Because error messages cause the commit operation to fail, they are used for problems that affect network traffic. An error message is best used as a response to configuration settings that you want to disallow—for example, when required statements are omitted from the configuration.

Junos OS generates system log messages (also called syslog messages) to record events that occur on the device, including the following:

- Routine operations, such as creation of an OSPF protocol adjacency or a user login into the configuration database
- Failure and error conditions, such as failure to access a configuration file or unexpected closure of a connection to a child or peer process
- Emergency or critical conditions, such as device power-down due to excessive temperature

Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred. The [System Log Explorer](#) provides more detailed information about system log messages.

With commit scripts, you can cause custom system log messages to be generated in response to particular events that you define. For example, if a configuration rule is broken, a custom message can be generated to record this occurrence. If the commit script corrects the configuration, a custom message can indicate that corrective action was taken.

**Related  
Documentation**

- [Example: Generating a Custom Error Message on page 279](#)
- [Example: Generating a Custom System Log Message on page 283](#)
- [Example: Generating a Custom Warning Message on page 276](#)
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)
- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages on page 274](#)

---

## Generating a Custom Warning, Error, or System Log Message in Commit Scripts

---

Junos OS commit scripts can generate custom warning, error, or system log messages during a commit operation to alert you when the configuration does not comply with custom configuration rules. Generating an error also causes the commit operation to fail. To generate a custom warning, error, or system log message in a commit script:

1. Include the appropriate commit script boilerplate from “[Required Boilerplate for Commit Scripts](#)” on page 251. It is reproduced here for convenience:

### XSLT Boilerplate

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <!-- ... insert your code here ... -->
  </xsl:template>
</xsl:stylesheet>
```

### SLAX Boilerplate

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  /*
   * insert your code here
  */
}
```

### Python Boilerplate

```
from junos import Junos_Configuration
import jcs

if __name__ == '__main__':
    # insert your code here
```

- At the position indicated by the comment “*insert your code here*,” include one or more XSLT programming instructions or their SLAX or Python equivalents. Commonly used XSLT constructs include the following:

- **<xsl:choose>** **<xsl:when>** **<xsl:otherwise>**—Conditional construct that causes different instructions to be processed in different circumstances. The **<xsl:choose>** instruction contains one or more **<xsl:when>** elements, each of which tests an XPath expression. If the test evaluates as true, the XSLT processor executes the instructions in the **<xsl:when>** element. The XSLT processor processes only the instructions contained in the first **<xsl:when>** element whose **test** attribute evaluates as true. If none of the **<xsl:when>** elements’ **test** attributes evaluate as true, the content of the **<xsl:otherwise>** element, if there is one, is processed.
- **<xsl:for-each select="xpath-expression">**—Programming instruction that tells the XSLT processor to gather together a set of nodes and process them one by one. The nodes are selected by the Extensible Markup Language (XML) Path Language (XPath) expression in the **select** attribute. Each of the nodes is then processed according to the instructions contained in the **<xsl:for-each>** instruction. Code inside an **<xsl:for-each>** instruction is evaluated recursively for each node that matches the XPath expression. The context is moved to the node during each pass.

- **<xsl:if test="xpath-expression">**—Conditional construct that causes instructions to be processed if the XPath expression in the **test** attribute evaluates to **true**.

For example, the following programming instruction evaluates as true when the **host-name** statement is not included at the **[edit system]** hierarchy level:

```
<xsl:if test="not(system/host-name)">
```

In SLAX, the **if** construct looks like this:

```
if (not(system/host-name))
```

Similarly in Python:

```
if not(Junos_Configuration.xpath("./system/host-name")):
```

3. Include the appropriate constructs to generate a warning, error, or system log message.

In SLAX and XSLT scripts, include **<xnm:warning>**, **<xnm:error>**, or **<syslog>** elements with a **<message>** child element that specifies the content of the message. For warning and error messages, you can include several other child elements, such as the **jcs:edit-path** and **jcs:statement** templates, which cause the warning or error message to include the relevant configuration hierarchy and statement information, as shown in the following examples.

In Python scripts, include **jcs.emit\_warning()**, **jcs.emit\_error()**, or **jcs.syslog()** functions, and include the message string in the argument list.

This **<xnm:warning>** element:

```
<xnm:warning>
  <xsl:call-template name="jcs:edit-path">
    <xsl:with-param name="dot" select="chassis"/>
  </xsl:call-template>
  <message>IP source-route processing is not enabled.</message>
</xnm:warning>
```

emits this output when you issue the **commit** command:

```
[edit]
user@host# commit

[edit chassis]
warning: IP source-route processing is not enabled.
commit complete
```

This **<xnm:error>** element:

```
<xnm:error>
  <xsl:call-template name="jcs:edit-path"/>
  <xsl:call-template name="jcs:statement"/>
  <message>Missing a description for this T1 interface.</message>
</xnm:error>
```

emits this output when you issue the **commit** command:

```
[edit]
user@host# commit

[edit interfaces interface t1-0/0/0]
'interface t1-0/0/0;'
Missing a description for this T1 interface.
error: 1 error reported by commit scripts
```



error: commit script failure



**NOTE:** In SLAX and XSLT scripts, if you are including a warning message in conjunction with a script-generated configuration change, you can generate the warning by including the `message` parameter with the `jcs:emit-change` template. The `message` parameter causes the `jcs:emit-change` template to call the `<xnm:warning>` template, which sends a warning notification to the CLI. (For more information, see [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts”](#) on page 287.)

For system log messages, the only supported child element is `<message>`:

```
<syslog>
  <message>syslog-string</message>
</syslog>
```

For a description of all the XSLT tags and attributes you can include, see [“SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages”](#) on page 274.

For SLAX versions of these constructs, see [“Example: Generating a Custom Warning Message”](#) on page 276, [“Example: Generating a Custom Error Message”](#) on page 279, and [“Example: Generating a Custom System Log Message”](#) on page 283.

4. Save the script with a meaningful name.
5. Copy the script to either the `/var/db/scripts/commit` directory on the hard disk or the `/config/scripts/commit` directory on the flash drive. For information about setting the storage location for commit scripts, see [“Storing and Enabling Scripts”](#) on page 661 and [“Storing Scripts in Flash Memory”](#) on page 664.

If the device has dual Routing Engines and you want the script to take effect on both of them, you must copy the script to the `/var/db/scripts/commit` or the `/config/scripts/commit` directory on both Routing Engines. The `commit synchronize` command does not copy scripts between Routing Engines.

6. Enable the script by including the `file filename` statement at the `[edit system scripts commit]` hierarchy level.

```
[edit system scripts commit]
user@host# set file filename
```

7. If the script is written in Python, configure the `language python` statement.

```
[edit]
user@host# set system script language python
```

8. Issue the `commit` command to commit the configuration.

```
[edit]
user@host# commit and-quit
```

#### Related Documentation

- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages](#) on page 274

- [Example: Generating a Custom Error Message on page 279](#)
- [Example: Generating a Custom System Log Message on page 283](#)
- [Example: Generating a Custom Warning Message on page 276](#)

## SLAX and XSLT Commit Script Tag Elements to Use When Generating Messages

Junos OS commit scripts can generate custom warning, error, or system log messages during a commit operation to alert you when the configuration does not comply with custom configuration rules. [Table 33 on page 274](#) summarizes the tag elements that you can include in a custom warning, error, or system log message in SLAX and XSLT commit scripts. For examples, see [“Example: Generating a Custom Warning Message” on page 276](#), [“Example: Generating a Custom Error Message” on page 279](#), and [“Example: Generating a Custom System Log Message” on page 283](#).

**Table 33: Tags and Attributes for Creating Custom Warning, Error, and System Log Messages**

Data Item, XML Element, or Attribute	Required or Supported	Description
<b>Container Tags and Attributes</b>		
<code>&lt;syslog&gt;</code>	Required for system log messages	Indicates that a system log message is going to be recorded.
<code>&lt;xnm:error&gt;</code>	Required for error messages	Indicates that the server has encountered a problem while processing the client application's request.
<code>&lt;xnm:warning&gt;</code>	Required for warning messages	Indicates that the server has encountered a problem while processing the client application's request.
<code>xmlns url</code>	Supported in warning and error messages	Names the XML namespace for the contents of the tag element. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code> , where <b>version</b> is a string such as 1.1.
<code>xmlns:xnm url</code>	Required for warning and error messages. The <code>xmlns:xnm</code> element is included in the script boilerplate, which sets the namespace globally.	Names the XML namespace for child tag elements that have the <b>xnm:</b> prefix on their names. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code> , where <b>version</b> is a string such as 1.1.
<b>Content Tags</b>		
<code>&lt;column&gt;</code>	Supported in warning and error messages only	Identifies the element that caused the error by specifying its position as the number of characters after the first character in the line specified by the <code>&lt;line-number&gt;</code> tag element in the configuration file that was being loaded (which is named in the <code>&lt;filename&gt;</code> tag element). We recommend combining the <code>&lt;column&gt;</code> tag with the <code>&lt;line-number&gt;</code> and <code>&lt;filename&gt;</code> tags.

Table 33: Tags and Attributes for Creating Custom Warning, Error, and System Log Messages (*continued*)

Data Item, XML Element, or Attribute	Required or Supported	Description
<code>&lt;database-status-information&gt;</code>	Supported in error messages only	Provides information about the users currently editing the configuration.
<code>&lt;edit-path&gt;</code>	Supported in warning and error messages only	Specifies the level in the configuration hierarchy where the problem occurred, using the CLI configuration mode banner. We recommend combining the <code>&lt;edit-path&gt;</code> tag with the <code>&lt;statement&gt;</code> tag.
<code>&lt;filename&gt;</code>	Supported in warning and error messages only	Names the configuration file that was being loaded.
<code>&lt;line-number&gt;</code>	Supported in warning and error messages only	Specifies the line number where the error occurred in the configuration file that was being loaded, which is named by the <code>&lt;filename&gt;</code> tag element. We recommend combining the <code>&lt;line-number&gt;</code> tag with the <code>&lt;column&gt;</code> and <code>&lt;filename&gt;</code> tags.
<code>&lt;message&gt;</code>	Required in warning, error, and system log messages	Describes the warning, error, or system log message in a natural-language text string.
<code>&lt;parse/&gt;</code>	Supported in error messages only	Indicates that there was a syntactic error in the request submitted by the client application.
<code>&lt;reason&gt;</code>	Supported in warning and error messages only	Describes the reason for the warning or error message.
<code>&lt;re-name&gt;</code>	Supported in warning and error messages only	Names the Routing Engine on which the process named by the <code>&lt;source-daemon&gt;</code> tag element is running.
<code>&lt;source-daemon&gt;</code>	Supported in warning and error messages only	Names the Junos OS module that was processing the request in which the warning or error message occurred.
<code>&lt;statement&gt;</code>	Supported in warning and error messages only	Specifies the configuration statement in effect when the problem occurred. We recommend combining the <code>&lt;statement&gt;</code> tag with the <code>&lt;edit-path&gt;</code> tag.
<code>&lt;token&gt;</code>	Supported in warning and error messages only	Names the element in the request that caused the warning or error message.
<code>&lt;xsl:call-template name="jcs:edit-path"&gt;</code>	Supported in warning and error messages only	<p>Emits an <code>&lt;edit-path&gt;</code> element, which specifies the CLI configuration mode edit path in effect when the warning or error was generated.</p> <p>If the problem is not at the current position in the XML hierarchy, you can alter the edit path by passing the <code>dot</code> parameter. For example, <code>&lt;xsl:param name="dot" select="system/ports/console"/&gt;</code> changes the edit path to <code>[edit system ports console]</code>.</p>

**Table 33: Tags and Attributes for Creating Custom Warning, Error, and System Log Messages** (*continued*)

Data Item, XML Element, or Attribute	Required or Supported	Description
<code>&lt;xsl:call-template name="jcs:statement"&gt;</code>	Supported in warning and error messages only	<p>Emits a <code>&lt;statement&gt;</code> element, which describes the configuration statement in effect when the warning or error was generated.</p> <p>If the problem is not at the current position in the XML hierarchy, you can alter the statement by passing the <b>dot</b> parameter. For example, <code>&lt;xsl:with-param name="dot" select="system/ports/console/type"/&gt;</code> changes the statement to <b>type</b>.</p>

### Example: Generating a Custom Warning Message

Junos OS commit scripts can generate custom warning messages during a commit operation to alert you when the configuration does not comply with custom configuration rules. The commit process is not affected by warnings. This example commit script generates a custom warning message when a specific statement is not included in the device configuration.

- [Requirements on page 276](#)
- [Overview and Commit Script on page 276](#)
- [Configuration on page 277](#)
- [Verification on page 278](#)

### Requirements

Junos OS Release 16.1R3 or later release when using a Python script.

### Overview and Commit Script

Using a commit script, write a custom warning message that appears when the **source-route** statement is not included at the **[edit chassis]** hierarchy level.

The script is shown in XSLT, SLAX, and Python.

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:if test="not(chassis/source-route)">
      <xnm:warning>
        <xsl:call-template name="jcs:edit-path">
          <xsl:with-param name="dot" select="chassis"/>
        </xsl:call-template>
        <message>IP source-route processing is not enabled.</message>
      </xnm:warning>
    </xsl:if>
  </xsl:template>
</xsl:stylesheet>
```

```

        </xsl:if>
    </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    if (not(chassis/source-route)) {
        <xnm:warning> {
            call jcs:edit-path($dot = chassis);
            <message> "IP source-route processing is not enabled.";
        }
    }
}

```

**Python Syntax**

```

from junos import Junos_Configuration
import jcs

def main():
    root = Junos_Configuration
    if not(root.xpath("./chassis/source-route")):
        jcs.emit_warning("IP source-route processing is not enabled.")

if __name__ == '__main__':
    main()

```

## Configuration

**Step-by-Step Procedure**

Download, enable, and test the script. To test that a commit script generates a warning message correctly, make sure that the candidate configuration contains the condition that elicits the warning. For this example, ensure that the **source-route** statement is not included at the **[edit chassis]** hierarchy level.

To test the example in this topic:

1. Copy the script into a text file, name the file **source-route.xsl**, **source-route.slax**, or **source-route.py** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file** statement at the **[edit system scripts commit]** hierarchy level and the appropriate filename, for example:

```

[edit]
user@host# set system scripts commit file source-route.xsl

```

3. If the script is written in Python, configure the **language python** statement.

```
[edit]
user@host# set system scripts language python
```

4. If the **source-route** statement is included at the **[edit chassis]** hierarchy level, issue the **delete chassis source-route** configuration mode command.

```
[edit]
user@host# delete chassis source-route
```

5. Issue the **commit and-quit** command.

```
[edit]
user@host# commit and-quit
```

## Verification

### Verifying Script Execution

---

**Purpose** Verify the warning message generated by the commit script.

**Action** Review the output of the **commit** command. The commit script generates a warning message when the **source-route** statement is not included at the **[edit chassis]** hierarchy level of the configuration. The warning does not affect the commit process.

```
[edit]
user@host# commit
[edit chassis]
warning: IP source-route processing is not enabled.
commit complete
```

To display the XML-formatted version of the warning message, issue the **commit check | display xml** command.

```
[edit]
user@host# commit check | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <commit-results>
    <routing-engine junos:style="normal">
      <name>re0</name>
      <xnm:warning>
        <edit-path>
          [edit chassis]
        </edit-path>
        <message>
          IP source-route processing is not enabled.
        </message>
      </xnm:warning>
      <commit-check-success/>
    </routing-engine>
  </commit-results>
</rpc-reply>
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command.

```
[edit]
user@host# commit check | display detail
2009-06-15 14:40:29 PDT: reading commit script configuration
2009-06-15 14:40:29 PDT: testing commit script configuration
2009-06-15 14:40:29 PDT: opening commit script
'/var/db/scripts/commit/source-route-warning.xml'
2009-06-15 14:40:29 PDT: reading commit script 'source-route-warning.xml'
2009-06-15 14:40:29 PDT: running commit script 'source-route-warning.xml'
2009-06-15 14:40:29 PDT: processing commit script 'source-route-warning.xml'
[edit chassis]
warning: IP source-route processing is not enabled.
2009-06-15 14:40:29 PDT: no errors from source-route-warning.xml
2009-06-15 14:40:29 PDT: saving commit script changes
2009-06-15 14:40:29 PDT: summary: changes 0, transients 0 (allowed), syslog 0
2009-06-15 14:40:29 PDT: no commit script changes
2009-06-15 14:40:29 PDT: exporting juniper.conf
2009-06-15 14:40:29 PDT: expanding groups
2009-06-15 14:40:29 PDT: finished expanding groups
2009-06-15 14:40:29 PDT: setup foreign files
2009-06-15 14:40:29 PDT: propagating foreign files
2009-06-15 14:40:30 PDT: complete foreign files
2009-06-15 14:40:30 PDT: daemons checking new configuration
configuration check succeeds
```

#### Related Documentation

- [Example: Generating a Custom Error Message on page 279](#)
- [Example: Generating a Custom System Log Message on page 283](#)
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)

---

## Example: Generating a Custom Error Message

Junos OS commit scripts can generate custom error messages during a commit operation to alert you when the configuration violates custom configuration rules. Emitting an error message causes the commit to fail. This example commit script generates a custom error message when a specific statement is not included in the device configuration, thereby halting the commit operation.

- [Requirements on page 279](#)
- [Overview and Commit Script on page 279](#)
- [Configuration on page 281](#)
- [Verification on page 281](#)

### Requirements

Junos OS Release 16.1R3 or later release when using a Python script.

### Overview and Commit Script

Using a commit script, write a custom error message that appears when the **description** statement is not included at the **[edit interfaces t1-fpc/pic/port]** hierarchy level:

The script is shown in XSLT, SLAX, and Python.

**XSLT Syntax**

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:variable name="interface" select="interfaces/interface"/>
    <xsl:for-each select="$interface[starts-with(name, 't1-')]">
      <xsl:variable name="ifname" select="."/>
      <xsl:if test="not(description)">
        <xnm:error>
          <xsl:call-template name="jcs:edit-path"/>
          <xsl:call-template name="jcs:statement"/>
          <message>Missing a description for this T1
interface.</message>
        </xnm:error>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $interface = interfaces/interface;
  for-each ($interface[starts-with(name, 't1-')]) {
    var $ifname = .;
    if (not(description)) {
      <xnm:error> {
        call jcs:edit-path();
        call jcs:statement();
        <message> "Missing a description for this T1 interface.";
      }
    }
  }
}
```

**Python Syntax**

```
from junos import Junos_Configuration
import jcs

def main():
    root = Junos_Configuration
    for element in root.xpath("../interfaces/ \
interface[starts-with(name,'t1-')]" ):
        # Missing description
        if element.find('description') is None:
            # Emit error message to console
            jcs.emit_error("Missing a description for this T1 interface: " \
+ element.find('name').text)

if __name__ == '__main__':
    main()
```



## Configuration

**Step-by-Step Procedure** Download, enable, and test the script: To test that a commit script generates an error message correctly, make sure that the candidate configuration contains the condition that elicits the error. For this example, ensure that the configuration for a T1 interface does not include the **description** statement.

To test the example in this topic:

1. Copy the script into a text file, name the file **description.xml**, **description.slax**, or **description.py** as appropriate, and copy it to the `/var/db/scripts/commit/` directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file** statement at the `[edit system scripts commit]` hierarchy level and the appropriate filename, for example:

```
[edit]
user@host# set system scripts commit file description.xml
```

3. If the script is written in Python, configure the **language python** statement.

```
[edit]
user@host# set system scripts language python
```

4. If the configuration for every T1 interface includes the **description** statement, delete the description for an existing T1 interface for testing purposes.

```
[edit]
user@host# delete interfaces t1-0/0/1 description
```

5. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying Script Execution

**Purpose** Verify the error message generated by the commit script.

**Action** Review the output of the **commit** command. The commit script generates an error message for each T1 interface that does not include a **description** statement. Any error causes the commit process to fail.

```
[edit]
user@host# commit
[edit interfaces interface t1-0/0/1]
'description'
Missing a description for this T1 interface.
```

```
[edit interfaces interface t1-0/0/2]
'description'
Missing a description for this T1 interface.
error: 2 errors reported by commit scripts
error: commit script failure
```

To display the XML-formatted version of the error message, issue the **commit check | display xml** command.

```
[edit interfaces t1-0/0/1]
user@host# commit check | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <commit-results>
    <routing-engine junos:style="normal">
      <name>re0</name>
      <xnm:error>
        <edit-path>
          [edit interfaces interface t1-0/0/1]
        </edit-path>
        <statement>
          description
        </statement>
        <message>
          Missing a description for this T1 interface.
        </message>
      </xnm:error>
      <xnm:error>
        <edit-path>
          [edit interfaces interface t1-0/0/2]
        </edit-path>
        <statement>
          description
        </statement>
        <message>
          Missing a description for this T1 interface.
        </message>
      </xnm:error>
      <xnm:error xmlns="http://xml.juniper.net/xnm/1.1/xnm"
        xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm">
        <message>
          2 errors reported by commit scripts
        </message>
      </xnm:error>
      <xnm:error xmlns="http://xml.juniper.net/xnm/1.1/xnm"
        xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm">
        <message>
          commit script failure
        </message>
      </xnm:error>
    </routing-engine>
  </commit-results>
</cli>
  <banner>[edit interfaces]</banner>
</cli>
</rpc-reply>
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command.

```
[edit interfaces t1-0/0/1]
```

```

user@host# commit check | display detail
2009-06-15 15:56:09 PDT: reading commit script configuration
2009-06-15 15:56:09 PDT: testing commit script configuration
2009-06-15 15:56:09 PDT: opening commit script '/var/db/scripts/commit/error.xml'
2009-06-15 15:56:09 PDT: reading commit script 'error.xml'
2009-06-15 15:56:09 PDT: running commit script 'error.xml'
2009-06-15 15:56:09 PDT: processing commit script 'error.xml'
[edit interfaces interface t1-0/0/1]
  'description'
    Missing a description for this T1 interface.
[edit interfaces interface t1-0/0/2]
  'description'
    Missing a description for this T1 interface.
2009-06-15 15:56:09 PDT: 2 errors from script 'error.xml'
error: 2 errors reported by commit scripts
error: commit script failure

```

#### Related Documentation

- [Example: Generating a Custom System Log Message on page 283](#)
- [Example: Generating a Custom Warning Message on page 276](#)
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)

## Example: Generating a Custom System Log Message

Junos OS commit scripts can generate custom system log messages during a commit operation to alert you when the configuration does not comply with custom configuration rules. The commit process is not affected by generating system log messages. This example commit script generates a custom system log message when a specific statement is not included in the device configuration.

- [Requirements on page 283](#)
- [Overview and Commit Script on page 283](#)
- [Configuration on page 285](#)
- [Verification on page 285](#)

### Requirements

Junos OS Release 16.1R3 or later when using a Python script.

### Overview and Commit Script

Using a commit script, write a custom system log message that appears when the **read-write** statement is not included at the **[edit snmp community *community-name* authorization]** hierarchy level.

The script is shown in XSLT, SLAX, and Python.

#### XSLT Syntax

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">

```

```
<xsl:import href="../import/junos.xsl"/>

<xsl:template match="configuration">
  <xsl:for-each select="snmp/community">
    <xsl:if test="not(authorization) or (authorization != 'read-write')">
      <xsl:variable name="community">
        <xsl:call-template name="jcs:edit-path"/>
      </xsl:variable>
      <xsl:variable name="message" select="concat('SNMP community does not
have read-write access: ', $community)"/>
      <syslog>
        <message>
          <xsl:value-of select="$message"/>
        </message>
      </syslog>
    </xsl:if>
  </xsl:for-each>
</xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (snmp/community) {
    if ( not(authorization) or (authorization != "read-write")) {
      var $community = call jcs:edit-path();
      var $message = "SNMP community does not have read-write access:
" _ $community;
      <syslog> {
        <message> $message;
      }
    }
  }
}
```

**Python Syntax**

```
from junos import Junos_Configuration
import jcs

def main():
    root = Junos_Configuration
    for element in root.xpath("./snmp/community"):
        if element.find("authorization") is None or
element.find("authorization").text != 'read-write':
            jcs.syslog("172", "SNMP community does not have read-write access:
" +
                        element.find('name').text)

if __name__ == '__main__':
    main()
```

## Configuration

**Step-by-Step Procedure** Download, enable, and test the script. To test that a commit script generates a system log message correctly, make sure that the candidate configuration contains the condition that elicits the system log message. For this example, ensure that the **read-write** statement is not included at the **[edit snmp community *community-name* authorization]** hierarchy level.

To test the example in this topic:

1. Copy the script into a text file, name the file **read-write.xml**, **read-write.slax**, or **read-write.py** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file** statement at the **[edit system scripts commit]** hierarchy level and the appropriate filename, for example:

```
[edit]
user@host# set system scripts commit file read-write.xml
```

3. If the script is written in Python, configure the **language python** statement.

```
[edit]
user@host# set system scripts language python
```

4. (Optional) To test the condition, if the **read-write** statement is included at the **[edit snmp community *community-name* authorization]** hierarchy level for every community, temporarily delete the authorization for an existing SNMP community.

```
[edit]
user@host# delete snmp community community-name authorization read-write
```

5. Issue the following command to verify that system logging is configured to write to a file (a commonly used file name is **messages**):

```
[edit]
user@host# show system syslog
```

For information about system log configuration, see the [System Log Explorer](#).

6. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying Script Execution

**Purpose** Verify the system log message generated by the commit script.

**Action** System log messages are generated during a commit operation but not during a commit check operation. This means you cannot use the **commit check | display xml** or **commit check | display detail** configuration mode commands to verify the output of system log messages. When the commit operation completes, inspect the system log file. The default directory for log files is **/var/log/**. View the log file by issuing the **show log filename** operational mode command. For example, if messages are logged to the **messages** file, issue the following command:

```
user@host> show log messages
```

System log entries generated by commit scripts have the following format:

```
timestamp host-name cscript: message
```

Since the **read-write** statement was not included at the **[edit snmp community community-name authorization]** hierarchy level, the commit script should generate the "SNMP community does not have read-write access" message in the system log file.

```
Jun 3 14:34:37 host-name cscript: SNMP community does not have read-write access:
[edit snmp community community-name]
```

**Related  
Documentation**

- [Example: Generating a Custom Error Message on page 279](#)
- [Example: Generating a Custom Warning Message on page 276](#)
- [Generating a Custom Warning, Error, or System Log Message in Commit Scripts on page 270](#)

## CHAPTER 18

# Generating Persistent or Transient Configuration Changes Using Commit Scripts

- [Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts on page 287](#)
- [Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292](#)
- [Generating a Persistent or Transient Configuration Change in Python Commit Scripts on page 297](#)
- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Persistent and Transient Configuration Changes on page 299](#)
- [Removing a Persistent or Transient Configuration Change Using Commit Scripts on page 300](#)
- [Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts on page 301](#)

## Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts

---

Junos OS commit scripts enforce custom configuration rules. When a candidate configuration includes statements that you have decided must not be included in your configuration, or when the candidate configuration omits statements that you have decided are required, commit scripts can automatically change the configuration and thereby correct the problem.

- [Differences Between Persistent and Transient Changes on page 287](#)
- [Interaction of Configuration Changes and Configuration Groups on page 290](#)
- [Tag Elements and Templates for Generating Changes on page 291](#)

## Differences Between Persistent and Transient Changes

Configuration changes made by commit scripts can be *persistent* or *transient*.

A persistent change remains in the candidate configuration and affects routing operations until you explicitly delete it, even if you subsequently remove or disable the commit script that generated the change and reissue the **commit** command. In other words, removing the commit script does not cause a persistent change to be removed from the configuration.

A transient change, in contrast, is made in the *checkout configuration* but not in the candidate configuration. The checkout configuration is the configuration database that is inspected for standard Junos OS syntax just before it is copied to become the active configuration on the device. If you subsequently remove or disable the commit script that made the change and reissue the **commit** command, the change is no longer made to the checkout configuration and so does not affect the active configuration. In other words, removing the commit script effectively removes a transient change from the configuration.

A common use for transient changes is to eliminate the need to repeatedly configure and display well-known policies, thus allowing these policies to be enforced implicitly. For example, if MPLS must be enabled on every interface with an International Organization for Standardization (ISO) protocol enabled, the change can be transient, so that the repetitive or redundant configuration data need not be carried or displayed in the candidate configuration. Furthermore, transient changes allow you to write script instructions that apply the change only if a set of conditions is met.

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute to a configuration element produces the same behavior as a **replace:** tag in a **load replace** operation.

By default, Junos OS merges the incoming configuration and the candidate configuration. New statements and hierarchies are added, and conflicting statements are overridden. When generating a persistent or transient change, if you add the **replace="replace"** attribute to a configuration element, Junos OS replaces the existing configuration element with the incoming configuration element. If the **replace="replace"** attribute is added to a configuration element, but there is no existing element of the same name in the current configuration, the incoming configuration element is added into the configuration. Elements that do not have the **replace** attribute are merged into the configuration.

Persistent and transient changes are loaded before the standard Junos OS validation checks are performed. This means any configuration changes introduced by a commit script are validated for correct syntax. If the syntax is correct, the new configuration becomes the active, operational device configuration.

Protected elements in the configuration hierarchy cannot be modified or deleted by either a persistent or a transient change. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made, and proceeds with the commit.

Persistent and transient changes have several important differences, as described in [Table 34 on page 289](#).



Table 34: Differences Between Persistent and Transient Changes

Persistent Changes	Transient Changes
<p>You can represent a persistent change in commit scripts by using the <b>content</b> parameter in conjunction with a <b>tag</b> parameter that is set to 'change' inside a call to the <b>jcs:emit-change</b> template in SLAX and XSLT scripts or a call to the <b>jcs.emit_change</b> method in Python scripts.</p> <p>SLAX and XSLT commit scripts can also represent a persistent change by using the <b>&lt;change&gt;</b> tag.</p>	<p>You can represent a transient change in commit scripts with the <b>content</b> parameter in conjunction with the a <b>tag</b> parameter that is set to 'transient-change' inside a call to the <b>jcs:emit-change</b> template in SLAX and XSLT scripts or a call to the <b>jcs.emit_change</b> method in Python scripts.</p> <p>SLAX and XSLT commit scripts can also represent a transient change by using the <b>&lt;transient-change&gt;</b> tag.</p>
<p>You can use persistent changes to perform any Junos XML protocol operation, such as activate, deactivate, delete, insert (reorder), comment (annotate), and replace sections of the configuration.</p>	<p>Like persistent changes, you can use transient changes to perform any Junos XML protocol operation. However, some Junos XML protocol operations do not make sense to use with transient changes, such as generating comments and inactive settings.</p>
<p>Persistent changes are always loaded during the commit process if no errors are generated by any commit scripts or by the standard Junos OS validity check.</p>	<p>For transient changes to be loaded, you must include the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level. If you enable a commit script that generates transient changes and you do not include the <b>allow-transients</b> statement in the configuration, the CLI generates an error message and the commit operation fails. You cannot use a commit script to generate the <b>allow-transients</b> statement.</p> <p>Like persistent changes, transient changes must pass the standard Junos OS validity check.</p>
<p>Persistent changes work like the <b>load replace</b> configuration mode command, and the change is added to the candidate configuration.</p> <p>When generating a persistent change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the candidate configuration with the incoming configuration element. If there is no existing element of the same name in the candidate configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p>	<p>Transient changes work like the <b>load replace</b> configuration mode command, and the change is added to the checkout configuration.</p> <p>When generating a transient change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the checkout configuration with the incoming configuration element. If there is no existing element of the same name in the checkout configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p> <p>Transient changes are not copied to the candidate configuration. For this reason, transient changes are not saved in the configuration if the associated commit script is deleted or deactivated.</p>
<p>After a persistent change is committed, the software treats it like a change you make by directly editing and committing the candidate configuration.</p> <p>After the persistent changes are copied to the candidate configuration, they are copied to the checkout configuration. If the changes pass the standard Junos OS validity checks, the changes are propagated to the switch, router, or security device components.</p>	<p>Each time a transient change is committed, the software updates the checkout configuration database. After the transient changes pass the standard Junos OS validity checks, the changes are propagated to the device components.</p>

Table 34: Differences Between Persistent and Transient Changes (*continued*)

Persistent Changes	Transient Changes
<p>After committing a script that causes a persistent change to be generated, you can view the persistent change by issuing the <b>show</b> configuration mode command:</p> <pre>user@host# show</pre> <p>This command displays persistent changes only, not transient changes.</p>	<p>After committing a script that causes a transient change to be generated, you can view the transient change by issuing the <b>show   display commit-scripts</b> configuration mode command:</p> <pre>user@host# show   display commit-scripts</pre> <p>This command displays both persistent and transient changes.</p>
<p>Persistent changes must conform to your custom configuration design rules as dictated by commit scripts.</p> <p>This does not become apparent until after a second commit operation because persistent changes are not evaluated by commit script rules on the current commit operation. The subsequent commit operation fails if the persistent changes do not conform to the rules imposed by the commit scripts configured during the first commit operation.</p>	<p>Transient changes are never tested by and do not need to conform to your custom rules. This is caused by the order of operations in the Junos OS commit model, which is explained in detail in <a href="#">“Commit Scripts and the Junos OS Commit Model” on page 247</a>.</p>
<p>A persistent change remains in the configuration even if you delete, disable, or deactivate the commit script instructions that generated the change.</p>	<p>If you delete, disable, or deactivate the commit script instructions that generate a transient change, the change is removed from the configuration after the next commit operation. In short, if the associated instructions or the entire commit script is removed, the transient change is also removed.</p>
<p>As with direct CLI configuration, you can remove a persistent change by rolling back to a previous configuration that did not include the change and issuing the <b>commit</b> command. However, if you do not disable or deactivate the associated commit script, and the problem that originally caused the change to be generated still exists, the change is automatically regenerated when you issue another <b>commit</b> command.</p>	<p>You cannot remove a transient change by rolling back to a previous configuration.</p>
<p>You can alter persistent changes directly by editing the configuration using the CLI.</p>	<p>You cannot directly alter or delete a transient change by using the Junos OS CLI, because the change is not in the candidate configuration.</p> <p>To alter the contents of a transient change, you must alter the statements in the commit script that generates the transient change.</p>

## Interaction of Configuration Changes and Configuration Groups

Any configuration change you can make by directly editing the configuration using the Junos OS command-line interface (CLI) can also be generated by a commit script as a persistent or transient change. This includes values specified at a specific hierarchy level or in configuration groups. As with direct CLI configuration, values specified in the *target* override values inherited from a configuration group. The target is the statement to which you apply a configuration group by including the **apply-groups** statement.

If you define persistent or transient changes as belonging to a configuration group, the configuration groups are applied in the order you specify in the **apply-groups** statements, which you can include at any hierarchy level except the top level. You can also disable inheritance of a configuration group by including the **apply-groups-except** statement at any hierarchy level except the top level.



**CAUTION:** Each commit script inspects the postinheritance view of the configuration. If a candidate configuration contains a configuration group, be careful when using a commit script to change the related target configuration, because doing so might alter the intended inheritance from the configuration group.

Also be careful when using a commit script to change a configuration group, because the configuration group might be generated by an application that performs a load replace operation on the group during each commit operation.

For more information about configuration groups, see the *CLI User Guide*.

## Tag Elements and Templates for Generating Changes

To generate persistent or transient changes in commit scripts, SLAX and XSLT scripts can use the **jcs:emit-change** template, and Python scripts can use the **jcs.emit\_change** method. The **jcs:emit-change** template and **jcs.emit\_change** method implicitly include **<change>** and **<transient-change>** XML elements. SLAX and XSLT scripts can also generate changes by including the **<change>** and **<transient-change>** elements directly in the commit script. Using the **jcs:emit-change** template in SLAX and XSLT scripts allows you to set the hierarchical context of the change once rather than multiple times. In Python scripts, the **jcs.emit\_change** method requires that the configuration data for the requested change include the full configuration path representing all levels of the configuration hierarchy formatted as an XML string.

The **<change>** and **<transient-change>** elements are similar to the **<load-configuration>** operation defined by the Junos XML management protocol. The possible contents of the **<change>** and **<transient-change>** elements are the same as the contents of the **<configuration>** tag element used in the Junos XML protocol operation **<load-configuration>**. For complete details about the **<load-configuration>** element, see the *Junos XML Management Protocol Developer Guide*.

### Related Documentation

- [Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292](#)
- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Persistent and Transient Configuration Changes on page 299](#)

## Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts

Junos OS commit scripts enforce custom configuration rules and can automatically change the configuration when it does not comply with your custom configuration rules. To generate a persistent or transient change in SLAX and XSLT commit scripts:

1. At the start of the script, include the Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) boilerplate from [“Required Boilerplate for Commit Scripts” on page 251](#). It is reproduced here for convenience:

### XSLT Boilerplate

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <!-- ... insert your code here ... -->
  </xsl:template>
</xsl:stylesheet>
```

### SLAX Boilerplate

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  /*
   * insert your code here
   */
}
```

2. At the position indicated by the comment *“insert your code here,”* include one or more XSLT programming instructions or their SLAX equivalents. Commonly used XSLT constructs include the following:
  - **<xsl:choose>** **<xsl:when>** **<xsl:otherwise>**—Conditional construct that causes different instructions to be processed in different circumstances. The **<xsl:choose>** instruction contains one or more **<xsl:when>** elements, each of which tests an XPath expression. If the test evaluates as true, the XSLT processor executes the instructions in the **<xsl:when>** element. The XSLT processor processes only the instructions contained in the first **<xsl:when>** element whose **test** attribute evaluates as true. If none of the **<xsl:when>** elements’ **test** attributes evaluate as true, the content of the **<xsl:otherwise>** element, if there is one, is processed.
  - **<xsl:for-each select="xpath-expression">**—Programming instruction that tells the XSLT processor to gather together a set of nodes and process them one by one.

The nodes are selected by the Extensible Markup Language (XML) Path Language (XPath) expression in the **select** attribute. Each of the nodes is then processed according to the instructions contained in the **<xsl:for-each>** instruction. Code inside an **<xsl:for-each>** instruction is evaluated recursively for each node that matches the XPath expression. The context is moved to the node during each pass.

- **<xsl:if test="xpath-expression">**—Conditional construct that causes instructions to be processed if the XPath expression in the **test** attribute evaluates to **true**.

For example, the following XSLT programming instructions select each SONET/SDH interface that does not have the MPLS protocol family enabled:

```
<xsl:for-each select="interfaces/interface[starts-with(name, 'so-')]/unit">
  <xsl:if test="not(family/mp1s)">
```

In SLAX, the **for-each** and **if** constructs look like this:

```
for-each (interfaces/interface[starts-with(name, 'so-')]/unit) {
  if (not(family/mp1s)) {
```

For more information about how to use programming instructions, including examples and pseudocode, see [“XSLT Programming Instructions Overview” on page 28](#). For information about writing scripts in SLAX instead of XSLT, see [“SLAX Overview” on page 57](#).

### 3. Include instructions for changing the configuration.

There are two ways to generate a persistent change and two ways to generate a transient change:

- To generate a persistent change, you can either reference the **jcs:emit-change** template or include a **<change>** element.
- To generate a transient change, you can either reference the **jcs:emit-change** template and pass in the **tag** parameter with **'transient-change'** selected or include a **<transient-change>** element.

The **jcs:emit-change** template allows for more efficient, less error-prone scripting because you can define the content of the change without specifying the complete XML hierarchy for the affected statement. Instead, the XML hierarchy is defined in the XPath expression contained in the script's programming instruction.

Consider the following examples. Both of the persistent change examples have the same result, even though they place the **unit** statement in different locations in the **<xsl:for-each>** and **<xsl:if>** programming instructions. In both cases, the script searches for SONET/SDH interfaces that do not have the MPLS protocol family enabled, adds the **family mpls** statement at the **[edit interfaces so-fpc/pic/port unit logical-unit-number]** hierarchy level, and emits a warning message stating that the configuration has been changed. Likewise, both of the transient change examples have the same result. They both set Point-to-Point Protocol (PPP) encapsulation on all SONET/SDH interface that have IP version 4 (IPv4) enabled.

### Persistent Change Generated with the `jcs:emit-change` Template

In this example, the content of the persistent change (contained in the `content` parameter) is specified without including the complete XML hierarchy. Instead, the XPath expression in the `<xsl:for-each>` programming instruction sets the context for the change.

The message parameter is also included. This parameter causes the `jcs:emit-change` template to call the `<xnm:warning>` template, which sends a warning notification to the CLI. The message parameter automatically includes the current hierarchy information in the warning message. (For more information about the parameters available with the `jcs:emit-change` template, see [emit-change Template \(SLAX and XSLT\)](#) and [emit\\_change\(Python\)](#).)

```
<xsl:for-each select="interfaces/interface[starts-with(name, 'so-')]/unit">
  <xsl:if test="not(family/mpls)">
    <xsl:call-template name="jcs:emit-change">
      <xsl:with-param name="content">
        <family>
          <mpls/>
        </family>
      </xsl:with-param>
      <xsl:with-param name="message">
        <xsl:text>Adding 'family mpls' to SONET interface.</xsl:text>
      </xsl:with-param>
    </xsl:call-template>
  </xsl:if>
</xsl:for-each>
```

### Persistent Change Generated with the `<change>` Element

In this example, the complete XML hierarchy leading to the affected statement must be included as child elements of the `<change>` element.

This example includes the current hierarchy information in the warning message by referencing the `jcs:edit-path` and `jcs:statement` templates. For more information about warning messages, see [“Overview of Generating Custom Warning, Error, and System Log Messages” on page 269](#).

```
<xsl:for-each select="interfaces/interface[starts-with(name, 'so-')]">
  <xsl:if test="not(unit/family/mpls)">
    <change>
      <interfaces>
        <interface>
          <name><xsl:value-of select="name"/></name>
          <unit>
            <name><xsl:value-of select="unit/name"/></name>
            <family>
              <mpls/>
            </family>
          </unit>
        </interface>
      </interfaces>
    </change>
    <xnm:warning>
      <xsl:call-template name="jcs:edit-path"/>
      <xsl:call-template name="jcs:statement">
        <xsl:with-param name="dot" select="unit/name"/>
      </xsl:call-template>
      <message>Adding 'family mpls' to SONET interface.</message>
```

```

        </xnm:warning>
    </xsl:if>
</xsl:for-each>

```

### Transient Change Generated with the `jcs:emit-change` Template

In this example, the content of the transient change (contained in the **content** parameter) is specified without including the complete XML hierarchy. Instead, the XPath expression in the `<xsl:for-each>` programming instruction sets the context of the change. The **and** operator in the XPath expression means both operands are **true** when converted to Booleans; the second operand is not evaluated if the first operand is **false**.

The tag parameter is included with 'transient-change' selected. Without the **tag** parameter, the `jcs:emit-change` template generates a persistent change by default. (For more information about the parameters available with the `jcs:emit-change` template, see [emit-change Template \(SLAX and XSLT\)](#) and [emit\\_change\(Python\)](#).)

```

<xsl:for-each select="interfaces/interface[starts-with(name, 'so-') \
    and unit/family/inet]">
    <xsl:call-template name="jcs:emit-change">
        <xsl:with-param name="tag" select="'transient-change'"/>
        <xsl:with-param name="content">
            <encapsulation>ppp</encapsulation>
        </xsl:with-param>
    </xsl:call-template>
</xsl:for-each>

```

### Transient Change Generated with the `<transient-change>` Element

In this example, the complete XML hierarchy leading to the affected statement must be included as child elements of the `<transient-change>` element.

```

<xsl:for-each select="interfaces/interface[starts-with(name, 'so-') \
    and unit/family/inet]">
    <transient-change>
        <interfaces>
            <interface>
                <name><xsl:value-of select="name"/></name>
                <encapsulation>ppp</encapsulation>
            </interface>
        </interfaces>
    </transient-change>
</xsl:for-each>

```

4. Save the script with a meaningful name.
5. Copy the script to either the `/var/db/scripts/commit` directory on the device hard disk or the `/config/scripts/commit` directory on the flash drive. For information about setting the storage location for commit scripts, see [“Storing Scripts in Flash Memory” on page 664](#).

If the device has dual Routing Engines and you want the script to take effect on both of them, you must copy the script to the `/var/db/scripts/commit` or the `/config/scripts/commit` directory on both Routing Engines. The **commit synchronize** command does not copy scripts between Routing Engines.

6. Enable the script by including the **file filename** statement at the **[edit system scripts commit]** hierarchy level.

```
[edit system scripts commit]
user@host# set file filename
```

7. If the script makes transient changes, include the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level.

```
[edit system scripts commit]
user@host# set allow-transients
```

If all the commit scripts run without errors, any transient changes are loaded into the checkout configuration, but not to the candidate configuration. Any persistent changes are loaded into the candidate configuration. The commit process then continues by validating the configuration and propagating changes to the affected processes on the device.

To display the configuration with both persistent and transient changes applied, issue the **show | display commit-scripts** configuration mode command.

```
[edit]
user@host# show | display commit-scripts
```

To display the configuration with only persistent changes applied, issue the **show | display commit-scripts no-transients** configuration mode command.

```
[edit]
user@host# show | display commit-scripts no-transients
```

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute to a configuration element produces the same behavior as a **replace:** tag in a **load replace** operation. Both persistent and transient changes are loaded into the configuration with the **load replace** behavior. However, persistent changes are loaded into the candidate configuration, and transient changes are loaded into the checkout configuration.

#### Related Documentation

- [SLAX and XSLT Commit Script Tag Elements to Use When Generating Persistent and Transient Configuration Changes on page 299](#)
- [Removing a Persistent or Transient Configuration Change Using Commit Scripts on page 300](#)
- [Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts on page 301](#)
- [Generating a Persistent or Transient Configuration Change in Python Commit Scripts on page 297](#)
- [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)



## Generating a Persistent or Transient Configuration Change in Python Commit Scripts

Junos OS commit scripts enforce custom configuration rules and can automatically change the configuration when it does not comply with your custom configuration rules. To generate a persistent or transient change using Python commit scripts:

1. At the start of the script, include the Python boilerplate from “[Required Boilerplate for Commit Scripts](#)” on page 251, which is reproduced here for convenience:

```
from junos import Junos_Configuration
import jcs

if __name__ == '__main__':
    # insert your code here
```

2. Include one or more programming instructions that test for your custom configuration rules.

For example, the following code selects each SONET/SDH interface that does not have the MPLS protocol family enabled:

```
# Get configuration root object
root = Junos_Configuration

for element in root.xpath("./interfaces/ \
    interface[starts-with(name, 'so-')]"):
    if element.find('unit/family/mp1s') is None:
```

3. Create an XML string that instructs Junos OS how to modify the configuration.

This example enables the MPLS protocol family for the selected interfaces.

```
if_name = element.find('name').text
unit_name = element.find('unit/name').text
change_xml = """
<interfaces>
  <interface>
    <name>{0}</name>
    <unit>
      <name>{1}</name>
      <family>
        <mp1s>
        </mp1s>
      </family>
    </unit>
  </interface>
</interfaces>
""".format(if_name, unit_name).strip()
```

4. Include any additional required or optional code. This example generates a warning message that is displayed on the CLI when the commit scripts updates a SONET/SDH interface.

```
jcs.emit_warning("Adding 'family mp1s' to SONET interface: " +
    if_name)
```

5. To generate the persistent or transient change, call the `jcs.emit_change` method, and specify the type of change, either 'change' or 'transient-change', in the argument list.

```
jcs.emit_change(change_xml, "change", "xml")
```

6. Save the script with a meaningful name.
7. Copy the script to either the `/var/db/scripts/commit` directory on the device hard disk or the `/config/scripts/commit` directory on the flash drive.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

8. Enable the script by including the `file filename` statement at the `[edit system scripts commit]` hierarchy level.

```
[edit]
user@host# set system scripts commit file filename
```

9. If the script generates any transient changes, configure the `allow-transients` statement at the `[edit system scripts commit]` hierarchy level.

```
[edit]
user@host# set system scripts commit allow-transients
```

10. Enable Python automation scripts on the device.

```
[edit system scripts]
user@host# set language python
```

11. Issue the `commit` command to commit the configuration.

```
user@host# commit
```

The resulting script searches for SONET/SDH interfaces that do not have the MPLS protocol family enabled, adds the `family mpls` statement at the `[edit interfaces so-fpc/pic/port unit logical-unit-number]` hierarchy level as a persistent change, and emits a warning message stating that the configuration has been changed.

```
from junos import Junos_Configuration
import jcs

def main():
    # Get configuration root object
    root = Junos_Configuration

    for element in root.xpath("./interfaces/ \
        interface[starts-with(name,'so-')]"):
        if element.find('unit/family/mpls') is None:
            if_name = element.find('name').text
            unit_name = element.find('unit/name').text
            change_xml = """
            <interfaces>
            <interface>
            <name>{0}</name>
            <unit>
            <name>{1}</name>
```

```

        <family>
            <mpls>
            </mpls>
        </family>
    </unit>
</interface>
</interfaces>
"".format(if_name, unit_name).strip()
jcs.emit_change(change_xml, "change", "xml")
jcs.emit_warning("Adding 'family mpls' to SONET interface: " +
if_name)

if __name__ == '__main__':
    main()

```

If all enabled commit scripts run without errors, any persistent changes are loaded into the candidate configuration, and any transient changes are loaded into the checkout configuration, but not to the candidate configuration. The commit process then continues by validating the configuration and propagating changes to the affected processes on the device.

To display the configuration with both persistent and transient changes applied, issue the **show | display commit-scripts** configuration mode command.

```
[edit]
user@host# show | display commit-scripts
```

To display the configuration with only persistent changes applied, issue the **show | display commit-scripts no-transients** configuration mode command.

```
[edit]
user@host# show | display commit-scripts no-transients
```

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute to a configuration element produces the same behavior as a **replace:** tag in a **load replace** operation. Both persistent and transient changes are loaded into the configuration with the **load replace** behavior. However, persistent changes are loaded into the candidate configuration, and transient changes are loaded into the checkout configuration.

#### Related Documentation

- [Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292](#)
- [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)

## SLAX and XSLT Commit Script Tag Elements to Use When Generating Persistent and Transient Configuration Changes

Junos OS commit scripts enforce custom configuration rules and can automatically change the configuration when it does not comply with your custom configuration rules. [Table 35 on page 300](#) summarizes the tag elements that you can include in SLAX and XSLT commit scripts to generate persistent and transient changes. To see how data values are supplied within a script, see [“Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts” on page 301](#).

Table 35: Tags and Attributes for Creating Configuration Changes in SLAX and XSLT Commit Scripts

Data Item, XML Element, or Attribute	Description
<b>Container Tags</b>	
<code>&lt;change&gt;</code>	Request that the Junos XML protocol server load configuration data into the candidate configuration.
<code>&lt;transient-change&gt;</code>	Request that the Junos XML protocol server load configuration data into the checkout configuration.
<b>Content Tags</b>	
<code>&lt;jcs:emit-change&gt;</code>	This is a template in the file <code>junos.xsl</code> . This template converts the contents of the <code>&lt;xsl:with-param&gt;</code> element into a <code>&lt;change&gt;</code> request.
<code>&lt;xsl:with-param name="content"&gt;</code>	You use the <code>content</code> parameter with the <code>jcs:emit-change</code> template. It allows you to include the content of the change, relative to <code>dot</code> .
<code>&lt;xsl:with-param name="tag" select="'transient-change'"/&gt;</code>	<p>Convert the contents of the <code>content</code> parameter into a <code>&lt;transient-change&gt;</code> request.</p> <p>You use the <code>tag</code> parameter with the <code>jcs:emit-change</code> template.</p> <p>By default, the <code>jcs:emit-change</code> template converts the contents of the <code>content</code> parameter into a <code>&lt;change&gt;</code> (persistent change) request.</p>

## Removing a Persistent or Transient Configuration Change Using Commit Scripts

After a commit script changes the configuration, you can remove the change and return the configuration to its previous state.

For persistent changes only, you can undo the configuration change by issuing the **delete**, **deactivate**, or **rollback** configuration mode command and committing the configuration. For both persistent and transient changes, you must remove, delete, or deactivate the associated commit script, or else the commit script regenerates the change during a subsequent commit operation.

Deleting the **file filename** statement from the configuration effectively disables the functionality associated with the corresponding commit script. Deactivating the statement adds the **inactive:** tag to the statement, effectively commenting out the statement from the configuration. Statements marked as inactive do not take effect when you issue the **commit** command.

To reverse the effect of a commit script and prevent the script from running again:

1. For persistent changes only, delete or deactivate the statement that was added by the commit script:

```
[edit]
user@host# delete (statement | identifier)
- OR -
user@host# deactivate (statement | identifier)
```

Alternatively, you can roll back the configuration to a candidate that does not contain the statement.

```
[edit]
user@host# rollback number
```

2. Either delete or deactivate the commit script, or remove or comment out the section of code that generates the unwanted change. To delete or deactivate the script, issue one of the following commands.

```
[edit]
user@host# delete system scripts commit file filename
- OR -
user@host# deactivate system scripts commit file filename
```

3. Issue the **commit** command:

```
[edit]
user@host# commit
```

4. If you are deleting the reference to the script from the configuration, you can also remove the file from commit scripts storage directory (either **/var/db/scripts/commit** on the hard disk or **/config/scripts/commit** on the flash drive). To do this, exit configuration mode and issue the **file delete** operational mode command:

```
[edit]
user@host# exit

user@host> file delete /var/db/scripts/commit/filename
- OR -
user@host> file delete /config/scripts/commit/filename
```

- |                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts on page 287</a></li><li>• <a href="#">Storing Scripts in Flash Memory on page 664</a></li></ul> |
|------------------------------|---|

---

## Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts

- [Example: Generating a Persistent Change on page 301](#)
- [Example: Generating a Transient Change on page 306](#)

### Example: Generating a Persistent Change

Junos OS commit scripts enable users to customize the validation process of their configurations in accordance with their own practices and policies and enforce custom configuration rules during the commit process. This example demonstrates a commit script that generates a persistent change that adds the **family mpls** statement in the configuration of SONET/SDH interfaces when the statement is not already included in the configuration. If you do not explicitly configure the MPLS protocol family on an interface, the interface is not enabled for MPLS applications.

- [Requirements on page 302](#)
- [Overview and Commit Script on page 302](#)

- [Configuration on page 304](#)
- [Verification on page 305](#)

## Requirements

This example uses the following hardware and software components:

- Device running Junos OS with one or more SONET/SDH interfaces.
- Junos OS Release 16.1R3 or later release when using a Python script.

## Overview and Commit Script

The commit script in this example finds all SONET/SDH interfaces that have a logical interface configured but that do not have the **family mpls** statement configured. For these interfaces, the script adds the **family mpls** statement to the interface configuration as a persistent change at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. The script is shown in SLAX, XSLT, and Python.

The SLAX and XSLT versions of the commit script generate the persistent change by using the **jcs:emit-change** template, which is a helper template contained in the **junos.xsl** import file. The **tag** parameter of the **jcs:emit-change** template is omitted, which directs the script to emit the change as a persistent change. The **content** parameter of the **jcs:emit-change** template includes the configuration statements to be added as a persistent change. The **message** parameter of the **jcs:emit-change** template includes the warning message to be displayed in the CLI, notifying you that the configuration has been changed.

The Python version of the commit script generates the persistent change by using the **jcs.emit\_change()** function, which is imported from the **jcs** module. The Python script indicates that this is a persistent change by passing in the positional argument 'change'.

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name,
'so-')]/unit">
      <xsl:if test="not(family/mpls)">
        <xsl:call-template name="jcs:emit-change">
          <xsl:with-param name="message">
            <xsl:text>Adding 'family mpls' to SONET/SDH
interface.</xsl:text>
          </xsl:with-param>
          <xsl:with-param name="content">
            <family>
              <mpls/>
            </family>
          </xsl:with-param>
        </xsl:call-template>
      </xsl:if>
```

```

        </xsl:for-each>
    </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (interfaces/interface[starts-with(name, 'so-')]/unit) {
    if (not(family/mps)) {
      call jcs:emit-change() {
        with $message = {
          expr "Adding 'family mpls' to SONET/SDH interface.";
        }
        with $content = {
          <family> {
            <mps>;
          }
        }
      }
    }
  }
}

```

**Python Syntax**

```

from junos import Junos_Configuration
import jcs

def main():
    # Get configuration root object
    root = Junos_Configuration

    for element in root.xpath("./interfaces/ \
interface[starts-with(name,'so-')/unit]"):
        if element.find('family/mps') is None:
            if_name = element.find('..name').text
            unit_name = element.find('name').text
            change_xml = """
            <interfaces>
            <interface>
            <name>{0}</name>
            <unit>
            <name>{1}</name>
            <family>
            <mps>
            </mps>
            </family>
            </unit>
            </interface>
            </interfaces>
            """.format(if_name, unit_name).strip()
            jcs.emit_change(change_xml, "change", "xml")
            jcs.emit_warning("Adding 'family mpls' to SONET/SDH interface: "
+ if_name)

if __name__ == '__main__':
    main()

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script.

1. Copy the script into a text file, name the file **mpls.xml**, **mpls.slax**, or **mpls.py** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file** statement at the **[edit system scripts commit]** hierarchy level and the appropriate filename, for example:

```
[edit]
user@host# set system scripts commit file mpls.xml
```

3. If the script is written in Python, configure the **language python** statement.

```
[edit]
user@host# set system scripts language python
```

4. To test that the commit script generates the persistent change correctly, make sure that the configuration contains the condition that elicits the change. To test this script, ensure that the **family mpls** statement is not included at the **[edit interfaces so-fpc/pic/port unit logical-unit-number]** hierarchy level for at least one SONET/SDH interface.
5. Issue the **commit check** command to preview a trace of commit script processing to verify that the script will add the persistent change to the candidate configuration. The **commit check** command verifies the syntax of the configuration prior to a commit, but it does not commit the changes.

The commit script in this example produces a message for each change it makes. Use the **commit check** command to preview these messages to determine whether the script will update the configuration with the **family mpls** statement for the appropriate interfaces.

Issue the **commit check | display xml** command to display the XML-formatted version of the message. The sample output indicates that the script will add the **family mpls** statement to the so-2/3/4.0 interface configuration during the commit operation.

```
[edit]
user@host# commit check | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.2R1/junos">
  <commit-results>
    <routing-engine junos:style="normal">
      <name>re0</name>
      <xnm:warning xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm">
        <edit-path>
          [edit interfaces interface so-2/3/4 unit 0]
        </edit-path>
        <message>
```



```

        Adding 'family mpls' to SONET/SDH interface.
    </message>
    </xnm:warning>
    <commit-check-success/>
  </routing-engine>
</commit-results>
</rpc-reply>

```

6. To display a detailed trace of commit script processing, issue the **commit check | display detail** command. In the sample output, there is one persistent change that will be loaded into the configuration during the commit operation.

[edit]

```

user@host# commit check | display detail
2011-06-17 14:17:35 PDT: reading commit script configuration
2011-06-17 14:17:35 PDT: testing commit script configuration
2011-06-17 14:17:35 PDT: opening commit script
'/var/db/scripts/commit/mps.xml'
2011-06-17 14:17:35 PDT: reading commit script 'mps.xml'
2011-06-17 14:17:35 PDT: running commit script 'mps.xml'
2011-06-17 14:17:35 PDT: processing commit script 'mps.xml'
2011-06-17 14:17:35 PDT: no errors from mps.xml
2011-06-17 14:17:35 PDT: saving commit script changes for script mps.xml
2011-06-17 14:17:35 PDT: summary of script mps.xml: changes 1, transients
0, syslog 0
2011-06-17 14:17:35 PDT: start loading commit script changes
2011-06-17 14:17:35 PDT: loading commit script changes into real db
2011-06-17 14:17:35 PDT: finished commit script changes into real db
2011-06-17 14:17:35 PDT: no transient commit script changes
2011-06-17 14:17:35 PDT: finished loading commit script changes
2011-06-17 14:17:35 PDT: copying juniper.db to juniper.data+
2011-06-17 14:17:35 PDT: finished copying juniper.db to juniper.data+
...
configuration check succeeds

```

7. After verifying that the script produces the correct changes, issue the **commit** command to start the commit operation and execute the script.

```
user@host# commit
```

## Verification

### Verifying the Configuration

**Purpose** Verify that the correct changes are integrated into the configuration.

**Action** After executing the commit operation, view the configuration by issuing the **show interfaces** configuration mode command. If the MPLS protocol family is not enabled on one or more SONET/SDH interfaces before the script runs, the output is similar to the following:

```
[edit]
user@host# show interfaces
... other configured interface types ...
so-2/3/4 {
  unit 0 {
    family mpls; # Added by persistent change
  }
}
... other configured interface types ...
```

## Example: Generating a Transient Change

This example uses a commit script to set PPP encapsulation on all SONET/SDH interfaces with the IPv4 protocol family enabled. The changes are added as transient changes.

- [Requirements on page 306](#)
- [Overview and Commit Script on page 306](#)
- [Configuration on page 308](#)
- [Verification on page 309](#)
- [Troubleshooting on page 310](#)

### Requirements

---

This example uses the following hardware and software components:

- Device running Junos OS with one or more SONET/SDH interfaces.
- Junos OS Release 16.1R3 or later release when using a Python script.

### Overview and Commit Script

---

The commit script in this example finds all SONET/SDH interfaces with the IPv4 protocol family enabled in the configuration and adds the **encapsulation ppp** statement to the interface configuration. The commit script generates a transient change, which adds the change to the checkout configuration but not the candidate configuration. The script is shown in SLAX, XSLT, and Python.

The SLAX and XSLT versions of the commit script generate the transient change by using the **jcs:emit-change** template, which is a helper template contained in the **junos.xsl** import file. The **tag** parameter of the **jcs:emit-change** template has the value **transient-change**, which directs the script to emit the change as a transient change rather than a persistent change. The **content** parameter of the **jcs:emit-change** template includes the configuration statements to be added as a transient change.

The Python version of the commit script generates the transient change by using the **jcs.emit\_change()** function, which is imported from the **jcs** module. The Python script

indicates that this is a transient change by passing in the positional argument 'transient-change'.

**XSLT Syntax**

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name, 'so-')
      and unit/family/inet]">
      <xsl:call-template name="jcs:emit-change">
        <xsl:with-param name="tag" select="'transient-change'"/>
        <xsl:with-param name="content">
          <encapsulation>ppp</encapsulation>
        </xsl:with-param>
      </xsl:call-template>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (interfaces/interface[starts-with(name, 'so-') and
unit/family/inet]) {
    call jcs:emit-change($tag = 'transient-change') {
      with $content = {
        <encapsulation> "ppp";
      }
    }
  }
}
```

**Python Syntax**

```
from junos import Junos_Configuration
import jcs

def main():
    # Get configuration root object
    root = Junos_Configuration

    for element in root.xpath("./interfaces/ \
interface[starts-with(name,'so-') and unit/family/inet]"):
        if_name = element.find('name').text
        change_xml = ""
        <interfaces>
        <interface>
        <name>{0}</name>
        <encapsulation>ppp</encapsulation>
        </interface>
        </interfaces>
        "".format(if_name).strip()
        jcs.emit_change(change_xml, "transient-change", "xml")
```

```
jcs.emit_warning("Adding 'ppp' encapsulation to SONET interface: " +
if_name)
jcs.emit_warning(change_xml)

if __name__ == '__main__':
    main()
```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script.

1. Copy the script into a text file, name the file **encap-ppp.xml**, **encap-ppp.slax**, or **encap-ppp.py** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file** statement at the **[edit system scripts commit]** hierarchy level and the appropriate filename, for example:

```
[edit]
user@host# set system scripts commit file encap-ppp.xml
```

3. Include the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level. This enables commit scripts to load transient changes into the checkout configuration.

```
[edit]
user@host# set system scripts commit allow-transients
```

4. If the script is written in Python, configure the **language python** statement.

```
[edit]
user@host# set system scripts language python
```

5. To test that the commit script generates the transient change correctly, make sure that the configuration contains the condition that elicits the change. Ensure that the **encapsulation ppp** statement is not included at the **[edit interfaces so-fpc/pic/port]** hierarchy level for at least one SONET/SDH interface.
6. Issue the **commit check** command to preview a trace of commit script processing to verify that the script will add the transient change to the checkout configuration. The **commit check** command verifies the syntax of the configuration prior to a commit, but it does not commit the changes.

Issue the **commit check | display detail** command to display a detailed trace of commit script processing. In the sample output, there are two transient changes that are loaded into the checkout configuration.

```
[edit]
user@host# commit check | display detail
2011-06-15 12:07:30 PDT: reading commit script configuration
```

```
2011-06-15 12:07:30 PDT: testing commit script configuration
2011-06-15 12:07:30 PDT: opening commit script
'/var/db/scripts/commit/encap-ppp.xml'
2011-06-15 12:07:30 PDT: reading commit script 'encap-ppp.xml'
2011-06-15 12:07:30 PDT: running commit script 'encap-ppp.xml'
2011-06-15 12:07:30 PDT: processing commit script 'encap-ppp.xml'
2011-06-15 12:07:30 PDT: no errors from encap-ppp.xml
2011-06-15 12:07:30 PDT: saving commit script changes for script encap-ppp.xml
2011-06-15 12:07:30 PDT: summary of script encap-ppp.xml: changes 0,
transients 2 (allowed), syslog 0
2011-06-15 12:07:30 PDT: start loading commit script changes
2011-06-15 12:07:30 PDT: no commit script changes
2011-06-15 12:07:30 PDT: updating transient changes into transient tree
2011-06-15 12:07:30 PDT: finished loading commit script changes
2011-06-15 12:07:30 PDT: copying juniper.db to juniper.data+
2011-06-15 12:07:30 PDT: finished copying juniper.db to juniper.data+
2011-06-15 12:07:30 PDT: exporting juniper.conf
2011-06-15 12:07:30 PDT: merging transient changes
...
configuration check succeeds
```

7. After verifying that the script produces the correct changes, issue the **commit** command to start the commit operation and execute the script.

```
user@host# commit
```

---

## Verification

### *Verifying the Configuration*

**Purpose** Verify that the correct changes are integrated into the checkout configuration. If there are one or more SONET/SDH interfaces with the IPv4 protocol family enabled, you should see the **encapsulation ppp** statement added as a transient change to the interface hierarchy.

**Action** To view the configuration with transient changes, issue the **show interfaces | display commit-scripts** configuration mode command. The **show interfaces | display commit-scripts** command displays all the statements that are in the configuration, including statements that are generated by transient changes. If there are one or more SONET/SDH interfaces with the IPv4 protocol family enabled, the output is similar to the following:

```
[edit]

user@host# show interfaces | display commit-scripts
... other configured interface types ...
so-1/2/3 {
    mtu 576;
    encapsulation ppp; /* Added by transient change. */
    unit 0 {
        family inet {
            address 10.0.0.3/32;
        }
    }
}
so-1/2/4 {
    encapsulation ppp; /* Added by transient change. */
    unit 0 {
        family inet {
            address 10.0.0.4/32;
        }
    }
}
so-2/3/4 {
    encapsulation cisco-hdlc; # Not affected by this script, because IPv4 protocol
                                # family is not configured on this interface.
    unit 0 {
        family mpls;
    }
}
... other configured interface types ...
```

---

## Troubleshooting

### Troubleshooting Commit Errors

**Problem** The CLI generates an invalid transient change error, and the commit fails.

```
user@host# commit check
error: invalid transient change generated by commit script: encap-ppp.xml
warning: 1 transient change was generated without [system scripts commit
allow-transients]
error: 1 error reported by commit scripts
error: commit script failure
```

**Solution** You must configure the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level to enable commit scripts to load transient changes into the checkout configuration.

Configure the following statement to allow transient changes:

```
[edit]
user@host# set system scripts commit allow-transients
```

**Related  
Documentation**

- [Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292](#)
- [Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts on page 287](#)
- [Removing a Persistent or Transient Configuration Change Using Commit Scripts on page 300](#)
- [emit-change Template \(SLAX and XSLT\) and emit\\_change\(Python\) on page 230](#)





# Creating Custom Configuration Syntax with Commit Script Macros

- [Overview of Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements on page 320](#)
- [Example: Creating Custom Configuration Syntax with Commit Script Macros on page 322](#)

## Overview of Creating Custom Configuration Syntax with Commit Script Macros

---

Using commit script macros, you can create a custom configuration language based on simplified syntax that is relevant to your network design. This means you can use your own aliases for frequently used configuration statements.

Commit scripts generally impose restrictions on the Junos OS configuration and automatically correct configuration mistakes when they occur (as discussed in [“Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts” on page 287](#)). However, macros are useful for an entirely different reason. Commit scripts that contain macros do not generally correct configuration mistakes, nor do they necessarily restrict configuration. Instead, they provide a way to simplify and speed configuration tasks, thereby preventing mistakes from occurring at all.

For a detailed example of how macros can save time and effort, see [“Example: Automatically Configuring Logical Interfaces and IP Addresses” on page 343](#).

### Related Documentation

- [Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements on page 320](#)
- [Example: Creating Custom Configuration Syntax with Commit Script Macros on page 322](#)


## Creating Custom Configuration Syntax with Commit Script Macros

---

Commit script macros enable you to create custom configuration syntax and expand it into standard Junos OS configuration statements. Your custom syntax serves as input

to a commit script. The output of the commit script is standard Junos OS configuration syntax, as shown in [Figure 7 on page 314](#). The standard Junos OS statements are added to the configuration to cause your intended operational changes.

**Figure 7: Macro Input and Output**



Macros use either persistent or transient change elements to expand your custom syntax into standard Junos OS configuration statements. If you use persistent changes, both the custom syntax and the standard Junos OS syntax appear in the candidate configuration. If you use transient changes, the custom syntax appears in the candidate configuration, and the standard Junos OS syntax is copied to the checkout configuration only.

This section discusses the following topics:

- [Creating a Custom Syntax on page 314](#)
- [<data> Element on page 315](#)
- [Expanding the Custom Syntax on page 316](#)
- [Other Ways to Use Macros on page 319](#)

## Creating a Custom Syntax

Macros work by locating **apply-macro** statements in the candidate configuration and using the values specified in the **apply-macro** statement as parameters to a set of instructions defined in a commit script. In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides to the script the data necessary to generate a complex configuration.

To enter custom syntax, you include the **apply-macro** statement at any hierarchy level and specify any data that you want inside the **apply-macro** statement, for example:

```

apply-macro macro-name {
    parameter-name parameter-value;
}
  
```

You can include the **apply-macro** statement at any level of the configuration hierarchy. In this sense, the **apply-macro** statement is similar to the **apply-groups** statement. Each **apply-macro** statement must be uniquely named, relative to other **apply-macro** statements at the same hierarchy level.

An **apply-macro** statement can contain a set of parameters with optional values. The corresponding commit script can refer to the macro name, its parameters, or the parameters' values. When the script inspects the configuration and finds the data, the script performs the actions specified by the corresponding persistent or transient change.

For example, given the following configuration stanza, you can write script instructions to generate a standard configuration based on the name of the parameter:

```

protocols {
    mpls {
        apply-macro blue-type-lsp {
  
```

```

        color blue;
    }
}

```

The following `<xsl:for-each>` programming instruction finds **apply-macro** statements at the **[edit protocols mpls]** hierarchy level that contain a parameter named **color**:

```
<xsl:for-each select="protocols/mppls/apply-macro[data/name = 'color']">
```

The following instruction creates a variable named **color** and assigns to the variable the value of the **color** parameter, which in this case is **blue**:

```
<xsl:variable name="color" select="data[name = 'color']/value"/>
```

The following instruction adds the **admin-groups** statement to the configuration and assigns the value of the **color** variable to the group name:

```

<transient-change>
  <protocols>
    <mppls>
      <admin-groups>
        <name>
          <xsl:value-of select="$color"/>
        </name>
      </admin-groups>
    </mppls>
  </protocols>
</transient-change>

```

The resulting configuration statements are as follows:

```

protocols {
  mpls {
    admin-groups {
      blue;
    }
  }
}

```

## <data> Element

In the XML rendering of the custom syntax within an **apply-macro** statement, parameters and their values are contained in **<name>** and **<value>** elements, respectively. The **<name>** and **<value>** elements are sibling children of the **<data>** element. For example, the **apply-macro blue-type-lsp** statement contains six parameters, as follows:

```

[edit protocols mpls]
apply-macro blue-type-lsp {
  10.1.1.1;
  10.2.2.2;
  10.3.3.3;
  10.4.4.4;
  color blue;
  group-value 0;
}

```

The parameters and values are rendered in Junos XML tag elements as follows:

```
[edit protocols mpls]
```

```
user@host# show | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <configuration>
    <protocols>
      <mpls>
        <apply-macro>
          <name>blue-type-lsp</name>
          <data>
            <name>10.1.1.1</name>
          </data>
          <data>
            <name>10.2.2.2</name>
          </data>
          <data>
            <name>10.3.3.3</name>
          </data>
          <data>
            <name>10.4.4.4</name>
          </data>
          <data>
            <name>color</name>
            <value>blue</value>
          </data>
          <data>
            <name>group-value</name>
            <value>0</value>
          </data>
        </apply-macro>
      </mpls>
    </protocols>
  </configuration>
</rpc-reply>
```

When you write commit script macros, referring to the **<data>**, **<name>**, and **<value>** elements enables you to extract and manipulate the parameters contained in **apply-macro** statements. For example, in the following **select** attribute, the XPath expression extracts the text contained in the **<value>** element that is a child of a **<data>** element that also contains a **<name>** child element with the text **color**. The variable declaration assigns the text of the **<value>** element to a variable named **color**.

```
<xsl:variable name="color" select="data[name = 'color']/value"/>
```

The SLAX equivalent is:

```
var $color = ./data[name='color']/value;
```

The Python equivalent, which assumes that **element** has selected an **apply-macro** element, is:

```
color = element.find("data[name='color']/value").text
```

## Expanding the Custom Syntax

In the corresponding commit script, you include one or more programming instructions that inspect the configuration for the **apply-macro** statement at a specified hierarchy level. Optionally, you can use the **data/name** expression to select a parameter in the **apply-macro** statement.

```
<xsl:for-each select="xpath-expression/apply-macro[data/name =  
'parameter-name']">
```

For example, the following XSLT programming instruction selects every **apply-macro** statement that contains the **color** parameter and that appears at the **[edit protocols mpls]** hierarchy level:

```
<xsl:for-each select="protocols/mppls/apply-macro[data/name = 'color']">
```

The SLAX equivalent is:

```
for-each (protocols/mppls/apply-macro[data/name = 'color']) {
```

The Python equivalent, which spans multiple lines for readability, is:

```
for element in Junos_Configuration.xpath \
    ("./protocols/mppls/apply-macro[data/name='color']"):
```

When expanding macros, a particularly useful programming instruction in XSLT scripts is the **<xsl:value-of>** instruction. This instruction selects a parameter value and uses it to build option values for Junos OS statements. For example, the following instruction concatenates the value of the **color** variable, the text **-lsp-**, and the current context node (represented by **"."**) to build a name for an LSP.

```
<label-switched-path>
  <name>
    <xsl:value-of select="concat($color, '-lsp-', .)"/>
  </name>
</label-switched-path>
```

SLAX uses the underscore (**\_**) to concatenate values.

```
<label-switched-path> {
  <name> $color _ '-lsp-' _ .;
}
```

When the script includes instructions to find the necessary data, you can provide content for a persistent or transient change that uses the data to construct a standard Junos OS configuration.

The following transient change creates an administration group and adds the **label-switched-path** statement to the configuration. The label-switched path is assigned a name that concatenates the value of the **color** variable, the text **-lsp-**, and the currently selected IP address represented by the period (**"."**). The transient change also adds the **to** statement and assigns the currently selected IP address. Finally, the transient change adds the **admin-group include-any** statement and assigns the value of the **color** variable.

```
<transient-change>
  <protocols>
    <mppls>
      <admin-groups>
        <name><xsl:value-of select="$color"/></name>
        <group-value><xsl:value-of select="$group-value"/></group-value>
      </admin-groups>
      <xsl:for-each select="data[not(value)]/name">
        <label-switched-path>
          <name><xsl:value-of select="concat($color, '-lsp-',
.)"/></name>
          <to><xsl:value-of select="."/></to>
          <admin-group>
            <include-any><xsl:value-of
select="$color"/></include-any>
          </admin-group>
        </label-switched-path>
```

```

        </xsl:for-each>
      </mpls>
    </protocols>
  </transient-change>

```

The SLAX equivalent is:

```

<transient-change> {
  <protocols> {
    <mpls> {
      <admin-groups> {
        <name> $color;
        <group-value> $group-value;
      }
      for-each (data[not(value)]/name) {
        <label-switched-path> {
          <name> $color _ '-lsp-' _ .;
          <to> .;
          <admin-group> {
            <include-any> $color;
          }
        }
      }
    }
  }
}

```

Similarly in Python:

```

lsp_config = ""
for element2 in element.xpath("data[not(value)]/name"):
    lsp_config = lsp_config + """
    <label-switched-path>
      <name>{0}-lsp-{1}</name>
      <to>{2}</to>
      <admin-group>
        <include-any>{3}</include-any>
      </admin-group>
    </label-switched-path>
    """.format(color, element2.text, element2.text, color)

change_xml = """
<protocols>
  <mpls>
    <admin-groups>
      <name>{0}</name>
      <group-value>{1}</group-value>
    </admin-groups>
    {2}
  </mpls>
</protocols>
""".format(color, group_value, lsp_config).strip()

jcs.emit_change(change_xml, "transient-change", "xml")

```



**NOTE:** The example shown here is partial. For a full example, see [“Example: Creating Custom Configuration Syntax with Commit Script Macros” on page 322](#).

After committing the configuration, the script runs, and the resulting full configuration looks like this:

```
[edit]
protocols {
  mpls {
    admin-groups {
      blue 0;
    }
    label-switched-path blue-lsp-10.1.1.1 {
      to 10.1.1.1;
      admin-group include-any blue;
    }
    label-switched-path blue-lsp-10.2.2.2 {
      to 10.2.2.2;
      admin-group include-any blue;
    }
    label-switched-path blue-lsp-10.3.3.3 {
      to 10.3.3.3;
      admin-group include-any blue;
    }
    label-switched-path blue-lsp-10.4.4.4 {
      to 10.4.4.4;
      admin-group include-any blue;
    }
  }
}
```

The previous example demonstrates how you can use a simplified custom syntax to configure label-switched paths (LSPs). If your network design requires a large number of LSPs to be configured, using a commit script macro can save time, ensure consistency, and prevent configuration errors.

## Other Ways to Use Macros

The example discussed in [“Creating a Custom Syntax” on page 314](#) shows a macro that uses transient changes to create the intended operational impact. Alternatively, you can create a commit script that uses persistent changes to add the standard Junos OS statements to the candidate configuration and delete your custom syntax entirely. This way, a network operator who might be unfamiliar with your custom syntax can view the configuration file and see the full configuration rendered as standard Junos OS statements. Still, because the commit script macro remains in effect, you can quickly and easily create a complex configuration using your custom syntax.

In addition to the type of application discussed in [“Creating a Custom Syntax” on page 314](#), you can also use macros to prevent a commit script from performing a task. For example, a basic commit script that automatically adds MPLS configuration to interfaces can make an exception for interfaces you explicitly tag as not requiring MPLS, by testing for the presence of an **apply-macro** statement named **no-mpls**. For an example of this use of macros, see [“Example: Controlling LDP Configuration” on page 370](#).

You can use the **apply-macro** statement as a place to store external data. The commit script does not inspect the **apply-macro** statement, so the **apply-macro** statement has

no operational impact on the device, but the data can be carried in the configuration file to be used by external applications.

**Related Documentation**

- [Overview of Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements on page 320](#)
- [Example: Creating Custom Configuration Syntax with Commit Script Macros on page 322](#)

---

## Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements

---

Commit script macros enable you to expand custom configuration syntax into standard Junos OS configuration statements. By itself, the custom syntax in an **apply-macro** statement has no operational impact on the device. To give meaning to your syntax, there must be a corresponding commit script that uses the syntax as data for generating related Junos OS statements.

To write such a script:

1. At the start of the script, include the appropriate commit script boilerplate from [“Required Boilerplate for Commit Scripts” on page 251](#). It is reproduced here for convenience:

### XSLT Boilerplate

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <!-- ... insert your code here ... -->
  </xsl:template>
</xsl:stylesheet>
```

### SLAX Boilerplate

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  /*
   * insert your code here
   */
}
```



**Python Boilerplate**

```

from junos import Junos_Configuration
import jcs

if __name__ == '__main__':
    # insert your code here

```

2. At the position indicated by the comment “*insert your code here*,” include programming instructions that inspect the configuration for the **apply-macro** statement at a specified hierarchy level and change the configuration to include standard Junos OS syntax.

For an example that uses both types of instructions and includes a line-by-line analysis of the XSLT syntax, see [“Example: Creating Custom Configuration Syntax with Commit Script Macros” on page 322](#).

3. Save the script with a meaningful name.
4. Copy the script to either the `/var/db/scripts/commit` directory on the hard disk or the `/config/scripts/commit` directory on the flash drive.

For information about setting the storage location for commit scripts, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

If the device has dual Routing Engines and you want the script to take effect on both of them, you must copy the script to the `/var/db/scripts/commit` or the `/config/scripts/commit` directory on both Routing Engines. The **commit synchronize** command does not copy scripts between Routing Engines.

5. Enable the script by including the **file filename** statement at the **[edit system scripts commit]** hierarchy level.

```

[edit system scripts commit]
user@host# set file filename

```

6. If the script makes transient changes, include the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level.

```

[edit system scripts commit]
user@host# set allow-transients

```

If all the commit scripts run without errors, any persistent changes are loaded into the candidate configuration, and any transient changes are loaded into the checkout configuration, but not to the candidate configuration. The commit process then continues by validating the configuration and propagating changes to the affected processes on the device running Junos OS.

**Related Documentation**

- [Overview of Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
- [Example: Creating Custom Configuration Syntax with Commit Script Macros on page 322](#)

## Example: Creating Custom Configuration Syntax with Commit Script Macros

A Junos OS configuration can contain **apply-macro** statements with custom configuration syntax. By itself, the **apply-macro** statement has no operational impact on the device. Commit script macros process the custom configuration syntax and expand it into standard Junos OS configuration statements, which are then added as a persistent or transient change. This example demonstrates how to use commit script macros to inspect **apply-macro** statements and generate Junos OS configuration statements.

- [Requirements on page 322](#)
- [Overview and Commit Script on page 322](#)
- [Configuration on page 327](#)
- [Verification on page 328](#)

### Requirements

This example uses the following hardware and software components:

- Device running Junos OS.
- Junos OS Release 16.1R3 or later release when using a Python script.

### Overview and Commit Script

[Table 36 on page 322](#) shows a macro containing custom syntax and the corresponding expansion to standard Junos OS command-line interface (CLI) syntax.

**Table 36: Sample Macro and Junos OS CLI Expansion**

Custom Macro Syntax	Expanded Junos OS CLI Syntax
<pre> protocols {   mpls {     apply-macro       blue-type-lsp {         10.1.1.1;         10.2.2.2;         10.3.3.3;         10.4.4.4;         color blue;         group-value 0;       }   } } </pre>	<pre> protocols {   mpls {     admin-groups {       blue 0;     }     label-switched-path blue-lsp-10.1.1.1 {       to 10.1.1.1;       admin-group include-any blue;     }     label-switched-path blue-lsp-10.2.2.2 {       to 10.2.2.2;       admin-group include-any blue;     }     label-switched-path blue-lsp-10.3.3.3 {       to 10.3.3.3;       admin-group include-any blue;     }     label-switched-path blue-lsp-10.4.4.4 {       to 10.4.4.4;       admin-group include-any blue;     }   } } </pre>

In this example, the Junos OS management (mgd) process inspects the configuration, looking for **apply-macro** statements. For each **apply-macro** statement with the **color** parameter included at the **[edit protocols mpls]** hierarchy level, the script generates a transient change, using the data provided within the **apply-macro** statement to expand the macro into a standard Junos OS administrative group for LSPs.

For this example to work, an **apply-macro** statement must be included at the **[edit protocols mpls]** hierarchy level with a set of addresses, a **color**, and a **group-value** parameter. The commit script converts each address to an LSP configuration, and the script converts the **color** parameter into an administrative group.

Following are the commit script instructions that expand the macro in [Table 36 on page 322](#) and a line-by-line explanation of the script:

XSLT Syntax	<pre> 1  &lt;?xml version="1.0" standalone="yes"?&gt; 2  &lt;xsl:stylesheet version="1.0" 3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform" 4      xmlns:junos="http://xml.juniper.net/junos/*/junos" 5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm" 6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"&gt; 7      &lt;xsl:import href="../import/junos.xsl"/&gt;  8      &lt;xsl:template match="configuration"&gt; 9          &lt;xsl:variable name="mpls" select="protocols/mpls"/&gt; 10         &lt;xsl:for-each select="\$mpls/apply-macro[data/name = 'color']"&gt; 11             &lt;xsl:variable name="color" select="data[name = 'color']/value"/&gt; 12             &lt;xsl:variable name="group-value" select="data[name = \ 'group-value']/value"/&gt; 13             &lt;transient-change&gt; 14                 &lt;protocols&gt; 15                     &lt;mpls&gt; 16                         &lt;admin-groups&gt; 17                             &lt;name&gt; 18                                 &lt;xsl:value-of select="\$color"/&gt; 19                             &lt;/name&gt; 20                             &lt;group-value&gt; 21                                 &lt;xsl:value-of select="\$group-value"/&gt; 22                             &lt;/group-value&gt; 23                         &lt;/admin-groups&gt; 24                         &lt;xsl:for-each select="data[not(value)]/name"&gt; 25                             &lt;label-switched-path&gt; 26                                 &lt;name&gt; 27                                     &lt;xsl:value-of select="concat(\$color, '-lsp-', .)"/&gt; 28                                 &lt;/name&gt; 29                                 &lt;to&gt;&lt;xsl:value-of select="."/&gt;&lt;/to&gt; 30                             &lt;admin-group&gt; 31                                 &lt;include-any&gt; 32                                     &lt;xsl:value-of select="\$color"/&gt; 33                                 &lt;/include-any&gt; 34                             &lt;/admin-group&gt; 35                         &lt;/label-switched-path&gt; 36                     &lt;/xsl:for-each&gt; 37                 &lt;/mpls&gt; 38             &lt;/protocols&gt; 39         &lt;/transient-change&gt; 40     &lt;/xsl:for-each&gt; 41 &lt;/xsl:template&gt; </pre>
-------------	---

```
42    </xsl:stylesheet>
```

Lines 1 through 8 (and Lines 43 and 44) are the boilerplate that you include in every XSLT commit script. For brevity, Lines 1 through 8 are omitted here.

Line 9 assigns the **[edit protocols mpls]** hierarchy level to a variable called **mpls**.

```
9      <xsl:variable name="mpls" select="protocols/mppls"/>
```

Line 10 selects every **apply-macro** statement at the **[edit protocols mpls]** hierarchy level that contains the **color** parameter. The sample configuration in [Table 36 on page 322](#) contains only one **apply-macro** statement. Therefore, this **<xsl:for-each>** programming instruction takes effect only once.

```
10     <xsl:for-each select="$mpls/apply-macro[data/name = 'color']">
```

Line 11 assigns the value of the **color** parameter, in this case **blue**, to a variable called **color**.

```
11         <xsl:variable name="color" select="data[name =  
'color']/value"/>
```

Line 12 assigns the value of the **group-value** parameter, in this case **0**, to a variable called **group-value**.

```
12         <xsl:variable name="group-value" select="data[name = \  
'group-value']/value"/>
```

Lines 13 through 15 generate a transient change at the **[edit protocols mpls]** hierarchy level.

```
13     <transient-change>  
14         <protocols>  
15             <mpls>
```

Lines 16 through 23 add the **admin-groups** statement to the configuration and assign the value of the **color** variable to the group name and the value of the **group-value** variable to the group value.

```
16         <admin-groups>  
17             <name>  
18                 <xsl:value-of select="$color"/>  
19             </name>  
20             <group-value>  
21                 <xsl:value-of select="$group-value"/>  
22             </group-value>  
23         </admin-groups>
```

The resulting configuration statements are as follows:

```
admin-groups {  
    blue 0;  
}
```

Line 24 selects the name of every parameter that does not have a value assigned to it, which in this case are the four IP addresses. This **<xsl:for-each>** programming instruction uses recursion through the macro and selects each IP address in turn. The **color** and **group-value** parameters each have a value assigned (**blue** and **0**, respectively), so this line does not apply to them.

```
24     <xsl:for-each select="data[not(value)]/name">
```

Line 25 adds the **label-switched-path** statement in the configuration.

```
25                                <label-switched-path>
```

Lines 26 through 28 assign the **label-switched-path** a name that concatenates the value of the **color** variable, the text **-lsp-**, and the current IP address currently selected by Line 24 (represented by the **"."**).

```
26                                <name>
27                                <xsl:value-of select="concat($color,
28                                '-lsp-', .)"/>
28                                </name>
```

Line 29 adds the **to** statement to the configuration and sets its value to the IP address currently selected by Line 24.

```
29                                <to><xsl:value-of select="."/></to>
```

Lines 30 through 34 add the **admin-group include-any** statement to the configuration and set its value to the value of the **color** variable.

```
30                                <admin-group>
31                                <include-any>
32                                <xsl:value-of select="$color"/>
33                                </include-any>
34                                </admin-group>
```

The resulting configuration statements (for one pass) are as follows:

```
label-switched-path blue-lsp-10.1.1.1 {
  to 10.1.1.1;
  admin-group include-any blue;
}
```

Lines 35 through 42 are closing tags.

```
35                                </label-switched-path>
36                                </xsl:for-each>
37                                </mpls>
38                                </protocols>
39                                </transient-change>
40                                </xsl:for-each>
41                                </xsl:template>
42                                </xsl:stylesheet>
```

**SLAX Syntax** The equivalent SLAX script is:

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $mpls = protocols/mpls;
  for-each ($mpls/apply-macro[data/name = 'color']) {
    var $color = data[name = 'color']/value;
    var $group-value = data[name='group-value']/value;
    <transient-change> {
      <protocols> {
        <mpls> {
          <admin-groups> {
            <name> $color;
```

```

    <group-value> $group-value;
  }
  for-each (data[not(value)]/name) {
    <label-switched-path> {
      <name> $color _ '-lsp-' _ .;
      <to> .;
      <admin-group> {
        <include-any> $color;
      }
    }
  }
}
}
}
}
}
}

```

**Python Syntax** The following Python syntax produces the identical configuration changes:

```

from junos import Junos_Configuration
import jcs

def main():
    # Get configuration root object
    root = Junos_Configuration

    for element in root.xpath("./protocols/mpls/ \
        apply-macro[data/name='color']"):
        color = element.find("data[name='color']/value").text
        group_value = element.find("data[name='group-value']/value").text

        lsp_config = ""
        for element2 in element.xpath("data[not(value)]/name"):
            lsp_config = lsp_config + """
            <label-switched-path>
                <name>{0}-lsp-{1}</name>
                <to>{1}</to>
                <admin-group>
                    <include-any>{0}</include-any>
                </admin-group>
            </label-switched-path>
            """.format(color, element2.text)

        change_xml = """
        <protocols>
            <mpls>
                <admin-groups>
                    <name>{0}</name>
                    <group-value>{1}</group-value>
                </admin-groups>
                {2}
            </mpls>
        </protocols>
        """.format(color, group_value, lsp_config).strip()

        jcs.emit_change(change_xml, "transient-change", "xml")

if __name__ == '__main__':
    main()

```

For more information about this example, see [“Example: Configuring Administrative Groups for LSPs” on page 350](#).

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file `lsp-admin.xml`, `lsp-admin.slax`, or `lsp-admin.py` as appropriate, and copy it to the `/var/db/scripts/commit/` directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. If the script is written in Python, configure the `language python` statement.

[edit]

user@host# `set system scripts language python`

3. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard. If you are using the SLAX or Python version of the script, update the filename at the `[edit system scripts commit file]` hierarchy level.

```
system {
  scripts {
    commit {
      allow-transients;
      file lsp-admin.xml;
    }
  }
}
protocols {
  mpls {
    apply-macro blue-type-lsp {
      10.1.1.1;
      10.2.2.2;
      10.3.3.3;
      10.4.4.4;
      color blue;
      group-value 0;
    }
  }
}
```

4. In configuration mode, issue the `load merge terminal` command to merge the stanzas into your device configuration.

[edit]

user@host# `load merge terminal`

[Type ^D at a new line to end input]

... *Paste the contents of the clipboard here* ...

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
5. Issue the **commit** command to commit the configuration.  
  
user@host# commit

## Verification

---

### Verifying Script Execution

---

**Purpose** Verify that the script behaves as expected.

**Action** To display the configuration statements created by the script, issue the **show protocols mpls | display commit-scripts** command.

```
[edit]
user@host# show protocols mpls | display commit-scripts
apply-macro blue-type-lsp {
  10.1.1.1;
  10.2.2.2;
  10.3.3.3;
  10.4.4.4;
  color blue;
  group-value 0;
}
admin-groups {
  blue 0;
}
label-switched-path blue-lsp-10.1.1.1 {
  to 10.1.1.1;
  admin-group include-any blue;
}
label-switched-path blue-lsp-10.2.2.2 {
  to 10.2.2.2;
  admin-group include-any blue;
}
label-switched-path blue-lsp-10.3.3.3 {
  to 10.3.3.3;
  admin-group include-any blue;
}
label-switched-path blue-lsp-10.4.4.4 {
  to 10.4.4.4;
  admin-group include-any blue;
}
```

- Related Documentation**
- [Overview of Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)
  - [Creating Custom Configuration Syntax with Commit Script Macros on page 313](#)



- [Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements on page 320](#)



# Commit Script Examples

- [Example: Adding a Final then accept Term to a Firewall on page 331](#)
- [Example: Adding T1 Interfaces to a RIP Group on page 336](#)
- [Example: Assigning a Classifier on page 339](#)
- [Example: Automatically Configuring Logical Interfaces and IP Addresses on page 343](#)
- [Example: Configuring Administrative Groups for LSPs on page 350](#)
- [Example: Configuring a Default Encapsulation Type on page 354](#)
- [Example: Configuring Dual Routing Engines on page 357](#)
- [Example: Configuring an Interior Gateway Protocol on an Interface on page 362](#)
- [Example: Controlling IS-IS and MPLS Interfaces on page 366](#)
- [Example: Controlling LDP Configuration on page 370](#)
- [Example: Creating a Complex Configuration Based on a Simple Interface Configuration on page 375](#)
- [Example: Imposing a Minimum MTU Setting on page 382](#)
- [Example: Limiting the Number of ATM Virtual Circuits on page 384](#)
- [Example: Limiting the Number of E1 Interfaces on page 388](#)
- [Example: Loading a Base Configuration on page 397](#)
- [Example: Prepending a Global Policy on page 411](#)
- [Example: Preventing Import of the Full Routing Table on page 416](#)
- [Example: Requiring Internal Clocking on T1 Interfaces on page 419](#)
- [Example: Requiring and Restricting Configuration Statements on page 422](#)

---

## Example: Adding a Final then accept Term to a Firewall

This commit script example adds a **then accept** statement to any firewall filter that does not already end with an explicit **then accept** statement.

- [Requirements on page 332](#)
- [Overview and Commit Script on page 332](#)
- [Configuration on page 334](#)
- [Verification on page 335](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

Each firewall filter in Junos OS has an implicit discard action at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {  
  then discard;  
}
```

As a result, if a packet matches none of the terms in the filter, it is discarded. In some cases, you might want to override the default by adding a last term to accept all packets that do not match a firewall filter's series of match conditions. In this example, the commit script adds a final **then accept** statement to any firewall filter that does not already end with an explicit **then accept** statement.

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>  
<xsl:stylesheet version="1.0"  
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"  
  xmlns:junos="http://xml.juniper.net/junos/*/junos"  
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"  
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">  
  <xsl:import href="../import/junos.xml"/>  
  
  <xsl:template match="configuration">  
    <xsl:apply-templates select="firewall/filter | firewall/family/inet  
      | firewall/family/inet6" mode="filter"/>  
  </xsl:template>  
  <xsl:template match="filter" mode="filter">  
    <xsl:param name="last" select="term[position() = last()]" />  
    <xsl:comment>  
      <xsl:text>Found </xsl:text>  
      <xsl:value-of select="name" />  
      <xsl:text>; last </xsl:text>  
      <xsl:value-of select="$last/name" />  
    </xsl:comment>  
    <xsl:if test="$last and ($last/from or $last/to or not($last/then/accept))">  
      <xnm:warning>  
        <xsl:call-template name="jcs:edit-path" />  
        <message>  
          <xsl:text>filter is missing final 'then accept' rule</xsl:text>  
        </message>  
      </xnm:warning>  
      <xsl:call-template name="jcs:emit-change">  
        <xsl:with-param name="content">  
          <term>  
            <name>very-last</name>  
            <junos:comment>  
              <xsl:text>This term was added by a commit script</xsl:text>  
            </junos:comment>
```

```

        <then>
            <accept/>
        </then>
    </term>
</xsl:with-param>
</xsl:call-template>
</xsl:if>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    apply-templates firewall/filter | firewall/family/inet |
firewall/family/inet6 {
    mode "filter";
    }
}
match filter {
    mode "filter";
    param $last = term[position() = last()];
    <xsl:comment> {
        expr "Found ";
        expr name;
        expr "; last ";
        expr $last/name;
    }
    if ($last and ($last/from or $last/to or not($last/then/accept))) {
        <xnm:warning> {
            call jcs:edit-path();
            <message> "filter is missing final 'then accept' rule";
        }
        call jcs:emit-change() {
            with $content = {
                <term> {
                    <name> "very-last";
                    <junos:comment> "This term was added by a commit script";
                    <then> {
                        <accept>;
                    }
                }
            }
        }
    }
}
}
}
}
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **add-accept.xml** or **add-accept.slax** as appropriate, and copy it to the `/var/db/scripts/commit/` directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **add-accept.slax**.

```
system {
  scripts {
    commit {
      file add-accept.xml;
    }
  }
}
firewall {
  policer sgt-friday {
    if-exceeding {
      bandwidth-percent 10;
      burst-size-limit 250k;
    }
    then discard;
  }
}
family inet {
  filter test {
    term one {
      from {
        interface t1-0/0/0;
      }
      then {
        count ten-network;
        discard;
      }
    }
    term two {
      from {
        forwarding-class assured-forwarding;
      }
      then discard;
    }
  }
}
}
interfaces {
  t1-0/0/0 {
    unit 0 {
      family inet {
        policer output sgt-friday;
        filter input test;
      }
    }
  }
}
```

```
}
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The script requires that all firewall filters end with an explicit **then accept** statement. The sample configuration stanzas include the **test** filter with two terms but do not include an explicit **then accept** statement. When you issue the **commit** command, the script adds the missing **then accept** statement and commits the configuration. When you issue the **commit** command, the following output appears:

```
[edit]
user@host# commit
[edit firewall family inet filter test]
warning: filter is missing final 'then accept' rule
commit complete
```

In configuration mode, issue the **show firewall** command to review the modified configuration. The following output appears:

```
[edit]
user@host# show firewall
policer sgt-friday {
  if-exceeding {
    bandwidth-percent 10;
    burst-size-limit 250k;
  }
  then discard;
}
family inet {
  filter test {
    term one {
      from {
        interface t1-0/0/0;
```

```
    }
    then {
        count ten-network;
        discard;
    }
}
term two {
    from {
        forwarding-class assured-forwarding;
    }
    then {
        discard;
    }
}
term very-last {
    then accept; /* This term was added by a commit script */
}
}
```

---

## Example: Adding T1 Interfaces to a RIP Group

This example shows how to use commit scripts to decrease the amount of manual configuration, specifically how to add every T1 interface configured at the **[edit interfaces]** hierarchy level to the **[edit protocols rip group test]** hierarchy level.

- [Requirements on page 336](#)
- [Overview and Commit Script on page 336](#)
- [Configuration on page 338](#)
- [Verification on page 339](#)

### Requirements

This example uses a device running Junos OS with T1 interfaces.

### Overview and Commit Script

If you want to enable RIP on an interface, you must make changes at both the **[edit interfaces]** and **[edit protocols rip]** hierarchy levels. This example shows how to use commit scripts to add every T1 interface configured at the **[edit interfaces]** hierarchy level to the **[edit protocols rip group test]** hierarchy level. This example includes no error, warning, or system log messages. The changes to the configuration are made silently.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../../../import/junos.xsl"/>

  <xsl:template match="configuration">
```



```

<xsl:variable name="all-t1"
  select="interfaces/interface[starts-with(name, 't1-')]">
<xsl:if test="$all-t1">
  <change>
    <protocols>
      <rip>
        <group>
          <name>test</name>
          <xsl:for-each select="$all-t1">
            <xsl:variable name="ifname" select="concat(name,
'.0')"/>
            <neighbor>
              <name><xsl:value-of select="$ifname"/></name>
            </neighbor>
          </xsl:for-each>
        </group>
      </rip>
    </protocols>
  </change>
</xsl:if>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $all-t1 = interfaces/interface[starts-with(name, 't1-')];
  if ($all-t1) {
    <change> {
      <protocols> {
        <rip> {
          <group> {
            <name> "test";
            for-each ($all-t1) {
              var $ifname = name _ '.0';
              <neighbor> {
                <name> $ifname;
              }
            }
          }
        }
      }
    }
  }
}

```

## Configuration

**Step-by-Step Procedure** To download, enable, and test the script:

1. Copy the script into a text file, name the file **rip-t1.xml** or **rip-t1.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **rip-t1.slax**.

```
system {
  scripts {
    commit {
      file rip-t1.xml;
    }
  }
}
interfaces {
  t1-0/0/0 {
    unit 0 {
      family iso;
    }
  }
  t1-0/0/1 {
    unit 0 {
      family iso;
    }
  }
  t1-0/0/2 {
    unit 0 {
      family iso;
    }
  }
  t1-0/0/3 {
    unit 0 {
      family iso;
    }
  }
  t1-0/1/0 {
    unit 0 {
      family iso;
    }
  }
  t1-0/1/1 {
    unit 0 {
      family iso;
    }
  }
  t1-0/1/2 {
    unit 0 {
      family iso;
    }
  }
}
```

```

t1-0/1/3 {
    unit 0 {
        family iso;
    }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying the Configuration

**Purpose** Verify that the script behaves as expected.

**Action** Issue the **show protocols rip group test** command. All T1 interfaces should now appear under the **[edit protocols rip group test]** hierarchy level.

```

[edit]
user@host# show protocols rip group test
neighbor t1-0/0/0.0;
neighbor t1-0/0/1.0;
neighbor t1-0/0/2.0;
neighbor t1-0/0/3.0;
neighbor t1-0/1/0.0;
neighbor t1-0/1/1.0;
neighbor t1-0/1/2.0;
neighbor t1-0/1/3.0;

```

## Example: Assigning a Classifier

For each interface configured with the IPv4 protocol family, this commit script automatically assigns a specified classifier, which associates incoming packets with a forwarding class and loss priority as well as assigns packets to an output queue.

- [Requirements on page 340](#)
- [Overview and Commit Script on page 340](#)

- [Configuration on page 341](#)
- [Verification on page 342](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

In the Junos OS class of service (CoS), classifiers allow you to associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues. After you configure a classifier, you must assign it to an input interface.

For each interface configured with the IPv4 protocol family, this script automatically assigns a specified classifier called **fc-q3**. The **fc-q3** classifier must be configured at the **[edit class-of-service]** hierarchy level.

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:variable name="cos-all" select="class-of-service"/>
    <xsl:for-each
      select="interfaces/interface[contains(name,
        '/' )]/unit[family/inet]">
      <xsl:variable name="ifname" select="../name"/>
      <xsl:variable name="unit" select="name"/>
      <xsl:variable name="cos"
        select="$cos-all/interfaces[name = $ifname]"/>
      <xsl:if test="not($cos/unit[name = $unit])">
        <xsl:call-template name="jcs:emit-change">
          <xsl:with-param name="message">
            <xsl:text>Adding CoS forwarding class for </xsl:text>
            <xsl:value-of select="concat($ifname, '.', $unit)"/>
          </xsl:with-param>
          <xsl:with-param name="dot" select="$cos-all"/>
          <xsl:with-param name="content">
            <interfaces>
              <name><xsl:value-of select="$ifname"/></name>
              <unit>
                <name><xsl:value-of select="$unit"/></name>
                <forwarding-class>fc-q3</forwarding-class>
              </unit>
            </interfaces>
          </xsl:with-param>
        </xsl:call-template>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $cos-all = class-of-service;
  for-each (interfaces/interface[contains(name, '/')]/unit[family/inet]) {
    var $ifname = ../name;
    var $unit = name;
    var $cos = $cos-all/interfaces[name = $ifname];
    if (not($cos/unit[name = $unit])) {
      call jcs:emit-change($dot = $cos-all) {
        with $message = {
          expr "Adding CoS forwarding class for ";
          expr $ifname _ '.' _ $unit;
        }
        with $content = {
          <interfaces> {
            <name> $ifname;
            <unit> {
              <name> $unit;
              <forwarding-class> "fc-q3";
            }
          }
        }
      }
    }
  }
}

```

**Configuration****Step-by-Step  
Procedure**

To download, enable, and test the script:

1. Copy the script into a text file, name the file **classifier.xsl** or **classifier.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **classifier.slax**.

```

system {
  scripts {
    commit {
      file classifier.xsl;
    }
  }
}
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet {
        address 10.168.16.2/24;
      }
    }
  }
}

```

```
    }  
  }  
  class-of-service {  
    forwarding-classes {  
      queue 3 fc-q3;  
    }  
    classifiers {  
      inet-precedence fc-q3 {  
        forwarding-class fc-q3 {  
          loss-priority low code-points 010;  
        }  
      }  
    }  
  }  
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]  
user@host# load merge terminal  
[Type ^D at a new line to end input]  
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

---

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. In the test configuration stanzas, the fe-0/0/0.0 interface is configured with the **family inet** statement. Because the interface is configured with the IPv4 protocol family, the script automatically assigns the **fc-q3** classifier to the interface, which is indicated in the **commit** command output.

```
[edit]  
user@host# commit  
[edit interfaces interface fe-0/0/0 unit 0]  
warning: Adding CoS forwarding class for fe-0/0/0.0  
commit complete
```

View the configuration to verify that the script-generated changes are present. Issue the **show class-of-service** configuration mode command. The output shows that the fe-0/0/0.0 interface has been assigned the **fc-q3** classifier:

```
[edit]  
user@host# show class-of-service  
classifiers {
```

```

    inet-precedence fc-q3 {
        forwarding-class fc-q3 {
            loss-priority low code-points 010;
        }
    }
}
forwarding-classes {
    queue 3 fc-q3;
}
interfaces {
    fe-0/0/0 {
        unit 0 {
            forwarding-class fc-q3; # Added by commit script
        }
    }
}

```

## Example: Automatically Configuring Logical Interfaces and IP Addresses

Every interface you configure requires at least one logical unit and one IP address. Asynchronous Transfer Mode (ATM) interfaces also require a virtual circuit identifier (VCI) for each logical interface. If you need to configure multiple logical units on an interface, you can use a commit script and macro to complete the task quickly and with no errors.

- [Requirements on page 343](#)
- [Overview and Commit Script on page 343](#)
- [Configuration on page 348](#)
- [Verification on page 349](#)

### Requirements

This example uses a device running Junos OS with physical ATM interfaces.

### Overview and Commit Script

The following commit script expands an **apply-macro** statement that provides the name of a physical ATM interface and a set of parameters that specify how to configure a number of logical units on the interface. The units and VCI numbers are numbered sequentially from the **unit** variable to the **max** variable and are given IP addresses starting at the **address** variable. To loop through the logical units, Extensible Stylesheet Language Transformations (XSLT) uses recursion, which is implemented in the **<emit-interface>** template. Calculation of the next address is performed in the **<next-address>** template.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
    xmlns:junos="http://xml.juniper.net/junos/*/junos"
    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
    <xsl:import href="../import/junos.xsl"/>

```

```

<xsl:template match="configuration">
  <xsl:for-each select="interfaces/apply-macro">
    <xsl:variable name="device" select="name"/>
    <xsl:variable name="address" select="data[name='address']/value"/>
    <xsl:variable name="max" select="data[name='max']/value"/>
    <xsl:variable name="unit" select="data[name='unit']/value"/>
    <xsl:variable name="real-max">
      <xsl:choose>
        <xsl:when test="string-length($max) > 0">
          <xsl:value-of select="$max"/>
        </xsl:when>
        <xsl:otherwise>0</xsl:otherwise>
      </xsl:choose>
    </xsl:variable>
    <xsl:variable name="real-unit">
      <xsl:choose>
        <xsl:when test="string-length($unit) > 0">
          <xsl:value-of select="$unit"/>
        </xsl:when>
        <xsl:when test="contains($device, '.')">
          <xsl:value-of select="substring-after($device, '.')" />
        </xsl:when>
        <xsl:otherwise>0</xsl:otherwise>
      </xsl:choose>
    </xsl:variable>
    <xsl:variable name="real-device">
      <xsl:choose>
        <xsl:when test="contains($device, '.')">
          <xsl:value-of select="substring-before($device, '.')" />
        </xsl:when>
        <xsl:otherwise><xsl:value-of
select="$device"/></xsl:otherwise>
      </xsl:choose>
    </xsl:variable>
    <transient-change>
      <interfaces>
        <interface>
          <name><xsl:value-of select="$real-device"/></name>
          <xsl:call-template name="emit-interface">
            <xsl:with-param name="address" select="$address"/>
            <xsl:with-param name="unit" select="$real-unit"/>
            <xsl:with-param name="max" select="$real-max"/>
          </xsl:call-template>
        </interface>
      </interfaces>
    </transient-change>
  </xsl:for-each>
</xsl:template>
<xsl:template name="emit-interface">
  <xsl:param name="$max"/>
  <xsl:param name="$unit"/>
  <xsl:param name="$address"/>
  <unit>
    <name><xsl:value-of select="$unit"/></name>
    <vci><xsl:value-of select="$unit"/></vci>
    <family>
      <inet>
        <address><xsl:value-of select="$address"/></address>
      </inet>
    </family>
  </unit>
</xsl:template>

```



```

</unit>
<xsl:if test="$max > $unit">
  <xsl:call-template name="emit-interface">
    <xsl:with-param name="address">
      <xsl:call-template name="next-address">
        <xsl:with-param name="address" select="$address"/>
      </xsl:call-template>
    </xsl:with-param>
    <xsl:with-param name="unit" select="$unit + 1"/>
    <xsl:with-param name="max" select="$max"/>
  </xsl:call-template>
</xsl:if>
</xsl:template>
<xsl:template name="next-address">
  <xsl:param name="address"/>
  <xsl:variable name="arg-prefix" select="substring-after($address,
'/' )"/>
  <xsl:variable name="arg-addr" select="substring-before($address,
'/' )"/>
  <xsl:variable name="addr">
    <xsl:choose>
      <xsl:when test="string-length($arg-addr) > 0">
        <xsl:value-of select="$arg-addr"/>
      </xsl:when>
      <xsl:otherwise>
        <xsl:value-of select="$address"/>
      </xsl:otherwise>
    </xsl:choose>
  </xsl:variable>
  <xsl:variable name="prefix">
    <xsl:choose>
      <xsl:when test="string-length($arg-prefix) > 0">
        <xsl:value-of select="$arg-prefix"/>
      </xsl:when>
      <xsl:otherwise>32</xsl:otherwise>
    </xsl:choose>
  </xsl:variable>
  <xsl:variable name="a1" select="substring-before($addr, '.')"/>
  <xsl:variable name="a234" select="substring-after($addr, '.')"/>
  <xsl:variable name="a2" select="substring-before($a234, '.')"/>
  <xsl:variable name="a34" select="substring-after($a234, '.')"/>
  <xsl:variable name="a3" select="substring-before($a34, '.')"/>
  <xsl:variable name="a4" select="substring-after($a34, '.')"/>
  <xsl:variable name="r3">
    <xsl:choose>
      <xsl:when test="$a4 < 255">
        <xsl:value-of select="$a3"/>
      </xsl:when>
      <xsl:otherwise>
        <xsl:value-of select="$a3 + 1"/>
      </xsl:otherwise>
    </xsl:choose>
  </xsl:variable>
  <xsl:variable name="r4">
    <xsl:choose>
      <xsl:when test="$a4 < 255">
        <xsl:value-of select="$a4 + 1"/>
      </xsl:when>
      <xsl:otherwise>
        <xsl:value-of select="0"/>
      </xsl:otherwise>
    </xsl:choose>
  </xsl:variable>

```

```

        </xsl:choose>
      </xsl:variable>
      <xsl:value-of select="$a1"/>
      <xsl:text>.</xsl:text>
      <xsl:value-of select="$a2"/>
      <xsl:text>.</xsl:text>
      <xsl:value-of select="$r3"/>
      <xsl:text>.</xsl:text>
      <xsl:value-of select="$r4"/>
      <xsl:text></xsl:text>
      <xsl:value-of select="$prefix"/>
    </xsl:template>
  </xsl:stylesheet>

```

## SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (interfaces/apply-macro) {
    var $device = name;
    var $address = data[name='address']/value;
    var $max = data[name='max']/value;
    var $unit = data[name='unit']/value;
    var $real-max = {
      if (string-length($max) > 0) {
        expr $max;
      } else {
        expr "0";
      }
    }
    var $real-unit = {
      if (string-length($unit) > 0) {
        expr $unit;
      } else if (contains($device, '.')) {
        expr substring-after($device, '.');
      } else {
        expr "0";
      }
    }
    var $real-device = {
      if (contains($device, '.')) {
        expr substring-before($device, '.');
      } else {
        expr $device;
      }
    }
    <transient-change> {
      <interfaces> {
        <interface> {
          <name> $real-device;
          call emit-interface($address, $unit = $real-unit, $max =
$real-max);
        }
      }
    }
  }
}
emit-interface ($max, $unit, $address) {

```

```

<unit> {
  <name> $unit;
  <vci> $unit;
  <family> {
    <inet> {
      <address> $address;
    }
  }
}
if ($max > $unit) {
  call emit-interface($unit = $unit + 1, $max) {
    with $address = {
      call next-address($address);
    }
  }
}
}
next-address ($address) {
  var $arg-prefix = substring-after($address, '/');
  var $arg-addr = substring-before($address, '/');
  var $addr = {
    if (string-length($arg-addr) > 0) {
      expr $arg-addr;
    } else {
      expr $address;
    }
  }
  var $prefix = {
    if (string-length($arg-prefix) > 0) {
      expr $arg-prefix;
    } else {
      expr "32";
    }
  }
  var $a1 = substring-before($addr, '.');
  var $a234 = substring-after($addr, '.');
  var $a2 = substring-before($a234, '.');
  var $a34 = substring-after($a234, '.');
  var $a3 = substring-before($a34, '.');
  var $a4 = substring-after($a34, '.');
  var $r3 = {
    if ($a4 < 255) {
      expr $a3;
    } else {
      expr $a3 + 1;
    }
  }
  var $r4 = {
    if ($a4 < 255) {
      expr $a4 + 1;
    } else {
      expr 0;
    }
  }
  expr $a1;
  expr ".";
  expr $a2;
  expr ".";
  expr $r3;
  expr ".";
  expr $r4;
}

```

```
    expr "/";  
    expr $prefix;  
}
```

## Configuration

**Step-by-Step Procedure** To download, enable, and run the script:

1. Copy the script into a text file, name the file **atm-logical.xml** or **atm-logical.slax** as appropriate, and download it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **atm-logical.slax**.

```
system {  
  scripts {  
    commit {  
      allow-transients;  
      file atm-logical.xml;  
    }  
  }  
}  
interfaces {  
  apply-macro at-1/2/3 {  
    address 10.12.13.14/20;  
    max 200;  
    unit 32;  
  }  
  at-1/2/3 {  
    atm-options {  
      pic-type atm2;  
      vpi 0;  
    }  
  }  
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]  
user@host# load merge terminal  
[Type ^D at a new line to end input]  
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

**Purpose** Verify that the correct changes are integrated into the configuration.

**Action** Before you commit the configuration, you can verify that the commit script will produce the correct results by issuing the **show interfaces at-1/2/3 | display commit-scripts** configuration mode command. After you commit the configuration, you can review the active configuration by issuing the **show configuration interfaces at-1/2/3** operational mode command. The following output appears:

```
atm-options {
  pic-type atm2;
  vpi 0;
}
unit 32 {
  vci 32;
  family inet {
    address 10.12.13.14/20;
  }
}
unit 33 {
  vci 33;
  family inet {
    address 10.12.13.15/20;
  }
}
unit 34 {
  vci 34;
  family inet {
    address 10.12.13.16/20;
  }
}
unit 35 {
  vci 35;
  family inet {
    address 10.12.13.17/20;
  }
}
... Logical units 36 through 199 are omitted for brevity ...
unit 200 {
  vci 200 ;
  family inet {
    address 10.12.13.182/20;
  }
}
```

**Meaning** The **| display commit-scripts** option displays the configuration data after all commit scripts have been applied. The output includes both persistent and transient changes. If the appropriate **unit** and **vci** are configured on each ATM interface, the commit script executes successfully during a commit operation. After you commit the configuration, you can review the active configuration by issuing the **show configuration interfaces at-1/2/3** operational mode command.

## Example: Configuring Administrative Groups for LSPs

Administrative groups, also known as link coloring or resource classes, are manually assigned attributes that describe the color of links. Links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based label-switched path (LSP) setups.

This commit script example searches for **apply-macro** statements with the **color** parameter included at the **[edit protocols mpls]** hierarchy level. For each **apply-macro** statement, the script uses the data provided to generate a transient change and expand the macro into a standard Junos OS administrative group for LSPs.

- [Requirements on page 350](#)
- [Overview and Commit Script on page 350](#)
- [Configuration on page 352](#)
- [Verification on page 353](#)

### Requirements

This example uses a device running Junos OS.

### Overview and Commit Script

In this example, the Junos OS management process (mgd) inspects the configuration, looking for **apply-macro** statements. For each **apply-macro** statement with the **color** parameter included at the **[edit protocols mpls]** hierarchy level, the script generates a transient change, using the data provided within the **apply-macro** statement to expand the macro into a standard Junos OS administrative group for LSPs.

For this example to work, an **apply-macro** statement must be included at the **[edit protocols mpls]** hierarchy level with a set of addresses, a **color** parameter, and a **group-value** parameter. The commit script converts each address to an LSP configuration and converts the **color** parameter into an administrative group.

For a line-by-line explanation of this script, see [“Example: Creating Custom Configuration Syntax with Commit Script Macros” on page 322](#).

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:variable name="mpls" select="protocols/mpls"/>
    <xsl:for-each select="$mpls/apply-macro[data/name = 'color']">
      <xsl:variable name="color" select="data[name = 'color']/value"/>
      <xsl:for-each select="$mpls/apply-macro[data/name = 'group-value']">
        <xsl:variable name="group-value" select="data[name =
```

```

        'group-value']/value"/>
    <transient-change>
    <protocols>
    <mpls>
    <admin-groups>
    <name>
    <xsl:value-of select="$color"/>
    </name>
    <group-value>
    <xsl:value-of select="$group-value"/>
    </group-value>
    </admin-groups>
    <xsl:for-each select="data[not(value)]/name">
    <label-switched-path>
    <name>
    <xsl:value-of select="concat($color,
'-lsp-', .)"/>
    </name>
    <to><xsl:value-of select="."/></to>
    <admin-group>
    <include-any>
    <xsl:value-of select="$color"/>
    </include-any>
    </admin-group>
    </label-switched-path>
    </xsl:for-each>
    </mpls>
    </protocols>
    </transient-change>
    </xsl:for-each>
    </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $mpls = protocols/mpls;
  for-each ($mpls/apply-macro[data/name = 'color']) {
    var $color = data[name = 'color']/value;
    for-each ($mpls/apply-macro[data/name = 'group-value']) {
      var $group-value = data[name = 'group-value']/value;
      <transient-change> {
        <protocols> {
          <mpls> {
            <admin-groups> {
              <name> $color;
              <group-value> $group-value;
            }
            for-each (data[not(value)]/name) {
              <label-switched-path> {
                <name> $color _ '-lsp-' _ .;
                <to> .;
                <admin-group> {
                  <include-any> $color;
                }
              }
            }
          }
        }
      }
    }
  }
}

```

```

    }
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **lsp-admin.xml** or **lsp-admin.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **lsp-admin.slax**.

```

system {
  scripts {
    commit {
      allow-transients;
      file lsp-admin.xml;
    }
  }
}
protocols {
  mpls {
    apply-macro blue-type-lsp {
      10.1.1.1;
      10.2.2.2;
      10.3.3.3;
      10.4.4.4;
      color blue;
      group-value 0;
    }
  }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.



4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

<b>Purpose</b>	Verify that the script behaves as expected.
<b>Action</b>	Issue the <b>show protocols mpls   display commit-scripts</b> configuration mode command and review the output. Adding the <b>  display commit-scripts</b> option allows you to see the configuration statements that are generated by transient changes.
<b>With Script-Generated Changes</b>	<p>When you issue the <b>show protocols mpls   display commit-scripts</b> configuration mode command, the following output appears:</p> <pre>[edit] user@host# show protocols mpls   display commit-scripts apply-macro blue-type-lsp {   10.1.1.1;   10.2.2.2;   10.3.3.3;   10.4.4.4;   color blue;   group-value 0; } admin-groups {   blue 0; } label-switched-path blue-lsp-10.1.1.1 {   to 10.1.1.1;   admin-group include-any blue; } label-switched-path blue-lsp-10.2.2.2 {   to 10.2.2.2;   admin-group include-any blue; } label-switched-path blue-lsp-10.3.3.3 {   to 10.3.3.3;   admin-group include-any blue; } label-switched-path blue-lsp-10.4.4.4 {   to 10.4.4.4;   admin-group include-any blue; }</pre>
<b>Without Script-Generated Changes</b>	<p>The output of the <b>show protocols mpls   display commit-scripts no-transients</b> configuration mode command excludes the <b>label-switched-path</b> statements:</p> <pre>[edit] user@host# show protocols mpls   display commit-scripts no-transients apply-macro blue-type-lsp {   10.1.1.1;   10.2.2.2;</pre>

```
10.3.3.3;  
10.4.4.4;  
color blue;  
group-value 0;  
}
```

When you issue the **show protocols mpls** command without the piped **display commit-scripts no-transients** command, you see the same output because this script does not generate any persistent changes:

```
[edit]  
user@host# show protocols mpls  
apply-macro blue-type-lsp {  
  10.1.1.1;  
  10.2.2.2;  
  10.3.3.3;  
  10.4.4.4;  
  color blue;  
  group-value 0;  
}
```

---

## Example: Configuring a Default Encapsulation Type

This commit script example configures default Cisco HDLC encapsulation on SONET/SDH interfaces not configured as aggregate interfaces.

- [Requirements on page 354](#)
- [Overview and Commit Script on page 354](#)
- [Configuration on page 355](#)
- [Verification on page 356](#)

### Requirements

This example uses a device running Junos OS with SONET/SDH interfaces.

### Overview and Commit Script

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You do not need to configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

This example configures default Cisco HDLC encapsulation on SONET/SDH interfaces not configured as aggregate interfaces. The **tag** variable is passed to the **jcs:emit-change** template as **transient-change**, so this change is not copied to the candidate configuration.

Simply including configuration groups in the configuration does not enable you to test whether the **aggregate** statement is included for an interface at the **[edit interfaces interface-name sonet-options]** hierarchy level. A commit script can perform this test and set the encapsulation only on nonaggregated interfaces. The script written to perform this test has the following syntax:

**XSLT Syntax**

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name, 'so-')
      and not(sonet-options/aggregate)]">
      <xsl:call-template name="jcs:emit-change">
        <xsl:with-param name="tag" select="'transient-change'"/>
        <xsl:with-param name="content">
          <encapsulation>cisco-hdlc</encapsulation>
        </xsl:with-param>
      </xsl:call-template>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (interfaces/interface[starts-with(name, 'so-') and
    not(sonet-options/aggregate)]) {
    call jcs:emit-change($tag = 'transient-change') {
      with $content = {
        <encapsulation> "cisco-hdlc";
      }
    }
  }
}
```

**Configuration****Step-by-Step  
Procedure**

To download, enable, and test the script:

1. Copy the script into a text file, name the file **so-encap.xsl** or **so-encap.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **so-encap.slax**.

```
system {
  scripts {
    commit {
      allow-transients;
      file so-encap.xsl;
    }
  }
}
```

```
interfaces {
  so-1/2/2 {
    sonet-options {
      aggregate as0;
    }
  }
  so-1/2/3 {
    unit 0 {
      family inet {
        address 10.0.0.3/32;
      }
    }
  }
  so-1/2/4 {
    unit 0 {
      family inet {
        address 10.0.0.4/32;
      }
    }
  }
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

---

**Purpose** Verify that the script behaves as expected.

**Action** When you issue the **commit** command, the commit script tests for SONET/SDH interfaces that are not configured as aggregate interfaces and sets the default encapsulation type on the nonaggregated interfaces to Cisco HDLC encapsulation. This is implemented as a **transient-change**. Even though the transient changes are in effect, they are not, by default, displayed in the normal output of the **show interfaces** command.

```
[edit]
user@host# show interfaces
so-1/2/2 {
```

```

sonet-options {
    aggregate as0;
}
}
so-1/2/3 {
    unit 0 {
        family inet {
            address 10.0.0.3/32;
        }
    }
}
so-1/2/4 {
    unit 0 {
        family inet {
            address 10.0.0.4/32;
        }
    }
}
}

```

To view the configuration with the transient changes, issue the **show interfaces | display commit-scripts** command:

```

[edit]
user@host# show interfaces | display commit-scripts
so-1/2/2 {
    sonet-options { # The presence of these statements prevents the
        aggregate as0; # transient change from affecting this interface.
    }
}
so-1/2/3 {
    encapsulation cisco-hdlc; # Added by transient change.
    unit 0 {
        family inet {
            address 10.0.0.3/32;
        }
    }
}
so-1/2/4 {
    encapsulation cisco-hdlc; # Added by transient change.
    unit 0 {
        family inet {
            address 10.0.0.4/32;
        }
    }
}
}

```

## Example: Configuring Dual Routing Engines

If your device has redundant (also called *dual*) Routing Engines, your Junos OS configuration can be complex. This example shows how you can use commit scripts to simplify and control the configuration of dual Routing Engine platforms.

- [Requirements on page 358](#)
- [Overview and Commit Script on page 358](#)

- [Configuration on page 361](#)
- [Verification on page 362](#)

## Requirements

This example uses a device running Junos OS with dual Routing Engines.

## Overview and Commit Script

Junos OS supports two special configuration groups: **re0** and **re1**. When these groups are applied using the **apply-groups [ re0 re1 ]** statement, they take effect if the Routing Engine name matches the group name. Statements included at the **[edit groups re0]** hierarchy level are inherited only on the Routing Engine named RE0, and statements included at the **[edit groups re1]** hierarchy level are inherited only on the Routing Engine named RE1.

This example includes two commit scripts. The first script, **dual-re.xsl**, generates a warning if the **system host-name** statement, any IP version 4 (IPv4) interface address, or the fxp0 interface configuration is configured in the target configuration instead of in a configuration group.

The second script, **dual-re2.xsl**, first checks whether the hostname configuration is configured and then checks whether it is configured in a configuration group. The **otherwise** construct generates an error message if the hostname is not configured at all. The first **when** construct allows the script to do nothing if the hostname is already configured in a configuration group. The second **when** construct takes effect when the hostname is configured in the target configuration. In this case, the script generates a transient change that places the hostname configuration into the **re0** and **re1** configuration groups, copies the configured hostname into those groups, concatenates each group hostname with **-RE0** and **-RE1**, and deactivates the hostname in the target configuration so the configuration group hostnames can be inherited.

The example scripts are shown in both XSLT and SLAX syntax:

### XSLT Syntax: dual-re.xsl Script

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each select="system/host-name |
      interfaces/interface/unit/family/inet/address |
      interfaces/interface[name = 'fxp0']">
      <xsl:if test="not(@junos:group) or not(starts-with(@junos:group,
        're'))">
        <xnm:warning>
          <xsl:call-template name="jcs:edit-path">
            <xsl:with-param name="dot" select=".."/>
          </xsl:call-template>
          <xsl:call-template name="jcs:statement"/>
          <message>
            <xsl:text>statement should not be in target</xsl:text>
            <xsl:text> configuration on dual RE system</xsl:text>
          </message>
        </xnm:warning>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

```

        </xnm:warning>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>

```

**XSLT Syntax:  
dual-re2.xsl Script**

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:variable name="hn" select="system/host-name"/>
    <xsl:choose>
      <xsl:when test="$hn@junos:group"/>
      <xsl:when test="$hn">
        <transient-change>
          <groups>
            <name>re0</name>
            <system>
              <host-name>
                <xsl:value-of select="concat($hn, '-RE0')"/>
              </host-name>
            </system>
          </groups>
          <groups>
            <name>re1</name>
            <system>
              <host-name>
                <xsl:value-of select="concat($hn, '-RE1')"/>
              </host-name>
            </system>
          </groups>
          <system>
            <host-name inactive="inactive"/>
          </system>
        </transient-change>
      </xsl:when>
      <xsl:otherwise>
        <xnm:error>
          <message>Missing [system host-name]</message>
        </xnm:error>
      </xsl:otherwise>
    </xsl:choose>
  </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax:  
dual-re.xsl Script**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (system/host-name | interfaces/interface/unit/family/inet/address
    |
      interfaces/interface[name = 'fxp0']) {
    if (not(@junos:group) or not(starts-with(@junos:group, 're'))) {

```

```

        <xnm:warning> {
            call jcs:edit-path($dot = ..);
            call jcs:statement();
            <message> {
                expr "statement should not be in target";
                expr " configuration on dual RE system";
            }
        }
    }
}

```

### SLAX Syntax: dual-re2.xsl Script

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    var $hn = system/host-name;
    if ($hn@junos:group) {
    }
    else if ($hn) {
        <transient-change> {
            <groups> {
                <name> "re0";
                <system> {
                    <host-name> $hn _ '-RE0';
                }
            }
            <groups> {
                <name> "re1";
                <system> {
                    <host-name> $hn _ '-RE1';
                }
            }
            <system> {
                <host-name inactive="inactive">;
            }
        }
    }
    else {
        <xnm:error> {
            <message> "Missing [system host-name]";
        }
    }
}
}

```



## Configuration

### Step-by-Step Procedure

To download, enable, and run the scripts:

1. Copy the scripts into two text files, name the files **dual-re.xml** and **dual-re2.xml** or **dual-re.slax** and **dual-re2.slax** as appropriate, and copy them to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filenames at the **[edit system scripts commit file]** hierarchy level to **dual-re.slax** and **dual-re2.slax**.

```

groups {
  re0 {
    interfaces {
      fxp0 {
        unit 0 {
          family inet {
            address 10.0.0.1/24;
          }
        }
      }
    }
  }
}
apply-groups re0;
system {
  host-name router1;
  scripts {
    commit {
      file dual-re.xml;
      file dual-re2.xml;
    }
  }
}
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.220.1/30;
      }
    }
  }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.  
  
user@host# **commit**

## Verification

---

### Verifying the Commit Script Changes

<b>Purpose</b>	Verify that the script behaves as expected.
<b>Action</b>	Review the output of the <b>commit</b> command. After the commit operation completes, the device hostname is changed to router1-RE0.  <pre>[edit] user@host# <b>commit</b> [edit system] 'host-name router1;' warning: statement should not be in target configuration on dual RE system [edit interfaces interface fe-0/0/0 unit 0 family inet] 'address 192.168.220.1/30;' warning: statement should not be in target configuration on dual RE system commit complete</pre>

---

## Example: Configuring an Interior Gateway Protocol on an Interface

This commit script example uses a macro to automatically include an interface at the **[edit protocols]** hierarchy level and to configure the proper interior gateway protocol (IGP) on the interface.

- [Requirements on page 362](#)
- [Overview and Commit Script on page 362](#)
- [Configuration on page 365](#)
- [Verification on page 366](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

When you add a new interface to an OSPF or IS-IS domain, you must configure the interface at multiple hierarchy levels, including **[edit interfaces]** and **[edit protocols]**. This example uses a commit script and macro to automatically include the interface at the **[edit protocols]** hierarchy level and to configure the proper IGP on the interface, either OSPF or IS-IS, depending on the content of an **apply-macro** statement that you include

in the interface configuration. This macro allows you to perform more configuration tasks at a single hierarchy level.

In this example, the Junos OS management (mgd) process inspects the configuration, looking for **apply-macro** statements. For each **apply-macro ifclass** statement included at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level, the script tests whether the **role** parameter is defined as **cpe**. If so, the script checks the **igp** parameter.

If the **igp** parameter is defined as **isis**, the script includes the relevant interface name at the **[edit protocols isis interface]** hierarchy level.

If the **igp** parameter is defined as **ospf**, the script includes the relevant interface name at the **[edit protocols ospf area address interface]** hierarchy level. For OSPF, the script references the **area** parameter to determine the correct subnet address of the area.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each
      select="interfaces/interface/unit/apply-macro[name = 'ifclass']">
      <xsl:variable name="role" select="data[name='role']/value"/>
      <xsl:variable name="igp" select="data[name='igp']/value"/>
      <xsl:variable name="ifname">
        <xsl:value-of select="../../name"/>
        <xsl:text>.</xsl:text>
        <xsl:value-of select="../../name"/>
      </xsl:variable>
      <xsl:choose>
        <xsl:when test="$role = 'cpe'">
          <change>
            <xsl:choose>
              <xsl:when test="$igp = 'isis'">
                <protocols>
                  <isis>
                    <interface>
                      <name><xsl:value-of select="$ifname"/></name>
                    </interface>
                  </isis>
                </protocols>
              </xsl:when>
              <xsl:when test="$igp = 'ospf'">
                <protocols>
                  <ospf>
                    <area>
                      <name>
                        <xsl:value-of select="data[name='area']/value"/>
                      </name>
                    </area>
                  </ospf>
                </protocols>
              </xsl:when>
            </xsl:choose>
          </change>
        </xsl:when>
      </xsl:choose>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

```

        </name>
        <interface>
            <name><xsl:value-of select="$ifname"/></name>
        </interface>
    </area>
</ospf>
</protocols>
</xsl:when>
</xsl:choose>
</change>
</xsl:when>
</xsl:choose>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>

```

## SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    for-each (interfaces/interface/unit/apply-macro[name = 'ifclass']) {
        var $role = data[name='role']/value;
        var $igp = data[name='igp']/value;
        var $ifname = {
            expr ../../name;
            expr ".";
            expr ../name;
        }
        if ($role = 'cpe') {
            <change> {
                if ($igp = 'isis') {
                    <protocols> {
                        <isis> {
                            <interface> {
                                <name> $ifname;
                            }
                        }
                    }
                }
            }
        }
        else if ($igp = 'ospf') {
            <protocols> {
                <ospf> {
                    <area> {
                        <name> data[name='area']/value;
                        <interface> {
                            <name> $ifname;
                        }
                    }
                }
            }
        }
    }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **if-class.xml** or **if-class.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **if-class.slax**.

```

system {
  scripts {
    commit {
      file if-class.xml;
    }
  }
}
interfaces {
  so-1/2/3 {
    unit 0 {
      apply-macro ifclass {
        area 10.4.0.0;
        igp ospf;
        role cpe;
      }
    }
  }
  t3-0/0/0 {
    unit 0 {
      apply-macro ifclass {
        igp isis;
        role cpe;
      }
    }
  }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying the Configuration

---

**Purpose** Verify that the script behaves as expected.

**Action** View the configuration to verify that the manual changes and the script-generated changes are present.

When you issue the **show interfaces** configuration mode command, the changes added by the sample configuration stanzas should be present in the configuration.

```
[edit]
user@host# show interfaces
t3-0/0/0 {
  unit 0 {
    apply-macro ifclass {
      igp isis;
      role cpe;
    }
  }
}
so-1/2/3 {
  unit 0 {
    apply-macro ifclass {
      area 10.4.0.0;
      igp ospf;
      role cpe;
    }
  }
}
```

When you issue the **show protocols** configuration mode command, the script-generated changes should be present in the configuration.

```
[edit]
user@host# show protocols
isis {
  interface t3-0/0/0.0;
}
ospf {
  area 10.4.0.0 {
    interface so-1/2/3.0;
  }
}
```

### Example: Controlling IS-IS and MPLS Interfaces

---

This example shows how to use commit scripts to decrease the amount of manual configuration.

- [Requirements on page 367](#)
- [Overview and Commit Script on page 367](#)

- [Configuration on page 369](#)
- [Verification on page 370](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

If you want to enable MPLS on an interface, you must make changes at both the **[edit interfaces]** and **[edit protocols mpls]** hierarchy levels. This example shows how to use commit scripts to decrease the amount of manual configuration.

This example performs two related tasks. If an interface has **[family iso]** configured but not **[family mpls]**, a configuration change is made (using the **jcs:emit-change** template) to enable MPLS. MPLS is not valid on loopback interfaces (**loX**), so this script ignores loopback interfaces. Secondly, if the interface is not configured at the **[edit protocols mpls]** hierarchy level, a change is made to add the interface. Both changes are accompanied by appropriate warning messages.

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:variable name="mpls" select="protocols/mpls"/>
    <xsl:for-each select="interfaces/interface[not(starts-with(name,'lo'))]
      /unit[family/iso]">
      <xsl:variable name="ifname" select="concat(../name, '.', name)"/>
      <xsl:if test="not(family/mpls)">
        <xsl:call-template name="jcs:emit-change">
          <xsl:with-param name="message">
            <xsl:text>Adding 'family mpls' to ISO-enabled
interface</xsl:text>
          </xsl:with-param>
          <xsl:with-param name="content">
            <family>
              <mpls/>
            </family>
          </xsl:with-param>
        </xsl:call-template>
      </xsl:if>
      <xsl:if test="$mpls and not($mpls/interface[name = $ifname])">
        <xsl:call-template name="jcs:emit-change">
          <xsl:with-param name="message">
            <xsl:text>Adding ISO-enabled interface </xsl:text>
            <xsl:value-of select="$ifname"/>
            <xsl:text> to [protocols mpls]</xsl:text>
          </xsl:with-param>
          <xsl:with-param name="dot" select="$mpls"/>
          <xsl:with-param name="content">
            <interface>
```

```

        <name>
          <xsl:value-of select="$ifname"/>
        </name>
      </interface>
    </xsl:with-param>
  </xsl:call-template>
</xsl:if>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $mpls = protocols/mpls;
  for-each (interfaces/interface[not(starts-with(name,
"lo"))]/unit[family/iso]) {
    var $ifname = ../name _ '.' _ name;
    if (not(family/mpls)) {
      call jcs:emit-change() {
        with $message = {
          expr "Adding 'family mpls' to ISO-enabled interface";
        }
        with $content = {
          <family> {
            <mpls>;
          }
        }
      }
    }
    if ($mpls and not($mpls/interface[name = $ifname])) {
      call jcs:emit-change($dot = $mpls) {
        with $message = {
          expr "Adding ISO-enabled interface ";
          expr $ifname;
          expr " to [protocols mpls]";
        }
        with $content = {
          <interface> {
            <name> $ifname;
          }
        }
      }
    }
  }
}

```



## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **iso.xml** or **iso.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **iso.slax**.

```
system {
  scripts {
    commit {
      file iso.xml;
    }
  }
}
interfaces {
  lo0 {
    unit 0 {
      family iso;
    }
  }
  so-1/2/3 {
    unit 0 {
      family iso;
    }
  }
  so-1/3/2 {
    unit 0 {
      family iso;
    }
  }
}
protocols {
  mpls {
    enable;
  }
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.

4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

---

### Verifying the Configuration

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command.

```
[edit]
user@host# commit
[edit interfaces interface so-1/2/3 unit 0]
  warning: Adding 'family mpls' to ISO-enabled interface
[edit interfaces interface so-1/2/3 unit 0]
  warning: Adding ISO-enabled interface so-1/2/3.0 to [protocols mpls]
[edit interfaces interface so-1/3/2 unit 0]
  warning: Adding 'family mpls' to ISO-enabled interface
[edit interfaces interface so-1/3/2 unit 0]
  warning: Adding ISO-enabled interface so-1/3/2.0 to [protocols mpls]
commit complete
```

Issue the **show interfaces** command. Confirm that the loopback interface is not altered and that the SONET/SDH interfaces are altered.

```
[edit]
user@host# show interfaces
so-1/2/3 {
  unit 0 {
    family iso;
    family mpls;
  }
}
so-1/3/2 {
  unit 0 {
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family iso;
  }
}
```

---

## Example: Controlling LDP Configuration

This commit script example generates a warning on LDP-enabled devices for any interfaces that are configured at either the **[edit protocols ospf]** or **[edit protocols isis]** hierarchy level but are not configured at the **[edit protocols ldp]** hierarchy level. A second test ensures that all LDP-enabled interfaces are configured for an interior gateway

protocol (IGP). The example also provides instructions for excluding a particular interface from the commit script LDP test.

- [Requirements on page 371](#)
- [Overview and Commit Script on page 371](#)
- [Configuration on page 373](#)
- [Verification on page 374](#)

## Requirements

This example uses a router running Junos OS.

## Overview and Commit Script

If you want to enable LDP on an interface, you must configure the interface at both the **[edit protocols routing-protocol-name]** and **[edit protocols ldp]** hierarchy levels. This example shows how to use commit scripts to ensure that the interface is configured at both levels.

This example tests for interfaces that are configured at either the **[edit protocols ospf]** or **[edit protocols isis]** hierarchy level but not at the **[edit protocols ldp]** hierarchy level. If LDP is not enabled on the device, there is no problem. Otherwise, a warning is generated with the message that the interface does not have LDP enabled.

In case you want some interfaces to be exempt from the LDP test, this script allows you to tag those interfaces as not requiring LDP by including the **apply-macro no-ldp** statement at the **[edit protocols isis interface interface-name]** or **[edit protocols ospf area area-id interface interface-name]** hierarchy level. For example:

```
[edit]
protocols {
  isis {
    interface so-0/1/2.0 {
      apply-macro no-ldp;
    }
  }
}
```

If the **apply-macro no-ldp** statement is included, the warning is not generated.

A second test ensures that all LDP-enabled interfaces are configured for an interior gateway protocol (IGP). As for LDP, you can exempt some interfaces from the test by including the **apply-macro no-igp** statement at the **[edit protocols ldp interface interface-name]** hierarchy level. If that statement is not included and no IGP is configured, a warning is generated.

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

```

<xsl:import href="../../import/junos.xsl"/>

<xsl:template match="configuration">
  <xsl:variable name="ldp" select="protocols/ldp"/>
  <xsl:variable name="isis" select="protocols/isis"/>
  <xsl:variable name="ospf" select="protocols/ospf"/>
  <xsl:if test="$ldp">
    <xsl:for-each select="$isis/interface/name |
      $ospf/area/interface/name">
      <xsl:variable name="ifname" select="."/>
      <xsl:if test="not(../apply-macro[name = 'no-ldp'])
        and not($ldp/interface[name = $ifname])">
        <xnm:warning>
          <xsl:call-template name="jcs:edit-path"/>
          <xsl:call-template name="jcs:statement"/>
          <message>ldp not enabled for this interface</message>
        </xnm:warning>
      </xsl:if>
    </xsl:for-each>
    <xsl:for-each select="protocols/ldp/interface/name">
      <xsl:variable name="ifname" select="."/>
      <xsl:if test="not(apply-macro[name = 'no-igp'])
        and not($isis/interface[name = $ifname])
        and not($ospf/area/interface[name = $ifname])">
        <xnm:warning>
          <xsl:call-template name="jcs:edit-path"/>
          <xsl:call-template name="jcs:statement"/>
          <message>
            <xsl:text>ldp-enabled interface does not have
              <xsl:text>an IGP configured</xsl:text>
          </message>
        </xnm:warning>
      </xsl:if>
    </xsl:for-each>
  </xsl:if>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../../import/junos.xsl";

apply-macro no-ldp;
match configuration {
  var $ldp = protocols/ldp;
  var $isis = protocols/isis;
  var $ospf = protocols/ospf;
  if ($ldp) {
    for-each ($isis/interface/name | $ospf/area/interface/name) {
      var $ifname = .;
      if (not(../apply-macro[name = 'no-ldp']) and not($ldp/interface[name =
        $ifname])) {
        <xnm:warning> {
          call jcs:edit-path();
          call jcs:statement();
          <message> "ldp not enabled for this interface";

```

```

    }
  }
}
for-each (protocols/ldp/interface/name) {
  var $ifname = .;
  if (not(apply-macro[name = 'no-igp']) and not($isis/interface[name =
    $ifname]) and not($ospf/area/interface[name = $ifname])) {
    <xnm:warning> {
      call jcs:edit-path();
      call jcs:statement();
      <message> {
        expr "ldp-enabled interface does not have ";
        expr "an IGP configured";
      }
    }
  }
}
}
}
}
}
}

```

## Configuration

**Step-by-Step Procedure** To download, enable, and test the script:

1. Copy the script into a text file, name the file **ldp.xml** or **ldp.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **ldp.slax**.

```

system {
  scripts {
    commit {
      file ldp.xml;
    }
  }
}
protocols {
  isis {
    interface so-1/2/2.0 {
      apply-macro no-ldp;
    }
    interface so-1/2/3.0;
  }
  ospf {
    area 10.4.0.0 {
      interface ge-3/2/1.0;
      interface ge-2/2/1.0;
    }
  }
  ldp {
    interface ge-1/2/1.0;
    interface ge-2/2/1.0;
  }
}

```

```
}  
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]  
user@host# load merge terminal  
[Type ^D at a new line to end input]  
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

---

### Verifying the Script Execution

---

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The sample configuration stanzas enable LDP on the device and configure the so-1/2/2 and so-1/2/3 interfaces at the **[edit protocols isis]** hierarchy level and the ge-3/2/1 and ge-2/2/1 interfaces at the **[edit protocols ospf]** hierarchy level.

Because ge-2/2/1 is also configured at the **[edit protocols ldp]** hierarchy level, the script does not issue a warning message for this interface during the commit operation. The configuration includes the **apply-macro no-ldp** statement under the so-1/2/2 interface, so the script does not test this interface or issue a warning message for it, even though it is not configured at the **[edit protocols ldp]** hierarchy.

Neither so-1/2/3 nor ge-3/2/1 is configured at the **[edit protocols ldp]** hierarchy level as required by the commit script, so a warning is issued for both interfaces. The ge-1/2/1 interface is configured at the **[edit protocols ldp]** hierarchy. However, it is not configured for an IGP, so the commit script also issues a warning for the ge-1/2/1 interface.

```
[edit]  
user@host# commit  
  
[edit protocols ospf area 10.4.0.0 interface so-1/2/3.0]  
  'interface so-1/2/3.0;  
    warning: LDP not enabled for this interface  
[edit protocols ospf area 10.4.0.0 interface ge-3/2/1.0]  
  'interface ge-3/2/1.0;  
    warning: LDP not enabled for this interface  
[edit protocols ldp interface ge-1/2/1.0]  
  'interface ge-1/2/1.0;  
    warning: LDP-enabled interface does not have an IGP configured  
commit complete
```

## Example: Creating a Complex Configuration Based on a Simple Interface Configuration

This commit script example uses a macro to automatically expand a simple interface configuration.

- [Requirements on page 375](#)
- [Overview and Commit Script on page 375](#)
- [Configuration on page 379](#)
- [Verification on page 380](#)

### Requirements

This example uses a device running Junos OS.

### Overview and Commit Script

This example uses a commit script and macro to automatically expand a simple interface configuration by generating a transient change that assigns a default encapsulation type, configures multiple routing protocols on the interface, and applies multiple configuration groups. The Junos OS management (mgd) process inspects the configuration, looking for **apply-macro params** statements included at the **[edit interfaces *interface-name*]** hierarchy level.

When the script finds an **apply-macro params** statement, it performs the following actions:

- Applies the **interface-details** configuration group to the interface.
- Includes the value of the **description** parameter at the **[edit interfaces *interface-name* description]** hierarchy level.
- Includes the value of the **encapsulation** parameter at the **[edit interfaces *interface-name* encapsulation]** hierarchy level. If the **encapsulation** parameter is not included in the **apply-macro params** statement, the script sets the encapsulation to **cisco-hdlc** as the default.
- Sets the logical unit number to **0** and tests whether the **inet-address** parameter is included in the **apply-macro params** statement. If it is, the script includes the value of the **inet-address** parameter at the **[edit interfaces *interface-name* unit 0 family inet address]** hierarchy level.
- Includes the interface name at the **[edit protocols rsvp interface]** hierarchy level.
- Includes the **level 1 enable** and **metric** statements at the **[edit protocols isis interface *interface-name*]** hierarchy level.
- Includes the **level 2 enable** and **metric** statements at the **[edit protocols isis interface *interface-name*]** hierarchy level.
- Tests whether the **isis-level-1** or **isis-level-1-metric** parameter is included in the **apply-macro params** statement. If one or both of these parameters are included, the script includes the **level 1** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level. If the **isis-level-1** parameter is included, the script also includes the value

of the `isis-level-1` parameter (`enable` or `disable`) at the `[edit protocols isis interface interface-name level 1]` hierarchy level. If the `isis-level-1-metric` parameter is included, the script also includes the value of the `isis-level-1-metric` parameter at the `[edit protocols isis interface interface-name level 1 metric]` hierarchy level.

- Tests whether the `isis-level-2` or `isis-level-2-metric` parameter is included in the `apply-macro params` statement. If one or both of these parameters are included, the script includes the `level 2` statement at the `[edit protocols isis interface interface-name level 2]` hierarchy level. If the `isis-level-2` parameter is included, the script also includes the value of the `isis-level-2` parameter (`enable` or `disable`) at the `[edit protocols isis interface interface-name level 2]` hierarchy level. If the `isis-level-2-metric` parameter is included, the script also includes the value of the `isis-level-2-metric` parameter at the `[edit protocols isis interface interface-name level 2 metric]` hierarchy level.
- Includes the interface name at the `[edit protocols ldp interface]` hierarchy level.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
<xsl:import href="../import/junos.xsl"/>

<xsl:template match="configuration">
  <xsl:variable name="top" select="."/>
  <xsl:for-each select="interfaces/interface/apply-macro[name = 'params']">
    <xsl:variable name="description"
      select="data[name = 'description']/value"/>
    <xsl:variable name="inet-address"
      select="data[name = 'inet-address']/value"/>
    <xsl:variable name="encapsulation"
      select="data[name = 'encapsulation']/value"/>
    <xsl:variable name="isis-level-1"
      select="data[name = 'isis-level-1']/value"/>
    <xsl:variable name="isis-level-1-metric"
      select="data[name = 'isis-level-1-metric']/value"/>
    <xsl:variable name="isis-level-2"
      select="data[name = 'isis-level-2']/value"/>
    <xsl:variable name="isis-level-2-metric"
      select="data[name = 'isis-level-2-metric']/value"/>
    <xsl:variable name="ifname" select="concat(..name, '.0')"/>
    <transient-change>
      <interfaces>
        <interface>
          <name><xsl:value-of select="..name"/></name>
          <apply-groups>
            <name>interface-details</name>
          </apply-groups>
          <xsl:if test="$description">
            <description>
              <xsl:value-of select="$description"/>
            </description>
          </if>
        </interface>
      </interfaces>
    </transient-change>
  </for-each>
</template>
```



```

</xsl:if>
<encapsulation>
  <xsl:choose>
    <xsl:when test="string-length($encapsulation) > 0">
      <xsl:value-of select="$encapsulation"/>
    </xsl:when>
    <xsl:otherwise>cisco-hdlc</xsl:otherwise>
  </xsl:choose>
</encapsulation>
<unit>
  <name>0</name>
  <xsl:if test="string-length($inet-address) > 0">
    <family>
      <inet>
        <address>
          <xsl:value-of select="$inet-address"/>
        </address>
      </inet>
    </family>
  </xsl:if>
</unit>
</interface>
</interfaces>
<protocols>
  <rsvp>
    <interface>
      <name><xsl:value-of select="$ifname"/></name>
    </interface>
  </rsvp>
  <isis>
    <interface>
      <name><xsl:value-of select="$ifname"/></name>
      <xsl:if test="$isis-level-1 or $isis-level-1-metric">
        <level>
          <name>1</name>
          <xsl:if test="$isis-level-1">
            <xsl:element name="{ $isis-level-1 }"/>
          </xsl:if>
          <xsl:if test="$isis-level-1-metric">
            <metric>
              <xsl:value-of select="$isis-level-1-metric"/>
            </metric>
          </xsl:if>
        </level>
      </xsl:if>
      <xsl:if test="$isis-level-2 or $isis-level-2-metric">
        <level>
          <name>2</name>
          <xsl:if test="$isis-level-2">
            <xsl:element name="{ $isis-level-2 }"/>
          </xsl:if>
          <xsl:if test="$isis-level-2-metric">
            <metric>
              <xsl:value-of select="$isis-level-2-metric"/>
            </metric>
          </xsl:if>
        </level>
      </xsl:if>
    </interface>
  </isis>
</protocols>

```

```

        </level>
      </xsl:if>
    </interface>
  </isis>
  <ldp>
    <interface>
      <name><xsl:value-of select="$ifname"/></name>
    </interface>
  </ldp>
</protocols>
</transient-change>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>

```

## SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  var $top = .;
  for-each (interfaces/interface/apply-macro[name = 'params']) {
    var $description = data[name = 'description']/value;
    var $inet-address = data[name = 'inet-address']/value;
    var $encapsulation = data[name = 'encapsulation']/value;
    var $isis-level-1 = data[name = 'isis-level-1']/value;
    var $isis-level-1-metric = data[name = 'isis-level-1-metric']/value;
    var $isis-level-2 = data[name = 'isis-level-2']/value;
    var $isis-level-2-metric = data[name = 'isis-level-2-metric']/value;
    var $ifname = ../name - '.0';
    <transient-change> {
      <interfaces> {
        <interface> {
          <name> ../name;
          <apply-groups> {
            <name> "interface-details";
          }
          if ($description) {
            <description> $description;
          }
          <encapsulation> {
            if (string-length($encapsulation) > 0) {
              expr $encapsulation;
            } else {
              expr "cisco-hdlc";
            }
          }
          <unit> {
            <name> "0";
            if (string-length($inet-address) > 0) {
              <family> {
                <inet> {
                  <address> $inet-address;
                }
              }
            }
          }
        }
      }
    }
  }
}

```

## Configuration

To download, enable, and test the script:

- If you are using the SLAX version of the script, change the filename at the `[edit system scripts commit file]` hierarchy level to `if-params.slax`.

Copyright © 2016, Juniper Networks, Inc.

```
        allow-transients;
        file if-params.xml;
    }
}
groups {
    interface-details {
        interfaces {
            <so-*/*/*> {
                clocking internal;
            }
        }
    }
}
interfaces {
    so-1/2/3 {
        apply-macro params {
            description "Link to Hoverville";
            encapsulation ppp;
            inet-address 10.1.2.3/28;
            isis-level-1 enable;
            isis-level-1-metric 50;
            isis-level-2-metric 85;
        }
    }
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Configuration

---

**Purpose** Verify that the script behaves as expected.

**Action** Issue the **show interfaces | display commit-scripts | display inheritance** configuration mode command. The **| display commit-scripts** option displays all the statements that are in the configuration, including statements that are generated by transient changes. The **| display inheritance** option displays inherited configuration data and information about the source group from which the configuration has been inherited. This option also shows interface ranges configuration data in expanded format and information about the source interface-range from which the configuration has been expanded. You should see the following output:

```
[edit]
user@host# show interfaces | display commit-scripts | display inheritance
so-1/2/3 {
  apply-macro params {
    clocking internal;
    description "Link to Hoverville";
    encapsulation ppp;
    inet-address 10.1.2.3/28;
    isis-level-1 enable;
    isis-level-1-metric 50;
    isis-level-2-metric 85;
  }
  description "Link to Hoverville";
  ##
  ## 'internal' was inherited from group 'interface-details'
  ##
  clocking internal;
  encapsulation ppp;
  unit 0 {
    family inet {
      address 10.1.2.3/28;
    }
  }
}
```

Issue the **show protocols | display commit-scripts** configuration mode command. You should see the following output:

```
[edit]
user@host# show protocols | display commit-scripts
rsvp {
  interface so-1/2/3.0;
}
isis {
  interface so-1/2/3.0 {
    level 1 {
      enable;
      metric 50;
    }
    level 2 metric 85;
  }
}
ldp {
  interface so-1/2/3.0;
}
```

## Example: Imposing a Minimum MTU Setting

The maximum transmission unit (MTU) is the greatest amount of data or packet size (in bytes) that can be transferred in one physical frame on a network. In this example, a commit script tests the MTU of SONET/SDH interfaces. If the MTU is less than a specified minimum value, the commit script reports the error and causes the commit operation to fail.

- [Requirements on page 382](#)
- [Overview and Commit Script on page 382](#)
- [Configuration on page 383](#)
- [Verification on page 384](#)

### Requirements

This example uses a device running Junos OS with SONET/SDH interfaces.

### Overview and Commit Script

This example tests the MTU of SONET/SDH interfaces, reports when the MTU is less than the value of the **min-mtu** parameter, here set to 2048, and causes the commit operation to fail. The **for** loop selects all SONET/SDH interfaces that start with so- and that have an MTU statement that is defined and less than the value of **min-mtu**. For the selected interfaces, the script generates an error, which includes the location of the interface in the configuration hierarchy and the MTU configured for that interface.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:param name="min-mtu" select="2048"/>
  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name, 'so-')
      and mtu and mtu < $min-mtu]">
      <xnm:error>
        <xsl:call-template name="jcs:edit-path"/>
        <xsl:call-template name="jcs:statement">
          <xsl:with-param name="dot" select="mtu"/>
        </xsl:call-template>
        <message>
          <xsl:text>SONET interfaces must have a minimum MTU of
            <xsl:value-of select="$min-mtu"/>
            <xsl:text>.</xsl:text>
          </message>
        </xnm:error>
      </xsl:for-each>
    </xsl:template>
  </xsl:stylesheet>
```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

param $min-mtu = 2048;
match configuration {
  for-each (interfaces/interface[starts-with(name, 'so-') and mtu and
    mtu < $min-mtu]) {
    <xnm:error> {
      call jcs:edit-path();
      call jcs:statement($dot = mtu);
      <message> {
        expr "SONET interfaces must have a minimum MTU of ";
        expr $min-mtu;
        expr ".";
      }
    }
  }
}

```

**Configuration****Step-by-Step  
Procedure**

To download, enable, and test the script:

1. Copy the script into a text file, name the file **so-mtu.xsl** or **so-mtu.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **so-mtu.slax**.

```

system {
  scripts {
    commit {
      file so-mtu.xsl;
    }
  }
}
interfaces {
  so-1/2/2 {
    mtu 2048;
  }
  so-1/2/3 {
    mtu 576;
  }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.  
  
user@host# **commit**

## Verification

---

### Verifying the Commit Script Output

- Purpose** Verify that the script behaves as expected.
- Action** Review the output of the **commit** command. The sample configuration stanzas configure two SONET/SDH interfaces so-1/2/2 and so-1/2/3. The so-1/2/3 interface is configured with an MTU of 576, so the script generates an error message, and the commit operation fails. The following output appears after issuing a **commit** command:
- ```
[edit]
user@host# commit
[edit interfaces interface so-1/2/3]
  'mtu 576;'
      SONET interfaces must have a minimum MTU of 2048.
error: 1 error reported by commit scripts
error: commit script failure
```

---

## Example: Limiting the Number of ATM Virtual Circuits

This commit script example limits the number of Asynchronous Transfer Mode (ATM) virtual circuits (VCs) configured on an ATM interface.

- [Requirements on page 384](#)
- [Overview and Commit Script on page 384](#)
- [Configuration on page 386](#)
- [Verification on page 387](#)

## Requirements

This example uses a device running Junos OS with an ATM interface.

## Overview and Commit Script

For each ATM interface, the set of corresponding VCs is selected. The number of those VCs, as determined by the built-in Extensible Stylesheet Language Transformations (XSLT) **count()** function, cannot exceed the limit set by the global variable **limit**. If there are more ATM VCs than **limit**, a commit error is generated, and the commit operation fails.

The example script is shown in both XSLT and SLAX syntax:



**XSLT Syntax**

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:param name="limit" select="10"/>
  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name, 'at-')] ">
      <xsl:variable name="count" select="count(unit)"/>
      <xsl:if test="$count > $limit">
        <xnm:error>
          <edit-path>[edit interfaces]</edit-path>
          <statement><xsl:value-of select="name"/></statement>
          <message>
            <xsl:text>ATM VC limit exceeded; </xsl:text>
            <xsl:value-of select="$count"/>
            <xsl:text> are configured but only </xsl:text>
            <xsl:value-of select="$limit"/>
            <xsl:text> are allowed.</xsl:text>
          </message>
        </xnm:error>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

param $limit = 10;
match configuration {
  for-each (interfaces/interface[starts-with(name, 'at-')]) {
    var $count = count(unit);
    if ($count > $limit) {
      <xnm:error> {
        <edit-path> "[edit interfaces]";
        <statement> name;
        <message> {
          expr "ATM VC limit exceeded; ";
          expr $count;
          expr " are configured but only ";
          expr $limit;
          expr " are allowed.";
        }
      }
    }
  }
}

```

## Configuration

**Step-by-Step Procedure** To download, enable, and test the script:

1. Copy the script into a text file, name the file **atm-vc-limit.xsl** or **atm-vc-limit.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **atm-vc-limit.slax**.

```
system {
  scripts {
    commit {
      file atm-vc-limit.xsl;
    }
  }
}
interfaces {
  at-1/2/3 {
    unit 15 {
      family inet {
        address 10.12.13.15/20;
      }
    }
    unit 16 {
      family inet {
        address 10.12.13.16/20;
      }
    }
    unit 17 {
      family inet {
        address 10.12.13.17/20;
      }
    }
    unit 18 {
      family inet {
        address 10.12.13.18/20;
      }
    }
    unit 19 {
      family inet {
        address 10.12.13.19/20;
      }
    }
    unit 20 {
      family inet {
        address 10.12.13.20/20;
      }
    }
    unit 21 {
      family inet {
        address 10.12.13.21/20;
      }
    }
  }
}
```

```

    }
    unit 22 {
        family inet {
            address 10.12.13.22/20;
        }
    }
    unit 23 {
        family inet {
            address 10.12.13.23/20;
        }
    }
    unit 24 {
        family inet {
            address 10.12.13.24/20;
        }
    }
    unit 25 {
        family inet {
            address 10.12.13.25/20;
        }
    }
    unit 26 {
        family inet {
            address 10.12.13.26/20;
        }
    }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying the Commit Script Output

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The sample configuration stanzas configure 12 virtual circuits on the ATM interface atm-1/2/3. Because the commit script only allows 10 ATM VCs to be configured on any ATM interface, the script generates an error, and the commit operation fails. The following output appears after issuing a **commit** command:

```
[edit]
user@host# commit
[edit interfaces]
  'at-1/2/3'
    ATM VC limit exceeded; 12 are configured but only 10 are allowed.
error: 1 error reported by commit scripts
error: commit script failure
```

---

## Example: Limiting the Number of E1 Interfaces

This commit script example limits the number of E1 interfaces configured on a Channelized STM1 Intelligent Queuing (IQ) PIC to avoid contention issues with per-unit-schedulers.

- [Requirements on page 388](#)
- [Overview and Commit Script on page 388](#)
- [Configuration on page 390](#)
- [Verification on page 397](#)

### Requirements

This example uses a device running Junos OS with a Channelized STM1 Intelligent Queuing (IQ) PIC.

### Overview and Commit Script

The following script ensures that there are no more than 16 E1 interfaces configured on a channelized STM1 IQ interface. For each channelized STM1 interface (cstm1-), the set of corresponding E1 interfaces is selected. The number of those interfaces, as determined by the built-in Extensible Stylesheet Language Transformations (XSLT) **count()** function, cannot exceed the limit set by the global parameter **limit**. If there are more E1 interfaces than **limit**, a commit error is generated, and the commit operation fails.

The example script is shown in both XSLT and SLAX syntax:

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>XSLT Syntax</b> | <pre>&lt;?xml version="1.0" standalone="yes"?&gt; &lt;xsl:stylesheet version="1.0"   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"   xmlns:junos="http://xml.juniper.net/junos/*/junos"   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"&gt;   &lt;xsl:import href="../import/junos.xml"/&gt;    &lt;xsl:param name="limit" select="16"/&gt;   &lt;xsl:template match="configuration"&gt;     &lt;xsl:variable name="interfaces" select="interfaces"/&gt;     &lt;xsl:for-each select="\$interfaces/interface[starts-with(name, 'cstm1-')]"&gt;       &lt;xsl:variable name="triple" select="substring-after(name, 'cstm1-')"/&gt;       &lt;xsl:variable name="e1name" select="concat('e1-', \$triple)"/&gt;       &lt;xsl:variable name="count"</pre> |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

```

        select="count($interfaces/interface[starts-with(name, $elname)])"/>
<xsl:if test="$count > $limit">
  <xnm:error>
    <edit-path>[edit interfaces]</edit-path>
    <statement><xsl:value-of select="name"/></statement>
    <message>
      <xsl:text>E1 interface limit exceeded on CSTM1 IQ PIC. </xsl:text>
      <xsl:value-of select="$count"/>
      <xsl:text> E1 interfaces are configured, but only </xsl:text>
      <xsl:value-of select="$limit"/>
      <xsl:text> are allowed.</xsl:text>
    </message>
  </xnm:error>
</xsl:if>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

param $limit = 16;
match configuration {
  var $interfaces = interfaces;
  for-each ($interfaces/interface[starts-with(name, 'cstm1-')]) {
    var $triple = substring-after(name, 'cstm1-');
    var $elname = 'e1-' _ $triple;
    var $count = count($interfaces/interface[starts-with(name, $elname)]);
    if ($count > $limit) {
      <xnm:error> {
        <edit-path> "[edit interfaces]";
        <statement> name;
        <message> {
          expr "E1 interface limit exceeded on CSTM1 IQ PIC. ";
          expr $count;
          expr " E1 interfaces are configured, but only ";
          expr $limit;
          expr " are allowed.";
        }
      }
    }
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **e1-limit.xml** or **e1-limit.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **e1-limit.slax**.

```
system {
  scripts {
    commit {
      file e1-limit.xml;
    }
  }
}
interfaces {
  cau4-0/1/0 {
    partition 1 interface-type ce1;
    partition 2-18 interface-type e1;
  }
  cstm1-0/1/0 {
    no-partition interface-type cau4;
  }
  ce1-0/1/0:1 {
    clocking internal;
    e1-options {
      framing g704;
    }
    partition 1 timeslots 1-4 interface-type ds;
  }
  ds-0/1/0:1:1 {
    no-keepalives;
    dce;
    encapsulation frame-relay;
    lmi {
      lmi-type ansi;
    }
    unit 100 {
      point-to-point;
      dlci 100;
      family inet {
        address 10.0.0.0/31;
      }
    }
  }
  e1-0/1/0:2 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
```

```
e1-options {
    framing g704;
}
lmi {
    lmi-type ansi;
}
unit 100 {
    point-to-point;
    dlci 100;
    family inet {
        address 10.0.0.2/31;
    }
}
}
e1-0/1/0:3 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.4/31;
        }
    }
}
}
e1-0/1/0:4 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.6/31;
        }
    }
}
}
e1-0/1/0:5 {
    no-keepalives;
```

```
per-unit-scheduler;
dce;
clocking internal;
encapsulation frame-relay;
e1-options {
    framing g704;
}
lmi {
    lmi-type ansi;
}
unit 100 {
    point-to-point;
    dlci 100;
    family inet {
        address 10.0.0.8/31;
    }
}
}
e1-0/1/0:6 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.10/31;
        }
    }
}
e1-0/1/0:7 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.12/31;
        }
    }
}
```



```
    }  
  }  
  e1-0/1/0:8 {  
    no-keepalives;  
    per-unit-scheduler;  
    dce;  
    clocking internal;  
    encapsulation frame-relay;  
    e1-options {  
      framing g704;  
    }  
    lmi {  
      lmi-type ansi;  
    }  
    unit 100 {  
      point-to-point;  
      dlci 100;  
      family inet {  
        address 10.0.0.14/31;  
      }  
    }  
  }  
  e1-0/1/0:9 {  
    no-keepalives;  
    per-unit-scheduler;  
    dce;  
    clocking internal;  
    encapsulation frame-relay;  
    e1-options {  
      framing g704;  
    }  
    lmi {  
      lmi-type ansi;  
    }  
    unit 100 {  
      point-to-point;  
      dlci 100;  
      family inet {  
        address 10.0.0.16/31;  
      }  
    }  
  }  
  e1-0/1/0:10 {  
    no-keepalives;  
    per-unit-scheduler;  
    dce;  
    clocking internal;  
    encapsulation frame-relay;  
    e1-options {  
      framing g704;  
    }  
    lmi {  
      lmi-type ansi;  
    }  
    unit 100 {  
      point-to-point;
```

```
        dlc1 100;
        family inet {
            address 10.0.0.18/31;
        }
    }
}
e1-0/1/0:11 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlc1 100;
        family inet {
            address 10.0.0.20/31;
        }
    }
}
e1-0/1/0:12 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlc1 100;
        family inet {
            address 10.0.0.22/31;
        }
    }
}
e1-0/1/0:13 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
```

```
    lmi-type ansi;
  }
  unit 100 {
    point-to-point;
    dlci 100;
    family inet {
      address 10.0.0.24/31;
    }
  }
}
e1-0/1/0:14 {
  no-keepalives;
  per-unit-scheduler;
  dce;
  clocking internal;
  encapsulation frame-relay;
  e1-options {
    framing g704;
  }
  lmi {
    lmi-type ansi;
  }
  unit 100 {
    point-to-point;
    dlci 100;
    family inet {
      address 10.0.0.26/31;
    }
  }
}
e1-0/1/0:15 {
  no-keepalives;
  per-unit-scheduler;
  dce;
  clocking internal;
  encapsulation frame-relay;
  e1-options {
    framing g704;
  }
  lmi {
    lmi-type ansi;
  }
  unit 100 {
    point-to-point;
    dlci 100;
    family inet {
      address 10.0.0.28/31;
    }
  }
}
e1-0/1/0:16 {
  no-keepalives;
  per-unit-scheduler;
  dce;
  clocking internal;
  encapsulation frame-relay;
```

```
e1-options {
    framing g704;
}
lmi {
    lmi-type ansi;
}
unit 100 {
    point-to-point;
    dlci 100;
    family inet {
        address 10.0.0.30/31;
    }
}
}
e1-0/1/0:17 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.32/31;
        }
    }
}
e1-0/1/0:18 {
    no-keepalives;
    per-unit-scheduler;
    dce;
    clocking internal;
    encapsulation frame-relay;
    e1-options {
        framing g704;
    }
    lmi {
        lmi-type ansi;
    }
    unit 100 {
        point-to-point;
        dlci 100;
        family inet {
            address 10.0.0.34/31;
        }
    }
}
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Commit Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The sample configuration stanzas channelize a cstm1-0/1/0 interface into 17 E1 interfaces, so the script generates an error, and the commit operation fails. The following output appears after issuing a **commit** command:

```
[edit]
user@host# commit
[edit interfaces]
'cstm1-0/1/0'
E1 interface limit exceeded on CSTM1 IQ PIC.
17 E1 interfaces are configured, but only 16 are allowed.
error: 1 error reported by commit scripts
error: commit script failure
```

## Example: Loading a Base Configuration

This commit script example sets up a sample base configuration on a device running Junos OS.

- [Requirements on page 397](#)
- [Overview and Commit Script on page 398](#)
- [Configuration on page 410](#)
- [Verification on page 411](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

This script is a macro that sets up a device running Junos OS with a sample base configuration. With minimal manual user input, the script automatically configures:

- A device hostname
- Authentication services
- A superuser login
- System log settings
- Some SNMP settings
- System services, such as FTP and Telnet
- Static routes and a policy to redistribute the static routes
- Configuration groups **re0** and **re1**
- An address for the management Ethernet interface (fxp0)
- The loopback interface (lo0) with the device ID as the loopback address

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xml"/>

  <xsl:variable name="macro-name" select="'config-system.xml'"/>
  <xsl:template match="configuration">
    <xsl:variable name="rid" select="routing-options/router-id"/>
    <xsl:for-each select="apply-macro[name = 'config-system']">
      <xsl:variable name="hostname" select="data[name =
        'host-name']/value"/>
      <xsl:variable name="fxp0-addr" select="data[name =
        'mgmt-address']/value"/>
      <xsl:variable name="backup-router" select="data[name =
        'backup-router']/value"/>
      <xsl:variable name="bkup-rtr">
        <xsl:choose>
          <xsl:when test="$backup-router">
            <xsl:value-of select="$backup-router"/>
          </xsl:when>
          <xsl:otherwise>
            <xsl:variable name="fxp01" select="substring-before($fxp0-addr,
              '.')"/>
            <xsl:variable name="fxp02"
              select="substring-before(substring-after($fxp0-addr, '.'), '.')"/>
            <xsl:variable name="fxp03"
              select="substring-before(substring-after(substring-after(
                $fxp0-addr, '.'), '.'), '.')"/>
```

```

<xsl:variable name="plen" select="substring-after($fxp0-addr, '/')"/>
<xsl:choose>
  <xsl:when test="$plen = 22">
    <xsl:value-of select="concat($fxp01, '.', $fxp02, '.', $fxp03 div
      4 * 4 + 3, '.254')"/>
  </xsl:when>
  <xsl:when test="$plen = 24">
    <xsl:value-of select="concat($fxp01, '.', $fxp02, '.', $fxp03,
      '.254')"/>
  </xsl:when>
</xsl:choose>
</xsl:otherwise>
</xsl:choose>
</xsl:variable>
<xsl:choose>
  <xsl:when test="not($rid) or not($hostname) or not($fxp0-addr)">
    <xnm:error>
      <message>
        Must set router ID, host-name and mgmt-address to use this script.
      </message>
    </xnm:error>
  </xsl:when>
  <xsl:otherwise>
    <transient-change>
      <system>
        <!-- Set the following -->
        <domain-name>your-domain.net</domain-name>
        <domain-search>domain.net</domain-search>
        <backup-router>
          <address><xsl:value-of select="$bkup-rtr"/></address>
        </backup-router>
        <time-zone>America/Los_Angeles</time-zone>
        <authentication-order>radius</authentication-order>
        <authentication-order>password</authentication-order>
        <root-authentication>
          <encrypted-password>
            $ABC123
          </encrypted-password>
        </root-authentication>
        <name-server>
          <name>192.168.5.68</name>
        </name-server>
        <name-server>
          <name>172.17.28.100</name>
        </name-server>
        <radius-server>
          <name>192.168.170.241</name>
          <secret>
            $ABC123
          </secret>
        </radius-server>
        <radius-server>
          <name>192.168.4.240</name>
          <secret>
            $ABC123
          </secret>

```

```
</radius-server>
<login>
  <class>
    <permissions>all</permissions>
  </class>
  <user>
    <name>johnny</name>
    <uid>928</uid>
    <class>superuser</class>
    <authentication>
      <encrypted-password>
        $ABC123
      </encrypted-password>
    </authentication>
  </user>
</login>
<services>
  <finger/>
  <ftp/>
  <ssh/>
  <telnet/>
  <xnm-clear-text/>
</services>
<syslog>
  <user>
    <name>*</name>
    <contents>
      <name>any</name>
      <emergency/>
    </contents>
  </user>
  <host>
    <name>host1</name>
    <contents>
      <name>any</name>
      <notice/>
    </contents>
    <contents>
      <name>interactive-commands</name>
      <any/>
    </contents>
  </host>
  <file>
    <name>messages</name>
    <contents>
      <name>any</name>
      <notice/>
    </contents>
    <contents>
      <name>any</name>
      <warning/>
    </contents>
    <contents>
      <name>authorization</name>
      <info/>
    </contents>
```



```

    <archive>
      <world-readable/>
    </archive>
  </file>
  <file>
    <name>security</name>
    <contents>
      <name>interactive-commands</name>
      <any/>
    </contents>
    <archive>
      <world-readable/>
    </archive>
  </file>
</syslog>
<processes>
  <routing>
    <undocumented><enable/></undocumented>
  </routing>
  <snmp>
    <undocumented><enable/></undocumented>
  </snmp>
  <ntp>
    <undocumented><enable/></undocumented>
  </ntp>
  <inet-process>
    <undocumented><enable/></undocumented>
  </inet-process>
  <mib-process>
    <undocumented><enable/></undocumented>
  </mib-process>
  <undocumented><management><enable/>
  </undocumented></management>
  <watchdog>
    <enable/>
  </watchdog>
</processes>
<ntp>
  <boot-server>domain.net</boot-server>
  <server>
    <name>domainr.net</name>
  </server>
</ntp>
</system>
<snmp>
  <location>Software lab</location>
  <contact>Michael Landon</contact>
  <interface>fxp0.0</interface>
  <community>
    <name>public</name>
    <authorization>read-only</authorization>
  <clients>
    <name>0.0.0.0/0</name>
    <restrict/>
  </clients>
  <clients>

```

```
<name>192.168.1.252/32</name>
</clients>
<clients>
  <name>10.197.169.222/32</name>
</clients>
<clients>
  <name>10.197.169.188/32</name>
</clients>
<clients>
  <name>10.197.169.193/32</name>
</clients>
<clients>
  <name>192.168.65.46/32</name>
</clients>
<clients>
  <name>10.209.152.0/23</name>
</clients>
</community>
<community>
  <name>private</name>
  <authorization>read-write</authorization>
  <clients>
    <name>0.0.0.0/0</name>
    <restrict/>
  </clients>
  <clients>
    <name>10.197.169.188/32</name>
  </clients>
</community>
</snmp>
<routing-options>
  <static>
    <junos:comment>/* safety precaution */</junos:comment>
    <route>
      <name>0.0.0.0/0</name>
      <discard/>
      <retain/>
      <no-readvertise/>
    </route>
    <junos:comment>/* corporate net */</junos:comment>
    <route>
      <name>172.16.0.0/12</name>
      <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
      <retain/>
      <no-readvertise/>
    </route>
    <junos:comment>/* lab nets */</junos:comment>
    <route>
      <name>192.168.0.0/16</name>
      <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
      <retain/>
      <no-readvertise/>
    </route>
    <junos:comment>/* reflector */</junos:comment>
    <route>
      <name>10.17.136.192/32</name>
```

```

    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
  <junos:comment>/* another lab1*/</junos:comment>
  <route>
    <name>10.10.0.0/16</name>
    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
  <junos:comment>/* ssh servers */</junos:comment>
  <route>
    <name>10.17.136.0/24</name>
    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
  <junos:comment>/* Workstations */</junos:comment>
  <route>
    <name>10.150.0.0/16</name>
    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
  <junos:comment>/* Hosts */</junos:comment>
  <route>
    <name>10.157.64.0/19</name>
    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
  <junos:comment>/* Build Servers */</junos:comment>
  <route>
    <name>10.10.0.0/16</name>
    <next-hop><xsl:value-of select="$bkup-rtr"/></next-hop>
    <retain/>
    <no-readvertise/>
  </route>
</static>
</routing-options>
<policy-options>
  <policy-statement>
    <name>redist</name>
    <from>
      <protocol>static</protocol>
    </from>
    <then>
      <accept/>
    </then>
  </policy-statement>
</policy-options>
<apply-groups>re0</apply-groups>
<apply-groups>re1</apply-groups>
<groups>
  <name>re0</name>

```

```
<system>
  <host-name>
    <xsl:value-of select="$hostname"/></host-name>
  </system>
</interfaces>
<interface>
  <name>fxp0</name>
  <unit>
    <name>0</name>
    <family>
      <inet>
        <address>
          <name>
            <xsl:value-of select="$fxp0-addr"/>
          </name>
        </address>
      </inet>
    </family>
  </unit>
</interface>
</interfaces>
</groups>
<groups>
  <name>re1</name>
</groups>
<interfaces>
  <interface>
    <name>lo0</name>
    <unit>
      <name>0</name>
      <family>
        <inet>
          <address>
            <name><xsl:value-of select="$rid"/></name>
          </address>
        </inet>
      </family>
    </unit>
  </interface>
</interfaces>
</transient-change>
</xsl:otherwise>
</xsl:choose>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>
```

**SLAX Syntax**

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";

var $macro-name = 'config-system.xml';
match configuration {
  var $rid = routing-options/router-id;
```

```

for-each (apply-macro[name = 'config-system']) {
  var $hostname = data[name = 'host-name']/value;
  var $fxp0-addr = data[name = 'mgmt-address']/value;
  var $backup-router = data[name = 'backup-router']/value;
  var $bkup-rtr = {
    if ($backup-router) {
      expr $backup-router;
    }
    else {
      var $fxp01 = substring-before($fxp0-addr, '.');
      var $fxp02 = substring-before(substring-after($fxp0-addr, '.'), '.');
      var $fxp03 = substring-before(substring-after(substring-after(
        $fxp0-addr, '.'), '.'), '.');
      var $plen = substring-after($fxp0-addr, '/');
      if ($plen = 22) {
        expr $fxp01 _ '.' _ $fxp02 _ '.' _ $fxp03 div 4 * 4 + 3 _ '.254';
      }
      else if ($plen = 24) {
        expr $fxp01 _ '.' _ $fxp02 _ '.' _ $fxp03 _ '.254';
      }
    }
  }
}
if (not($rid) or not($hostname) or not($fxp0-addr)) {
  <xn timer> {
    <message> "Must set router ID, host-name, and mgmt-address to use
      this script.";
  }
}
else {
  <transient-change> {
    <system> {
      /* Set the following */
      <domain-name> "your-domain.net";
      <domain-search> "domain.net";
      <backup-router> {
        <address> $bkup-rtr;
      }
      <time-zone> "America/Los_Angeles";
      <authentication-order> "radius";
      <authentication-order> "password";
      <root-authentication> {
        <encrypted-password>
          "$ABC123";
      }
      <name-server> {
        <name> "192.168.5.68";
      }
      <name-server> {
        <name> "172.17.28.100";
      }
      <radius-server> {
        <name> "192.168.170.241";
        <secret> "$ABC123";
      }
      <radius-server> {
        <name> "192.168.4.240";

```

```
<secret> "$ABC123";
}
<login> {
  <class> {
    <permissions> "all";
  }
  <user> {
    <name> "johnny";
    <uid> "928";
    <class> "superuser";
    <authentication> {
      <encrypted-password> "$ABC123";
    }
  }
}
}
<services> {
  <finger>;
  <ftp>;
  <ssh>;
  <telnet>;
  <xnm-clear-text>;
}
<syslog> {
  <user> {
    <name> "*";
    <contents> {
      <name> "any";
      <emergency>;
    }
  }
}
<host> {
  <name> "host1";
  <contents> {
    <name> "any";
    <notice>;
  }
  <contents> {
    <name> "interactive-commands";
    <any>;
  }
}
}
<file> {
  <name> "messages";
  <contents> {
    <name> "any";
    <notice>;
  }
  <contents> {
    <name> "any";
    <warning>;
  }
  <contents> {
    <name> "authorization";
    <info>;
  }
}
<archive> {
```

```

        <world-readable>;
    }
}
<file> {
    <name> "security";
    <contents> {
        <name> "interactive-commands";
        <any>;
    }
    <archive> {
        <world-readable>;
    }
}
}
<processes> {
    <routing> {
        <undocumented><enable>;
    }
    <snmp> {
        <undocumented><enable>;
    }
    <ntp> {
        <undocumented><enable>;
    }
    <inet-process> {
        <undocumented> <enable>;
    }
    <mib-process> {
        <undocumented> <enable>;
    }
    <undocumented><management> {
        <enable>;
    }
    <watchdog> {
        <enable>;
    }
    <ntp> {
        <boot-server> "domain.net";
        <server> {
            <name> "domainr.net";
        }
    }
}
<snmp> {
    <location> "Software lab";
    <contact> "Michael Landon";
    <interface> "fxp0.0";
    <community> {
        <name> "public";
        <authorization> "read-only";
        <clients> {
            <name> "0.0.0.0/0";
            <restrict>;
        }
        <clients> {
            <name> "192.168.1.252/32";

```

```
}
<clients> {
  <name> "10.197.169.222/32";
}
<clients> {
  <name> "10.197.169.188/32";
}
<clients> {
  <name> "10.197.169.193/32";
}
<clients> {
  <name> "192.168.65.46/32";
}
<clients> {
  <name> "10.209.152.0/23";
}
}
<community> {
  <name> "private";
  <authorization> "read-write";
  <clients> {
    <name> "0.0.0.0/0";
    <restrict>;
  }
  <clients> {
    <name> "10.197.169.188/32";
  }
}
}
<routing-options> {
  <static> {
    <junos:comment> "/* safety precaution */";
    <route> {
      <name> "0.0.0.0/0";
      <discard>;
      <retain>;
      <no-readvertise>;
    }
    <junos:comment> "/* corporate net */";
    <route> {
      <name> "172.16.0.0/12";
      <next-hop> $bkup-rtr;
      <retain>;
      <no-readvertise>;
    }
    <junos:comment> "/* lab nets */";
    <route> {
      <name> "192.168.0.0/16";
      <next-hop> $bkup-rtr;
      <retain>;
      <no-readvertise>;
    }
    <junos:comment> "/* reflector */";
    <route> {
      <name> "10.17.136.192/32";
      <next-hop> $bkup-rtr;
```



```

    <retain>;
    <no-readvertise>;
  }
  <junos:comment> "/* another lab1*/";
  <route> {
    <name> "10.10.0.0/16";
    <next-hop> $bkup-rtr;
    <retain>;
    <no-readvertise>;
  }
  <junos:comment> "/* ssh servers */";
  <route> {
    <name> "10.17.136.0/24";
    <next-hop> $bkup-rtr;
    <retain>;
    <no-readvertise>;
  }
  <junos:comment> "/* Workstations */";
  <route> {
    <name> "10.150.0.0/16";
    <next-hop> $bkup-rtr;
    <retain>;
    <no-readvertise>;
  }
  <junos:comment> "/* Hosts */";
  <route> {
    <name> "10.157.64.0/19";
    <next-hop> $bkup-rtr;
    <retain>;
    <no-readvertise>;
  }
  <junos:comment> "/* Build Servers */";
  <route> {
    <name> "10.10.0.0/16";
    <next-hop> $bkup-rtr;
    <retain>;
    <no-readvertise>;
  }
}
}
<policy-options> {
  <policy-statement> {
    <name> "redist";
    <from> {
      <protocol> "static";
    }
    <then> {
      <accept>;
    }
  }
}
}
<apply-groups> "re0";
<apply-groups> "re1";
<groups> {
  <name> "re0";
  <system> {

```

```

    }
    <interfaces> {
        <interface> {
            <name> "fxp0";
            <unit> {
                <name> "0";
                <family> {
                    <inet> {
                        <address> {
                            <name> $fxp0-addr;
                        }
                    }
                }
            }
        }
    }
    <groups> {
        <name> "re1";
    }
    <interfaces> {
        <interface> {
            <name> "lo0";
            <unit> {
                <name> "0";
                <family> {
                    <inet> {
                        <address> {
                            <name> $rid;
                        }
                    }
                }
            }
        }
    }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **config-system.xml** or **config-system.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **config-system.slax**.

```
system {
  scripts {
```

```

        commit {
            allow-transients;
            file config-system.xml;
        }
    }
    apply-macro config-system {
        host-name test;
        mgmt-address 10.0.0.1/32;
        backup-router 10.0.0.2;
    }

```

The **host-name** and **mgmt-address** statements are mandatory. The **backup-router** statement is optional. You can substitute a hostname, a management Ethernet (fxp0) IP address, and a backup router IP address that are appropriate for your device.

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying the Configuration

**Purpose** Verify that the script behaves as expected.

**Action** After committing the configuration, issue the **show | display commit-scripts** configuration mode command to view the device base configuration.

```

user@host# show | display commit-scripts
...

```

## Example: Prepending a Global Policy

This commit script example ensures that a BGP global import policy is applied to all your BGP imports before any other import policies are applied.

- [Requirements on page 412](#)
- [Overview and Commit Script on page 412](#)

- [Configuration on page 414](#)
- [Verification on page 415](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Commit Script

For most configuration objects, the order in which the object or its children is created is not significant, because the Junos OS configuration management software stores and displays configuration objects in predetermined positions in the configuration hierarchy. However, some configuration objects—such as routing policies and firewall filters—consist of elements that must be processed and analyzed sequentially in order to produce the intended routing behavior.

This example commit script ensures that a BGP global import policy is applied to all your BGP imports before any other import policies are applied.

This example automatically prepends the **bgp\_global\_import** policy in front of any other BGP import policies. If the **bgp\_global\_import** policy statement is not included in the configuration, an error message is generated, and the commit operation fails.

Otherwise, the commit script uses the **insert="before"** Junos XML protocol attribute and the **position()** XSLT function to control the position of the global BGP policy in relation to any other applied policies. The **insert="before"** attribute inserts the **bgp\_global\_import** policy in front of the first preexisting BGP import policy.

If there is no preexisting default BGP import policy, the global policy is included in the configuration.

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:if
      test="not(policy-options/policy-statement[name='bgp_global_import'])">
      <xnm:error>
        <message>Policy error: Policy bgp_global_import
required</message>
      </xnm:error>
    </xsl:if>
    <xsl:for-each select="protocols/bgp | protocols/bgp/group |
      protocols/bgp/group/neighbor">
      <xsl:variable name="first" select="import[position() = 1]"/>
      <xsl:if test="$first">
        <xsl:call-template name="jcs:emit-change">
          <xsl:with-param name="tag" select="'transient-change'"/>
          <xsl:with-param name="content">
```

```

        <import insert="before"
            name="{ $first }">bgp_global_import</import>
        </xsl:with-param>
    </xsl:call-template>
</xsl:if>
</xsl:for-each>
<xsl:for-each select="protocols/bgp">
    <xsl:if test="not(import)">
        <xsl:call-template name="jcs:emit-change">
            <xsl:with-param name="tag" select="'transient-change'"/>
            <xsl:with-param name="content">
                <import>bgp_global_import</import>
            </xsl:with-param>
        </xsl:call-template>
    </xsl:if>
</xsl:for-each>
</xsl:template>
</xsl:stylesheet>

```

## SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
    if (not(policy-options/policy-statement[name='bgp_global_import'])) {
        <xnm:error> {
            <message> "Policy error: Policy bgp_global_import required";
        }
    }
    for-each (protocols/bgp | protocols/bgp/group |
        protocols/bgp/group/neighbor) {
        var $first = import[position() = 1];
        if ($first) {
            call jcs:emit-change($tag = 'transient-change') {
                with $content = {
                    <import insert="before" name="{ $first }">
"bgp_global_import";
                }
            }
        }
    }
    for-each (protocols/bgp) {
        if (not(import)) {
            call jcs:emit-change($tag = 'transient-change') {
                with $content = {
                    <import> "bgp_global_import";
                }
            }
        }
    }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **bgp-global-import.xsl** or **bgp-global-import.slax** as appropriate, and copy it to the `/var/db/scripts/commit/` directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the `[edit system scripts commit file]` hierarchy level to **bgp-global-import.slax**.

```
system {
  scripts {
    commit {
      allow-transients;
      file bgp-global-import.xsl;
    }
  }
}
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.16.2/24;
      }
      family inet6 {
        address 2002:18a5:e996:beef::2/64;
      }
    }
  }
}
routing-options {
  autonomous-system 64500;
}
protocols {
  bgp {
    group fish {
      neighbor 192.168.16.4 {
        import [ blue green ];
        peer-as 64501;
      }
      neighbor 192.168.16.6 {
        peer-as 64502;
      }
    }
  }
}
policy-options {
  policy-statement blue {
    from protocol bgp;
    then accept;
  }
}
```

```

    policy-statement green {
        then accept;
    }
    policy-statement bgp_global_import {
        then accept;
    }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying the Configuration

**Purpose** Verify that the script behaves as expected.

**Action** When you issue the **show protocols** configuration mode command, the **bgp\_global\_import** import policy is not displayed, because it is added as a transient change.

```

[edit]
user@host# show protocols
bgp {
  group fish {
    neighbor 192.168.16.4 {
      import [ blue green ];
      peer-as 64501;
    }
    neighbor 192.168.16.6 {
      peer-as 64502;
    }
  }
}

```

The commit script adds the **import bgp\_global\_import** statement at the **[edit protocols bgp]** hierarchy level and prepends the **bgp\_global\_import** policy to the 192.168.16.4 neighbor policy chain. Issue the **show protocols | display commit-scripts** to view all configuration statements including transient changes.

```

[edit]
user@host# show protocols | display commit-scripts

```

```
bgp {
  import bgp_global_import;
  group fish {
    neighbor 192.168.16.4 {
      import [ bgp_global_import blue green ];
      peer-as 64501;
    }
    neighbor 192.168.16.6 {
      peer-as 64502;
    }
  }
}
```

After you add a policy to the 192.168.16.6 neighbor, which previously had no policies applied, the **bgp\_global\_import** policy is prepended. Issue the **show protocols | display commit-scripts** command to view all configuration statements including transient changes.

```
[edit]
user@host# set protocols bgp group fish neighbor 192.168.16.6 import green
```

```
[edit]
user@host# show protocols | display commit-scripts
bgp {
  import bgp_global_import;
  group fish {
    neighbor 192.168.16.4 {
      import [ bgp_global_import blue green ];
      peer-as 64501;
    }
    neighbor 192.168.16.6 {
      import [ bgp_global_import green ];
      peer-as 64502;
    }
  }
}
```

---

## Example: Preventing Import of the Full Routing Table

In the Junos OS routing policy, if you configure a policy with no match conditions and a terminating action of **then accept**, and then apply the policy to a routing protocol, the protocol imports the entire routing table. This example shows how to use a commit script to prevent this scenario.

- [Requirements on page 416](#)
- [Overview and Commit Script on page 417](#)
- [Configuration on page 418](#)
- [Verification on page 419](#)

### Requirements

This example uses a device running Junos OS.



## Overview and Commit Script

This example inspects the **import** statements configured at the **[edit protocols ospf]** and **[edit protocols isis]** hierarchy levels to determine if any of the named policies contain a **then accept** term with no match conditions. The script protects against importing the full routing table into these interior gateway protocols (IGPs).

The example script is shown in both XSLT and SLAX syntax:

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:param name="po"
    select="commit-script-input/configuration/policy-options"/>
  <xsl:template match="configuration">
    <xsl:apply-templates select="protocols/ospf/import"/>
    <xsl:apply-templates select="protocols/isis/import"/>
  </xsl:template>
  <xsl:template match="import">
    <xsl:param name="test" select="."/>
    <xsl:for-each select="$po/policy-statement[name=$test]">
      <xsl:choose>
        <xsl:when test="then/accept and not(to) and not(from)">
          <xnm:error>
            <xsl:call-template name="jcs:edit-path">
              <xsl:with-param name="dot" select="$test"/>
            </xsl:call-template>
            <xsl:call-template name="jcs:statement">
              <xsl:with-param name="dot" select="$test"/>
            </xsl:call-template>
            <message>policy contains bare 'then accept'</message>
          </xnm:error>
        </xsl:when>
      </xsl:choose>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>
```

### SLAX Syntax

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

param $po = commit-script-input/configuration/policy-options;
match configuration {
  apply-templates protocols/ospf/import;
  apply-templates protocols/isis/import;
}
match import {
  param $test = .;
  for-each ($po/policy-statement[name=$test]) {
    if (then/accept and not(to) and not(from)) {
      <xnm:error> {
```

```
        call jcs:edit-path($dot = $test);
        call jcs:statement($dot = $test);
        <message> "policy contains bare 'then accept'";
    }
}
}
```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **import.xml** or **import.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **import.slax**.

```
system {
  scripts {
    commit {
      file import.xml;
    }
  }
}
protocols {
  ospf {
    import bad-news;
  }
}
policy-options {
  policy-statement bad-news {
    then accept;
  }
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying the Commit Script Execution

- Purpose** Verify that the script behaves as expected.
- Action** Review the output of the **commit** command. The sample configuration configures an **import** statement at the **[edit protocols ospf]** hierarchy level. Because the policy contains a **then accept** term with no match conditions, the script generates an error, and the commit operation fails. The following output appears after issuing a **commit** command:

```
[edit]
user@host# commit
[edit protocols ospf]
  'import bad-news;'
    policy contains bare 'then accept'
error: 1 error reported by commit scripts
error: commit script failure
```

### Example: Requiring Internal Clocking on T1 Interfaces

This example shows how to use a commit script to require that T1 interfaces be configured with internal clocking.

- [Requirements on page 419](#)
- [Overview and Commit Script on page 419](#)
- [Configuration on page 420](#)
- [Verification on page 421](#)

## Requirements

This example uses a device running Junos OS with T1 interfaces.

### Overview and Commit Script

This commit script ensures that T1 interfaces are explicitly configured to use internal clocking. If the **clocking** statement is not included in the configuration, or if the **clocking external** statement is included, an error message is generated, and the configuration is not committed.

The example script is shown in both XSLT and SLAX syntax:

**XSLT Syntax**

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:for-each select="interfaces/interface[starts-with(name, 't1-')]">

      <xsl:variable name="clock-source">
```

```

        <xsl:value-of select="clocking"/>
      </xsl:variable>
      <xsl:if test="not($clock-source = 'internal')">
      <!-- or xsl:if test="$clock-source != 'internal'" -->
        <xnm:error>
          <xsl:call-template name="jcs:edit-path"/>
          <xsl:call-template name="jcs:statement">
            <xsl:with-param name="dot" select="clocking"/>
          </xsl:call-template>
          <message>
            This T1 interface should have internal clocking.
          </message>
        </xnm:error>
      </xsl:if>
    </xsl:for-each>
  </xsl:template>
</xsl:stylesheet>

```

### SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
  for-each (interfaces/interface[starts-with(name, 't1-')]) {
    var $clock-source = {
      expr clocking;
    }
    if (not($clock-source = 'internal')) {
      <xnm:error> {
        call jcs:edit-path();
        call jcs:statement($dot = clocking);
        <message> "This T1 interface should have internal clocking.";
      }
    }
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the script into a text file, name the file **clocking-error.xsl** or **clocking-error.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **clocking-error.slax**.

```

system {
  scripts {
    commit {
      file clocking-error.xsl;
    }
  }
}

```

```

interfaces {
  t1-0/0/0 {
    clocking external;
  }
  t1-0/0/1 {
    unit 0;
  }
}

```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```

[edit]
user@host# load merge terminal
[Type ^D at a new line to end input]
... Paste the contents of the clipboard here ...

```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
  - b. Press Enter.
  - c. Press Ctrl+d.
4. Issue the **commit** command to commit the configuration.

```

user@host# commit

```

## Verification

### Verifying Commit Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The sample configuration stanzas configure two T1 interfaces t1-0/0/0 and t1-0/0/1. Interface t1-0/0/0 is configured with the **clocking external** statement, and interface t1-0/0/1 does not include any **clocking** statement. The script generates an error, and the commit operation fails. The following output appears after issuing a **commit** command:

```

[edit]
user@host# commit
[edit interfaces interface t1-0/0/0]
'clocking external;'
This T1 interface should have internal clocking.
[edit interfaces interface t1-0/0/1]
';'
This T1 interface should have internal clocking.
error: 2 errors reported by commit scripts
error: commit script failure

```

## Example: Requiring and Restricting Configuration Statements

Junos OS commit scripts enforce custom configuration rules. When a candidate configuration is committed, it is inspected by each active commit script. This example uses a commit script to specify required and prohibited configuration statements.

- [Requirements on page 422](#)
- [Overview and Commit Script on page 422](#)
- [Configuration on page 424](#)
- [Verification on page 425](#)

### Requirements

This example uses a device running Junos OS that has the Ethernet management interface fxp0.

### Overview and Commit Script

This example shows how to use a commit script to specify required and prohibited configuration statements. The following commit script ensures that the Ethernet management interface (fxp0) is configured and detects when the interface is improperly disabled. The script also detects when the **bgp** statement is not included at the **[edit protocols]** hierarchy level. In all cases, the script generates an error message, and the commit operation fails.

The example script is shown in both XSLT and SLAX syntax:

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:template match="configuration">
    <xsl:call-template name="error-if-missing">
      <xsl:with-param name="must"
        select="interfaces/interface[name='fxp0']/
          unit[name='0']/family/inet/address"/>
      <xsl:with-param name="statement"
        select="'interfaces fxp0 unit 0 family inet address'"/>
    </xsl:call-template>
    <xsl:call-template name="error-if-present">
      <xsl:with-param name="must"
        select="interfaces/interface[name='fxp0']/disable
          | interfaces/interface[name='fxp0']/
          unit[name='0']/disable"/>
      <xsl:with-param name="message">
        <xsl:text>The fxp0 interface is disabled.</xsl:text>
      </xsl:with-param>
    </xsl:call-template>
    <xsl:call-template name="error-if-missing">
      <xsl:with-param name="must" select="protocols/bgp"/>
      <xsl:with-param name="statement" select="'protocols bgp'"/>
    </xsl:call-template>
  </xsl:template>
```

```

        </xsl:call-template>
    </xsl:template>
    <xsl:template name="error-if-missing">
        <xsl:param name="must"/>
        <xsl:param name="statement" select="'unknown'"/>
        <xsl:param name="message"
            select="'missing mandatory configuration statement'"/>
        <xsl:if test="not($must)">
            <xnm:error>
                <edit-path><xsl:copy-of select="$statement"/></edit-path>
                <message><xsl:copy-of select="$message"/></message>
            </xnm:error>
        </xsl:if>
    </xsl:template>
    <xsl:template name="error-if-present">
        <xsl:param name="must" select="1"/> <!-- error if param missing -->
        <xsl:param name="message" select="'invalid configuration statement'"/>
        <xsl:for-each select="$must">
            <xnm:error>
                <xsl:call-template name="jcs:edit-path"/>
                <xsl:call-template name="jcs:statement"/>
                <message><xsl:copy-of select="$message"/></message>
            </xnm:error>
        </xsl:for-each>
    </xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";

match configuration {
    call error-if-missing($must =
        interfaces/interface[name='fxp0']/unit[name='0']/family/inet/address,
        $statement = 'interfaces fxp0 unit 0 family inet address');
    call error-if-present($must = interfaces/interface[name='fxp0']/disable |
        interfaces/interface[name='fxp0']/unit[name='0']/disable) {
        with $message = {
            expr "The fxp0 interface is disabled.";
        }
    }
    call error-if-missing($must = protocols/bgp, $statement = 'protocols bgp');
}
error-if-missing ($must, $statement = 'unknown', $message =
    'missing mandatory configuration statement') {
    if (not($must)) {
        <xnm:error> {
            <edit-path> {
                copy-of $statement;
            }
            <message> {
                copy-of $message;
            }
        }
    }
}
}

```

```
error-if-present ($must = 1, $message = 'invalid configuration statement') {  
  for-each ($must) {  
    <xnm:error> {  
      call jcs:edit-path();  
      call jcs:statement();  
      <message> {  
        copy-of $message;  
      }  
    }  
  }  
}
```

## Configuration

**Step-by-Step Procedure** To download, enable, and test the script:

1. Copy the script into a text file, name the file **no-nukes.xml** or **no-nukes.slax** as appropriate, and copy it to the **/var/db/scripts/commit/** directory on the device.
2. Select the following test configuration stanzas, and press Ctrl+c to copy them to the clipboard.

If you are using the SLAX version of the script, change the filename at the **[edit system scripts commit file]** hierarchy level to **no-nukes.slax**.

```
system {  
  scripts {  
    commit {  
      file no-nukes.xml;  
    }  
  }  
}  
interfaces {  
  fxp0 {  
    disable;  
    unit 0 {  
      family inet {  
        address 10.0.0.1/24;  
      }  
    }  
  }  
}
```

3. In configuration mode, issue the **load merge terminal** command to merge the stanzas into your device configuration.

```
[edit]  
user@host# load merge terminal  
[Type ^D at a new line to end input]  
... Paste the contents of the clipboard here ...
```

- a. At the prompt, paste the contents of the clipboard by using the mouse and the paste icon.
- b. Press Enter.
- c. Press Ctrl+d.



4. Issue the **commit** command to commit the configuration.

```
user@host# commit
```

## Verification

### Verifying Commit Script Execution

---

**Purpose** Verify that the script behaves as expected.

**Action** Review the output of the **commit** command. The script requires that the Ethernet management interface (fxp0) is configured and enabled and that the **bgp** statement is included at the **[edit protocols]** hierarchy level. The sample configuration stanzas include the fxp0 interface but disable it. In addition, the **bgp** statement is not configured at the **[edit protocols]** hierarchy level. When you run the script, it generates an error, and the commit operation fails. The following output appears after issuing a **commit** command:

```
[edit]
user@host# commit
[edit interfaces interface fxp0 disable]
'disable;'
The fxp0 interface is disabled.
protocols bgp
missing mandatory configuration statement
error: 2 errors reported by commit scripts
error: commit script failure
```



# Junos XML and XSLT Tag Elements Used in Commit Scripts

## <change> (SLAX and XSLT)

---

<b>Usage</b>	<pre>&lt;change&gt;   &lt;!-- tag elements representing configuration statements to load --&gt; &lt;/change&gt;</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	In Junos OS commit scripts, request that the Junos XML protocol server load configuration data into the candidate configuration by enclosing the configuration data within an opening <b>&lt;change&gt;</b> tag and closing <b>&lt;/change&gt;</b> tag. Inside the <b>&lt;change&gt;</b> element, include the configuration data as Junos XML tag elements.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">&lt;transient-change&gt; (SLAX and XSLT) on page 428</a></li> <li>• <a href="#">Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts on page 287</a></li> <li>• <a href="#">Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292</a></li> <li>• <a href="#">Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts on page 301</a></li> <li>• <a href="#">Creating Custom Configuration Syntax with Commit Script Macros on page 313</a></li> </ul>

## <syslog> (Junos XML)

---

<b>Usage</b>	<pre>&lt;syslog xmlns="namespace-URL" xmlns:xnm="namespace-URL"&gt;   &lt;message&gt;syslog-message &lt;/message&gt; &lt;/syslog&gt;</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Indicate that the commit script should record the given message in the system log file.

<b>Attributes</b>	<b>xmlns</b> —(Optional) Names the Extensible Markup Language (XML) namespace for the contents of the tag element. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code> , where <b>version</b> is a string such as 1.1.
	<b>xmlns:xnm</b> —(Optional) Names the XML namespace for child tag elements that have the <b>xnm:</b> prefix on their names. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code> , where <b>version</b> is a string such as 1.1.
<b>Contents</b>	<b>&lt;message&gt;</b> —Specifies the content of the system log message in a natural-language text string.
<b>Usage Guidelines</b>	See <a href="#">“Generating a Custom Warning, Error, or System Log Message in Commit Scripts”</a> on page 270.

---

## <transient-change> (SLAX and XSLT)

---

<b>Usage</b>	<pre>&lt;transient-change&gt;   &lt;!-- tag elements representing configuration statements to load --&gt; &lt;/transient-change&gt;</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	In Junos OS commit scripts, request that the Junos XML protocol server load configuration data into the checkout configuration by enclosing the configuration data within an opening <b>&lt;transient-change&gt;</b> and closing <b>&lt;/transient-change&gt;</b> tag. Inside the <b>&lt;transient-change&gt;</b> element, include the configuration data as Junos XML tag elements.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">&lt;change&gt; (SLAX and XSLT)</a> on page 427</li><li>• <a href="#">Overview of Generating Persistent or Transient Configuration Changes Using Commit Scripts</a> on page 287</li><li>• <a href="#">Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts</a> on page 292</li><li>• <a href="#">Example: Generating Persistent and Transient Configuration Changes Using Commit Scripts</a> on page 301</li><li>• <a href="#">Creating Custom Configuration Syntax with Commit Script Macros</a> on page 313</li></ul>

---

## xnm:error (Junos XML)

---

<b>Usage</b>	<pre>&lt;xnm:error xmlns="namespace-URL" xmlns:xnm="namespace-URL"&gt;   &lt;parse/&gt;   &lt;source-daemon&gt;module-name&lt;/source-daemon&gt;   &lt;filename&gt;filename&lt;/filename&gt;   &lt;line-number&gt;line-number&lt;/line-number&gt;   &lt;column&gt;column-number&lt;/column&gt;   &lt;token&gt;input-token-id&lt;/token&gt;   &lt;edit-path&gt;edit-path-name&lt;/edit-path&gt;   &lt;statement&gt;statement-string&lt;/statement&gt;   &lt;message&gt;error-string&lt;/message&gt;</pre>
--------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

```

    <re-name>re-name-string</re-name>
    <database-status-information>user</database-status-information>
    <reason>reason-string</reason>
  </xnm:error>

```

<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Indicate that the commit script has detected an error in the configuration and has caused the commit operation to fail. The child tag elements described in the Contents section detail the nature of the error.
<b>Attributes</b>	<p><b>xm:ns</b>—Names the XML namespace for the contents of the tag element. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code>, where <i>version</i> is a string such as 1.1.</p> <p><b>xm:ns:xnm</b>—Names the XML namespace for child tag elements that have the <b>xnm:</b> prefix on their names. The value is a URL of the form <code>http://xml.juniper.net/xnm/version/xnm</code>, where <i>version</i> is a string such as 1.1.</p>
<b>Contents</b>	<p><b>&lt;column&gt;</b>—Identifies the element that caused the error by specifying its position as the number of characters after the first character in the line specified by the <b>&lt;line-number&gt;</b> tag element in the configuration file that was being loaded (which is named in the <b>&lt;filename&gt;</b> tag element).</p> <p><b>&lt;database-status-information&gt;</b>—Provides information about the users currently editing the configuration.</p> <p><b>&lt;edit-path&gt;</b>—Specifies the command-line interface (CLI) configuration mode edit path in effect when the error occurred (provided only during loading of a configuration file).</p> <p><b>&lt;filename&gt;</b>—Names the configuration file that was being loaded.</p> <p><b>&lt;line-number&gt;</b>—Specifies the line number where the error occurred in the configuration file that was being loaded, which is named by the <b>&lt;filename&gt;</b> tag element.</p> <p><b>&lt;message&gt;</b>—Describes the error in a natural-language text string.</p> <p><b>&lt;parse/&gt;</b>—Indicates that there was a syntactic error in the request submitted by the client application.</p> <p><b>&lt;re-name&gt;</b>—Names the Routing Engine on which the <b>&lt;source-daemon&gt;</b> is running.</p> <p><b>&lt;reason&gt;</b>—Describes the reason for the error.</p> <p><b>&lt;source-daemon&gt;</b>—Names the Junos OS module that was processing the request in which the error occurred.</p> <p><b>&lt;statement&gt;</b>—Specifies the configuration statement in effect when the problem occurred.</p> <p><b>&lt;token&gt;</b>—Names the element in the request that caused the error.</p>

**Usage Guidelines** See [“Generating a Custom Warning, Error, or System Log Message in Commit Scripts”](#) on page 270.

**Related Documentation**

- [xnm:warning \(Junos XML\)](#) on page 430

---

## xnm:warning (Junos XML)

---

**Usage**

```
<xnm:warning xmlns="namespace-URL" xmlns:xnm="namespace-URL">
  <source-daemon>module-name</source-daemon>
  <filename>filename</filename>
  <line-number>line-number</line-number>
  <column>column-number</column>
  <token>input-token-id</token>
  <edit-path>edit-path-name</edit-path>
  <statement>statement-name</statement>
  <message>error-string</message>
  <reason>reason-string</reason>
</xnm:warning>
```

**Release Information** Statement introduced in Junos OS Release 7.4.

**Description** Indicate that the commit script has encountered a problem with the configuration and pass a warning message to the Junos OS CLI or Junos XML protocol client application. The child tag elements described in the Contents section detail the nature of the warning.

**Attributes** **xmlns**—Names the XML namespace for the contents of the tag element. The value is a URL of the form <http://xml.juniper.net/xnm/version/xnm>, where *version* is a string such as 1.1.

**xmlns:xnm**—Names the XML namespace for child tag elements that have the **xnm:** prefix on their names. The value is a URL of the form <http://xml.juniper.net/xnm/version/xnm>, where *version* is a string such as 1.1.

**Contents** **<column>**—Identifies the element that caused the warning by specifying its position as the number of characters after the first character in the line specified by the **<line-number>** tag element in the configuration file that was being loaded (which is named in the **<filename>** tag element).

**<edit-path>**—Specifies the CLI configuration mode edit path in effect when the problem occurred (provided only during loading of a configuration file).

**<filename>**—Names the configuration file that was being loaded.

**<line-number>**—Specifies the line number where the problem occurred in the configuration file that was being loaded, which is named by the **<filename>** tag element.

**<message>**—Describes the warning in a natural-language text string.

**<reason>**—Describes the reason for the warning.

**<source-daemon>**—Names the Junos OS module that was processing the request in which the problem occurred.

**<statement>**—Names the configuration statement in effect when the problem occurred.

**<token>**—Names which element in the request caused the warning.

**Usage Guidelines** See [“Generating a Custom Warning, Error, or System Log Message in Commit Scripts” on page 270](#)

**Related Documentation**

- [xnm:error \(Junos XML\) on page 428.](#)





# Troubleshooting Commit Scripts

- [Displaying Commit Script Output on page 433](#)
- [Tracing Commit Script Processing on page 434](#)
- [Troubleshooting Commit Scripts on page 438](#)

## Displaying Commit Script Output

Table 37 on page 433 summarizes the Junos OS command-line interface (CLI) commands you can use to monitor and troubleshoot commit scripts. For more information about the `cscript.log` file, see [“Tracing Commit Script Processing” on page 434](#).



**NOTE:** Tracing commit script processing, including the `cscript.log` file, is not supported on the QFX3000-G QFabric system.

**Table 37: Commit Script Configuration and Operational Mode Commands**

Task	Command
<b>Configuration Mode Commands</b>	
Display errors and warnings generated by commit scripts.	<code>commit</code> or <code>commit check</code>
Display detailed information about the commit operation and commit script execution.	<code>commit   display detail</code>
Display the underlying Extensible Markup Language (XML) data.	<code>commit   display xml</code>
Display the postinheritance contents of the configuration database. This view includes transient changes, but does not include changes made in configuration groups.	<code>show   display commit-scripts</code>
Display the postinheritance contents of the configuration database. This view excludes transient changes.	<code>show   display commit-scripts no-transients</code>
Display the postinheritance configuration in XML format. This is the configuration format that each commit script receives as input.	<code>show   display commit-scripts view</code>
Viewing the configuration in XML format can be helpful when you are writing XML Path Language (XPath) expressions and configuration element tags.	

Table 37: Commit Script Configuration and Operational Mode Commands (*continued*)

Task	Command
Display the postinheritance configuration in XML format, but exclude transient changes.	<code>show   display commit-scripts view   display commit-scripts no-transients</code>
Display all configuration groups data, including script-generated changes to the groups.	<code>show groups   display commit-scripts</code>
Display a particular configuration group, including script-generated changes to the group.	<code>show groups <i>group-name</i>   display commit-scripts</code>
<b>Operational Mode Commands</b>	
Display logging data associated with all commit script processing.	<code>show log cscript.log</code>
Display processing for only the most recent commit operation.	<code>show log cscript.log   last</code>
Display processing for script errors.	<code>show log cscript.log   match error</code>
Display processing for a particular script.	<code>show log cscript.log   match <i>filename</i></code>

**Related Documentation**

- [Tracing Commit Script Processing on page 434](#)

## Tracing Commit Script Processing

Commit script tracing operations track all commit script operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

The default operation of commit script tracing is to log important events in a file called **cscript.log** located in the **/var/log** directory on the device. When the file **cscript.log** reaches 128 kilobytes (KB), it is renamed with a number 0 through 9 (in ascending order) appended to the end of the file and then compressed. For example, the log file is saved as **cscript.log.0.gz**, then **cscript.log.1.gz** until there are 10 trace files. Then the oldest trace file (**cscript.log.9.gz**) is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)

This section discusses the following topics:

- [Minimum Configuration for Tracing for Commit Script Operations on page 434](#)
- [Configuring Tracing of Commit Scripts on page 436](#)

### Minimum Configuration for Tracing for Commit Script Operations

If no commit script trace options are configured, the simplest way to view the trace output of a commit script is to configure the **output** trace flag and issue the **show log cscript.log | last** command. To do this, perform the following steps:

- 1. If you have not done so already, enable a commit script by including the **file** statement at the **[edit system scripts commit]** hierarchy level:

```
[edit system scripts commit]
user@host# set file filename
```

- 2. Enable trace options by including the **traceoptions flag output** statement at the **[edit system scripts commit]** hierarchy level:

```
[edit system scripts commit]
user@host# set traceoptions flag output
```

- 3. Issue the **commit** command:

```
[edit]
user@host# commit
```

- 4. Display the resulting trace messages recorded in the file **/var/log/cscript.log**. At the end of the log is the output generated by the commit script you enabled in Step 1. To display the end of the log, issue the **show log cscript.log | last** operational mode command:

```
[edit]
user@host# run show log cscript.log | last
```

Table 38 on page 435 summarizes useful filtering commands that display selected portions of the **cscript.log** file.

Table 38: Commit Script Tracing Operational Mode Commands

Task	Command
Display logging data associated with all script processing.	<b>show log cscript.log</b>
Display script processing for only the most recent commit operation.	<b>show log cscript.log   last</b>
Display processing for script errors.	<b>show log cscript.log   match error</b>
Display script processing for a particular script.	<b>show log cscript.log   match filename</b>

Example: Minimum Configuration for Enabling Traceoptions for Commit Scripts

Display the trace output for the commit script file **source-route.xml**:

```
[edit]
system {
  scripts {
    commit {
      file source-route.xml;
      traceoptions flag output;
    }
  }
}
```

```
}  
[edit]  
user@host# commit  
[edit]  
user@host# run show log cscript.log | last  
Jun 20 10:21:24 summary: changes 0, transients 0 (allowed), syslog 0  
Jun 20 10:24:15 commit script processing begins  
Jun 20 10:24:15 reading commit script configuration  
Jun 20 10:24:15 testing commit script configuration  
Jun 20 10:24:15 opening commit script '/var/db/scripts/commit/source-route.xml'  
Jun 20 10:24:15 script file '/var/db/scripts/commit/source-route.xml': size=699;  
md5 = d947972b429d17ce97fe987d94add6fd  
Jun 20 10:24:15 reading commit script 'source-route.xml'  
Jun 20 10:24:15 running commit script 'source-route.xml'  
Jun 20 10:24:15 processing commit script 'source-route.xml'  
Jun 20 10:24:15 results of 'source-route.xml'  
Jun 20 10:24:15 begin dump  
<commit-script-output xmlns:junos="http://xml.juniper.net/junos/*/junos"  
xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"  
xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">  
  <xnm:warning>  
    <edit-path>[edit chassis]</edit-path>  
    <message>IP source-route processing is not enabled.</message>  
  </xnm:warning>  
</commit-script-output>Jun 20 10:24:15 end dump  
Jun 20 10:24:15 no errors from source-route.xml  
Jun 20 10:24:15 saving commit script changes  
Jun 20 10:24:15 summary: changes 0, transients 0 (allowed), syslog 0
```

## Configuring Tracing of Commit Scripts

You cannot change the directory (**/var/log**) to which trace files are written. However, you can customize other trace file settings by including the following statements at the **[edit system scripts commit traceoptions]** hierarchy level:

```
[edit system scripts commit traceoptions]  
file <filename> <files number> <size size> <world-readable | no-world-readable>;  
flag all;  
flag events;  
flag input;  
flag offline;  
flag output;  
flag rpc;  
flag xslt;  
no-remote-trace;
```

These statements are described in the following sections:

- [Configuring the Commit Script Log Filename on page 437](#)
- [Configuring the Number and Size of Commit Script Log Files on page 437](#)
- [Configuring Access to Commit Script Log Files on page 437](#)
- [Configuring the Commit Script Trace Operations on page 438](#)

### Configuring the Commit Script Log Filename

By default, the name of the file that records trace output is **cscript.log**. You can specify a different name by including the **file** statement at the **[edit system scripts commit traceoptions]** hierarchy level:

```
[edit system scripts commit traceoptions]
file filename;
```

### Configuring the Number and Size of Commit Script Log Files

By default, when the trace file reaches 128 KB in size, it is renamed and compressed to **filename.0.gz**, then **filename.1.gz**, and so on, until there are 10 trace files. Then the oldest trace file (**filename.9.gz**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit system scripts commit traceoptions file <filename>]** hierarchy level:

```
[edit system scripts commit traceoptions file <filename>]
files number size size;
```

For example, set the maximum file size to 640 KB and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 640 KB, it is renamed and compressed to **filename.0.gz**, and a new file called **filename** is created. When **filename** reaches 640 KB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is renamed and compressed to **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19.gz**) is overwritten.

The number of files can range from 2 through 1000 files. The file size can range from 10 KB through 1 gigabyte (GB).



#### NOTE:

If you set either a maximum file size or a maximum number of trace files, you also must specify the other parameter and a filename.

### Configuring Access to Commit Script Log Files

By default, access to the commit script log file is restricted to the owner. You can manually configure access by including the **world-readable** or **no-world-readable** statement at the **[edit system scripts commit traceoptions file <filename>]** hierarchy level.

```
[edit system scripts commit traceoptions file <filename>]
(world-readable | no-world-readable);
```

The **no-world-readable** statement restricts commit script log access to the owner. The **world-readable** statement enables unrestricted access to the commit script log file.

## Configuring the Commit Script Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit system scripts commit traceoptions]** hierarchy level:

```
[edit system scripts commit traceoptions]
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
flag xslt;
```

Table 39 on page 438 describes the meaning of the commit script tracing flags.

Table 39: Commit Script Tracing Flags

Flag	Description	Default Setting
all	Trace all commit script operations.	Off
events	Trace important commit script events.	On
input	Trace commit script input data.	Off
offline	Generate data for offline development.	Off
output	Trace commit script output data.	Off
rpc	Trace commit script RPCs.	Off
xslt	Trace the Extensible Stylesheet Language Transformations (XSLT) library.	Off

## Troubleshooting Commit Scripts

After you enable a commit script and issue a **commit** command, the commit script takes effect immediately.

Table 40 on page 438 describes some common problems that might occur.

Table 40: Troubleshooting Commit Scripts

Problem	Solution
The output of the <b>commit check   display detail</b> command does not reference the expected commit scripts.	Make sure you have enabled all the scripts by including the <b>file</b> statement for each one at the <b>[edit system scripts commit]</b> hierarchy level.

Table 40: Troubleshooting Commit Scripts (*continued*)

Problem	Solution
<p>The output contains the error message:</p> <pre>error: could not open commit script: /var/db/scripts/commit/<i>filename</i>: No such file or directory</pre>	<p>Make sure the commit script is present in the <code>/var/db/scripts/commit/</code> directory on your switch, router, or security device.</p>
<p>The following error and warning messages appear:</p> <pre>error: invalid transient change generated by commit script: <i>filename</i> warning: 1 transient change was generated without [system scripts commit allow-transients]</pre>	<p>One of your commit scripts contains instructions to generate a transient change, but you have not enabled transient changes.</p> <p>To rectify this problem, take one of the following actions:</p> <ul style="list-style-type: none"> <li>• Include the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level.</li> <li>• Remove the code that generates a transient change from the indicated script.</li> <li>• Disable the script in the configuration.</li> </ul>
<p>An expected action does not occur.</p> <p>For example, a warning message does not appear even though the configuration contains the problem that is supposed to evoke the warning message.</p>	<ol style="list-style-type: none"> <li>1. Make sure you have enabled the script. Scripts are ignored if they are not enabled. To enable a script, include the <b>file <i>filename</i></b> statement at the <b>[edit system scripts commit]</b> hierarchy level.</li> <li>2. Make sure you have included the required boilerplate in your script. For more information, see <a href="#">“Required Boilerplate for Commit Scripts” on page 251</a>.</li> <li>3. Make sure that the Extensible Markup Language Path (XPath) expressions in the script contain valid Junos OS command-line interface (CLI) statements expressed as Junos XML protocol tag elements.  You can verify the XML hierarchy by checking the <i>Junos XML API Configuration Developer Reference</i> or by issuing the <b>show configuration   display xml</b> operational mode command.</li> <li>4. Make sure that the programming instructions in the script are referencing the correct context node.  If you nest one instruction inside another, the outer instruction changes the context node, so the inner instruction must be relative to the outer.  For example, the <b>&lt;xsl:for-each&gt;</b> instruction contains an XPath expression, which changes the context node. So the nested <b>&lt;xsl:if&gt;</b> instruction uses an XPath expression that is relative to the <b>interfaces/interface[starts-with(name, 't1-')]</b> XPath expression.   <pre>&lt;xsl:for-each select="interfaces/   interface[starts-with(name, 't1-')]"&gt; &lt;xsl:if test="not(description)"&gt;</pre> </li> </ol>





## PART 8

# Op Scripts

- [Op Scripts Overview on page 443](#)
- [Creating and Executing Op Scripts on page 445](#)
- [Op Script Examples on page 467](#)
- [Provisioning Services Using Service Template Automation on page 515](#)
- [Troubleshooting Op Scripts on page 529](#)



## CHAPTER 23

# Op Scripts Overview

- [Op Script Overview on page 443](#)
- [How Op Scripts Work on page 444](#)

## Op Script Overview

---

Junos OS operation (op) scripts automate network and device management and troubleshooting. Op scripts can perform any function available through the remote procedure calls (RPCs) supported by either the Junos XML management protocol or the Junos Extensible Markup Language (XML) API. Op scripts can be executed manually in the CLI or upon user login, or they can be called from another script. They are executed by the Junos OS management process (mgd).

Op scripts enable you to:

- Create custom operational mode commands
- Execute a series of operational mode commands
- Customize the output of operational mode commands
- Shorten troubleshooting time by gathering operational information and iteratively narrowing down the cause of a network problem
- Perform controlled configuration changes
- Monitor the overall status of a device by creating a general operation script that periodically checks network warning parameters, such as high CPU usage.

Op scripts are based on the Junos XML management protocol and the Junos XML API. The Junos XML management protocol is an XML based RPC mechanism, and the Junos XML API is an XML representation of Junos configuration statements and operational mode commands. For additional information, see [“Junos XML Management Protocol and Junos XML API Overview” on page 11](#).

Op scripts can be written in Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX). Op scripts use XML Path Language (XPath) to locate the operational objects to be inspected and automation script constructs to specify the actions to perform on the operational objects. The actions can change the configuration or output or execute additional commands based on the output.

- Related Documentation**
- [Junos XML Management Protocol and Junos XML API Overview on page 11](#)
  - [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
  - [SLAX Overview on page 57](#)
  - [XSLT Overview on page 17](#)
  - [XPath Overview on page 19](#)

---

## How Op Scripts Work

Op scripts execute Junos OS operational commands and inspect the resulting output. After inspection, op scripts can manipulate the output or automatically correct errors within the device running Junos OS based on this output.

You enable op scripts by listing the filenames of one or more op script files within the **[edit system scripts op]** hierarchy level. To execute local op scripts, you must add the files to the appropriate op script directory on the device. For more information about op script file directories, see [“Storing and Enabling Scripts” on page 661](#). Once added to the device, op scripts are invoked from the command line, using the **op filename** command.

You can also store and execute op scripts from a remote site. Remote op scripts are invoked from the command line using the **op url url** command. For more information about executing remote op scripts, see [“Executing an Op Script from a Remote Site” on page 462](#).

You can use op scripts to generate changes to the device configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

- Related Documentation**
- [Op Script Overview on page 443](#)
  - [Executing an Op Script on page 461](#)
  - [Executing an Op Script from a Remote Site on page 462](#)

# Creating and Executing Op Scripts

- [Required Boilerplate for Op Scripts on page 445](#)
- [Mapping Operational Mode Commands and Output Fields to Junos XML Notation on page 447](#)
- [Using RPCs and Operational Mode Commands in Op Scripts on page 448](#)
- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
- [Configuring Help Text for Op Scripts on page 457](#)
- [Defining Commands to Allow in an Op Script on page 458](#)
- [Enabling an Op Script and Defining a Script Alias on page 459](#)
- [Configuring Checksum Hashes for an Op Script on page 460](#)
- [Executing an Op Script on page 461](#)
- [Executing an Op Script from a Remote Site on page 462](#)
- [Disabling an Op Script on page 464](#)

## Required Boilerplate for Op Scripts

---

Junos OS op scripts can be written in Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python. Op scripts must include the necessary boilerplate required for that script language for both basic script functionality as well as any optional functionality used within the script such as the Junos OS extension functions and named templates. This topic provides standard boilerplate that can be used in XSLT, SLAX, and Python scripts.

SLAX and XSLT op scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

### XSLT Boilerplate for Op Scripts

The XSLT op script boilerplate is as follows:

```

1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
```

```

6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7      <xsl:import href="../import/junos.xsl"/>

8      <xsl:template match="/">
9          <op-script-results>
10             <!-- ... insert your code here ... -->
11             </op-script-results>
12         </xsl:template>
13         <!-- ... insert additional template definitions here ... -->
14     </xsl:stylesheet>

```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI), which marks this file as XML and specifies the version of XML as 1.0. The XML PI, if present, must be the first non-comment token in the script file.

```
1      <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2      <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in operation scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```

3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">

```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of Junos OS (in the file `/usr/libdata/cscript/import/junos.xsl`). The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in [“Understanding Named Templates in Junos OS Automation Scripts” on page 225](#).

```
7      <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `</>` element. The `<xsl:template match="/">` element is the root element and represents the top level of the XML hierarchy. All XPath expressions in the script must start at the top level. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see [“XPath Overview” on page 19](#).

```
8      <xsl:template match="/">
```

After the `<xsl:template match="/">` tag element, the `<op-script-results>` and `</op-script-results>` container tags must be the top-level child tags, as shown in Lines 9 and 10.

```

9          <op-script-results>
10             <!-- ... insert your code here ... -->
11             </op-script-results>

```

Line 11 closes the template.

```
11         </xsl:template>
```

Between Line 11 and Line 12, you can define additional XSLT templates that are called from within the `<xsl:template match="/">` template.

Line 12 closes the style sheet and the op script.

```
12    </xsl:stylesheet>
```

### SLAX Boilerplate for Op Scripts

The corresponding SLAX op script boilerplate is as follows:

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match / {
    <op-script-results> {
        /*
            * Insert your code here
        */
    }
}
```

### Python Boilerplate for Op Scripts

Python op scripts do not have a required boilerplate, but they must import any objects that are used in the script. Python automation scripts import the **Junos\_Context** dictionary when the script must access information about the script execution environment. Python automation scripts that include the **import jcs** statement can use Junos OS extension functions and Junos OS named template functionality in the script. If the script uses Junos PyEZ functionality, it must import the necessary classes from the **jnpr.junos** module.

```
from junos import Junos_Context
from jnpr.junos import Device
import jcs

if __name__ == '__main__':
```

Python automation scripts do not need to include an interpreter directive line (**#!/usr/bin/env python**) at the start of the script. However, the program will still execute correctly if one is present.

### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)

## Mapping Operational Mode Commands and Output Fields to Junos XML Notation

In SLAX and XSLT op scripts, you use tag elements from the Junos XML API to represent operational mode commands and output fields in the scripts. For the Junos XML equivalent of commands and output fields, consult the *Junos XML API Operational Developer Reference*.

You can also display the Junos XML tag elements for operational mode command output by directing the output from the command to the **| display xml** command:

```
user@host> command-string | display xml
```

For example:

```
user@host> show interfaces terse | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <interface-information
xmlns="http://xml.juniper.net/junos/10.0R10/junos-interface" junos:style="terse">
    <physical-interface>
      <name>dsc</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    </physical-interface>
    <physical-interface>
      <name>fxp0</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    <logical-interface>
      <name>fxp0.0</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    ...
```

**Related Documentation**

- [Configuring Help Text for Op Scripts on page 457](#)
- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
- [Using RPCs and Operational Mode Commands in Op Scripts on page 448](#)

---

## Using RPCs and Operational Mode Commands in Op Scripts

Most Junos OS operational mode commands have XML equivalents. Op scripts can execute these XML commands on a local or remote device using the remote procedure call (RPC) protocol. All operational mode commands that have XML equivalents are listed in the *Junos XML API Operational Developer Reference*.

Use of RPCs and operational mode commands in op scripts is discussed in more detail in the following sections:

- [Using RPCs in Op Scripts on page 448](#)
- [Displaying the RPC Tags for a Command on page 450](#)
- [Using Operational Mode Commands in Op Scripts on page 451](#)

### Using RPCs in Op Scripts

To use an RPC in a SLAX or XSLT op script, include the RPC in a variable declaration, and then invoke the RPC using the `jcs:invoke()` or `jcs:execute()` extension function with RPC variable as an argument. The `jcs:invoke()` function executes the RPC on the local device. The `jcs:execute()` function, in conjunction with a connection handle, executes the RPC on a remote device.

The following snippet, which invokes an RPC on the local device, is expanded and fully described in “[Example: Customizing Output of the show interfaces terse Command Using an Op Script](#)” on page 477:

**XSLT Syntax**

```
<xsl:variable name="rpc">
  <get-interface-information/> # Junos RPC for the show interfaces command
</xsl:variable>
<xsl:variable name="out" select="jcs:invoke($rpc)"/>
```



...

**SLAX Syntax**      `var $rpc = <get-interface-information>;`  
                          `var $out = jcs:invoke($rpc);`

The following snippet invokes the same RPC on a remote device:

**XSLT Syntax**      `<xsl:variable name="rpc">`  
                          `<get-interface-information/> # Junos RPC for the show interfaces command`  
                          `</xsl:variable>`  
                          `<xsl:variable name="connection" select="jcs:open('198.51.100.1', 'bsmith',`  
                          `'test123')"/>`  
                          `<xsl:variable name="out" select="jcs:execute($connection, $rpc)"/>`  
                          `<xsl:value-of select="jcs:close($connection)"/>`  
                          ...

**SLAX Syntax**      `var $rpc = <get-interface-information>;`  
                          `var $connection = jcs:open('198.51.100.1', 'bsmith', 'test123');`  
                          `var $out = jcs:execute($connection, $rpc);`  
                          `expr jcs:close($connection);`

In Python op scripts, RPCs are easy to execute using [Junos PyEZ](#) APIs. Each instance of the Junos PyEZ **Device** class has an **rpc** property that enables you to execute any RPC available through the Junos XML API. After establishing a session with a local or remote device, you can execute the RPC by appending the **rpc** property and RPC method name to the device instance. The return value is an XML object starting at the first element under the **<rpc-reply>** tag.

To execute the RPC on the local device, create the **Device** instance using an empty argument list. To execute the RPC on a remote device, create an instance of **Device** using the appropriate arguments to connect to that device.

The following snippet invokes an RPC on the local device:

**Python Syntax**      `jdev = Device()`  
                          `jdev.open()`  
                          `rsp = jdev.rpc.get_interface_information()`  
                          `jdev.close()`

The following snippet invokes the same RPC on a remote device:

**Python Syntax**      `jdev = Device(host='198.51.100.1', user='bsmith', passwd='test123')`  
                          `jdev.open()`  
                          `rsp = jdev.rpc.get_interface_information()`  
                          `jdev.close()`

To execute an RPC on a remote device, an SSH session must be established. In order for the script to establish the connection, you must either configure the SSH host key information for the remote device on the local device where the script will be executed, or the SSH host key information for the remote device must exist in the known hosts file of the user executing the script. For each remote device where an RPC is executed, configure the SSH host key information with one of the following methods:

- To configure SSH known hosts on the local device, include the **host** statement, and specify hostname and host key options for the remote device at the **[edit security ssh-known-hosts]** hierarchy level of the configuration.
- To manually retrieve SSH host key information, issue the **set security ssh-known-hosts fetch-from-server hostname** configuration mode command to instruct Junos OS to connect to the remote device and add the key.

```
user@host# set security ssh-known-hosts fetch-from-server router2
The authenticity of host 'router2 (198.51.100.1)' can't be established.
RSA key fingerprint is 30:18:99:7a:3c:ed:40:04:0f:fd:c1:57:7e:6b:f3:90.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'router2,198.51.100.1' (RSA) to the list of known
hosts.
```

- To manually import SSH host key information from a file, use the **set security ssh-known-hosts load-key-file filename** configuration mode command and specify the known-hosts file.

```
user@host# set security ssh-known-hosts load-key-file /var/tmp/known_hosts
Import SSH host keys from trusted source /var/tmp/known_hosts ? [yes,no] (no)
yes
```

- Alternatively, the user executing the script can log in to the local device, SSH to the remote device, and then manually accept the host key, which is added to that user's known hosts file. In the following example, root is logged in to router1. In order to execute a remote RPC on router2, root adds the host key of router2 by issuing the **ssh router2** operational mode command and manually accepting the key.

```
root@router1> ssh router2
The authenticity of host 'router2 (198.51.100.1)' can't be established.
RSA key fingerprint is 30:18:99:7a:3c:ed:40:04:0f:fd:c1:57:7e:6b:f3:90.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'router2,198.51.100.1' (RSA) to the list of known
hosts.
```

## Displaying the RPC Tags for a Command

You can display the RPC XML tags for operational mode commands in the CLI of the device. To display the RPC XML tags for a command, enter **display xml rpc** after the pipe symbol (|).

The following example displays the RPC tags for the **show route** command:

```
user@host> show route | display xml rpc
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.110/junos">
  <rpc>
    <get-route-information>
    </get-route-information>
  </rpc>
</cli>
  <banner></banner>
</cli>
</rpc-reply>
```

SLAX and XSLT scripts can execute RPCs using the RPC XML tags. Python scripts must convert the RPC tags and command options into a format suitable for Python. For more information about using Junos PyEZ to execute RPCs and about mapping RPC tags to

the corresponding Python method and method arguments, see [Using Junos PyEZ to Execute RPCs on Devices Running Junos OS](#).

## Using Operational Mode Commands in Op Scripts

Some operational mode commands do not have XML equivalents. SLAX and XSLT scripts can execute commands that have no XML equivalent using the `<command>` element. Python scripts can execute these commands by using the Junos PyEZ `cli()` method defined in the `Device` class.

If a command is not listed in the *Junos XML API Operational Developer Reference*, the command does not have an XML equivalent. Another way to determine whether a command has an XML equivalent is to issue the command followed by the `| display xml` command, for example:

```
user@host> operational-mode-command | display xml
```

If the output includes only tag elements like `<output>`, `<cli>`, and `<banner>`, the command might not have an XML equivalent. In the following example, the output indicates that the `show host` command has no XML equivalent:

```
user@host> show host hostname | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <output>
    ...
  </output>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```



**NOTE:** For some commands that have an XML equivalent, the output of the piped `| display xml` command does not include tag elements other than `<output>`, `<cli>`, and `<banner>` only because the relevant feature is not configured. For example, the `show services cos statistics forwarding-class` command has an XML equivalent that returns output in the `<service-cos-forwarding-class-statistics>` response tag, but if the configuration does not include any statements at the `[edit class-of-service]` hierarchy level, then there is no actual data for the `show services cos statistics forwarding-class | display xml` command to display. The output is similar to this:

```
user@host> show services cos statistics forwarding-class | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/8.3I0/junos">
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```

For this reason, the information in the *Junos XML API Operational Developer Reference* is usually more reliable.

SLAX and XSLT op scripts can include commands that have no XML equivalent. Use the `<command>`, `<xsl:value-of>`, and `<output>` elements in the script, as shown in the following code snippet. This snippet is expanded and fully described in “[Example: Displaying DNS Hostname Information Using an Op Script](#)” on page 487.

```
<xsl:variable name="query">
  <command>
    <xsl:value-of select="concat('show host ', $hostname)"/>
  </command>
</xsl:variable>
<xsl:variable name="result" select="jcs:invoke($query)"/>
<xsl:variable name="host" select="$result"/>
<output>
  <xsl:value-of select="concat('Name: ', $host)"/>
</output>
...
```

Python op scripts can execute commands that have no XML equivalent by using Junos PyEZ APIs. The `cli()` method defined in the **Device** class executes an operational mode command and returns the output in text format. For example:

```
from jnpr.junos import Device

def main():

    jdev = Device()
    jdev.open()
    res = jdev.cli('show host hostname', warning=False)
    print res
    jdev.close()

if __name__ == "__main__":
    main()
```

You can also specify **format='xml'** to return the output formatted as Junos OS XML elements. For more information about the Junos PyEZ `cli()` method, see [http://junos-pyez.readthedocs.org/en/latest/\\_modules/jnpr/junos/device.html#Device.cli](http://junos-pyez.readthedocs.org/en/latest/_modules/jnpr/junos/device.html#Device.cli).

#### Related Documentation

- [open\(\) Function \(SLAX and XSLT\) on page 209](#)
- [execute\(\) Function \(SLAX and XSLT\) on page 198](#)
- [Mapping Operational Mode Commands and Output Fields to Junos XML Notation on page 447](#)
- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
- [Junos PyEZ](#)

---

## Declaring and Using Command-Line Arguments in Op Scripts

Junos OS op scripts can accept command-line arguments when you invoke the script. You can include declarations in the op script or statements in the configuration such that a user executing the script can see the list of possible arguments when using the context-sensitive help in the CLI. The script must also include any necessary declarations

and code to process those arguments. The following sections detail how to define the help text and use the arguments in an op script:

- [Declaring Op Script Command-Line Arguments on page 453](#)
- [Using Command-Line Arguments in Op Scripts on page 454](#)
- [Example: Declaring Arguments in XSLT Scripts on page 455](#)
- [Example: Declaring and Using Arguments in Python Scripts on page 456](#)

## Declaring Op Script Command-Line Arguments

There are two ways to define the list of expected op scripts arguments that will be displayed when using context-sensitive help in the CLI: by including declarations in the script or by including statements in the Junos OS configuration. *Script-generated* and *configuration-generated* arguments have the same operational impact.

To declare arguments within SLAX or XSLT scripts, declare a global variable named **arguments**, containing **<argument>** tag elements. Within each **<argument>** tag element, include the required **<name>** tag element and the optional **<description>** tag element.

**XSLT Syntax**

```
<xsl:variable name="arguments">
  <argument>
    <name>name</name>
    <description>name</description>
  </argument>
</xsl:variable>
```

**SLAX Syntax**

```
var $arguments = {
  <argument> {
    <name> "name";
    <description> "descriptive-text";
  }
}
```

To declare arguments within a Python script, declare a global dictionary named **arguments** containing name-value pairs that map to the argument name and argument description.

**Python Syntax**

```
# Define arguments dictionary
arguments = {'name': 'description', 'name2': 'description2'}
```

Alternatively, you can declare op script arguments in the configuration. To declare arguments in the configuration, include the **arguments** statement at the **[edit system scripts op file filename]** hierarchy level.

```
[edit system scripts op file filename]
arguments {
  argument-name {
    description descriptive-text;
  }
}
```

After you declare an argument in either the script or the configuration, you can use the context-sensitive help to list available arguments. If you include the optional argument description, the text of the description appears in the CLI as a help-text string to describe the purpose of the argument, as discussed in [“Configuring Help Text for Op Scripts” on page 457](#).

```
user@host> op filename ?
Possible completions:
  argument-name      description
  argument-name      description
```

You can create a hidden argument for an op script by not including the argument declarations in the op script or configuration.



**NOTE:** If you configure arguments by including the `arguments` statement in the configuration, any arguments that you declare directly in the script are still available, but are not listed among the Possible completions when you issue the `op filename ?` command.

If you declare all arguments in the script (and none in the configuration), then the arguments do appear in the Possible completions list. This is because the management (mgd) process populates the Possible completions list by first checking the configuration for arguments. The mgd process looks in the script for arguments only if no arguments are found in the configuration. Thus, if arguments are declared in the configuration, the arguments declared in the script become hidden in the CLI.

## Using Command-Line Arguments in Op Scripts

To use command-line arguments in SLAX or XSLT op scripts, you must include a parameter declaration for each argument. The parameter name must be identical to the name provided in the argument declaration. The op script assigns the value for each script argument to the corresponding parameter, which can then be referenced throughout the script.

**XSLT Syntax**      `<xsl:param name="name" />`

**SLAX Syntax**      `param $name;`

Python op scripts include the **arguments** dictionary to define the arguments that appear in the CLI help for that op script. To process and use the command-line arguments, Python op scripts can use any valid means in the Python language. The following example uses the Python **argparse** module to process the script arguments. Note that, in this case, the argument names match those in the **arguments** dictionary.

**Python Syntax**

```
import argparse

arguments = {'arg1': 'description1', 'arg2': 'description2'}

def main():
    parser = argparse.ArgumentParser(description='This is a demo script.')

    #Define the arguments accepted by parser
    # which use the key names defined in the arguments dictionary
    for key in arguments:
        parser.add_argument('-' + key), required=True, help=arguments[key])
    args = parser.parse_args()

    # Extract the value
    print args.arg1
```

```

print args.arg2

if __name__ == '__main__':
    main()

```

When you execute an op script, for each argument you include on the command-line, you must specify a corresponding value. To do this, include an **argument-name** and an **argument-value** when you execute the script with the **op filename** command.

```
user@host> op filename argument-name argument-value
```



**NOTE:** If you specify an argument that the script does not recognize, the script ignores the argument.

### Example: Declaring Arguments in XSLT Scripts

Declare two arguments named **interface** and **protocol**. Execute the script, specifying the ge-0/2/0.0 interface and the **inet** protocol as values for the arguments. For either method, you must declare corresponding script parameters.

```

<xsl:param name="interface"/>
<xsl:param name="protocol"/>

```

The following examples show how to declare the arguments in the XSLT script and in the configuration:

#### Declaring Arguments in the Op Script (script1)

```

<xsl:variable name="arguments">
  <argument>
    <name>interface</name>
    <description>Name of interface to display</description>
  </argument>
  <argument>
    <name>protocol</name>
    <description>Protocol to display (inet, inet6)</description>
  </argument>
</xsl:variable>

```

#### Declaring Arguments in the Configuration

```

[edit system scripts op]
file script1 {
  arguments {
    interface {
      description "Name of interface to display";
    }
    protocol {
      description "Protocol to display (inet, inet6)";
    }
  }
}

```

### Executing the Script

```
user@host> op script1 interface ge-0/2/0.0 protocol inet
```

## Example: Declaring and Using Arguments in Python Scripts

Declare two arguments named **interface** and **protocol** in the Python op script. Execute the script, specifying the ge-0/2/0.0 interface and the **inet** protocol as values for the arguments.

### Declaring Arguments in the Op Script (script1.py)

```
from jnpr.junos import Device
import argparse

# Define arguments dictionary
arguments = {'interface': 'Name of interface to display', 'protocol': 'Protocol
to display (inet, inet6)'}

def main():
    parser = argparse.ArgumentParser()

    for key in arguments:
        parser.add_argument('--' + key, required=True, help=arguments[key])
    args = parser.parse_args()

    dev = Device()
    dev.open()

    res = dev.rpc.get_interface_information(interface_name=args.interface,
    terse=True)
    print args.interface + " status: " + \
        res.findtext("logical-interface/oper-status")

    for elem in res.xpath("//address-family \
[normalize-space(address-family-name)=$protocol]", \
protocol=args.protocol):
        if (elem.find("interface-address/ifa-local") is not None):
            print "inet address " + \
                elem.find("interface-address/ifa-local").text

    dev.close()

if __name__ == '__main__':
    main()
```

### Executing the Script

```
user@host> op script1.py interface ge-0/2/0.0 protocol inet
```

Alternatively, instead of including the **arguments** dictionary in the Python op script, you can include the arguments in the configuration exactly as you would for SLAX and XSLT scripts.

### Related Documentation

- [Example: Importing Files Using an Op Script on page 505](#)
- [Configuring Help Text for Op Scripts on page 457](#)
- [Mapping Operational Mode Commands and Output Fields to Junos XML Notation on page 447](#)
- [Using RPCs and Operational Mode Commands in Op Scripts on page 448](#)



## Configuring Help Text for Op Scripts

You can provide help text to describe an op script and its arguments when the `?` is used to list possible completions in the CLI. To configure help text, include the **description** statement:

```
description descriptive-text;
```

You can include this statement at the following hierarchy levels:

- [edit system scripts op file *filename*]
- [edit system scripts op file *filename* arguments *argument-name*]

The following examples show the configuration and the resulting output.

### Examples: Configuring Help Text for Op Scripts

Configure help text for a script and display the resulting output:

```
[edit system scripts op]
user@host# set file interface.xml description "Test the interface"
user@host# commit
...
[edit system scripts op]
user@host# set file ?
Possible completions:
<name>      Local filename of the script file
interface.xml  Test the interface
```

Configure help text for a script's command-line arguments and display the resulting output:

```
[edit system scripts op file interface.xml arguments]
user@host# set t1 description "Search for T1 interfaces"
user@host# set t3 description "Search for T3 interfaces"
user@host# commit

[edit system scripts op file interface.xml arguments]
user@host# set ?
Possible completions:
<name>      Name of the argument
t1          Search for T1 interfaces
t3          Search for T3 interfaces
```

#### Related Documentation

- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
- [Mapping Operational Mode Commands and Output Fields to Junos XML Notation on page 447](#)
- [Using RPCs and Operational Mode Commands in Op Scripts on page 448](#)

## Defining Commands to Allow in an Op Script

---

Operation (op) scripts automate operational tasks and network troubleshooting on devices running Junos OS. Op scripts can execute operational mode commands within the script. By default, when a user executes a script, the system does not permit the user to execute commands within a script for which their login class does not normally have permission to execute. Starting in Junos OS Release 14.2, you can configure commands that a particular op script is allowed to execute. The permission to execute commands within a script applies to all users, and the commands that you specify are executed, even if the user who executes the script does not have permissions to execute the commands.



**NOTE:** The `allow-commands` statement is only supported for op scripts that are local to the device. Remote op scripts that are executed using the `op url` command do not support executing unauthorized commands even when you configure the `allow-commands` statement.

In the following example, the op script `sam.slax` uses the `date` command, which is not permitted for user `user1`, who has view permissions.

```
user1@device> op sam
error: permission denied: date
```

To configure the `sam.slax` op script to execute the `date` command (must be a user in the Junos OS `super-user` login class):

```
[edit system scripts op file sam.slax]
admin@device# set allow-commands date
admin@device# commit
```

User `user1` can now successfully execute the op script.

```
user1@device> op sam
Tue Aug 19 11:11:33 EDT 2014
```

To define the commands to allow in an op script:

1. Navigate to the op script where you want to allow commands.



**NOTE:** Only users who belong to the Junos OS `super-user` login class can configure op scripts.

```
[edit]
admin@device# edit system scripts op file filename
```

2. Define the commands to allow.

```
[edit system scripts op file filename]
admin@device# set allow-commands "regular-expression"
```

3. Commit the configuration.

```
[edit system scripts op]
admin@device# commit
```

#### Related Documentation

- [Enabling an Op Script and Defining a Script Alias on page 459](#)
- [Disabling an Op Script on page 464](#)
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies](#)
- [allow-commands on page 744](#)

## Enabling an Op Script and Defining a Script Alias

Operation (op) scripts are stored on a device's hard disk in the `/var/db/scripts/op` directory or on the flash drive in the `/config/scripts/op` directory. Only users in the Junos OS **super-user** login class can access and edit files in these directories. For information about setting the storage location for scripts, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

To prevent the execution of unauthorized Python code on devices running Junos OS, unsigned Python scripts must meet certain requirements before you can execute the scripts on a device. Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user. For detailed information about the requirements for executing Python automation scripts on devices running Junos OS, see [“Requirements for Executing Python Automation Scripts on Devices Running Junos OS” on page 169](#).



**NOTE:** If the device has dual Routing Engines and you want to enable an op script to execute on both Routing Engines, you must copy the script to the `/var/db/scripts/op` or `/config/scripts/op` directory on both Routing Engines. The `commit synchronize` command does not automatically copy scripts between Routing Engines.

You must enable an op script before it can be executed. To enable an op script, include the `file filename` statement at the `[edit system scripts op]` hierarchy level, specifying the name of the file containing the op script. Only users who belong to the Junos OS **super-user** login class can enable op scripts.

```
[edit system scripts op]
user@host# set file filename
```

SLAX and Python scripts must include the `.slax` or `.py` filename extension, respectively, in both the actual script name and the filename in the configuration. XSLT scripts do not require a filename extension, but we strongly recommend that you append the `.xsl` extension. Whether or not you choose to include the `.xsl` extension on the file, the filename that you add at the `[edit system scripts op file]` hierarchy level must exactly match the

filename of the script in the directory. For example, if the XSLT script filename is **script1.xml**, then you must include **script1.xml** in the configuration hierarchy to enable the script; likewise, if the XSLT script filename is **script1**, then you must include **script1** in the configuration hierarchy.

To determine which op scripts are currently enabled on the device, use the **show** command to display the files configured at the **[edit system scripts op]** hierarchy level. To ensure that the enabled files are on the device, list the contents of the **/var/run/scripts/op/** directory using the **file list /var/run/scripts/op** operational mode command.

Optionally, you can define an alias for an op script. To define the alias, include the **command** statement at the **[edit system scripts op file filename]** hierarchy level.

```
[edit system scripts op file filename]
user@host# set command filename-alias
```

To execute the script, you can specify either the script filename or the alias.

```
user@host> op (filename | filename-alias)
```

By default, you cannot execute unsigned Python scripts on devices running Junos OS. To enable the execution of unsigned Python automation scripts that meet the requirements outlined in “[Requirements for Executing Python Automation Scripts on Devices Running Junos OS](#)” on page 169, you must configure the **language python** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set language python
```

#### Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user.

#### Related Documentation

- [Defining Commands to Allow in an Op Script on page 458](#)
- [Disabling an Op Script on page 464](#)

## Configuring Checksum Hashes for an Op Script

You can configure one or more checksum hashes that can be used to verify the integrity of an op script before the script runs on the switch, router, or security device.

To configure a checksum hash:

1. Create the script.
2. Place the script in the **/var/db/scripts/op** directory on the device.
3. Run the script through one or more hash functions to calculate hash values.

Junos OS supports MD5, SHA-1, and SHA-256 hash functions.

```

user@host> file checksum md5 /var/db/scripts/op/script1.slax
MD5 (/var/db/scripts/op/script1.slax) = 3af7884eb56e2d4489c2e49b26a39a97

user@host> file checksum sha1 /var/db/scripts/op/script1.slax
SHA1 (/var/db/scripts/op/script1.slax) =
00dc690fb08fb049577d012486c9a6dad34212c0

user@host> file checksum sha-256 /var/db/scripts/op/script1.slax
SHA256 (/var/db/scripts/op/script1.slax) =
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738

```

4. Configure the script with one or more hash values.

```

[edit system scripts op]
user@host# set file script1.slax checksum md5 3af7884eb56e2d4489c2e49b26a39a97

[edit system scripts op]
user@host# set file script1.slax checksum
sha-1 00dc690fb08fb049577d012486c9a6dad34212c0

[edit system scripts op]
user@host# set file script1.slax checksum
sha-256 150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738

```

During the execution of the script, Junos OS recalculates the checksum value using the configured hash and verifies that the calculated value matches the configured value. If the values differ, the execution of the script fails. When you configure multiple checksum values with different hash algorithms, all the configured values must match the calculated values; otherwise, the script execution fails.



**NOTE:** If the op script is stored remotely, do not include the checksum statement in the configuration. You can verify the script's integrity before it runs by specifying the hash value on the command line when you run the op command with the <url> option and the <key> option.

#### Related Documentation

- [Executing an Op Script from a Remote Site on page 462](#)
- [Configuring Checksum Hashes for a Commit Script on page 264](#)
- [Configuring Checksum Hashes for an Event Script on page 628](#)
- [Configuring Checksum Hashes for an SNMP Script on page 647](#)

## Executing an Op Script

Unlike commit scripts, operation (op) scripts do not execute during a commit operation. When you issue the **commit** command, op scripts configured at the **[edit system scripts op]** hierarchy level are placed into system memory and enabled for execution. After the commit operation completes, you can execute an op script from the CLI by issuing the **op** command in operational mode. You also can configure the device to execute an op script automatically when a member of a specific Junos OS login class logs in to the CLI.

- [Executing an Op Script by Issuing the op Command on page 462](#)
- [Executing an Op Script at Login on page 462](#)

## Executing an Op Script by Issuing the op Command

To execute an op script from the CLI, issue the **op** command, and specify a script filename, a script alias as defined by the **command** statement at the **[edit system scripts op file filename]** hierarchy level, or a URL.

```
user@host> op (filename | filename-alias | url url)
```

For more information about executing op scripts from a remote site using **op url**, see [“Executing an Op Script from a Remote Site” on page 462](#).



**NOTE:** In order to execute Python op scripts from a remote site, you must configure the **allow-url-for-python** statement at the **[edit system scripts op]** hierarchy level.

## Executing an Op Script at Login

You can configure an op script to execute automatically when any user belonging to a designated Junos OS login class logs in to the CLI. To associate an op script with a login class, include the **login-script filename** statement at the **[edit system login class class-name]** hierarchy level

```
[edit system login]
class class-name {
  login-script filename;
}
```

The following example configures the **super-user-login.slax** op script to execute when any user who belongs to the **super-user** class logs in to the CLI (provided that the script has been enabled as discussed in [“Enabling an Op Script and Defining a Script Alias” on page 459](#)).

```
[edit system login]
class super-user {
  login-script super-user-login.slax;
}
```

- Related Documentation**
- [login-script \(Login\)](#)
  - [Executing an Op Script from a Remote Site on page 462](#)

---

## Executing an Op Script from a Remote Site

As an alternative to storing operation (op) scripts locally on the device, you can store op scripts at a remote site. You then execute a remote op script by specifying the URL as an argument to the **op** command when you execute the script on the command line. You can execute SLAX and XSLT op scripts from a remote site by default. To execute Python op scripts from a remote site, you must first configure the **allow-url-for-python** statement at the **[edit system scripts op]** hierarchy level. Because you cannot guarantee that scripts executed from remote sites are secure, we recommend that you only authorize trusted users to execute scripts using the **op url** command.



**NOTE:** Statements configured under the `[edit system scripts op]` hierarchy level are only enforced for op scripts that are local to the device. Thus, even if you configure memory allocation, script dampening, traceoptions, or other op script-specific statements within that hierarchy, Junos OS does not apply the configuration when you execute a remote script using the `op url` command.

To execute an op script from a remote site:

1. Create the script.
2. (Optional) Store the script temporarily in the `/var/tmp` directory on the device, and run the script through one or more hash functions to calculate hash values.

Junos OS supports MD5, SHA-1, and SHA-256 hash functions.

```
user@host> file checksum md5 /var/tmp/script1.slax
MD5 (/var/tmp/script1.slax) = 3af7884eb56e2d4489c2e49b26a39a97

user@host> file checksum sha1 /var/tmp/script1.slax
SHA1 (/var/tmp/script1.slax) = 00dc690fb08fb049577d012486c9a6dad34212c0

user@host> file checksum sha-256 /var/tmp/script1.slax
SHA256 (/var/tmp/script1.slax) =
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

3. Enable the script on the device.

```
[edit]
user@host# set system scripts op file script1.slax
user@host# commit
```

4. For Python scripts, configure the `language python` and `allow-url-for-python` statements.

```
user@host# set system scripts language python
user@host# set system scripts op allow-url-for-python
```

5. Place the script on the remote server.
6. Provide the script URL and the optional hash values to the administrators who will execute the script.
7. Execute the script by running the `op url` command and specifying the URL that points to the remote file.

```
user@host> op url https://www.juniper.net/scripts/script1.slax
key md5 3af7884eb56e2d4489c2e49b26a39a97
```

This example shows how to include the `key` option and the MD5 checksum information.

To prevent the execution of any op scripts from remote sites, configure the `no-allow-url` statement at the `[edit system scripts op]` hierarchy level.

```
user@host# set system scripts op no-allow-url
```

When you configure the `no-allow-url` statement, issuing the `op url` operational mode command generates an error. This statement takes precedence when the `allow-url-for-python` statement is also present in the configuration.

**Related Documentation**

- [Executing an Op Script on page 461](#)
- [Configuring Checksum Hashes for an Op Script on page 460](#)
- [no-allow-url \(Op Scripts\) on page 757](#)
- [allow-url-for-python \(Op Scripts\) on page 745](#)

---

## Disabling an Op Script

You can disable an op script by deleting or deactivating the **file *filename*** statement at the **[edit system scripts op]** hierarchy in the configuration. To determine which op scripts are active on the device, issue the **show configuration system scripts op** operational mode command. The command output lists the enabled op scripts.

To delete an op script from the configuration:

1. From configuration mode in the CLI, enter the following command:

```
[edit]
user@host# delete system scripts op file filename
```

2. Commit the configuration:

```
user@host# commit
```

The **file** statement is removed from the configuration for the specified op script, and the **op** operational mode command no longer lists the op script filename as a valid completion.

To deactivate an op script in the configuration:

1. From configuration mode in the CLI, enter the following command:

```
[edit]
user@host# deactivate system scripts op file filename
```

2. Commit the configuration:

```
user@host# commit
```

The filename of the deactivated script remains in the configuration, but it is flagged with **inactive**. For example:

```
[edit system scripts op]
user@host# show

inactive: file script1.xml;
file script2.xml;
file script3.xml;
```



**NOTE:** You can reactivate an op script using the **activate system scripts op file *filename*** command.

---



Alternatively, you can delete the script from the `/var/db/scripts/op` directory on a device's hard disk or from the `/config/scripts/op` directory on the flash drive. Only users in the Junos OS **super-user** login class can access and edit files in these directories.

If you delete a script, you should also remove the **file** statement at the **[edit system scripts op]** hierarchy level in the configuration. If you delete an op script, but the **file** statement remains in the configuration, the CLI lists this script as a valid completion for the **op** command, but Junos OS issues an invalid filename error when the script is executed.

If you deactivate or delete the **file** statement for an op script in the configuration, you must enable the script again in order to execute it.

**Related  
Documentation**

- [Enabling an Op Script and Defining a Script Alias on page 459](#)
- [Defining Commands to Allow in an Op Script on page 458](#)



## CHAPTER 25

# Op Script Examples

- [Example: Changing the Configuration Using SLAX and XSLT Op Scripts on page 467](#)
- [Example: Changing the Configuration Using Python Op Scripts on page 472](#)
- [Example: Customizing Output of the show interfaces terse Command Using an Op Script on page 477](#)
- [Example: Displaying DNS Hostname Information Using an Op Script on page 487](#)
- [Example: Finding LSPs to Multiple Destinations Using an Op Script on page 491](#)
- [Example: Restarting an FPC Using an Op Script on page 495](#)
- [Example: Exporting Files Using an Op Script on page 498](#)
- [Example: Importing Files Using an Op Script on page 505](#)
- [Example: Searching Files Using an Op Script on page 510](#)

### Example: Changing the Configuration Using SLAX and XSLT Op Scripts

---

This example explains how to make structured changes to the Junos OS configuration using a SLAX op script.

- [Requirements on page 467](#)
- [Overview and Op Script on page 467](#)
- [Device Configuration on page 470](#)
- [Verification on page 471](#)

#### Requirements

This example uses a device running Junos OS.

#### Overview and Op Script

SLAX and XSLT op scripts can be used to make structured changes to the Junos OS configuration using the **jcs:load-configuration** template, which is included in the import file **junos.xsl**. Experienced users, who are familiar with Junos OS, can write scripts that prompt for the relevant configuration information and modify the configuration accordingly. This enables users who have less experience with Junos OS to safely modify the configuration using the script.

When called, the **jcs:load-configuration** template performs the following actions:

1. Locks the configuration database
2. Loads the configuration changes
3. Commits the configuration
4. Unlocks the configuration database

The **jcs:load-configuration** template makes changes to the configuration in **configure exclusive** mode. In this mode, Junos OS locks the candidate *global* configuration for as long as the script accesses the shared database and makes changes to the configuration without interference from other users.

If another user is currently editing the configuration in **configure exclusive** mode or if the database is already locked when the template is called, the call fails. In addition, if there are existing, uncommitted changes to the configuration when the template is called, the commit fails. If the template call is successful but the commit fails, Junos OS discards the uncommitted changes and rolls back the configuration.

You provide arguments to the **jcs:load-configuration** template to specify how to integrate the changes into the existing configuration, how to customize the commit operation, what changes to make to the configuration, and which connection handle to use. Starting in Junos OS Release 12.2, you can include the rollback parameter to return the configuration to a previously committed configuration, or you can supply a null configuration for the configuration parameter. If you supply a null configuration to **jcs:load-configuration**, the template performs a simple commit of the candidate configuration. The XSLT and SLAX syntax for the template call is:

```
<xsl:call-template name="jcs:load-configuration">
  <xsl:with-param name="action" select="(merge | override | replace)"/>
  <xsl:with-param name="commit-options" select="node-set"/>
  <xsl:with-param name="configuration" select="configuration-data"/>
  <xsl:with-param name="connection" select="connection-handle"/>
  <xsl:with-param name="rollback" select="number"/>
</xsl:call-template>

call jcs:load-configuration($action="(merge | override | replace)",
  $commit-options=node-set, $configuration=configuration-data,
  $connection=connection-handle, $rollback=number);
```

The following sample SLAX op script demonstrates how to use the **jcs:load-configuration** template to disable an interface on a device running Junos OS. All of the values required for the **jcs:load-configuration** template are defined as variables, which are then passed into the template as arguments.

In this example, the **usage** variable is initialized with a general description of the function of the script. When the script is run, the usage description is output to the CLI using a call to the **jcs:output()** function. This allows the user to verify that he is using the script for the correct purpose.

The script calls the **jcs:get-input()** function and prompts the user to enter the name of the interface that should be disabled. The interface name is stored in the **interface** variable. The configuration data that includes the changes to the configuration are stored in the

variable **config-changes**. This is the value used for the **configuration** parameter of the **jcs:load-configuration** template. This variable includes the Junos XML API tags for the configuration statements that are to be modified. The variable **interface**, which is supplied by the user, designates the name of the interface to disable.

The **load-action** variable is initialized to **merge**, which merges the configuration changes in the **disable** variable with the candidate configuration. This is the equivalent of the CLI configuration mode command **load merge**. Other load options include **replace** and **override**.

The **options** variable uses the **:=** operator to create a node-set, which is passed to the template as the value of the **commit-options** parameter. This example includes the **log** tag to add the description of the commit to the commit log file for future reference.

The call to the **jcs:open()** function opens a connection with the Junos OS management process (mgd) and returns a connection handle that is stored in the **conn\_handle** variable. All of the defined variables are passed as arguments to the **jcs:load-configuration** template at the time that it is called.

#### SLAX Syntax

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

import "../import/junos.xsl";

match / {
<op-script-results> {

    var $usage = "This script disables the interface specified by the user." _
                "The script modifies the candidate configuration to disable " _
                "the interface and commits the configuration to activate it.";
    var $temp = jcs:output($usage);

    var $interface = jcs:get-input("Enter interface to disable: ");

    var $config-changes = {
        <configuration> {
            <interfaces> {
                <interface> {
                    <name> $interface;
                    <disable>;
                }
            }
        }
    }

    var $load-action = "merge";

    var $options := {
        <commit-options> {
            <log> "disabling interface " _ $interface;
        }
    }
}
```

```
var $conn_handle = jcs:open();

var $results := {
  call jcs:load-configuration( $action=$load-action, $commit-options=$options,
    $configuration=$config-changes, $connection=$conn_handle);
}

$close-results = jcs:close($conn_handle);
}
```

The `:=` operator copies the results of the `jcs:load-configuration` template call to a temporary variable and runs the `node-set` function on that variable. The resulting node-set is then stored in the `results` variable. The `:=` operator ensures that the `results` variable is a node-set rather than a result tree fragment so that the script can access the contents. The `jcs:close()` function closes the connection.

By default, the `jcs:load-configuration` template does not output messages to the CLI. To quickly view any issues with the commit, you should add code to the script to output any error or warning messages that are generated as a result of the `jcs:load-configuration` template call.

```
if ($results//xnm:error) {
  for-each ($results//xnm:error) {
    <output> message;
  }
}
if ($results//xnm:warning) {
  for-each ($results//xnm:warning) {
    <output> message;
  }
}
```

## Device Configuration

To download, enable, and test the script:

1. Copy the script into a text file, name the file **config-change.slax**, and copy it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **config-change.slax**.

```
[edit system scripts op]
user@host# set file config-change.slax
```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```
[edit]
user@host# commit and-quit
```

4. Before running the script, issue the **show interfaces interface-name** operational mode command and record the current state of the interface that will be disabled by the script.
5. Execute the op script by issuing the **op config-change** operational mode command.

```
user@host> op config-change
```

This script disables the interface specified by the user. The script modifies the candidate configuration to disable the interface and commits the configuration to activate it.

```
Enter interface to disable: so-0/0/0
```

## Verification

- [Verifying the Commit on page 471](#)
- [Verifying the Configuration Changes on page 471](#)

### Verifying the Commit

**Purpose** Verify that the commit succeeded.

**Action** You should include code in your script that parses the node-set returned by the `jcs:load-configuration` template for any errors or warnings. This allows you to more easily determine whether the commit succeeded. If there are no warning or error messages, you can verify the success of the commit in several ways.

- Check the commit log to verify that the commit was successful. If you included the `log` option in the `commit-options` parameter, the message should be visible in the commit log along with the commit information.

```
user@host> show system commit
0   2010-09-22 17:08:17 PDT by user via junoscript
    disabling interface so-0/0/0
```

- Check the syslog message file to verify that the commit operation was logged. In this case, you also see an `SNMP_TRAP_LINK_DOWN` message for the disabled interface so-0/0/0. Depending on your configuration settings for traceoptions, this message might or might not appear in your log file.

```
user@host> show log messages | last
Sep 22 17:08:13 host file[7319]: UI_COMMIT: User 'user' requested 'commit'
operation (comment: disabling interface so-0/0/0)
Sep 22 17:08:16 host xntpd[1386]: ntpd exiting on signal 1
Sep 22 17:08:16 host xntpd[1386]: ntpd 4.2.0-a Fri Jun 25 13:48:13 UTC 2010
(1)
Sep 22 17:08:16 host mib2d[1434]: SNMP_TRAP_LINK_DOWN: ifIndex 526,
ifAdminStatus down(2), ifOperStatus down(2), ifName so-0/0/0
```

### Verifying the Configuration Changes

**Purpose** Verify that the correct changes are integrated into the configuration.

**Action** • Display the configuration and verify that the changes are visible for the specified interface.

```
user@host> show configuration interfaces so-0/0/0
disable;
```

- For this example, you also can issue the **show interfaces *interface-name*** operational mode command to check that the interface was disabled. In this case, the output captured *before* the interface was disabled shows that the interface is **Enabled**.

```
user@host> show interfaces so-0/0/0
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 526
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2010-09-14 10:33:25 PDT (1w1d 06:27 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
```

The output captured *after* running the script to disable the interface shows that the interface is now **Administratively down**.

```
user@host> show interfaces so-0/0/0
Physical interface: so-0/0/0, Administratively down, Physical link is Up
  Interface index: 128, SNMP ifIndex: 526
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2010-09-14 10:33:25 PDT (1w1d 06:40 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
```

#### Related Documentation

- [jcs:load-configuration Template on page 235](#)
- [Storing and Enabling Scripts on page 661](#)
- [get-input\(\) Function \(SLAX and XSLT\) and get\\_input\(\) \(Python\) on page 202](#)
- [open\(\) Function \(SLAX and XSLT\) on page 209](#)
- [Example: Changing the Configuration Using Python Op Scripts on page 472](#)

---

## Example: Changing the Configuration Using Python Op Scripts

Op scripts enable you to make controlled changes to the Junos OS configuration. Op scripts are advantageous, because they can gather operational information about a device and update the configuration based on that information. Experienced users who are familiar with Junos OS can write op scripts that prompt for the relevant configuration information and modify the configuration accordingly. This enables users who have less experience with Junos OS to safely modify the configuration using the script. This example



demonstrates how to make changes to the Junos OS configuration using a Python op script that leverages Junos PyEZ APIs.

- [Requirements on page 473](#)
- [Overview and Op Script on page 473](#)
- [Device Configuration on page 475](#)
- [Verification on page 476](#)

## Requirements

This example uses the following hardware and software components:

- Device running Junos OS Release 16.1R3 or later release that includes the Python extensions package.

## Overview and Op Script

Python op scripts can make changes to the Junos OS configuration using the [Junos PyEZ `jnpr.junos.utils.config.Config`](#) utility. The Junos PyEZ **Config** utility provides instance methods to lock the configuration, load the configuration data and specify how to integrate it into the configuration, commit the configuration, and unlock the configuration. For more information about using Junos PyEZ to configure devices running Junos OS, see [Using Junos PyEZ to Configure Devices Running Junos OS](#). The Python op script in this example demonstrates how to update the configuration to disable an interface on the local device.

The Python op script imports the **Device** class, which handles the connection with the device running Junos OS, the **Config** class, which is used to perform configuration mode commands on the target device, the **`jnpr.junos.exception`** module, which contains exceptions encountered when managing devices running Junos OS, and the **jcs** module, which permits the script to execute supported extension functions. This example binds the **Config** instance to the **Device** instance rather than create a standalone variable for the instance of the **Config** class.

In this example, the **usage** variable is initialized with a general description of the script's function. When the script is executed, the script outputs the usage description on the CLI so that the user can verify the purpose for that script.

The script calls the **`jcs.get_input()`** extension function, which prompts the user to enter the name of the interface to disable, and stores the interface name in the **interface** variable. The **config\_xml** variable is an XML string that defines the configuration changes.

The script does not supply a host parameter when creating the **Device** instance, which causes the **open()** method to establish a connection with the local device. The **Config** utility methods then lock the candidate configuration, load the configuration changes into the candidate configuration as a **load merge** operation, commit the candidate configuration, and then unlock it. The **dev.close()** method closes the connection.

The Python program includes code for handling exceptions such as **LockError** for errors that occur when locking the configuration, **ConfigLoadError** for errors that occur when

loading the configuration data into the candidate configuration, and **CommitError** for errors that occur during the commit operation.

**Python Script**

```
from jnpr.junos import Device
from jnpr.junos.utils.config import Config
from jnpr.junos.exception import *
import jcs

def main():

    usage = """
    This script disables the interface specified by the user.
    The script modifies the candidate configuration to disable
    the interface and commits the configuration to activate it.
    """
    print usage

    interface = jcs.get_input("Enter interface to disable: ")
    if not interface:
        print "invalid interface"
        sys.exit(1)

    config_xml = """
    <configuration>
      <interfaces>
        <interface>
          <name>{0}</name>
          <disable/>
        </interface>
      </interfaces>
    </configuration>
    """.format(interface).strip()

    dev = Device()
    dev.open()
    dev.bind( cu=Config )

    # Lock the configuration, load configuration changes, and commit
    print "Locking the configuration"
    try:
        dev.cu.lock()
    except LockError:
        print "Error: Unable to lock configuration"
        dev.close()
        return

    print "Loading configuration changes"
    try:
        dev.cu.load(config_xml, format="xml", merge=True)
    except ConfigLoadError as err:
        print err
        print "Unable to load configuration changes: "
        print "Unlocking the configuration"
        try:
            dev.cu.unlock()
        except UnlockError:
            print "Error: Unable to unlock configuration"
        dev.close()
        return

    print "Committing the configuration"
```

```

try:
    dev.cu.commit()
except CommitError:
    print "Error: Unable to commit configuration"
    print "Unlocking the configuration"
    try:
        dev.cu.unlock()
    except UnlockError:
        print "Error: Unable to unlock configuration"
    dev.close()
    return

print "Unlocking the configuration"
try:
    dev.cu.unlock()
except UnlockError:
    print "Error: Unable to unlock configuration"

dev.close()

if __name__ == "__main__":
    main()

```

## Device Configuration

To download, enable, and test the script:

1. Copy the script into a text file, name the file **config-change.py**, and copy it to the **/var/db/scripts/op/** directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. In configuration mode, include the **file config-change.py** statement at the **[edit system scripts op]** hierarchy level.

```

[edit system scripts op]
user@host# set file config-change.py

```

3. Enable Python automation scripts on the device.

```

[edit system scripts]
user@host# set language python

```

4. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```

[edit]
user@host# commit and-quit

```

5. Before running the script, issue the **show interfaces interface-name** operational mode command and record the current state of the interface that will be disabled by the script.
6. Execute the op script by issuing the **op config-change.py** operational mode command.

```
user@host> op config-change.py
This script disables the interface specified by the user.
The script modifies the candidate configuration to disable
the interface and commits the configuration to activate it.
Enter interface to disable: so-0/0/0
```

## Verification

- [Verifying the Commit on page 476](#)
- [Verifying the Configuration Changes on page 476](#)

---

### Verifying the Commit

**Purpose** Verify that the commit succeeded.

**Action** You should include code in your script that catches any warnings or errors associated with changing and committing the configuration. This enables you to more easily determine whether the commit succeeded. If there are no warning or error messages, you can verify the success of the commit in several ways.

- Check the commit log to verify that the commit was successful.

```
user@host> show system commit
0   2010-09-22 17:08:17 PDT by user via netconf
```

- Check the syslog message file to verify that the commit operation was logged. In this case, you also see an **SNMP\_TRAP\_LINK\_DOWN** message for the disabled interface. Depending on your configuration settings for traceoptions, this message might or might not appear in your log file.

```
user@host> show log messages | last
Sep 22 17:08:13 host file[7319]: UI_COMMIT: User 'user' requested 'commit'
operation
Sep 22 17:08:16 host xntpd[1386]: ntpd exiting on signal 1
Sep 22 17:08:16 host xntpd[1386]: ntpd 4.2.0-a Fri Jun 25 13:48:13 UTC 2010
(1)
Sep 22 17:08:16 host mib2d[1434]: SNMP_TRAP_LINK_DOWN: ifIndex 526,
ifAdminStatus down(2), ifOperStatus down(2), ifName so-0/0/0
```

---

### Verifying the Configuration Changes

**Purpose** Verify that the correct changes are integrated into the configuration.

**Action** • Display the configuration and verify that the changes are visible for the specified interface.

```
user@host> show configuration interfaces so-0/0/0
disable;
```

- For this example, you also can issue the **show interfaces *interface-name* operational** mode command to check that the interface was disabled. In this case, the output captured *before* the interface was disabled shows that the interface is **Enabled**.

```
user@host> show interfaces so-0/0/0
```

```
Physical interface: so-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 526
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2010-09-14 10:33:25 PDT (1w1d 06:27 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
```

The output captured *after* running the script to disable the interface shows that the interface is now **Administratively down**.

```
user@host> show interfaces so-0/0/0
Physical interface: so-0/0/0, Administratively down, Physical link is Up
  Interface index: 128, SNMP ifIndex: 526
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Down Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  CoS queues     : 4 supported, 4 maximum usable queues
  Last flapped   : 2010-09-14 10:33:25 PDT (1w1d 06:40 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
```

#### Related Documentation

- [Junos PyEZ](#)
- [Example: Changing the Configuration Using SLAX and XSLT Op Scripts on page 467](#)
- [Storing and Enabling Scripts on page 661](#)

## Example: Customizing Output of the show interfaces terse Command Using an Op Script

This example uses an op script to customize the output of the **show interfaces terse** command. A line-by-line explanation of the XSLT script is provided.

- [Requirements on page 477](#)
- [Overview and Op Script on page 478](#)
- [Configuration on page 486](#)
- [Verification on page 486](#)

### Requirements

This example uses a device running Junos OS.

## Overview and Op Script

By default, the layout of the **show interfaces terse** command looks like this:

```
user@host> show interfaces terse
Interface      Admin Link Proto  Local                      Remote
dsc             up   up
fxp0            up   up
fxp0.0          up   up   inet   192.168.71.246/21
fxp1            up   up
fxp1.0          up   up   inet   10.0.0.4/8
               inet6  fe80::200:ff:fe00:4/64
               tnp    fc00::10:0:0:4/64
               4
gre             up   up
ipip            up   up
lo0             up   up
lo0.0           up   up   inet   127.0.0.1                --> 0/0
lo0.16385       up   up   inet   fe80::2a0:a5ff:fe12:2f04
               inet6
lsi             up   up
mtun            up   up
pimd            up   up
pime            up   up
tap             up   up
```

In Junos XML, the output fields are represented as follows:

```
user@host> show interfaces terse | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <interface-information
    xmlns="http://xml.juniper.net/junos/10.0R1/junos-interface" junos:style="terse">
    <physical-interface>
      <name>dsc</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    </physical-interface>
    <physical-interface>
      <name>fxp0</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    <logical-interface>
      <name>fxp0.0</name>
      <admin-status>up</admin-status>
      <oper-status>up</oper-status>
    ... Remainder of output omitted for brevity ...
```

**XSLT Syntax** The following script customizes the output of the **show interfaces terse** command. A line-by-line explanation of the script is provided.

```
1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4    xmlns:junos="http://xml.juniper.net/junos/*/junos"
5    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7    <xsl:import href="../import/junos.xsl"/>
8
9    <xsl:variable name="arguments">
      <argument>
```

```

10         <name>interface</name>
11         <description>Name of interface to display</description>
12     </argument>
13     <argument>
14         <name>protocol</name>
15         <description>Protocol to display (inet, inet6)</description>
16     </argument>
17 </xsl:variable>
18 <xsl:param name="interface"/>
19 <xsl:param name="protocol"/>
20 <xsl:template match="/">
21     <op-script-results>
22         <xsl:variable name="rpc">
23             <get-interface-information>
24                 <terse/>
25                 <xsl:if test="$interface">
26                     <interface-name>
27                         <xsl:value-of select="$interface"/>
28                     </interface-name>
29                 </xsl:if>
30             </get-interface-information>
31         </xsl:variable>
32         <xsl:variable name="out" select="jcs:invoke($rpc)"/>
33         <interface-information junos:style="terse">
34             <xsl:choose>
35                 <xsl:when test="$protocol='inet' or
$protocol='inet6'
                                     or $protocol='mpls' or
$protocol='tnp'">
36                     <xsl:for-each select="$out/physical-interface/
logical-interface[address-family/address-family-name =
$protocol]">
37                         <xsl:call-template name="intf"/>
38                     </xsl:for-each>
39                 </xsl:when>
40                 <xsl:when test="$protocol">
41                     <xnm:error>
42                         <message>
43                             <xsl:text>invalid protocol: </xsl:text>
44                             <xsl:value-of select="$protocol"/>
45                         </message>
46                     </xnm:error>
47                 </xsl:when>
48                 <xsl:otherwise>
49                     <xsl:for-each
select="$out/physical-interface/logical-interface">
50                         <xsl:call-template name="intf"/>
51                     </xsl:for-each>
52                 </xsl:otherwise>
53             </xsl:choose>
54         </interface-information>
55     </op-script-results>
56 </xsl:template>
57 <xsl:template name="intf">
58     <xsl:variable name="status">
59         <xsl:choose>
60             <xsl:when test="admin-status='up' and oper-status='up'">
61                 <xsl:text> </xsl:text>
62             </xsl:when>
63             <xsl:when test="admin-status='down'">
64                 <xsl:text>offline</xsl:text>

```

```

65             </xsl:when>
66             <xsl:when test="oper-status='down' and
../admin-status='down'">
67                 <xsl:text>p-offline</xsl:text>
68             </xsl:when>
69             <xsl:when test="oper-status='down' and
../oper-status='down'">
70                 <xsl:text>p-down</xsl:text>
71             </xsl:when>
72             <xsl:when test="oper-status='down'">
73                 <xsl:text>down</xsl:text>
74             </xsl:when>
75             <xsl:otherwise>
76                 <xsl:value-of select="concat(oper-status, '/',
admin-status)"/>
77             </xsl:otherwise>
78         </xsl:choose>
79     </xsl:variable>
80     <xsl:variable name="desc">
81         <xsl:choose>
82             <xsl:when test="description">
83                 <xsl:value-of select="description"/>
84             </xsl:when>
85             <xsl:when test="../description">
86                 <xsl:value-of select="../description"/>
87             </xsl:when>
88         </xsl:choose>
89     </xsl:variable>
90     <logical-interface>
91         <name><xsl:value-of select="name"/></name>
92         <xsl:if test="string-length($desc)">
93             <admin-status><xsl:value-of
select="$desc"/></admin-status>
94         </xsl:if>
95         <admin-status><xsl:value-of
select="$status"/></admin-status>
96         <xsl:choose>
97             <xsl:when test="$protocol">
98                 <xsl:copy-of
select="address-family[address-family-name =
$protocol]"/>
99             </xsl:when>
100            <xsl:otherwise>
101                <xsl:copy-of select="address-family"/>
102            </xsl:otherwise>
103        </xsl:choose>
104    </logical-interface>
105 </xsl:template>
106 </xsl:stylesheet>

```

**Line-by-Line Explanation** Lines 1 through 7, Line 20, and Lines 105 and 106 are the boilerplate that you include in every op script. For more information, see ["Required Boilerplate for Op Scripts" on page 445](#).

```

1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7  <xsl:import href="../import/junos.xsl"/>
...

```



```

20      <xsl:template match="/">
      ...
105    </xsl:template>
106  </xsl:stylesheet>

```

Lines 8 through 17 declare a variable called **arguments**, containing two arguments to the script: **interface** and **protocol**. This variable declaration causes **interface** and **protocol** to appear in the command-line interface (CLI) as available arguments to the script.

```

8      <xsl:variable name="arguments">
9        <argument>
10         <name>interface</name>
11         <description>Name of interface to display</description>
12       </argument>
13       <argument>
14         <name>protocol</name>
15         <description>Protocol to display (inet, inet6)</description>
16       </argument>
17     </xsl:variable>

```

Lines 18 and 19 declare two parameters to the script, corresponding to the arguments created in Lines 8 through 17. The parameter names must exactly match the argument names.

```

18     <xsl:param name="interface"/>
19     <xsl:param name="protocol"/>

```

Lines 20 through 31 declare a variable named **rpc**. The **show interfaces terse** command is assigned to the **rpc** variable. If you include the **interface** argument when you execute the script, the value of the argument (the interface name) is passed into the script.

```

20     <xsl:template match="/">
21       <op-script-results>
22         <xsl:variable name="rpc">
23           <get-interface-information>
24             <terse/>
25             <xsl:if test="$interface">
26               <interface-name>
27                 <xsl:value-of select="$interface"/>
28               </interface-name>
29             </xsl:if>
30           </get-interface-information>
31         </xsl:variable>

```

Line 32 declares a variable named **out** and applies to it the execution of the **rpc** variable (**show interfaces terse** command).

```

32       <xsl:variable name="out" select="jcs:invoke($rpc)"/>

```

Line 33 specifies that the output level of the **show interfaces** command being modified is **terse** (as opposed to **extensive**, **detail**, and so on).

```

33       <interface-information junos:style="terse">

```

Lines 34 through 39 specify that if you include the **protocol** argument when you execute the script and if the protocol value that you specify is **inet**, **inet6**, **mpls**, or **tnp**, the **intf** template is applied to each instance of that protocol type in the output.

```

34         <xsl:choose>
35         <xsl:when test="$protocol='inet' or
$protocol='inet6'
or $protocol='mpls' or $protocol='tnp'">

```

```

36                <xsl:for-each select="$out/physical-interface/
logical-interface[address-family/address-family-name =
$protocol]">
37                    <xsl:call-template name="intf"/>
38                </xsl:for-each>
39            </xsl:when>

```

Lines 40 through 47 specify that if you include the **protocol** argument when you execute the script and if the protocol value that you specify is something other than **inet**, **inet6**, **mpls**, or **tnp**, an error message is generated.

```

40                <xsl:when test="$protocol">
41                    <xnm:error>
42                        <message>
43                            <xsl:text>invalid protocol: </xsl:text>
44                            <xsl:value-of select="$protocol"/>
45                        </message>
46                    </xnm:error>
47                </xsl:when>

```

Lines 48 through 52 specify that if you do not include the **protocol** argument when you execute the script, the **intf** template is applied to each logical interface in the output.

```

48                <xsl:otherwise>
49                    <xsl:for-each
select="$out/physical-interface/logical-interface">
50                        <xsl:call-template name="intf"/>
51                    </xsl:for-each>
52                </xsl:otherwise>

```

Lines 53 through 56 are closing tags.

```

53            </xsl:choose>
54        </interface-information>
55    </op-script-results>
56</xsl:template>

```

Line 57 opens the **intf** template. This template customizes the output of the **show interfaces terse** command.

```

57    <xsl:template name="intf">

```

Line 58 declares a variable called **status**, the purpose of which is to specify how the interface status is reported. Lines 59 through 78 contain a **<xsl:choose>** instruction that populates the **status** variable by considering all the possible states. As always in XSLT, the first **<xsl:when>** instruction that evaluates as TRUE is executed, and the remainder are ignored. Each **<xsl:when>** instruction is explained separately.

```

58        <xsl:variable name="status">
59            <xsl:choose>

```

Lines 60 through 62 specify that if **admin-status** is 'up' and **oper-status** is 'up', no output is generated. In this case, the **status** variable remains empty.

```

60                <xsl:when test="admin-status='up' and oper-status='up'">
61                    <xsl:text> </xsl:text>
62                </xsl:when>

```

Lines 63 through 65 specify that if **admin-status** is 'down', the **status** variable contains the text **offline**.

```

63                <xsl:when test="admin-status='down'">
64                    <xsl:text>offline</xsl:text>

```

```
65                </xsl:when>
```

Lines 66 through 68 specify that if **oper-status** is 'down' and the physical interface **admin-status** is 'down', the **status** variable contains the text **p-offline**. (../ selects the physical interface.)

```
66                <xsl:when test="oper-status='down' and
67                ../admin-status='down'">
67                    <xsl:text>p-offline</xsl:text>
68                </xsl:when>
```

Lines 69 through 71 specify that if **oper-status** is 'down' and the physical interface **oper-status** is 'down', the **status** variable contains the text **p-down**. (../ selects the physical interface.)

```
69                <xsl:when test="oper-status='down' and
70                ../oper-status='down'">
70                    <xsl:text>p-down</xsl:text>
71                </xsl:when>
```

Lines 72 through 74 specify that if **oper-status** is 'down', the **status** variable contains the text **down**.

```
72                <xsl:when test="oper-status='down'">
73                    <xsl:text>down</xsl:text>
74                </xsl:when>
```

Lines 75 through 77 specify that if none of the test cases are true, the **status** variable contains **oper-status** and **admin-status** concatenated with a slash as a separator.

```
75                <xsl:otherwise>
76                    <xsl:value-of select="concat(oper-status, '/',
77                admin-status)"/>
77                </xsl:otherwise>
```

Lines 78 and 79 are closing tags.

```
78                </xsl:choose>
79            </xsl:variable>
```

Lines 80 through 89 define a variable called **desc**. An **<xsl:choose>** instruction populates the variable by selecting the most specific interface description available. If a logical interface description is included in the configuration, it is used to populate the **desc** variable. If not, the physical interface description is used. If no physical interface description is included in the configuration, the variable remains empty. As always in XSLT, the first **<xsl:when>** instruction that evaluates as TRUE is executed, and the remainder are ignored.

```
80            <xsl:variable name="desc">
81                <xsl:choose>
82                    <xsl:when test="description">
83                        <xsl:value-of select="description"/>
84                    </xsl:when>
85                    <xsl:when test="../description">
86                        <xsl:value-of select="../description"/>
87                    </xsl:when>
88                </xsl:choose>
89            </xsl:variable>
```

The remainder of the script specifies how the operational mode output is displayed.

Lines 90 and 91 specify that the logical interface name is displayed first in the output.

```
90                <logical-interface>
```

```
91          <name><xsl:value-of select="name"/></name>
```

Lines 92 through 94 test whether the **desc** variable has a nonzero number of characters. If the number of characters is more than zero, the interface description is displayed in the standard location of the **admin-status** field. (In standard output, the **admin-status** field is displayed on the second line.)

```
92          <xsl:if test="string-length($desc)">
93              <admin-status><xsl:value-of
select="$desc"/></admin-status>
94          </xsl:if>
```

Line 95 specifies that the interface status as defined in the **status** variable is displayed next.

```
95          <admin-status><xsl:value-of
select="$status"/></admin-status>
```

Lines 96 through 103 specify that if you include the **protocol** argument when you execute the script, only interfaces with that protocol configured are displayed. If you do not include the **protocol** argument, all interfaces are displayed.

```
96          <xsl:choose>
97              <xsl:when test="$protocol">
98                  <xsl:copy-of
select="address-family[address-family-name =
$protocol]"/>
99              </xsl:when>
100             <xsl:otherwise>
101                 <xsl:copy-of select="address-family"/>
102             </xsl:otherwise>
103         </xsl:choose>
```

Lines 104 through 106 are closing tags.

```
104         </logical-interface>
105     </xsl:template>
106 </xsl:stylesheet>
```

**SLAX Syntax** The SLAX version of the script is as follows:

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
    <argument> {
        <name> "interface";
        <description> "Name of interface to display";
    }
    <argument> {
        <name> "protocol";
        <description> "Protocol to display (inet, inet6)";
    }
}
param $interface;
param $protocol;
match / {
    <op-script-results> {
        var $rpc = {
            <get-interface-information> {
```

```

        <terse>;
        if ($interface) {
            <interface-name> $interface;
        }
    }
}
var $out = jcs:invoke($rpc);
<interface-information junos:style="terse"> {
    if ($protocol='inet' or $protocol='inet6' or $protocol='mpls' or
        $protocol='tnp') {
        for-each ($out/physical-interface/
            logical-interface[address-family/address-family-name =
$protocol]) {
            call intf();
        }
    } else if ($protocol) {
        <xnm:error> {
            <message> {
                expr "invalid protocol: ";
                expr $protocol;
            }
        }
    } else {
        for-each ($out/physical-interface/logical-interface) {
            call intf();
        }
    }
}
}
}
intf () {
    var $status = {
        if (admin-status='up' and oper-status='up') {
        } else if (admin-status='down') {
            expr "offline";
        } else if (oper-status='down' and ../admin-status='down') {
            expr "p-offline";
        } else if (oper-status='down' and ../oper-status='down') {
            expr "p-down";
        } else if (oper-status='down') {
            expr "down";
        } else {
            expr oper-status _ '/' _ admin-status;
        }
    }
    var $desc = {
        if (description) {
            expr description;
        } else if (../description) {
            expr ../description;
        }
    }
    <logical-interface> {
        <name> name;
        if (string-length($desc)) {
            <admin-status> $desc;
        }
        <admin-status> $status;
        if ($protocol) {
            copy-of address-family[address-family-name = $protocol];
        } else {

```

```
        copy-of address-family;  
    }  
}
```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **interface.xml** or **interface.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **interface.xml** or **interface.slax** as appropriate.

```
[edit system scripts op]  
user@host# set file interface.(slax | xml)
```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```
[edit]  
user@host# commit and-quit
```

4. Execute the op script by issuing the **op interface** operational mode command.

## Verification

### Verifying the Commit Script Output

---

**Purpose** Verify that the script behaves as expected.

**Action** Issue the **show interfaces terse** and **op interface** operational commands and compare the output. The **show interfaces terse** command displays the standard output. The **op interface** command displays the customized output.

```
user@host> show interfaces terse
Interface      Admin Link Proto  Local              Remote
dsc            up    up
fxp0           up    up
fxp0.0         up    up    inet   192.168.71.246/21
fxp1           up    up
fxp1.0         up    up    inet   10.0.0.4/8
               inet6  fe80::200:ff:fe00:4/64
               tnp    fc00::10:0:0:4/64
               4
gre            up    up
ipip           up    up
lo0            up    up
lo0.0          up    up    inet   127.0.0.1          --> 0/0
lo0.16385      up    up    inet   fe80::2a0:a5ff:fe12:2f04
```

```
user@host> op interface
Interface      Admin Link Proto  Local              Remote
fxp0.0         This is the Ethernet Management interface.
               inet   192.168.71.246/21
fxp1.0         inet   10.0.0.4/8
               inet6  fe80::200:ff:fe00:4/64
               fc00::10:0:0:4/64
               tnp    4
lo0.0          inet   127.0.0.1          --> 0/0
lo0.16385      inet   fe80::2a0:a5ff:fe12:2f04-->
```

Issue the **op interface** operational command for different hierarchy levels and review the output. For example:

```
user@host> op interface interface fxp0
Interface      Admin Link Proto  Local              Remote
fxp0.0         This is the Ethernet Management interface.
               inet   192.168.71.246/21
```

```
user@host> op interface protocol inet
Interface      Admin Link Proto  Local              Remote
fxp0.0         This is the Ethernet Management interface.
               inet   192.168.71.246/21
fxp1.0         inet   10.0.0.4/8
lo0.0          inet   127.0.0.1          --> 0/0
lo0.16385      inet
```

## Example: Displaying DNS Hostname Information Using an Op Script

This example uses an op script to display Domain Name System (DNS) information for a device in your network.

- [Requirements on page 488](#)
- [Overview and Op Script on page 488](#)
- [Configuration on page 490](#)
- [Verification on page 491](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Op Script

This script displays DNS information for a device in your network. The script offers a slight improvement over the **show host *hostname*** command because you do not need to enter a hostname or IP address to view DNS information for the device you are currently using.

There is no Junos Extensible Markup Language (XML) equivalent for the **show host *hostname*** command. Therefore, this script uses the **show host *hostname*** command directly rather than using a remote procedure call (RPC).

The script is provided in two distinct versions, one using the **<xsl:choose>** element and the other using the **jcs:first-of()** function. Both versions accept the same argument and produce the same output. Each version is shown in both XSLT and SLAX syntax.

### XSLT Syntax Using the <xsl:choose> Element

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:variable name="arguments">
    <argument>
      <name>dns</name>
      <description>Name or IP address of a host</description>
    </argument>
  </xsl:variable>
  <xsl:param name="dns"/>
  <xsl:template match="/">
    <op-script-results>
      <xsl:variable name="query">
        <xsl:choose>
          <xsl:when test="$dns">
            <command>
              <xsl:value-of select="concat('show host ', $dns)"/>
            </command>
          </xsl:when>
          <xsl:when test="$hostname">
            <command>
              <xsl:value-of select="concat('show host ',
$hostname)"/>
            </command>
          </xsl:when>
        </xsl:choose>
```



```

        </xsl:variable>
        <xsl:variable name="result" select="jcs:invoke($query)"/>
        <xsl:variable name="host" select="$result"/>
        <output>
            <xsl:value-of select="concat('Name: ', $host)"/>
        </output>
    </op-script-results>
</xsl:template>
</xsl:stylesheet>

```

### XSLT Syntax Using the jcs:first-of() Function

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
    xmlns:junos="http://xml.juniper.net/junos/*/junos"
    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
    <xsl:import href="../import/junos.xsl"/>

    <xsl:variable name="arguments">
        <argument>
            <name>dns</name>
            <description>Name or IP address of a host</description>
        </argument>
    </xsl:variable>
    <xsl:param name="dns"/>
    <xsl:template match="/">
        <op-script-results>
            <xsl:variable name="target" select="jcs:first-of($dns, $hostname)"/>
            <xsl:variable name="query">
                <command>
                    <xsl:value-of select="concat('show host ', $target)"/>
                </command>
            </xsl:variable>
            <xsl:variable name="result" select="jcs:invoke($query)"/>
            <xsl:variable name="host" select="$result"/>
            <output>
                <xsl:value-of select="concat('Name: ', $host)"/>
            </output>
        </op-script-results>
    </xsl:template>
</xsl:stylesheet>

```

### SLAX Syntax Using the <xsl:choose> Element

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
    <argument> {
        <name> "dns";
        <description> "Name or IP address of a host";
    }
}
param $dns;
match / {
    <op-script-results> {
        var $query = {
            if ($dns) {
                <command> 'show host ' _ $dns;
            } else if ($hostname) {

```

```

        <command> 'show host ' _ $hostname;
    }
}
var $result = jcs:invoke($query);
var $host = $result;
<output> 'Name: ' _ $host;
}
}

```

### SLAX Syntax Using the jcs:first-of() Function

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "dns";
    <description> "Name or IP address of a host";
  }
}
param $dns;
match / {
  <op-script-results> {
    var $target = jcs:first-of($dns, $hostname);
    var $query = {
      <command> 'show host ' _ $target;
    }
    var $result = jcs:invoke($query);
    var $host = $result;
    <output> 'Name: ' _ $host;
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **hostname.xsl** or **hostname.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.

2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **hostname.xsl** or **hostname.slax** as appropriate.

```

[edit system scripts op]
user@host# set file hostname.(slax | xsl)

```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```

[edit]
user@host# commit and-quit

```

4. Execute the op script by issuing the **op hostname <dns (hostname | address)>** operational mode command.

## Verification

### Verifying the Commit Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** When you issue the **op hostname** operational mode command without the **dns** option, DNS information is displayed for the local device:

```
user@host1> op hostname
Name:
host1 has address 10.168.71.246
```

When you issue the **op hostname dns *hostname*** command, DNS information is displayed for the specified device:

```
user@host1> op hostname dns router1
Name:
router1 has address 10.168.71.249
```

When you issue the **op hostname dns *address*** command, DNS information is displayed for the specified address:

```
user@host1> op hostname dns 10.168.71.249
Name:
249.71.168.10.IN-ADDR.ARPA domain name pointer router1
```

## Example: Finding LSPs to Multiple Destinations Using an Op Script

This example uses an op script to check for label-switched paths (LSPs) to multiple destinations.

- [Requirements on page 491](#)
- [Overview and Op Script on page 491](#)
- [Configuration on page 495](#)
- [Verification on page 495](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Op Script

The following example script, which is shown in both XSLT and SLAX, checks for LSPs to multiple destinations. The script takes one mandatory command-line argument, the address specifying the LSP endpoint. The address argument can include an optional prefix length. If no address is specified, the script generates an error message and halts execution.

The **get-configuration** variable stores the remote procedure call (RPC) to retrieve the **[edit protocols mpls]** hierarchy level of the device's committed configuration. This

configuration is stored in the **config** variable. The **get-route-information** variable stores the RPC equivalent of the **show route address terse** operational mode command, where the value of the **destination** tag specifies *address*. The script sets this value to the address specified by the user on the command line. The script invokes the **get-route-information** RPC and stores the output in the **rpc-out** variable. If **rpc-out** does not contain any errors, the script examines all host route entries present at the **route-table/rt/rt-destination** node.

For each host route entry, if an LSP to the destination is configured in the active configuration, the script generates a "Found" message with the destination address and corresponding LSP name in the output. If an LSP to the destination is not configured, the output generates a "Missing" message containing the destination address and hostname.

#### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0" version="1.0">
  <xsl:import href="../../import/junos.xml"/>

  <xsl:variable name="arguments">
    <argument>
      <name>address</name>
      <description>LSP endpoint</description>
    </argument>
  </xsl:variable>
  <xsl:param name="address"/>
  <xsl:template match="/">
    <op-script-output>
      <xsl:choose>
        <xsl:when test="$address = ''">
          <xnm:error>
            <message>missing mandatory argument 'address'</message>
          </xnm:error>
        </xsl:when>
        <xsl:otherwise>
          <xsl:variable name="get-configuration">
            <get-configuration database="committed">
              <configuration>
                <protocols>
                  <mpls/>
                </protocols>
              </configuration>
            </get-configuration>
          </xsl:variable>
          <xsl:variable name="config"
            select="jcs:invoke($get-configuration)"/>
          <xsl:variable name="mpls" select="$config/protocols/mpls"/>
          <xsl:variable name="get-route-information">
            <get-route-information>
              <terse/>
              <destination>
                <xsl:value-of select="$address"/>
              </destination>
            </get-route-information>
          </xsl:variable>
        </xsl:otherwise>
      </xsl:choose>
    </op-script-output>
  </xsl:template>
</xsl:stylesheet>
```

```

    </get-route-information>
  </xsl:variable>
  <xsl:variable name="rpc-out"
    select="jcs:invoke($get-route-information)"/>
  <xsl:choose>
    <xsl:when test="$rpc-out//xnm:error">
      <xsl:copy-of select="$rpc-out//xnm:error"/>
    </xsl:when>
    <xsl:otherwise>
      <xsl:for-each select="$rpc-out/route-table/rt/rt-destination">
        <xsl:choose>
          <xsl:when test="contains(.,'/32')">
            <xsl:variable name="dest"
              select="substring-before(.,'/')"/>
            <xsl:variable name="lsp"
              select="$mpls/label-switched-path[to = $dest]"/>
            <xsl:choose>
              <xsl:when test="$lsp">
                <output>
                  <xsl:value-of select="concat('Found: ', $dest,
                    ' (', $lsp/to, ') --&gt; ', $lsp/name)"/>
                </output>
              </xsl:when>
              <xsl:otherwise>
                <xsl:variable name="name"
                  select="jcs:hostname($dest)"/>
                <output>
                  <xsl:value-of select="concat('Name: ', $name)"/>
                </output>
                <output>
                  <xsl:value-of select="concat('Missing: ', $dest)"/>
                </output>
              </xsl:otherwise>
            </xsl:choose>
          </xsl:when>
          <xsl:otherwise>
            <output>
              <xsl:value-of select="concat('Not a host route: ', .)"/>
            </output>
          </xsl:otherwise>
        </xsl:choose>
      </xsl:for-each>
    </xsl:otherwise>
  </xsl:choose>
</op-script-output>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

```

```

import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "address";
    <description> "LSP endpoint";
  }
}
param $address;
match / {
  <op-script-output> {
    if ($address = '') {
      <xnm:error> {
        <message> "missing mandatory argument 'address'";
      }
    } else {
      var $get-configuration = {
        <get-configuration database="committed"> {
          <configuration> {
            <protocols> {
              <mpls>;
            }
          }
        }
      }
      var $config = jcs:invoke($get-configuration);
      var $mpls = $config/protocols/mpls;
      var $get-route-information = {
        <get-route-information> {
          <terse>;
          <destination> $address;
        }
      }
      var $rpc-out = jcs:invoke($get-route-information);
      if ($rpc-out//xnm:error) {
        copy-of $rpc-out//xnm:error;
      } else {
        for-each ($rpc-out/route-table/rt/rt-destination) {
          if (contains(.,'/32')) {
            var $dest = substring-before(.,'/');
            var $lsp = $mpls/label-switched-path[to = $dest];
            if ($lsp) {
              <output> 'Found: ' _ $dest _ ' (' _ $lsp/to _
                ' _ $lsp/name;
            } else {
              var $name = jcs:hostname($dest);
              <output> 'Name: ' _ $name;
              <output> 'Missing: ' _ $dest;
            }
          } else {
            <output> 'Not a host route: ' _ .;
          }
        }
      }
    }
  }
}

```

## Configuration

- Step-by-Step Procedure** To download, enable, and test the script:
1. Copy the XSLT or SLAX script into a text file, name the file **lsp.xml** or **lsp.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.
  2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **lsp.xml** or **lsp.slax** as appropriate.
 

```
[edit system scripts op]
user@host# set file lsp.(slax | xml)
```
  3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.
 

```
[edit]
user@host# commit and-quit
```
  4. Execute the op script by issuing the **op lsp address address** operational mode command.

## Verification

### Verifying Script Execution

- Purpose** Verify that the script behaves as expected.
- Action** Issue the **op lsp address address** operational mode command to execute the script. The output varies depending on the configuration.
- ```
user@R4> op lsp address 10.168.215.0/24
Found: 192.168.215.1 (192.168.215.1) --> R4>R1
Found: 192.168.215.2 (192.168.215.2) --> R4>R2
Name: R3
Missing: 10.168.215.3
Name: R5
Missing: 10.168.215.4
Name: R6
Missing: 10.168.215.5
```

## Example: Restarting an FPC Using an Op Script

This example uses an op script to restart a Flexible PIC Concentrator (FPC).

- [Requirements on page 496](#)
- [Overview and Op Script on page 496](#)
- [Configuration on page 497](#)
- [Verification on page 497](#)

## Requirements

This example uses a device running Junos OS that contains a Flexible PIC Concentrator (FPC) or equivalent component.

## Overview and Op Script

The following script, which is shown in both XSLT and SLAX formats, restarts an FPC given the slot number in which the FPC resides. The user provides the slot number in the command-line interface (CLI) when the script is invoked. The script stores the slot number as the value of the parameter **slot** and constructs the **request chassis fpc** command string to include the slot number of the FPC to restart. There is no Junos Extensible Markup Language (XML) equivalent for the **request chassis** commands. Therefore, this script invokes the **request chassis fpc** command directly rather than using a remote procedure call (RPC).

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
  <xsl:import href="../import/junos.xsl"/>

  <xsl:variable name="arguments">
    <argument>
      <name>slot</name>
      <description>Slot number of the FPC</description>
    </argument>
  </xsl:variable>
  <xsl:param name="slot"/>
  <xsl:template match="/">
    <op-script-results>
      <xsl:variable name="restart">
        <command>
          <xsl:value-of select="concat('request chassis fpc slot ', $slot, '
            restart')"/>
        </command>
      </xsl:variable>
      <xsl:variable name="result" select="jcs:invoke($restart)"/>
      <output>
        <xsl:text>Restarting the FPC in slot </xsl:text>
        <xsl:value-of select="$slot"/>
        <xsl:text>. </xsl:text>
        <xsl:text>To verify, issue the "show chassis fpc" command.</xsl:text>
      </output>
    </op-script-results>
  </xsl:template>
</xsl:stylesheet>
```

### SLAX Syntax

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
```



```

import "../import/junos.xml";

var $arguments = {
  <argument> {
    <name> "slot";
    <description> "Slot number of the FPC";
  }
}
param $slot;
match / {
  <op-script-results> {
    var $restart = {
      <command> 'request chassis fpc slot ' _ $slot _ ' restart';
    }
    var $result = jcs:invoke($restart);
    <output> {
      expr "Restarting the FPC in slot ";
      expr $slot;
      expr ". ";
      expr "To verify, issue the \"show chassis fpc\" command.";
    }
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **restart-fpc.xml** or **restart-fpc.slax** as appropriate, and download it to the **/var/db/scripts/op/** directory on the device.

Only users who belong to the Junos OS **super-user** login class can access and edit files in this directory.

2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **restart-fpc.xml** or **restart-fpc.slax** as appropriate.

```

[edit system scripts op]
user@host# set file restart-fpc.(slax | xml)

```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```

[edit]
user@host# commit and-quit

```

4. Execute the op script by issuing the **op restart-fpc slot slot-number** operational mode command.

## Verification

### Verifying Op Script Execution

**Purpose** Verify that the FPC has been restarted and is currently online.

**Action** Execute the op script by issuing the **op filename** operational mode command. Supply the **slot** number of the FPC as an argument.

```
user@host> op restart-fpc slot 0
```

When you execute the script, you should see output similar to the following:

Restarting the FPC in slot 0. To verify, issue the "show chassis fpc" command.

Issue the **show chassis fpc detail fpc-slot** operational mode command.

```
user@host> show chassis fpc detail 0
```

```
Slot 0 information:
  State                               Online
  Temperature                         36 degrees C / 96 degrees F
  Total CPU DRAM                      1024 MB
  Total RLDRAM                        256 MB
  Total DDR DRAM                      4096 MB
  Start time:                         2009-08-11 21:20:30 PDT
  Uptime:                             0 hours, 1 minutes, 50 seconds
  Max Power Consumption               335 Watts
```

**Meaning** The **show chassis fpc detail** command output displays the state, start time, uptime, and characteristics for the FPC. Verify that the FPC was restarted by checking the start time and uptime of the FPC. Verify the status of the restarted FPC by checking the state. If the status is **Present**, the FPC is coming up but is not yet online. If the status is **Online**, the FPC is online and running.

---

## Example: Exporting Files Using an Op Script

The op script in this example uses the Junos XML protocol **file-put** operation to write to a file on a remote server and on the local device.

- [Requirements on page 498](#)
- [Overview and Op Script on page 498](#)
- [Configuration on page 502](#)
- [Verification on page 503](#)

### Requirements

This example uses a device running Junos OS.

### Overview and Op Script

The Junos XML protocol **file-put** operation creates a file and writes the specified contents to that file. The basic syntax for using the **file-put** command is as follows:

```
<rpc>
  <file-put>
    <delete-if-exist />
    <encoding>value</encoding>
    <filename>value</filename>
    <permission>value</permission>
```

```

        <file-contents>file</file-contents>
    </file-put>
</rpc>

```

The following tag elements are used with the **file-put** command. These tags can be placed in any order with the exception of **file-contents**. The **file-contents** tag element must be the last tag in list.

- **delete-if-exist**—(Optional) If included, any existing file is overwritten. If the tag is omitted, an error is returned if an existing file is encountered.
- **encoding**—(Mandatory) Specifies the type of encoding used. You can use **ASCII** or **base64** encoding.
- **filename**—(Mandatory) Within this tag, you include the full or relative path and filename of the file to create. When you use a relative path, the specified path is relative to the user's home directory. If the specified directory does not exist, the system returns a "directory does not exist" error.
- **permission**—(Optional) Sets the file's UNIX permission on the remote server. For example, to apply read/write access for the user, and read access to others, you would set the permission value to 0644. For a full explanation of UNIX permissions, see the **chmod** command.
- **file-contents**—(Mandatory) The ASCII or base64 encoded file contents to export. This must be the last tag in the list.

#### XSLT Syntax

The following sample script executes a Junos XML API request and exports the results to a file on a remote device and a file on the local device. The script takes three arguments: the IP address or hostname of the remote device, the filename, and the file encoding. The **arguments** variable is declared at the global level of the script so that the argument names and descriptions are visible in the command-line interface (CLI).

The script invokes the Junos XML API **<get-software-information>** request on the local device and stores the result in the **result** variable. The script declares the **fileput** variable, which contains the remote procedure call (RPC) for the **file-put** operation. The command-line arguments define the values for the **filename** and **encoding** tag elements. If the mandatory argument **myhost** is missing, the script issues an error and halts execution. Otherwise, the script prompts for the username and password that will be used to connect to the remote device.

If connection to the remote device is successful, the script executes the RPC within the context of the connection handle. The output of the **file-put** operation, which is the result of the **jcs:execute()** function, is stored in the **out** variable. If the operation encounters an error, the script prints the error to the CLI. If the **file-put** operation is successful, the contents specified by the **file-contents** tag element are exported to the specified file on the remote device. The connection to the remote host is then closed. The script also exports the contents to an identical file on the local device.

The sample script includes the optional tag elements **permission** and **delete-if-exist** for the **file-put** operation. By including the **delete-if-exist** tag, the script overwrites any existing file of the same name on the remote and local hosts. In this example, the **permission** tag is set to **0644**.

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet
    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"

```

```
xmlns:junos="http://xml.juniper.net/junos/*/junos"
xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0" version="1.0">

<xsl:import href="../import/junos.xsl"/>

<xsl:variable name="arguments">
  <argument>
    <name>myhost</name>
    <description>IP address or hostname of the remote host</description>
  </argument>
  <argument>
    <name>filename</name>
    <description>name of destination file</description>
  </argument>
  <argument>
    <name>encoding</name>
    <description>ascii or base64</description>
  </argument>
</xsl:variable>

<xsl:param name="myhost"/>
<xsl:param name="filename"/>
<xsl:param name="encoding"/>

<xsl:template match="/">
  <op-script-results>

    <xsl:variable name="rpc">
      <get-software-information/>
    </xsl:variable>
    <xsl:variable name="result" select="jcs:invoke($rpc)"/>

    <xsl:variable name="fileput">
      <file-put>
        <filename>
          <xsl:value-of select="$filename"/>
        </filename>
        <encoding>
          <xsl:value-of select="$encoding"/>
        </encoding>
        <permission>0644</permission>
        <delete-if-exist/>
        <file-contents>
          <xsl:value-of select="$result"/>
        </file-contents>
      </file-put>
    </xsl:variable>

    <xsl:choose>
      <xsl:when test="$myhost = ''">
        <xnm:error>
          <message>missing mandatory argument 'myhost'</message>
        </xnm:error>
      </xsl:when>
      <xsl:otherwise>
        <xsl:variable name="username" select="jcs:get-input('Enter username:
')"/>
        <xsl:variable name="pw" select="jcs:get-secret('Enter password:
')"/>
        <xsl:variable name="connect" select="jcs:open($myhost, $username,
```

```

$pw)"/>
    <xsl:choose>
    <xsl:when test="$connect">
        <output>Connected to host. Exporting file... </output>
        <xsl:variable name="out" select="jcs:execute($connect,
$fileput)"/>
        <xsl:choose>
        <xsl:when test="$out//xnm:error">
            <xsl:copy-of select="($out//xnm:error)"/>
        </xsl:when>
        <xsl:otherwise>
            <output>
                <xsl:value-of select="$out"/>
            </output>
        </xsl:otherwise>
        </xsl:choose>
        <xsl:value-of select="jcs:close($connect)"/>
    </xsl:when>
    <xsl:otherwise>
        <output>No connection to host.</output>
    </xsl:otherwise>
    </xsl:choose>
    </xsl:otherwise>
</xsl:choose>

    <xsl:variable name="local-out" select="jcs:invoke($fileput)"/>
    <output>
        <xsl:value-of select="concat('Saving file on local host\n',
$local-out)"/>
    </output>
</op-script-results>
</xsl:template>
</xsl:stylesheet>

```

## SLAX Syntax

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
    <argument> {
        <name> "myhost";
        <description> "IP address or hostname of the remote host";
    }
    <argument> {
        <name> "filename";
        <description> "name of destination file";
    }
    <argument> {
        <name> "encoding";
        <description> "ascii or base64";
    }
}

param $myhost;
param $filename;
param $encoding;

match / {
    <op-script-results> {

```

```

var $rpc = <get-software-information>;
var $result = jcs:invoke($rpc);

var $fileput = {
  <file-put> {
    <filename>$filename;
    <encoding>$encoding;
    <permission>'0644';
    <delete-if-exist>;
    <file-contents>$result;
  }
}

if ($myhost = '') {
  <xnm:error> {
    <message> "missing mandatory argument 'myhost'";
  }
}
else {
  var $username = jcs:get-input("Enter username: ");
  var $pw = jcs:get-secret("Enter password: ");
  var $connect = jcs:open($myhost, $username, $pw);

  if ($connect) {
    <output> "Connected to host. Exporting file... \n";
    var $out = jcs:execute($connect, $fileput);
    if ($out//xnm:error) {
      copy-of ($out//xnm:error);
    }
    else {
      <output> $out;
    }
    expr jcs:close($connect);
  }
  else {
    <output> "No connection to host.";
  }
}

var $local-out = jcs:invoke($fileput);
<output> "Saving file on local host\n" _ $local-out;
}
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **export.xml** or **export.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **export.xml** or **export.slax** as appropriate.

```

[edit system scripts op]
user@host# set file export.(slax | xml)

```

3. Issue the **commit and-quit** command.

```
[edit]
user@host# commit and-quit
```

4. Execute the op script by issuing the **op export** operational mode command and include any necessary arguments.

## Verification

- [Verifying the Op Script Arguments on page 503](#)
- [Verifying Op Script Execution on page 503](#)

### Verifying the Op Script Arguments

**Purpose** Verify that the argument names and descriptions show up in the CLI.

**Action** Issue the **op export ?** operational mode command. The CLI lists the possible completions for the script arguments based on the definitions within the global **arguments** variable in the script.

```
user@host> op export ?
Possible completions:
<[Enter]>      Execute this command
<name>         Argument name
detail         Display detailed output
encoding       ascii or base64
filename       name of destination file
myhost         IP address or hostname of the remote host
|             Pipe through a command
```

### Verifying Op Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** Issue the **op export myhost *host* encoding *encoding* filename *file*** operational mode command, and include the appropriate username and password when prompted. If script execution is successful, the result of the <**get-software-information**> RPC request is written to the file on the remote device and also on the local device. For example:

```
root@host> op export myhost router1 encoding ascii filename /var/log/host-version.txt
Enter username: root
Enter password:
Connected to host. Exporting file...
```

```
/var/log/host-version.txt
Saving file on local host
```

```
/var/log/host-version.txt
```

If you fail to supply the IP address or hostname of the remote device in the command-line arguments, the script issues an error and halts execution.

```
root@host> op export
error: missing mandatory argument 'myhost'
```

If you omit the **delete-if-exist** child tag of the **file-put** operation, and the specified file already exists, the script reports an error.

```
root@host> op export myhost router1 encoding ascii filename /var/log/host-version.txt
Enter username: root
Enter password:
Connected to host. Exporting file...
```

```
Destination file exists
Saving file on local host
```

```
Destination file exists
```

If you execute the script and include a directory path that does not exist on either the remote or the local host, the script reports an error.

```
root@host> op export myhost router1 encoding ascii filename /var/test/host-version.txt
Enter username: root
Enter password:
Connected to host. Exporting file...
```

```
Destination directory does not exist: /var/test
Saving file on local host
```

```
Destination directory does not exist: /var/test
```

**Related  
Documentation**

- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
- [Example: Importing Files Using an Op Script on page 505](#)



---

## Example: Importing Files Using an Op Script

---

The op script in this example uses the Junos XML protocol **file-get** operation to read the contents of a file from a remote server.

- [Requirements on page 505](#)
- [Overview and Op Script on page 505](#)
- [Configuration on page 508](#)
- [Verification on page 509](#)

### Requirements

This example uses a device running Junos OS.

### Overview and Op Script

The Junos XML protocol **file-get** operation reads the contents of a file. The basic syntax for using the **file-get** command is as follows:

```
<rpc>
  <file-get>
    <filename>value</filename>
    <encoding>value</encoding>
  </file-get>
</rpc>
```

The following tag elements are used with the **file-get** command.

- **encoding**—(Mandatory) Specifies the type of encoding used. You can use **ASCII**, **base64**, or **raw** encoding.
- **filename**—(Mandatory) Within this tag, you include the full or relative path and filename of the file to import. When you use a relative path, the specified path is relative to the **/var/tmp/** directory if the **file-get** operation is executed locally. If the operation is executed remotely within the context of a connection handle, the path is relative to the user's home directory.



**NOTE:** When you use ASCII encoding, the **file-get** operation converts any control characters in the imported file to the Unicode character 'SECTION SIGN' (U+00A7).

---

**XSLT Syntax** The following sample script connects to a remote device and reads the contents of the specified file. The script takes three arguments: the IP address or hostname of the remote device, the filename, and the file encoding. The **arguments** variable is declared at the global level of the script so that the argument names and descriptions are visible in the command-line interface (CLI).

The script declares the **fileget** variable, which contains the remote procedure call (RPC) for the **file-get** operation. The command-line arguments define the values for the **filename** and **encoding** tag elements. If the mandatory argument **myhost** is missing, the script issues an error and halts execution. Otherwise, the script prompts for the username and password that will be used to connect to the remote device.

If connection to the remote device is successful, the script executes the RPC within the context of the connection handle. The output of the **file-get** operation, which is the result of the **jcs:execute()** function, is stored in the **out** variable. If the operation encounters an error, the script prints the error to the CLI. If the **file-get** operation is successful, the contents of the file are stored in the **out** variable, which is printed to the CLI. The connection to the remote host is then closed.

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0" version="1.0">
  <xsl:import href="../../import/junos.xsl"/>

  <xsl:variable name="arguments">
    <argument>
      <name>myhost</name>
      <description>IP address or hostname of the remote host</description>
    </argument>
    <argument>
      <name>filename</name>
      <description>name of file</description>
    </argument>
    <argument>
      <name>encoding</name>
      <description>ascii, base64, or raw</description>
    </argument>
  </xsl:variable>

  <xsl:param name="myhost"/>
  <xsl:param name="filename"/>
  <xsl:param name="encoding"/>

  <xsl:template match="/">
    <op-script-results>
      <xsl:variable name="fileget">
        <file-get>
          <filename>
            <xsl:value-of select="$filename"/>
          </filename>
          <encoding>
            <xsl:value-of select="$encoding"/>
          </encoding>
        </file-get>
      </xsl:variable>
    </op-script-results>
  </template>
</xsl:stylesheet>
```

```

    </file-get>
  </xsl:variable>
  <xsl:choose>
    <xsl:when test="$myhost = ''">
      <xnm:error>
        <message>missing mandatory argument 'myhost'</message>
      </xnm:error>
    </xsl:when>
    <xsl:otherwise>
      <xsl:variable name="username" select="jcs:get-input('Enter username: ')" />

      <xsl:variable name="pw" select="jcs:get-secret('Enter password: ')" />
      <xsl:variable name="connect" select="jcs:open($myhost, $username, $pw)" />

      <xsl:choose>
        <xsl:when test="$connect">
          <output>Connected to host. Reading file...
          </output>
          <xsl:variable name="out" select="jcs:execute($connect, $fileget)" />
          <xsl:choose>
            <xsl:when test="$out//xnm:error">
              <xsl:copy-of select="$out//xnm:error" />
            </xsl:when>
            <xsl:otherwise>
              <output>
                <xsl:value-of select="concat('File contents: ', $out)" />
              </output>
            </xsl:otherwise>
          </xsl:choose>
          <xsl:value-of select="jcs:close($connect)" />
        </xsl:when>
        <xsl:otherwise>
          <output>No connection to host.</output>
        </xsl:otherwise>
      </xsl:choose>
    </xsl:otherwise>
  </xsl:choose>
</op-script-results>
</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "myhost";
    <description> "IP address or hostname of the remote host";
  }
  <argument> {
    <name> "filename";
    <description> "name of file";
  }
  <argument> {

```

```

        <name> "encoding";
        <description> "ascii, base64, or raw";
    }
}

param $myhost;
param $filename;
param $encoding;

match / {
    <op-script-results> {
        var $fileget = {
            <file-get> {
                <filename>$filename;
                <encoding>$encoding;
            }
        }

        if ($myhost = '') {
            <xnm:error> {
                <message> "missing mandatory argument 'myhost'";
            }
        }
        else {
            var $username = jcs:get-input("Enter username: ");
            var $pw = jcs:get-secret("Enter password: ");
            var $connect = jcs:open($myhost, $username, $pw);

            if ($connect) {
                <output> "Connected to host. Reading file... \n";
                var $out = jcs:execute($connect, $fileget);
                if ($out//xnm:error) {
                    copy-of $out//xnm:error;
                }
                else {
                    <output> "File contents: " _ $out;
                }
                expr jcs:close($connect);
            }
            else {
                <output> "No connection to host.";
            }
        }
    }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **import.xml** or **import.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **import.xml** or **import.slax** as appropriate.

```

[edit system scripts op]
user@host# set file import.(slax | xml)

```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```
[edit]
user@host# commit and-quit
```

4. Execute the op script by issuing the **op import** operational mode command and include any necessary arguments.

## Verification

### Verifying the Script Arguments

**Purpose** Verify that the argument names and descriptions show up in the CLI.

**Action** Issue the **op import ?** operational mode command. The CLI lists the possible completions for the script arguments based on the definitions within the global **arguments** variable in the script.

```
user@host> op import ?
Possible completions:
<[Enter]>          Execute this command
<name>             Argument name
detail             Display detailed output
encoding           ascii, base64, or raw
filename           name of file
myhost             IP address or hostname of the remote host
|                 Pipe through a command
```

### Verifying Op Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** Issue the **op import myhost *host* encoding *encoding* filename *file*** operational mode command, and include the appropriate username and password when prompted. If script execution is successful, the contents of the requested file are displayed. For example:

```
root@host> op import myhost router1 encoding ascii filename /var/db/scripts/op/test.slax
Enter username: root
Enter password:
Connected to host. Reading file...
File contents:

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";
...
```

If you fail to supply the IP address or hostname of the remote device in the command-line arguments, the script issues an error and halts execution.

```
root@host> op import
error: missing mandatory argument 'myhost'
```

Also, if the specified path or file does not exist, the script issues an error.

```
root@host> op import myhost router1 encoding ascii filename /var/db/scripts/op/test1.slax
Enter username: root
Enter password:
Connected to host. Reading file...
File contents:
```

```
Failed to open file (/var/db/scripts/op/test1.slax): No such file or directory
```

- Related Documentation**
- [Declaring and Using Command-Line Arguments in Op Scripts on page 452](#)
  - [Example: Exporting Files Using an Op Script on page 498](#)

---

## Example: Searching Files Using an Op Script

This sample script searches a file on a device running Junos OS for lines matching a given regular expression. The example uses the **jcs:grep** template in an op script.

- [Requirements on page 510](#)
- [Overview and Op Script on page 511](#)
- [Configuration on page 513](#)
- [Verification on page 513](#)

## Requirements

This example uses a device running Junos OS.

## Overview and Op Script

The **jcs:grep** template searches an ASCII file for lines matching a regular expression. The template resides in the **junos.xml** import file, which is included with the standard Junos OS installation available on all switches, routers, and security devices running Junos OS. To use the **jcs:grep** template in a script, you must import the **junos.xml** file into the script and map the **jcs** prefix to the namespace identified by the URI <http://xml.juniper.net/junos/commit-scripts/1.0>.

In this example, all values required for the **jcs:grep** template are defined as global parameters. The values for the parameters are passed into the script as command-line arguments. The following script defines two parameters, **filename** and **pattern**, which store the values of the input file path and the regular expression. If you omit either argument when you execute the script, the script generates an error and halts execution. Otherwise, the script calls the **jcs:grep** template and passes in the supplied arguments.

If the regular expression contains a syntax error, the **jcs:grep** template generates an **error: regex error** message for each line in the file. If the regular expression syntax is valid, the template parses the input file. For each match, the template adds a **<match>** element, which contains **<input>** and **<output>** child tags, to the result tree. The template writes the matching string to the **<output>** child element and writes the corresponding matching line to the **<input>** child element:

```
<match> {
  <input>
  <output>
}
```

In the SLAX script, the **:=** operator copies the results of the **jcs:grep** template call to a temporary variable and runs the **node-set** function on that variable. The **:=** operator ensures that the **results** variable is a node-set rather than a result tree fragment so that the script can access the contents. The XSLT script explicitly calls out the equivalent steps. The script then loops through all resulting input elements and prints each match.

### XSLT Syntax

```
<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0" version="1.0">
  <xsl:import href="../import/junos.xml"/>

  <xsl:variable name="arguments">
    <argument>
      <name>filename</name>
      <description>name of file in which to search for the specified pattern
      </description>
    </argument>
    <argument>
      <name>pattern</name>
      <description>regular expression</description>
    </argument>
  </xsl:variable>

  <xsl:param name="filename"/>
```

```

<xsl:param name="pattern"/>

<xsl:template match="/">

  <op-script-results>
    <xsl:choose>
      <xsl:when test="$filename = "">
        <xnm:error>
          <message>missing mandatory argument 'filename'</message>
        </xnm:error>
      </xsl:when>
      <xsl:when test="$pattern = ';'>
        <xnm:error>
          <message>missing mandatory argument 'pattern'</message>
        </xnm:error>
      </xsl:when>
      <xsl:otherwise>
        <xsl:variable name="results-temp">
          <xsl:call-template name="jcs:grep">
            <xsl:with-param name="filename" select="$filename"/>
            <xsl:with-param name="pattern" select="$pattern"/>
          </xsl:call-template>
        </xsl:variable>
        <xsl:variable xmlns:ext="http://xmlsoft.org/XSLT/namespace"
          name="results" select="ext:node-set($results-temp)"/>
        <output>
          <xsl:value-of select="concat('Search for ', $pattern, ' in ', $filename)"/>
        </output>
        <xsl:for-each select="$results//input">
          <output>
            <xsl:value-of select="."/>
          </output>
        </xsl:for-each>
      </xsl:otherwise>
    </xsl:choose>
  </op-script-results>

</xsl:template>
</xsl:stylesheet>

```

**SLAX Syntax**

```

version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

import "../import/junos.xsl";

var $arguments = {
  <argument> {
    <name> "filename";
    <description> "name of file in which to search for the specified
pattern";
  }
  <argument> {
    <name> "pattern";
    <description> "regular expression";
  }
}

```



```

}

param $filename;
param $pattern;

match / {
  <op-script-results> {
    if ($filename = '') {
      <xnm:error> {
        <message> "missing mandatory argument 'filename'";
      }
    }
    else if ($pattern = '') {
      <xnm:error> {
        <message> "missing mandatory argument 'pattern'";
      }
    }
    else {
      var $results := { call jcs:grep($filename, $pattern); }

      <output> "Search for " _ $pattern _ " in " _ $filename;
      for-each ($results//input) {
        <output> .;
      }
    }
  }
}

```

## Configuration

### Step-by-Step Procedure

To download, enable, and run the script:

1. Copy the XSLT or SLAX script into a text file, name the file **grep.xml** or **grep.slax** as appropriate, and download it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **grep.xml** or **grep.slax** as appropriate.

```

[edit system scripts op]
user@host# set file grep.(slax | xml)

```

3. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```

[edit]
user@host# commit and-quit

```

4. Execute the op script by issuing the **op grep filename filename pattern pattern** operational mode command.

## Verification

### Verifying the Script Arguments

**Purpose** Verify that the argument names and descriptions appear in the command-line interface (CLI) help.

**Action** Issue the **op grep ?** operational mode command. The CLI lists the possible completions for the script arguments based on the definitions within the global variable **arguments** in the script.

```
user@host> op grep
Possible completions:
  <[Enter]>      Execute this command
  <name>         Argument name
  detail        Display detailed output
  filename      name of file in which to search for the specified pattern

  pattern       regular expression
  |            Pipe through a command
```

---

### Verifying Op Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** If you issue the **op grep** command, but you fail to supply either the filename or the regex pattern, the script issues an error message and halts execution. For example:

```
user@host> op grep filename /var/log/messages
error: missing mandatory argument 'pattern'

user@host> op grep pattern SNMP_TRAP_LINK_DOWN
error: missing mandatory argument 'filename'
```

When you issue the **op grep filename *filename* pattern *pattern*** command, the script lists all lines from the input file that match the regular expression.

```
user@host> op grep filename /var/log/messages pattern SNMP_TRAP_LINK_DOWN
Search for SNMP_TRAP_LINK_DOWN in /var/log/messages
Feb 24 09:04:00 host mib2d[1325]: SNMP_TRAP_LINK_DOWN: ifIndex 543, ifAdminStatus
down(2), ifOperStatus down(2), ifName
1t-0/1/0.9
Feb 24 09:04:00 host mib2d[1325]: SNMP_TRAP_LINK_DOWN: ifIndex 542, ifAdminStatus
down(2), ifOperStatus down(2), ifName
1t-0/1/0.10
```

- Related Documentation**
- [SLAX Templates Overview on page 66](#)
  - [jcs:grep Template on page 234](#)
  - [regex\(\) Function \(SLAX and XSLT\) on page 216](#)

# Provisioning Services Using Service Template Automation

- [Service Template Automation Overview on page 515](#)
- [Example: Configuring Service Template Automation on page 516](#)

## Service Template Automation Overview

---

Starting in Junos OS Release 12.3, you can use service template automation (STA) to provision services such as VPLS VLAN, Layer 2 and Layer 3 VPNs, and IPsec across similar platforms running Junos OS. Service template automation uses the **service-builder.slax** op script to transform a user-defined service template definition into a uniform API, which you can then use to configure and provision services on similar platforms running Junos OS. This enables you to create a service template on one device, generalize the parameters, and then quickly and uniformly provision that service on other devices. This decreases the time required to configure the same service on multiple devices, and reduces configuration errors associated with manually configuring each device.

The following process outlines how to use service template automation to provision services:

1. Create a service template definition.
2. Execute the **service-builder.slax** script and define service-specific instance parameters.
3. Generate the service interface, which automatically builds the required interface (API) from the template.
4. Enable the service interface on each device where the service is required.
5. Provision systems by invoking the service interface using NETCONF and supplying the service parameter values.

You create a new service template by configuring the hierarchies for the actual service to be provisioned on a device running Junos OS. Service template hierarchies are configured at the **[edit groups]** hierarchy level. When creating the service template:

- Do not include **apply-groups** or **apply-macro** statements.

- Do not include any statements that are supported on the current device that are not also supported on the devices where the service will be provisioned (for example dual Routing Engine versus single Routing Engine).
- Commit the configuration. The service template group configuration is read from the committed configuration.

Once you create the basic service template definition, you invoke the **service-builder.slax** op script. The script reads the service template information from the committed configuration and uses an interactive interface to help you build and generate the service API. You have the option to parameterize every variable in the service template or only selected variables. For each selected variable, you create a generic service template parameter. The **service-builder.slax** script guides you through the creation and configuration of each parameter.

After you define the service template parameters, you generate the service interface. This creates a platform-specific service op script. If the **load-scripts-from-flash** statement is configured, the generated service script is stored in the **/config/scripts/op** directory in flash memory. Otherwise, the generated script is stored in the **/var/db/scripts/op** directory on the hard disk.

To enable the service interface on a device, you enable the generated service script in the configuration as you would any op script. You can enable the service interface on the local device using the **service-builder.slax** script or by manually updating the configuration. To enable the service interface on a similar platform, you must copy the generated service script to the corresponding directory on the new device and enable the service script in the configuration.

To provision the service on a device, invoke the service interface using NETCONF, and supply the necessary values for each parameter. Alternatively, you can invoke the service interface in the CLI by executing the service script and supplying the necessary values for each parameter as command-line arguments to the script. You can direct the service script to create a new service configuration, or update or delete an existing service configuration. The service script makes the changes to the candidate configuration and then commits the configuration. The service script does not support the context-sensitive help and auto-completion features available in the Junos OS CLI.

**Related  
Documentation**

- [Example: Configuring Service Template Automation on page 516](#)

---

## Example: Configuring Service Template Automation

---

This example shows how to use service template automation to provision services across similar platforms running Junos OS.

- [Requirements on page 517](#)
- [Overview on page 517](#)
- [Configuration on page 518](#)
- [Verification on page 526](#)
- [Troubleshooting on page 527](#)

## Requirements

- Two MX Series devices running Junos OS Release 12.3 or later.

## Overview

This example uses service template automation to provision services on an MX Series router. To use the service template automation **service-builder.slax** script, you must first copy the script to the `/var/db/scripts/op` or `/config/scripts/op` directory and enable the script on the device.

The following process outlines how to use service template automation to provision services:

1. Create a service template definition.
2. Execute the **service-builder.slax** script, and define service-specific instance parameters.
3. Generate the service interface.
4. Enable the service interface on each device where the service is required.
5. Provision systems by invoking the service interface using NETCONF and supplying the service parameter values.

This example creates a new VPN service interface on an MX Series device running Junos OS Release 12.3 and provisions the service on a second MX Series device running Junos OS Release 12.3. You configure service template definitions under the **[edit groups]** hierarchy level. For this example, the service name is **vpn-service**, and the template group name is **vpn-service-template-group**. The **load merge terminal** configuration mode command loads the service template configuration hierarchies into the candidate configuration, which is then committed.

Once you create the initial service template, you execute the **service-builder.slax** script. The script prompts for the service name and the template group name, and then reads the service template configuration from the committed configuration.

The **service-builder.slax** script interface consists of two menus: **Main Menu** and **Hierarchies Menu**. Within the **Main Menu**, you can review the variables defined in the service template configuration, or you can build or enable the service API. The **Build Service API** menu option displays the **Hierarchies Menu**, which steps you through the parameterization of the variables. The default is to parameterize every variable, or you can choose to parameterize selected variables. If you must exit the **service-builder.slax** script while building the service API, you must finish configuring all the parameters for the current hierarchy in order to save that hierarchy configuration when you exit using the **Quit** option. Then you can finish configuring any incomplete hierarchies at a later time. This example parameterizes two variables: the interface name and the interface description. After the parameters are specified, the service builder script generates the service script.

The **Enable Service API** menu option enables the service script on the local device. To enable the service script on the second MX Series device, the generated service script is copied to the `/var/db/scripts/op` directory on the second device, and the script is enabled

in the configuration. If the **load-scripts-from-flash** statement is configured, the script must be copied to the corresponding directory on the flash drive instead.

NETCONF is used to provision the service on the remote MX Series device. The NETCONF remote procedure call (RPC) action depends on whether the service is a new service or an existing service. Supported actions include **create**, **update**, and **delete**. This example creates a new service. If the given service is already provisioned on the device and you are updating or deleting the service parameters, you can alter the RPC to perform these actions.

## Configuration

- [Storing and Enabling the Service Builder Script on page 518](#)
- [Configuring the Service Template Definition on page 519](#)
- [Configuring and Generating the Service Interface on page 520](#)
- [Verifying the Service Interface on page 522](#)
- [Enabling the Service Interface on page 523](#)
- [Provisioning the Service Using NETCONF on page 524](#)
- [Updating or Deleting Services Using NETCONF on page 525](#)

---

### Storing and Enabling the Service Builder Script

#### Step-by-Step Procedure

The Junos OS installation includes the **service-builder.slax** script, which is stored in the **/usr/libexec/scripts/op/** directory on the device. To use the **service-builder.slax** script, you must first copy it to the **op** scripts directory and enable it in the configuration. Only users in the Junos OS superuser login class can access and edit files in these directories.

1. Copy the **service-builder.slax** script to the **/var/db/scripts/op** directory on the hard disk or the **/config/scripts/op** directory on the flash drive.  

```
user@host> file copy /usr/libexec/scripts/op/service-builder.slax /var/db/scripts/op
```
2. Verify that the script is in the correct directory by using the **file list** operational mode command.  

```
user@host> file list /var/db/scripts/op
/var/db/scripts/op:
service-builder.slax*
```
3. Enable the script in the configuration.  

```
[edit]
user@host# set system scripts op file service-builder.slax
```
4. If you store scripts in and load them from flash, configure the **load-scripts-from-flash** statement, if it is not already configured.  

```
[edit]
user@host# set system scripts load-scripts-from-flash
```
5. Commit the configuration.  

```
[edit]
user@host# commit
```

## Configuring the Service Template Definition

### Step-by-Step Procedure

To create a new service template on a device running Junos OS:

1. Select a service name.

This example uses **vpn-service**.

2. In configuration mode, create a new group, which will contain the hierarchies for the actual service to be provisioned.

```
[edit]
user@host# set groups vpn-service-template-group
```

3. Configure the hierarchies for the service.

For this example, the pre-constructed configuration hierarchies are loaded into the candidate configuration using the **load merge terminal** command.

```
[edit]
user@host# load merge terminal
groups {
  vpn-service-template-group {
    interfaces {
      ge-2/2/6 {
        description "connected to customer3-site-1";
        unit 0 {
          family bridge {
            interface-mode access;
            vlan-id 300;
          }
        }
      }
    }
    protocols {
      rstp {
        interface ge-2/3/0;
      }
      mvrp {
        interface ge-2/3/0;
      }
    }
    bridge-domains {
      bd {
        vlan-id-list 100;
      }
    }
  }
}
[Ctrl+D]
```

4. Verify that the configuration syntax is correct.

```
[edit]
user@host# commit check
configuration check succeeds
```

5. Commit the configuration.

```
[edit]
user@host# commit
```

---

### Configuring and Generating the Service Interface

---

#### Step-by-Step Procedure

To configure and generate the service interface:

1. In operational mode, execute the **service-builder.slax** script, which starts an interactive Service Builder session.

```
user@host> op service-builder
Welcome to Service Builder Script: (v1.0)
-
Enter the service name :
```

2. Enter the service name that was defined in “[Configuring the Service Template Definition](#)” on page 519.

```
Enter the service name : vpn-service
```

3. Enter the group name under which the service hierarchies are configured.

This example uses the group name **vpn-service-template-group**. The script reads the configuration specified in the **vpn-service-template-group** hierarchy and then displays the main menu.

```
Enter the group name : vpn-service-template-group
.. reading [edit group vpn-service-template-group] ..
```

```
[Op Script Builder - Main Menu]
```

```
-----
1. Show Variables
2. Build Service API
3. Enable Service API
Q. Quit
-----
```

```
Enter Selection:>
```

4. (Optional) To review the service template variables that you can parameterize, select the **Show Variables** option.

The script translates the template definition in the candidate configuration into a general parameter list grouped by hierarchy level.

```
[Op Script Builder - Main Menu]
```

```
-----
1. Show Variables
2. Build Service API
3. Enable Service API
Q. Quit
-----
```

```
Enter Selection:> 1
```

```
List of variables under each hierarchy to parameterize:
-
```



```

1. [ edit groups vpn-service-template-group interfaces ]
   1.1. interface/name
   1.2. interface/description
   1.3. interface/unit/name
   1.4. interface/unit/family/bridge/interface-mode
   1.5. interface/unit/family/bridge/vlan-id
-
2. [ edit groups vpn-service-template-group protocols ]
   2.1. rstp/interface/name
   2.2. mvrp/interface/name
-
3. [ edit groups vpn-service-template-group bridge-domains ]
   3.1. domain/name
   3.2. domain/vlan-id-list
-

```

5. To build the Service API, select the **Build Service API** option.

```
[Op Script Builder - Main Menu]
```

```

-----
1. Show Variables
2. Build Service API
3. Enable Service API
Q. Quit
-----

```

```
Enter Selection:> 2
```

6. From the **Hierarchies Menu**, enter the menu selections for the hierarchies that have variables you want to parameterize, or press Enter to select all hierarchies.

```
[Op Script Builder - Hierarchies Menu]
```

```

-----
1. interfaces
2. protocols
3. bridge-domains
Q. Quit
-----

```

```
Please enter multiple selections separated by a comma (,) only.
```

```
Enter Selection:> [default:all] 1
```

7. From the variables list, enter the menu selections for the variables you want to parameterize for the service interface, or press Enter to parameterize all variables within that hierarchy.

```
List of variables to parameterize: ...
```

```
[ edit groups vpn-service-template-group interfaces ]
```

```

-----
1. interfaces/interface/name
2. interfaces/interface/description
3. interfaces/interface/unit/name
4. interfaces/interface/unit/family/bridge/interface-mode
5. interfaces/interface/unit/family/bridge/vlan-id
Q. Quit
-----

```

```
Please enter multiple selections separated by a comma (,) only.
```

```
Enter Selection:> [default:all] 1,2
```

8. Configure the selected parameters.

The system prompts for the required information. This example configures the interface name parameter as **ifname** and the interface description parameter as **ifdesc**.

```
Enter parameter name for: 1.interfaces/interface/name
*****
[ edit groups vpn-service-template-group interfaces ]
Name for this parameter? ifname
Do you want to revise 'ifname'? (yes/no)[no]: no
Enter parameter name for: 2.interfaces/interface/description
*****
[ edit groups vpn-service-template-group interfaces ]
Name for this parameter? ifdesc
Do you want to revise 'ifdesc'? (yes/no)[no]: no
```

9. Configure the selected parameters at each hierarchy level.

The script iterates over each selected hierarchy and the specified parameters. If you must exit the **service-builder.slax** script while building the service API, you must finish configuring all the parameters for the current hierarchy in order to save that hierarchy configuration when you exit using the **Quit** option.

10. Generate the service interface, which creates the service script.

Once all parameters are configured, the script automatically prompts you to generate the service interface. Press Enter or type yes to generate the service interface.

```
Do you want to commit the previously selected options to create vpn-service
script? (yes/no)[yes]: yes
Created service script: /var/db/scripts/op/vpn-service.slax
```

---

### Verifying the Service Interface

**Purpose** Verify the creation of the service script. If the **load-scripts-from-flash** statement is configured, the generated file is stored in flash memory. Otherwise, the generated file is stored on the hard disk.

**Action** Issue the **file list** operational mode command. For this example, the **vpn-service.slax** script should be present in the **/var/db/scripts/op** directory. The **service-builder.slax** script also generates the **utility.slax** script in the **/var/db/scripts/op** directory and the **vpn-service-builder-info.xml** file in the **/var/db/scripts/lib** directory. These files are used by the **service-builder.slax** script and should not be deleted.

```
user@host> file list /var/db/scripts/op
/var/db/scripts/op:
service-builder.slax
utility.slax
vpn-service.slax

user@host> file list /var/db/scripts/lib
/var/db/scripts/lib:
vpn-service-builder-info.xml
```

### Enabling the Service Interface

#### Step-by-Step Procedure

To enable the service interface on a remote device:

1. Copy the generated service script to the device where you are provisioning the new service.

If the **load-scripts-from-flash** statement is not configured, copy the service script to the **/var/db/scripts/op** directory on the second device. Otherwise, the script must be copied to the corresponding directory on the flash drive instead.

Copy the service script to the **/var/db/scripts/op** directory on the second device.

2. Enable the op script in the configuration.

```
[edit]
user@host2# set system scripts op file vpn-service.slax
```

3. Commit the configuration.

```
[edit]
user@host2# commit
commit complete
```

4. In operational mode, verify that the script is enabled and that the service parameters display as arguments for the script.

```
user@host2> op vpn-service ?
Possible completions:
<[Enter]>    Execute this command
<name>      Argument name
action      Please enter either create/delete/update
detail      Display detailed output
ifdesc      Text description of interface
ifname      Name of interface
service-id  Service Name
|           Pipe through a command
```

## Provisioning the Service Using NETCONF

### Step-by-Step Procedure

To provision the service:

1. If it is not already configured, configure NETCONF service over SSH on any devices where you are provisioning the new service.

```
[edit]
user@host2# set system services netconf ssh
user@host2# commit
```

2. From a configuration management server, establish a NETCONF session with the device where you are provisioning the service.

```
%ssh -p 830 -s user@host2 netconf
user@host2's password:
```

```
<!-- user user, class super-user -->
<hello>
  <capabilities>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0</capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:candidate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:validate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file
    </capability>
    <capability>http://xml.juniper.net/netconf/junos/1.0</capability>
    <capability>http://xml.juniper.net/dmi/system/1.0</capability>
  </capabilities>
  <session-id>28898</session-id>
</hello>
]]>]]>
```

3. If you are provisioning a new service on the device, enter a remote procedure call (RPC) that calls the service op script using the **create** action, and include values for all parameters that require configuring.

The value for the **service-id** parameter should be identical to the service name.

```
<rpc>
  <op-script>
    <script>vpn-service</script>
    <action>create</action>
    <service-id>vpn-service</service-id>
    <ifname>ge-2/0/5,ge-2/0/6</ifname>
    <ifdesc>connected to customer1-site-1,connected to
customer3-site-2</ifdesc>
  </op-script>
</rpc>
```

## Updating or Deleting Services Using NETCONF

### Step-by-Step Procedure

To update or delete an existing service:

1. If it is not already configured, configure NETCONF service over SSH on any devices where you are updating or deleting the service.

```
[edit]
user@host2# set system services netconf ssh
user@host2# commit
```

2. From a configuration management server, establish a NETCONF session with the device where you are provisioning the service.

```
%ssh -p 830 -s user@host2 netconf
user@host2's password:
```

```
<!-- user user, class super-user -->
<hello>
  <capabilities>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0</capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:candidate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:validate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file
    </capability>
    <capability>http://xml.juniper.net/netconf/junos/1.0</capability>
    <capability>http://xml.juniper.net/dmi/system/1.0</capability>
  </capabilities>
  <session-id>28898</session-id>
</hello>
]]>]]>
```

3. If the given service is already provisioned on the device and you are updating the service, enter an RPC that calls the service op script using the **update** action, and include values for all parameters that require updating.

```
<rpc>
  <op-script>
    <script>vpn-service</script>
    <action>update</action>
    <service-id>vpn-service</service-id>
    <ifname>ge-2/0/5</ifname>
    <ifdesc>connected to customer1-site-2</ifdesc>
  </op-script>
</rpc>
```

4. If the given service is already provisioned on the device and you are deleting some or all of the service parameters, enter an RPC that calls the service op script using the **delete** action, and include any parameters that need to be deleted.

```
<rpc>
  <op-script>
    <script>vpn-service</script>
    <action>update</action>
    <service-id>vpn-service</service-id>
    <ifname>ge-2/0/6</ifname>
  </op-script>
</rpc>
```

## Verification

Confirm that the configuration is updated.

- [Verifying the Service Configuration on page 526](#)
- [Verifying the Service Configuration on page 526](#)

---

### Verifying the Service Configuration

**Purpose** Verify that the commit is successful.

**Action** Issue the **show system commit** operational mode command to view the recent commits. The most recent commit entry shows that a commit was made through the NETCONF server by **user**.

```
user@host2> show system commit

0   2012-05-21 12:15:08 PDT by user via junoscript
1   2012-05-18 09:47:40 PDT by user via other
...
```

---

### Verifying the Service Configuration

**Purpose** Verify that the service configuration is present in the active configuration.

**Action** Issue the **show configuration | compare rollback *num*** operational mode command to view configuration changes.

```
user@host2> show configuration | compare rollback 1

[edit interfaces]
+   ge-2/0/5 {
+       description "connected to customer1-site-1";
+   }
+   ge-2/0/6 {
+       description "connected to customer3-site-2";
+   }
```

**Meaning** A comparison of the current configuration with the previous configuration shows that the interfaces and interface descriptions were added to the configuration.

## Troubleshooting

- [Troubleshooting a Failed Commit on page 527](#)
- [Troubleshooting a Failed Attempt to Delete Service Parameters on page 527](#)

---

### Troubleshooting a Failed Commit

**Problem** You see the following message when creating, updating, or deleting a service on a device through a NETCONF session:

```
<output>
configuration database modified
</output>
```

The configuration has previously uncommitted changes, and the service script cannot commit the service configuration changes.

**Solution** Commit the previous changes or roll back the configuration as appropriate, and then resubmit the service configuration changes.

---

### Troubleshooting a Failed Attempt to Delete Service Parameters

**Problem** You see the following message when deleting a service parameter on a device through NETCONF:

```
<xnm:error xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm">
  <source-daemon>
    op-script
  </source-daemon>
  <message>
    xsi:attribute: Cannot add attributes to an element if children have been
    already added to the element.
  </message>
</xnm:error>
```

**Solution** The RPC might include both the parameter and a child element. Remove the child element from the RPC.

**Related Documentation**

- [Service Template Automation Overview on page 515](#)





# Troubleshooting Op Scripts

- [Tracing Op Script Processing on page 529](#)

## Tracing Op Script Processing

---

Op script tracing operations track all op script operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

The default operation of op script tracing is to log important events in a file called **op-script.log** located in the **/var/log** directory. When the file **op-script.log** reaches 128 kilobytes (KB), it is renamed with a number 0 through 9 (in ascending order) appended to the end of the file and then compressed. The resulting files are **op-script.log.0.gz**, then **op-script.log.1.gz**, until there are 10 trace files. Then the oldest trace file (**op-script.log.9.gz**) is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)

This section discusses the following topics:

- [Minimum Configuration for Enabling Traceoptions for Op Scripts on page 529](#)
- [Configuring Tracing of Local Op Scripts on page 530](#)

## Minimum Configuration for Enabling Traceoptions for Op Scripts

If no op script trace options are configured, the simplest way to view the trace output of an op script is to configure the **output** trace flag and issue the **show log op-script.log | last** command. To do this, perform the following steps:

1. If you have not done so already, enable an op script by including the **file** statement at the **[edit system scripts op]** hierarchy level:

```
[edit system scripts op]
user@host# set file filename
```

2. Enable trace options by including the **traceoptions flag output** statement at the **[edit system scripts op]** hierarchy level:

```
[edit system scripts op]
user@host# set traceoptions flag output
```

- Issue the **commit** command:

```
[edit]
user@host# commit
```

- Display the resulting trace messages recorded in the file `/var/log/op-script.log` file. At the end of the log is the output generated by the op script you enabled in Step 1. To display the end of the log, issue the **show log op-script.log | last** operational mode command:

```
[edit]
user@host# run show log op-script.log | last
```

Table 41 on page 530 summarizes useful filtering commands that display selected portions of the **op-script.log** file.

**Table 41: Op Script Tracing Operational Mode Commands**

Task	Command
Display logging data associated with all op script processing.	<b>show log op-script.log</b>
Display processing for only the most recent operation.	<b>show log op-script.log   last</b>
Display processing for script errors.	<b>show log op-script.log   match error</b>
Display processing for a particular script.	<b>show log op-script.log   match <i>filename</i></b>

#### Example: Minimum Configuration for Enabling Traceoptions for Op Scripts

Display the trace output of the op script file **source-route.xml**:

```
[edit]
system {
  scripts {
    op {
      file source-route.xml;
      traceoptions flag output;
    }
  }
}

[edit]
user@host# commit
[edit]
user@host# run show log op-script.log | last
```

## Configuring Tracing of Local Op Scripts

You cannot change the directory (`/var/log`) to which trace files are written. However, you can customize other trace file settings for local op scripts by including the following statements at the **[edit system scripts op traceoptions]** hierarchy level:

```
[edit system scripts op traceoptions]
file <filename> <files number> <size size> <world-readable | no-world-readable>;
```

```
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
flag xslt;
no-remote-trace;
```



**NOTE:** You can only customize tracing operations for local op scripts. For remote op scripts (op scripts executed using the `op url` command), Junos OS always logs the default trace events in the default op script log file.

These statements are described in the following sections:

- [Configuring the Op Script Log Filename on page 531](#)
- [Configuring the Number and Size of Op Script Log Files on page 531](#)
- [Configuring Access to Op Script Log Files on page 532](#)
- [Configuring the Op Script Trace Operations on page 532](#)

### Configuring the Op Script Log Filename

By default, the name of the file that records trace output for all op scripts is **op-script.log**. You can record the trace output for local op scripts in a different file by including the `file` statement at the `[edit system scripts op traceoptions]` hierarchy level:

```
[edit system scripts op traceoptions]
file filename;
```



**NOTE:** The `file` statement does not affect the trace output location for remote op scripts, which always use the default log file **op-script.log**.

### Configuring the Number and Size of Op Script Log Files

By default, when the trace file reaches 128 KB in size, it is renamed and compressed to **filename.0.gz**, then **filename.1.gz**, and so on, until there are 10 trace files. Then the oldest trace file (**filename.9.gz**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the `[edit system scripts op traceoptions file <filename>]` hierarchy level:

```
[edit system scripts op traceoptions file <filename>]
files number size size;
```

For example, set the maximum file size to 640 KB and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 640 KB, it is renamed and compressed to **filename.0.gz**, and a new file called **filename** is created. When **filename** reaches 640 KB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is renamed and compressed to **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19.gz**) is overwritten.

The number of files can range from 2 through 1000 files. The file size can range from 10 KB through 1 gigabyte (GB).



**NOTE:**

If you set either a maximum file size or a maximum number of trace files, you also must specify the other parameter and a filename.

### Configuring Access to Op Script Log Files

By default, access to the op script log file is restricted to the owner. You can manually configure access by including the **world-readable** or **no-world-readable** statement at the **[edit system scripts op traceoptions file <filename>]** hierarchy level.

```
[edit system scripts op traceoptions file <filename>]
(world-readable | no-world-readable);
```

The **no-world-readable** statement restricts op script log access to the owner. The **world-readable** statement enables unrestricted access to the op script log file.

### Configuring the Op Script Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged for local op scripts by including the following statements at the **[edit system scripts op traceoptions]** hierarchy level:

```
[edit system scripts op traceoptions]
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
flag xslt;
```



**NOTE:** The flag statement does not affect the trace output for remote op scripts, which only log default trace events.

Table 42 on page 532 describes the meaning of the op script tracing flags.

**Table 42: Op Script Tracing Flags**

Flag	Description	Default Setting
all	Trace all operations.	Off
events	Trace important events.	On
input	Trace op script input data.	Off
offline	Generate data for offline development.	Off

Table 42: Op Script Tracing Flags (*continued*)

Flag	Description	Default Setting
<b>output</b>	Trace op script output data.	Off
<b>rpc</b>	Trace op script RPCs.	Off
<b>xslt</b>	Trace the Extensible Stylesheet Language Transformations (XSLT) library.	Off



## PART 9

# Event Policies and Event Scripts

- [Event Policy Overview on page 537](#)
- [Triggering an Event Policy on page 541](#)
- [Configuring Event Policy Actions on page 555](#)
- [Configuring Event Policy File Archiving on page 591](#)
- [Configuring Event Policy Privileges on page 609](#)
- [Event Scripts Overview on page 613](#)
- [Creating and Executing Event Scripts on page 615](#)
- [Troubleshooting Event Policies and Event Scripts on page 631](#)





# Event Policy Overview

- [Event Notifications and Policies Overview on page 537](#)
- [How Event Policies Work on page 538](#)

## Event Notifications and Policies Overview


To diagnose a fault or error condition on a device, you need relevant information about the state of the platform. You can derive state information from *event notifications*. Event notifications are system log messages and SNMP traps. A Junos OS process called the *event process* (eventd) receives event notifications—henceforth simply called *events*—from other Junos OS processes.

Timely diagnosis and intervention can correct error conditions and keep the device in operation. After the eventd process receives events, *event policies* instruct the eventd process to select specific events, correlate the events, and perform a set of actions. These actions can either help you diagnose a fault or take corrective action. For example, the eventd process can upload device files to a given destination and issue operational mode commands.

Events can originate as SNMP traps or system log messages. The event process receives event messages from other Junos OS processes, such as the routing protocol process (rpd) and the management process (mgd). Depending on the custom event policy you configure, eventd listens for specific events and in response to these events might create a log file, invoke a Junos OS command, or invoke an event script. When an event script is invoked, event details are passed to the event script in the form of XML inputs.

[Figure 8 on page 537](#) shows how the event process (eventd) interacts with other Junos OS processes.

**Figure 8: Interaction of eventd Process with Other Junos OS Processes**



## How Event Policies Work

---

An event policy is an if-then-else construct. It defines actions to be executed by the eventd process on receipt of one or more events. You can configure multiple policies to be processed for an event. The policies are executed in the order in which they appear in the configuration. For each policy, you can configure multiple actions. The actions are also executed in the order in which they appear in the configuration.

To view a list of the events that can be referenced in an event policy, issue the **help syslog ?** command.

```
user@host> help syslog ?
Possible completions:
<syslog-tag>      System log tag
ACCT_ACCOUNTING_FERROR  Error occurred during file processing
ACCT_ACCOUNTING_FOPEN_ERROR  Open operation failed on file
...
```

You can filter the output of a search by using the pipe (|) symbol. The following example lists the filters that can be used with the pipe symbol:

```
user@host> help syslog | ?
Possible completions:
count             Count occurrences
display           Show additional kinds of information
except            Show only text that does not match a pattern
find              Search for first occurrence of pattern
hold              Hold text without exiting the --More-- prompt
last              Display end of output only
match             Show only text that matches a pattern
no-more           Don't paginate output
request           Make system-level requests
resolve           Resolve IP addresses
save              Save output text to file
trim              Trim specified number of columns from start of line
```

For more information about using the pipe symbol, see the *CLI User Guide*.

You can also list events as you configure the event policy. To view a partial list of the events that can be referenced in an event policy, issue the **set event-options policy *policy-name* events ?** configuration mode command.

```
[edit]
user@host# set event-options policy policy-name events ?
Possible completions:
<event>
[          Open a set of values
acct_accounting_ferror
acct_accounting_fopen_error
...
```

Some of the system log messages that you can reference in an event policy are not listed in the output of the **set event-options policy *policy-name* events ?** command. For information about referencing these system log messages in your event policies, see [“Using Nonstandard System Log Messages to Trigger Event Policies” on page 552](#).

You can configure an event policy to trigger for a single event or for two or more correlated events. For more information about correlating events, see [“Using Correlated Events to Trigger an Event Policy” on page 541](#)

In addition, you can reference internally generated events, which are discussed in [“Generating Internal Events to Trigger Event Policies” on page 551](#).

In response to events, the eventd process can execute the following actions:

- Ignore the event—Do not generate a system log message for this event and do not process any further policy instructions for this event.
- Upload a file—Upload a file to a specified destination. You can specify a transfer delay, so that, on receipt of an event, the upload of the file begins after the configured transfer delay. For example, to upload a core file, a transfer delay can ensure that the core file has been completely generated before the upload begins.
- Execute Junos OS operational mode commands—Execute commands on receipt of an event. The XML or text output of these commands is stored in a file, which is then uploaded to a specified URL. You can include variables in the command that allow data from the triggering event to be automatically included in the command syntax.
- Execute Junos OS configuration mode commands—Execute commands to modify the configuration on receipt of an event. Starting in Junos OS Release 12.1, you can configure an event policy to modify the configuration using Junos OS configuration mode commands and then commit the updated configuration.
- Execute Junos OS event scripts—Execute event scripts on receipt of an event. Event scripts are Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python scripts that you write to perform any function available through Junos XML or Junos XML protocol remote procedure calls (RPCs). Additionally, you can pass to an event script a set of arguments that you define. A script can build and run an operational mode command, receive the command output, inspect the output, and determine the next appropriate action. This process can be repeated until the source of the problem is determined. The output of the scripts is stored in a file, which is then uploaded to a specified URL. You can include variables in the arguments to the scripts that allow data from the triggering event to be incorporated into the script.
- Raise an SNMP trap.



## CHAPTER 29

# Triggering an Event Policy

- [Using Correlated Events to Trigger an Event Policy on page 541](#)
- [Representing the Correlating Event in an Event Policy on page 544](#)
- [Example: Representing the Correlating Event in an Event Policy on page 545](#)
- [Example: Correlating Events Based on Receipt of Other Events Within a Specified Time Interval on page 546](#)
- [Example: Correlating Events Based on Event Attributes on page 546](#)
- [Triggering an Event Policy Based on Event Count on page 547](#)
- [Example: Triggering a Policy Based on Event Count on page 547](#)
- [Using Regular Expressions to Refine the Set of Events That Trigger a Policy on page 549](#)
- [Example: Controlling Event Policy Using a Regular Expression on page 550](#)
- [Generating Internal Events to Trigger Event Policies on page 551](#)
- [Example: Generating Internal Events to Trigger Event Policies on page 552](#)
- [Using Nonstandard System Log Messages to Trigger Event Policies on page 552](#)
- [Example: Using Nonstandard System Log Messages to Trigger an Event Policy on page 553](#)

## Using Correlated Events to Trigger an Event Policy

---

You can configure a policy that correlates two or more events. If the correlated events occur as specified, they cause particular actions to be taken. For example, you might want to issue certain operational mode commands when a `UI_CONFIGURATION_ERROR` event is generated within five minutes (300 seconds) after a `UI_COMMIT_PROGRESS` event. As another example, you might want to upload a particular file if a `DCD_INTERFACE_DOWN` event is generated two times within a 60-second interval.

To configure a policy that correlates events, include the following statements at the **[edit event-options]** hierarchy level:

```
[edit event-options]
policy policy-name {
  events [ events ];
  within seconds {
    events [ events ];
    not events [ events ];
    trigger (on | after | until) event-count;
```

```

}
attributes-match {
    event1.attribute-name equals event2.attribute-name;
    event.attribute-name matches regular-expression;
    event1.attribute-name starts-with event2.attribute-name;
}
then {
    ...
}
}

```

In the **events** statement, you can list multiple events. To view a list of the events that can be referenced in an event policy, issue the **set event-options policy *policy-name* events ?** configuration mode command:

```

user@host# set event-options policy policy-name events ?
Possible completions:
<event>
[                Open a set of values
acct_accounting_ferror
acct_accounting_fopen_error
...

```

Some of the system log messages that you can reference in an event policy are not listed in the output of the **set event-options policy *policy-name* events ?** command. For information about referencing these system log messages in your event policies, see [“Using Nonstandard System Log Messages to Trigger Event Policies” on page 552](#).

In addition, you can reference internally generated events, which are discussed in [“Generating Internal Events to Trigger Event Policies” on page 551](#).

The actions configured in the **then** statement are executed only if certain conditions are met, which you specify in the **within** and **attributes-match** statements.

You can configure a policy that is executed only if a specified event occurs within a specified time interval after another event. You do this by including the **within *seconds* events** statement. The policy is executed only if one or more of the events in the first **events** statement occur within a configured number of seconds after one or more of the events in the **within *seconds* events** statement. The number of seconds can be from 60 through 604,800. The **not** statement causes the policy to be executed only if the events do not occur within the configured time interval.

For example, the following policy is executed if **event3**, **event4**, or **event5** occurs within 60 seconds after **event1** or **event2** occurs:

```

[edit event-options]
policy 1 {
    events [ event3 event4 event5 ];
    within 60 events [ event1 event2 ];
    then {
        ...
    }
}

```

The **attributes-match** statement correlates two events as follows:

- ***event1.attribute-name equals event2.attribute-name***—Execute the policy only if the specified attribute of ***event1*** equals the specified attribute of ***event2***.
- ***event.attribute-name matches regular-expression***—Execute the policy only if the specified attribute of ***event*** matches a regular expression. For more information, see [“Using Regular Expressions to Refine the Set of Events That Trigger a Policy”](#) on page 549.
- ***event1.attribute-name starts-with event2.attribute-name***—Execute the policy only if the specified attribute of ***event1*** starts with the specified attribute of ***event2***.

If the **attributes-match** statement includes the **equals** or **starts-with** options, or if it includes a **matches** option that includes a clause for an event that is not specified at the **[edit event-options policy policy-name events]** hierarchy level, you must include one or more **within** statements in the same policy configuration.

Starting in Junos OS Release 11.1, you can use event policy variables within the **attributes-match** statement to differentiate between a trigger event attribute and a correlated event attribute. The double dollar sign (**\$\$**) notation represents the event that is triggering a policy, and **{\$\$.attribute-name}** resolves to the value of the attribute of the triggering event. Triggering events are those that you configure at the **[edit event-options policy policy-name events]** hierarchy level. For correlating events, the single dollar sign with the event name (**\$event**) notation represents the most recent event that matches the event name, and **{\$event.attribute-name}** resolves to the value of the attribute associated with that event.

In the following example, the policy will execute the actions under the **then** statement if four or more commits are performed within a 5-minute period, and the username of one or more of the correlated events is the same as the username of the trigger event.

```
policy multiple-commits {
  events ui_commit;
  attributes-match {
    {$$.user-name} equals {$ui_commit.user-name};
  }
  within 300 {
    trigger after 3;
    events ui_commit;
  }
  then ...
}
```

To view a list of all event attributes that you can reference, issue the **help syslog event** operational mode command. The output of this command shows the event attributes in angle brackets (<>). The following output shows that three attributes can be referenced for the **ACCT\_ACCOUNTING\_SMALL\_FILE\_SIZE** event: **filename**, **file-size**, and **record-size**.

```
user@host> help syslog ACCT_ACCOUNTING_SMALL_FILE_SIZE
Name:          ACCT_ACCOUNTING_SMALL_FILE_SIZE
Message:       File <filename> size (<file-size>) is smaller than record size
               (<record-size>)
```

You can filter the output of a search by using the pipe (|) symbol. The following example lists the filters that can be used with the pipe symbol:

```
user@host> help syslog | ?
```

## Possible completions:

count	Count occurrences
display	Show additional kinds of information
except	Show only text that does not match a pattern
find	Search for first occurrence of pattern
hold	Hold text without exiting the --More-- prompt
last	Display end of output only
match	Show only text that matches a pattern
no-more	Don't paginate output
request	Make system-level requests
resolve	Resolve IP addresses
save	Save output text to file
trim	Trim specified number of columns from start of line

For more information about using the pipe symbol, see the *CLI User Guide*.

Another way to view the attributes you can reference is by issuing the **set attributes-match event?** command at the **[edit event-options policy policy-name]** hierarchy level, as shown in the following example:

```
[edit event-options policy p1]
user@host# set attributes-match acct_accounting_small_file_size?
Possible completions:
<from-event-attribute> First attribute to compare
acct_accounting_small_file_size.filename
acct_accounting_small_file_size.filesize
acct_accounting_small_file_size.record-size
```



**NOTE:** In this **set** command, there is no space between the event name and the question mark (?).

**Related Documentation**

- [Representing the Correlating Event in an Event Policy on page 544](#)
- [Triggering an Event Policy Based on Event Count on page 547](#)
- [Using Regular Expressions to Refine the Set of Events That Trigger a Policy on page 549](#)
- [attributes-match on page 771](#)
- [policy \(Event Policy\) on page 790](#)
- [not \(Event Policy\) on page 787](#)
- [then \(Event Policy\) on page 797](#)
- [within on page 806](#)

---

## Representing the Correlating Event in an Event Policy

As described in “[Configuring an Event Policy to Execute Operational Mode Commands](#)” on [page 555](#), the double dollar sign (\$\$) notation represents the event that is triggering a policy. Triggering events are those that you configure at the **[edit event-options policy policy-name events]** hierarchy level.



As described in [“Using Correlated Events to Trigger an Event Policy” on page 541](#), you can configure a policy that is executed only if a specified event occurs within a specified time interval after another event. You do this by including the **within seconds events** statement at the **[edit event-options policy policy-name]** hierarchy level:

```
[edit event-options policy policy-name ]
events [ events ];
within seconds events [ events ];
```

The policy is executed only if one or more of the events at the **[edit event-options policy policy-name events]** hierarchy level occur within a configured number of seconds after one or more of the events in the **within seconds events** statement.

For correlating events, the single dollar sign with the event name (**\$event**) notation represents the most recent event that matches the event name. The dollar sign with the asterisk (**\$\***) notation represents the most recent event that matches any of the correlating events.

For a configuration example, see [“Example: Representing the Correlating Event in an Event Policy” on page 545](#).

### Example: Representing the Correlating Event in an Event Policy

```
[edit event-options]
policy p1 {
  events [ e1 e2 e3 ];
  within 60 events [ e4 e5 e6 ];
  then {
    execute-commands {
      commands {
        "show interfaces {${$.interface-name}";
        "show interfaces {$e4.interface-name}";
        "show interfaces {$*.interface-name}";
      }
      output-filename command-output.txt;
      destination some-dest;
    }
  }
}
```

In the **show interfaces {\${\$.interface-name}** command, the value of the **interface-name** attribute of event **e1**, **e2**, or **e3** is substituted for the **{\${\$.interface-name}** variable.

In the **show interfaces {\$e4.interface-name}** command, the value of the **interface-name** attribute of the most recent **e4** event is substituted for the **{\$e4.interface-name}** variable.

In the **show interfaces {\$\*.interface-name}** command, the value of the **interface-name** attribute of the most recent **e4**, **e5**, or **e6** event is substituted for the **{\$\*.interface-name}** variable. If one of **e4**, **e5**, or **e6** occurs within 60 seconds of **e1**, **e2**, or **e3**, the value of the **interface-name** attribute for that correlating event (**e4**, **e5**, or **e6**) is substituted for the **{\$\*.interface-name}** variable. If the correlating event does not have an **interface-name** attribute, the software does not execute the **show interfaces {\$\*.interface-name}** command.

If both **e4** and **e5** occur within 60 seconds of **e1**, then the value of the **interface-name** attribute for **e4** is substituted for the **{\${interface-name}}** variable. This is because the event process (eventd) searches for correlating events in sequential order as configured in the **within** statement. In this case, the order is **e4 > e5 > e6**.

## Example: Correlating Events Based on Receipt of Other Events Within a Specified Time Interval

---

In the following policy, a set of commands is issued and the output is logged and saved to a given location. The policy is executed if **event3**, **event4**, or **event5** occurs within 60 seconds after **event1** or **event2** occurs. The pseudocode for the policy is as follows:

```
if this event is (event3 or event4 or event5)
and
(event1 or event2 has been received within the last 60 seconds)
then {
  run a set of commands;
  log the output of these commands to a location;
}
```

Specify two archive sites in the configuration. The device attempts to transfer to the first archive site in the list, moving to the next site only if the transfer fails.

```
[edit event-options]
policy 1 {
  events [ event3 event4 event5 ];
  within 60 events [ event1 event2 ];
  then {
    execute-commands {
      commands {
        "command";
      }
      output-filename my_cmd_out;
      destination policy-1-command-dest;
    }
  }
}
destinations {
  policy-1-command-dest {
    archive-sites {
      scp://robot@my.big.com/a/b;
      scp://robot@my.little.com/a/b;
    }
  }
}
```

## Example: Correlating Events Based on Event Attributes

---

In the following policy, the two events are correlated only if two of their parameter values match. Matching on attributes of both events ensures that the two events are related. In this case, the interface addresses must match and the physical interface (ifd) names must match.

The **RPD\_KRT\_IFDCHANGE** error occurs when the routing protocol process (rpd) sends a request to the kernel to change the state of an interface and the request fails. The

**RPD\_RDISC\_NOMULTI** error occurs when an interface is configured for router discovery but the interface does not support IP multicast operations as required.

In this example, **RPD\_RDISC\_NOMULTI.interface-name** might be **so-0/0/0.0**, and **RPD\_KRT\_IFDCHANGE.ifd-index** might be **so-0/0/0**.

```
[edit event-options]
policy 1 {
  events rpd_rdisc_nomulti;
  within 500 events rpd_krt_ifdchange;
  attributes-match {
    rpd_rdisc_nomulti.interface-address equals rpd_krt_ifdchange.address;
    rpd_rdisc_nomulti.interface-name starts-with rpd_krt_ifdchange.ifd-index;
  }
  then {
    ... actions ...
  }
}
```

## Triggering an Event Policy Based on Event Count

You can configure an event policy to be triggered if an event or set of events occurs a specified number of times within a specified time period.

To do this, include the optional **trigger** statement at the **[edit event-options policy *policy-name* within *seconds*]** hierarchy level:

```
[edit event-options policy policy-name within seconds]
  trigger (after | on | until) event-count;
```

The software counts the number of times the triggering event occurs. A triggering event can be any event configured at the **[edit event-options policy *policy-name* events]** hierarchy level. You can configure the following options:

- **after *event-count***—The policy is executed when the number of matching events received equals *event-count* plus one.
- **on *event-count***—The policy is executed when the number of matching events received equals *event-count*.
- **until *event-count***—The policy is executed each time a matching event is received and stops being executed when the number of matching events received equals *event-count*.

For a configuration example, see “[Example: Triggering a Policy Based on Event Count](#)” on page 547.

## Example: Triggering a Policy Based on Event Count

This section discusses two examples.



**NOTE:** The **RADIUS\_LOGIN\_FAIL**, **TELNET\_LOGIN\_FAIL**, and **SSH\_LOGIN\_FAIL** events are not actual Junos OS events. They are illustrative for these examples.

**Example 1** Configure an event policy called **login**. The **login** policy is executed if five login failure events (**RADIUS\_LOGIN\_FAIL**, **TELNET\_LOGIN\_FAIL**, or **SSH\_LOGIN\_FAIL**) are generated within 120 seconds. Take action by executing the **login-fail.xsl** event script, which disables the user account.

```
[edit event-options]
policy login {
  events [ RADIUS_LOGIN_FAIL TELNET_LOGIN_FAIL SSH_LOGIN_FAIL ];
  within 120 {
    trigger after 4;
  }
  then {
    event-script login-fail.xsl {
      destination some-dest;
    }
  }
}
```

Table 43 on page 548 shows how events add to the count.

**Table 43: Event Count Triggers Policy**

Event Number	Event	Time	Count	Order
1	RADIUS_LOGIN_FAIL	00:00:00	1	[1]
2	TELNET_LOGIN_FAIL	00:00:20	2	[1 2]
3	RADIUS_LOGIN_FAIL	00:02:05	2	[2 3]
4	SSH_LOGIN_FAIL	00:02:40	2	[3 4]
5	TELNET_LOGIN_FAIL	00:02:55	3	[3 4 5]
6	TELNET_LOGIN_FAIL	00:03:01	4	[3 4 5 6]
7	RADIUS_LOGIN_FAIL	00:03:55	5	[3 4 5 6 7]

The columns in Table 43 on page 548 mean the following:

- Event number—Event sequence number.
- Event—Policy login events received by the event process (eventd).
- Time—Time (in **hh:mm:ss** format) when eventd receives the event.
- Count—The number of events received by eventd within the last 120 seconds.
- Order—Order of events as received by eventd within the last 120 seconds.

At time 00:03:55, the value of count is more than 4; therefore, the **login** policy executes the **login-fail.xsl** script.

**Example 2** Configure an event policy called **login**. The **login** policy is executed if five login failure events (**RADIUS\_LOGIN\_FAIL**, **TELNET\_LOGIN\_FAIL**, or **SSH\_LOGIN\_FAIL**) are generated within 120 seconds from username **roger**. Take action by executing the **login-fail.xml** event script, which disables the **roger** user account.

```
[edit event-options]
policy p2 {
  events [ RADIUS_LOGIN_FAIL TELNET_LOGIN_FAIL SSH_LOGIN_FAIL ];
  within 120 {
    trigger after 4;
  }
  attributes-match {
    RADIUS_LOGIN_FAIL.username matches roger;
    TELNET_LOGIN_FAIL.username matches roger;
  }
  then {
    event-script login-fail.xml {
      destination some-dest;
    }
  }
}
```

## Using Regular Expressions to Refine the Set of Events That Trigger a Policy

You can use regular expression matching to specify more exactly which events cause a policy to be executed.

To specify the text string that must appear in an event attribute for the policy to be executed, include the **matches** statement at the **[edit event-options policy *policy-name* attributes-match]** hierarchy level, and specify the regular expression that the event attribute must match:

```
[edit event-options policy policy-name attributes-match]
event.attribute-name matches regular-expression;
```

When you specify the regular expression, use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions. Explaining regular expression syntax is beyond the scope of this document. [Table 44 on page 549](#) specifies which character or characters are matched by some of the regular expression operators that you can use in the **matches** statement. In the descriptions, the term *term* refers to either a single alphanumeric character or a set of characters enclosed in square brackets, parentheses, or braces.



**NOTE:** The **matches** statement is not case-sensitive.

**Table 44: Regular Expression Operators for the matches Statement**

Operator	Matches
. (period)	One instance of any character except the space.
* (asterisk)	Zero or more instances of the immediately preceding term.

**Table 44: Regular Expression Operators for the matches Statement (*continued*)**

Operator	Matches
+ (plus sign)	One or more instances of the immediately preceding term.
? (question mark)	Zero or one instance of the immediately preceding term.
(pipe)	One of the terms that appear on either side of the pipe operator.
! (exclamation point)	Any string except the one specified by the expression, when the exclamation point appears at the start of the expression. Use of the exclamation point is specific to Junos OS.
^ (caret)	The start of a line, when the caret appears outside square brackets.  One instance of any character that does not follow it within square brackets, when the caret is the first character inside square brackets.
\$ (dollar sign)	The end of a line.
[ ] (paired square brackets)	One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen ( - ) to separate the beginning and ending characters of the range. For example, <b>[a-z0-9]</b> matches any letter or number.
( ) (paired parentheses)	One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.

**Related Documentation**

- [Example: Controlling Event Policy Using a Regular Expression on page 550](#)

## Example: Controlling Event Policy Using a Regular Expression

The following policy is executed only if the **interface-name** attribute in both traps (**SNMP\_TRAP\_LINK\_DOWN** and **SNMP\_TRAP\_LINK\_UP**) match each other and the **interface-name** attribute in the **SNMP\_TRAP\_LINK\_DOWN** trap starts with letter *t*. This means the policy is executed only for T1 (t1-) and T3 (t3-) interfaces. The policy is not executed when the **eventd** process receives traps from other interfaces.



**NOTE:** In system log files, the message tags appear in all uppercase letters. In the command-line interface (CLI), the message tags appear in all lowercase letters.

```
[edit event-options]
policy pol6 {
  events snmp_trap_link_down;
  within 120 events snmp_trap_link_up;
  attributes-match {
    snmp_trap_link_up.interface-name equals snmp_trap_link_down.interface-name;
```

```

    snmp_trap_link_down.interface-name matches "^t";
  }
  then {
    execute-commands {
      commands {
        "show interfaces {${$.interface-name}";
        "show configuration interfaces {${$.interface-name}";
      }
      output-filename config.txt;
      destination bsd2;
      output-format text;
    }
  }
}

```

## Generating Internal Events to Trigger Event Policies

*Internal events* are events that you create to trigger a policy to be executed. They are not generated by Junos OS processes, and they do not have any associated system log messages. You can generate an internal event based on a time interval or the time of day.

To generate an event, include the following statements at the **[edit event-options]** hierarchy level:

```

[edit event-options]
generate-event {
  event-name (time-interval seconds | time-of-day hh:mm:ss) <no-drift>;
}

```

In the **time-interval** statement, configure a frequency, in seconds, with which to repeatedly generate an event. The time interval can range from 60 through 2,592,000 seconds.

In the **time-of-day** statement, configure a time of day for the event to occur. Use the format **hh:mm:ss**.



**NOTE:** If you modify the system time by issuing the **set date** operational mode command, we recommend that you also issue the **commit full** or the **restart event-process** command. Otherwise, an internal event based on the time of day might not be generated at the configured time.

For example, if you configure an internal event to be generated at 15:55:00, and then you modify the system time from 15:47:17 to 15:53:00, the event is generated when the system time is approximately 16:00 instead of at the configured time, 15:55:00. You can correct this problem by issuing the **commit full** or the **restart event-process** command.

Starting in Junos OS Release 14.1, you can include the **no-drift** option when generating internal events. When you configure **no-drift**, Junos OS does not propagate the delay caused in triggering an event to the triggering of the next event.

You can configure up to 10 internal events. If you attempt to commit a configuration with more than 10 internal events, Junos OS generates an error, and the commit fails.

For configuration examples, see [“Example: Generating Internal Events to Trigger Event Policies” on page 552](#).

**Related  
Documentation**

- [Example: Generating Internal Events to Trigger Event Policies on page 552](#)
- [no-drift \(Generate Events\) on page 786](#)

---

## Example: Generating Internal Events to Trigger Event Policies

The following two examples generate an internal event. In the first example, the configuration generates an internal event every hour. In the second example, the configuration generates an event every night at midnight.

In the following example, the internal event called **EVERY-ONE-HOUR** is generated every hour (3600 seconds). If 3601 seconds pass and the event has not been generated, certain actions are taken.

```
[edit event-options]
generate-event {
  every-one-hour-time-interval 3600;
}
policy check-heartbeat {
  events every-one-hour;
  within 3601 not events every-one-hour;
  then {
    ... actions ...
  }
}
```

In the following example, the internal event called **IT-IS-MIDNIGHT** is generated at 12:00 AM every night (00:00:00). When the eventd process receives the **IT-IS-MIDNIGHT** event, certain actions are taken.

```
[edit event-options]
generate-event {
  it-is-midnight time-of-day 00:00:00;
}
policy midnight-chores {
  events it-is-midnight;
  then {
    ... actions ...
  }
}
```

---

## Using Nonstandard System Log Messages to Trigger Event Policies

Some of the system log messages that you can reference in an event policy are not listed in the output of the **set event-options policy *policy-name* events ?** command. These system log messages have an event ID and a **message** attribute. Event IDs are based on the origin of the message, as shown in [Table 45 on page 553](#).



Table 45: Event ID by System Log Message Origin

Event IDs	Origin
SYSTEM	Messages from Junos OS daemons and utilities
KERNEL	Messages from the Junos OS kernel
PIC	Messages from physical interface cards (PICs)
PFE	Messages from the Packet Forwarding Engine
LCC	On a TX Matrix router, messages from a line-card chassis (LCC)
SCC	On a TX Matrix router, messages from a switch-card chassis (SCC)

To base your event policy on the event types shown in [Table 45 on page 553](#), include the **events** *event-id* statement and the **attributes-match** statement with the *event-id.message* matches "message" attribute at the [edit event-options policy *policy-name*] hierarchy level:

```
[edit event-options policy policy-name]
events event-id;
attributes-match {
  event-id.message matches "message";
}
```

For a configuration example, see “[Example: Using Nonstandard System Log Messages to Trigger an Event Policy](#)” on page 553.

### Example: Using Nonstandard System Log Messages to Trigger an Event Policy

Reference a **KERNEL** system log message in an event policy. The **raise-trap** action in the **then** statement is executed only if a **KERNEL** event containing a message that matches "exited on signal 11" occurs.

```
[edit event-options]
policy kernel-policy {
  events KERNEL;
  attributes-match {
    KERNEL.message matches "exited on signal 11";
  }
  then {
    raise-trap;
  }
}
```



# Configuring Event Policy Actions

- [Configuring an Event Policy to Execute Operational Mode Commands on page 555](#)
- [Configuring an Event Policy to Change the Configuration on page 558](#)
- [Example: Changing the Configuration Using an Event Policy on page 559](#)
- [Example: Changing the Interface Configuration in Response to an Event on page 566](#)
- [Executing Event Scripts in an Event Policy on page 573](#)
- [Configuring an Event Policy to Pass Arguments to an Event Script on page 576](#)
- [Configuring Event Policies to Ignore an Event on page 578](#)
- [Example: Ignoring Events Based on Receipt of Other Events on page 579](#)
- [Overview of Using Event Policies to Raise SNMP Traps on page 580](#)
- [Example: Raising an SNMP Trap in Response to an Event on page 581](#)
- [Understanding the Event System Log Priority in an Event Policy on page 583](#)
- [Example: Configuring the Event System Log Priority in an Event Policy on page 584](#)
- [Example: Limiting Event Script Output Based on a Specific Event Type on page 588](#)

## Configuring an Event Policy to Execute Operational Mode Commands

---

*Operational mode commands* request that the device running Junos OS perform an operation or provide diagnostic output. They allow you to view statistics and information about a device's current operating status. They also allow you to take corrective actions, such as restarting software processes, taking a PIC offline and back online, switching to redundant interfaces, and adjusting Label Switching Protocol (LSP) bandwidth. For more information about operational mode commands, see the [CLI Explorer](#).

You can configure an event policy that executes operational mode commands and uploads the output of those commands to a specified location for analysis.

To configure such a policy, include the following statements at the **[edit event-options]** hierarchy level:

```
[edit event-options]
policy policy-name {
  events [ events ];
  then {
    execute-commands {
      commands {
```

```

        "command";
    }
    output-filename filename;
    output-format (text | xml);
    destination destination-name;
}
}
}

```

In the **events** statement, you can list multiple events. If one or more of the listed events occurs, the eventd process executes the operational mode commands configured for the **commands** statement. Enclose each command in quotation marks (" "). The eventd process issues the commands in the order in which they appear in the configuration. For example, in the following configuration, the execution of **policy1** causes the **show interfaces** command to be issued first, followed by the **show chassis alarms** command:

```

[edit event-options policy policy1 then execute-commands]
user@host# show
commands {
    "show interfaces";
    "show chassis alarms";
}

```

You can include variables in the command to allow data from the triggering event to be automatically included in the command syntax. The eventd process replaces each variable with values contained in the event that triggers the policy. You can use command variables of the following forms:

- **{{\$.attribute-name}}**—The double dollar sign (\$\$) notation represents the event that is triggering a policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event. For example, **{{\$.interface-name}}** resolves to the interface name associated with the triggering event.
- **{\$event.attribute-name}**—The single dollar sign with the event name (**\$event**) notation represents the most recent event that matches **event**. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event. For example, when a policy issues the **show interfaces {\$COSD\_CHAS\_SCHED\_MAP\_INVALID.interface-name}** command, the **{\$COSD\_CHAS\_SCHED\_MAP\_INVALID.interface-name}** variable resolves to the interface name associated with the most recent **COSD\_CHAS\_SCHED\_MAP\_INVALID** event cached by the event process.

For a given event, you can view a list of event attributes that you can reference in an operational mode command by issuing the **help syslog event** command:

```
user@host> help syslog event
```

For example, in the following command output, text in angle brackets (< >) shows that **classifier-type** is an attribute of the **cosd\_unknown\_classifier** event:

```

user@host> help syslog cosd_unknown_classifier
Name:          COSD_UNKNOWN_CLASSIFIER
Message:       rtsock classifier type <classifier-type> is invalid
...

```

Another way to view a list of event attributes is to issue the **set attributes-match event?** configuration mode command at the **[edit event-options policy *policy-name*]** hierarchy level:

```
[edit event-options policy policy-name]
user@host# set attributes-match event ?
```

For example, in the following command output, the **event.attribute** list shows that **classifier-type** is an attribute of the **cosd\_unknown\_classifier** event:

```
[edit event-options policy policy-name]
user@host# set attributes-match cosd_unknown_classifier?
Possible completions:
<from-event-attribute> First attribute to compare
cosd_unknown_classifier.classifier-type
```



**NOTE:** In this **set** command, there is no space between the event name and the question mark (?).

To view a list of all event attributes that you can reference, issue the **set attributes-match ?** configuration mode command at the **[edit event-options policy *policy-name*]** hierarchy level:

```
[edit event-options policy policy-name]
user@host# set attributes-match ?
Possible completions:
<from-event-attribute> First attribute to compare
acct_accounting_ferror
acct_accounting_fopen_error
...
```

When the eventd process executes the commands, it uploads the file with the command output to the location specified in the **destination** statement. In the **destination** statement, include a destination name that is configured at the **[edit event-options destinations]** hierarchy level. For more information, see [“Example: Defining Destinations for File Archiving by Event Policies” on page 592](#).

In the **output-filename** statement, define a descriptive string that will be included in the filename. Each uploaded file also includes the hostname and timestamp in the filename to ensure that the each filename is unique. If a policy is triggered multiple times in a 1-second period, an index number is appended to the filename to ensure that the filenames are still unique. The index number range is 001 through 999.

When eventd executes the commands defined at the **[edit event-options policy *policy-name* then execute-commands *commands*]** hierarchy level and generates a file containing the command output, the name of the file depends on the version of Junos OS running on the device. Prior to Junos OS Release 14.1R3, the filename has the following naming convention:

***hostname\_output-filename\_YYYYMMDD\_HHMMSS\_index-number***

Starting in Junos OS Release 14.1R3, the filename places the **output-filename** string after the timestamp.

***hostname\_YYYYMMDD\_HHMMSS\_output-filename\_index-number***

For example, on a device named `r1` running Junos OS Release 14.1R3 or a later release, if you configure the **output-filename** statement as **ifl-events**, and this event policy is triggered three times in 1 second, the files are named:

- `r1_20060623_132333_ifl-events`
- `r1_20060623_132333_ifl-events_001`
- `r1_20060623_132333_ifl-events_002`

By default, the command output format is Junos Extensible Markup Language (XML). Configure the **output-format text** statement to format the command output as ASCII text.

#### Release History Table

Release	Description
14.1R3	When eventd executes the commands defined at the <b>[edit event-options policy <i>policy-name</i> then execute-commands commands]</b> hierarchy level and generates a file containing the command output, the name of the file depends on the version of Junos OS running on the device.

#### Related Documentation

- [Example: Correlating Events Based on Receipt of Other Events Within a Specified Time Interval on page 546](#)

## Configuring an Event Policy to Change the Configuration

An event policy performs actions in response to specific events. You can configure custom event policies in the Junos OS configuration that listen for a specific event or correlated events and then execute an action, which might include creating a log file, invoking Junos OS commands, or executing an event script.

At times, it might be necessary to modify the configuration in response to a particular event. Prior to Junos OS Release 12.1, an event policy invoked an event script to execute configuration changes. Starting in Junos OS Release 12.1, you can configure an event policy to modify the configuration using Junos OS configuration mode commands and then commit the updated configuration. For example, upon receipt of an `SNMP_TRAP_LINK_DOWN` or `SNMP_TRAP_LINK_UP` event for a given interface, the event policy action might modify the configuration of a static route to adjust its metric or modify its next hop.

You configure event policy actions at the **[edit event-options policy *policy-name* then]** hierarchy level. To modify the configuration through an event policy, configure the **change-configuration** statement and specify the configuration mode commands that are executed upon receipt of the configured event or events. Enclose each command in quotation marks (`" "`), and specify the complete statement path to the element, identifier, or value as you do in configuration mode when issuing commands at the **[edit]** hierarchy level.

The event process (eventd) executes the configuration commands in the order in which they appear in the event policy configuration. The commands update the candidate configuration, which is then committed, provided that no commit errors occur.

You can configure the **commit-options** child statement to customize the event policy commit operation. You can commit the changes on a single Routing Engine or configure the **synchronize** option to synchronize the commit on both Routing Engines. When you configure the **synchronize** option, the Routing Engine on which you execute this command copies and loads its candidate configuration to the other Routing Engine. Both Routing Engines perform a syntax check on the candidate configuration file. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. By default, the **synchronize** option does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines and ignore any warnings by configuring the **force** option.

Additionally, if you are testing or troubleshooting an event policy, you can configure the **check** commit option to verify the candidate configuration syntax without committing the changes. On dual control plane systems, when the **check synchronize** statement is configured, the candidate configuration on one control plane is copied to the other control plane, and the system verifies that both candidate configurations are syntactically correct. The **check** statement and the other **commit-options** statements are mutually exclusive.

The change configuration action might fail while acquiring a lock on the configuration. Configure the **retry** statement to have the system attempt the change configuration event policy action a specified number of times if the first attempt fails. Configure the **user-name** statement to execute the configuration changes and commit under the privileges of a specific user. If you do not specify a username, the action is executed as user **root**.

#### Related Documentation

- [Example: Changing the Interface Configuration in Response to an Event on page 566](#)
- [Example: Changing the Configuration Using an Event Policy on page 559](#)
- [Event Notifications and Policies Overview on page 537](#)
- [change-configuration \(Event Policy\) on page 772](#)
- [commands \(Event Policy Change Configuration\) on page 773](#)
- [commit-options on page 775](#)
- [retry \(Event Policy\) on page 793](#)

### Example: Changing the Configuration Using an Event Policy

It might be necessary to modify the configuration in response to a particular event. Starting in Junos OS Release 12.1, you can configure an event policy to make and commit configuration changes when the event policy is triggered by one or more specific events.

This example simulates an `SNMP_TRAP_LINK_DOWN` event for a specific interface. Upon receipt of the event, the event policy modifies the configuration of a static route to use a new next-hop IP address through a different exit interface.

- [Requirements on page 560](#)
- [Overview on page 560](#)
- [Configuration on page 560](#)
- [Verification on page 563](#)
- [Troubleshooting on page 565](#)

## Requirements

- Routing, switching, or security device running Junos OS Release 12.1 or later.

## Overview

You can configure an event policy action to modify the configuration when the policy is triggered by a single event or correlated events. Suppose you have a static route to the 10.1.10.0/24 network with a next-hop IP address of 10.1.2.1 through the exit interface `ge-0/3/1`. At some point, this interface goes down, triggering an `SNMP_TRAP_LINK_DOWN` event.

This example creates an event policy named `update-on-snmp-trap-link-down`. The event policy is configured so that the `eventd` process listens for an `SNMP_TRAP_LINK_DOWN` event associated with the interface `ge-0/3/1.0`. If the interface goes down, the event policy executes a change configuration action. The event policy configuration commands remove the static route through the `ge-0/3/1` exit interface and create a new static route to the same target network with a next-hop IP address of 10.1.3.1 through the exit interface `ge-0/2/1`. The commands are executed in the order in which they appear in the event policy.

The event policy change configuration commit operation is executed under the username `bsmith` with a commit comment specifying that the change was made through the associated event policy. The retry count is set to 5 and the retry interval is set to 4 seconds. If the initial attempt to issue the configuration change fails, the system attempts the configuration change 5 additional times and waits 4 seconds between each attempt.

Although not presented here, you might have a second, similar event policy that executes a change configuration action to update the static route when the interface comes back up. In that case the policy would trigger on the `SNMP_TRAP_LINK_UP` event for the same interface.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level:

```
set event-options policy update-on-snmp-trap-link-down events snmp_trap_link_down
```



```

set event-options policy update-on-snmp-trap-link-down attributes-match
  snmp_trap_link_down.interface-name matches ge-0/3/1.0
set event-options policy update-on-snmp-trap-link-down then change-configuration
  retry count 5
set event-options policy update-on-snmp-trap-link-down then change-configuration
  retry interval 4
set event-options policy update-on-snmp-trap-link-down then change-configuration
  commands "delete routing-options static route 10.1.10.0/24 next-hop"
set event-options policy update-on-snmp-trap-link-down then change-configuration
  commands "set routing-options static route 10.1.10.0/24 next-hop 10.1.3.1"
set event-options policy update-on-snmp-trap-link-down then change-configuration
  user-name bsmith
set event-options policy update-on-snmp-trap-link-down then change-configuration
  commit-options log "updating configuration from event policy
  update-on-snmp-trap-link-down"
set routing-options static route 10.1.10.0/24 next-hop 10.1.2.1
set system syslog file syslog-event-daemon-warning daemon warning

```

### Configuring the Event Policy

#### Step-by-Step Procedure

1. Create and name the event policy.  

```

[edit]
bsmith@R1# edit event-options policy update-on-snmp-trap-link-down

```
2. Configure the **events** statement so that the event policy triggers on the SNMP\_TRAP\_LINK\_DOWN event.  

Set the **attributes-match** statement so that the policy triggers only if the SNMP\_TRAP\_LINK\_DOWN event occurs for the ge-0/3/1.0 interface.

```

[edit event-options policy update-on-snmp-trap-link-down]
bsmith@R1# set events snmp_trap_link_down
bsmith@R1# set attributes-match snmp_trap_link_down.interface-name matches
ge-0/3/1.0

```
3. Specify the configuration mode commands that are executed if the ge-0/3/1 interface goes down.  

Configure each command on a single line, enclose the command string in quotes, and specify the complete statement path.

```

[edit event-options policy update-on-snmp-trap-link-down then
change-configuration]
bsmith@R1# set commands "delete routing-options static route 10.1.10.0/24
next-hop"
bsmith@R1# set commands "set routing-options static route 10.1.10.0/24 next-hop
10.1.3.1"

```
4. Configure the commit options.  

Configure the **log** option with a comment describing the configuration changes. The comment is added to the commit logs after a successful commit operation is made through the associated event policy.

```

[edit event-options policy update-on-snmp-trap-link-down then
change-configuration]

```

```
bsmith@R1# set commit-options log "updating configuration from event policy
update-on-snmp-trap-link-down"
```

If you have dual Routing Engines, configure the **synchronize** option to commit the configuration on both Routing Engines. Include the **force** option to force the commit on the other Routing Engine, ignoring any warnings. This example does not configure the **synchronize** and **force** options.

5. (Optional) Configure the retry count and retry interval.

In this example, **count** is set to 5 and the **interval** is 4 seconds.

```
[edit event-options policy update-on-snmp-trap-link-down then
change-configuration]
bsmith@R1# set retry count 5 interval 4
```

6. (Optional) Configure the username under whose privileges the configuration changes and commit are made.

If you do not specify a username, the action is executed as user root.

```
[edit event-options policy update-on-snmp-trap-link-down then
change-configuration]
bsmith@R1# set user-name bsmith
```

7. Configure a new log file at the **[edit system syslog]** hierarchy level to record syslog events of facility **daemon** and severity **warning**.

This captures the SNMP\_TRAP\_LINK\_DOWN events.

```
[edit system syslog]
bsmith@R1# set file syslog-event-daemon-warning daemon warning
```

8. To test this example, configure a static route to the 10.1.10.0/24 network with a next hop IP address of 10.1.2.1.

```
[edit]
bsmith@R1# set routing-options static route 10.1.10.0/24 next-hop 10.1.2.1
```

9. Commit the configuration.

```
bsmith@R1# commit
```

10. Review the **[edit routing-options static]** hierarchy level of the configuration before disabling the ge-0/3/1 interface, and note the next hop IP address.

```
bsmith@R1> show configuration routing-options static
...
route 10.1.10.0/24 next-hop 10.1.2.1;
...
```

11. To manually test the event policy, take the ge-0/3/1 interface temporarily offline to generate the SNMP\_TRAP\_LINK\_DOWN event.

```
[edit]
bsmith@R1# set interfaces ge-0/3/1 disable
bsmith@R1# commit
```

## Results

```
[edit]
event-options {
  policy update-on-snmp-trap-link-down {
    events snmp_trap_link_down;
    attributes-match {
      snmp_trap_link_down.interface-name matches ge-0/3/1.0;
    }
    then {
      change-configuration {
        retry count 5 interval 4;
        commands {
          "delete routing-options static route 10.1.10.0/24 next-hop";
          "set routing-options static route 10.1.10.0/24 next-hop 10.1.3.1";
        }
        user-name bsmith;
        commit-options {
          log "updating configuration from event policy update-on-snmp-trap-link-down";
        }
      }
    }
  }
}
routing-options {
  static {
    route 10.1.10.0/24 next-hop 10.1.2.1;
  }
}
system {
  syslog {
    file syslog-event-daemon-warning {
      daemon warning;
    }
  }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Status of the Interface on page 563](#)
- [Verifying the Commit on page 564](#)
- [Verifying the Configuration Changes on page 565](#)

### Verifying the Status of the Interface

**Purpose** Verify that the ge-0/3/1 interface is down and that it triggered the SNMP\_TRAP\_LINK\_DOWN event.

**Action** Issue the **show interfaces ge-0/3/1** operational mode command. The command output shows that the interface is administratively offline.

```
bsmith@R1> show interfaces ge-0/3/1
Physical interface: ge-0/3/1, Administratively down, Physical link is Down
<output omitted>
```

Review the contents of the system log file configured in Step 7. The output shows that the ge-0/3/1.0 interface went down and generated an SNMP\_TRAP\_LINK\_DOWN event.

```
bsmith@R1> show log syslog-event-daemon-warning
Oct 10 18:00:57 R1 mib2d[1371]: SNMP_TRAP_LINK_DOWN: ifIndex 531, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0
```

### Verifying the Commit

**Purpose** Verify that the event policy commit operation was successful by reviewing the commit log and the messages log file.

**Action** Issue the **show system commit** operational mode command to view the commit log. In this example, the log confirms that the configuration was committed through the event policy under the privileges of user bsmith at the given date and time.

```
bsmith@R1> show system commit
0   2011-10-10 18:01:03 PDT by bsmith via junoscript
    updating configuration from event policy update-on-snmp-trap-link-down
1   2011-09-02 14:16:44 PDT by admin via netconf
2   2011-07-08 14:33:46 PDT by root via other
```

Review the **messages** log file. Upon receipt of the SNMP\_TRAP\_LINK\_DOWN event, Junos OS executed the configured event policy action to modify and commit the configuration. The commit operation occurred under the privileges of user bsmith.

```
bsmith@R1> show log messages | last 20
...
Oct 10 18:00:57 R1 mib2d[1371]: SNMP_TRAP_LINK_DOWN: ifIndex 531, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0
Oct 10 18:00:59 R1 file[17575]: UI_COMMIT: User 'bsmith' requested 'commit'
operation (comment: updating configuration from event policy
update-on-snmp-trap-link-down)
Oct 10 18:01:03 R1 eventd: EVENTD_CONFIG_CHANGE_SUCCESS: Configuration change
successful: while executing policy update-on-snmp-trap-link-down with user bsmith
privileges
```



**NOTE:** If you configure a different log file, review the file specific to your configuration.

**Meaning** The output from the **show system commit** operational mode command and the **messages** log file verify that the commit operation, which was made through the event policy under the privileges of the user bsmith, was successful. The **show system commit** output and

**messages** log file reference the commit comment specified in the **log** statement at the **[edit event-options policy update-on-snmp-trap-link-down then change-configuration commit-options]** hierarchy level.

### Verifying the Configuration Changes

- Purpose** Verify the configuration changes by reviewing the **[edit routing-options static]** hierarchy level of the configuration after disabling the ge-0/3/1 interface.
- Action** Issue the following operational mode command:
- ```
bsmith@R1> show configuration routing-options static
...
route 10.1.10.0/24 next-hop 10.1.3.1;
...
```
- Meaning** The configured next hop has been modified by the event policy to the new IP address 10.1.3.1, which has its route through the exit interface ge-0/2/1.

## Troubleshooting

- [Troubleshooting Commit Errors on page 565](#)

### Troubleshooting Commit Errors

- Problem** The triggered event policy does not make the specified configuration changes, and the logs verify that the commit was unsuccessful.
- ```
bsmith@R1> show log messages | last 20
...
Oct 10 17:48:59 R1 mib2d[1371]: SNMP_TRAP_LINK_DOWN: ifIndex 531, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0
Oct 10 17:49:01 R1 file[17142]: UI_LOAD_EVENT: User 'bsmith' is performing a
'rollback'
Oct 10 17:49:01 R1 eventd: EVENTD_CONFIG_CHANGE_FAILED: Configuration change
failed: rpc to management daemon failed while executing policy
update-on-snmp-trap-link-down with user bsmith privileges
```
- A failed commit might occur if the configuration is locked or if the configuration mode commands have the incorrect syntax or order.
- Solution** Check the configuration mode commands at the **[edit event-options policy update-on-snmp-trap-link-down then change-configuration commands]** hierarchy level, and verify that the syntax and the order of execution are correct.
- Additionally, increase the retry count and interval options so that if the configuration is locked, the event policy attempts the configuration changes a specified number of times after the first failed instance.
- Related Documentation**
- [Configuring an Event Policy to Change the Configuration on page 558](#)
  - [Example: Changing the Interface Configuration in Response to an Event on page 566](#)

- [change-configuration \(Event Policy\) on page 772](#)
- [commands \(Event Policy Change Configuration\) on page 773](#)
- [commit-options on page 775](#)
- [retry \(Event Policy\) on page 793](#)

---

## Example: Changing the Interface Configuration in Response to an Event

It might be necessary to modify the configuration in response to a particular event. Starting in Junos OS Release 12.1, you can configure an event policy to make and commit configuration changes when the event policy is triggered by one or more specific events.

This example uses a real-time performance monitoring (RPM) probe to generate PING\_TEST\_FAILED events for a given interface. Upon receipt of the first instance of three PING\_TEST\_FAILED events within a 60-second period from the configured RPM probe, the event policy executes a change configuration event policy action that modifies the configuration to administratively disable the specified interface. This type of action might be necessary if you have an unstable, flapping interface that is consistently affecting network performance.

- [Requirements on page 566](#)
- [Overview on page 566](#)
- [Configuration on page 567](#)
- [Verification on page 571](#)

### Requirements

- Routing, switching, or security device running Junos OS Release 12.1 or later.

### Overview

This example creates an event policy named `disable-interface-on-ping-failure`. The event policy is configured so that the `eventd` process listens for PING\_TEST\_FAILED events generated by a specific RPM probe and associated with the `ge-0/3/1` interface. If three PING\_TEST\_FAILED events occur for the given interface within a 60-second interval, the event policy executes a change configuration action. The event policy configuration commands administratively disable the interface.

To test the event policy, the example configures an RPM probe that pings the IP address associated with the `ge-0/3/1` interface. In this example, the `ge-0/3/1.0` interface is configured with the IPv4 address 10.1.4.1/26. By default, one probe is sent per test, and the example uses a 5-second pause between tests. After three successive probe losses, the RPM probe generates a PING\_TEST\_FAILED event. Because multiple RPM tests could be running simultaneously, the event policy matches the **owner-name** and **test-name** attributes of the received PING\_TEST\_FAILED events to the RPM probe owner name and test name. When the RPM probe generates three PING\_TEST\_FAILED events in a 60-second interval, it triggers the event policy, which disables the interface.

This event policy also demonstrates how to restrict the execution of the same configuration change multiple times because of occurrences of the same event or correlated events. In this example, the **within 60 trigger on 3** statement specifies that the configuration change is only triggered on the third occurrence of a PING\_TEST\_FAILED event within a 60-second interval. The **trigger until 4** statement specifies that subsequent occurrences of the PING\_TEST\_FAILED event should not cause the event policy to re-trigger.



**NOTE:** If you only configure the **trigger on 3** condition, the commit operation might go into a loop. The combination of **trigger on 3** and **trigger until 4** prevents the event policy from repeatedly making the same configuration change.

## Configuration

### Configuring the RPM Probe

#### CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set services rpm probe icmp-ping-probe test ping-probe-test probe-type icmp-ping
set services rpm probe icmp-ping-probe test ping-probe-test target address 10.1.4.1
set services rpm probe icmp-ping-probe test ping-probe-test test-interval 5
set services rpm probe icmp-ping-probe test ping-probe-test thresholds successive-loss 3
set system syslog file syslog-event-daemon-info daemon info
```

#### Step-by-Step Procedure

To configure the RPM probe, which creates the PING\_TEST\_FAILED events for this example:

1. Create an RPM probe named ping-probe-test with owner icmp-ping-probe.

```
[edit services rpm]
bsmith@R1# set probe icmp-ping-probe test ping-probe-test
```

2. Configure the RPM probe to send ICMP echo requests.

```
[edit services rpm probe icmp-ping-probe test ping-probe-test]
bsmith@R1# set probe-type icmp-ping
```

Configure the RPM probe to send the ICMP echo requests to the ge-0/3/1 interface at IP address 10.1.4.1.

```
[edit services rpm probe icmp-ping-probe test ping-probe-test]
bsmith@R1# set target address 10.1.4.1
```

3. Configure a 5-second pause between test windows.

```
[edit services rpm probe icmp-ping-probe test ping-probe-test]
bsmith@R1# set test-interval 5
```

4. Configure the RPM probe threshold so that the PING\_TEST\_FAILED event is triggered after three successive probe losses.

```
[edit services rpm probe icmp-ping-probe test ping-probe-test]
```

```
bsmith@R1# set thresholds successive-loss 3
```

5. Configure a new log file at the **[edit system syslog]** hierarchy level to record syslog events of facility **daemon** and severity **info**.

This captures the events sent during the probe tests.

```
[edit system syslog]
bsmith@R1# set file syslog-event-daemon-info daemon info
```

6. Commit the configuration.

```
bsmith@R1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show services** and **show system syslog** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
services {
  rpm {
    probe icmp-ping-probe {
      test ping-probe-test {
        probe-type icmp-ping;
        target address 10.1.4.1;
        test-interval 5;
        thresholds {
          successive-loss 3;
        }
      }
    }
  }
}
system {
  syslog {
    file syslog-event-daemon-info {
      daemon info;
    }
  }
}
```

---

### Configuring the Event Policy

#### CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set event-options policy disable-on-ping-failure events ping_test_failed
set event-options policy disable-on-ping-failure within 60 trigger on
set event-options policy disable-on-ping-failure within 60 trigger 3
set event-options policy disable-on-ping-failure within 65 trigger until
set event-options policy disable-on-ping-failure within 65 trigger 4
set event-options policy disable-on-ping-failure attributes-match
ping_test_failed.test-owner matches icmp-ping-probe
```



```

set event-options policy disable-on-ping-failure attributes-match
ping_test_failed.test-name matches ping-probe-test
set event-options policy disable-on-ping-failure then change-configuration commands
"set interfaces ge-0/3/1 disable"
set event-options policy disable-on-ping-failure then change-configuration user-name
bsmith
set event-options policy disable-on-ping-failure then change-configuration
commit-options log "updating configuration from event policy disable-on-ping-failure"

```

#### Step-by-Step Procedure

1. Create and name the event-policy.  

```

[edit]
bsmith@R1# edit event-options policy disable-interface-on-ping-failure

```
2. Configure the event policy to match on the PING\_TEST\_FAILED event.  

```

[edit event-options policy disable-interface-on-ping-failure]
bsmith@R1# set events ping_test_failed

```
3. Configure the policy to trigger when three PING\_TEST\_FAILED events occur within 60 seconds.  

```

[edit event-options policy disable-interface-on-ping-failure]
bsmith@R1# set within 60 trigger on 3

```
4. Configure the **within 65 trigger until 4** statement to prevent the policy from re-triggering if more than three PING\_TEST\_FAILED events occur.  

```

[edit event-options policy disable-interface-on-ping-failure]
bsmith@R1# set within 65 trigger until 4

```
5. Configure the **attributes-match** statement so that the event policy is only triggered by the PING\_TEST\_FAILED events generated by the associated RPM probe.  

```

[edit event-options policy disable-interface-on-ping-failure]
bsmith@R1# set attributes-match ping_test_failed.test-owner matches
icmp-ping-probe
bsmith@R1# set attributes-match ping_test_failed.test-name matches
ping-probe-test

```
6. Specify the configuration mode commands that are executed if the event policy is triggered.  

Configure each command on a single line, enclose the command string in quotes, and specify the complete statement path.

```

[edit event-options policy disable-interface-on-ping-failure then
change-configuration]
bsmith@R1# set commands "set interfaces ge-0/3/1 disable"

```
7. Configure the **log** option with a comment describing the configuration changes.  

The comment is added to the commit logs after a successful commit operation is made through the associated event policy.

```

[edit event-options policy disable-interface-on-ping-failure then
change-configuration]
bsmith@R1# set commit-options log "updating configuration from event policy
disable-interface-on-ping-failure"

```

8. (Optional) If you have dual Routing Engines, configure the **synchronize** option to commit the configuration on both Routing Engines. Include the **force** option to force the commit on the other Routing Engine, ignoring any warnings. This example does not configure the **synchronize** and **force** options.
9. (Optional) Configure the username under whose privileges the configuration changes and commit are made.

If you do not specify a username, the action is executed as user **root**.

```
[edit event-options policy disable-interface-on-ping-failure then
 change-configuration]
bsmith@R1# set user-name bsmith
```

10. Review the output of the **show interfaces ge-0/3/1** operational mode command before the configuration change takes place.



**NOTE:** The interface should be enabled.

```
bsmith@R1> show interfaces ge-0/3/1
Physical interface: ge-0/3/1, Enabled, Physical link is Up
Interface index: 142, SNMP ifIndex: 531
...
```

11. Commit the configuration.

```
bsmith@R1# commit
```

## Results

From configuration mode, confirm your configuration by entering the **show event-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit event-options]
policy disable-interface-on-ping-failure {
  events ping_test_failed;
  within 60 {
    trigger on 3;
  }
  within 65 {
    trigger until 4;
  }
  attributes-match {
    ping_test_failed.test-owner matches icmp-ping-probe;
    ping_test_failed.test-name matches ping-probe-test;
  }
  then {
    change-configuration {
      commands {
        "set interfaces ge-0/3/1 disable";
      }
      user-name bsmith;
      commit-options {
```

```

        log "updating configuration from event policy disable-interface-on-ping-failure";
    }
}
}
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Events on page 571](#)
- [Verifying the Commit on page 571](#)
- [Verifying the Configuration Changes on page 572](#)
- [Verifying the Status of the Interface on page 572](#)

### Verifying the Events

**Purpose** To manually test the event policy, take the ge-0/3/1 interface offline until three PING\_TEST\_FAILED events are generated.

**Action** Review the configured syslog file. Verify that the RPM probe generates a PING\_TEST\_FAILED event after successive lost probes.

```

bsmith@R1> show log syslog-event-daemon-info
Oct  7 15:48:54  R1 rmopd[1345]: PING_TEST_COMPLETED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test
Oct  7 15:49:54  R1 rmopd[1345]: PING_TEST_COMPLETED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test
...
Oct  7 15:52:54  R1 rmopd[1345]: RMOPD_ICMP_SENDSMSG_FAILURE: sendmsg(ICMP): No
route to host
Oct  7 15:52:54  R1 rmopd[1345]: PING_PROBE_FAILED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test
Oct  7 15:52:54  R1 rmopd[1345]: PING_TEST_FAILED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test
Oct  7 15:52:57  R1 rmopd[1345]: PING_TEST_FAILED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test
Oct  7 15:53:00  R1 rmopd[1345]: PING_TEST_FAILED: pingCtlOwnerIndex =
icmp-ping-probe, pingCtlTestName = ping-probe-test

```

### Verifying the Commit

**Purpose** Verify that the event policy commit operation was successful by reviewing the commit log and the messages log file.

**Action** Issue the **show system commit** operational mode command to view the commit log.

```
bsmith@R1> show system commit
0   2011-10-07 15:53:00 PDT by bsmith via junoscript
    updating configuration from event policy disable-interface-on-ping-failure
1   2011-09-02 14:16:44 PDT by admin via netconf
2   2011-07-08 14:33:46 PDT by root via other
```

Review the messages log file.

```
bsmith@R1> show log messages | last 20
Oct  7 15:52:54  R1 rmopd[1345]: RMOPD_ICMP_SENDSMSG_FAILURE: sendmsg(ICMP): No
route to host
Oct  7 15:53:00  R1 file[9972]: UI_COMMIT: User 'bsmith' requested 'commit'
operation (comment: updating configuration from event policy
disable-interface-on-ping-failure)
Oct  7 15:53:02  R1 eventd: EVENTD_CONFIG_CHANGE_SUCCESS: Configuration change
successful: while executing policy disable-interface-on-ping-failure with user
bsmith privileges
```

**Meaning** The output from the **show system commit** operational mode command and the **messages** log file verify that Junos OS executed the configured event policy action to modify and commit the configuration. The commit operation, which was made through the event policy under the privileges of the user bsmith at the given date and time, was successful. The **show system commit** output and **messages** log file reference the commit comment specified in the **log** statement at the **[edit event-options policy disable-interface-on-ping-failure then change-configuration commit-options]** hierarchy level.

---

### Verifying the Configuration Changes

**Purpose** Verify the configuration changes by reviewing the **[edit interfaces ge-0/3/1]** hierarchy level of the configuration.

**Action**

```
bsmith@R1> show configuration interfaces ge-0/3/1
disable;
unit 0 {
    family inet {
        address 10.1.4.1/26;
    }
}
```

**Meaning** The ge-0/3/1 configuration hierarchy was modified through the event policy to add the **disable** statement.

---

### Verifying the Status of the Interface

**Purpose** Review the output of the **show interfaces ge-0/3/1** operational mode command after the configuration change takes place.

**Action** Issue the **show interfaces ge-0/3/1** operational mode command. After the event policy configuration change action disables the interface, the status changes from "Enabled" to "Administratively down".

```
bsmith@R1> show interfaces ge-0/3/1
Physical interface: ge-0/3/1, Administratively down, Physical link is Down
Interface index: 142, SNMP ifIndex: 531
```

- Related Documentation**
- [Configuring an Event Policy to Change the Configuration on page 558](#)
  - [Example: Changing the Configuration Using an Event Policy on page 559](#)
  - [change-configuration \(Event Policy\) on page 772](#)
  - [commands \(Event Policy Change Configuration\) on page 773](#)
  - [commit-options on page 775](#)
  - [retry \(Event Policy\) on page 793](#)

## Executing Event Scripts in an Event Policy

*Event scripts* are Extensible Stylesheet Language Transformations (XSLT) scripts, Stylesheet Language Alternative Syntax (SLAX) scripts, or Python scripts that an event policy can execute when it is triggered. Event scripts can perform any function available through Junos XML or Junos XML protocol remote procedure calls (RPCs). Additionally, you can pass to an event script a set of arguments that you define.

Event scripts can build and run an operational mode command, receive the command output, inspect the output, and determine the next appropriate action. This process can be repeated until the source of the problem is determined. The script can then report the source of the problem to you on the CLI or automatically change the device configuration.

You can configure an event policy that executes event scripts and uploads the output of those scripts to a specified location for analysis. To configure such a policy, include the following statements at the **[edit event-options]** hierarchy level:

```
[edit event-options]
policy policy-name {
  events [ events ];
  then {
    event-script filename {
      arguments {
        argument-name argument-value;
      }
      destination destination-name;
      output-filename filename;
      output-format (text | xml);
    }
  }
}
```

In the **events** statement, you can list multiple events. If one or more of the listed events occurs, the eventd process executes the actions configured under the **then** hierarchy.

The event policy actions can include executing one or more event scripts, which are configured by including the **event-script filename** statement. The eventd process runs the scripts in the order in which they appear in the configuration. The scripts that you reference in the **event-script** statement must be located in the **/var/db/scripts/event** directory on the device's hard disk or the **/config/scripts/event/** directory in flash memory. Furthermore, the event scripts must be enabled at the **[edit event-options event-script file]** hierarchy level. For more information, see [“Storing and Enabling Scripts” on page 661](#).



**NOTE:** If the scripts are located in flash memory, you must configure the **load-scripts-from-flash** statement at the **[edit system scripts]** hierarchy level.



**NOTE:** For detailed information about the requirements and restrictions when executing Python automation scripts on devices running Junos OS, see [“Requirements for Executing Python Automation Scripts on Devices Running Junos OS” on page 169](#).

You can configure the **arguments** statement to pass arguments to an event script as name/value pairs. The argument values can include variables that contain information about the triggering event or other received events. The event script can then reference this information during execution. For detailed information about passing arguments to event scripts, see [“Configuring an Event Policy to Pass Arguments to an Event Script” on page 576](#).

You can configure event scripts to write their output to a file. When the eventd process executes the scripts, it uploads the file with the script output to the location specified in the **destination** statement. In the **destination** statement, include a destination name that is configured at the **[edit event-options destinations]** hierarchy level. For more information, see [“Example: Defining Destinations for File Archiving by Event Policies” on page 592](#).

In the **output-filename** statement, define a descriptive string that will be included in the filename. Each uploaded file also includes the hostname and timestamp in the filename to ensure that the each filename is unique. If a policy is triggered multiple times in a 1-second period, an index number is appended to the filename to ensure that the filenames are still unique. The index number range is 001 through 999.

When an event policy executes an event script and the script generates an output file, the name of the file depends on the version of Junos OS running on the device. Prior to Junos OS Release 14.1R3, the filename has the following naming convention:

*hostname\_output-filename\_YYYYMMDD\_HHMMSS\_index-number*

Starting in Junos OS Release 14.1R3, the filename places the **output-filename** string after the timestamp.

*hostname\_YYYYMMDD\_HHMMSS\_output-filename\_index-number*

For example, on a device named `r1` running Junos OS Release 14.1R3 or a later release, if you configure the **output-filename** statement as **ifl-events**, and this event policy is triggered three times in 1 second, the files are named:

- `r1_20060623_132333_ifl-events`
- `r1_20060623_132333_ifl-events_001`
- `r1_20060623_132333_ifl-events_002`

For the **destination** and **output-filename** statements, there are four configuration scenarios:

- You can omit the **destination** and **output-filename** statements. This option makes sense when the event script has no output. For example, the event script might execute only **request** commands, which have no output.
- You can include both the **destination** and **output-filename** statements. If you include the **output-filename** statement in the configuration, you must also include the **destination** statement in the configuration. In this case, the script output is redirected to the output file, and the file is sent to the destination specified in the configuration.
- You can include the **destination** statement in the configuration. You can omit the **output-filename** statement in the configuration and specify an output filename in the event script instead. The script output is sent to the destination specified in the configuration. If you do not include the **destination** statement in the configuration, the script output is not uploaded.

In this scenario, the event policy extracts the filename from the event script. The event script writes the output filename as **STDOUT**. The XML syntax to use in the event script is:

```
<output>
  <event-script-output-filename>filename</event-script-output-filename>
</output>
```

The **<event-script-output-filename>** element must be the first child tag within the **<output>** parent tag.

On a device named **device2**, configure an event script action with a destination **host**, and omit the **output-filename** statement. Define the destination **host** as `ftp://user@device1/tmp`.

In the **script1.xml** event script, write the following output to **STDOUT**:

```
<event-script-output-filename>/var/cmd.txt</event-script-output-filename>
```

Configure the **policy1** event policy as follows:

```
[edit event-options]
policy policy1 {
  then {
    event-script script1.xml {
      destination host;
    }
  }
}
destinations {
  host {
```

```

archive-sites {
  "ftp://user@device1//tmp" password "$ABC123"; ## SECRET-DATA***
}
}

```

In this example, the `/var/cmd.txt` file resides on device2. The event policy uses the File Transfer Protocol (FTP) to upload this file to the `/tmp` directory on device1.

The event policy reads the output filename `/var/cmd.txt` from **STDOUT**. Then the event policy uploads the `/var/cmd.txt` file to the configured destination, which is the `/tmp` directory on device1. The event policy renames the `/var/cmd.txt` file as **device2\_YYYYMMDD\_HHMMSS\_cmd.txt\_index-number**.

- You can include the **destination** and **output-filename** statements and also specify an output filename directly within the event script. If you do this, the output filename specified in the configuration overrides the output filename specified in the event script.

The default and only format for event script output is Junos Extensible Markup Language (XML).



**NOTE:** Event script output is always emitted in XML format even if you configure the `output-format text` statement under the `[edit event-options policy policy-name then event-script filename]` hierarchy.

#### Release History Table

Release	Description
14.1R3	When an event policy executes an event script and the script generates an output file, the name of the file depends on the version of Junos OS running on the device.

#### Related Documentation

- [Configuring an Event Policy to Pass Arguments to an Event Script on page 576](#)

## Configuring an Event Policy to Pass Arguments to an Event Script

When an event policy invokes an event script, the policy can pass arguments to the script. The following sections outline how to configure the arguments in the event policy and use the arguments within the event script:

- [Configuring Event Script Arguments in an Event Policy on page 576](#)
- [Using Arguments in an Event Script on page 578](#)

### Configuring Event Script Arguments in an Event Policy

You configure the arguments that an event policy passes to an event script within the **then** clause of the policy under the **event-script filename arguments** hierarchy. You can configure any number of arguments for each invoked event script.

[edit **event-options** policy *policy-name* then]



```
event-script filename {
  arguments {
    argument-name argument-value;
  }
}
```

You include arguments to the script as name/value pairs. The argument values can include variables that contain information about the triggering event or other received events. The event script can then reference this information during execution. You can use variables of the following forms:

- **`{{attribute-name}}`**—The double dollar sign (`$$`) notation represents the event that is triggering a policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event. For example, **`{{interface-name}}`** resolves to the interface name associated with the triggering event.
- **`$event.attribute-name`**—The single dollar sign with the event name (**`$event`**) notation represents the most recent event that matches **`event`**. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event. For example, **`$(COSD_CHAS_SCHED_MAP_INVALID.interface-name)`** resolves to the interface name associated with the most recent **`COSD_CHAS_SCHED_MAP_INVALID`** event cached by the `eventd` process.

For a given event, you can view a list of event attributes that you can reference by issuing the **`help syslog event`** command.

```
user@host> help syslog event
```

For example, in the following command output, text in angle brackets (`< >`) shows attributes of the **`COSD_CHASSIS_SCHEDULER_MAP_INVALID`** event:

```
user@host> help syslog COSD_CHASSIS_SCHEDULER_MAP_INVALID
Name:      COSD_CHASSIS_SCHEDULER_MAP_INVALID
Message:    Chassis scheduler map incorrectly applied to interface
<interface-name>: <error-message>
...
```

Another way to view a list of event attributes is to issue the **`set attributes-match event?`** configuration mode command at the **`[edit event-options policy policy-name]`** hierarchy level.

```
[edit event-options policy policy-name]
user@host# set attributes-match event?
```

For example, in the following command output, the **`event.attribute`** list shows that **`error-message`** and **`interface-name`** are attributes of the **`cosd_chassis_scheduler_map_invalid`** event:

```
[edit event-options policy p1]
user@host# set attributes-match cosd_chassis_scheduler_map_invalid?
Possible completions:
<from-event-attribute> First attribute to compare
cosd_chassis_scheduler_map_invalid.error-message
cosd_chassis_scheduler_map_invalid.interface-name
```

In this **`set`** command, there is no space between the event name and the question mark (`?`).

To view a list of all event attributes that you can reference, issue the **set attributes-match ?** configuration mode command at the **[edit event-options policy *policy-name*]** hierarchy level.

```
[edit event-options policy policy-name]  
user@host# set attributes-match ?  
Possible completions:  
  <from-event-attribute> First attribute to compare  
  acct_accounting_error  
  acct_accounting_fopen_error  
  ...
```

## Using Arguments in an Event Script

When an event policy invokes an event script, the event script can reference any of the arguments that are passed in by the policy. The names of the arguments in the event script must match the names of the arguments configured for that event script under the **[edit event-options policy *policy-name* then event-scripts *filename* arguments]** hierarchy in the configuration.

To use the arguments within a SLAX or XSLT event script, you must include a parameter declaration for each argument. The event script assigns the value for each script argument to the corresponding parameter of the same name, which can then be referenced throughout the script.

**XSLT Syntax**      `<xsl:param name="argument-name" />`

**SLAX Syntax**      `param $argument-name;`

To use the arguments within a Python event script, you can use any valid means in the Python language. The following example uses the Python **argparse** module to process the script arguments. A **parser.add\_argument** statement must be included for each argument passed to the script.

**Python Syntax**

```
import argparse  
  
def main():  
    parser = argparse.ArgumentParser()  
    parser.add_argument('-argument-name', required=True)  
    args = parser.parse_args()  
  
    # to use the argument reference args.argument-name  
  
if __name__ == '__main__':  
    main()
```

---

## Configuring Event Policies to Ignore an Event

You can modify a policy to cause particular events to be ignored or to cause all events to be ignored during a particular time interval, to allow for maintenance for example. To configure such a policy, include the following statements at the **[edit event-options]** hierarchy level:

```
[edit event-options]  
policy policy-name {  
  events [events];
```

```

then {
  ignore;
}
}

```

In the **events** statement, you can list multiple events. To view a list of the events that can be referenced in an event policy, issue the **set event-options policy *policy-name* events ?** configuration mode command:

```

[edit]
user@host# set event-options policy policy-name events ?
Possible completions:
  <event>
  [                               Open a set of values
  acct_accounting_ferror
  acct_accounting_fopen_error
  ...

```

Some of the system log messages that you can reference in an event policy are not listed in the output of the **set event-options policy *policy-name* events ?** command. For information about referencing these system log messages in your event policies, see [“Using Nonstandard System Log Messages to Trigger Event Policies” on page 552](#).

In addition, you can reference internally generated events, which are discussed in [“Generating Internal Events to Trigger Event Policies” on page 551](#).

If one or more of the listed events occur, a system log message for the event is not generated, and no further policies associated with this event are processed. If you include the **ignore** statement in a policy configuration, you cannot configure any other actions in the policy.

## Example: Ignoring Events Based on Receipt of Other Events

In the following policy, if any of **event1**, **event2**, or **event3** has occurred, and either **event4** or **event5** has occurred within the last 600 seconds, and **event6** has not occurred within the last 800 seconds, then the event that triggered the policy (**event1**, **event2**, or **event3**) is ignored, meaning system log messages are not created.

```

[edit event-options]
policy 1 {
  events [ event1 event2 event3 ];
  within 600 events [ event4 event5 ];
  within 800 not events event6;
  then {
    ignore;
  }
}

```

Sometimes events are generated repeatedly within a short period of time. In this case, it is redundant to execute a policy multiple times, once for each instance of the event. Event dampening allows you to slow down the execution of policies by ignoring instances of an event that occur within a specified time after another instance of the same event.

In the following example, an action is taken only if the eventd process has not received another instance of the event within the past 60 seconds. If an instance of the event has been received within the last 5 seconds, the policy is not executed and a system log message for the event is not created again.

```
[edit event-options]
policy dampen-policy {
  events event1;
  within 60 events event1;
  then {
    ignore;
  }
}
policy policy {
  events event1;
  then {
    ... actions ...
  }
}
```

---

## Overview of Using Event Policies to Raise SNMP Traps

SNMP *traps* enable an agent to notify a network management system (NMS) of significant events by way of an unsolicited SNMP message. You can configure an event policy action that raises traps for events based on system log messages. If one or more of the listed events occur, the system log message for the event is converted into a trap. This enables notification of an SNMP trap-based application when an important system log message occurs. You can convert any system log message (for which there are no corresponding traps) into a trap. This is helpful if you use NMS traps rather than system log messages to monitor your network.

To configure an event policy that raises a trap on receipt of an event, include the following statements at the **[edit event-options policy *policy-name*]** hierarchy level:

```
[edit event-options policy policy-name]
events [ events ];
then {
  raise-trap;
}
```

The Juniper Networks enterprise-specific System Log MIB, whose object identifier is {jnxMibs 35}, provides support for this feature.

### Related Documentation

- *Configuring SNMP on a Device Running Junos OS*
- [Example: Raising an SNMP Trap in Response to an Event on page 581](#)

## Example: Raising an SNMP Trap in Response to an Event

This example configures an event policy to raise a trap and to execute an event script in response to an event:

- [Requirements on page 581](#)
- [Overview on page 581](#)
- [Configuration on page 581](#)

### Requirements

A device running Junos OS, which is configured for SNMP.

### Overview

The following example configures the event policy **raise-trap-on-ospf-nbrdown** to trigger on the **RPD\_OSPF\_NBRDOWN** event, which indicates a terminated OSPF adjacency with a neighboring router. The event policy action raises a trap in response to the event. The device sends a notification to the SNMP manager, if one is configured under the **[edit snmp]** hierarchy level.

Additionally, the event policy executes the event script **ospf.xml** in response to this event and provides the affected interface as an argument to the script. The **\$\$rpd\_ospf\_nbrdown.interface-name** argument resolves to the interface name associated with the triggering event.

The event script output is recorded in the file **ospf-out**, and the output file is uploaded to the destination **mgmt-archives**, which is configured at the **[edit event-options destinations]** hierarchy level. To invoke an event script in an event policy, the event script must be present in the **/var/db/scripts/event** directory on the hard disk, and it must be enabled in the configuration.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set event-options policy raise-trap-on-ospf-nbrdown events rpd_ospf_nbrdown
set event-options policy raise-trap-on-ospf-nbrdown then raise-trap
set event-options policy raise-trap-on-ospf-nbrdown then event-script ospf.xml arguments
  interface "${$rpd_ospf_nbrdown.interface-name}"
set event-options policy raise-trap-on-ospf-nbrdown then event-script ospf.xml
  output-filename ospf-out
set event-options policy raise-trap-on-ospf-nbrdown then event-script ospf.xml destination
  mgmt-archives
```

## Configuring the Event Policy

### Step-by-Step Procedure

To configure an event policy that raises a trap on receipt of an event and optionally executes an event script:

1. Create and name the event-policy.

```
[edit]
user@R1# edit event-options policy raise-trap-on-ospf-nbrdown
```

2. Configure the event policy to match on the desired event, which in this example is the **RPD\_OSPF\_NBRDOWN** event.

```
[edit event-options policy raise-trap-on-ospf-nbrdown]
user@R1# set events rpd_ospf_nbrdown
```

3. Configure the event policy action to raise an SNMP trap in response to the event.

```
[edit event-options policy raise-trap-on-ospf-nbrdown]
user@R1# set then raise-trap
```

4. (Optional) Configure additional actions to take in response to the event.

This example executes an event script and uploads the associated output file to a predefined destination.

```
[edit event-options policy raise-trap-on-ospf-nbrdown]
user@R1# set then event-script ospf.xsl arguments interface
    {$$rpd_ospf_nbrdown.interface-name}
user@R1# set then event-script ospf.xsl output-filename ospf-out destination
    mgmt-archives
```

5. Commit the configuration.

```
user@R1# commit
```

### Results

```
[edit event-options]
policy raise-trap-on-ospf-nbrdown {
  events rpd_ospf_nbrdown;
  then {
    event-script ospf.xsl {
      arguments {
        interface "{$$rpd_ospf_nbrdown.interface-name}";
      }
      output-filename ospf-out;
      destination mgmt-archives;
    }
    raise-trap;
  }
}
```

### Related Documentation

- *Configuring SNMP on a Device Running Junos OS*
- [Overview of Using Event Policies to Raise SNMP Traps on page 580](#)

---

## Understanding the Event System Log Priority in an Event Policy

---

Starting in Junos OS Release 12.1, you can configure an event policy to override the default system log priority of a triggering event so that the system logs the event with a different facility type, severity level, or both. To override the priority of the triggering event, configure the **priority-override** statement at the **[edit event-options policy policy-name then]** hierarchy level. To override the facility type with which the triggering event is logged, include the **facility** statement and the new facility type. To override the severity level with which the triggering event is logged, include the **severity** statement and the new severity level.

Junos OS processes generate system log messages, or event notifications, to record the events that occur on a routing, switching, or security platform. Each system log message identifies the Junos OS process that generated the message and describes the operation or error that occurred. The Junos OS event process (eventd) receives the event notifications, and configured event policies instruct the eventd process to perform a set of actions upon receipt of specific events or correlated events.

Each system log message belongs to a facility, which groups messages that either are generated by the same source (such as a software process) or concern a similar condition or activity (such as authentication attempts). Each message is also preassigned a severity level, which indicates how seriously the triggering event affects the functions of the routing, switching, or security platform. A message's facility and severity level are together referred to as its priority. For more information about facility and severity levels, see *Junos OS System Logging Facilities and Message Severity Levels*.

When you configure logging on a device for a specific facility and destination, you also specify a severity level. Messages from that facility that are rated at the configured severity level or higher are logged. To log related events with different severity levels in the same log file, you must filter events using the lowest severity level of any of the events from that facility to be logged. This can result in unwieldy log files that are difficult and time-consuming to parse.

For example, Junos OS logs the protocol UP and DOWN events with different severity levels. Both the SNMP\_TRAP\_LINK\_DOWN and SNMP\_TRAP\_LINK\_UP events have a facility of 'daemon', but the SNMP\_TRAP\_LINK\_DOWN event has a severity level of 'warning', and the SNMP\_TRAP\_LINK\_UP event has a severity level of 'info'. Normally, when you configure a system log file, you must filter events to that file using the lower severity level of 'info' in order to log both of the events.

The event policy **priority-override** statement enables you to customize the priority of the triggering event so that it is logged using a different facility type and severity level. Suppose you configure a system log file to filter events of facility 'daemon' and severity 'notice', and you have event policies that trigger on the RPD\_ISIS\_ADJDOWN and RPD\_ISIS\_ADJUP events. When the system generates an RPD\_ISIS\_ADJDOWN message reporting that the IS-IS adjacency with a neighboring router was terminated, this message is logged. However, if the system subsequently generates an RPD\_ISIS\_ADJUP event notification reporting that the IS-IS adjacency has been restored, by default, the message is not logged, because it has a lower severity level of 'info'. In the event policy that triggers on the RPD\_ISIS\_ADJUP event, you can configure the associated priority so that the triggering

RPD\_ISIS\_ADJUP event is logged with a severity level of 'notice' and is captured in the configured log file.



**NOTE:** Event policies are executed in the order in which they appear in the configuration. When you configure multiple event policies to override the priority of the same event, the event is logged based on the priority set by the last executed event policy to change it.

#### Related Documentation

- [Example: Configuring the Event System Log Priority in an Event Policy on page 584](#)
- *Junos OS System Logging Facilities and Message Severity Levels*
- *Specifying the Facility and Severity of Messages to Include in the Log*
- [facility on page 784](#)
- [priority-override \(Event Policy\) on page 792](#)
- [severity \(Event Policy\) on page 795](#)

---

## Example: Configuring the Event System Log Priority in an Event Policy

---

It is necessary to log events when monitoring, managing, and troubleshooting routing, switching, and security devices. You can configure an event policy to override the priority of its triggering event so that it is logged based on a different facility type and severity level. This enables the event to be logged even if the system filters events to the destination log file using a different facility type or a higher severity level.

This example simulates an SNMP\_TRAP\_LINK\_UP event for a specific interface. Upon receipt of the event, the event policy overrides the severity level of the event so that it is captured in the configured log file.

- [Requirements on page 584](#)
- [Overview on page 584](#)
- [Configuration on page 585](#)
- [Verification on page 587](#)

### Requirements

- Routing, switching, or security device running Junos OS Release 12.1 or later.
- Interface is configured and active. This example uses the ge-0/3/1.0 interface.

### Overview

This example configures two log files to capture events of facility 'daemon'. One log file is configured to filter for events of severity 'warning' or higher, and the second log file is configured to filter for events of severity 'info' or higher.



The configured event policy triggers on the SNMP\_TRAP\_LINK\_UP event for interface ge-0/3/1.0. The example generates an SNMP\_TRAP\_LINK\_DOWN event followed by an SNMP\_TRAP\_LINK\_UP event for the ge-0/3/1.0 interface. The SNMP\_TRAP\_LINK\_DOWN event, which has a severity level of 'warning' is captured in both configured log files. Upon receipt of the SNMP\_TRAP\_LINK\_UP event, the event policy overrides the severity level of the event to 'warning' so that it is also captured in the log file that filters for events of severity 'warning'. By default, if the event policy does not override the severity level of this event, it is only captured in the log file that filters for the severity level 'info'.

## Configuration

- [Configuring the Log Files on page 585](#)
- [Verifying the Default System Log Priority of the Events on page 586](#)
- [Configuring the Event Policy on page 586](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
set system syslog file syslog-event-daemon-info daemon info
set system syslog file syslog-event-daemon-warning daemon warning
set event-options policy log-on-snmp-trap-link-up events snmp_trap_link_up
set event-options policy log-on-snmp-trap-link-up attributes-match
  snmp_trap_link_up.interface-name matches ge-0/3/1.0
set event-options policy log-on-snmp-trap-link-up then priority-override severity warning
```

### Configuring the Log Files

#### Step-by-Step Procedure

1. Configure two log files at the **[edit system syslog]** hierarchy level to record events of facility **daemon**.

Configure one log to record events of severity 'info' or higher and one log file to record events of severity 'warning' or higher.

```
[edit system syslog]
bsmith@R1# set file syslog-event-daemon-info daemon info
bsmith@R1# set file syslog-event-daemon-warning daemon warning
```

2. Commit the configuration.

```
bsmith@R1# commit
```

3. To manually test the logging of the events, take the ge-0/3/1.0 logical interface temporarily offline, and then bring it back up.

This generates an SNMP\_TRAP\_LINK\_DOWN event followed by an SNMP\_TRAP\_LINK\_UP event.

```
[edit]
bsmith@R1# set interfaces ge-0/3/1 unit 0 disable
bsmith@R1# commit
bsmith@R1# delete interfaces ge-0/3/1 unit 0 disable
bsmith@R1# commit
```

```

Results    [edit]
           system {
             syslog {
               file syslog-event-daemon-info {
                 daemon info;
               }
               file syslog-event-daemon-warning {
                 daemon warning;
               }
             }
           }

```

---

### Verifying the Default System Log Priority of the Events

**Purpose** Verify that the system generated the SNMP\_TRAP\_LINK\_DOWN and SNMP\_TRAP\_LINK\_UP events for the ge-0/3/1.0 interface, and note where each event is logged.

**Action** Review the contents of the **syslog-event-daemon-info** file configured in Step 1 of the previous procedure. The output shows that the ge-0/3/1.0 interface was brought down and back up and generated an SNMP\_TRAP\_LINK\_DOWN event followed by an SNMP\_TRAP\_LINK\_UP event.

```

bsmith@R1> show log syslog-event-daemon-info
Oct 24 13:22:17 R1 mib2d[1394]: SNMP_TRAP_LINK_DOWN: ifIndex 539, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0
...
Oct 24 13:22:29 R1 mib2d[1394]: SNMP_TRAP_LINK_UP: ifIndex 539, ifAdminStatus
up(1), ifOperStatus up(1), ifName ge-0/3/1.0

```

Review the contents of the **syslog-event-daemon-warning** file configured in Step 1 of the previous procedure. Because the severity level of the SNMP\_TRAP\_LINK\_UP event is 'info', it does not appear in a log file that is configured to only record events of severity 'warning' or higher. By default, this system log file captures the SNMP\_TRAP\_LINK\_DOWN events, but does not capture the SNMP\_TRAP\_LINK\_UP events.

```

bsmith@R1> show log syslog-event-daemon-warning
Oct 24 13:22:17 R1 mib2d[1394]: SNMP_TRAP_LINK_DOWN: ifIndex 539, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0

```

**Meaning** Because the SNMP\_TRAP\_LINK\_UP event has a default severity of 'info', it is not forwarded to log files that are configured to capture events of higher severity.

---

### Configuring the Event Policy

**Step-by-Step Procedure** 1. Create and name the event-policy.

```

[edit]
bsmith@R1# edit event-options policy log-on-snmp-trap-link-up

```

2. Configure the **events** statement.

For this example, the event policy triggers on the SNMP\_TRAP\_LINK\_UP event. Set the **attributes-match** statement so that the policy triggers only if the SNMP\_TRAP\_LINK\_UP event occurs for the ge-0/3/1.0 interface.

```
[edit event-options policy log-on-snmp-trap-link-up]
bsmith@R1# set events snmp_trap_link_up
bsmith@R1# set attributes-match snmp_trap_link_up.interface-name matches
ge-0/3/1.0
```

3. Configure the **priority-override** event policy action, and include the **severity** statement with a value of **warning**.

```
[edit event-options policy log-on-snmp-trap-link-up]
bsmith@R1# set then priority-override severity warning
```

4. Commit the configuration.

```
bsmith@R1# commit
```

5. To manually test the event policy, take the ge-0/3/1.0 logical interface temporarily offline, and then bring it back up. This generates an SNMP\_TRAP\_LINK\_DOWN event followed by an SNMP\_TRAP\_LINK\_UP event.

```
[edit]
bsmith@R1# set interfaces ge-0/3/1 unit 0 disable
bsmith@R1# commit
bsmith@R1# delete interfaces ge-0/3/1 unit 0 disable
bsmith@R1# commit
```

**Results**

```
[edit]
event-options {
  policy log-on-snmp-trap-link-up {
    events snmp_trap_link_up;
    attributes-match {
      snmp_trap_link_up.interface-name matches ge-0/3/1.0;
    }
    then {
      priority-override {
        severity warning;
      }
    }
  }
}
```

## Verification

Confirm that the configuration is working properly.

### Verifying the Configured System Log Priority of the Events

**Purpose** Verify that the system generated the SNMP\_TRAP\_LINK\_DOWN and SNMP\_TRAP\_LINK\_UP events for the ge-0/3/1.0 interface, and note where each event is logged.

**Action** Review the contents of the `syslog-event-daemon-warning` file. Because the event policy overrides the severity level of the `SNMP_TRAP_LINK_UP` event, it now appears in the log file that is configured to only record events of severity 'warning' or higher. By default, this system log file captures the `SNMP_TRAP_LINK_DOWN` events, but does not capture the `SNMP_TRAP_LINK_UP` events.

```
bsmith@R1> show log syslog-event-daemon-warning
Oct 24 13:29:48 R1 mib2d[1394]: SNMP_TRAP_LINK_DOWN: ifIndex 539, ifAdminStatus
down(2), ifOperStatus down(2), ifName ge-0/3/1.0
Oct 24 13:30:02 R1 mib2d[1394]: SNMP_TRAP_LINK_UP: ifIndex 539, ifAdminStatus
up(1), ifOperStatus up(1), ifName ge-0/3/1.0
```

**Meaning** Although the `SNMP_TRAP_LINK_UP` event has a severity of 'info', configuring the **priority-override** statement with a severity of 'warning' causes the event to be forwarded to the system logs with the configured severity level. The event can be captured in logs that filter for a different facility type and a higher severity level.

- Related Documentation**
- [Understanding the Event System Log Priority in an Event Policy on page 583](#)
  - [Junos OS System Logging Facilities and Message Severity Levels](#)
  - [Specifying the Facility and Severity of Messages to Include in the Log](#)
  - [facility on page 784](#)
  - [priority-override \(Event Policy\) on page 792](#)
  - [severity \(Event Policy\) on page 795](#)

---

## Example: Limiting Event Script Output Based on a Specific Event Type

In situations where an event policy is triggered by multiple event types, you can limit the number of events that trigger the event script. For example, the following event policy triggers the `event-details.slax` event script whenever a `ui_login_event` or `ui_logout_event` occurs.

```
event-options {
  policy event-detail {
    events [ ui_login_event ui_logout_event ];
    then {
      event-script event-details.slax {
        output-filename systemlog;
        destination /tmp;
      }
    }
  }
}
```

The `event-details.slax` event script writes a log file only when the `ui_login_event` event occurs.

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
```

```

ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

var $event-definition = {
  <event-options> {
    <policy> {
      <namex> "event-detail";
      <eventsx> "ui_login_event";
      <thenx> {
        <event-scriptx> {
          <namex> "event_detail.slax";
          <output-filenamex> "foo";
          <destinationx> {
            <namex> "foo";
          }
        }
      }
    }
  }
}

match / {
  <event-script-resultsx> {
    <event-triggered-this-policyx> {
      expr event-script-input/trigger-event/id;
    }
    <type-of-eventx> {
      expr event-script-input/trigger-event/type;
    }
    <process-namex> {
      expr event-script-input/trigger-event/attribute-list/attribute/name;
    }
  }
}

```



# Configuring Event Policy File Archiving

- [Event Policy File Archiving Overview on page 591](#)
- [Example: Defining Destinations for File Archiving by Event Policies on page 592](#)
- [Example: Configuring an Event Policy to Upload Files on page 595](#)
- [Configuring the Delay Before Files Are Uploaded by an Event Policy on page 602](#)
- [Example: Configuring the Delay Before Files Are Uploaded by an Event Policy on page 603](#)
- [Configuring an Event Policy to Retry the File Upload Action on page 605](#)
- [Example: Configuring an Event Policy to Retry the File Upload Action on page 606](#)

## Event Policy File Archiving Overview

---

Various types of files are useful in diagnosing events. When an event policy action generates output files, you can archive the files for later analysis. Similarly, you might want to archive system files, including system log files, core files, and configuration files, from the time an event occurs.

When an event occurs, you can upload relevant files to a specified location for analysis. To archive files from an event policy, configure one or more *destinations* specifying the archive sites to which the files are uploaded. You then reference the configured destinations within event policies.

A transfer delay allows you to specify the number of seconds the event process (eventd) waits before beginning to upload a file or multiple files. A transfer delay helps ensure that a large file, such as a core file, is completely generated before the upload begins.

You can associate transfer delays with a destination and with an event policy action. If you associate a transfer delay with a destination, the transfer delay applies to all file upload actions that use that destination. You can also assign a transfer delay to a single event policy action. For example, you might have multiple event policy actions that use the same destination, and for some of these event policy actions, you want a transfer delay, and for other event policy actions you want no transfer delay. If you configure a transfer delay for a destination, and you also configure a transfer delay for the event policy action, the resulting transfer delay is the sum of the two delays.

Transient network problems can cause a file upload operation to fail. By default, if the file upload operation fails for any reason, the event policy does not retry the upload operation. However, you can configure an event policy to retry the file upload operation

a specified number of times if the initial upload fails. You can also configure the time interval between each retry attempt.

**Related  
Documentation**

- [Example: Defining Destinations for File Archiving by Event Policies on page 592](#)
- [Example: Configuring an Event Policy to Upload Files on page 595](#)
- [Configuring the Delay Before Files Are Uploaded by an Event Policy on page 602](#)
- [Configuring an Event Policy to Retry the File Upload Action on page 605](#)

---

## Example: Defining Destinations for File Archiving by Event Policies

This example configures an archive site for event policies. Event policy actions that reference the configured destination upload specified files to that site.

- [Requirements on page 592](#)
- [Overview on page 592](#)
- [Configuration on page 594](#)
- [Verification on page 594](#)

### Requirements

This example uses a device running Junos OS. No additional configuration beyond device initialization is required before configuring this example.

### Overview

When an event policy action generates output files, you can archive the files for later analysis. Similarly, you might want to archive system files, including system log files, core files, and configuration files, from the time an event occurs.

When an event occurs, you can upload relevant files to a specified location for analysis. To archive files from an event policy, configure one or more *destinations* specifying the archive sites to which the files are uploaded. You then reference the configured destinations within event policies.

To define a destination archive site, include the **destinations** statement at the **[edit event-options]** hierarchy level.

```
[edit event-options]
destinations {
  destination-name {
    archive-sites {
      url <password password>;
    }
    transfer-delay seconds;
  }
}
```

The *destination-name* is a user-defined identifier, which is referenced by event policies. You can define multiple destinations with different archive sites.



For each destination, configure one or more archive site URIs, which are the actual sites to which the files are uploaded. If you specify multiple archive site URIs for a given destination, the device attempts to transfer the files to the first archive site in the list and only uses subsequent sites in the list if the transfer to the first site fails. Optionally, you can specify a plain-text password for login into an archive site.

Specify the archive site URI as a file URI, an active or passive FTP URI, or a Secure Copy (SCP) URI. Local device directories are also supported (for example, `/var/tmp`). When you specify the archive site URI, do not add a forward slash (/) to the end of the URI.

```
file:<://host>/path
ftp://username@host:<port>url-path
pasvftp://username@host:<port>url-path
scp://username@host:<port>url-path
<path>/<filename>
```

When an event policy action uploads files, the name of the file depends on the version of Junos OS running on the device. Prior to Junos OS Release 14.1R3, the filename has the following naming convention:

***hostname\_output-filename\_YYYYMMDD\_HHMMSS\_index-number***

Starting in Junos OS Release 14.1R3, the filename has the following naming convention:

***hostname\_YYYYMMDD\_HHMMSS\_output-filename\_index-number***

The ***output-filename*** string is either the name of an existing file or the value configured for the **output-filename** statement within the event policy **then** clause under the hierarchy for the appropriate event policy action. The ***index-number*** string, which ranges from 001 to 999, is appended to the filename in the event that the policy is triggered multiple times in a 1-second period.

The **transfer-delay** statement allows you to specify the number of seconds the event process (eventd) waits before beginning to upload a file or multiple files to that destination. A transfer delay allows you to ensure that a large file, such as a core file, is completely generated before the upload begins. For more information, see [“Configuring the Delay Before Files Are Uploaded by an Event Policy” on page 602](#).

This example configures a new archive destination named `mgmt-archives`, which can be referenced in event policies for file archiving. The example configures two archive sites for this destination. The first site is the Secure Copy URI

`"scp://username@example.com/test"` for which a password is configured. The second site is a directory on the local device. The device attempts to transfer to the first archive site in the list, moving to the next site only if the transfer to the first site fails. The example configures a transfer delay of five seconds for all files uploaded to the `mgmt-archives` archive site.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
set event-options destinations mgmt-archives archive-sites
  "scp://username@example.com/test" password PaSsWoRd
set event-options destinations mgmt-archives archive-sites /var/log
set event-options destinations mgmt-archives transfer-delay 5
```

**Step-by-Step Procedure** Configure a new archive destination named mgmt-archives that can be referenced by event-policies.

1. Configure the identifier and associated archive sites for each destination.

The device transfers to the first archive site in the list, moving to the next site only if the transfer to the first site fails.

```
[edit event-options destinations]
user@host# set mgmt-archives archive-sites scp://username@example.com/test
user@host# set mgmt-archives archive-sites /var/log
```

2. If authentication is required to access any of the archive sites, configure the required plain-text password for that site.

```
[edit event-options destinations]
user@host# set mgmt-archives archive-sites scp://username@example.com/test
  password PaSsWoRd
```

3. (Optional) Configure the transfer delay associated with each destination. The mgmt-archives destination has a transfer delay of five seconds.

```
[edit event-options destinations]
user@host# set mgmt-archives transfer-delay 5
```

4. Commit the configuration.

```
user@host# commit
```

5. You can reference configured destinations in an event policy. For information about referencing destinations in event policies, see [“Example: Configuring an Event Policy to Upload Files” on page 595](#) and [“Configuring an Event Policy to Execute Operational Mode Commands” on page 555](#).

## Verification

---

### Verifying the Configuration

**Purpose** Issue the **show configuration event-options** operational mode command to review the resulting configuration.

**Action**

```
user@host> show configuration event-options
destinations {
  mgmt-archives {
```

```

transfer-delay 5;
archive-sites {
    "scp://username@example.com/test" password "$ABC123"; ## SECRET-DATA
    /var/log;
}
}
}

```

**Meaning** In the sample output, the mgmt-archives destination has two archive sites and a transfer delay of five seconds. You can now reference this destination in event policies. When you reference the mgmt-archives destination in an event policy, specified files are uploaded to the first archive site after a five second delay. If the transfer to the first archive fails, the device attempts to upload the files to the **/var/log** archive site. For more information about referencing destinations in event policies, see [“Example: Configuring an Event Policy to Upload Files” on page 595](#).

Note that although the plain-text password is visible when you configure it, the configuration displays the encrypted password.

**Related Documentation**

- [Example: Configuring an Event Policy to Upload Files on page 595](#)
- [Configuring the Delay Before Files Are Uploaded by an Event Policy on page 602](#)
- [Configuring an Event Policy to Retry the File Upload Action on page 605](#)
- [destinations \(Event Policy\) on page 777](#)

## Example: Configuring an Event Policy to Upload Files

This example configures event policy actions that upload relevant files to a specified location for analysis.

- [Requirements on page 595](#)
- [Overview on page 595](#)
- [Configuration on page 598](#)
- [Verification on page 601](#)

### Requirements

Before you begin:

- Configure the destinations that you will reference in the event policy. See [“Example: Defining Destinations for File Archiving by Event Policies” on page 592](#).
- Configure the general event policy and triggering events.

### Overview

When an event policy action generates output files, you can archive the files for later analysis. Similarly, you might want to archive system files, including system log files, core files, and configuration files, from the time an event occurs. You can configure an event

policy to upload existing system files or to upload the output files generated from an invoked event-script or command at the time an event occurs. This section outlines the configuration required for uploading each of these file types using an event policy.

When you configure an event policy to upload files, the relevant files are uploaded to the location referenced in the **destination** statement configured for that event policy action. You must specify a destination name that is configured at the **[edit event-options destinations]** hierarchy level.

To upload system files to a configured archive site, configure the **upload** statement at the **[edit event-options policy *policy-name* then]** hierarchy level. If the configured events occur, the eventd process executes the upload action.

```
[edit event-options policy policy-name then]
upload filename (filename | committed) destination destination-name {
  retry-count number retry-interval seconds;
  transfer-delay seconds;
  user-name username;
}
```

The **upload filename committed destination *destination-name*** statement uploads the committed configuration file.

If desired, you can include multiple **upload** statements, one for each type of file to be archived. In the **filename** statement, specify a file or multiple files to be uploaded. You can specify multiple files with one **filename** configuration statement (sometimes called *filename globbing*). For example, to upload all files that are located in the **/var/log** directory and that start with the string "messages", include the following statement:

```
upload filename /var/log/messages* destination destination-name;
```

When an event policy executes commands in response to an event, you can write the command output to a file. To configure an event policy to upload the generated output file to a configured archive site, include the following statements at the **[edit event-options policy *policy-name* then]** hierarchy level:

```
[edit event-options policy policy-name then]
execute-commands {
  destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
  }
  output-filename filename;
}
```

When an event policy executes an event script in response to an event, you can write the script output to a file. To configure an event policy to upload the generated output file to a configured archive site, include the following statements at the **[edit event-options policy *policy-name* then]** hierarchy level:

```
[edit event-options policy policy-name then]
event-script filename {
  destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
  }
  output-filename filename;
```

```
}
```

The **transfer-delay** statement listed in each hierarchy defines the time interval that the system waits before uploading the files specified by that event policy action. If you have also configured a transfer delay for the destination at the **[edit event-options destinations destination-name]** hierarchy level, the total transfer delay is the sum of the two delays. For more detailed information about transfer delays, see [“Configuring the Delay Before Files Are Uploaded by an Event Policy” on page 602](#).

If the first upload attempt fails, **retry-count** specifies the number of additional times the system attempts to upload the file. The **retry-interval** specifies the time interval that the system waits between upload attempts. For more information, see [“Configuring an Event Policy to Retry the File Upload Action” on page 605](#).

The **output-filename** statement listed in each hierarchy is a descriptive string that is included in the filename.

When an event policy action uploads files, each uploaded file includes the hostname and timestamp in the filename to ensure that each filename is unique. If a policy is triggered multiple times in a 1-second period, an index number is appended to the filename to ensure that the filenames are still unique. The index number range is 001 through 999. The name of the file depends on the version of Junos OS running on the device. Prior to Junos OS Release 14.1R3, the filename has the following naming convention:

*hostname\_output-filename\_YYYYMMDD\_HHMMSS\_index-number*

Starting in Junos OS Release 14.1R3, the filename has the following naming convention:

*hostname\_YYYYMMDD\_HHMMSS\_output-filename\_index-number*

The **output-filename** string is either the name of an existing file or the value configured for the **output-filename** statement within the event policy **then** clause under the hierarchy for the appropriate event policy action.

For example, if you have an event policy action with output filename **rpdc-messages** on device **r1** running Junos OS Release 14.1R3 or a later release, and this event policy is executed three times in 1 second, the files are named as follows:

- **r1\_20070623\_132333\_rpd-messages**
- **r1\_20070623\_132333\_rpd-messages\_001**
- **r1\_20070623\_132333\_rpd-messages\_002**

In this example, **policy1** consists of the following statements, where **e1** is the triggering event. The example then configures the event policy to upload a log file and the committed configuration file as well as the output files generated from the **execute-commands** and **event-script** actions.

```
[edit event-options policy policy1]
events e1;
then {
  execute-commands {
    commands {
      "show interfaces brief ge-*";
    }
  }
}
```

```

    }
    event-script event-script1;
  }

```

## Configuration

- [Uploading System Files on page 598](#)
- [Uploading Command Output Files on page 599](#)
- [Uploading Event Script Output Files on page 600](#)
- [Results on page 601](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```

set event-options policy policy1 then upload filename /var/log/messages destination
  mgmt-archives transfer-delay 4
set event-options policy policy1 then upload filename /var/log/messages destination
  mgmt-archives retry-count 5 retry-interval 4
set event-options policy policy1 then upload filename /var/log/messages destination
  mgmt-archives user-name admin
set event-options policy policy1 then upload filename /var/log/messages destination
  mgmt-server
set event-options policy policy1 then upload filename committed destination
  mgmt-archives
set event-options policy policy1 then execute-commands output-filename ge-interfaces
set event-options policy policy1 then execute-commands destination mgmt-archives
  transfer-delay 5
set event-options policy policy1 then execute-commands destination mgmt-archives
  retry-count 5 retry-interval 4
set event-options policy policy1 then event-script event-script1 output-filename
  policy1-script-output
set event-options policy policy1 then event-script event-script1 destination mgmt-archives
  transfer-delay 5
set event-options policy policy1 then event-script event-script1 destination mgmt-archives
  retry-count 5 retry-interval 4

```

### Uploading System Files

#### Step-by-Step Procedure

Configure the event policy **policy1** to upload the system file **/var/log/messages** to the archive sites **mgmt-archives** and **mgmt-server**. Additionally, upload the committed configuration to the archive site **mgmt-archives**. The destination archive sites should already be configured at the **[edit event-options destinations]** hierarchy level

1. Configure the **upload** statement, and include the file to archive and the destination archive site.

```

[edit event-options policy policy1 then]
bsmith@R1# set upload filename /var/log/messages destination mgmt-archives
bsmith@R1# set upload filename /var/log/messages destination mgmt-server

```

2. To upload the committed configuration file, specify the filename value as **committed**.

```

[edit event-options policy policy1 then]

```

```
bsmith@R1# set upload filename committed destination mgmt-archives
```

3. (Optional) Configure the transfer delay associated with each file and destination.

The following configuration mode command sets the transfer delay for the **/var/log/messages** file to 4 seconds when uploaded to the mgmt-archives destination. If you have also configured a transfer delay for the destination, the total delay is the sum of the two delays.

```
[edit event-options policy policy1 then]
bsmith@R1# set upload filename /var/log/messages destination mgmt-archives
transfer-delay 4
```

4. (Optional) Configure the retry count and retry interval associated with a file and destination.

In this example, if the **/var/log/messages** file fails to upload to the mgmt-archives site, the system attempts the upload up to 5 more times and waits 4 seconds in between each attempt.

```
[edit event-options policy policy1 then]
bsmith@R1# set upload filename /var/log/messages destination mgmt-archives
retry-count 5 retry-interval 4
```

5. (Optional) Configure the username associated with a file and destination. The system uploads the file using the privileges of the specified user.

```
[edit event-options policy policy1 then]
bsmith@R1# set upload filename /var/log/messages destination mgmt-archives
user-name admin
```

6. Commit the configuration.

```
[edit event-options policy policy1 then]
bsmith@R1# commit
```

### Uploading Command Output Files

#### Step-by-Step Procedure

When the event policy invokes the **execute-commands** action, the command output can be written to a file. Configure the event policy **policy1** to write command output to a file and upload the generated file to the destination mgmt-archives, which is already configured at the **[edit event-options destinations]** hierarchy level.

1. Configure the filename of the generated output file.

```
[edit event-options policy policy1 then]
bsmith@R1# set execute-commands output-filename ge-interfaces
```

2. Configure the **destination** statement to upload the generated file to the desired archive site.

```
[edit event-options policy policy1 then]
bsmith@R1# set execute-commands destination mgmt-archives
```

3. (Optional) Configure the transfer delay for each destination.

The following command sets the transfer delay for files uploaded to the mgmt-archives destination to 5 seconds.

```
[edit event-options policy policy1 then]
bsmith@R1# set execute-commands destination mgmt-archives transfer-delay 5
```

4. (Optional) Configure the retry count and retry interval associated with each destination.

In this example, if the output file fails to upload to the mgmt-archives site, the system attempts the upload up to 5 more times and waits 4 seconds in between each attempt.

```
[edit event-options policy policy1 then]
bsmith@R1# set execute-commands destination mgmt-archives retry-count 5
retry-interval 4
```

5. Commit the configuration.

```
[edit event-options policy policy1 then]
bsmith@R1# commit
```

---

### Uploading Event Script Output Files

#### Step-by-Step Procedure

When the event policy invokes an event script, the script output can be written to a file. Configure the event policy **policy1** to write event-script output to a file and upload the generated file to the destination mgmt-archives, which is already configured at the **[edit event-options destinations]** hierarchy level. In this example, the event policy invokes an event script named **event-script1**.

1. Configure the filename of the generated output file.

```
[edit event-options policy policy1 then]
bsmith@R1# set event-script event-script1 output-filename policy1-script-output
```

2. Configure the **destination** statement to upload the generated file to the desired archive site.

```
[edit event-options policy policy1 then]
bsmith@R1# set event-script event-script1 destination mgmt-archives
```

3. (Optional) Configure the transfer delay for each destination.

The following command sets the transfer delay for files uploaded to the mgmt-archives destination to 5 seconds.

```
[edit event-options policy policy1 then]
bsmith@R1# set event-script event-script1 destination mgmt-archives
transfer-delay 5
```



4. (Optional) Configure the retry count and retry interval associated with each destination.

In this example, if the output file fails to upload to the mgmt-archives site, the system attempts the upload up to 5 more times and waits 4 seconds in between each attempt.

```
[edit event-options policy policy1 then]
bsmith@R1#set event-script event-script1 destination mgmt-archives retry-count 5
retry-interval 4
```

5. Commit the configuration.

```
[edit event-options policy policy1 then]
bsmith@R1# commit
```

## Results

```
[edit event-options policy policy1 then]
upload filename /var/log/messages destination mgmt-archives {
  user-name admin;
  transfer-delay 4;
  retry-count 5 retry-interval 4;
}
upload filename /var/log/messages destination mgmt-server
upload filename committed destination mgmt-archives;
execute-commands {
  commands {
    "show interfaces brief ge-*";
  }
  output-filename ge-interfaces;
  destination mgmt-archives {
    transfer-delay 5;
    retry-count 5 retry-interval 4;
  }
}
event-script event-script1 {
  output-filename policy1-script-output;
  destination mgmt-archives {
    transfer-delay 5;
    retry-count 5 retry-interval 4;
  }
}
```

## Verification

### Verifying the Upload

**Purpose** When the configured event triggers the event policy, the system uploads the generated output files and the specified system files to the URL defined in the mgmt-archives destination. On the destination server, verify that all files have been uploaded.

**Action** On the destination server, verify that all uploaded files are present.

```
% 1s
```

```
R1_20111209_213452_ge-interfaces
R1_20111209_213409_juniper.conf.gz
R1_20111209_212941_messages
R1_20111209_212619_policy1-script-output
```

**Meaning** Note that the filename format for each file includes the device name, the filename, and the date and time stamp.

If all of the uploaded files are present, the event policy and upload actions are working correctly. If none of the files are uploaded, verify that the destination is configured and that the archive site URL and any required password is entered correctly. For information about configuring destinations, see [“Example: Defining Destinations for File Archiving by Event Policies” on page 592](#). If a portion of the files are missing, configure a longer transfer delay and increase the retry count and retry interval for those files.

- Related Documentation**
- [Example: Defining Destinations for File Archiving by Event Policies on page 592](#)
  - [Configuring the Delay Before Files Are Uploaded by an Event Policy on page 602](#)
  - [Configuring an Event Policy to Retry the File Upload Action on page 605](#)

---

## Configuring the Delay Before Files Are Uploaded by an Event Policy

You can configure an event policy to upload existing system files or to upload the output files generated from an invoked event-script or command at the time an event occurs. A transfer delay allows you to specify the number of seconds the event process (eventd) waits before beginning to upload a file or multiple files. A transfer delay allows you to ensure that a large file, such as a core file, is completely generated before the upload begins.

As described in [“Example: Defining Destinations for File Archiving by Event Policies” on page 592](#), you can associate a transfer delay with a destination. If you associate a transfer delay with a destination, the transfer delay applies to all file upload actions that use the destination.

In the following example, the **some-dest** destination is common for both event policies, **policy1** and **policy2**. A transfer delay of 2 seconds is associated with the **some-dest** destination and applies to uploading the output files to the destination for both event policies.

```
[edit event-options]
policy policy1 {
  events e1;
  then {
    execute-commands {
      commands {
        "show version";
      }
      output-filename command-output.txt;
      destination some-dest;
    }
  }
}
policy policy2 {
```

```

events e2;
then {
    event-script bar.xml {
        output-filename event-script-output.txt;
        destination some-dest;
    }
}
destinations {
    some-dest {
        transfer-delay 2;
        archive-sites {
            "scp://robot@my.big.com/foo/moo" password "password";
            "scp://robot@my.little.com/foo/moo" password "password";
        }
    }
}

```

Suppose you have multiple event policy actions that use the same destination. For some of these event policy actions, you want a transfer delay, and for other event policy actions you want no transfer delay. To assign a transfer delay to a single event policy action, include the optional **transfer-delay** statement for each action:

**transfer-delay** *seconds*;

You can include this statement at the following hierarchy levels:

- [edit **event-options policy** *policy-name* then **event-script** *filename* **destination** *destination-name*]
- [edit **event-options policy** *policy-name* then **execute-commands** **destination** *destination-name*]
- [edit **event-options policy** *policy-name* then **upload** *filename* (*filename* | committed) **destination** *destination-name*]

If you configure a transfer delay at the [edit **event-options destinations** *destination-name*] hierarchy level, and you also configure a transfer delay for the event policy action, the resulting transfer delay is the sum of the two:

Total transfer-delay =  
transfer-delay (destination) + transfer-delay (event-policy-action)

#### Related Documentation

- [Example: Configuring the Delay Before Files Are Uploaded by an Event Policy on page 603](#)

## Example: Configuring the Delay Before Files Are Uploaded by an Event Policy

When an event policy action generates output files, you can archive the files for later analysis. Similarly, you might want to archive system files, including system log files, core files, and configuration files, from the time an event occurs. You can configure an event policy to upload relevant files to a specified location for analysis. By default, the event policy immediately uploads the files. However, you can configure a transfer delay, which causes the event process (eventd) to wait a specified number of seconds before beginning to upload one or more files. The transfer delay helps ensure that a large file, such as a

core file, is completely generated before the upload begins. The following event policies are configured with a transfer delay:

**Example 1** Configure two event policies, **policy1** and **policy2**. The **policy1** event policy has a 5-second transfer-delay when uploading the **process.core** file to the **some-dest** destination. The **policy2** event policy has no transfer delay when uploading the **process.core** file to the same destination.

```
[edit event-options]
policy policy1 {
  events e1;
  then {
    upload filename process.core destination some-dest {
      transfer-delay 5;
    }
  }
}
policy policy2 {
  events e2;
  then {
    upload filename process.core destination some-dest;
  }
}
destinations {
  some-dest {
    archive-sites {
      "scp://robot@my.little.com/foo/moo" password "password";
      "scp://robot@my.big.com/foo/moo" password "password";
    }
  }
}
```

**Example 2** The **policy1** event policy has a 7-second (5 seconds + 2 seconds) transfer delay when uploading the **process.core** file to the destination. The **policy2** event policy has a 2-second transfer delay when uploading the **process.core** file to the destination.

```
[edit event-options]
policy policy1 {
  events e1;
  then {
    upload filename process.core destination some-dest {
      transfer-delay 5;
    }
  }
}
policy policy2 {
  events e2;
  then {
    upload filename process.core destination some-dest;
  }
}
destinations {
  some-dest {
    transfer-delay 2;
    archive-sites {
      "scp://robot@my.little.com/foo/moo" password "password";
    }
  }
}
```

```

        "scp://robot@my.big.com/foo/moo" password "password";
    }
}

```

**Example 3** The **policy1** event-policy is executed with **user1** privileges and uploads the **process.core** file after a transfer delay of 7 seconds (5 seconds + 2 seconds). The **policy2** event policy is executed with **root** privileges and uploads the **process.core** file after a transfer delay of 6 seconds (4 seconds + 2 seconds).

```

[edit event-options]
policy policy1 {
  events e1;
  then {
    upload filename process.core destination some-dest {
      transfer-delay 5;
      user-name user1;
    }
  }
}
policy policy2 {
  events e2;
  then {
    upload filename process.core destination some-dest {
      transfer-delay 4;
    }
  }
}
destinations {
  some-dest {
    transfer-delay 2;
    archive-sites {
      "scp://robot@my.little.com/foo/moo" password "password";
      "scp://robot@my.big.com/foo/moo" password "password";
    }
  }
}

```

## Configuring an Event Policy to Retry the File Upload Action

You can configure an event policy to upload existing system files or to upload the output files generated from an invoked event-script or command at the time an event occurs. Transient network problems can cause a file upload operation to fail. When this happens, you might want to retry the file upload operation. By default, if the file upload operation fails for any reason, the event policy does not retry the upload operation.

To configure the policy to retry a file upload operation, include the optional **retry-count** and **retry-interval** statements:

```
retry-count number retry-interval seconds;
```

You can include these statements at the following hierarchy levels:

- **[edit event-options policy policy-name then event-script filename destination destination-name]**

- [edit **event-options policy** *policy-name* then **execute-commands** *destination destination-name*]
- [edit **event-options policy** *policy-name* then **upload** *filename (filename | committed)* *destination destination-name*]

The **retry-count** statement sets the number of times the policy retries the upload operation if the upload fails. The default value for the **retry-count** statement is 0 and the maximum is 10.

If you include the **retry-count** statement, you can also include the **retry-interval** statement, which sets the time interval (in seconds) between each retry.

For a configuration example, see “[Example: Configuring an Event Policy to Retry the File Upload Action](#)” on page 606.

---

### Example: Configuring an Event Policy to Retry the File Upload Action

---

When an event policy action generates output files, you can archive the files for later analysis. Similarly, you might want to archive system files, including system log files, core files, and configuration files, from the time an event occurs. You can configure an event policy to upload relevant files to a specified location for analysis. By default, if the file upload operation fails for any reason, the event policy does not retry the upload operation. However, you can configure an event policy to retry the file upload operation a specified number of times if the initial upload fails. You can also configure the time interval between each retry attempt. The following event policies are configured to retry the file upload operation:

- Example 1** Configure a policy that retries the file upload operation two times with a time interval of 5 seconds between retries:

```
event-options {
  policy p1 {
    events e1;
    then {
      execute-commands {
        commands {
          command1;
        }
        output-filename command-output.txt;
        destination some-dest {
          retry-count 2 retry-interval 5;
        }
      }
    }
  }
}
```

- Example 2** Configure a transfer delay of 10 seconds and retry the file upload operation two times with a time interval of 5 seconds between retries:

```
event-options {
  policy p2 {
    events e1;
```

```
    then {
      execute-commands {
        commands {
          command1;
        }
        output-filename command-output.txt;
        destination some-dest {
          retry-count 2 retry-interval 5;
          transfer-delay 10;
        }
      }
    }
  }
}
```

The transfer delay is in operation for the first upload attempt only. The policy uploads the **command-output.txt** file after a 10-second transfer delay. If the event process (eventd) detects failure of the upload operation, eventd retries the upload operation after 5 seconds. The failure detection time can be in the range from 60 to 90 seconds, depending on the transmission protocol, such as FTP.

The following sequence describes the file upload operation with two failed retransmissions:

1. Policy triggers upload operation.
2. Transmission delay of 10 seconds.
3. Policy tries to upload the output file.
4. Policy detects transmission failure.
5. Retry interval of 5 seconds.
6. Policy tries to upload the output file.
7. Policy detects transmission failure.
8. Retry interval of 5 seconds.
9. Policy tries to upload the output file.
10. Policy detects transmission failure.
11. Policy declares the failure of the file upload operation.





# Configuring Event Policy Privileges

- [Changing the User Privilege Level for an Event Policy Action on page 609](#)
- [Example: Associating an Optional User with an Event Policy Action on page 610](#)

## Changing the User Privilege Level for an Event Policy Action

---

Only superusers can configure event policies. By default, event policy actions—such as executing operational mode commands, uploading files, and executing SLAX and XSLT event scripts—are executed by user **root**, because the event process (eventd) runs with root privileges.



**NOTE:** To prevent the execution of unauthorized Python code on devices running Junos OS, by default, Junos OS executes Python event scripts using the access privileges of the generic, unprivileged user and group **nobody**.

In some cases, you might want an event policy action to be executed with restricted privileges. For example, suppose you configure an event policy that executes a script if an interface goes down. The script includes remote procedure calls (RPCs) to change the device configuration if certain conditions are present. If you do not want the script to change the configuration, you can execute the script with a restricted user profile. When the script is executed with a user profile that disallows configuration changes, the RPCs to change the configuration fail.

You can associate a user with each action in an event policy. If a user is not associated with an event policy action, then the action is executed as user **root** by default.

To specify the user under whose privileges an action is executed, configure the **user-name** statement.

**user-name** *username*;

You can include this statement at the following hierarchy levels:

- `[edit event-options policy policy-name then event-script filename]`



**NOTE:** The `user-name` statement only applies to SLAX and XSLT event scripts. This statement has no effect when configured for Python event scripts.

- [edit event-options policy *policy-name* then execute-commands]
- [edit event-options policy *policy-name* then upload filename (*filename* | committed) destination *destination-name*]

By default, Junos OS executes Python event scripts with the access privileges of the generic, unprivileged user and group **nobody**. Starting in Junos OS Release 16.1R3, you can specify the user under whose access privileges a Python event script will execute. To execute a Python event script under the access privileges of a specific user, configure the `python-script-user username` statement at the [edit event-options event-script file *filename*] hierarchy level.

```
[edit event-options event-script file filename]
user@host# set python-script-user username
```



**NOTE:** To enable a user who does not belong to the file's user or group class to execute an unsigned Python automation script, the script's file permissions must include read permission for others.



**NOTE:** The username that you specify for the `user-name` and `python-script-user` statements must be configured at the [edit system login] hierarchy level. For more information, see the *Junos OS System Basics Configuration Guide*.

Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, you can specify the user under whose access privileges a Python event script will execute.

#### Related Documentation

- [Example: Associating an Optional User with an Event Policy Action on page 610](#)

## Example: Associating an Optional User with an Event Policy Action

Configure two event policies, **policy1** and **policy2**.

In **policy1**, associate user **user1** with the `execute-commands` action. The `execute-commands` action is executed with **user1** privileges.

In **policy2**, do not explicitly associate a user with the **event-script** action. The **event-script** action is executed with root privileges.

```
[edit system]
login {
  user user1 {
    class operator;
  }
}
[edit event-options]
policy p1 {
  events e1;
  then {
    execute-commands {
      commands {
        "show version";
      }
      user-name user1;
      output-filename command-output.txt;
      destination some-dest;
    }
  }
}
policy p2 {
  events e2;
  then {
    event-script script.xml {
      output-filename event-script-output.txt;
      destination some-dest;
    }
  }
}
```

#### Related Documentation

- [Changing the User Privilege Level for an Event Policy Action on page 609](#)



## CHAPTER 33

# Event Scripts Overview

- [Event Scripts Overview on page 613](#)

### Event Scripts Overview

---

Junos OS event scripts are triggered automatically by defined event policies in response to a system event and can instruct Junos OS to take immediate action. Event scripts automate network and device management and troubleshooting. Event scripts can perform functions available through the remote procedure calls (RPCs) supported by either Junos XML management protocol or the Junos Extensible Markup Language (XML) API. Event scripts are executed by the event process (eventd).

Event scripts enable you to:

- Automatically diagnose and fix problems in the network
- Monitor the overall status of a device.
- Run automatically as part of an event policy that detects periodic error conditions
- Change the configuration in response to a problem

Event scripts are based on the Junos XML management protocol and the Junos XML API, which are discussed in [“Junos XML Management Protocol and Junos XML API Overview” on page 11](#). Event scripts can be written in Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX) scripting language. Event scripts use XML Path Language (XPath) to locate the operational objects to be inspected and automation script constructs to specify the actions to perform on the operational objects. The actions can change the output or execute additional commands based on the output.

Event scripts are invoked within an event policy. To use event scripts in an event policy, event scripts must be enabled at the **[edit event-options event-script file]** hierarchy level and must be configured under the **then** hierarchy for the event policy. When an event policy is triggered, the policy forwards the event details to the event scripts and executes the scripts in the order in which they are configured. These scripts contain instructions that execute operational mode commands, inspect the output automatically, and perform any necessary actions. For information about event policies, see [“Event Notifications and Policies Overview” on page 537](#) and [“Executing Event Scripts in an Event Policy” on page 573](#).

You can use event scripts to generate changes to the device configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

**Related  
Documentation**

- [Junos XML Management Protocol and Junos XML API Overview on page 11](#)
- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
- [SLAX Overview on page 57](#)
- [XSLT Overview on page 17](#)
- [XPath Overview on page 19](#)
- [Executing Event Scripts in an Event Policy on page 573](#)

# Creating and Executing Event Scripts

- [Required Boilerplate for Event Scripts on page 615](#)
- [Using Event and Remote Execution Details in Event Scripts on page 617](#)
- [Using RPCs and Operational Mode Commands in Event Scripts on page 619](#)
- [Enabling and Executing Event Scripts on page 626](#)
- [Replacing an Event Script on page 628](#)
- [Configuring Checksum Hashes for an Event Script on page 628](#)

## Required Boilerplate for Event Scripts

---

Junos OS event scripts can be written in Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python. Event scripts must include the necessary boilerplate required for that script language for both basic script functionality as well as any optional functionality used within the script such as the Junos OS extension functions and named templates. This topic provides standard boilerplate that can be used in XSLT, SLAX, and Python scripts.

SLAX and XSLT event scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

### XSLT Boilerplate for Event Scripts

The XSLT event script boilerplate is as follows:

```

1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3      xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4      xmlns:junos="http://xml.juniper.net/junos/*/junos"
5      xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6      xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7      <xsl:import href="../import/junos.xsl"/>

8      <xsl:template match="configuration">
9          <event-script-results>
10             <!-- ... Insert your code here ... -->
11             </event-script-results>
12          </xsl:template>
13          <!-- ... insert additional template definitions here ... -->

```

```
12    </xsl:stylesheet>
```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI). This PI specifies that the code is written in XML using version 1.0. The XML PI, if present, must be the first non-comment token in the script file.

```
1    <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2    <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in event scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```
3        xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4        xmlns:junos="http://xml.juniper.net/junos/*/junos"
5        xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6        xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of Junos OS (in the file `/usr/libdata/cscript/import/junos.xsl`). The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in [“Understanding Named Templates in Junos OS Automation Scripts” on page 225](#).

```
7        <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `</>` element. The `<xsl:template match="/">` element is the root element and represents the top level of the XML hierarchy. All XPath expressions in the script must start at the top level. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see [“XPath Overview” on page 19](#).

```
8        <xsl:template match="/">
```

After the `<xsl:template match="/">` tag element, the `<event-script-results>` and `</event-script-results>` container tags must be the top-level child tags, as shown in Lines 9 and 10.

```
9            <event-script-results>
10                <!-- ... insert your code here ... -->
11            </event-script-results>
```

Line 11 closes the template.

```
11        </xsl:template>
```

Between Line 11 and Line 12, you can define additional XSLT templates that are called from within the `<xsl:template match="/">` template.

Line 12 closes the style sheet and the event script.

```
12    </xsl:stylesheet>
```

#### SLAX Boilerplate for Event Scripts

The corresponding SLAX op script boilerplate is as follows:

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
```



```

ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";

match / {
    <event-script-results> {
        /*
         * Insert your code here
         */
    }
}

```

### Python Boilerplate for Event Scripts

Python event scripts do not have a required boilerplate, but they must import any objects that are used in the script. Python automation scripts import the **Junos\_Context** dictionary when the script must access information about the script execution environment. The **Junos\_Trigger\_Event** and **Junos\_Received\_Events** objects are required when the event script must access details about the events that triggered the corresponding event policy. The **Junos\_Remote\_Execution\_Details** generator function is required to access remote execution details configured for an event script at the **[edit event-options event-script file filename remote-execution]** hierarchy level. Python automation scripts that include the **import jcs** statement can use Junos OS extension functions and Junos OS named template functionality in the script. And if the script uses Junos PyEZ functionality, it must import the necessary classes from the **jnpr.junos** module.

```

from junos import Junos_Context
from junos import Junos_Trigger_Event
from junos import Junos_Received_Events
from junos import Junos_Remote_Execution_Details
from jnpr.junos import Device
import jcs

if __name__ == '__main__':

```

Python automation scripts do not need to include an interpreter directive line (**#!/usr/bin/env python**) at the start of the script. However, the program will still execute correctly if one is present.

### Related Documentation

- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)
- [Using Event and Remote Execution Details in Event Scripts on page 617](#)

## Using Event and Remote Execution Details in Event Scripts

Event policy actions can include executing one or more event scripts. When an event policy executes an event script, the event process forwards the event details to the script. These event details can be captured, evaluated, and sent to log files as required. In addition, any configured remote execution details are also forwarded to the event script.

Two types of event details are returned: triggered events and received events. *Triggered events* record the details of the event that triggered the policy. *Received events* record the details of events that happened before the triggering event. Trigger event details are always forwarded to event scripts. Received event details are only present when an event policy is triggered for correlated events.

Remote execution details, which include the hostname, username, and passphrase for one or more remote hosts, enable an event script to invoke remote procedure calls on remote hosts without encoding the connection information directly in the event script. You configure remote execution details at the **[edit event-options event-script file filename remote-execution]** hierarchy level. When you include remote execution details in the configuration instead of in the individual event scripts, the information is captured in a single location and the passphrase is encrypted. This is not the case in an event script file.

Event details and remote execution details are forwarded to SLAX and XSLT event scripts as XML in the following format:

```
<event-script-input>
  <junos-context>
    ...
  </junos-context>
  <trigger-event>
    <id>event-id</id>
    <type>event-type</type>
    <generation-time>timestamp</generation-time>
    <process>
      <name>process-name</name>
      <pid>pid</pid>
    </process>
    <hostname>hostname</hostname>
    <message>message-string</message>
    <facility>facility-string</facility>
    <severity>severity-string</severity>
    <attribute-list>
      <attribute>
        <name>attribute-name</name>
        <value>attribute-value</value>
      </attribute>
    </attribute-list>
  </trigger-event>
  <received-events>
    <received-event>
      <id>event-id</id>
      <type>event-type</type>
      <generation-time>timestamp</generation-time>
      <process>
        <name>process-name</name>
        <pid>pid</pid>
      </process>
      <hostname>hostname</hostname>
      <facility>facility-string</facility>
      <severity>severity-string</severity>
      <attribute-list>
        <attribute>
          <name>attribute-name</name>
          <value>attribute-value</value>
        </attribute>
      </attribute-list>
    </received-event>
  </received-events>
  <remote-execution-details>
    <remote-execution-detail>
      <remote-hostname>hostname</remote-hostname>
      <username>username</username>
```

```

        <passphrase>passphrase</passphrase>
    </remote-execution-detail>
</remote-execution-details>
</event-script-input>

```

For information about the `<junos-context>` element, see [“Global Parameters and Variables in Junos OS Automation Scripts” on page 175](#).

Python event scripts must import the `Junos_Trigger_Event` and `Junos_Received_Events` objects to access details about the trigger event and received events. `Junos_Trigger_Event` and `Junos_Received_Events` are `lxml.etree.Element` objects and contain the same hierarchy and tag names as the `<trigger-event>` and `<received-events>` hierarchies in the SLAX and XSLT script input.

Python event scripts can extract the necessary event details from the objects using `lxml` methods such as `xpath()` and `find()`, `findall()`, and `findtext()`. For example:

```
Junos_Trigger_Event.xpath('//trigger-event/process/name')[0].text
```

Python event scripts must import `Junos_Remote_Execution_Details` to access the remote execution details configured at the `[edit event-options event-script file filename remote-execution]` hierarchy level. `Junos_Remote_Execution_Details` is a generator function that produces a sequence of remote devices, which makes it easy to iterate over multiple configured hosts. You can reference the hostname, username, and passphrase for a configured remote host by using the `host`, `user`, and `passwd` properties as in the following code:

```

from junos import Junos_Remote_Execution_Details

for remote in Junos_Remote_Execution_Details():
    hostname = remote.host
    username = remote.user
    passphrase = remote.passwd

```

#### Related Documentation

- [Using RPCs and Operational Mode Commands in Event Scripts on page 619](#)
- [Global Parameters and Variables in Junos OS Automation Scripts on page 175](#)
- [Example: Limiting Event Script Output Based on a Specific Event Type on page 588](#)

## Using RPCs and Operational Mode Commands in Event Scripts

Most Junos OS operational mode commands have XML equivalents. Event scripts can execute these XML commands on a local or remote device using the *remote procedure call* (RPC) protocol. All operational mode commands that have XML equivalents are listed in the *Junos XML API Operational Developer Reference*.

SLAX and XSLT scripts execute RPCs on a local or remote device by using the `jcs:invoke()` or `jcs:execute()` extension functions, respectively. In Python scripts, RPCs are easy to execute using [Junos PyEZ](#) APIs. Each instance of the `Junos PyEZ Device` class has an `rpc` property that enables you to execute any RPC available through the Junos XML API. After establishing a session with a local or remote device, you can execute the RPC by appending the `rpc` property and RPC method name to the device instance. The return value is an XML object starting at the first element under the `<rpc-reply>` tag.

Use of RPCs and operational mode commands in event scripts is discussed in more detail in the following sections:

- [Executing RPCs on a Local Device on page 620](#)
- [Executing RPCs on a Remote Device on page 620](#)
- [Displaying the RPC Tags for a Command on page 623](#)
- [Using Operational Mode Commands in Event Scripts on page 624](#)

## Executing RPCs on a Local Device

In a SLAX or XSLT event script, to execute an RPC on the local device, include the RPC in a variable declaration, and call the **jcs:invoke()** extension function with the RPC variable as an argument. The following snippet invokes an RPC on the local device:

**XSLT Syntax**

```
<xsl:variable name="rpc">
  <get-interface-information/> # Junos RPC for the show interfaces command
</xsl:variable>
<xsl:variable name="out" select="jcs:invoke($rpc)"/>
...
```

**SLAX Syntax**

```
var $rpc = <get-interface-information>;
var $out = jcs:invoke($rpc);
```

In a Python event script, to execute an RPC on the local device, create the **Device** instance using an empty argument list, and append the **rpc** property and the RPC method name and argument list to the device instance.

**Python Syntax**

```
jdev = Device()
jdev.open()
rsp = jdev.rpc.get_interface_information()
jdev.close()
```



**NOTE:** When you create the **Device** instance using an empty argument list to connect to the local device, Junos OS uses the access privileges of the user configured at the [edit event-options event-script file *filename* python-script-user] hierarchy level. If the **python-script-user** statement is omitted, Junos OS uses the access privileges of the generic, unprivileged user and group **nobody**.

---

## Executing RPCs on a Remote Device

In a SLAX or XSLT event script, to execute an RPC on a remote device, first include the RPC in a variable declaration, and create a connection handle using the **jcs:open()** extension function with the arguments required to connect to the remote device. Then call the **jcs:execute()** extension function and include the connection handle and RPC variable as arguments. For example:

**XSLT Syntax**

```
<xsl:variable name="rpc">
  <get-interface-information/>
</xsl:variable>
<xsl:variable name="connection"
  select="jcs:open('198.51.100.1', 'bsmith', 'test123')"/>
```

```
<xsl:variable name="out" select="jcs:execute($connection, $rpc)"/>
<xsl:value-of select="jcs:close($connection)"/>
...
```

**SLAX Syntax**

```
var $rpc = <get-interface-information>;
var $connection = jcs:open("198.51.100.1", "bsmith", "test123");
var $out = jcs:execute($connection, $rpc);
expr jcs:close($connection);
...
```

In a Python event script, to execute an RPC on a remote device, first create an instance of **Device** using the arguments required to connect to the remote device. Then execute the RPC by appending the **rpc** property and the RPC method name and argument list to the device instance.

**Python Syntax**

```
jdev = Device(host='198.51.100.1', user='bsmith', passwd='test123')
jdev.open()
rsp = jdev.rpc.get_interface_information()
jdev.close()
```



**NOTE:** Junos OS connects to and executes operations on the remote device using the access privileges of the user specified in the `Device()` argument list, even if a different user is configured for the `python-script-user` statement at the `[edit event-options event-script file filename]` hierarchy level.

To avoid adding the remote connection details directly into an event script, you can specify remote execution details for each event script that executes RPCs on a remote device at the `[edit event-options event-script file filename remote-execution]` hierarchy level. We recommend adding the remote execution details to the configuration instead of directly in the event script, because all of the information is available in a single location, and the passphrase is encrypted in the configuration.

For each remote device where an RPC is executed, configure the device hostname and the corresponding username and passphrase.

```
[edit event-options event-script file filename]
remote-execution {
  remote-hostname {
    username username;
    passphrase passphrase;
  }
}
```

The remote hostnames and their corresponding username and passphrase, in addition to the event details, are passed as input to the event script when it is executed by an event policy. For more information about the details that are forwarded to the event script, see [“Using Event and Remote Execution Details in Event Scripts” on page 617](#).

An event script references the remote execution details in the argument list of the function used to create the connection to the remote host. Once the connection has been established, the script can execute RPCs on that device.

In Python event scripts, you reference the remote execution details in the argument list of the Junos PyEZ **Device()** instance. The following code iterates over the remote execution

details for all hosts configured for that event script and connects to and executes the same RPC on each host.

**Python Syntax**

```
from junos import Junos_Remote_Execution_Details
from jnpr.junos import Device

def main()
    for remote in Junos_Remote_Execution_Details():
        hostname = remote.host
        username = remote.user
        passphrase = remote.passwd

        jdev = Device(host=hostname, user=username, passwd=passphrase)
        jdev.open()

        inv = jdev.rpc.get_interface_information()
        #process RPC information...
        jdev.close()

if __name__ == "__main__":    main()
```

In SLAX or XSLT scripts, create a connection to the remote host by using the **jcs:open()** function and reference the remote execution details in the argument list. For example:

**XSLT Syntax**

```
<xsl:variable name="rpc">
    <get-interface-information/>
</xsl:variable>
<xsl:for-each select="event-script-input/remote-execution-details">
    <xsl:variable name="d" select="remote-execution-detail"/>
    <xsl:variable name="connection"
        select="jcs:open($d/remote-hostname,$d/username,$d/passphrase)"/>
    <xsl:variable name="out" select="jcs:execute($connection, $rpc)"/>
    <xsl:value-of select="jcs:close($connection)"/>
    ...
</xsl:for-each>
```

**SLAX Syntax**

```
var $rpc = <get-interface-information>;
for-each (event-script-input/remote-execution-details) {
    var $d = remote-execution-detail;
    var $connection = jcs:open($d/remote-hostname,$d/username,$d/passphrase);
    var $out = jcs:execute($connection, $rpc);
    expr jcs:close($connection);
    ...
}
```

To execute an RPC on a remote device, an SSH session must be established. In order for the script to establish the connection, you must either configure the SSH host key information for the remote device on the local device where the script will be executed, or the SSH host key information for the remote device must exist in the known hosts file of the user executing the script. For each remote device where the RPC is executed, configure the SSH host key information using one of the following methods:

- To configure SSH known hosts on the local device, include the **host** statement, and specify hostname and host key options for the remote device at the **[edit security ssh-known-hosts]** hierarchy level of the configuration.
- To manually retrieve SSH host key information, issue the **set security ssh-known-hosts fetch-from-server hostname** configuration mode command to instruct Junos OS to connect to the remote device and add the key.

```
user@host# set security ssh-known-hosts fetch-from-server router2
The authenticity of host 'router2 (198.51.100.1)' can't be established.
RSA key fingerprint is 30:18:99:7a:3c:ed:40:04:0f:fd:c1:57:7e:6b:f3:90.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'router2,198.51.100.1' (RSA) to the list of known
hosts.
```

- To manually import SSH host key information from a file, use the **set security ssh-known-hosts load-key-file filename** configuration mode command and specify the known-hosts file.

```
user@host# set security ssh-known-hosts load-key-file /var/tmp/known_hosts
Import SSH host keys from trusted source /var/tmp/known_hosts ? [yes,no] (no)
yes
```

- Alternatively, the user executing the script can log in to the local device, SSH to the remote device, and then manually accept the host key, which is added to that user's known hosts file. In the following example, root is logged in to router1. In order to execute a remote RPC on router2, root adds the host key of router2 by issuing the **ssh router2** operational mode command and manually accepting the key.

```
root@router1> ssh router2
The authenticity of host 'router2 (198.51.100.1)' can't be established.
RSA key fingerprint is 30:18:99:7a:3c:ed:40:04:0f:fd:c1:57:7e:6b:f3:90.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'router2,198.51.100.1' (RSA) to the list of known
hosts.
```

After configuring the required SSH host key and obtaining a connection handle to the remote device, the event script can execute RPCs with the **jcs:execute()** extension function on that remote device.

## Displaying the RPC Tags for a Command

You can display the RPC XML tags for operational mode commands in the CLI of the device. To display the RPC XML tags for a command, enter **display xml rpc** after the pipe symbol (|).

The following example displays the RPC tags for the **show route** command:

```
user@host> show route | display xml rpc
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.1I0/junos">
  <rpc>
    <get-route-information>
    </get-route-information>
  </rpc>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```

SLAX and XSLT scripts can execute RPCs using the RPC XML tags. Python scripts must convert the RPC tags and command options into a format suitable for Python. For more information about using Junos PyEZ to execute RPCs and about mapping RPC tags to the corresponding Python method and method arguments, see [Using Junos PyEZ to Execute RPCs on Devices Running Junos OS](#).

## Using Operational Mode Commands in Event Scripts

Some operational mode commands do not have XML equivalents. SLAX and XSLT scripts can execute commands that have no XML equivalent by using the `<command>` element. Python scripts can execute these commands by using the Junos PyEZ `cli()` method defined in the `Device` class.

If a command is not listed in the *Junos XML API Operational Developer Reference*, the command does not have an XML equivalent. Another way to determine whether a command has an XML equivalent is to issue the command followed by the `| display xml` command.

```
user@host> operational-mode-command | display xml
```

If the output includes only tag elements like `<output>`, `<cli>`, and `<banner>`, the command might not have an XML equivalent. In the following example, the output indicates that the `show host` command has no XML equivalent:

```
user@host> show host hostname | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.0R1/junos">
  <output>
    ...
  </output>
  <cli>
    <banner></banner>
  </cli>
</rpc-reply>
```





**NOTE:** For some commands that have an XML equivalent, the output of the piped `| display xml` command does not include tag elements other than `<output>`, `<cli>`, and `<banner>` only because the relevant feature is not configured. For example, the `show services cos statistics forwarding-class` command has an XML equivalent that returns output in the `<service-cos-forwarding-class-statistics>` response tag, but if the configuration does not include any statements at the `[edit class-of-service]` hierarchy level, then there is no actual data for the `show services cos statistics forwarding-class | display xml` command to display. The output is similar to this:

```
user@host> show services cos statistics forwarding-class | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/8.3I0/junos">

  <cli>

    <banner></banner>

  </cli>
</rpc-reply>
```

For this reason, the information in the *Junos XML API Operational Developer Reference* is usually more reliable.

SLAX and XSLT event scripts can execute commands that have no XML equivalent. Use the `<command>`, `<xsl:value-of>`, and `<output>` elements in the script, as shown in the following code snippet. This snippet is expanded and fully described in ["Example: Displaying DNS Hostname Information Using an Op Script" on page 487](#).

```
<xsl:variable name="query">
  <command>
    <xsl:value-of select="concat('show host ', $hostname)"/>
  </command>
</xsl:variable>
<xsl:variable name="result" select="jcs:invoke($query)"/>
<xsl:variable name="host" select="$result"/>
<output>
  <xsl:value-of select="concat('Name: ', $host)"/>
</output>
...
```

Python event scripts can execute commands that have no XML equivalent by using Junos PyEZ APIs. The `cli()` method defined in the `Device` class executes an operational mode command and returns the output in text format. For example:

```
from jnpr.junos import Device

def main():
    jdev = Device()
    jdev.open()
    res = jdev.cli('show host hostname', warning=False)
    print res
    jdev.close()

if __name__ == "__main__":
    main()
```

You can also specify **format='xml'** to return the output formatted as Junos OS XML elements. For more information about the Junos PyEZ **cli** method, see [http://junos-pyez.readthedocs.org/en/latest/\\_modules/jnpr/junos/device.html#Device.cli](http://junos-pyez.readthedocs.org/en/latest/_modules/jnpr/junos/device.html#Device.cli).

**Related  
Documentation**

- [Using Event and Remote Execution Details in Event Scripts on page 617](#)
- [open\(\) Function \(SLAX and XSLT\) on page 209](#)
- [execute\(\) Function \(SLAX and XSLT\) on page 198](#)
- [Junos PyEZ](#)

---

## Enabling and Executing Event Scripts

---

Event scripts are stored on a device's hard disk in the **/var/db/scripts/event** directory or on the flash drive in the **/config/scripts/event** directory. Only users in the Junos OS **super-user** login class can access and edit files in these directories. For information about setting the storage location for scripts, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

To prevent the execution of unauthorized Python code on devices running Junos OS, unsigned Python scripts must meet certain requirements before you can execute the scripts on a device. Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user. For detailed information about the requirements for executing Python automation scripts on devices running Junos OS, see [“Requirements for Executing Python Automation Scripts on Devices Running Junos OS” on page 169](#).



**NOTE:** If the device has dual Routing Engines and you want to enable an event script to execute on both Routing Engines, you must copy the script to the **/var/db/scripts/event** or **/config/scripts/event** directory on both Routing Engines. The **commit synchronize** command does not automatically copy scripts between Routing Engines.

You must enable an event script before it can be executed. To enable an event script, include the **file filename** statement at the **[edit event-options events-script]** hierarchy level, specifying the name of the file containing the event script. Only users who belong to the Junos **super-user** login class can enable event scripts.

```
[edit event-options event-script]
user@host# set file filename
```

SLAX and Python scripts must include the **.slax** or **.py** filename extension, respectively, in both the actual script name and the filename in the configuration. XSLT scripts do not require a filename extension, but we strongly recommend that you append the **.xsl** extension. Whether or not you choose to include the **.xsl** extension on the file, the filename that you add at the **[edit event-options event-script file]** hierarchy level must exactly match the filename of the script in the directory. For example, if the XSLT script filename

is **script1.xsl**, then you must include **script1.xsl** in the configuration hierarchy to enable the script; likewise, if the XSLT script filename is **script1**, then you must include **script1** in the configuration hierarchy.

To determine which event scripts are currently enabled on the device, use the **show** command to display the files configured at the **[edit event-options event-script]** hierarchy level. To ensure that the enabled files are on the device, list the contents of the **/var/run/scripts/event/** directory using the **file list /var/run/scripts/event** operational mode command.

By default, you cannot execute unsigned Python scripts on devices running Junos OS. To enable the execution of unsigned Python automation scripts that meet the requirements outlined in “[Requirements for Executing Python Automation Scripts on Devices Running Junos OS](#)” on page 169, you must configure the **language python** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set language python
```

By default, Junos OS executes Python event scripts with the access privileges of the generic, unprivileged user and group **nobody**. Starting in Junos OS Release 16.1R3, you can specify the user under whose access privileges the Python script will execute. To execute a Python event script under the access privileges of a specific user, configure the **python-script-user username** statement at the **[edit event-options event-script file filename]** hierarchy level.

```
[edit event-options event-script file filename]
user@host# set python-script-user username
```



**NOTE:** To enable a user who does not belong to the file's user or group class to execute an unsigned Python automation script, the script's file permissions must include read permission for others.

When you issue the **commit** command, event scripts configured at the **[edit event-options event-script]** hierarchy level are placed into system memory and enabled for execution. After the commit operation completes, an event policy can execute an event script in response to an event notification.

Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user.

**Related Documentation**

- [Event Scripts Overview on page 613](#)
- [Executing Event Scripts in an Event Policy on page 573](#)

## Replacing an Event Script

---

You can update or replace an existing event script without changing the device's configuration or disrupting operations. Follow these steps:

1. Edit the existing event script or write a new one.
2. Copy the script to the `/var/db/scripts/event` directory on the hard disk or the `/config/scripts/event` directory on the flash drive; for information about setting the storage location for scripts, see [“Storing Scripts in Flash Memory” on page 664](#). Only users who belong to the Junos **super-user** login class can alter files in these directories.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.



**NOTE:** If the device has dual Routing Engines, remember to copy the script to the `/var/db/scripts/event` or `/config/scripts/event` directory on both Routing Engines. The `commit synchronize` command does not automatically copy scripts between Routing Engines.

3. Issue the `request system scripts event-scripts reload` operational mode command.

```
user@host> request system scripts event-scripts reload
```

All event scripts are reloaded into the eventd process' memory.

### Related Documentation

- [Storing Scripts in Flash Memory on page 664](#)

## Configuring Checksum Hashes for an Event Script

---

You can configure one or more checksum hashes that can be used to verify the integrity of an event script before the script runs on the switch, router, or security device.

To configure a checksum hash:

1. Create the script.
2. Place the script in the `/var/db/scripts/event` directory on the device.
3. Run the script through one or more hash functions to calculate hash values.

Junos OS supports MD5, SHA-1, and SHA-256 hash functions.

```
user@host> file checksum md5 /var/db/scripts/event/script1.slax
MD5 (/var/db/scripts/event/script1.slax) = 3af7884eb56e2d4489c2e49b26a39a97

user@host> file checksum sha1 /var/db/scripts/event/script1.slax
SHA1 (/var/db/scripts/event/script1.slax) =
00dc690fb08fb049577d012486c9a6dad34212c0
```

```
user@host> file checksum sha-256 /var/db/scripts/event/script1.slax
SHA256 (/var/db/scripts/event/script1.slax) =
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

4. Configure the script.

```
[edit event-options event-script]
user@host# set file script1.slax checksum
md5 3af7884eb56e2d4489c2e49b26a39a97

[edit event-options event-script]
user@host# set file script1.slax checksum
sha-1 00dc690fb08fb049577d012486c9a6dad34212c0

[edit event-options event-script]
user@host# set file script1.slax checksum
sha-256 150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

During the execution of the script, Junos OS recalculates the checksum value using the configured hash and verifies that the calculated value matches the configured value. If the values differ, the execution of the script fails. When you configure multiple checksum values with different hash algorithms, all the configured values must match the calculated values; otherwise, the script execution fails and the event policy fails.

- Related Documentation**
- [Configuring Checksum Hashes for a Commit Script on page 264](#)
  - [Configuring Checksum Hashes for an Op Script on page 460](#)
  - [Configuring Checksum Hashes for an SNMP Script on page 647](#)



# Troubleshooting Event Policies and Event Scripts

- [Tracing Event Policy Processing on page 631](#)
- [Tracing Event Script Processing on page 634](#)

## Tracing Event Policy Processing

---

Event policy tracing operations track all event policy operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, no events are traced. If you include the **traceoptions** statement at the **[edit event-options]** hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **eventd** located in the **/var/log** directory.
- When the file **eventd** reaches 128 kilobytes (KB), it is renamed and compressed to **eventd.0.gz**, then **eventd.1.gz**, and so on, until there are three trace files. Then the oldest trace file (**eventd.2.gz**) is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)
- Log files can be accessed only by the user who configures the tracing operation.

You cannot change the directory (**/var/log**) to which trace files are written. However, you can customize the other trace file settings by including the following statements at the **[edit event-options traceoptions]** hierarchy level:

```
[edit event-options traceoptions]
file <filename> <files number> <match regular-expression> <size size> <world-readable |
  no-world-readable>;
flag all;
flag configuration;
flag database;
flag events;
flag policy;
flag server;
flag syslog;
flag timer-events;
no-remote-trace;
```

These statements are described in the following sections:

- [Configuring the Event Policy Log Filename on page 632](#)
- [Configuring the Number and Size of Event Policy Log Files on page 632](#)
- [Configuring Access to the Log File on page 632](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 633](#)
- [Configuring the Trace Operations on page 633](#)

## Configuring the Event Policy Log Filename

By default, the name of the file that records trace output is **eventd**. You can specify a different name by including the **file** statement at the **[edit event-options traceoptions]** hierarchy level:

```
[edit event-options traceoptions]  
file filename;
```

## Configuring the Number and Size of Event Policy Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed **filename.0**, then **filename.1**, and so on, until there are three trace files. Then the oldest trace file (**filename.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit event-options traceoptions file <filename>]** hierarchy level:

```
[edit event-options traceoptions file <filename>]  
files number size size;
```

For example, set the maximum file size to 2 MB and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 2 MB, **filename** is renamed and compressed to **filename.0.gz** and a new file called **filename** is created.

When **filename** reaches 2 MB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is renamed and compressed to **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19.gz**) is overwritten.

The number of files can range from 2 through 1000 files. The file size can range from 10 KB through 1 gigabyte (GB).

## Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the **world-readable** statement at the **[edit event-options traceoptions file <filename>]** hierarchy level:

```
[edit event-options traceoptions file <filename>]  
world-readable;
```

To explicitly set the default behavior, include the **no-world-readable** statement at the **[edit event-options traceoptions file <filename>]** hierarchy level:

```
[edit event-options traceoptions file <filename>]
```



```
no-world-readable;
```

## Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** statement at the **[edit event-options traceoptions file <filename>]** hierarchy level and specifying a regular expression to be matched:

```
[edit event-options traceoptions file <filename>]
match regular-expression;
```

## Configuring the Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit event-options traceoptions]** hierarchy level:

```
[edit event-options traceoptions]
flag all;
flag configuration;
flag database;
flag events;
flag policy;
flag server;
flag syslog;
flag timer-events;
```

Table 46 on page 633 describes the meaning of the event policy tracing flags.

**Table 46: Event Policy Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>configuration</b>	Log reading of configuration at the <b>[edit event-options]</b> hierarchy level.	Off
<b>events</b>	Trace important events.	Off
<b>database</b>	Log events involving storage and retrieval in events database.	Off
<b>policy</b>	Log policy processing.	Off
<b>server</b>	Log communication with processes that are generating events.	Off
<b>syslogd</b>	Log syslog related traces	Off
<b>timer-events</b>	Log internally generated events.	Off

To display the end of the log, issue the **show log eventd | last** operational mode command:

```
[edit]
user@host# run show log eventd | last
```

**Related  
Documentation**

- [Understanding the Event System Log Priority in an Event Policy on page 583](#)
- [Tracing Event Script Processing on page 634](#)
- [traceoptions on page 800](#)

---

## Tracing Event Script Processing

Event script tracing operations track all event script operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

The default operation of event script tracing is to log important events in a file called **escript.log** located in the **/var/log** directory. When the file **escript.log** reaches 128 kilobytes (KB), it is renamed with a number 0 through 9 (in ascending order) appended to the end of the file and then compressed. The resulting files are **escript.log.0.gz**, then **escript.log.1.gz**, until there are 10 trace files. Then the oldest trace file (**escript.log.9.gz**) is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)

This section discusses the following topics:

- [Minimum Configuration for Enabling Traceoptions for Event Scripts on page 634](#)
- [Configuring Tracing of Event Scripts on page 635](#)

### Minimum Configuration for Enabling Traceoptions for Event Scripts

If no event script trace options are configured, the simplest way to view the trace output of an event script is to configure the **output** trace flag and issue the **show log escript.log | last** command. To do this, perform the following steps:

1. If you have not done so already, enable an event script by including the **file** statement at the **[edit event-options event-script]** hierarchy level:

```
[edit event-options event-script]
user@host# set file filename
```

2. Enable trace options by including the **traceoptions flag output** statement at the **[edit event-options event-script]** hierarchy level:

```
[edit event-options event-script]
user@host# set traceoptions flag output
```

3. Issue the **commit** command:

```
[edit]
user@host# commit
```

4. Display the resulting trace messages recorded in the `/var/log/escrpt.log` file. At the end of the log is the output generated by the event script you enabled in Step 1 after a configured event policy is triggered and invokes the script. To display the end of the log, issue the **show log escrpt.log | last** operational mode command:

```
[edit]
user@host# run show log escrpt.log | last
```

Table 47 on page 635 summarizes useful filtering commands that display selected portions of the `escrpt.log` file.

**Table 47: Event Script Tracing Operational Mode Commands**

Task	Command
Display logging data associated with all event script processing.	<b>show log escrpt.log</b>
Display processing for only the most recent operation.	<b>show log escrpt.log   last</b>
Display processing for script errors.	<b>show log escrpt.log   match error</b>
Display processing for a particular script.	<b>show log escrpt.log   match <i>filename</i></b>

#### Example: Minimum Configuration for Enabling Traceoptions for Event Scripts

Display the trace output of the event script file `source-route.xml`:

```
[edit]
event-options {
  event-script {
    file source-route.xml;
    traceoptions flag output;
  }
}

[edit]
user@host# commit
[edit]
user@host# run show log escrpt.log | last
```

## Configuring Tracing of Event Scripts

You cannot change the directory (`/var/log`) to which trace files are written. However, you can customize other trace file settings by including the following statements at the **[edit event-options event-script traceoptions]** hierarchy level:

```
[edit event-options event-script traceoptions]
file <filename> <files number> <size size> <world-readable | no-world-readable>;
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
```

```
flag xslt;  
no-remote-trace;
```

These statements are described in the following sections:

- [Configuring the Event Script Log Filename on page 636](#)
- [Configuring the Number and Size of Event Script Log Files on page 636](#)
- [Configuring Access to Event Script Log Files on page 637](#)
- [Configuring the Event Script Trace Operations on page 637](#)

---

### Configuring the Event Script Log Filename

By default, the name of the file that records trace output is **escript.log**. You can specify a different name by including the **file** statement at the **[edit event-options event-script traceoptions]** hierarchy level:

```
[edit event-options event-script traceoptions]  
file filename;
```

---

### Configuring the Number and Size of Event Script Log Files

By default, when the trace file reaches 128 KB in size, it is renamed and compressed to **filename.0.gz**, then **filename.1.gz**, and so on, until there are 10 trace files. Then the oldest trace file (**filename.9.gz**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit event-options event-script traceoptions file <filename>]** hierarchy level:

```
[edit event-options event-script traceoptions file <filename>]  
files number size size;
```

For example, set the maximum file size to 640 KB and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 640 KB, it is renamed and compressed to **filename.0.gz**, and a new file called **filename** is created. When **filename** reaches 640 KB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is renamed and compressed to **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19.gz**) is overwritten.

The number of files can range from 2 through 1000 files. The file size can range from 10 KB through 1 gigabyte (GB).



#### NOTE:

If you set either a maximum file size or a maximum number of trace files, you also must specify the other parameter and a filename.

---

### Configuring Access to Event Script Log Files

By default, access to the event script log file is restricted to the owner. You can manually configure access by including the **world-readable** or **no-world-readable** statement at the **[edit event-options event-script traceoptions file <filename>]** hierarchy level.

```
[edit event-options event-script traceoptions file <filename>]
(world-readable | no-world-readable);
```

The **no-world-readable** statement restricts event script log access to the owner. The **world-readable** statement enables unrestricted access to the event script log file.

### Configuring the Event Script Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit event-options event-script traceoptions]** hierarchy level:

```
[edit event-options event-script traceoptions]
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
flag xslt;
```

Table 48 on page 637 describes the meaning of the event script tracing flags.

**Table 48: Event Script Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>events</b>	Trace important events.	On
<b>input</b>	Trace event script input data.	Off
<b>offline</b>	Generate data for offline development.	Off
<b>output</b>	Trace event script output data.	Off
<b>rpc</b>	Trace event script RPCs.	Off
<b>xslt</b>	Trace the Extensible Stylesheet Language Transformations (XSLT) library.	Off

#### Related Documentation

- [Enabling and Executing Event Scripts on page 626](#)
- [Tracing Event Policy Processing on page 631](#)
- [traceoptions on page 823](#)



## PART 10

# SNMP Scripts

- [SNMP Scripts Overview on page 641](#)
- [Creating and Executing SNMP Scripts on page 643](#)
- [SNMP Script Example on page 649](#)
- [Troubleshooting SNMP Scripts on page 655](#)





## CHAPTER 36

# SNMP Scripts Overview

- [SNMP Scripts Overview on page 641](#)

## SNMP Scripts Overview

Junos OS SNMP scripts, which are supported in Junos OS Release 15.1 and later releases, provide the flexibility to support custom MIBs until they are implemented in the Junos operating system (Junos OS). SNMP scripts are triggered automatically when the SNMP manager requests information from the SNMP agent for an object identifier (OID) that is mapped to an SNMP script for an unsupported OID. The script acts like an SNMP subagent, and the system sends the return value from the script to the network management system (NMS).

[Table 49 on page 641](#) lists SNMP script tasks and where to find details.

**Table 49: SNMP Script Tasks**

SNMP Script Task	Find Details Here
Getting started with the SLAX script language	<ul style="list-style-type: none"><li>• <a href="#">SLAX Overview on page 57</a></li><li>• <a href="#">Understanding Extension Functions in Junos OS Automation Scripts on page 183</a></li></ul>
Developing SNMP scripts	<ul style="list-style-type: none"><li>• <a href="#">Required Boilerplate for SNMP Scripts on page 643</a></li><li>• <a href="#">Example: Processing Unsupported OIDs with an SNMP Script on page 649</a></li></ul>
Configuring CLI for SNMP scripts	<ul style="list-style-type: none"><li>• <a href="#">Enabling an SNMP Script on page 646</a></li><li>• <a href="#">snmp</a></li></ul>
Debugging SNMP scripts	<ul style="list-style-type: none"><li>• <a href="#">Tracing SNMP Script Processing on page 655</a></li></ul>
Managing SNMP scripts	<ul style="list-style-type: none"><li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li><li>• <a href="#">Using an Alternate Source Location for a Script on page 673</a></li><li>• <a href="#">Synchronizing Scripts Between Routing Engines on page 702</a></li><li>• <a href="#">Storing Scripts in Flash Memory on page 664</a></li><li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li></ul>
Adding security for SNMP scripts	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li></ul>

SNMP scripts are based on the Junos XML management protocol and the Junos XML API, which are discussed in [“Junos XML Management Protocol and Junos XML API Overview” on page 11](#). You can write SNMP scripts in Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX). SNMP scripts written in SLAX and XSLT use XPath to locate the operational objects to be inspected and XSLT constructs to specify the actions to perform on the located operational objects. The actions can change the output or execute additional commands based on the output.

**Related  
Documentation**

- [Required Boilerplate for SNMP Scripts on page 643](#)
- [SLAX Overview on page 57](#)
- [XPath Overview on page 19](#)
- [XSLT Overview on page 17](#)

# Creating and Executing SNMP Scripts

- [Required Boilerplate for SNMP Scripts on page 643](#)
- [Enabling an SNMP Script on page 646](#)
- [Configuring Checksum Hashes for an SNMP Script on page 647](#)

## Required Boilerplate for SNMP Scripts

---

Junos OS SNMP scripts can be written in Extensible Stylesheet Language Transformations (XSLT), Stylesheet Language Alternative Syntax (SLAX), or Python. SNMP scripts must include the necessary boilerplate required for that script language for both basic script functionality, as well as any optional functionality used within the script such as the Junos OS extension functions and named templates. This topic provides standard boilerplate that can be used in XSLT, SLAX, and Python scripts.

### SLAX Boilerplate for SNMP Scripts

The SLAX SNMP script boilerplate is as follows:

```

1  version 1.0;

2  ns junos = "http://xml.juniper.net/junos/*/junos";
3  ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
4  ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
5  ns dyn = "http://exslt.org/dynamic";
6  ns snmp extension = "http://exslt.org/functions";

7  match / {
8    <snmp-script-results> {
9      var $snmp-action = snmp-script-input/snmp-action;
10     var $snmp-oid = snmp-script-input/snmp-oid;
11     <!-- ... insert your code here ... -->
12     <snmp-oid> $snmp-oid;
13     <snmp-type> $snmp-type;
14     <snmp-value> $snmp-value;
15   }
16 }
```



**NOTE:** SLAX and XSLT SNMP scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

Line 1 specifies the version of XML as 1.0.

```
1    version 1.0;
```

Lines 2 through 6 list all the namespace mappings commonly used in SNMP scripts. Not all of these prefixes are used in this example. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script. These namespace mappings contain named templates that you can call in your scripts. These named templates are discussed in [“Understanding Named Templates in Junos OS Automation Scripts” on page 225](#).

Line 5 and line 6 have EXSLT namespace mappings. SNMP extension functions are defined in the namespace with the associated URI <http://exslt.org/functions>. Line 6 registers the **snmp** extension namespace with the EXSLT functions namespace, allowing you to define customized functions using **snmp** as a prefix within your SLAX script. For more information about the EXSLT namespace, see <http://exslt.org/func/index.html>.

```
2    ns junos = "http://xml.juniper.net/junos/*/junos";
3    ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
4    ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
5    ns dyn = "http://exslt.org/dynamic";
6    ns snmp extension = "http://exslt.org/functions";
```

Line 7 defines an unnamed template, **match /**, that represents the top level of the configuration hierarchy. All XML Path Language (XPath) expressions in the script must start at the top-level element in the configuration hierarchy. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see [“XPath Overview” on page 19](#).

```
7    match / {
```

After the **match /** tag element, the **<snmp-script-results>** container tag must be the top-level child tag, as shown in Line 8. The value of this container is returned to the OID requester.

```
8        <snmp-script-results> {
```

Lines 9 and 10 define variables based on the corresponding elements that you can use in your code to determine whether the action is get or get-next, and the value of the OID.

```
9            var $snmp-action = snmp-script-input/snmp-action;
10           var $snmp-oid = snmp-script-input/snmp-oid;
```

Between Line 10 and Line 11, you can define additional code that includes XSLT templates that are called from within the **match /** template.

Lines 11 through 13 define the values returned by the SNMP script to the OID requester. The value of `<snmp-oid>` is taken from the input value of `snmp-script-input/snmp-oid`. For the SNMP script feature, the following object identifier types for `<snmp-type>` are supported:

- Counter32
- Counter64
- Integer32
- Unsigned32
- Octet String

You set the `<snmp-value>` to the return value from the script.

```
11         <snmp-oid> $snmp-oid;
12         <snmp-type> $snmp-type;
13         <snmp-value> $snmp-value;
```

#### XSLT Boilerplate for SNMP Scripts

The corresponding XSLT boilerplate is:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"
  xmlns:dyn="http://exslt.org/dynamic"
  xmlns:snmp="http://exslt.org/functions"
  version="1.0" extension-element-prefixes="snmp">

  <xsl:template match="/">
    <snmp-script-results>
      <xsl:variable name="snmp-action"
        select="snmp-script-input/snmp-action"/>
      <xsl:variable name="snmp-oid" select="snmp-script-input/snmp-oid"/>

      <!-- Insert your code here -->
      <snmp-oid>
        <xsl:value-of select="$snmp-oid"/>
      </snmp-oid>
      <snmp-type>
        <xsl:value-of select="$snmp-type"/>
      </snmp-type>
      <snmp-value>
        <xsl:value-of select="$snmp-value"/>
      </snmp-value>
    </snmp-script-results>
  </xsl:template>
</xsl:stylesheet>
```

#### Python Boilerplate for SNMP Scripts

Python SNMP scripts do not have a required boilerplate, but they must import any objects that are used in the script. Python SNMP scripts must include the `import jcs` statement in order to use the `get_snmp_action()`, `get_snmp_oid()`, and `emit_snmp_attributes()` functions that retrieve the action and OID values passed to the script and return the data for the requested MIB object.

```
import jcs

if __name__ == '__main__':
```

Python automation scripts do not need to include an interpreter directive line (`#!/usr/bin/env python`) at the start of the script. However, the program will still execute correctly if one is present.

**Related  
Documentation**

- [Example: Processing Unsupported OIDs with an SNMP Script on page 649](#)
- [Understanding Extension Functions in Junos OS Automation Scripts on page 183](#)
- [Understanding Named Templates in Junos OS Automation Scripts on page 225](#)
- [SLAX Templates Overview on page 66](#)
- [XSLT Templates Overview on page 22](#)

---

## Enabling an SNMP Script

SNMP scripts are stored on a device's hard drive in the `/var/db/scripts/snmp` directory or on the flash drive in the `/config/scripts/snmp` directory. Only users in the Junos OS **super-user** login class can access and edit files in these directories. For information about setting the storage location for scripts, see [“Storing and Enabling Scripts” on page 661](#) and [“Storing Scripts in Flash Memory” on page 664](#).

To prevent the execution of unauthorized Python code on devices running Junos OS, unsigned Python scripts must meet certain requirements before you can execute the scripts on a device. Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user. For detailed information about the requirements for executing Python automation scripts on devices running Junos OS, see [“Requirements for Executing Python Automation Scripts on Devices Running Junos OS” on page 169](#).



**NOTE:** If the device has dual Routing Engines and you want to enable an SNMP script to execute on both Routing Engines, you must copy the script to the `/var/db/scripts/snmp` or `/config/scripts/snmp` directory on both Routing Engines. The `commit synchronize` command does not automatically copy scripts between Routing Engines.

You must enable an SNMP script before it can be executed. To enable an SNMP script, include the `file filename` statement at the `[edit system scripts snmp]` hierarchy level, specifying the name of a file containing an SNMP script. Only users who belong to the Junos OS **super-user** login class can enable SNMP scripts.

```
[edit system scripts snmp]
user@host# set file filename
```

SLAX and Python scripts must include the `.slax` or `.py` filename extension, respectively, in both the actual script name and the filename in the configuration. XSLT scripts do not require a filename extension, but we strongly recommend that you append the `.xsl` extension. Whether or not you choose to include the `.xsl` extension on the file, the filename that you add at the `[edit system scripts op file]` hierarchy level must exactly match the

filename of the script in the directory. For example, if the XSLT script filename is **script1.xml**, then you must include **script1.xml** in the configuration hierarchy to enable the script; likewise, if the XSLT script filename is **script1**, then you must include **script1** in the configuration hierarchy.

To determine which SNMP scripts are currently enabled on the device, use the **show** command to display the files configured at the **[edit system scripts snmp]** hierarchy level. To ensure that the enabled files are on the device, list the contents of the **/var/run/scripts/snmp/** directory using the **file list /var/run/scripts/snmp** operational mode command.

By default, you cannot execute unsigned Python scripts on devices running Junos OS. To enable the execution of unsigned Python automation scripts that meet the requirements outlined in “[Requirements for Executing Python Automation Scripts on Devices Running Junos OS](#)” on page 169, you must configure the **language python** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set language python
```

By default, Junos OS executes Python SNMP scripts with the access privileges of the generic, unprivileged user and group **nobody**. Starting in Junos OS Release 16.1R3, you can specify the user under whose access privileges the Python script will execute. To execute a Python SNMP script under the access privileges of a specific user, configure the **python-script-user username** statement at the **[edit system scripts snmp file filename]** hierarchy level.

```
[edit system scripts snmp file filename]
user@host# set python-script-user username
```



**NOTE:** To enable a user who does not belong to the file's user or group class to execute an unsigned Python automation script, the script's file permissions must include read permission for others.

Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user.

## Configuring Checksum Hashes for an SNMP Script

You can configure one or more checksum hashes that can be used to verify the integrity of an SNMP script before the script runs on the switch, router, or security device.

To configure a checksum hash:

1. Create the script.
2. Place the script in the **/var/db/scripts/snmp** directory on the device.

3. Run the script through one or more hash functions to calculate hash values.

Junos OS supports MD5, SHA-1, and SHA-256 hash functions.

```
user@host> file checksum md5 /var/db/scripts/snmp/script1.slax
MD5 (/var/db/scripts/snmp/script1.slax) = 3af7884eb56e2d4489c2e49b26a39a97

user@host> file checksum sha1 /var/db/scripts/snmp/script1.slax
SHA1 (/var/db/scripts/snmp/script1.slax) =
00dc690fb08fb049577d012486c9a6dad34212c0

user@host> file checksum sha-256 /var/db/scripts/snmp/script1.slax
SHA256 (/var/db/scripts/snmp/script1.slax) =
150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

4. Configure the script with one or more hash values.

```
[edit system scripts snmp]
user@host# set file script1.slax checksum md5 3af7884eb56e2d4489c2e49b26a39a97

[edit system scripts snmp]
user@host# set file script1.slax checksum
sha-1 00dc690fb08fb049577d012486c9a6dad34212c0

[edit system scripts snmp]
user@host# set file script1.slax checksum
sha-256 150bf53383769f3bfedd41fe73320777f208d4fda81230cb27b8738
```

During the execution of the script, Junos OS recalculates the checksum value using the configured hash and verifies that the calculated value matches the configured value. If the values differ, the execution of the script fails. When you configure multiple checksum values with different hash algorithms, all the configured values must match the calculated values; otherwise, the script execution fails.

- Related Documentation**
- [Configuring Checksum Hashes for an Event Script on page 628](#)
  - [Configuring Checksum Hashes for an Op Script on page 460](#)
  - [Configuring Checksum Hashes for a Commit Script on page 264](#)



# SNMP Script Example

- [Example: Processing Unsupported OIDs with an SNMP Script on page 649](#)

## Example: Processing Unsupported OIDs with an SNMP Script

---

This sample SNMP script shows how to process object identifiers (OIDs) that are not supported by Junos OS on a device running Junos OS.

- [Requirements on page 649](#)
- [Overview and SNMP Script on page 649](#)
- [Configuration on page 651](#)
- [Verification on page 652](#)

### Requirements

Junos OS Release 15.1 or later when using SLAX SNMP scripts.

Junos OS Release 16.1R3 or later when using Python SNMP scripts.

### Overview and SNMP Script

In this example, two unsupported OIDs are matched and processed. The script returns the value for the requested object by using the `<snmp-script-results>` element in the SLAX script and the `jcs.emit_snmp_attributes()` function in the Python script. The `syslog()` function is called to log the requested SNMP action and OID in the system log file. For more information, see [syslog\(\) Function \(Python, SLAX, and XSLT\)](#).

#### SLAX Syntax

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns dyn = "http://exslt.org/dynamic";
ns snmp extension = "http://exslt.org/functions";

match / {

    var $snmp-action = snmp-script-input/snmp-action;
    var $snmp-oid = snmp-script-input/snmp-oid;

    expr jcs:syslog(8, "snmp-action = ", $snmp-action, " snmp-oid = ",
$snmp-oid);
```

```
if ($snmp-action == 'get') {

    if($snmp-oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1') {

        <snmp-script-results> {
            <snmp-oid> $snmp-oid;
            <snmp-type> "Integer32";
            <snmp-value> "10";
        }
    }

    else if($snmp-oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1') {

        <snmp-script-results> {
            <snmp-oid> $snmp-oid;
            <snmp-type> "Integer32";
            <snmp-value> "211";
        }
    }
}

else if ($snmp-action == 'get-next') {

    if($snmp-oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1') {

        <snmp-script-results> {
            <snmp-oid> ".1.3.6.1.4.1.2636.13.61.1.9.1.1.1";
            <snmp-type> "Integer32";
            <snmp-value> "211";
        }
    }

    else if ($snmp-oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1'){

        <snmp-script-results> {
            <snmp-oid> ".1.3.6.1.4.1.2636.13.61.1.9.1.1.2";
            <snmp-type> "Integer32";
            <snmp-value> "429";
        }
    }
}
}
```

**Python Syntax**

```
import jcs

def main():

    snmp_action = jcs.get_snmp_action()
    snmp_oid = jcs.get_snmp_oid()

    jcs.syslog("8", "snmp_action = ", snmp_action, " snmp_oid = ", snmp_oid)

    if snmp_action == 'get':
        if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "10")
        elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
            jcs.emit_snmp_attributes(snmp_oid, "Integer32", "211")

    elif snmp_action == 'get-next':
        if snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1':
```

```
jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.1",
"Integer32", "211")
elif snmp_oid == '.1.3.6.1.4.1.2636.13.61.1.9.1.1.1':
jcs.emit_snmp_attributes(".1.3.6.1.4.1.2636.13.61.1.9.1.1.2",
"Integer32", "429")

if __name__ == '__main__':
    main()
```

## Configuration

**Step-by-Step Procedure** To download and enable the script:



**NOTE:** You can create SNMP scripts in Python, SLAX, or XSLT. You can use the [request system scripts convert](#) command to convert between SLAX and XSLT.

1. Copy the script into a text file, name the file `snmp_sample.slax` or `snmp_sample.py`, as appropriate, and download it to the `/var/db/scripts/snmp/` directory on the device.



**NOTE:** Unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file.

2. Enable the SNMP script.

In configuration mode, configure the `file filename` statement with the appropriate filename and extension for your script language at the `[edit system scripts snmp]` hierarchy level. Configure two OIDs that will trigger the script. Set the priority of the second OID higher than the default of 127 (1 is the highest priority and 255 the lowest priority).

```
[edit system scripts]
user@host# set snmp file sample_snmp.slax oid .1.3.6.1.4.1.2636.13.61.1.9.1.1
user@host# set snmp file sample_snmp.slax oid .1.3.6.1.4.1.2636.13.61.1.9.1.1.1
user@host# set snmp file sample_snmp.slax oid .1.3.6.1.4.1.2636.13.61.1.9.1.1.1 priority
120
```

3. If the script is written in Python, configure the `language python` statement.

```
[edit system scripts]
user@host# set language python
```

4. If the script is written in Python, configure the user under whose access privileges the script executes.

```
[edit system scripts]
user@host# set snmp file sample_snmp.py python-script-user username
```



**NOTE:** If you do not configure the `python-script-user` statement, then by default, Junos OS executes Python SNMP scripts under the access privileges of the user and group `nobody`.

5. Issue the **commit and-quit** command to commit the configuration and to return to operational mode.

```
[edit]  
user@host# commit and-quit
```

**Results** From configuration mode, confirm your configuration by entering the **show system scripts snmp** command.

```
file sample_snmp.slax {  
  oid .1.3.6.1.4.1.2636.13.61.1.9.1.1;  
  oid .1.3.6.1.4.1.2636.13.61.1.9.1.1.1 {  
    priority 120;  
  }  
}
```

If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

To ensure that the enabled files are on the device, list the contents of the `/var/run/scripts/snmp/` directory using the **file list /var/run/scripts/snmp** operational mode command.

## Verification

---

### Verifying the Script Execution

**Purpose** Verify that the SNMP script functions as expected.

**Action** Issue the **show snmp mib get** command to generate an SNMP request.

```
user@host> show snmp mib get .1.3.6.1.4.1.2636.13.61.1.9.1.1  
host cscript: snmp-action = get snmp-oid = .1.3.6.1.4.1.2636.13.61.1.9.1.1.1  
juniperMIB.13.61.1.9.1.1.1 = 211
```

The system log file contains the following ERRMSG tag after script registration:  
SNMPD\_SCRIPT\_REGISTERED.

The system log file contains the following message after script execution:

```
host cscript: snmp-action = get snmp-oid = .1.3.6.1.4.1.2636.13.61.1.9.1.1
```

**Related Documentation**

- [SLAX Templates Overview on page 66](#)
- [syslog\(\) Function \(Python, SLAX, and XSLT\) on page 220](#)
- `show snmp mib`

- [snmp on page 833](#)



# Troubleshooting SNMP Scripts

- [Tracing SNMP Script Processing on page 655](#)

## Tracing SNMP Script Processing

---

SNMP script tracing operations track all SNMP script operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

The default operation of SNMP script tracing is to log important events in a file called **snmp.log** located in the **/var/log** directory. When the file **snmp.log** reaches 128 kilobytes (KB), it is renamed with a number 0 through 9 (in ascending order) appended to the end of the file and then compressed. The resulting files are **snmp.log.0.gz**, then **snmp.log.1.gz**, until there are 10 trace files. Then the oldest trace file (**snmp.log.9.gz**) is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)

This section discusses the following topics:

- [Minimum Configuration for Enabling Traceoptions for SNMP Scripts on page 655](#)
- [Configuring Tracing of SNMP Scripts on page 656](#)

## Minimum Configuration for Enabling Traceoptions for SNMP Scripts

If no SNMP script trace options are configured, the simplest way to view the trace output of an SNMP script is to configure the **output** trace flag and issue the **show log snmp.log | last** command. To do this, perform the following steps:

1. If you have not done so already, enable an SNMP script by including the **file** statement at the **[edit system scripts snmp]** hierarchy level.

```
[edit system scripts snmp]  
user@host# set file filename
```

2. Enable trace options by including the **traceoptions flag output** statement at the **[edit system scripts snmp]** hierarchy level.

```
[edit system scripts snmp]  
user@host# set traceoptions flag output
```

3. Issue the **commit** command.

```
[edit]
```

```
user@host# commit
```

- To display the resulting trace messages recorded at the end of the `/var/log/snmp.log` file, issue the `show log snmp.log | last` operational mode command.

```
[edit]
```

```
user@host# run show log snmp.log | last
```

Table 41 on page 530 summarizes useful filtering commands that display selected portions of the `snmp.log` file.

**Table 50: SNMP Script Tracing Operational Mode Commands**

Task	Command
Display logging data associated with all SNMP script processing.	<code>show log snmp.log</code>
Display processing for only the most recent operation.	<code>show log snmp.log   last</code>
Display processing for script errors.	<code>show log snmp.log   match error</code>
Display processing for a particular script.	<code>show log snmp.log   match <i>filename</i></code>

#### Example: Minimum Configuration for Enabling Traceoptions for SNMP Scripts

Display the trace output of the SNMP script file `example_snmp.slax`:

```
[edit]
system {
  scripts {
    snmp {
      traceoptions {
        file example_snmp.slax;
        flag output;
      }
    }
  }
}

[edit]
user@host# commit
[edit]
user@host# run show log snmp.log | last
```

## Configuring Tracing of SNMP Scripts

You cannot change the directory (`/var/log`) to which trace files are written. However, you can customize other trace file settings by including the following statements at the `[edit system scripts snmp traceoptions]` hierarchy level:

```
[edit system scripts snmp traceoptions]
file <filename> <files number> <size size> <world-readable | no-world-readable>;
flag all;
flag events;
flag input;
flag offline;
```



```
flag output;
flag rpc;
flag xslt;
no-remote-trace;
```

These statements are described in the following sections:

- [Configuring the SNMP Script Log Filename on page 657](#)
- [Configuring the Number and Size of SNMP Script Log Files on page 657](#)
- [Configuring Access to SNMP Script Log Files on page 657](#)
- [Configuring the SNMP Script Trace Operations on page 658](#)

### Configuring the SNMP Script Log Filename

By default, the name of the file that records trace output is **snmp.log**. You can specify a different name by including the **file** statement at the **[edit system scripts snmp traceoptions]** hierarchy level.

```
[edit system scripts snmp traceoptions]
file filename;
```

### Configuring the Number and Size of SNMP Script Log Files

By default, when the trace file reaches 128 KB in size, it is renamed and compressed to **filename.0.gz**, then **filename.1.gz**, and so on, until there are 10 trace files. Then the oldest trace file (**filename.9.gz**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit system scripts snmp traceoptions file filename]** hierarchy level:

```
[edit system scripts snmp traceoptions file filename]
files number size size;
```

For example, set the maximum file size to 640 KB and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 640 KB, it is renamed and compressed to **filename.0.gz**, and a new file called **filename** is created. When **filename** reaches 640 KB, **filename.0.gz** is renamed **filename.1.gz**, and **filename** is renamed and compressed to **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19.gz**) is overwritten.

The number of files can range from 2 through 1000 files. The file size can range from 10 KB through 1 gigabyte (GB).



**NOTE:** If you set either a maximum file size or a maximum number of trace files, you also must specify the other parameter and a filename.

### Configuring Access to SNMP Script Log Files

By default, access to the SNMP script log file is restricted to the owner. You can manually configure access by including the **world-readable** or **no-world-readable** statement at the **[edit system scripts snmp traceoptions file filename]** hierarchy level.

```
[edit system scripts snmp traceoptions file filename]
```

(world-readable | no-world-readable);

The **no-world-readable** statement restricts SNMP script log access to the owner. The **world-readable** statement enables unrestricted access to the SNMP script log file.

### Configuring the SNMP Script Trace Operations

By default, only important events are logged. You can configure the trace operations to be logged by including the following statements at the **[edit system scripts snmp traceoptions]** hierarchy level:

```
[edit system scripts snmp traceoptions]
flag all;
flag events;
flag input;
flag offline;
flag output;
flag rpc;
flag xslt;
```

Table 42 on page 532 describes the meaning of the SNMP script tracing flags.

Table 51: SNMP Script Tracing Flags

Flag	Description	Default Setting
<b>all</b>	Trace all operations.	Off
<b>events</b>	Trace important events.	On
<b>input</b>	Trace SNMP script input data.	Off
<b>offline</b>	Generate data for offline development.	Off
<b>output</b>	Trace SNMP script output data.	Off
<b>rpc</b>	Trace SNMP script RPCs.	Off
<b>xslt</b>	Trace the Extensible Stylesheet Language Transformations (XSLT) library.	Off

## PART 11

# Managing Automation Scripts

- [Storing and Enabling Scripts on page 661](#)
- [Configuring a Remote Source for Scripts on page 667](#)
- [Configuring the Session Protocol for Scripts on page 679](#)
- [Controlling Execution of Scripts on page 693](#)
- [Synchronizing Scripts Between Routing Engines on page 701](#)
- [Converting Scripts Between SLAX and XSLT on page 711](#)



# Storing and Enabling Scripts

- [Storing and Enabling Scripts on page 661](#)
- [Storing Scripts in Flash Memory on page 664](#)
- [Storing and Using Imported Scripts and Script Functionality on page 665](#)

## Storing and Enabling Scripts

---

To use a Junos OS automation script on a switch, router, or security device, you must copy the script to the device and enable it in the configuration, as described in the following steps:

1. Create the script.
2. Copy the script to the appropriate directory on the device for that script type. Only users who belong to the Junos OS **super-user** login class can access and edit files in the script directories on a device running Junos OS.

By default, scripts are stored in and executed from the **/var/db/scripts** directory on the device's hard disk under the subdirectory appropriate to the script type. You can also store scripts in flash memory in the **/config/scripts** directory under the subdirectory appropriate to the script type.

**commit script**—Copy the script to the **/var/db/scripts/commit** directory on the hard disk or the **/config/scripts/commit** directory in flash memory.

**op script**—Copy the script to the **/var/db/scripts/op** directory on the hard disk or the **/config/scripts/op** directory in flash memory.

**event script**—Copy the script to the **/var/db/scripts/event** directory on the hard disk or the **/config/scripts/event** directory in flash memory.

**snmp script**—Copy the script to the **/var/db/scripts/snmp** directory on the hard disk or the **/config/scripts/snmp** directory in flash memory.



**NOTE:** If the device has dual Routing Engines and you want to enable the script to execute on both Routing Engines, you must copy it to the appropriate directory on both Routing Engines. The **commit synchronize** command does not automatically copy scripts between Routing Engines.



**NOTE:** Junos OS supports using symbolic links for files in the script directories, but the device will only execute the script at the target location if it is signed.

3. For unsigned Python scripts, ensure that the following requirements are met:

- File owner is either root or a user in the Junos OS **super-user** login class.
- Only the file owner has write permission for the file.
- The **language python** statement is configured at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set language python
```



**NOTE:** Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS super-user login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user.



**NOTE:** To enable a user who does not belong to the file's user or group class to execute an unsigned Python automation script, the script's file permissions must include read permission for others.

4. Enable the script by including the **file filename** statement at the appropriate hierarchy level for that script type.

**commit script**—Include the **file filename** statement at the **[edit system scripts commit]** hierarchy level. For instructions, see [“Controlling Execution of Commit Scripts During Commit Operations”](#) on page 261.

**op script**—Include the **file filename** statement at the **[edit system scripts op]** hierarchy level. For instructions, see [“Enabling an Op Script and Defining a Script Alias”](#) on page 459.

**event script**—Include the **file filename** statement at the **[edit event-options event-script]** hierarchy level. For instructions, see [“Enabling and Executing Event Scripts”](#) on page 626.

**snmp script**—Include the **file filename** statement at the **[edit system scripts snmp]** hierarchy level. For instructions, see [“Enabling an SNMP Script”](#) on page 646.

5. If you store scripts in and load them from flash memory, include the **load-scripts-from-flash** statement at the **[edit system scripts]** hierarchy level. For detailed information about storing scripts in flash memory, see [“Storing Scripts in Flash Memory”](#) on page 664.

```
[edit]
```

```
user@host# set system scripts load-scripts-from-flash
```

6. For Python event and SNMP scripts, configure the user under whose access privileges the script executes.

- For event scripts:

```
[edit event-options event-script]
```

```
user@host# set file filename python-script-user username
```

- For SNMP scripts:

```
[edit system scripts snmp]
```

```
user@host# set file filename python-script-user username
```



**NOTE:** If you do not configure the `python-script-user` statement, then by default, Junos OS executes Python event and SNMP scripts under the access privileges of the generic, unprivileged user and group `nobody`. Interactive Python scripts, such as `commit` and `op` scripts, run with the access privileges of the user who executes the command or operation that invokes the script.



**NOTE:** Starting in Junos OS Release 16.1R3, you can execute unsigned Python `commit`, `event`, `op`, and SNMP scripts using the access privileges of authorized users. In Junos OS Release 16.1R2 and earlier releases, all unsigned Python automation scripts are executed using the access privileges of the user and group `nobody`.

7. Issue the `commit` command.

```
[edit]
```

```
user@host# commit
```

Newly enabled `commit` scripts execute during the current `commit` operation. After the `commit` operation completes, enabled event scripts are loaded into memory and can be executed by an event policy that is triggered in response to system log events. For more information, see [“Executing Event Scripts in an Event Policy” on page 573](#). After the `commit` operation completes, `op` scripts can be executed on the device. For more information, see [“Executing an Op Script” on page 461](#).

Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must only be owned by the root user.
16.1R3	Starting in Junos OS Release 16.1R3, you can execute unsigned Python commit, event, op, and SNMP scripts using the access privileges of authorized users. In Junos OS Release 16.1R2 and earlier releases, all unsigned Python automation scripts are executed using the access privileges of the user and group <b>nobody</b> .

#### Related Documentation

- [Storing Scripts in Flash Memory on page 664](#)
- [Storing and Using Imported Scripts and Script Functionality on page 665](#)
- [Controlling Execution of Commit Scripts During Commit Operations on page 261](#)
- [load-scripts-from-flash \(Scripts\) on page 719](#)

## Storing Scripts in Flash Memory

By default, Junos OS automation scripts are stored in and executed from the `/var/db/scripts/` directory on the device's hard disk under the subdirectory appropriate to the script type. To store scripts on and load them from flash memory instead, include the **load-scripts-from-flash** statement at the **[edit system scripts]** hierarchy level:

```
[edit]
user@host# set system scripts load-scripts-from-flash
```

When you add the **load-scripts-from-flash** statement in the configuration, all commit, event, op, SNMP, and script library scripts are loaded from the `/config/scripts/` directory in flash memory under the subdirectory appropriate to the script type. You must manually move scripts from the hard disk to flash memory. They are not moved automatically. Similarly, if you delete the **load-scripts-from-flash** statement from the configuration, you must manually copy the scripts from the flash memory to the hard disk to ensure that the current versions of the scripts are executed. Changing the scripts' physical location has no effect on their operation.

The `/var/run/scripts/` directory always links to the directory from which the scripts are loaded. If you do not set the **load-scripts-from-flash** statement in the configuration, `/var/run/scripts/` points to the `/var/db/scripts/` directory on the device's hard disk. If you set the **load-scripts-from-flash** statement in the configuration, `/var/run/scripts/` points to the `/config/scripts/` directory in flash memory.

To view the scripts currently on the device, list the contents of `/var/run/scripts/type/`, where *type* is the subdirectory appropriate to the script type. In the following example, the **load-scripts-from-flash** statement is not configured. In this case, `/var/run/scripts/commit/` points to the `/var/db/scripts/commit/` directory on the hard disk. Listing the files for `/var/run/scripts/commit/` is identical to listing the files in the `/var/db/scripts/commit/` directory.



```
user@host> file list /var/run/scripts/commit
```

```
/var/run/scripts/commit:
commit-changes-load-replace.slax
commit-protect.slax
```

```
user@host> file list /var/db/scripts/commit
```

```
/var/db/scripts/commit:
commit-changes-load-replace.slax
commit-protect.slax
```

```
user@host> file list /config/scripts/commit
```

```
/config/scripts/commit:
```

#### Related Documentation

- [Storing and Enabling Scripts on page 661](#)
- [Storing and Using Imported Scripts and Script Functionality on page 665](#)
- [load-scripts-from-flash \(Scripts\) on page 719](#)

## Storing and Using Imported Scripts and Script Functionality

Starting in Junos OS Release 11.1, Junos OS provides a dedicated directory for script libraries, where users can store scripts and script functionality that then can be imported into any commit, event, op, or SNMP script. Upon installation, Junos OS creates the `/var/db/scripts/lib/` directory. Junos OS will not overwrite or erase any files in an existing `lib/` directory upon installation or upgrade.

If you store scripts in and run them from flash memory, both the executed scripts and the imported scripts must be present on the flash drive. When you configure the `load-scripts-from-flash` statement at the `[edit system scripts]` hierarchy level, Junos OS creates the `/config/scripts/lib/` directory. When you add or remove the `load-scripts-from-flash` statement in the configuration, you must manually move scripts and script libraries from the hard disk to the flash drive, or vice versa, as appropriate. They are not moved automatically.

Imported scripts must be stored in the `/var/db/scripts/lib/` directory on the hard disk, or if the `load-scripts-from-flash` statement is configured, in the `/config/db/scripts/lib/` directory on the flash drive. To import a script from the script library, include the `<xsl:import>` tag in the style sheet declaration of an XSLT script or the `import` statement in a SLAX script and specify the file location. The following sample code imports the `/var/db/scripts/lib/test.xml` file:

#### XSLT Syntax

```
<?xml version="1.0"?>
<xsl:stylesheet version="1.0">
  <xsl:import href="../lib/test.xml"/>
  ...
</xsl:stylesheet>
```

#### SLAX Syntax

```
version 1.0;
import "../lib/test.xml";
```

- Related Documentation**
- [Storing and Enabling Scripts on page 661](#)
  - [Storing Scripts in Flash Memory on page 664](#)

# Configuring a Remote Source for Scripts

- [Overview of Updating Scripts from a Remote Source on page 667](#)
- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [Example: Configuring and Refreshing from the Master Source for a Script on page 671](#)
- [Using an Alternate Source Location for a Script on page 673](#)

## Overview of Updating Scripts from a Remote Source

---

You can update the automation scripts on a device running Junos OS by retrieving a copy from a remote machine (which can be another device running Junos OS or a regular networked computer). This eases file management, because you can make changes to the master script in a single location and then update the copy on each device where the script is currently enabled. Each device continues to use its locally stored scripts, only updating a script when you issue the appropriate operational or configuration mode command.

For each script, you can configure the **source** statement and a URL at the hierarchy level where you configured the script to define the remote location that houses the master copy of that script. When you then issue the **set refresh** configuration mode command for a script, the device running Junos OS updates its local copy by retrieving the remote master copy from that URL.

You can also store a copy of a particular script at a remote location other than the master source. This is convenient when, for example, the master source cannot be accessed due to network issues or other problems. To refresh a single script or multiple scripts from the remote location, you issue the **set refresh-from** configuration mode command at the appropriate hierarchy level and specify the URL. You can also refresh a single script from a remote location using the **request system scripts refresh-from** operational mode command.

You can use the **set refresh** and **set refresh-from** commands to update either an individual script or multiple scripts of a given type on the device. When you issue the **set refresh** or **set refresh-from** command, the switch, router, or security device immediately attempts to connect to the appropriate remote source for each script. If successful, the device updates the local script with the remote source. If a problem occurs, a set of error messages is returned.

Issuing the **set refresh** or **set refresh-from** command does not add the **refresh** and **refresh-from** statements to the configuration. Thus, these commands behave like operational mode commands by executing an operation, instead of adding a statement to the configuration. The **refresh** and **refresh-from** statements are mutually exclusive.

If a device has dual Routing Engines and you want to update the script on both Routing Engines, you must issue the **set refresh** or **set refresh-from** command on each Routing Engine separately. Alternatively, starting in Junos OS Release 13.2, you can refresh the scripts on the requesting Routing Engine and then use either the **request system scripts synchronize** operational mode command to synchronize scripts to the other Routing Engine or the **commit synchronize scripts** configuration mode command to synchronize all scripts to the other Routing Engine when you commit and synchronize the configuration. If you use the **request system scripts refresh-from** operational mode command to refresh a script from a specific URL, include the **sync** option to refresh the script on both Routing Engines.



**CAUTION:** For commit scripts, we recommend that you do not automate the update function by including the refresh statement as a commit script change element. Even though this might seem like a good way to ensure that the most current commit script is always used, we recommend against it for the following reasons:

- Automated update means that the network must be operational for the commit operation to succeed. If the network goes down after you make a configuration error, you cannot recover quickly.
- If multiple commit scripts need to be updated during each commit operation, the network response time can slow down.
- Automated update is always the last action performed during a commit operation. Consequently, the updated commit script executes only during the next commit operation. This is because commit scripts are applied to the candidate configuration before the software copies any persistent changes generated by the scripts to the candidate configuration. In contrast, if you perform the update operation manually, the updated commit script takes effect as expected, that is, immediately after you commit the refresh statement in the configuration.
- If you automate the update operation, the **refresh-from** statement has no effect, because the **refresh-from** URL conflicts with and is overridden by the **source** statement URL. For information about the **refresh-from** statement, see [“Using an Alternate Source Location for a Script” on page 673](#).

---

**Related Documentation**

- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [Using an Alternate Source Location for a Script on page 673](#)

## Configuring and Using a Master Source Location for a Script

You can store a master copy of each script in a central repository. This eases file management because you can make changes to the master script in one place and then update the copy on each device where the script is currently enabled. This section discusses the following concepts:

- [Configuring the Master Source for a Script on page 669](#)
- [Updating a Script from the Master Source on page 669](#)

### Configuring the Master Source for a Script

To specify the location of the master source for a single script, configure the **source** statement and the URL of the master file. Including the **source** statement in the configuration does not affect the local copy of the script until you issue the **set refresh** command. At that point, the device retrieves the master copy from the specified URL and overwrites the local copy.

The hierarchy location for the **source** statement depends on the script type and filename.

```
[edit system scripts commit file filename]
user@R1# set source url

[edit system scripts op file filename]
user@R1# set source url

[edit event-options event-script file filename]
user@R1# set source url

[edit system scripts snmp file filename]
user@R1# set source url
```

Where

- *filename*—Name of the script.
- *url*—URL of the script's master source file. Specify the source as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.

### Updating a Script from the Master Source

If you configure a master source for one or more scripts on a device, you can refresh the scripts on that device using the **set refresh** configuration mode command. You can update a single script or all scripts of a given script type that have a master source location configured.

The update operation occurs as soon as you issue the **set refresh** command. When you issue the **set refresh** command, the switch, router, or security device immediately attempts to connect to the specified URL and retrieve a copy of the master file. The master copy overwrites the local script stored in the scripts directory on the device. If the **load-scripts-from-flash** statement is configured, the device updates the script on the flash drive instead of the script on the hard disk. If a master source is not defined for a

script, that script is not updated and a warning is issued. For commit scripts, the updated commit script is executed when you next issue the **commit** command.



**NOTE:** Issuing the **set refresh** command does not add the **refresh** statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

The **set refresh** command is unique in the Junos OS CLI in that it behaves like an operational mode command and yet it can be executed from within configuration mode. All other Junos OS CLI operational mode commands can only be executed from command mode. The functionality is provided in this manner as a convenience to users developing commit scripts.

If the device has dual Routing Engines and you want to update a script on both Routing Engines, you must issue the **set refresh** command on each Routing Engine separately. Alternatively, you can refresh the scripts on the requesting Routing Engine and then use either the **request system scripts synchronize** operational mode command to synchronize scripts to the other Routing Engine or the **commit synchronize scripts** configuration mode command to synchronize all scripts to the other Routing Engine when you commit and synchronize the configuration. The **commit synchronize** command does not cause the **refresh** statement to update scripts on both Routing Engines.

To update a single script from its master source, issue the **set refresh** command at the hierarchy level where the script is configured. The hierarchy location depends on the script type and filename as shown in the following examples. The **source** statement specifying the master source location must already be configured.

```
[edit system scripts commit file filename]
user@R1# set refresh

[edit system scripts op file filename]
user@R1# set refresh

[edit system scripts snmp file filename]
user@R1# set refresh

[edit event-options event-script file filename]
user@R1# set refresh
```

Where **filename** is the name of the script.

To update all enabled scripts of a given script type from their master source files, issue the **set refresh** command at the hierarchy level for that script type.

```
[edit system scripts commit]
user@R1# set refresh

[edit system scripts op]
user@R1# set refresh

[edit event-options event-script]
user@R1# set refresh
```

```
[edit system scripts snmp]
user@R1# set refresh
```

## Example: Configuring and Refreshing from the Master Source for a Script

The following example configures a master source file for an op script on a device running Junos OS. The remote source is defined as an HTTP URL. The example uses the master source to update the local copy of the script on the device.

- [Requirements on page 671](#)
- [Overview on page 671](#)
- [Configuration on page 671](#)
- [Verification on page 672](#)

### Requirements

- Routing, switching, or security device running Junos OS.

### Overview

You can store a master copy of each script in a central repository. You can make changes to the master script in one place and then update the local copy of the script on devices where the script is enabled.

This example enables the op script **iso.xml** on a device running Junos OS and then configures a master source location for the script. The remote source for the **iso.xml** file is the HTTP URL `http://my.example.com/pub/scripts/iso.xml`.

Once you configure the master source location, you refresh the local script by issuing the **set refresh** configuration mode command at the hierarchy level where you configured the script. In this example, you would issue the **set refresh** command at the **[edit system scripts op file iso.xml]** hierarchy level.

### Configuration

#### Step-by-Step Procedure

To download, enable, and configure the master source location for the script:

1. Copy the script to the `/var/db/scripts/op/` directory on the device.
2. In configuration mode, configure the **file** statement to enable the **iso.xml** script.
 

```
[edit system scripts op]
user@R1# set file iso.xml
```
3. To configure the master source for the **iso.xml** file, include the **source** statement and source location at the **[edit system scripts op file iso.xml]** hierarchy level.
 

```
[edit system scripts op file iso.xml]
user@R1# set source http://my.example.com/pub/scripts/iso.xml
```
4. Issue the **commit and-quit** command to commit the configuration and exit to operational mode.
 

```
[edit]
```

```
user@R1# commit and-quit
```

**Results**

```
system {
  scripts {
    op {
      file iso.xml {
        source http://my.example.com/pub/scripts/iso.xml;
      }
    }
  }
}
```

---

### Verifying the Script

**Purpose** Verify that the script is on the device and enabled in the configuration.

**Action** Issue the **file list** operational mode command to view the files in the specified directory. The **detail** option provides additional information such as permissions, file size, and modified date.

```
user@R1> file list /var/db/scripts/op detail
```

```
/var/db/scripts/op:
total 128
-rw-r--r--  1 root  admin  13897 Feb 10  2011 iso.xml
...
```

Issue the **show configuration system scripts op** operational mode command to list the op scripts currently enabled on the device.

```
user@R1> show configuration system scripts op
file iso.xml
```

---

### Refreshing the Script from the Master Source

**Step-by-Step Procedure** To refresh the local copy of the script from the master source file:

1. In configuration mode, issue the **set refresh** command at the **[edit system scripts op file iso.xml]** hierarchy level.

```
[edit system scripts op file iso.xml]
user@R1# set refresh
```

## Verification

---

### Verifying the Updated Script

**Purpose** After refreshing the script, verify that the local copy is updated.



**Action** Issue the **file list** operational mode command with the **detail** option to view the files in the specified directory. Verify that the modified date reflects the refreshed version.

```
user@R1> file list /var/db/scripts/op detail

/var/db/scripts/op:
total 128
-rw-r--r--  1 root  admin  14128 May 26  2011 iso.xsl
...
```

**Related Documentation**

- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [Using an Alternate Source Location for a Script on page 673](#)
- [refresh \(Commit Scripts\) on page 735](#)
- [refresh \(Op Scripts\) on page 758](#)
- [refresh \(Event Scripts\) on page 818](#)

## Using an Alternate Source Location for a Script

- [Refreshing a Script from an Alternate Location on page 673](#)
- [Example: Refreshing a Script from an Alternate Source on page 675](#)

### Refreshing a Script from an Alternate Location

In addition to updating a script from the master source defined by the **source** statement, you also can update a script from an alternate location using the **set refresh-from** configuration mode command or the **request system scripts refresh-from** operational mode command. This is convenient when, for example, the master source cannot be accessed due to network issues or other problems.

The update operation occurs as soon as you issue either the **set refresh-from** configuration mode command or the **request system scripts refresh-from** operational mode command. When you issue the command, the switch, router, or security device immediately attempts to connect to the specified URL and retrieve a copy of the file. The copy overwrites the local script stored in the scripts directory on the device. If a copy of the source is not available at the remote URL, that script is not updated and a warning is issued. For commit scripts, the updated commit script is executed when you next issue the **commit** command.

Issuing the **set refresh-from** command does not add the **refresh-from** statement to the configuration. Thus the **set refresh-from** command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.

If a device has dual Routing Engines and you want to update the script on both Routing Engines, you must issue the **set refresh-from** command on each Routing Engine separately. Alternatively, starting in Junos OS Release 13.2, you can refresh the scripts on the requesting Routing Engine and then use either the **request system scripts synchronize** operational mode command to synchronize scripts to the other Routing Engine or the **commit synchronize scripts** configuration mode command to synchronize all scripts to the other Routing Engine when you commit and synchronize the configuration. In

operational mode, you can also use the **request system scripts refresh-from** command with the **sync** option to refresh the script on both Routing Engines.

To update a single script from an alternate source, issue the **set refresh-from** command under the hierarchy level where the script is configured, and specify the location of the remote file. The hierarchy location depends on the script type and filename as shown in the following examples:

```
[edit system scripts commit file filename]
user@R1# set refresh-from url

[edit system scripts op file filename]
user@R1# set refresh-from url

[edit event-options event-script file filename]
user@R1# set refresh-from url
```

To update all enabled scripts of a given script type from an alternate source, issue the **set refresh-from** command at the hierarchy level for that script type, and specify the URL of the remote repository that houses the scripts.

```
[edit system scripts commit]
user@R1# set refresh-from url

[edit system scripts op]
user@R1# set refresh-from url

[edit event-options event-script]
user@R1# set refresh-from url
```

In operational mode, to update a single script from an alternate source, issue the **request system scripts refresh-from** command, and specify the script type, filename, and remote URL.

```
user@R1> request system scripts refresh-from (commit | event | op) file filename url
url <sync>
```

Where

- *filename*—Name of the script.
- *url*—URL of the remote script or directory. Specify the source as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.

If you request to refresh a script that does not exist at the remote site, the device generates an error message. For example:

```
user@host> request system scripts refresh-from op file nonexistent.slax url
http://host1.juniper.net/nonexistent.slax
refreshing 'nonexistent.slax' from 'http://host1.juniper.net/nonexistent.slax'
fetch-secure: http://host1.juniper.net/nonexistent.slax: Not Found
error: file-fetch failed
error: communication error: rpc failed (file-fetch)
error: error retrieving file http://host1.juniper.net/nonexistent.slax
```

When you issue the **set refresh-from** command, Junos OS creates a folder in the **/var/tmp** directory, which is used for the file transfer. After the transfer and refresh operations are complete, Junos OS deletes the temporary folder.

### Example: Refreshing a Script from an Alternate Source

The following example uses an alternate source location to update the local copy of the script on a device running Junos OS. The remote source is defined as an HTTP URL.

- [Requirements on page 675](#)
- [Overview on page 675](#)
- [Configuration on page 675](#)
- [Verification on page 676](#)

#### Requirements

- Routing, switching, or security device running Junos OS.

#### Overview

You can update a script from a location other than that of the master source. This is convenient when, for example, the master source cannot be accessed due to network issues or other problems. You can refresh a single script or all scripts of a given type from the alternate location.

This example enables the **op** script **iso.xml** on a device running Junos OS and then refreshes the script from a location other than the master source location. The remote source for the **iso.xml** file is the HTTP URL <http://my.example.com/pub/scripts2/iso.xml>.

You refresh the local script by issuing the **set refresh-from** configuration mode command at the hierarchy level where you configured the script. In this example, you would issue the **set refresh-from** command at the **[edit system scripts op file iso.xml]** hierarchy level.

#### Configuration

##### Step-by-Step Procedure

To download and enable the script:

1. Copy the script to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, configure the **file** statement to enable the **iso.xml** script.

```
[edit system scripts op]
user@R1# set file iso.xml
```

3. Issue the **commit and-quit** command to commit the configuration and exit to operational mode.

```
[edit]
user@R1# commit and-quit
```

##### Results

```
system {
  scripts {
    op {
      file iso.xml;
```

```
    }  
  }  
}
```

### *Verifying the Script*

**Purpose** Verify that the script is on the device and enabled in the configuration.

**Action** Issue the **file list** operational mode command to view the files in the specified directory. The **detail** option provides additional information such as permissions, file size, and modified date.

```
user@R1> file list /var/db/scripts/op detail
```

```
/var/db/scripts/op:  
total 128  
-rw-r--r--  1 root  admin  13897 Feb 10  2011 iso.xml  
...
```

Issue the **show configuration system scripts op** operational mode command to list the op scripts currently enabled on the device.

```
user@R1> show configuration system scripts op  
file iso.xml
```

### *Refreshing the Script from the Alternate Location*

**Step-by-Step Procedure** To refresh the local copy of the script from the alternate location:

- In configuration mode, issue the **set refresh-from** command at the **[edit system scripts op file iso.xml]** hierarchy level.

```
[edit system scripts op file iso.xml]  
user@R1# set refresh-from http://my.example.com/pub/scripts2/iso.xml
```

---

### **Verification**

#### *Verifying the Updated Script*

**Purpose** After refreshing the script, verify that the local copy is updated.

**Action** Issue the **file list** operational mode command with the **detail** option to view the files in the specified directory. Verify that the modified date reflects the refreshed version.

```
user@R1> file list /var/db/scripts/op detail  
  
/var/db/scripts/op:  
total 128  
-rw-r--r--  1 root  admin  14128 May 26  2011 iso.xml  
...
```

**Related Documentation**

- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [request system scripts refresh-from on page 849](#)

- [refresh-from \(Commit Scripts\) on page 736](#)
- [refresh-from \(Op Scripts\) on page 759](#)
- [refresh-from \(Event Scripts\) on page 819](#)



# Configuring the Session Protocol for Scripts

- [Understanding the Session Protocol in Automation Scripts on page 679](#)
- [Example: Specifying the Session Protocol for a Connection within Scripts on page 681](#)

## Understanding the Session Protocol in Automation Scripts

---

The Junos XML management protocol is a Juniper Networks proprietary protocol that is used to request information from and configure devices running Junos OS. The NETCONF XML management protocol is a standard that is used to request and change configuration information on a routing, switching, or security device. The NETCONF protocol is defined in [RFC 6241](#), *NETCONF Configuration Protocol*, which is available at <http://www.ietf.org/rfc/rfc6241.txt>.

The `jcs:open()` function, which is used in SLAX and XSLT scripts to establish a connection with a device, includes the option to create a session either with the Junos XML protocol server on devices running Junos OS or with the NETCONF server on devices where NETCONF service over SSH is enabled. Prior to Junos OS Release 11.4, the function supported only sessions with the Junos XML protocol server on devices running Junos OS. The additional support for NETCONF sessions enables automation scripts to configure and manage devices in a multi-vendor environment.

The `jcs:open()` function supports the following session protocol types:

- **junoscript**—Session with the Junos XML protocol server on a routing, switching, or security device running Junos OS. This session type supports the operations defined in the Junos XML protocol and the Junos XML API, which are used to configure devices running Junos OS or to request information about the device configuration or operation. This is the default session type.
- **netconf**—Session with the NETCONF XML protocol server on a routing, switching, or security device over an SSHv2 connection. The device to which the connection is made must be enabled for NETCONF service over SSH. NETCONF over SSH is described in [RFC 4742](#), *Using the NETCONF Configuration Protocol over Secure Shell (SSH)*, which is available at <http://www.ietf.org/rfc/rfc4742.txt>.
- **junos-netconf**—Proprietary session with the NETCONF XML protocol server over an SSHv2 connection on a routing, switching, or security device running Junos OS.

The NETCONF server on a device running Junos OS has the additional capabilities defined in <http://xml.juniper.net/netconf/junos/1.0>. The NETCONF server on these devices supports NETCONF XML protocol operations, most Junos XML protocol operations, and the tag elements defined in the Junos XML API. For **netconf** and **junos-netconf** sessions with devices running Junos OS, you should use only native NETCONF XML protocol operations and the extensions available in the Junos XML protocol for configuration functions as documented in the *NETCONF XML Management Protocol Developer Guide*.

The syntax for the **jcs:open()** function when specifying a session protocol is:

<b>SLAX Syntax</b>	<code>var \$connection = jcs:open(remote-hostname, session-options);</code>
<b>XSLT Syntax</b>	<code>&lt;xsl:variable name="connection" select="jcs:open(remote-hostname, session-options)"/&gt;</code>

The *session-options* parameter is an XML node-set that specifies the session type and connection parameters. The session type is one of three values: **junoscript**, **netconf**, or **junos-netconf**. If you do not specify a session type, the default is **junoscript**, which opens a session with the Junos XML protocol server on a device running Junos OS. The format of the node-set is:

```
var $session-options := {  
    <method> ("junoscript" | "netconf" | "junos-netconf");  
    <username> "username";  
    <passphrase> "passphrase";  
    <password> "password";  
    <port> "port-number";  
    <instance> "routing-instance-name";  
    <routing-instance> "routing-instance-name";  
}
```

If you do not specify a username and it is required for the connection, the script uses the local name of the user executing the script. The **<passphrase>** and **<password>** elements serve the same purpose. If you do not specify a passphrase or password element and it is required for authentication, you should be prompted for one during script execution by the device to which you are connecting.

Optionally, you can specify the server port number for **netconf** and **junos-netconf** sessions. The default NETCONF server port number is 830. If you do not specify a port number for a **netconf** or **junos-netconf** session, **jcs:open()** connects to the NETCONF server using port 830. However, if you specify a port number, **jcs:open()** connects to the given port instead. Specifying a port number has no impact on **junoscript** sessions, which are always established over SSH port 22.

To redirect the SSH connection to originate from within a specific routing instance, include the **instance** or **routing-instance** element and the routing instance name. The routing instance must be configured at the **[edit routing-instances]** hierarchy level. The remote device must be reachable either using the routing table for that routing instance or from one of the interfaces configured under that routing instance. The **instance** and **routing-instance** elements serve the same purpose.

To verify the protocol for a specific connection, call the **jcs:get-protocol(connection)** extension function and pass the connection handle as the argument. The function returns "junoscript", "netconf", or "junos-netconf", depending on the session type.



During session establishment with a NETCONF server, the client application and NETCONF server each emit a `<hello>` tag element to specify which operations, or *capabilities*, they support from among those defined in the NETCONF specification or published as proprietary extensions. In `netconf` and `junos-netconf` sessions, you can retrieve the session capabilities of the NETCONF server by calling the `jcs:get-hello(connection)` extension function.

For example, the NETCONF server on a typical device running Junos OS might return the following capabilities:

```
<hello>
  <capabilities>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0</capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:candidate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:validate:1.0
    </capability>
    <capability>
      urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file
    </capability>
    <capability>http://xml.juniper.net/netconf/junos/1.0</capability>
    <capability>http://xml.juniper.net/dmi/system/1.0</capability>
  </capabilities>
  <session-id>20826</session-id>
</hello>
```

#### Related Documentation

- [Example: Specifying the Session Protocol for a Connection within Scripts on page 681](#)
- *Junos XML Management Protocol Developer Guide*
- *NETCONF XML Management Protocol Developer Guide*
- [get-hello\(\) Function \(SLAX and XSLT\) on page 201](#)
- [get-protocol\(\) Function \(SLAX and XSLT\) on page 203](#)
- [open\(\) Function \(SLAX and XSLT\) on page 209](#)

### Example: Specifying the Session Protocol for a Connection within Scripts

The following example demonstrates how to specify the session protocol within a Junos automation script when creating a connection with a remote device. Specifically, the example `op` script establishes a NETCONF session with a remote device running Junos OS, retrieves and prints the NETCONF server capabilities, and then updates and commits the configuration on that device.

- [Requirements on page 682](#)
- [Overview and Script on page 682](#)
- [Configuration on page 687](#)

- [Verification on page 688](#)
- [Troubleshooting on page 690](#)

## Requirements

- Routing, switching, or security device running Junos OS Release 11.4 or later.
- Client application can log in to the device where the NETCONF server resides.
- NETCONF service over SSH is enabled on the device where the NETCONF server resides.

## Overview and Script

The `jcs:open()` extension function includes the option to create a session either with the Junos XML protocol server on devices running Junos OS or with the NETCONF server on devices where NETCONF service over SSH is enabled. In the following example, the script creates a connection and establishes a NETCONF session with a remote device running Junos OS. If the connection and session are successfully established, the script updates the configuration on the remote device to add the `ftp` statement to the `[edit system services]` hierarchy level. The script also retrieves and prints the session protocol and the capabilities of the NETCONF server.

The script takes one argument, `remote-host`, which is the IP address or hostname of the remote device. The `arguments` variable is declared at the global level of the script so that the argument name and description are visible in the command-line interface (CLI) when a user requires context-sensitive help.

The variable `netconf` is a node-set that specifies the session protocol and the connection parameters for the remote device. The value of the `<method>` element is set to “netconf” to establish a session with the NETCONF server over an SSHv2 connection. The `<username>` element specifies the username for the connection. If you do not specify a username and it is required for the connection, the script uses the local name of the user executing the script. In this example, the passphrase and port are not specified. If a passphrase is required for authentication, the remote device should prompt for one during script execution. The script establishes the session using the default NETCONF port 830.

If the connection and establishment of the NETCONF session are successful, the script executes remote procedure calls (RPCs). The RPCs contain the tag elements `<lock>`, `<edit-config>`, `<commit>`, and `<unlock>`, which are NETCONF operations to lock, edit, commit, and unlock the candidate configuration. The script stores the RPC for each task in a separate variable. The results for each RPC are also stored separately and parsed for errors. The script only executes each subsequent step if the previous step is successful. For example, if the script cannot lock the configuration, it does not execute the RPCs to edit, commit, or unlock the configuration.

The variable `rpc-edit-config` contains the tag element `<edit-config>`, which is a NETCONF operation to modify a configuration. The child element, `<config>`, includes the modified portion of the configuration that is merged with the candidate configuration on the device. If errors are encountered, the script calls the `copy-of` statement to copy the result tree fragment variable to the results tree so that the error message prints to the CLI during script execution.

**SLAX Syntax**

```

version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

var $arguments = {
  <argument> {
    <name> "remote-host";
    <description> "device hostname or IP address to which to connect";
  }
}
param $remote-host;

match / {

  <op-script-results> {

    var $netconf := {
      <method> "netconf";
      <username> "bsmith";
    }

    var $rpc-lock-config = {
      <lock> {
        <target> {
          <candidate>;
        }
      }
    }

    var $rpc-unlock-config = {
      <unlock> {
        <target> {
          <candidate>;
        }
      }
    }

    var $rpc-commit = {
      <commit>;
    }

    var $rpc-edit-config = {
      <edit-config> {
        <target> {
          <candidate>;
        }
        <default-operation> "merge";
        <config> {
          <configuration> {
            <system> {
              <services> {
                <ftp>;
              }
            }
          }
        }
      }
    }
  }
}

```

```

    }
  }
}
}
}
}

if ($remote-host = '') {
  <xnm:error> {
    <message> "missing mandatory argument 'remote-host'";
  }
}
else {

  var $connection = jcs:open($remote-host, $netconf);
  if ($connection) {

    /* request protocol and capabilities */
    var $protocol = jcs:get-protocol($connection);
    var $capabilities = jcs:get-hello($connection);

    <output> "\nSession protocol: " _ $protocol _ "\n";
    copy-of $capabilities;

    /* execute rpcs to lock, edit, commit, and unlock config */
    var $lock-reply = jcs:execute($connection, $rpc-lock-config);
    if ($lock-reply/.../rpc-error) {
      copy-of $lock-reply;
    }
    else {
      var $edit-config-reply = jcs:execute($connection, $rpc-edit-config);
      if ($edit-config-reply/.../rpc-error) {
        <output> "Configuration error: " _ $edit-config-reply/.../error-message/_
          _ "\nConfiguration not committed.\n";
        copy-of $edit-config-reply;
      }
      else {
        var $commit-reply = jcs:execute($connection, $rpc-commit);
        if ($commit-reply/.../rpc-error) {
          <output> "Commit error or warning: " _ $commit-reply/.../error-message/;
          copy-of $commit-reply;
        }
      }
      var $unlock-reply = jcs:execute($connection, $rpc-unlock-config);
    }

    expr jcs:close($connection);
  }
  else {
    <output> "\nNo connection - exiting script";
  }
}
}
}
}

```

**XSLT Syntax**

```

<?xml version="1.0" standalone="yes"?>
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"
  xmlns:ext="http://xmlsoft.org/XSLT/namespace" version="1.0">

  <xsl:variable name="arguments">
    <argument>
      <name>remote-host</name>
      <description>device hostname or IP address to which to connect</description>
    </argument>
  </xsl:variable>

  <xsl:param name="remote-host"/>

  <xsl:template match="/">
    <op-script-results>
      <xsl:variable name="netconf-temp-1">
        <method>netconf</method>
        <username>bsmith</username>
      </xsl:variable>
      <xsl:variable xmlns:ext="http://xmlsoft.org/XSLT/namespace"
        name="netconf" select="ext:node-set($netconf-temp-1)"/>

      <xsl:variable name="rpc-lock-config">
        <lock>
          <target>
            <candidate/>
          </target>
        </lock>
      </xsl:variable>

      <xsl:variable name="rpc-unlock-config">
        <unlock>
          <target>
            <candidate/>
          </target>
        </unlock>
      </xsl:variable>

      <xsl:variable name="rpc-commit">
        <commit/>
      </xsl:variable>

      <xsl:variable name="rpc-edit-config">
        <edit-config>
          <target>
            <candidate/>
          </target>
          <default-operation>merge</default-operation>
          <config>
            <configuration>
              <system>
                <services>
                  <ftp/>

```

```

        </services>
    </system>
</configuration>
</config>
</edit-config>
</xsl:variable>

<xsl:choose>
<xsl:when test="$remote-host = ''">
    <xnm:error>
        <message>missing mandatory argument 'remote-host'</message>
    </xnm:error>
</xsl:when>
<xsl:otherwise>
    <xsl:variable name="connection" select="jcs:open($remote-host, $netconf)"/>

    <xsl:choose>
    <xsl:when test="$connection">

        <!-- request protocol and capabilities -->
        <xsl:variable name="protocol" select="jcs:get-protocol($connection)"/>
        <xsl:variable name="capabilities" select="jcs:get-hello($connection)"/>
        <output>
            <xsl:value-of select="concat('&#10;Session protocol: ', $protocol, '&#10;')"/>
        </output>
        <xsl:copy-of select="$capabilities"/>

        <!-- execute rpcs -->
        <xsl:variable name="lock-reply"
            select="jcs:execute($connection, $rpc-lock-config)"/>
        <xsl:choose>
        <xsl:when test="$lock-reply/../../rpc-error">
            <xsl:copy-of select="$lock-reply"/>
        </xsl:when>
        <xsl:otherwise>
            <xsl:variable name="edit-config-reply"
                select="jcs:execute($connection, $rpc-edit-config)"/>
            <xsl:choose>
            <xsl:when test="$edit-config-reply/../../rpc-error">
                <output>
                    <xsl:value-of select="concat('Configuration error: ',
                        $edit-config-reply/../../error-message/,
                        '&#10;Configuration not committed.&#10;')"/>
                </output>
                <xsl:copy-of select="$edit-config-reply"/>
            </xsl:when>
            <xsl:otherwise>
                <xsl:variable name="commit-reply"
                    select="jcs:execute($connection, $rpc-commit)"/>
                <xsl:if test="$commit-reply/../../rpc-error">
                    <output>
                        <xsl:value-of select="concat('Commit error or warning: ',
                            $commit-reply/../../error-message./)"/>
                    </output>
                    <xsl:copy-of select="$commit-reply"/>
                </xsl:if>
            </xsl:otherwise>
            <xsl:copy-of select="$commit-reply"/>
        </xsl:if>
        </xsl:otherwise>
        </xsl:choose>
    </xsl:when>
    </xsl:choose>
</xsl:variable>

```

```

        </xsl:otherwise>
      </xsl:choose>
      <xsl:variable name="unlock-reply" select="jcs:execute($connection,
        $rpc-unlock-config)"/>
    </xsl:otherwise>
  </xsl:choose>

  <xsl:value-of select="jcs:close($connection)"/>
</xsl:when>
<xsl:otherwise>
  <output>No connection - exiting script</output>
</xsl:otherwise>
</xsl:choose>
</xsl:otherwise>
</xsl:choose>
</op-script-results>
</xsl:template>
</xsl:stylesheet>

```

## Configuration

### Step-by-Step Procedure

To download, enable, and test the script:

1. Copy the XSLT or SLAX script into a text file, name the file **netconf-session.xml** or **netconf-session.slax** as appropriate, and copy it to the **/var/db/scripts/op/** directory on the device.
2. In configuration mode, include the **file** statement at the **[edit system scripts op]** hierarchy level and **netconf-session.xml** or **netconf-session.slax** as appropriate.

```

[edit system scripts op]
bsmith@local-host# set file netconf-session.(slax | xml)

```

3. Issue the **commit and-quit** command.

```

[edit]
bsmith@local-host# commit and-quit

```

4. Execute the **op** script on the local device by issuing the **op netconf-session** operational mode command and include any necessary arguments.

In this example, the user, bsmith, is connecting to the remote device, fivestar. The remote device has dual routing engines, so the **commit** operation returns a warning that the **commit synchronize** command should be used to commit the new candidate configuration to both routing engines.

```

bsmith@local-host> op netconf-session remote-host fivestar
bsmith@fivestar's password:
Session protocol: netconf
Commit error or warning:
graceful-switchover is enabled, commit synchronize should be used

```

## Verification

Confirm that the device is working properly.

- [Verifying Op Script Execution on page 688](#)
- [Verifying the Configuration Changes on page 690](#)

---

### Verifying Op Script Execution

**Purpose** Verify that the script behaves as expected.

**Action** Review the script output in the CLI and in the op script log file. Take particular note of any errors that occurred during execution. The default op script log file is `/var/log/op-script.log`. If the log file is significantly lengthy, limit the display by appending the `| last number-of-lines` option to the `show log` command and specify the number of lines to print to the CLI. The output within the `<op-script-results>` element is relevant to the script execution.

```
bsmith@local-host> show log op-script.log | last 100
...output omitted for brevity...
<op-script-results xmlns:junos="http://xml.juniper.net/junos/*/junos"
  xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
  xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0"
  xmlns:ext="http://xmlsoft.org/XSLT/namespace">
<output>
Session protocol: netconf
</output>
<hello>
  <capabilities>
    <capability>urn:ietf:params:xml:ns:netconf:base:1.0</capability>
    <capability>urn:ietf:params:xml:ns:netconf:capability:candidate:1.0
    </capability>
    <capability>urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0
    </capability>
    <capability>urn:ietf:params:xml:ns:netconf:capability:validate:1.0</capability>

    <capability>
      urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file
    </capability>
    <capability>http://xml.juniper.net/netconf/junos/1.0</capability>
    <capability>http://xml.juniper.net/dmi/system/1.0</capability>
  </capabilities>
  <session-id>29087</session-id>
</hello>
  <output>Commit error or warning:
graceful-switchover is enabled, commit synchronize should be used
</output>
  <rpc-error>
<error-severity>warning</error-severity>
<error-message>
graceful-switchover is enabled, commit synchronize should be used
</error-message>
</rpc-error>
  <ok/>
</op-script-results>
```



You can also obtain more descriptive script output on a device running Junos OS by including the **| display xml** option when you execute an op script.

```
bsmith@local-host> op netconf-session remote-host fivestar | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.4D0/junos">
  <output>
    Session protocol: netconf
  </output>
  <hello>
    <capabilities>
      <capability>
        urn:ietf:params:xml:ns:netconf:base:1.0
      </capability>
      <capability>
        urn:ietf:params:xml:ns:netconf:capability:candidate:1.0
      </capability>
      <capability>
        urn:ietf:params:xml:ns:netconf:capability:confirmed-commit:1.0
      </capability>
      <capability>
        urn:ietf:params:xml:ns:netconf:capability:validate:1.0
      </capability>
      <capability>
        urn:ietf:params:xml:ns:netconf:capability:url:1.0?protocol=http,ftp,file
      </capability>
      <capability>
        http://xml.juniper.net/netconf/junos/1.0
      </capability>
      <capability>
        http://xml.juniper.net/dmi/system/1.0
      </capability>
    </capabilities>
    <session-id>
      29087
    </session-id>
  </hello>
  <output>
    Commit error or warning:
    graceful-switchover is enabled, commit synchronize should be used
  </output>
  <rpc-error>
    <error-severity>
      warning
    </error-severity>
    <error-message>
      graceful-switchover is enabled, commit synchronize should be used
    </error-message>
  </rpc-error>
  <ok/>
</op-script-results>
<cli>
  <banner></banner>
</cli>
</rpc-reply>
```

**Meaning** This example creates a NETCONF session on a remote device running Junos OS. The capabilities of the NETCONF server include both standard NETCONF operations and Juniper Networks proprietary extensions, which are defined in <http://xml.juniper.net/netconf/junos/1.0> and <http://xml.juniper.net/dmi/system/1.0>. The

RPC results for the **commit** operation include one warning, but the commit operation is still successful.

### Verifying the Configuration Changes

---

**Purpose** Verify that the commit was successful by viewing the configuration change and the commit log on the remote device.

**Action** On the remote device, execute the **show configuration system services** operational mode command to view the **[edit system services]** hierarchy level of the configuration. If the script is successful, the configuration includes the **ftp** statement.

```
bsmith@fivestar> show configuration system services
ftp;
netconf {
    ssh;
}
```

Additionally, you can review the commit log. On the remote device, execute the **show system commit** operational mode command to view the commit log. In this example, the log confirms that bsmith committed the candidate configuration in a NETCONF session at the given date and time.

```
bsmith@fivestar> show system commit
0   2011-07-11 12:04:01 PDT by bsmith via netconf
1   2011-07-08 15:16:33 PDT by root via cli
```

## Troubleshooting

- [Troubleshooting Connection Errors on page 690](#)
- [Troubleshooting Configuration Lock Errors on page 691](#)
- [Troubleshooting Configuration Syntax Errors on page 692](#)

### Troubleshooting Connection Errors

---

**Problem** The script generates the following error message:

```
hello packet:1:(0) Document is empty
hello packet:1:(0) Start tag expected, '<' not found
error: netconf: could not read hello
error: did not receive hello packet from server
error: Error in creating the session with "fivestar" server
No connection - exiting script
```

Potential causes for the connection error include:

- The device or interface to which you are connecting is down or unavailable.
- The script argument for the IP address or DNS name of the remote device is incorrect.
- The connection timeout value was exceeded before establishing the connection.

- The user authentication for the remote device is not valid or is entered incorrectly.
- You are trying to establish a NETCONF session, and NETCONF over SSH is not enabled on the device where the NETCONF server resides, or it is enabled on a different port.

**Solution** Ensure that the remote device is up and running and that the user has access to the device. Also verify that you supplied the correct argument for the IP address or DNS name of the remote device when executing the script.

For NETCONF sessions, ensure that you have enabled NETCONF over SSH on the device where the NETCONF server resides. Since the example program does not specify a specific port number for the NETCONF session, the session is established on the default NETCONF-over-SSH port, 830. To verify whether NETCONF over SSH is enabled on the default port for a device running Junos OS, enter the following operational mode command on the remote device:

```
bsmith@fivestar> show configuration system services
```

```
netconf {
  ssh;
}
```

If the **netconf** configuration hierarchy is absent on the remote device, issue the following statements in configuration mode to enable NETCONF over SSH on the default port:

```
[edit]
bsmith@fivestar# set system services netconf ssh
bsmith@fivestar# commit
```

If the **netconf** configuration hierarchy specifies a port other than the default port, include the port number in the XML node-set that you pass to the **jcs:open()** function. For example, the following device is configured for NETCONF over SSH on port 12345:

```
bsmith@fivestar> show configuration system services
```

```
netconf {
  ssh {
    port 12345;
  }
}
```

To create a NETCONF session on the alternate port, include the new port number in the XML node-set.

```
var $netconf := {
  <method> "netconf";
  <username> "bsmith";
  <port> "12345";
}
var $connection = jcs:open($remote-host, $netconf);
...
```

### Troubleshooting Configuration Lock Errors

**Problem** The script generates one of the following error messages:

```
configuration database locked by:
root terminal p0 (pid 24113) on since 2011-07-11 11:48:06 PDT, idle 00:07:59
```

```
Users currently editing the configuration:
  root terminal p1 (pid 24279) on since 2011-07-11 12:28:30 PDT
    {master}[edit]

configuration database modified
```

**Solution** Another user currently has a lock on the candidate configuration or has modified the candidate configuration but has not yet committed the configuration. Wait until the lock is released, and then execute the program.

---

### Troubleshooting Configuration Syntax Errors

---

**Problem** The following error message prints to the CLI:

```
Configuration error: syntax error
Configuration not committed.
```

Examine the result tree for additional information. In this case, the result tree shows the following error message:

```
<rpc-error>
  <error-severity>
    error
  </error-severity>
  <error-info>
    <bad-element>
      ftp2
    </bad-element>
  </error-info>
  <error-message>
    syntax error
  </error-message>
</rpc-error>
```

**Solution** The **<bad-element>** tag element indicates that the configuration statement is not valid. Correct the configuration hierarchy and run the script. In this example error, the user entered the tag **<ftp2>** instead of **<ftp>**. Since that is not an acceptable element in the configuration, the NETCONF server returns an error.

- Related Documentation**
- [Understanding the Session Protocol in Automation Scripts on page 679](#)
  - [Junos XML Management Protocol Developer Guide](#)
  - [NETCONF XML Management Protocol Developer Guide](#)
  - [get-hello\(\) Function \(SLAX and XSLT\) on page 201](#)
  - [get-protocol\(\) Function \(SLAX and XSLT\) on page 203](#)
  - [open\(\) Function \(SLAX and XSLT\) on page 209](#)

# Controlling Execution of Scripts

- [Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693](#)
- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)
- [Dampening Script Execution on page 697](#)

## Understanding Limits on Executed Event Policies and Memory Allocation for Scripts

By default, on a device running Junos OS, the maximum number of event policies that can run concurrently in the system is 15, and the maximum amount of memory allocated for the data segment portion of an executed script is half of the total available memory of the system, up to a maximum value of 128 MB. If a script requires more memory during execution than the set maximum limit, the script exits. If the system is running the maximum number of event policies, the system ignores any triggered event policy until such time that another policy finishes. The system logs the `EVENTD_POLICY_LIMIT_EXCEEDED` message for any triggered event policies that were not executed.

Starting in Junos OS Release 12.3, you can configure limits on the maximum number of concurrently running event policies and the maximum amount of memory allocated for the data segment for scripts of a given type.

Depending on the device and its function in the network, it might be necessary to configure larger or smaller limits on the number of event policies that can execute concurrently and the maximum amount of memory allocated to scripts. You might configure smaller limits on critical devices to ensure that priority processes are not adversely impacted, and that the device can perform all necessary functions in the network. Additionally, during normal device operation, you might want to allocate disproportionate amounts of memory to different script types. For example, a device might have a particular type of script that plays a vital role in its operation and requires a specific amount of memory to ensure proper execution.

To set the maximum number of event policies that can run concurrently on a device, configure the **max-policies** *policies* statement at the **[edit event-options]** hierarchy level. You can configure a maximum of 0 through 20 policies.

```
[edit]
event-options {
```

```
    max-policies policies;  
}
```

To set the maximum memory allocated to the data segment for scripts of a given type, configure the **max-datasize** *size* statement under the hierarchy appropriate for that script type, where *size* is the memory in bytes. To specify the memory in kilobytes, megabytes, or gigabytes, append **k**, **m**, or **g**, respectively, to the size. You can configure the memory in the range from 23,068,672 bytes (22 MB) through 1,073,741,824 bytes (1 GB).

```
[edit]  
event-options {  
  event-script {  
    max-datasize size;  
  }  
}  
system {  
  scripts {  
    commit {  
      max-datasize size;  
    }  
    op {  
      max-datasize size;  
    }  
    snmp {  
      max-datasize size;  
    }  
    translation {  
      max-datasize size;  
    }  
  }  
}
```

When the **max-datasize** statement is configured and a script executes, Junos OS sets the maximum memory limit for that script to the configured value irrespective of the total memory available on the system at the time of execution. If the script exceeds the maximum memory limit during execution, it exits gracefully.



**NOTE:** For **op** scripts, the **max-datasize** statement is only enforced for **op** scripts that are local to the device. If you execute an **op** script from a remote location using the **op url** command, Junos OS uses the default memory allocation settings.

---

**Related  
Documentation**

- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)
- [max-datasize on page 720](#)
- [max-policies on page 786](#)

## Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts

On devices running Junos OS, you can configure the maximum number of event policies that can run concurrently on the system and the maximum amount of memory allocated for the data segment for scripts of a given type. This example configures a limit on the number of event policies that the device can execute concurrently and also configures maximum memory limits for executed commit, event, op, and SNMP scripts.

- [Requirements on page 695](#)
- [Overview on page 695](#)
- [Configuration on page 695](#)
- [Verification on page 696](#)

### Requirements

A device running Junos OS Release 12.3 or later.

### Overview

This example configures a device running Junos OS to limit the number of event policies that can run simultaneously on that device to a maximum of 12 policies. Additionally, the example configures each script type with a maximum amount of memory that the system can allocate to the data segment portion of a script of that type. The device is configured to allocate 192 MB for each executed commit script and event script and 100 MB for each executed op and SNMP script.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
set system scripts commit max-datasize 192m
set system scripts op max-datasize 100m
set system scripts snmp max-datasize 100m
set event-options max-policies 12
set event-options event-script max-datasize 192m
```

#### Step-by-Step Procedure

1. Configure the maximum number of event policies that can execute concurrently.  

```
[edit]
user@host# set event-options max-policies 12
```
2. Configure the maximum memory allocated for the data segment for each executed commit script.  

```
[edit]
user@host# set system scripts commit max-datasize 192m
```

3. Configure the maximum memory allocated for the data segment for each executed op script.

```
[edit]
user@host# set system scripts op max-datasize 100m
```

4. Configure the maximum memory allocated for the data segment for each executed SNMP script.

```
[edit]
user@host# set system scripts snmp max-datasize 100m
```

5. Configure the maximum memory allocated for the data segment for each executed event script.

```
[edit]
user@host# set event-options event-script max-datasize 192m
```

6. Commit the configuration.

```
[edit]
user@host# commit
```

---

## Results

```
[edit]
event-options {
  event-script {
    max-datasize 192m;
  }
  max-policies 12;
}
system {
  scripts {
    commit {
      max-datasize 192m;
    }
    op {
      max-datasize 100m;
    }
    snmp {
      max-datasize 100m;
    }
  }
}
```

## Verification

Confirm that the configuration is working properly.

---

### Verifying the Limit on Concurrently Executing Event Policies

**Purpose** If the system is running the maximum number of event policies, the system ignores any triggered event policy until such time that another policy finishes. The system logs the EVENTD\_POLICY\_LIMIT\_EXCEEDED message for any triggered event policies that were not executed. By default, system log messages are recorded in the **messages** log file.



**Action** Review the configured log file to verify whether any policies were barred from execution, because the maximum limit was reached. You can narrow the output to include only the relevant error messages by appending `| match EVENTD_POLICY_LIMIT_EXCEEDED`.

```
user@R1> show log messages | match EVENTD_POLICY_LIMIT_EXCEEDED
Jun 11 17:02:42 R1 eventd[1177]: EVENTD_POLICY_LIMIT_EXCEEDED: Unable to execute
policy 'raise-trap' because current number of policies (12) exceeds system limit
(12)
[output omitted]
```

- Related Documentation**
- [Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693](#)
  - [max-datasize on page 720](#)
  - [max-policies on page 786](#)

## Dampening Script Execution

The Junos OS software automation feature enables you to create commit, op, event, and SNMP scripts to automate devices running Junos OS. These scripts can be CPU-intensive, potentially impacting other critical software processes such as the routing protocol process (rpd). Starting in Junos OS Release 14.1, you can configure Junos OS to dampen or slow down the execution of SLAX and XSLT commit, op, and event scripts. Although the scripts are processed more slowly, enabling script dampening helps to ensure that the other critical software processes can function normally.



**NOTE:** Script dampening is only supported for SLAX and XSLT scripts. Junos OS does not support script dampening for Python automation scripts.

The script dampening feature does not work for certain CPU-intensive remote procedure calls (RPCs), including the following:

- Executing the following RPC might cause the software process to consume numerous CPU cycles. This cannot be avoided as it is outside the control of the cscript process and up to the software process to optimize.

```
var $cmd = <command> "show snmp mib walk .1";
var $out = jcs:invoke($cmd);
```

- The cscript process might consume numerous CPU cycles for this particular RPC, as this involves parsing of the XML configuration returned by the management process (mgd). This cannot be avoided since a single line in the script triggers the processing.

```
var $config = jcs:invoke("get-configuration");
```

To enable script dampening:

1. Include the **dampen** statement.

```
dampen;
```

The **dampen** statement and the script dampening options configured under the **dampen-options** statement can be configured at various hierarchy levels, depending on the type of scripts you want to dampen and whether or not you want to apply dampening to specific scripts or all scripts.

- You can dampen all event scripts by configuring the **dampen** statement at the **[edit event-options event-script]** hierarchy level or dampen a specific event script by configuring the **dampen** statement at the **[edit event-options event-script file filename]** hierarchy level.
- You can dampen all op scripts by configuring the **dampen** statement at the **[edit system scripts op]** hierarchy level or dampen a specific op script by configuring the **dampen** statement at the **[edit system scripts op file filename]** hierarchy level.



**NOTE:** Script dampening is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the **op url** command, Junos OS does not dampen the script.

- You can dampen all commit scripts by configuring the **dampen** statement at the **[edit system scripts commit]** hierarchy level. You cannot dampen a specific commit script. This is a limitation of the cscript process which runs a single commit script at a time.
2. (Optional) Modify the behavior of the script dampening feature by configuring the following dampening options:
- **cpu-factor**—Script dampening is initiated when CPU use exceeds the value specified using this statement.
  - **line-interval**—Specifies the number of lines of script Junos OS executes before pausing.
  - **time-interval**—Specifies the time that script execution is paused.

To configure the script dampening options, configure the following statements. If you configure the **dampen** statement and do not configure values for these statements, the defaults are used. You can configure the **dampen-options** statements at all of the hierarchy levels at which you can configure the **dampen** statement.

If you configure the **dampen-options** statements at both a global hierarchy level (for example, at the **[edit event-options event-script]** hierarchy level) or for a specific script (for example, at the **[edit event-options event-script file filename]** hierarchy level), the file hierarchy level configuration takes precedence.

```
dampen-options {
  cpu-factor cpu-factor;
  line-interval line-interval;
  time-interval microseconds;
}
```

3. Configure the **traceoptions** statement to help debug issues related to the script dampening feature.

- For commit scripts, include the **events** flag in the **traceoptions** statement configuration.

```
[edit system scripts commit traceoptions]  
flag events;
```

- For event scripts, include the **events** flag in the **traceoptions** statement configuration.

```
[edit event-options event-script traceoptions]  
flag events;
```

- For op scripts, include the **events** flag in the **traceoptions** statement configuration.

```
[edit system scripts op traceoptions]  
flag events;
```

The following traceoptions file output is an example of the trace message logged when a script is paused by the script dampening feature:

```
Dec 11 21:40:40 cscript: paused for 100 microseconds
```

- Related Documentation
- [dampen on page 729](#)
  - [dampen-options on page 730](#)



# Synchronizing Scripts Between Routing Engines

- [Understanding Script Synchronization Between Routing Engines on page 701](#)
- [Synchronizing Scripts Between Routing Engines on page 702](#)
- [Example: Synchronizing Scripts Between Routing Engines on page 706](#)

## Understanding Script Synchronization Between Routing Engines

---

Starting in Junos OS Release 13.2, you can manually synchronize commit, event, lib, op, and SNMP scripts between Routing Engines on a device running Junos OS or configure the device to automatically synchronize scripts between Routing Engines when you commit and synchronize the configuration. When invoked, the device synchronizes the scripts from the Routing Engine on which you execute the request (the requesting Routing Engine) to the other Routing Engine (the responding Routing Engine).

In operational mode, you can manually synchronize scripts from the requesting Routing Engine to the responding Routing Engine using the **request system scripts synchronize** command. The command enables you to customize the scope of the synchronization. You can synchronize a single script, all scripts of a specific type, or all scripts on the device. You also have the option to synchronize scripts based on filename or on the timestamp of the file.

In configuration mode, you have the option to synchronize all scripts when you commit and synchronize the configuration. To synchronize scripts on a per-commit basis, use the **commit synchronize scripts** command when you commit and synchronize the configuration. Alternatively, you can configure the device to automatically synchronize scripts from the requesting Routing Engine to the responding Routing Engine every time you issue a **commit synchronize** command. To ensure that scripts are copied from the requesting Routing Engine to the responding Routing Engine during a **commit synchronize** operation, configure the **synchronize** statement at the **[edit system scripts]** hierarchy level.

When you synchronize the configuration and scripts, the device:

1. Performs a commit check on the requesting Routing Engine
2. Synchronizes scripts to the responding Routing Engine
3. Synchronizes the configuration to the responding Routing Engine

4. Performs a commit check on the responding Routing Engine
5. Commits the configuration on the responding Routing Engine
6. Commits the configuration on the requesting Routing Engine

This process ensures that any commit scripts that are required for a successful commit operation are present on the responding Routing Engine before committing the configuration. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

When synchronizing scripts, the device running Junos OS determines the script source and destination directories based on whether the **load-scripts-from-flash** statement is present in the configuration for each Routing Engine. If the **load-scripts-from-flash** statement is configured for the requesting Routing Engine, the device synchronizes the scripts that are in flash memory. Otherwise, the device synchronizes the scripts that are on the hard disk. If the **load-scripts-from-flash** statement is present in the final configuration for the responding Routing Engine, the scripts are synchronized to flash memory. Otherwise, the scripts are synchronized to the hard disk. The device synchronizes a script regardless of whether it is enabled in the configuration or has been updated since the last synchronization.

The **request system scripts refresh-from** operational mode command enables you to manually refresh a single script from a remote URL. Starting in Junos OS Release 13.2, you can synchronize the updated script to the other Routing Engine at the same time by including the **sync** option when you execute the command. When you execute the command, if the **load-scripts-from-flash** statement is configured on the Routing Engine, the script is refreshed in flash memory. Otherwise, the script is refreshed on the hard disk.

#### Related Documentation

- [Synchronizing Scripts Between Routing Engines on page 702](#)
- [Example: Synchronizing Scripts Between Routing Engines on page 706](#)
- [synchronize on page 722](#)
- [request system scripts synchronize on page 851](#)
- [request system scripts refresh-from on page 849](#)

---

## Synchronizing Scripts Between Routing Engines

Starting in Junos OS Release 13.2, you can manually synchronize commit, event, lib, op, and SNMP scripts between Routing Engines on a device running Junos OS or configure the device to automatically synchronize scripts between Routing Engines when you commit and synchronize the configuration.

If the **load-scripts-from-flash** statement is configured for the requesting Routing Engine, the device synchronizes the scripts that are in flash memory. Otherwise, the device synchronizes the scripts that are on the hard disk. If the **load-scripts-from-flash** statement

is present in the final configuration for the responding Routing Engine, the scripts are synchronized to flash memory. Otherwise, the scripts are synchronized to the hard disk. The device synchronizes a script regardless of whether it is enabled in the configuration or has been updated since the last synchronization.

The following sections outline the different methods for synchronizing scripts:

- [Configuring Script Synchronization Between Routing Engines for Commit Synchronize Operations on page 703](#)
- [Synchronizing Scripts Between Routing Engines on a Per-Commit Basis on page 704](#)
- [Synchronizing Scripts Between Routing Engines from Operational Mode on page 704](#)
- [Synchronizing a Script Between Routing Engines After a Refresh on page 705](#)

## Configuring Script Synchronization Between Routing Engines for Commit Synchronize Operations

You can configure a device running Junos OS to synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the other Routing Engine every time you issue the **commit synchronize** command to commit and synchronize the configuration.

To automatically synchronize scripts between Routing Engines during a **commit synchronize** operation:

1. Configure the **synchronize** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set synchronize
```

2. Commit and synchronize the configuration.

```
[edit system scripts]
user@host# commit synchronize
```

When you issue the first and subsequent **commit synchronize** commands, the device performs a commit check on the requesting Routing Engine, synchronizes all scripts to the other Routing Engine, synchronizes, performs a commit check, and commits the configuration on the responding Routing Engine, and finally commits the configuration on the requesting Routing Engine. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

Configuring the **synchronize** statement causes the device to synchronize all scripts even if they have not been updated since the last synchronization. If the device has a large number of scripts that are infrequently updated, it might be more suitable to synchronize scripts either manually using the **request system scripts synchronize** operational mode command or on a per-commit basis using the **commit synchronize scripts** command.

## Synchronizing Scripts Between Routing Engines on a Per-Commit Basis

You can synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the other Routing Engine on a device running Junos OS on a per-commit basis using the **commit synchronize scripts** command when you commit and synchronize the configuration. This is an alternative to configuring the device to synchronize scripts every time you execute a **commit synchronize** operation.

To synchronize scripts between Routing Engines on a per-commit basis:

1. Make all necessary changes to the configuration.
2. Issue the **commit synchronize scripts** command.

```
[edit]
user@host# commit synchronize scripts
```

When you issue the **commit synchronize scripts** command, the device performs a commit check on the requesting Routing Engine, synchronizes all scripts to the other Routing Engine, synchronizes, performs a commit check, and commits the configuration on the responding Routing Engine, and finally commits the configuration on the requesting Routing Engine. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

## Synchronizing Scripts Between Routing Engines from Operational Mode

You can manually synchronize scripts from the requesting Routing Engine to the other Routing Engine on a device running Junos OS by using the **request system scripts synchronize** operational mode command. You can synchronize a single script, all scripts of a specific type, or all scripts on the device. You also have the option to synchronize scripts based on the filename or on the timestamp of the file.

To manually synchronize scripts between Routing Engines, issue the **request system scripts synchronize** command with the desired options.

```
user@host> request system scripts synchronize (all | commit | event | lib | op | snmp)
<file filename> <newer-than time>
```

Specify **all** to synchronize all scripts present on the requesting Routing Engine to the responding Routing Engine. Specify **commit**, **event**, **lib**, **op**, or **snmp** to synchronize all scripts of the given type to the other Routing Engine. Include the **file** option or the **newer-than** option to narrow the scope to only synchronize scripts with the specified filename or date criteria. The format for the **newer-than** argument is YYYY-MM-DD.HH:MM:SS.

For example, the following command synchronizes all commit, event, lib, op, and SNMP scripts that have a timestamp newer than 2012-05-15:

```
user@host> request system scripts synchronize all newer-than 2012-05-15
```



The following command synchronizes a single op script with the filename **vpn-info.slax**.

```
user@host> request system scripts synchronize op file vpn-info.slax
```

A synchronize operation might fail if, for example, you request to synchronize a script that does not exist or if the responding Routing Engine cannot handle the request at that time, because it is performing other CPU-intensive operations. If the synchronize operation fails, the device generates an error message.

The following command requests to synchronize a single event script, but the script does not exist in the event scripts directory, so the device issues an error.

```
user@host> request system scripts synchronize event file nonexistent-file.slax

error: Invalid directory: No such file or directory
warning: No script will be pushed to other RE
```

The following command requests to synchronize a single event script, but the responding Routing Engine does not have the resources to perform the synchronization, so the device issues an error. The device also logs a UI\_SCRIPTS\_COPY\_FAILED error in the system log file with a severity level of error.

```
user@host> request system scripts synchronize event file ospf-neighbor.slax

error: Unable to copy scripts to re: re1
```

## Synchronizing a Script Between Routing Engines After a Refresh

You can manually refresh a single script from a remote URL and synchronize the updated script to the other Routing Engine on a device running Junos OS by using the **request system scripts refresh-from** operational mode command with the **sync** option.

To manually refresh a single script from a remote source and then synchronize the script to the other Routing Engine, issue the **request system scripts refresh-from** command with the **sync** option, and specify the script type, filename, and remote URL. Specify the URL as an HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.

```
user@host> request system scripts refresh-from (commit | event | op | snmp) file filename
url url sync
```

The system uses the script type to determine the directory on the device where the script resides. If the **load-scripts-from-flash** statement is present in the configuration for the Routing Engine, the system refreshes the script in flash memory. Otherwise, the system refreshes the script on the hard disk.

### Related Documentation

- [Understanding Script Synchronization Between Routing Engines on page 701](#)
- [Example: Synchronizing Scripts Between Routing Engines on page 706](#)
- [synchronize on page 722](#)
- [request system scripts synchronize on page 851](#)
- [request system scripts refresh-from on page 849](#)

## Example: Synchronizing Scripts Between Routing Engines

---

This example shows how to configure a device with dual Routing Engines running Junos OS to synchronize all commit, event, lib, op, and SNMP scripts between Routing Engines every time you execute a **commit synchronize** operation.

- [Requirements on page 706](#)
- [Overview on page 706](#)
- [Configuration on page 706](#)
- [Verification on page 707](#)
- [Troubleshooting on page 708](#)

### Requirements

A routing, switching, or security device with dual Routing Engines running Junos OS Release 13.2 (15.1 for SNMP scripts) or later is required.

### Overview

In this example, you configure a device with dual Routing Engines running Junos OS to synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the other Routing Engine whenever you execute a **commit synchronize** command to commit and synchronize the configuration. When configured, the device synchronizes all scripts regardless of whether they are enabled in the configuration.

In this example, the **load-scripts-from-flash** statement is not configured on the requesting Routing Engine. Thus, the device synchronizes the scripts that are on the hard disk of the requesting Routing Engine to the hard disk of the responding Routing Engine.



**NOTE:** On the hard disk, scripts are stored under the `/var/db/scripts` directory in the subdirectory appropriate to the script type. In flash memory, scripts are stored under the `/config/scripts` directory in the subdirectory appropriate to the script type. EX Series switches use the default directory `/config/db/scripts`.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following command, paste it in a text file, and then copy and paste the command into the CLI at the **[edit]** hierarchy level.

```
set system scripts synchronize
```

### Configuring Script Synchronization for Commit Synchronize Operations

---

- Step-by-Step Procedure** To automatically synchronize scripts between Routing Engines during a **commit synchronize** operation:
1. Configure the **synchronize** statement at the **[edit system scripts]** hierarchy level.

```
[edit system scripts]
user@host# set synchronize
```

2. Commit and synchronize the configuration.

```
[edit system scripts]
user@host# commit synchronize
re0:
configuration check succeeds
re1:
commit complete
re0:
commit complete
```

When you issue the first and subsequent **commit synchronize** commands, the device performs a commit check on the requesting Routing Engine, synchronizes all scripts to the other Routing Engine, synchronizes, performs a commit check, and commits the configuration on the responding Routing Engine, and finally commits the configuration on the requesting Routing Engine.

### Results

---

The resulting configuration is:

```
system {
  scripts {
    synchronize;
  }
}
```

## Verification

Confirm that the configuration is working properly and the synchronization is successful.

### Verifying Script Synchronization

---

- Purpose** Verify that the scripts present on the requesting Routing Engine are synchronized to the other Routing Engine.

In this example, the **load-scripts-from-flash** statement is not configured for the requesting Routing Engine. Therefore, the device synchronizes scripts from the **/var/db/scripts** directory on the requesting Routing Engine to the **/var/db/scripts** directory on the responding Routing Engine.

**Action** Use the **file list** operational mode command to view the files in the **/var/db/scripts** directory on each Routing Engine.

1. On the requesting Routing Engine, list the files under the **/var/db/scripts/** directory.

```
user@host> file list /var/db/scripts/* detail
/var/db/scripts/commit:
-rw-r--r--  1 remote wheel      1014 Jul 17 16:18 vpn-commit.slax

/var/db/scripts/event:

/var/db/scripts/lib:

/var/db/scripts/op:
-rw-r--r--  1 root  wheel      11485 Sep 21 2010 jcs-load-config-op.slax
```

2. Log in to the responding Routing Engine, and verify that the files are synchronized.

```
user@host> request routing-engine login other-routing-engine
user@host1> file list /var/db/scripts/* detail
/var/db/scripts/commit:
-rw-r--r--  1 remote wheel      1014 Jul 17 16:18 vpn-commit.slax

/var/db/scripts/event:

/var/db/scripts/lib:

/var/db/scripts/op:
-rw-r--r--  1 root  wheel      11485 Sep 21 2010 jcs-load-config-op.slax
```

**Meaning** The scripts present on each Routing Engine are identical indicating that the device successfully synchronized the scripts from the requesting Routing Engine to the responding Routing Engine.

## Troubleshooting

### Troubleshooting Script Synchronization Failure

---

**Problem** The device does not synchronize the scripts present on the requesting Routing Engine to the other Routing Engine.

**Solution** Verify the following:

- You configured the **synchronize** statement at the **[edit system scripts]** hierarchy level.
- You are viewing the correct directories on each Routing Engine.

If the **load-scripts-from-flash** statement is configured for the requesting Routing Engine, the device synchronizes scripts from flash memory on the requesting Routing Engine to flash memory on the responding Routing Engine.

- You executed a **commit synchronize** command when committing the configuration.

The device does not synchronize scripts for a **commit** operation, only for a **commit synchronize** operation.

- The commit check and commit operations for the requesting Routing Engine are successful.

If the commit check operation for the requesting Routing Engine fails, the process stops, and the scripts are not copied to the responding Routing Engine.

**Related  
Documentation**

- [Understanding Script Synchronization Between Routing Engines on page 701](#)
- [Synchronizing Scripts Between Routing Engines on page 702](#)
- [synchronize on page 722](#)



# Converting Scripts Between SLAX and XSLT

- [Converting Scripts Between SLAX and XSLT on page 711](#)

## Converting Scripts Between SLAX and XSLT

---

SLAX is a C-like alternative syntax to XSLT and can be viewed as a preprocessor for XSLT. Before Junos OS invokes the XSLT processor, the software converts any SLAX constructs in the script (such as **if/else if/else**) to equivalent XSLT constructs (such as **<xsl:choose>** and **<xsl:if>**). For more information about SLAX, see [“SLAX Overview” on page 57](#).

You can use the **request system scripts convert** operational mode command to convert a script or partial script input written in SLAX or XSLT into the alternate language. Users familiar with C and PERL can convert existing XSLT scripts to SLAX to more easily read and maintain the scripts. In addition, converting a script and studying the results facilitates learning the differences between the two languages.

The following sections explain how to convert a script from one language to the other:

- [Converting a Script from SLAX to XSLT on page 711](#)
- [Converting a Script from XSLT to SLAX on page 712](#)

## Converting a Script from SLAX to XSLT

To convert a SLAX script to XSLT, issue the **request system scripts convert slax-to-xslt** operational mode command, and specify the source file, the destination directory, and, optionally, a destination file. The source script is the basis for the new script. The source script is not overwritten by the new script. Starting in Junos OS Release 12.2, you can also include the **partial** option to convert partial script input.

The command syntax is:

```
user@host> request system scripts convert slax-to-xslt source source/filename destination  
destination/<filename> <partial>
```

The following three examples convert a script from SLAX to XSLT using a source and destination directory relevant to the default storage location for the type of script being converted:

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op/script1.xsl
conversion complete
```

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/event/script1.slax
destination /var/db/scripts/event/script1.xsl
conversion complete
```

```
user@host> request system scripts convert slax-to-xslt source
/var/db/scripts/commit/script1.slax destination /var/db/scripts/commit/script1.xsl
conversion complete
```

When you issue the **slax-to-xslt** conversion command, the **script1.slax** file remains unchanged in the source directory, and a new script called **script1.xsl** is added to the destination directory.

```
user@host> file list /var/db/scripts/op
script1.slax
script1.xsl
```

If you specify only the destination directory and do not specify a destination filename, the generated filename is **SLAX-Conversion-Temp** or **slax-temp** depending on the Junos OS release, with a randomly generated, five-character, alpha-numeric extension.

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op/
conversion complete
```

```
user@host> file list /var/db/scripts/op
SLAX-Conversion-Temp.S1hIr
script1.slax
```

## Converting a Script from XSLT to SLAX

To convert an XSLT script to SLAX, issue the **request system scripts convert xslt-to-slax** operational mode command, and specify the source file, the destination directory, and, optionally, a destination file. The source script is the basis for the new script. The source script is not overwritten by the new script.

The command syntax is:

```
user@host> request system scripts convert xslt-to-slax source source/filename destination
destination/<filename> <partial> <version (1.0 | 1.1)>
```

To convert partial script input, include the **partial** option in the command. The **version** option specifies the SLAX version that will be listed in the version statement of the generated script. Specify the version as either 1.0 or 1.1. The default is 1.1. The **partial** and **version** options are supported starting in Junos OS Release 12.2.



The following three examples convert a script from XSLT to SLAX using a source and destination directory relevant to the default storage location for the type of script being converted:

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/op/script1.xml
destination /var/db/scripts/op/script1.slax version 1.0
conversion complete
```

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/event/script1.xml
destination /var/db/scripts/event/script1.slax
conversion complete
```

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/commit/script1.xml
destination /var/db/scripts/commit/script1.slax
conversion complete
```

When you issue the **xslt-to-slax** conversion command, the **script1.xml** file remains unchanged in the source directory, and a new script called **script1.slax** is added to the destination directory.

```
user@host> file list /var/db/scripts/op
script1.slax
script1.xml
```

The SLAX script boilerplate lists the specified SLAX version. In this example, the version is 1.0.

```
user@host> file show /var/db/scripts/op/script1.slax
/* Machine Crafted with Care (tm) by slaxWriter */
version 1.0;
...
```

If you specify only the destination directory and do not specify a destination filename, the generated filename is **SLAX-Conversion-Temp** or **slax-temp** depending on the Junos OS release, with a randomly generated, five-character, alpha-numeric extension.

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/op/script1.xml
destination /var/db/scripts/op/
conversion complete
```

```
user@host> file list /var/db/scripts/op
slax-temp.Vosnd
script1.xml
```

**Related Documentation**

- [SLAX Overview on page 57](#)



## PART 12

# Configuration Statements

- [Configuration Statements Common to All Scripts on page 717](#)
- [Commit Script Configuration Statements on page 723](#)
- [Op Script Configuration Statements on page 743](#)
- [Event Policy Configuration Statements on page 767](#)
- [Event Script Configuration Statements on page 807](#)
- [SNMP Script Configuration Statements on page 825](#)
- [Translation Script Configuration Statements on page 837](#)



## CHAPTER 46

# Configuration Statements Common to All Scripts

- [language \(Scripts\) on page 718](#)
- [load-scripts-from-flash \(Scripts\) on page 719](#)
- [max-datasize on page 720](#)
- [synchronize \(Scripts\) on page 722](#)

## language (Scripts)

**Syntax** `language python;`

**Hierarchy Level** `[edit system scripts]`

**Release Information** Statement introduced in Junos OS Release 16.1.

**Description** Enable unsigned Python scripts to be executed on a device running Junos OS.

Configuring the **language python** statement enables you to execute unsigned Python commit, event, op, and SNMP automation scripts as well as execute unsigned Python scripts developed using the Juniper Extension Toolkit (JET) on devices running Junos OS. To prevent the execution of unauthorized Python code, unsigned Python scripts must be owned by either the root user or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file.



**NOTE:** Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS **super-user** login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must be owned by the root user.

Python commit, event, op, and SNMP scripts must be configured under the hierarchy level appropriate to the script type, and the filename must include the **.py** extension. To execute Python op scripts from a remote site, you must also configure the **allow-url-for-python** statement at the `[edit system scripts op]` hierarchy level.

**Default** If you do not include the **language python** statement, you cannot execute unsigned Python scripts on the device.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

### Release History Table

Release	Description
16.1R3	Starting in Junos OS Release 16.1R3, unsigned Python scripts must be owned by either root or a user in the Junos OS <b>super-user</b> login class, and only the file owner can have write permission for the file. Prior to Junos OS Release 16.1R3, unsigned Python scripts must be owned by the root user.

**Related Documentation**


- [Understanding Python Automation Scripts for Devices Running Junos OS on page 167](#)
- [Configuring the JET Application and its License on a Device Running Junos OS](#)
- [file \(System\)](#)

## load-scripts-from-flash (Scripts)

---

<b>Syntax</b>	load-scripts-from-flash;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5.
<b>Description</b>	Load commit, event, op, SNMP, and library scripts from the device's flash drive instead of the hard disk. To ensure that the current versions of the scripts are executed, you must manually move the scripts from the hard disk to the flash drive. They are not moved automatically.
<b>Default</b>	If you do not include the <b>load-scripts-from-flash</b> statement, scripts are executed from the device's hard disk.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Storing Scripts in Flash Memory on page 664</a></li><li>• <a href="#">scripts on page 737</a></li></ul>

## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p> </div> </div>	
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>



- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

## synchronize (Scripts)

---

<b>Syntax</b>	synchronize;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2.
<b>Description</b>	Synchronize all commit, event, lib, op, and SNMP scripts from the requesting Routing Engine to the responding Routing Engine when executing a <b>commit synchronize</b> operation. The device synchronizes all scripts regardless of whether they are enabled in the configuration or have been updated since the last synchronization.

When you issue the **commit synchronize** command, the device performs a commit check on the requesting Routing Engine, synchronizes all scripts to the other Routing Engine, synchronizes, performs a commit check, and commits the configuration on the responding Routing Engine, and finally commits the configuration on the requesting Routing Engine. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

If the **load-scripts-from-flash** statement is configured for the requesting Routing Engine, the device synchronizes the scripts from flash memory on the requesting Routing Engine to flash memory on the responding Routing Engine. Otherwise, the device synchronizes the scripts from the hard disk on the requesting Routing Engine to the hard disk on the responding Routing Engine.



**NOTE:** Configuring the **synchronize** statement causes the device to synchronize all scripts even if they have not been updated since the last synchronization. If the device has a large number of scripts that are infrequently updated, it might be more suitable to synchronize scripts either manually using the **request system scripts synchronize** operational mode command or on a per-commit basis using the **commit synchronize scripts** command.

---

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Script Synchronization Between Routing Engines on page 701</a></li><li>• <a href="#">Synchronizing Scripts Between Routing Engines on page 702</a></li><li>• <a href="#">Example: Synchronizing Scripts Between Routing Engines on page 706</a></li><li>• <i>commit</i></li><li>• <a href="#">request system scripts synchronize on page 851</a></li></ul>
------------------------------	--

## CHAPTER 47

# Commit Script Configuration Statements

- [allow-transients on page 724](#)
- [apply-macro on page 725](#)
- [checksum on page 726](#)
- [commit on page 727](#)
- [cpu-factor \(Dampen Options\) on page 728](#)
- [dampen on page 729](#)
- [dampen-options on page 730](#)
- [direct-access on page 730](#)
- [file \(Commit Scripts\) on page 731](#)
- [line-interval \(Dampen Options\) on page 732](#)
- [max-datasize on page 733](#)
- [optional on page 734](#)
- [refresh \(Commit Scripts\) on page 735](#)
- [refresh-from \(Commit Scripts\) on page 736](#)
- [scripts on page 737](#)
- [source \(Commit Scripts\) on page 739](#)
- [time-interval \(Dampen Options\) on page 740](#)
- [traceoptions \(Commit, Op, and SNMP Scripts\) on page 741](#)

## allow-transients

---

<b>Syntax</b>	allow-transients;
<b>Hierarchy Level</b>	[edit system scripts commit]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	For Junos OS commit scripts, enable transient configuration changes to be committed.
<b>Default</b>	Transient changes are disabled by default. If you do not include the <b>allow-transients</b> statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Generating a Persistent or Transient Configuration Change in SLAX and XSLT Commit Scripts on page 292</a></li><li>• <a href="#">Creating a Commit Script Macro to Read the Custom Syntax and Generate Related Configuration Statements on page 320</a></li></ul>

## apply-macro

<b>Syntax</b>	<pre>apply-macro <i>apply-macro-name</i> {     <i>parameter-name</i> <i>parameter-value</i>; }</pre>
<b>Hierarchy Level</b>	All hierarchy levels
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Use custom syntax in your configuration, which is expanded into standard Junos OS configuration statements during the commit process using the instructions defined in the corresponding commit script macro.</p> <p>Macros work by locating <b>apply-macro</b> statements that you include in the candidate configuration and using the values specified in the <b>apply-macro</b> statement as parameters to a set of instructions (the macro) defined in a commit script. The commit script alters your configuration from one that contains custom syntax into a full configuration containing standard Junos OS statements.</p> <p>In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides data (or <i>hooks</i>) that are used by commit script macros.</p> <p>You can include the <b>apply-macro</b> statement at any level of the configuration hierarchy. You can include multiple <b>apply-macro</b> statements at each level of the configuration hierarchy; however, each must have a unique name.</p>
<b>Options</b>	<p><b><i>apply-macro-name</i></b>—Name of the <b>apply-macro</b> statement.</p> <p><b><i>parameter-name</i></b>—One or more parameters. Parameters can be any text you want to include in your configuration.</p> <p><b><i>parameter-value</i></b>—A value that corresponds to the parameter name. Parameter values can be any text you want to include in your configuration.</p>
<b>Required Privilege Level</b>	configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Overview of Creating Custom Configuration Syntax with Commit Script Macros on page 313</a></li> </ul>

## checksum

---

Syntax	<code>checksum (md5   sha-256   sha1) hash;</code>
Hierarchy Level	[edit event-options <a href="#">event-script file filename</a> ], [edit system <a href="#">scripts commit file filename</a> ], [edit system <a href="#">scripts op file filename</a> ], [edit system <a href="#">scripts snmp file filename</a> ], [edit system extensions extension-service application file <i>filename</i> ]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced for the <a href="#">[edit system extensions extension-service application file filename]</a> hierarchy level in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.
Description	For Junos OS commit scripts, event scripts, op scripts, SNMP scripts, and scripts developed using the Juniper Extension Toolkit (JET) specify the MD5, SHA-1, or SHA-256 checksum hash. When Junos OS executes a local commit, event, op, or SNMP script, the system verifies the authenticity of the script by using the configured checksum hash.
Options	<b>md5 hash</b> —MD5 checksum of this script.  <b>sha-256 hash</b> —SHA-256 checksum of this script.  <b>sha1 hash</b> —SHA-1 checksum of this script.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li><li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li><li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li><li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li><li>• <a href="#">file checksum md5 on page 844</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha-256 on page 845</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha1 on page 846</a> command in the <i>System Basics and Services Command Reference</i></li></ul>

## commit

<b>Syntax</b>	<pre> commit {   allow-transients;   dampen {     dampen-options {       cpu-factor <i>cpu-factor</i>;       line-interval <i>line-interval</i>;       time-interval <i>time-interval</i>;     }   }   direct-access;   file <i>filename</i> {     checksum (md5   sha-256   sha1) <i>hash</i>;     optional;     refresh;     refresh-from <i>url</i>;     source <i>url</i>;   }   max-datasize <i>size</i>;   refresh;   refresh-from <i>url</i>;   traceoptions {     file &lt;<i>filename</i>&gt; &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace;   } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	For Junos OS commit scripts, configure the commit-time scripting mechanism.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Commit Script Overview on page 241</a></li> <li>• <a href="#">How Commit Scripts Work on page 243</a></li> <li>• <a href="#">Storing and Enabling Scripts on page 661</a></li> </ul>


## cpu-factor (Dampen Options)

---

<b>Syntax</b>	<code>cpu-factor <i>cpu-factor</i>;</code>
<b>Hierarchy Level</b>	[edit event-options event-script dampen <a href="#">dampen-options</a> ], [edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a> ], [edit system scripts commit dampen <a href="#">dampen-options</a> ], [edit system scripts op dampen <a href="#">dampen-options</a> ], [edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Slow down the operation of scripts executed on the device based on how heavily the CPU is being used.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <b>cpu-factor</b> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b>cpu-factor</b>—Specify the CPU use factor at which script execution is slowed down.</p> <p><b>Range:</b> 0 through 100</p> <p><b>Default:</b> 20</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>



## dampen

<b>Syntax</b>	<pre>dampen {   dampen-options {     cpu-factor <i>cpu-factor</i>;     line-interval <i>line-interval</i>;     time-interval <i>time-interval</i>;   } }</pre>
<b>Hierarchy Level</b>	<p>[edit event-options <a href="#">event-script</a>],  [edit event-options event-script <a href="#">file filename</a>],  [edit system scripts <a href="#">commit</a>],  [edit system scripts <a href="#">op</a>],  [edit system scripts op <a href="#">file filename</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Slow down the operation of SLAX and XSLT scripts executed on the device.</p> <p>To dampen all enabled scripts of a given script type, configure the statement at the hierarchy level for that script type. To dampen a single op or event script, configure the statement at the hierarchy level where the script is configured. Script dampening is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <b>op url</b> command, Junos OS does not dampen the script. You cannot dampen individual commit scripts.</p>
	<div>  <p><b>NOTE:</b> Junos OS does not support script dampening for Python automation scripts. If you configure script dampening for all enabled scripts of a given script type, Junos OS logs the following message in the system log file if it executes a Python op or event script: <b>dampen is not supported for python scripts</b>. No warning is logged for Python commit scripts.</p> </div>
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Dampening Script Execution on page 697</a></li> <li>• <a href="#">dampen-options on page 730</a></li> <li>• <a href="#">event-script on page 812</a></li> </ul>

## dampen-options

---

<b>Syntax</b>	<pre>dampen-options {     <b>cpu-factor</b> <i>cpu-factor</i>;     <b>line-interval</b> <i>line-interval</i>;     <b>time-interval</b> <i>time-interval</i>; }</pre>
<b>Hierarchy Level</b>	[edit event-options event-script <b>dampen</b> ], [edit event-options event-script file <i>filename</i> <b>dampen</b> ], [edit system scripts commit <b>dampen</b> ], [edit system scripts op <b>dampen</b> ], [edit system scripts op file <i>filename</i> <b>dampen</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Slow down the operation of automation scripts executed on the device.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>

## direct-access

---

<b>Syntax</b>	<pre>direct-access;</pre>
<b>Hierarchy Level</b>	[edit system <b>scripts commit</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify that commit scripts read input configurations directly from the database.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Executing Large Commit Scripts on page 265</a></li></ul>

## file (Commit Scripts)


<b>Syntax</b>	<pre>file <i>filename</i> {     checksum (md5   sha-256   sha1) <i>hash</i>;     optional;     refresh;     refresh-from <i>url</i>;     source <i>url</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Enable a Junos OS commit script that is located in the <code>/var/db/scripts/commit</code> directory on the hard disk or the <code>/config/scripts/commit</code> directory on the flash drive. During a commit operation, enabled commit scripts are executed in the order specified in the configuration.
<b>Options</b>	<p><b><i>filename</i></b>—Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX) file containing a commit script.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Storing and Enabling Scripts on page 661</a></li> <li>• <a href="#">Controlling Execution of Commit Scripts During Commit Operations on page 261</a></li> </ul>

## line-interval (Dampen Options)

---

<b>Syntax</b>	<code>line-interval <i>line-interval</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options event-script dampen <a href="#">dampen-options</a>],</code> <code>[edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts commit dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the number of lines of scripts to execute before pausing.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <b>line-interval</b> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the default values are used.</p>
<b>Options</b>	<p><b><i>line-interval</i></b>—Specify the line interval at which to pause.</p> <p><b>Range:</b> 1 through 4,294,967,295</p> <p><b>Default:</b> 10</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>


## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p> </div> </div>	
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>


- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

## optional

---


Syntax	optional;
Hierarchy Level	[edit system <a href="#">scripts commit</a> file <i>filename</i> ]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Allow a commit operation to succeed even if the commit script specified in the <b>file</b> statement is missing from the appropriate commit script directory on the device.  <div> <b>NOTE:</b> Commit scripts are stored on a device's hard disk in the <code>/var/db/scripts/commit</code> directory or on the flash drive in the <code>/config/scripts/commit</code> directory. On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/</code> directory on the Director device.</div>
Default	If you do not include the <b>optional</b> statement for a commit script, the commit operation fails if that script is absent from the commit script directory.
Required Privilege Level	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Controlling Execution of Commit Scripts During Commit Operations on page 261</a></li></ul>

## refresh (Commit Scripts)

<b>Syntax</b>	refresh;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ], [edit system <a href="#">scripts commit</a> file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<div>  <p><b>NOTE:</b> On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/</code> directory on the Director device.</p> </div>	
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li> <li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li> <li>• <a href="#">refresh-from (Commit Scripts) on page 736</a></li> <li>• <a href="#">source (Commit Scripts) on page 739</a></li> </ul>

## refresh-from (Commit Scripts)

---

Syntax	<code>refresh-from url;</code>
Hierarchy Level	<code>[edit system <a href="#">scripts commit</a>],</code> <code>[edit system <a href="#">scripts commit</a> file <i>filename</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Overwrite the local copy of all enabled commit scripts or a single enabled commit script with the copy located at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<div> <b>NOTE:</b> This statement is not supported on the QFabric system.</div>	
Options	<b>url</b> —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
Required Privilege Level	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Using an Alternate Source Location for a Script on page 673</a></li><li>• <a href="#">refresh (Commit Scripts) on page 735</a></li><li>• <a href="#">source (Commit Scripts) on page 739</a></li></ul>



## scripts

```

Syntax  scripts {
        commit {
            allow-transients;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
        }
        direct-access;
        file filename {
            checksum (md5 | sha-256 | sha1) hash;
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize size;
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    language python;
    load-scripts-from-flash;
    op {
        allow-url-for-python;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
    }
    file filename {
        allow-commands "regular-expression";
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        checksum (md5 | sha-256 | sha1) hash;
        command filename-alias;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
    }
}

```

```

    }
  }
  description descriptive-text;
  refresh;
  refresh-from url;
  source url;
}
max-datasize size;
no-allow-url
refresh;
refresh-from url;
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
snmp {
  file filename {
    checksum (md5 | sha-256 | sha1) hash;
    oid id {
      priority;
    }
    python-script-user username;
    refresh;
    refresh-from;
    source;
  }
  max-datasize size;
  refresh;
  refresh-from;
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
synchronize;
}

```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure scripting mechanisms for Junos OS commit, op, or SNMP scripts.



**NOTE:** The traceoptions statement is not supported on QFabric systems.

**Options** The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
 maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Commit Script Overview on page 241](#)
- [Op Script Overview on page 443](#)
- [Storing and Enabling Scripts on page 661](#)

## source (Commit Scripts)

**Syntax** `source url;`

**Hierarchy Level** [edit system **scripts** **commit file** *filename*]

**Release Information** Statement introduced in Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 11.1 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify the location of the master source file for a commit script. When you issue the **set refresh** configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the **load-scripts-from-flash** statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.



**NOTE:** On the QFabric system, commit scripts are stored in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/op/` directory on the Director device.

**Options** *url*—The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
 maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [Example: Configuring and Refreshing from the Master Source for a Script on page 671](#)
- [Overview of Updating Scripts from a Remote Source on page 667](#)
- [refresh \(Commit Scripts\) on page 735](#)
- [refresh-from \(Commit Scripts\) on page 736](#)

## time-interval (Dampen Options)

---

<b>Syntax</b>	<code>time-interval <i>microseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options event-script dampen <a href="#">dampen-options</a>],</code> <code>[edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts commit dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the length of time during which to pause the execution of one or more scripts.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b><i>microseconds</i></b>—Specify the time in microseconds to pause script execution.</p> <p><b>Range:</b> 1 through 4,294,967,295 microseconds</p> <p><b>Default:</b> 100 microseconds</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>

## traceoptions (Commit, Op, and SNMP Scripts)

<b>Syntax</b>	<pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;     flag flag;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ], [edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts snmp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define tracing operations for commit, op, and SNMP scripts.
<b>Default</b>	If you do not include this statement, the default operation during script execution is to log important events in the default trace file for that script type. When the file reaches 128 kilobytes (KB), it is renamed and compressed.



**NOTE:** For op scripts, you can only customize tracing operations for local op scripts. For remote op scripts (op scripts executed using the `op url` command), Junos OS always logs the default trace events in the default op script log file.

**Options** **file *filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`. By default, commit script process tracing output is placed in the file `cscript.log`, op script process tracing is placed in the file `op-script.log`, and SNMP script tracing is placed in the file `snmp.log`. If you include the **file** statement, you must specify a filename. To retain the default, you can specify `cscript.log`, `op-script.log`, or `snmp.log` as the filename.

**Default:**

- Commit scripts: `/var/log/cscript.log`
- Op scripts: `/var/log/op-script.log`
- SNMP scripts: `/var/log/snmp.log`

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

**Range:** 2 through 1000

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **events**—Log important events
- **input**—Log script input data
- **offline**—Generate data for offline development
- **output**—Log script output data
- **rpc**—Log script RPCs
- **xslt**—Log the XSLT library

**no-world-readable**—Restrict file access to owner. This is the default.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file.0.gz***. When ***trace-file*** again reaches its maximum size, ***trace-file.0.gz*** is renamed ***trace-file.1.gz*** and ***trace-file*** is renamed and compressed to ***trace-file.0.gz***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—Enable unrestricted file access.

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Commit Script Processing on page 434</a></li><li>• <a href="#">Tracing Op Script Processing on page 529</a></li><li>• <a href="#">Tracing SNMP Script Processing on page 655</a></li><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>
------------------------------	---


## CHAPTER 48

# Op Script Configuration Statements

- [allow-commands \(Op Scripts\) on page 744](#)
- [allow-url-for-python \(Op Scripts\) on page 745](#)
- [arguments \(Op Scripts\) on page 746](#)
- [checksum on page 747](#)
- [command on page 748](#)
- [cpu-factor \(Dampen Options\) on page 749](#)
- [dampen on page 750](#)
- [dampen-options on page 751](#)
- [description on page 751](#)
- [file \(Op Scripts\) on page 752](#)
- [line-interval \(Dampen Options\) on page 753](#)
- [max-datasize on page 754](#)
- [op on page 756](#)
- [no-allow-url \(Op Scripts\) on page 757](#)
- [refresh \(Op Scripts\) on page 758](#)
- [refresh-from \(Op Scripts\) on page 759](#)
- [scripts on page 760](#)
- [source \(Op Scripts\) on page 762](#)
- [time-interval \(Dampen Options\) on page 763](#)
- [traceoptions \(Commit, Op, and SNMP Scripts\) on page 764](#)

## allow-commands (Op Scripts)

---

<b>Syntax</b>	allow-commands <i>"regular-expression"</i> ;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2.
<b>Description</b>	<p>Specify a regular expression that defines which commands to explicitly allow during execution of a local op script.</p> <p>By default, when a user executes a script, the system does not permit the user to execute commands within a script for which their login class does not normally have permission to execute. You can explicitly enable the script to execute certain commands regardless of the permissions of the user invoking the script.</p>
	<div><b>NOTE:</b> Remote op scripts that are executed using the <code>op url</code> command do not support executing unauthorized commands even when you configure the <code>allow-commands</code> statement.</div>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Defining Commands to Allow in an Op Script on page 458</a></li><li>• <a href="#">Enabling an Op Script and Defining a Script Alias on page 459</a></li><li>• <a href="#">Disabling an Op Script on page 464</a></li><li>• <i>Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies</i></li></ul>



## allow-url-for-python (Op Scripts)

<b>Syntax</b>	allow-url-for-python;
<b>Hierarchy Level</b>	[edit system scripts op]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	<p>Enable the execution of Python op scripts from a remote site.</p> <p>When you configure the <b>language python</b> statement at the <b>[edit system scripts]</b> hierarchy level, you can execute local Python scripts on a device running Junos OS. In order to execute Python op scripts from a remote site, you must also configure the <b>allow-url-for-python</b> statement.</p> <p>To verify the integrity of the remote op script before it runs on the device, you can optionally configure a checksum hash for the script using the <b>file checksum (md5   sha1   sha-256)</b> operational mode command. When you run the <b>op url</b> command you can provide the hash value for that script by using the <b>key</b> option. If the calculated and specified values do not match, script execution fails.</p> <p>The <b>no-allow-url</b> statement takes precedence over the <b>allow-url-for-python</b> statement; if both statements are configured, you cannot execute Python op scripts from a remote site on the device.</p>
<b>Default</b>	If you do not include the <b>allow-url-for-python</b> statement, you cannot execute Python op scripts from a remote site, and the <b>op url</b> command generates an error.
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li> <li>• <a href="#">no-allow-url (Op Scripts) on page 757</a></li> </ul>

## arguments (Op Scripts)

---

<b>Syntax</b>	<pre>arguments {     <i>argument-name</i> {         <b>description</b> <i>descriptive-text</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit system <b>scripts op file</b> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6.
<b>Description</b>	Define the expected command-line arguments for an op script. Once configured, the list of possible arguments will be included in the context-sensitive help in the CLI.
<b>Options</b>	<p><b><i>argument-name</i></b>—The name of a command-line argument to an op script.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Declaring and Using Command-Line Arguments in Op Scripts on page 452</a></li><li>• <a href="#">Configuring Help Text for Op Scripts on page 457</a></li><li>• <a href="#">description on page 751</a></li></ul>

## checksum

<b>Syntax</b>	<code>checksum (md5   sha-256   sha1) <i>hash</i>;</code>
<b>Hierarchy Level</b>	[edit event-options <a href="#">event-script file <i>filename</i></a> ], [edit system <a href="#">scripts commit file <i>filename</i></a> ], [edit system <a href="#">scripts op file <i>filename</i></a> ], [edit system <a href="#">scripts snmp file <i>filename</i></a> ], [edit system extensions extension-service application file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced for the <a href="#">[edit system extensions extension-service application file <i>filename</i>]</a> hierarchy level in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.
<b>Description</b>	For Junos OS commit scripts, event scripts, op scripts, SNMP scripts, and scripts developed using the Juniper Extension Toolkit (JET) specify the MD5, SHA-1, or SHA-256 checksum hash. When Junos OS executes a local commit, event, op, or SNMP script, the system verifies the authenticity of the script by using the configured checksum hash.
<b>Options</b>	<b>md5 <i>hash</i></b> —MD5 checksum of this script.  <b>sha-256 <i>hash</i></b> —SHA-256 checksum of this script.  <b>sha1 <i>hash</i></b> —SHA-1 checksum of this script.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li> <li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li> <li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li> <li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li> <li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li> <li>• <a href="#">file checksum md5 on page 844</a> command in the <i>System Basics and Services Command Reference</i></li> <li>• <a href="#">file checksum sha-256 on page 845</a> command in the <i>System Basics and Services Command Reference</i></li> <li>• <a href="#">file checksum sha1 on page 846</a> command in the <i>System Basics and Services Command Reference</i></li> </ul>

## command

---


<b>Syntax</b>	<code>command <i>filename-alias</i>;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op file <i>filename</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure a filename alias for an op script file. This allows you to run the op script by referencing either the script filename or the filename alias.
<b>Options</b>	<i>filename-alias</i> —Alias for the op script file.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling an Op Script and Defining a Script Alias on page 459</a></li></ul>

## cpu-factor (Dampen Options)

<b>Syntax</b>	<code>cpu-factor <i>cpu-factor</i>;</code>
<b>Hierarchy Level</b>	<p>[edit event-options event-script dampen <a href="#">dampen-options</a>],  [edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],  [edit system scripts commit dampen <a href="#">dampen-options</a>],  [edit system scripts op dampen <a href="#">dampen-options</a>],  [edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Slow down the operation of scripts executed on the device based on how heavily the CPU is being used.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <b>cpu-factor</b> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b>cpu-factor</b>—Specify the CPU use factor at which script execution is slowed down.</p> <p><b>Range:</b> 0 through 100</p> <p><b>Default:</b> 20</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Dampening Script Execution on page 697</a></li> </ul>

## dampen

---

Syntax	<pre>dampen {   dampen-options {     cpu-factor <i>cpu-factor</i>;     line-interval <i>line-interval</i>;     time-interval <i>time-interval</i>;   } }</pre>
Hierarchy Level	[edit event-options <a href="#">event-script</a> ], [edit event-options event-script <a href="#">file filename</a> ], [edit system scripts <a href="#">commit</a> ], [edit system scripts <a href="#">op</a> ], [edit system scripts op <a href="#">file filename</a> ]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	<p>Slow down the operation of SLAX and XSLT scripts executed on the device.</p> <p>To dampen all enabled scripts of a given script type, configure the statement at the hierarchy level for that script type. To dampen a single op or event script, configure the statement at the hierarchy level where the script is configured. Script dampening is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <b>op url</b> command, Junos OS does not dampen the script. You cannot dampen individual commit scripts.</p> <div> <b>NOTE:</b> Junos OS does not support script dampening for Python automation scripts. If you configure script dampening for all enabled scripts of a given script type, Junos OS logs the following message in the system log file if it executes a Python op or event script: <b>dampen is not supported for python scripts</b>. No warning is logged for Python commit scripts.</div>
Options	The remaining statements are explained separately.
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li><li>• <a href="#">dampen-options on page 730</a></li><li>• <a href="#">event-script on page 812</a></li></ul>

## dampen-options

<b>Syntax</b>	<code>dampen-options {     <b>cpu-factor</b> <i>cpu-factor</i>;     <b>line-interval</b> <i>line-interval</i>;     <b>time-interval</b> <i>time-interval</i>; }</code>
<b>Hierarchy Level</b>	[edit event-options event-script <b>dampen</b> ], [edit event-options event-script file <i>filename</i> <b>dampen</b> ], [edit system scripts commit <b>dampen</b> ], [edit system scripts op <b>dampen</b> ], [edit system scripts op file <i>filename</i> <b>dampen</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Slow down the operation of automation scripts executed on the device.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Dampening Script Execution on page 697</a></li> </ul>

## description

<b>Syntax</b>	<code>description <i>description</i>;</code>
<b>Hierarchy Level</b>	[edit system <b>scripts op file</b> <i>filename</i> ] [edit system <b>scripts op file</b> <i>filename arguments</i> <i>argument-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Provide help text to describe an op script or its command-line arguments. The help text is included in the context-sensitive help in the command-line interface (CLI).
<b>Options</b>	<b>description</b> —At the script hierarchy level, specify the purpose of the script. At the arguments hierarchy level, provide a description for the command-line argument.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Help Text for Op Scripts on page 457</a></li> <li>• <a href="#">Declaring and Using Command-Line Arguments in Op Scripts on page 452</a></li> <li>• <a href="#">file (Op Scripts) on page 752</a></li> </ul>

## file (Op Scripts)

---

**Syntax**    file *filename* {  
              allow-commands "regular-expression";  
              arguments {  
                  argument-name {  
                      description *descriptive-text*;  
                  }  
              }  
              checksum (md5 | sha-256 | sha1) *hash*;  
              command *filename-alias*;  
              dampen {  
                  dampen-options {  
                      cpu-factor *cpu-factor*;  
                      line-interval *line-interval*;  
                      time-interval *time-interval*;  
                  }  
              }  
              description *descriptive-text*;  
              refresh;  
              refresh-from *url*;  
              source *url*;  
          }

**Hierarchy Level**    [edit system [scripts op](#)]

**Release Information**    Statement introduced in Junos OS Release 7.6.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Enable a Junos OS op script that is located in the `/var/db/scripts/op` directory on the hard disk or the `/config/scripts/op` directory on the flash drive. After enabling an op script, you can execute the script from the CLI using the **op** command.

**Options**    *filename*—Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX) file containing an op script.

The statements are explained separately.

**Required Privilege Level**    maintenance—To view this statement in the configuration.  
                                  maintenance-control—To add this statement to the configuration.

**Related Documentation**


- [Storing and Enabling Scripts on page 661](#)
- [Enabling an Op Script and Defining a Script Alias on page 459](#)



## line-interval (Dampen Options)

<b>Syntax</b>	<code>line-interval <i>line-interval</i>;</code>
<b>Hierarchy Level</b>	<p>[edit event-options event-script dampen <a href="#">dampen-options</a>],  [edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],  [edit system scripts commit dampen <a href="#">dampen-options</a>],  [edit system scripts op dampen <a href="#">dampen-options</a>],  [edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the number of lines of scripts to execute before pausing.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <b>line-interval</b> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the default values are used.</p>
<b>Options</b>	<p><b><i>line-interval</i></b>—Specify the line interval at which to pause.</p> <p><b>Range:</b> 1 through 4,294,967,295</p> <p><b>Default:</b> 10</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Dampening Script Execution on page 697</a></li> </ul>

## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
<div>  <p><b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p> </div>	
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>

- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

## op

```
Syntax  op {
        allow-url-for-python;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        file filename {
            allow-commands "regular-expression";
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
            checksum (md5 | sha-256 | sha1) hash;
            command filename-alias;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
            description descriptive-text;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize size;
        no-allow-url
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
```

**Hierarchy Level** [edit system [scripts](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure an operation scripting mechanism.

**Options** The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Storing and Enabling Scripts on page 661](#)

## no-allow-url (Op Scripts)

**Syntax** no-allow-url;

**Hierarchy Level** [edit system [scripts op](#)]

**Release Information** Statement introduced in Junos OS Release 10.0.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Prohibit the remote execution of Junos OS op scripts. When you include this configuration statement, the **op url** operational mode command does not support executing op scripts from a remote site and generates an error.

**Default** If you do not include the **no-allow-url** statement, authorized users can issue the **op url** command to execute op scripts from a remote site.



**NOTE:** You must configure the **allow-url-for-python** statement in order to execute Python automation scripts from a remote site.


**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**


- [file \(Op Scripts\) on page 752](#)
- [Executing an Op Script from a Remote Site on page 462](#)

## refresh (Op Scripts)

---

Syntax	refresh;
Hierarchy Level	[edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts op file filename</a> ]
Release Information	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 on the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Overwrite the local copy of all enabled op scripts or a single enabled op script with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p> <div> <b>NOTE:</b> On the QFabric system, op scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device.</div>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li><li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li><li>• <a href="#">refresh-from (Op Scripts) on page 759</a></li><li>• <a href="#">source (Op Scripts) on page 762</a></li></ul>

## refresh-from (Op Scripts)

<b>Syntax</b>	<code>refresh-from url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 on the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Overwrite the local copy of all enabled op scripts or a single enabled op script with the copy located at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<div>  <b>NOTE:</b> This statement is not supported on the QFabric system. </div>	
<b>Options</b>	<i>url</i> —Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using an Alternate Source Location for a Script on page 673</a></li> <li>• <a href="#">refresh (Op Scripts) on page 758</a></li> <li>• <a href="#">source (Op Scripts) on page 762</a></li> </ul>

## scripts

```

Syntax  scripts {
        commit {
            allow-transients;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
        }
        direct-access;
        file filename {
            checksum (md5 | sha-256 | sha1) hash;
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize size;
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    language python;
    load-scripts-from-flash;
    op {
        allow-url-for-python;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
    }
    file filename {
        allow-commands "regular-expression";
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        checksum (md5 | sha-256 | sha1) hash;
        command filename-alias;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
    }

```



```

    }
  }
  description descriptive-text;
  refresh;
  refresh-from url;
  source url;
}
max-datasize size;
no-allow-url
refresh;
refresh-from url;
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
snmp {
  file filename {
    checksum (md5 | sha-256 | sha1) hash;
    oid id {
      priority;
    }
    python-script-user username;
    refresh;
    refresh-from;
    source;
  }
  max-datasize size;
  refresh;
  refresh-from;
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
synchronize;
}

```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure scripting mechanisms for Junos OS commit, op, or SNMP scripts.



**NOTE:** The traceoptions statement is not supported on QFabric systems.

**Options** The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Commit Script Overview on page 241](#)
- [Op Script Overview on page 443](#)
- [Storing and Enabling Scripts on page 661](#)

---

## source (Op Scripts)

---

**Syntax** `source url;`

**Hierarchy Level** [edit system [scripts op file filename](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify the location of the master source file for an op script. When you issue the **set refresh** configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the **load-scripts-from-flash** statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.



**NOTE:** On the QFabric system, commit scripts are stored in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/op/` directory on the Director device.

---

**Options** *url*—Master source file for an op script specified as an HTTP URL, FTP URL, or scp-style remote file specification.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring and Using a Master Source Location for a Script on page 669](#)
- [Example: Configuring and Refreshing from the Master Source for a Script on page 671](#)
- [refresh \(Op Scripts\) on page 758](#)
- [refresh-from \(Op Scripts\) on page 759](#)

## time-interval (Dampen Options)

<b>Syntax</b>	<code>time-interval <i>microseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit event-options event-script dampen <a href="#">dampen-options</a>],  [edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],  [edit system scripts commit dampen <a href="#">dampen-options</a>],  [edit system scripts op dampen <a href="#">dampen-options</a>],  [edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the length of time during which to pause the execution of one or more scripts.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b><i>microseconds</i></b>—Specify the time in microseconds to pause script execution.</p> <p><b>Range:</b> 1 through 4,294,967,295 microseconds</p> <p><b>Default:</b> 100 microseconds</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Dampening Script Execution on page 697</a></li> </ul>

## traceoptions (Commit, Op, and SNMP Scripts)

<b>Syntax</b>	<pre>traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;     flag flag;     no-remote-trace; }</pre>
<b>Hierarchy Level</b>	<pre>[edit system scripts commit], [edit system scripts op], [edit system scripts snmp]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Define tracing operations for commit, op, and SNMP scripts.
<b>Default</b>	If you do not include this statement, the default operation during script execution is to log important events in the default trace file for that script type. When the file reaches 128 kilobytes (KB), it is renamed and compressed.



**NOTE:** For op scripts, you can only customize tracing operations for local op scripts. For remote op scripts (op scripts executed using the `op url` command), Junos OS always logs the default trace events in the default op script log file.

**Options** **file *filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`. By default, commit script process tracing output is placed in the file `cscript.log`, op script process tracing is placed in the file `op-script.log`, and SNMP script tracing is placed in the file `snmp.log`. If you include the **file** statement, you must specify a filename. To retain the default, you can specify `cscript.log`, `op-script.log`, or `snmp.log` as the filename.

**Default:**

- Commit scripts: `/var/log/cscript.log`
- Op scripts: `/var/log/op-script.log`
- SNMP scripts: `/var/log/snmp.log`

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file.0.gz***. When ***trace-file*** again reaches its maximum size, ***trace-file.0.gz*** is renamed ***trace-file.1.gz*** and ***trace-file*** is renamed and compressed to ***trace-file.0.gz***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

**Range:** 2 through 1000

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **events**—Log important events
- **input**—Log script input data
- **offline**—Generate data for offline development
- **output**—Log script output data
- **rpc**—Log script RPCs
- **xslt**—Log the XSLT library

**no-world-readable**—Restrict file access to owner. This is the default.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—Enable unrestricted file access.

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Tracing Commit Script Processing on page 434</a></li> <li>• <a href="#">Tracing Op Script Processing on page 529</a></li> <li>• <a href="#">Tracing SNMP Script Processing on page 655</a></li> <li>• <a href="#">Dampening Script Execution on page 697</a></li> </ul>
------------------------------	--



# Event Policy Configuration Statements

- [archive-sites \(Event Policy\) on page 769](#)
- [arguments \(Event Policy\) on page 770](#)
- [attributes-match on page 771](#)
- [change-configuration \(Event Policy\) on page 772](#)
- [commands \(Event Policy Change Configuration\) on page 773](#)
- [commands \(Event Policy Execute Commands\) on page 774](#)
- [commit-options on page 775](#)
- [destination \(Event Policy\) on page 776](#)
- [destinations \(Event Policy\) on page 777](#)
- [equals \(Event Policy\) on page 778](#)
- [event-options on page 779](#)
- [event-script \(Event Policy\) on page 781](#)
- [events \(Event Policy Trigger Events\) on page 782](#)
- [events \(Event Policy Correlated Events\) on page 782](#)
- [execute-commands on page 783](#)
- [facility on page 784](#)
- [generate-event on page 784](#)
- [ignore on page 785](#)
- [matches \(Event Policy\) on page 785](#)
- [max-policies on page 786](#)
- [no-drift \(Generate Events\) on page 786](#)
- [not \(Event Policy\) on page 787](#)
- [output-filename \(Event Policy\) on page 788](#)
- [output-format \(Event Policy\) on page 789](#)
- [policy \(Event Policy\) on page 790](#)
- [priority-override \(Event Policy\) on page 792](#)
- [raise-trap \(Event Policy\) on page 792](#)
- [retry \(Event Policy\) on page 793](#)

- [retry-count \(Event Policy\) on page 794](#)
- [severity \(Event Policy\) on page 795](#)
- [starts-with \(Event Policy\) on page 796](#)
- [then \(Event Policy\) on page 797](#)
- [time-interval \(Generate Events\) on page 798](#)
- [time-of-day on page 799](#)
- [traceoptions \(Event Options\) on page 800](#)
- [transfer-delay \(Event Policy\) on page 802](#)
- [trigger on page 803](#)
- [upload on page 804](#)
- [user-name \(Event Policy\) on page 805](#)
- [within on page 806](#)



## archive-sites (Event Policy)

<b>Syntax</b>	<pre>archive-sites {     url &lt;password password&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">destinations</a> <i>destination-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>For a given destination name, configure archive sites to which an event policy can transfer files. To archive files, you must reference the destination name in the event policy. If you specify more than one archive site for a particular destination name, the device attempts to transfer to the first archive site in the list, moving to the next site only if the transfer fails.</p>
<b>Options</b>	<p><b>url</b>—The archive destination specified as a file URI, an active or passive FTP URI, or a Secure Copy (SCP) URI. Local device directories are also supported (for example, <i>/var/tmp/</i>).</p> <pre>file:&lt;host&gt;/path ftp://username@host:&lt;port&gt;url-path pasvftp://username@host:&lt;port&gt;url-path scp://username@host:&lt;port&gt;url-path &lt;path&gt;/&lt;filename&gt;</pre> <p><b>password <i>password</i></b>—A plain-text password required for logging into the archive site.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Event Policy File Archiving Overview on page 591</a></li> <li>• <a href="#">Example: Defining Destinations for File Archiving by Event Policies on page 592</a></li> <li>• <a href="#">Example: Configuring an Event Policy to Upload Files on page 595</a></li> <li>• <a href="#">destinations (Event Policy) on page 777</a></li> </ul>

## arguments (Event Policy)

---

<b>Syntax</b>	<pre>arguments {     <i>argument-name</i> <i>argument-value</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">policy</a> <i>policy-name</i> <a href="#">then</a> <a href="#">event-script</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define the arguments that are passed to an event script when it is invoked from an event policy.
<b>Options</b>	<p><i>argument-name</i>—Name of the argument.</p> <p><i>argument-value</i>—String specifying the value of the argument.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Pass Arguments to an Event Script on page 576</a></li><li>• <a href="#">Executing Event Scripts in an Event Policy on page 573</a></li><li>• <a href="#">event-script (Event Policy) on page 781</a></li><li>• <a href="#">policy (Event Policy) on page 790</a></li></ul>

## attributes-match

<b>Syntax</b>	<pre>attributes-match {     event1.attribute-name equals event2.attribute-name;     event.attribute-name matches regular-expression;     event1.attribute-name starts-with event2.attribute-name; }</pre>
<b>Hierarchy Level</b>	[edit <b>event-options policy</b> <i>policy-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Execute an event policy only if the attributes of two events are correlated or if the attribute of one event matches a regular expression.</p> <p>If the <b>attributes-match</b> statement includes the <b>equals</b> or <b>starts-with</b> options, or if it includes a <b>matches</b> option that includes a clause for an event that is not specified at the [edit <b>event-options policy</b> <i>policy-name</i> <b>events</b>] hierarchy level, you must include one or more <b>within</b> statements in the same policy configuration.</p> <p>Starting in Junos OS Release 11.1, you can include event policy variables within the statement to differentiate between a trigger event attribute and a correlated event attribute. You can use variables of the following forms:</p> <ul style="list-style-type: none"> <li>• <b>`\${attribute-name}`</b>—The double dollar sign (\$\$) notation represents the event that is triggering a policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event.</li> <li>• <b>`\${event.attribute-name}`</b>—The single dollar sign with the event name (<b>`\${event}`</b>) notation represents the most recent event that matches <b>event</b>. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event.</li> </ul> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li> <li>• <a href="#">Using Regular Expressions to Refine the Set of Events That Trigger a Policy on page 549</a></li> <li>• <a href="#">equals (Event Policy) on page 778</a></li> <li>• <a href="#">matches (Event Policy) on page 785</a></li> <li>• <a href="#">starts-with (Event Policy) on page 796</a></li> <li>• <a href="#">within on page 806</a></li> </ul>

## change-configuration (Event Policy)

---

<b>Syntax</b>	<pre>change-configuration {   commands {     "command";   }   commit-options {     check &lt;synchronize&gt;;     force;     log "comment-string";     synchronize;   }   retry count <i>number</i> interval <i>seconds</i>;   user-name <i>username</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	<p>When the associated event policy is triggered, execute the Junos OS configuration mode commands to update the candidate configuration, and commit the changes.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Change the Configuration on page 558</a></li><li>• <a href="#">Example: Changing the Configuration Using an Event Policy on page 559</a></li><li>• <a href="#">Example: Changing the Interface Configuration in Response to an Event on page 566</a></li><li>• <a href="#">commands (Event Policy Change Configuration) on page 773</a></li></ul>

## commands (Event Policy Change Configuration)

<b>Syntax</b>	<pre>commands {   "command"; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name</a> then <a href="#">change-configuration</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify the configuration mode commands to be issued when the corresponding event policy is triggered. When the event policy is triggered upon receipt of the specified event or events, the event process (eventd) invokes the configured commands to update the candidate configuration, which is then committed, provided that no commit errors occur. The eventd process executes the configuration commands in the order in which they appear in the event policy configuration.
<b>Options</b>	<b>command</b> —Configuration mode command to be executed. Enclose each command in quotation marks (" "), and specify the complete statement path to the element, identifier, or value as you do in configuration mode when issuing commands at the <a href="#">[edit]</a> hierarchy level.
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Event Policy to Change the Configuration on page 558</a></li> <li>• <a href="#">Example: Changing the Configuration Using an Event Policy on page 559</a></li> <li>• <a href="#">Example: Changing the Interface Configuration in Response to an Event on page 566</a></li> <li>• <a href="#">change-configuration (Event Policy) on page 772</a></li> </ul>

## commands (Event Policy Execute Commands)

---

Syntax	<pre>commands {     "command"; }</pre>
Hierarchy Level	[edit <a href="#">event-options</a> <a href="#">policy</a> <i>policy-name</i> then <a href="#">execute-commands</a> ]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the operational mode commands to be issued when the corresponding event policy is triggered. When the event policy is triggered, the event process (eventd) executes the commands in the order in which they appear in the event policy configuration.
Options	<p><b>command</b>—Operational mode command to be executed. Enclose each command in quotation marks (" ").</p> <p>You can include variables in commands. The eventd process replaces each variable with values contained in the events that trigger the policy. You can use command variables of the following forms:</p> <ul style="list-style-type: none"><li>• <b>`\${attribute-name}</b>—The double dollar sign (\$\$) notation represents the event that is triggering a policy. When combined with an attribute name, the variable resolves to the value of the attribute associated with the triggering event.</li><li>• <b>`\${event.attribute-name}</b>—The dollar sign with the event name (<b>`\${event}`</b>) notation represents the most recent event that matches <b>event</b>. When combined with an attribute name, the variable resolves to the value of the attribute associated with that event.</li><li>• <b>`\${*attribute-name}</b>—The dollar sign with the asterisk (\$*) notation represents the most recent event that matches any of the correlating events. The variable resolves to the value of the attribute associated with most recent event that matches any of the correlated events specified in the policy configuration.</li></ul>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li><li>• <a href="#">Representing the Correlating Event in an Event Policy on page 544</a></li></ul>

## commit-options

<b>Syntax</b>	<pre>commit-options {   check &lt;synchronize&gt;;   force;   log "comment-string";   synchronize; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then change-configuration</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Customize the commit options for configuration updates made through an event policy. The <b>check</b> statement and the other <b>commit-options</b> statements are mutually exclusive.
<b>Options</b>	<p><b>check &lt;synchronize&gt;</b>—Verify that the candidate configuration is syntactically correct, but do not commit the changes. On dual control plane systems, when the <b>check synchronize</b> statement is configured, the candidate configuration on one control plane is copied to the other control plane, and the system verifies that both candidate configurations are syntactically correct. The <b>check</b> statement and the other <b>commit-options</b> statements are mutually exclusive.</p> <p><b>force</b>—Force the commit on the other Routing Engine, ignoring any warnings. By default, the <b>synchronize</b> command does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines by using the <b>force</b> option.</p> <p><b>log "comment-string"</b>—Include a comment describing changes to the committed configuration. Enclose the comment in quotation marks and include it on a single line. To view commit comments, issue the <b>show system commit</b> operational mode command.</p> <p><b>synchronize</b>—Synchronize the commit on both Routing Engines. The Routing Engine on which you execute this command copies and loads its candidate configuration to the other Routing Engine. Both Routing Engines perform a syntax check on the candidate configuration file. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Event Policy to Change the Configuration on page 558</a></li> <li>• <a href="#">Example: Changing the Configuration Using an Event Policy on page 559</a></li> <li>• <a href="#">Example: Changing the Interface Configuration in Response to an Event on page 566</a></li> <li>• <a href="#">change-configuration (Event Policy) on page 772</a></li> <li>• <a href="#">commands (Event Policy Change Configuration) on page 773</a></li> </ul>

- [retry \(Event Policy\) on page 793](#)
- [user-name \(Event Policy\) on page 805](#)

---

## destination (Event Policy)

---

Syntax	<pre>destination <i>destination-name</i> {     <i>retry-count count</i> retry-interval <i>seconds</i>;     <i>transfer-delay seconds</i>; }</pre>
Hierarchy Level	<pre>[edit event-options policy <i>policy-name</i> then event-script <i>filename</i>], [edit event-options policy <i>policy-name</i> then execute-commands]</pre>
Release Information	<p>Statement introduced in Junos OS Release 7.5.</p> <p>Support at the <code>[edit event-options policy <i>policy-name</i> then event-script <i>filename</i>]</code> hierarchy level introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Assign a location to which to upload command or script output for the specified event policy.
Options	<p><b><i>destination-name</i></b>—Name of a destination defined in the <b>destinations</b> statement at the <code>[edit event-options]</code> hierarchy level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li><li>• <a href="#">Executing Event Scripts in an Event Policy on page 573</a></li><li>• <a href="#">destinations (Event Policy) on page 777</a></li></ul>



## destinations (Event Policy)

<b>Syntax</b>	<pre> destinations {   destination-name {     archive-sites {       url &lt;password password&gt;;     }     transfer-delay seconds;   } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Define one or more destinations, each with a unique name. You can reference a destination within an event policy to indicate where to upload command output and various files, such as system log files and core files.
<b>Options</b>	<p><b><i>destination-name</i></b>—User-defined name of a destination.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Defining Destinations for File Archiving by Event Policies on page 592</a></li> </ul>

## **equals (Event Policy)**

---

<b>Syntax</b>	<i>event1.attribute-name equals event2.attribute-name;</i>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">policy</a> <i>policy-name</i> <a href="#">attributes-match</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an event policy to trigger only if the specified attribute of <b>event1</b> equals the specified attribute of <b>event2</b> .
<b>Options</b>	<i>event1.attribute-name</i> —Attribute of one event.  <i>event2.attribute-name</i> —Attribute of another event.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li></ul>

## event-options

```

Syntax  event-options {
        destinations {
            destination-name {
                archive-sites {
                    url <password password>;
                }
                transfer-delay seconds;
            }
        }
        event-script {
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
            file filename {
                checksum (md5 | sha-256 | sha1) hash;
                dampen {
                    dampen-options {
                        cpu-factor cpu-factor;
                        line-interval line-interval;
                        time-interval time-interval;
                    }
                }
                python-script-user username;
                refresh;
                refresh-from url;
                remote-execution {
                    remote-hostname {
                        passphrase user-password;
                        username user-login;
                    }
                }
                source url;
            }
            max-datasize size;
            optional;
            refresh;
            refresh-from url;
            traceoptions {
                file <filename> <files number> <size size> <world-readable | no-world-readable>;
                flag flag;
                no-remote-trace;
            }
        }
        generate-event {
            event-name (time-interval seconds | time-of-day hh:mm:ss) <no-drift>;
        }
        max-policies policies;
        policy policy-name {

```

```
attributes-match {
    event1.attribute-name equals event2.attribute-name;
    event.attribute-name matches regular-expression;
    event1.attribute-name starts-with event2.attribute-name;
}
events [events];
then {
    change-configuration {
        commands {
            "command";
        }
        commit-options {
            check <synchronize>;
            force;
            log "comment-string";
            synchronize;
        }
        retry count number interval seconds;
        user-name username;
    }
    event-script filename {
        arguments {
            argument-name argument-value;
        }
        destination destination-name {
            retry-count number retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    execute-commands {
        commands {
            "command";
        }
        destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    ignore;
    priority-override {
        facility facility-type;
        severity severity-level;
    }
    raise-trap;
    upload filename (filename | committed) destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
        user-name username;
    }
}
```

```

    within seconds {
        events [ events ];
        not events [ events ];
        trigger (after number | on number | until number);
    }
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
}

```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 7.5.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure event policies and event scripts.

The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

## event-script (Event Policy)

```

Syntax event-script filename {
    arguments {
        argument-name argument-value;
    }
    destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
}

```

**Hierarchy Level** [edit [event-options policy policy-name then](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure an event policy to execute one or more event scripts when triggered.

The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Executing Event Scripts in an Event Policy on page 573](#)
- [Configuring an Event Policy to Pass Arguments to an Event Script on page 576](#)

## events (Event Policy Trigger Events)

---

<b>Syntax</b>	<code>events [ <i>events</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">event-options policy</a> <i>policy-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure one or more events that will trigger the event policy. If one or more of the listed events occurs, the policy is executed.
<b>Options</b>	<code>[ <i>events</i> ]</code> —List of events. Events can be internally generated, or they can be generated by Junos OS processes.
<b>Required Privilege Level</b>	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li><li>• <a href="#">Example: Correlating Events Based on Event Attributes on page 546</a></li></ul>

## events (Event Policy Correlated Events)

---

<b>Syntax</b>	<code>events [ <i>events</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">event-options policy</a> <i>policy-name</i> <a href="#">within</a> <i>seconds</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure events that must occur within a specified time interval before one of the trigger events in order for the event policy to be triggered.
<b>Options</b>	<code>[ <i>events</i> ]</code> —List of events. Events can be internally generated, or they can be generated by Junos OS processes.
<b>Required Privilege Level</b>	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li></ul>

## execute-commands

<b>Syntax</b>	<pre> execute-commands {   commands {     "command";   }   destination destination-name {     retry-count count retry-interval seconds;     transfer-delay seconds;   }   output-filename filename;   output-format (text   xml);   user-name username; } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Specify the operational mode commands to be issued when the corresponding event policy is triggered. You can specify the format of the command output, and a filename and destination for the output file.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li> </ul>

## facility

---

<b>Syntax</b>	<code>facility <i>facility-type</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy <i>policy-name</i> then priority-override</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Within an event policy, override the default facility type of the triggering event so that the event is logged based on the configured facility type.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding the Event System Log Priority in an Event Policy on page 583</a></li><li>• <a href="#">Example: Configuring the Event System Log Priority in an Event Policy on page 584</a></li><li>• <a href="#">Junos OS System Logging Facilities and Message Severity Levels</a></li><li>• <a href="#">priority-override (Event Policy) on page 792</a></li><li>• <a href="#">severity (Event Policy) on page 795</a></li></ul>

## generate-event

---

<b>Syntax</b>	<pre>generate-event {     <i>event-name</i> ( <a href="#">time-interval seconds</a>   <a href="#">time-of-day hh:mm:ss</a> ) &lt;<a href="#">no-drift</a>&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Generate an internal event, based on a time interval or the time of day. You can configure up to 10 internal events.
<b>Options</b>	<b><i>event-name</i></b> —Name of an internally generated event.  The statements are explained separately.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Generating Internal Events to Trigger Event Policies on page 551</a></li><li>• <a href="#">no-drift on page 786</a></li><li>• <a href="#">time-interval (Generate Events) on page 798</a></li><li>• <a href="#">time-of-day on page 799</a></li></ul>



## ignore

---

<b>Syntax</b>	<code>ignore;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define a policy that ignores particular events. If one or more of the listed events occur, a system log message for the event is not generated, and no further policies associated with this event are processed. If you include the <b>ignore</b> statement in a policy configuration, you cannot configure any other actions in the policy.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Event Policies to Ignore an Event on page 578</a></li> </ul>

## matches (Event Policy)

---

<b>Syntax</b>	<code>event.attribute-name matches regular-expression;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name attributes-match</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure an event policy to trigger only if the specified attribute of <b>event</b> matches the regular expression.
<b>Options</b>	<p><b>event.attribute-name</b>—Event attribute to compare to a regular expression.</p> <p><b>regular-expression</b>—Regular expression to compare.</p>
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li> <li>• <a href="#">Using Regular Expressions to Refine the Set of Events That Trigger a Policy on page 549</a></li> </ul>

## max-policies

---

<b>Syntax</b>	max-policies <i>policies</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3.
<b>Description</b>	Define the maximum number of event policies that can run concurrently on the device.
<b>Options</b>	<i>policies</i> —Maximum number of event policies that can run concurrently. <b>Range:</b> 0 through 20 <b>Default:</b> 15
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">max-datasize on page 720</a></li><li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li><li>• <a href="#">Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695</a></li></ul>

## no-drift (Generate Events)

---

<b>Syntax</b>	no-drift;
<b>Hierarchy Level</b>	[edit event-options <a href="#">generate-event</a> <i>event-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Do not propagate the delay caused in triggering an event to the triggering of the next event.
<b>Default</b>	Disabled.
<b>Required Privilege Level</b>	maintenance - To view this statement in the configuration. maintenance-control - To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">generate-event on page 784</a></li></ul>

## not (Event Policy)

---


<b>Syntax</b>	not events [ <i>events</i> ];
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy</a> <i>policy-name</i> <a href="#">within seconds</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an event policy to trigger only if the events in the <b>not</b> statement do not occur within a specified time interval before one of the trigger events.
<b>Options</b>	[ <i>events</i> ]—List of events. Events can be internally generated, or they can be generated by Junos OS processes.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li></ul>

## output-filename (Event Policy)

---

<b>Syntax</b>	<code>output-filename <i>filename</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options policy <i>policy-name</i> then event-script <i>filename</i>],</code> <code>[edit event-options policy <i>policy-name</i> then execute-commands]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Support at the <code>[edit event-options policy <i>policy-name</i> then event-script <i>filename</i>]</code> hierarchy level introduced in Junos OS Release 7.6.
<b>Description</b>	<p>Define a descriptive string that will be included in the filename of the output file in which to write command or script output. For op scripts, this statement is optional.</p> <p>The name of the output file also includes the hostname and timestamp to ensure that the each filename is unique. If an event policy is triggered multiple times in a 1-second period, an index number is appended to the filename. The index number range is 001 through 999.</p> <p>Prior to Junos OS Release 14.1R3, the filename has the following naming convention:</p> <p style="text-align: center;"><i>hostname_filename_YYYYMMDD_HHMMSS_index-number</i></p> <p>Starting in Junos OS Release 14.1R3, the <i>filename</i> string is placed after the timestamp.</p> <p style="text-align: center;"><i>hostname_YYYYMMDD_HHMMSS_filename_index-number</i></p>
<b>Options</b>	<i>filename</i> —String included in the filename of the output file.
<b>Required Privilege Level</b>	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li><li>• <a href="#">Executing Event Scripts in an Event Policy on page 573</a></li></ul>

## output-format (Event Policy)

<b>Syntax</b>	output-format (text   xml);
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then event-script filename</a> ], [edit <a href="#">event-options policy policy-name then execute-commands</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Support at the [edit <a href="#">event-options policy policy-name then event-script filename</a> ] hierarchy level introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Specify the format (ASCII text or XML) for the output of event scripts or commands executed by an event policy.
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> Event script output is always emitted in XML format even if you configure the <code>output-format text</code> statement under the [edit <a href="#">event-options policy policy-name then event-script filename</a>] hierarchy.</p> </div> </div>	
<b>Default</b>	If you do not include the <b>output-format</b> statement, the output is emitted in XML format.
<b>Options</b>	<b>text</b> —Formatted ASCII text.  <b>xml</b> —Junos Extensible Markup Language (XML) tags.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li> <li>• <a href="#">Executing Event Scripts in an Event Policy on page 573</a></li> </ul>

## policy (Event Policy)

```
Syntax  policy policy-name {
        attributes-match {
            event1.attribute-name equals event2.attribute-name;
            event.attribute-name matches regular-expression;
            event1.attribute-name starts-with event2.attribute-name;
        }
        events [ events ];
        then {
            ... the then subhierarchy appears at the end of the [edit event-options policy policy-name]
            hierarchy level ...
        }
        within seconds {
            events [ events ];
            not events [ events ];
            trigger (on | after | until) event-count;
        }

        then {
            change-configuration {
                commands {
                    "command";
                }
                commit-options {
                    check <synchronize>;
                    force;
                    log "comment-string";
                    synchronize;
                }
                retry count number interval seconds;
                user-name username;
            }
            event-script filename {
                arguments {
                    argument-name argument-value;
                }
                destination destination-name {
                    retry-count count retry-interval seconds;
                    transfer-delay seconds;
                }
                output-filename filename;
                output-format (text | xml);
                user-name username;
            }
            execute-commands {
                commands {
                    "command";
                }
                destination destination-name {
                    retry-count count retry-interval seconds;
                    transfer-delay seconds;
                }
                output-filename filename;
            }
        }
    }
```

```

    output-format (text | xml);
    user-name username;
}
ignore;
priority-override {
    facility facility-type;
    severity severity-level;
}
raise-trap;
upload filename (filename | committed) destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
    user-name username;
}
}
}

```

<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	<p>Define an event policy to be processed by the eventd process. If you configure a policy, the <b>events</b> and <b>then</b> statements are mandatory.</p> <p>You can configure multiple policies to be processed for an event. The policies are executed in the order in which they appear in the configuration. If you configure more than one policy for an event, and if one of the policies is configured to ignore the event, no policies that follow the policy with the <b>ignore</b> statement are executed.</p>
<b>Default</b>	If you do not configure a policy for an event, the event is recorded in the system log.
<b>Options</b>	<p><b>policy-name</b>—Name of an event policy.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>

## priority-override (Event Policy)

---

<b>Syntax</b>	<pre>priority-override {     facility facility-type;     severity severity-level; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	<p>Configure an event policy to override the default system log priority of the triggering event so that the system logs the event with a different facility type, severity level, or both. If you configure multiple event policies to override the priority of the same event, the event is logged based on the priority set by the last executed event policy to change it.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding the Event System Log Priority in an Event Policy on page 583</a></li><li>• <a href="#">Example: Configuring the Event System Log Priority in an Event Policy on page 584</a></li><li>• <i>Junos OS System Logging Facilities and Message Severity Levels</i></li><li>• <a href="#">facility on page 784</a></li><li>• <a href="#">severity (Event Policy) on page 795</a></li></ul>

## raise-trap (Event Policy)

---

<b>Syntax</b>	<pre>raise-trap;</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name then</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	<p>Define an event policy that raises an SNMP trap in response to an event. If one or more of the listed events occur, the system log message for the event is converted into a trap. This enables an agent to notify a trap-based network management system (NMS) of significant events.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Overview of Using Event Policies to Raise SNMP Traps on page 580</a></li><li>• <a href="#">Example: Raising an SNMP Trap in Response to an Event on page 581</a></li></ul>



## retry (Event Policy)

<b>Syntax</b>	<code>retry count <i>number</i> interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy <i>policy-name</i> then change-configuration</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Specify the number of times that Junos OS attempts the <b>change-configuration</b> event policy action if the initial attempt fails while acquiring a lock on the configuration database. If you include the <b>retry</b> statement, you must configure both the <b>count</b> and <b>interval</b> statements.
<b>Default</b>	If you do not include the <b>retry</b> statement, and the <b>change-configuration</b> event policy action fails, the configuration changes specified in the event policy are not implemented or committed.
<b>Options</b>	<p><b>count <i>number</i></b>—The number of attempts to retry the <b>change-configuration</b> event policy action upon failure of the initial attempt.</p> <p><b>Range:</b> 0 through 10</p> <p><b>Default:</b> 0</p> <p><b>interval <i>seconds</i></b>—The time interval specified in seconds between retry attempts.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Event Policy to Change the Configuration on page 558</a></li> <li>• <a href="#">Example: Changing the Configuration Using an Event Policy on page 559</a></li> <li>• <a href="#">Example: Changing the Interface Configuration in Response to an Event on page 566</a></li> <li>• <a href="#">change-configuration (Event Policy) on page 772</a></li> <li>• <a href="#">commands (Event Policy Change Configuration) on page 773</a></li> </ul>

## retry-count (Event Policy)

---

<b>Syntax</b>	<code>retry-count <i>number</i> retry-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit <b>event-options</b> <b>policy</b> <i>policy-name</i> <b>then</b> <b>event-script</b> <i>filename</i> <b>destination</b> <i>destination-name</i>],</code> <code>[edit <b>event-options</b> <b>policy</b> <i>policy-name</i> <b>then</b> <b>execute-commands</b> <b>destination</b> <i>destination-name</i>],</code> <code>[edit <b>event-options</b> <b>policy</b> <i>policy-name</i> <b>then</b> <b>upload</b> <i>filename</i> (<i>filename</i>   committed) <b>destination</b> <i>destination-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an event policy to retry a file upload operation if the first attempt fails.
<b>Default</b>	If you do not include this statement, the file upload operation is attempted one time only.
<b>Options</b>	<i>number</i> —Number of retries.  <i>retry-interval seconds</i> —Length of time to wait between retries.
<b>Required Privilege Level</b>	<i>maintenance</i> —To view this statement in the configuration. <i>maintenance-control</i> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Event Policy to Retry the File Upload Action on page 605</a></li></ul>

## severity (Event Policy)

<b>Syntax</b>	<code>severity severity-level;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy policy-name</a> then <a href="#">priority-override</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Configure an event policy to override the preassigned severity level of a triggering event so that the event is logged based on the configured severity level.
<b>Options</b>	<b>severity-level</b> —Severity level logged for the triggering event. <a href="#">Table 52 on page 795</a> lists the possible severity levels.

**Table 52: System Log Message Severity Levels**

Severity Level	Description
emergency	System panic or other conditions that cause the routing platform to stop functioning
alert	Conditions that require immediate correction, such as a corrupted system database
critical	Critical conditions, such as hard disk errors
error	Error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels
warning	Conditions that warrant monitoring
notice	Conditions that are not errors but might warrant special handling
info	Events or non-error conditions of interest

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Event System Log Priority in an Event Policy on page 583</a></li> <li>• <a href="#">Example: Configuring the Event System Log Priority in an Event Policy on page 584</a></li> <li>• <i>Junos OS System Logging Facilities and Message Severity Levels</i></li> <li>• <a href="#">facility on page 784</a></li> <li>• <a href="#">priority-override (Event Policy) on page 792</a></li> </ul>

## starts-with (Event Policy)

---

<b>Syntax</b>	<code>event1.attribute-name starts-with event2.attribute-name;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">policy</a> <i>policy-name</i> <a href="#">attributes-match</a> <i>event1.attribute-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Configure an event policy to trigger only if the specified attribute of <b>event1</b> starts with the specified attribute of <b>event2</b> .
<b>Options</b>	<b>event1.attribute-name</b> —Attribute of one event. <b>event2.attribute-name</b> —Attribute of another event.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li></ul>

## then (Event Policy)

```
Syntax  then {
    change-configuration {
        commands {
            "command";
        }
        commit-options {
            check <synchronize>;
            force;
            log "comment-string";
            synchronize;
        }
        retry count number interval seconds;
        user-name username;
    }
    event-script filename {
        arguments {
            argument-name argument-value;
        }
        destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    execute-commands {
        commands {
            "command";
        }
        destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
    }
    ignore;
    priority-override {
        facility facility-type;
        severity severity-level;
    }
    raise-trap;
    upload filename (filename | committed) destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
        user-name username;
    }
}
```

Hierarchy Level [edit `event-options policy policy-name`]

<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define event policy actions that are executed if an event occurs. For each event policy, you can configure multiple actions.  The statements are explained separately.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring an Event Policy to Upload Files on page 595</a></li><li>• <a href="#">Configuring an Event Policy to Execute Operational Mode Commands on page 555</a></li><li>• <a href="#">Executing Event Scripts in an Event Policy on page 573</a></li><li>• <a href="#">Configuring Event Policies to Ignore an Event on page 578</a></li><li>• <a href="#">Overview of Using Event Policies to Raise SNMP Traps on page 580</a></li></ul>

---

## time-interval (Generate Events)

---

<b>Syntax</b>	time-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">generate-event</a> <i>event-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a frequency at which to generate a particular event.
<b>Options</b>	<b><i>seconds</i></b> —Time interval between internally generated events. <b>Range:</b> 60 through 2,592,000 seconds
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Generating Internal Events to Trigger Event Policies on page 551</a></li><li>• <a href="#">generate-event on page 784</a></li><li>• <a href="#">time-of-day on page 799</a></li></ul>

## time-of-day

---

<b>Syntax</b>	<code>time-of-day <i>hh:mm:ss</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options</a> <a href="#">generate-event</a> <i>event-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure a time of day at which to generate a particular event.
<b>Options</b>	<i>hh:mm:ss</i> —Time of day at which to generate an event.
<b>Required Privilege Level</b>	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Generating Internal Events to Trigger Event Policies on page 551</a></li><li>• <a href="#">generate-event on page 784</a></li><li>• <a href="#">time-interval (Generate Events) on page 798</a></li></ul>

## traceoptions (Event Options)

---

Syntax	<pre>traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;size size&gt; &lt;world-readable           no-world-readable&gt;;     flag flag;     no-remote-trace; }</pre>
Hierarchy Level	[edit <a href="#">event-options</a> ]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define tracing operations for event policies.
Default	If you do not include this statement, no event-policy-specific tracing operations are performed.
Options	<p><b>file filename</b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>. By default, event policy process tracing output is placed in the file <b>eventd</b>. If you include the <b>file</b> statement, you must specify a filename. To retain the default, you can specify <b>eventd</b> as the filename.</p> <p><b>Default:</b> <b>/var/log/eventd</b></p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed and compressed to <b>trace-file.0.gz</b>. When <b>trace-file</b> again reaches its maximum size, <b>trace-file.0.gz</b> is renamed <b>trace-file.1.gz</b> and <b>trace-file</b> is renamed and compressed to <b>trace-file.0.gz</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Log all operations</li><li>• <b>configuration</b>—Log reading of configuration at the <b>[edit event-options]</b> hierarchy level</li><li>• <b>events</b>—Log eventd processing</li><li>• <b>database</b>—Log events involving storage and retrieval in events database</li><li>• <b>server</b>—Log communication with processes that are generating events</li><li>• <b>timer-events</b>—Log internally generated events</li></ul>



**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-world-readable**—Restrict file access to owner. This is the default.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file.0.gz***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0.gz*** is renamed ***trace-file.1.gz*** and ***trace-file*** is renamed and compressed to ***trace-file.0.gz***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
---------------------------------	---

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Event Policy Processing on page 631</a></li></ul>
------------------------------	---

## transfer-delay (Event Policy)

---

Syntax	<code>transfer-delay seconds;</code>
Hierarchy Level	<code>[edit event-options destinations destination-name],</code> <code>[edit event-options policy policy-name then event-script filename</code> <code>destination destination-name],</code> <code>[edit event-options policy policy-name then execute-commands</code> <code>destination destination-name],</code> <code>[edit event-options policy policy-name then upload filename (filename   committed)</code> <code>destination destination-name]</code>
Release Information	Statement introduced in Junos OS Release 7.5. Support at the <code>[edit event-options policy policy-name then ...]</code> hierarchy levels introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a delay before transferring files. This allows the files to be completely generated before the upload starts. If you configure a transfer delay at the <code>[edit event-options destination destination-name]</code> hierarchy level and at one of the <code>[edit event-options policy policy-name then ...]</code> hierarchy levels, the resulting delay is the sum of the two delays.
Default	If you do not include this statement, there is no transfer delay.
Options	<i>seconds</i> —Duration of the delay before files are uploaded.
Required Privilege Level	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Example: Defining Destinations for File Archiving by Event Policies on page 592</a></li><li>• <a href="#">Configuring the Delay Before Files Are Uploaded by an Event Policy on page 602</a></li></ul>

## trigger

---


<b>Syntax</b>	<code>trigger (on   after   until) <i>event-count</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options policy</a> <i>policy-name</i> <a href="#">within seconds</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure an event policy to be triggered if an event or set of events occurs <i>event-count</i> times within a specified time period.
<b>Default</b>	If you do not include this statement, the policy is executed on receipt of the first configured event.
<b>Options</b>	<p><b>after <i>event-count</i></b>—The policy is executed when the number of matching events received equals <i>event-count</i> + 1.</p> <p><b>on <i>event-count</i></b>—The policy is executed when the number of matching events received equals <i>event-count</i>.</p> <p><b>until <i>event-count</i></b>—The policy is executed each time a matching event is received and stops being executed when the number of matching events received equals <i>event-count</i>.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Triggering an Event Policy Based on Event Count on page 547</a></li> </ul>

## upload

---

Syntax	<pre>upload filename (<i>filename</i>   committed) destination <i>destination-name</i> {     <i>retry-count</i> <i>count</i> retry-interval <i>seconds</i>;     <i>transfer-delay</i> <i>seconds</i>;     <i>user-name</i> <i>username</i>; }</pre>
Hierarchy Level	[edit <b>event-options</b> <i>policy</i> <i>policy-name</i> <b>then</b> ]
Release Information	Statement introduced in Junos OS Release 7.5. <b>committed</b> option for <b>filename</b> statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure an event policy to upload the committed configuration file or other specified files to a destination when triggered.
Options	<p><b>destination</b> <i>destination-name</i>—Name of the destination for the uploaded file. The destination must be defined in the <b>destinations</b> statement at the [edit <b>event-options</b>] hierarchy level.</p> <p><b>filename</b> (<i>filename</i>   <b>committed</b>)—Name of the file to upload. Specify either the word <b>committed</b> to upload the most recently committed configuration file, or the filename of another file.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">destinations (Event Policy) on page 777</a></li><li>• <a href="#">Example: Configuring an Event Policy to Upload Files on page 595</a></li></ul>

## user-name (Event Policy)

<b>Syntax</b>	<code>user-name username;</code>
<b>Hierarchy Level</b>	<code>[edit event-options policy policy-name then change-configuration],</code> <code>[edit event-options policy policy-name then event-script filename],</code> <code>[edit event-options policy policy-name then execute-commands],</code> <code>[edit event-options policy policy-name then upload filename (filename   committed)</code> <code>destination destination-name]</code>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support at the <code>[edit event-options policy policy-name then change-configuration]</code> hierarchy level introduced in Junos OS Release 12.1.</p>
<b>Description</b>	Associate a user with an action in an event policy. The event policy action is executed under the privileges of the associated user.
<div>  <p><b>NOTE:</b> When you configure the <code>user-name</code> statement for event scripts, it only applies to SLAX and XSLT scripts. This statement has no effect when configured for Python event scripts.</p> <p>To execute a Python event script under the access privileges of a specific user, configure the <code>python-script-user username</code> statement at the <code>[edit event-options event-script file filename]</code> hierarchy level. If you do not configure the <code>python-script-user</code> statement for a Python event script, Junos OS executes the script using the access privileges of the generic, unprivileged system account <code>nobody</code>.</p> </div>	
<b>Default</b>	If you do not associate a user with an action, the action is executed as user <code>root</code> .
<b>Options</b>	<code>username</code> —Username that is configured at the <code>[edit system login]</code> hierarchy level.
<b>Required Privilege Level</b>	<p><code>maintenance</code>—To view this statement in the configuration.</p> <p><code>maintenance-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Changing the User Privilege Level for an Event Policy Action on page 609</a></li> <li>• <a href="#">python-script-user (Event and SNMP Scripts) on page 817</a></li> </ul>

## within

---

<b>Syntax</b>	<pre>within seconds {     events [ events ];     not events [ events ];     trigger (after   on   until) event-count; }</pre>
<b>Hierarchy Level</b>	[edit <b>event-options policy</b> <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	<p>Configure events that must (or must not) occur within a specified time interval before one of the trigger events in order for the event policy to be triggered.</p> <p>The statements are explained separately.</p>
<b>Options</b>	<p><b>seconds</b>—Time interval between a correlated event and a trigger event.</p> <p><b>Range:</b> 60 through 604,800 seconds</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using Correlated Events to Trigger an Event Policy on page 541</a></li></ul>

## CHAPTER 50

# Event Script Configuration Statements

- [checksum](#) on page 808
- [cpu-factor \(Dampen Options\)](#) on page 809
- [dampen](#) on page 810
- [dampen-options](#) on page 811
- [event-script \(Event Options\)](#) on page 812
- [file \(Event Scripts\)](#) on page 813
- [line-interval \(Dampen Options\)](#) on page 814
- [max-datasize](#) on page 815
- [optional \(Event Scripts\)](#) on page 816
- [python-script-user \(Event and SNMP Scripts\)](#) on page 817
- [refresh \(Event Scripts\)](#) on page 818
- [refresh-from \(Event Scripts\)](#) on page 819
- [remote-execution](#) on page 820
- [source \(Event Scripts\)](#) on page 821
- [time-interval \(Dampen Options\)](#) on page 822
- [traceoptions \(Event Scripts\)](#) on page 823

## checksum

---


Syntax	<code>checksum (md5   sha-256   sha1) hash;</code>
Hierarchy Level	<code>[edit event-options event-script file filename],</code> <code>[edit system scripts commit file filename],</code> <code>[edit system scripts op file filename],</code> <code>[edit system scripts snmp file filename],</code> <code>[edit system extensions extension-service application file filename]</code>
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced for the <code>[edit system extensions extension-service application file filename]</code> hierarchy level in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.
Description	For Junos OS commit scripts, event scripts, op scripts, SNMP scripts, and scripts developed using the Juniper Extension Toolkit (JET) specify the MD5, SHA-1, or SHA-256 checksum hash. When Junos OS executes a local commit, event, op, or SNMP script, the system verifies the authenticity of the script by using the configured checksum hash.
Options	<code>md5 hash</code> —MD5 checksum of this script.  <code>sha-256 hash</code> —SHA-256 checksum of this script.  <code>sha1 hash</code> —SHA-1 checksum of this script.
Required Privilege Level	<code>maintenance</code> —To view this statement in the configuration. <code>maintenance-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li><li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li><li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li><li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li><li>• <a href="#">file checksum md5 on page 844</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha-256 on page 845</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha1 on page 846</a> command in the <i>System Basics and Services Command Reference</i></li></ul>



## cpu-factor (Dampen Options)

<b>Syntax</b>	<code>cpu-factor <i>cpu-factor</i>;</code>
<b>Hierarchy Level</b>	<p>[edit event-options event-script dampen <a href="#">dampen-options</a>],  [edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],  [edit system scripts commit dampen <a href="#">dampen-options</a>],  [edit system scripts op dampen <a href="#">dampen-options</a>],  [edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Slow down the operation of scripts executed on the device based on how heavily the CPU is being used.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <b>cpu-factor</b> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b>cpu-factor</b>—Specify the CPU use factor at which script execution is slowed down.</p> <p><b>Range:</b> 0 through 100</p> <p><b>Default:</b> 20</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Dampening Script Execution on page 697</a></li> </ul>

## dampen

<b>Syntax</b>	<pre>dampen {   dampen-options {     cpu-factor <i>cpu-factor</i>;     line-interval <i>line-interval</i>;     time-interval <i>time-interval</i>;   } }</pre>
<b>Hierarchy Level</b>	<p>[edit event-options <a href="#">event-script</a>],  [edit event-options event-script <a href="#">file filename</a>],  [edit system scripts <a href="#">commit</a>],  [edit system scripts <a href="#">op</a>],  [edit system scripts op <a href="#">file filename</a>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Slow down the operation of SLAX and XSLT scripts executed on the device.</p> <p>To dampen all enabled scripts of a given script type, configure the statement at the hierarchy level for that script type. To dampen a single op or event script, configure the statement at the hierarchy level where the script is configured. Script dampening is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <b>op url</b> command, Junos OS does not dampen the script. You cannot dampen individual commit scripts.</p>
	<div>  <p><b>NOTE:</b> Junos OS does not support script dampening for Python automation scripts. If you configure script dampening for all enabled scripts of a given script type, Junos OS logs the following message in the system log file if it executes a Python op or event script: <b>dampen is not supported for python scripts</b>. No warning is logged for Python commit scripts.</p> </div>
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.  maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Dampening Script Execution on page 697</a></li> <li>• <a href="#">dampen-options on page 730</a></li> <li>• <a href="#">event-script on page 812</a></li> </ul>

## dampen-options

---

<b>Syntax</b>	<pre>dampen-options {   <b>cpu-factor</b> <i>cpu-factor</i>;   <b>line-interval</b> <i>line-interval</i>;   <b>time-interval</b> <i>time-interval</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit event-options event-script <b>dampen</b>], [edit event-options event-script file <i>filename</i> <b>dampen</b>], [edit system scripts commit <b>dampen</b>], [edit system scripts op <b>dampen</b>], [edit system scripts op file <i>filename</i> <b>dampen</b>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	Slow down the operation of automation scripts executed on the device.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Dampening Script Execution on page 697</a></li> </ul>

## event-script (Event Options)

```
Syntax event-script {
    dampen {
        dampen-options {
            cpu-factor cpu-factor;
            line-interval line-interval;
            time-interval time-interval;
        }
    }
    file filename {
        checksum (md5 | sha-256 | sha1) hash;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        python-script-user username;
        refresh;
        refresh-from url;
        remote-execution {
            remote-hostname {
                passphrase user-password;
                username user-login;
            }
        }
        source url;
    }
    max-datasize size;
    optional;
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

**Hierarchy Level** [edit [event-options](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** For Junos OS event scripts, configure scripting mechanisms.  
  
The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

- Related Documentation**
- [Storing and Enabling Scripts on page 661](#)

## file (Event Scripts)

**Syntax**

```
file filename {
  checksum (md5 | sha-256 | sha1) hash;
  dampen {
    dampen-options {
      cpu-factor cpu-factor;
      line-interval line-interval;
      time-interval time-interval;
    }
  }
  python-script-user username;
  refresh;
  refresh-from url;
  remote-execution {
    remote-hostname {
      passphrase user-password;
      username user-login;
    }
  }
  source url;
}
```

**Hierarchy Level** [edit [event-options event-script](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Enable a Junos OS event script that is located in the `/var/db/scripts/event` directory on the hard disk or the `/config/scripts/event` directory on the flash drive. After enabling an event script, an event policy can execute the event script when triggered.

**Options** *filename*—Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX) file containing an event script.

The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.


- Related Documentation**
- [Storing and Enabling Scripts on page 661](#)
  - [Enabling and Executing Event Scripts on page 626](#)

## line-interval (Dampen Options)

---

<b>Syntax</b>	<code>line-interval <i>line-interval</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options event-script dampen <a href="#">dampen-options</a>],</code> <code>[edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts commit dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the number of lines of scripts to execute before pausing.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <b>line-interval</b> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the default values are used.</p>
<b>Options</b>	<p><b><i>line-interval</i></b>—Specify the line interval at which to pause.</p> <p><b>Range:</b> 1 through 4,294,967,295</p> <p><b>Default:</b> 10</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>

## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
	<p> <b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p>
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>

- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

---



## optional (Event Scripts)

---

<b>Syntax</b>	optional;
<b>Hierarchy Level</b>	[edit <a href="#">event-options event-script</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	<p>For Junos OS event scripts, globally indicate the physical presence of the scripts as optional to allow commits to proceed if they are missing. Scripts can be missing, for example, if the primary storage on the device running Junos OS becomes corrupted.</p> <p>If not set, the event process (eventd) will validate the physical existence of event scripts during the commit process. This means that if a configured script is not physically accessible, the commit will fail.</p> <p>If set, the event process (eventd) ignores missing event scripts during the commit process and proceeds with the commit. If the event script is present, but invalid, the commit fails.</p>
<b>Default</b>	By default, optional is not set. Event scripts are required to be accessible during commit operations.
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">source on page 821</a></li><li>• <a href="#">event-script on page 812</a></li></ul>



## python-script-user (Event and SNMP Scripts)

<b>Syntax</b>	<code>python-script-user <i>username</i>;</code>
<b>Hierarchy Level</b>	[edit event-options event-script <i>file filename</i> ], [edit system scripts snmp <i>file filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1R3 for ACX Series, MX Series, PTX Series, QFX Series, and T Series.
<b>Description</b>	Define the user under whose access privileges the Python event or SNMP script executes. If you do not associate a user with the script, then by default, Junos OS executes the script with the privileges of the *nix user and group <b>nobody</b> , which is the generic, unprivileged system account.
<div>  <p><b>NOTE:</b> You cannot configure Python event and SNMP scripts to execute with root access privileges.</p> </div>	
<div>  <p><b>NOTE:</b> Junos OS executes Python commit and op scripts under the access privileges of the user who executes the command or operation that invokes the script.</p> </div>	
<b>Default</b>	If you do not include the <b>python-script-user</b> statement, Junos OS executes the script under the access privileges of the *nix user and group <b>nobody</b> .
<b>Options</b>	<i>username</i> —Username that is configured at the [edit system login] hierarchy level.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Requirements for Executing Python Automation Scripts on Devices Running Junos OS on page 169</a></li> <li>• <a href="#">Understanding Python Automation Scripts for Devices Running Junos OS on page 167</a></li> </ul>

## refresh (Event Scripts)

---

<b>Syntax</b>	refresh;
<b>Hierarchy Level</b>	[edit <a href="#">event-options event-script</a> ], [edit <a href="#">event-options event-script file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches.
<b>Description</b>	<p>Overwrite the local copy of all enabled event scripts or a single enabled event script with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li><li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li><li>• <a href="#">refresh-from (Event Scripts) on page 819</a></li><li>• <a href="#">source (Event Scripts) on page 821</a></li></ul>

## refresh-from (Event Scripts)

<b>Syntax</b>	<code>refresh-from <i>url</i>;</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">event-options event-script</a>],</code> <code>[edit <a href="#">event-options event-script file</a> <i>filename</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches.
<b>Description</b>	<p>Overwrite the local copy of all enabled event scripts or a single enabled event script with the copy located at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from <i>url</i></b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<b>Options</b>	<b><i>url</i></b> —Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using an Alternate Source Location for a Script on page 673</a></li> <li>• <a href="#">refresh (Event Scripts) on page 818</a></li> <li>• <a href="#">source (Event Scripts) on page 821</a></li> </ul>

## remote-execution

---

<b>Syntax</b>	<pre>remote-execution {     remote-hostname {         passphrase <i>user-password</i>;         username <i>user-login</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options event-script file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6.
<b>Description</b>	<p>Configure connection information for a local or remote host, which the event script can use to connect to and execute RPCs on the device.</p> <p>Remote execution details, which include the hostname, username, and passphrase for one or more hosts, enable an event script to invoke remote procedure calls on a host without encoding the connection information directly in the event script. When you include remote execution details in the configuration instead of in individual event scripts, it has the advantage of capturing the information in a single location and the passphrase is encrypted.</p> <p>When an event policy invokes an event script, it passes any configured remote execution details to the script. An event script references the remote execution details in the argument list of the function used to create the connection to the remote host.</p>
<b>Options</b>	<p><b>passphrase <i>user-password</i></b>—User's password for the host.</p> <p><b>remote-hostname</b>—Name of the host with which the event script will communicate.</p> <p><b>username <i>username</i></b>—User's login name for the host.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using RPCs and Operational Mode Commands in Event Scripts on page 619</a></li><li>• <a href="#">Using Event and Remote Execution Details in Event Scripts on page 617</a></li></ul>

## source (Event Scripts)

---

<b>Syntax</b>	<code>source url;</code>
<b>Hierarchy Level</b>	[edit <a href="#">event-options event-script file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.6 for EX Series switches.
<b>Description</b>	Specify the location of the master source file for an event script. When you issue the <b>set refresh</b> configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard disk.
<b>Options</b>	<i>url</i> —Master source file for an event script specified as an HTTP URL, FTP URL, or scp-style remote file specification.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li> <li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li> <li>• <a href="#">refresh (Event Scripts) on page 818</a></li> <li>• <a href="#">refresh-from (Event Scripts) on page 819</a></li> </ul>

## time-interval (Dampen Options)

---

<b>Syntax</b>	<code>time-interval <i>microseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options event-script dampen <a href="#">dampen-options</a>],</code> <code>[edit event-options event-script file <i>filename</i> dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts commit dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op dampen <a href="#">dampen-options</a>],</code> <code>[edit system scripts op file <i>filename</i> dampen <a href="#">dampen-options</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1.
<b>Description</b>	<p>Specify the length of time during which to pause the execution of one or more scripts.</p> <p>Script dampening is initiated when CPU use exceeds the value specified using the <a href="#">cpu-factor</a> statement. Junos OS then executes the number of lines of script specified using the <a href="#">line-interval</a> statement and then pauses script execution for the time specified using the <a href="#">time-interval</a> statement. If you do not configure values for these statements, the defaults are used.</p>
<b>Options</b>	<p><b><i>microseconds</i></b>—Specify the time in microseconds to pause script execution.</p> <p><b>Range:</b> 1 through 4,294,967,295 microseconds</p> <p><b>Default:</b> 100 microseconds</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>

## traceoptions (Event Scripts)

<b>Syntax</b>	<pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;     flag flag;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">event-options event-script</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Define tracing operations for event scripts.
<b>Default</b>	If you do not include this statement, the default operation during script execution is to log important events in the default trace file <b>/var/log/escrpt.log</b> . When the file reaches 128 kilobytes (KB), it is renamed and compressed.
<b>Options</b>	<p><b>file filename</b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>. By default, event script process tracing output is placed in the file <b>escrpt.log</b>. If you include the <b>file</b> statement, you must specify a filename. To retain the default, you can specify <b>escrpt.log</b> as the filename.</p> <p><b>Default:</b> <b>/var/log/escrpt.log</b></p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed and compressed to <b>trace-file.0.gz</b>. When <b>trace-file</b> again reaches its maximum size, <b>trace-file.0.gz</b> is renamed <b>trace-file.1.gz</b> and <b>trace-file</b> is renamed and compressed to <b>trace-file.0.gz</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>flag flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Log all operations</li> <li>• <b>events</b>—Log important events</li> <li>• <b>input</b>—Log event script input data</li> <li>• <b>offline</b>—Generate data for offline development</li> <li>• <b>output</b>—Log event script output data</li> <li>• <b>rpc</b>—Log event script RPCs</li> <li>• <b>xslt</b>—Log the XSLT library</li> </ul>

**no-world-readable**—Restrict file access to owner. This is the default.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—Enable unrestricted file access.

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration.
	maintenance-control—To add this statement to the configuration.

<b>Related Documentation</b>	• <a href="#">Tracing Event Script Processing on page 634</a>
	• <a href="#">Dampening Script Execution on page 697</a>



## CHAPTER 51

# SNMP Script Configuration Statements

- [checksum](#) on page 826
- [file \(SNMP Scripts\)](#) on page 827
- [max-datasize](#) on page 828
- [oid](#) on page 829
- [python-script-user \(Event and SNMP Scripts\)](#) on page 830
- [refresh \(SNMP Scripts\)](#) on page 831
- [refresh-from \(SNMP Scripts\)](#) on page 832
- [snmp](#) on page 833
- [source \(SNMP Scripts\)](#) on page 834
- [traceoptions \(Commit, Op, and SNMP Scripts\)](#) on page 835

## checksum


---

<b>Syntax</b>	<code>checksum (md5   sha-256   sha1) <i>hash</i>;</code>
<b>Hierarchy Level</b>	<code>[edit event-options <a href="#">event-script file filename</a>],</code> <code>[edit system <a href="#">scripts commit file filename</a>],</code> <code>[edit system <a href="#">scripts op file filename</a>],</code> <code>[edit system <a href="#">scripts snmp file filename</a>],</code> <code>[edit system extensions extension-service application file <i>filename</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced for the <code>[edit system extensions extension-service application file <i>filename</i>]</code> hierarchy level in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.
<b>Description</b>	For Junos OS commit scripts, event scripts, op scripts, SNMP scripts, and scripts developed using the Juniper Extension Toolkit (JET) specify the MD5, SHA-1, or SHA-256 checksum hash. When Junos OS executes a local commit, event, op, or SNMP script, the system verifies the authenticity of the script by using the configured checksum hash.
<b>Options</b>	<b>md5 <i>hash</i></b> —MD5 checksum of this script.  <b>sha-256 <i>hash</i></b> —SHA-256 checksum of this script.  <b>sha1 <i>hash</i></b> —SHA-1 checksum of this script.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li><li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li><li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li><li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li><li>• <a href="#">file checksum md5 on page 844</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha-256 on page 845</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha1 on page 846</a> command in the <i>System Basics and Services Command Reference</i></li></ul>

## file (SNMP Scripts)

<b>Syntax</b>	<pre> file filename {   checksum (md5   sha-256   sha1) hash;   oid id {     priority;   }   python-script-user username;   refresh;   refresh-from;   source; } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts snmp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	For Junos OS SNMP scripts, enable an SNMP script that is located in the <code>/var/db/scripts/snmp</code> directory. If the <a href="#">load-scripts-from-flash</a> option is configured, the system reads SNMP scripts from the <code>/config/scripts/snmp</code> directory instead.
<b>Options</b>	<p><b>filename</b>—Name of a Python, Extensible Stylesheet Language Transformations (XSLT), or Stylesheet Language Alternative Syntax (SLAX) file containing an SNMP script.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">snmp on page 833</a></li> </ul>

## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
<div>  <p><b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p> </div>	
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>

- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

---



## oid

---

<b>Syntax</b>	<code>oid id {     priority; }</code>
<b>Hierarchy Level</b>	[edit system <code>scripts snmp file filename</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	For Junos OS SNMP scripts, specify the global object identifier (OID), associated with a subtree of object in the management information base (MIB), to be mapped to the SNMP script. Specify the OID using either a sequence of dotted integers or a subtree name.
<b>Options</b>	<b>priority</b> —(Optional) Override the default registration priority of the OID. The highest priority is 1; the lowest priority is 255. <b>Default:</b> 127 <b>Range:</b> 1 through 255
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Processing Unsupported OIDs with an SNMP Script on page 649</a></li></ul>

## python-script-user (Event and SNMP Scripts)

---

<b>Syntax</b>	<code>python-script-user <i>username</i>;</code>
<b>Hierarchy Level</b>	[edit event-options event-script <i>file filename</i> ], [edit system scripts snmp <i>file filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1R3 for ACX Series, MX Series, PTX Series, QFX Series, and T Series.
<b>Description</b>	Define the user under whose access privileges the Python event or SNMP script executes. If you do not associate a user with the script, then by default, Junos OS executes the script with the privileges of the *nix user and group <b>nobody</b> , which is the generic, unprivileged system account.
	<div> <b>NOTE:</b> You cannot configure Python event and SNMP scripts to execute with root access privileges.</div>
	<div> <b>NOTE:</b> Junos OS executes Python commit and op scripts under the access privileges of the user who executes the command or operation that invokes the script.</div>
<b>Default</b>	If you do not include the <b>python-script-user</b> statement, Junos OS executes the script under the access privileges of the *nix user and group <b>nobody</b> .
<b>Options</b>	<b>username</b> —Username that is configured at the [edit system login] hierarchy level.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Requirements for Executing Python Automation Scripts on Devices Running Junos OS on page 169</a></li><li>• <a href="#">Understanding Python Automation Scripts for Devices Running Junos OS on page 167</a></li></ul>


## refresh (SNMP Scripts)

---

<b>Syntax</b>	refresh;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts snmp</a> ], [edit system <a href="#">scripts snmp file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Overwrite the local copy of all enabled SNMP scripts or a single enabled SNMP script with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li><li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li><li>• <a href="#">snmp on page 833</a></li><li>• <a href="#">refresh-from (SNMP Scripts) on page 832</a></li></ul>

## refresh-from (SNMP Scripts)

---

Syntax	refresh-from <i>url</i> ;
Hierarchy Level	[edit system <a href="#">scripts snmp</a> ], [edit system <a href="#">scripts snmp file</a> <i>filename</i> ]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Overwrite the local copy of all enabled SNMP scripts or a single enabled SNMP script with the copy located at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from <i>url</i></b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus, the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<hr/> <div> <b>NOTE:</b> This statement is not supported on the QFabric system.</div> <hr/>	
Options	<i>url</i> —The source specified as an HTTP URL, FTP URL, or secure copy (scp)-style remote file specification.
Required Privilege Level	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Using an Alternate Source Location for a Script on page 673</a></li><li>• <a href="#">snmp on page 833</a></li><li>• <a href="#">refresh (SNMP Scripts) on page 831</a></li></ul>



## snmp

```
Syntax  snmp {
        file filename {
            checksum (md5 | sha-256 | sha1) hash;
            oid id {
                priority;
            }
            python-script-user username;
            refresh;
            refresh-from;
            source;
        }
        max-datasize size;
        refresh;
        refresh-from;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
```

**Hierarchy Level** [edit system [scripts](#)]

**Release Information** Statement introduced in Junos OS Release 15.1.

**Description** For Junos OS SNMP scripts, configure the SNMP scripting mechanism.  
The remaining statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- [Storing and Enabling Scripts on page 661](#)

## source (SNMP Scripts)

---

<b>Syntax</b>	<code>source url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts snmp file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specify the location of the master source file for an SNMP script. When you issue the <b>set refresh</b> configuration mode command at the same hierarchy level, the local copy of the script is overwritten by the version stored at the specified URL. If the <b>load-scripts-from-flash</b> statement is configured, the device refreshes the scripts on the flash drive instead of the hard drive.
<b>Options</b>	<b>url</b> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring and Using a Master Source Location for a Script on page 669</a></li><li>• <a href="#">Overview of Updating Scripts from a Remote Source on page 667</a></li><li>• <a href="#">refresh (SNMP Scripts) on page 831</a></li><li>• <a href="#">refresh-from (SNMP Scripts) on page 832</a></li></ul>

## traceoptions (Commit, Op, and SNMP Scripts)

<b>Syntax</b>	<pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;     flag flag;     no-remote-trace; } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ], [edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts snmp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Define tracing operations for commit, op, and SNMP scripts.
<b>Default</b>	If you do not include this statement, the default operation during script execution is to log important events in the default trace file for that script type. When the file reaches 128 kilobytes (KB), it is renamed and compressed.



**NOTE:** For op scripts, you can only customize tracing operations for local op scripts. For remote op scripts (op scripts executed using the `op url` command), Junos OS always logs the default trace events in the default op script log file.

**Options** **file *filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`. By default, commit script process tracing output is placed in the file `cscript.log`, op script process tracing is placed in the file `op-script.log`, and SNMP script tracing is placed in the file `snmp.log`. If you include the **file** statement, you must specify a filename. To retain the default, you can specify `cscript.log`, `op-script.log`, or `snmp.log` as the filename.

**Default:**

- Commit scripts: `/var/log/cscript.log`
- Op scripts: `/var/log/op-script.log`
- SNMP scripts: `/var/log/snmp.log`

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

**Range:** 2 through 1000

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **events**—Log important events
- **input**—Log script input data
- **offline**—Generate data for offline development
- **output**—Log script output data
- **rpc**—Log script RPCs
- **xslt**—Log the XSLT library

**no-world-readable**—Restrict file access to owner. This is the default.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file.0.gz***. When ***trace-file*** again reaches its maximum size, ***trace-file.0.gz*** is renamed ***trace-file.1.gz*** and ***trace-file*** is renamed and compressed to ***trace-file.0.gz***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—Enable unrestricted file access.

<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
---------------------------------	---


<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing Commit Script Processing on page 434</a></li><li>• <a href="#">Tracing Op Script Processing on page 529</a></li><li>• <a href="#">Tracing SNMP Script Processing on page 655</a></li><li>• <a href="#">Dampening Script Execution on page 697</a></li></ul>
------------------------------	---

## CHAPTER 52

# Translation Script Configuration Statements

- [max-datasize on page 838](#)
- [translation on page 839](#)

## max-datasize

<b>Syntax</b>	<code>max-datasize size;</code>
<b>Hierarchy Level</b>	<p>[edit <a href="#">event-options event-script</a>],  [edit system <a href="#">scripts commit</a>],  [edit system <a href="#">scripts op</a>],  [edit system <a href="#">scripts snmp</a>],  [edit system extension extension-service application],  [edit system <a href="#">scripts translation</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced for the [edit system extension extension-service application] hierarchy in Junos OS Release 16.1 for MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series.</p> <p>Statement introduced for the [edit system scripts translation] hierarchy in Junos OS Release 16.1.</p>
<b>Description</b>	<p>Maximum amount of memory allocated for the data segment during execution of a script of the configured type. Junos OS sets the maximum memory limit for the executing script to the configured value irrespective of the total memory available on the system at the time of execution. If the executing script exceeds the specified maximum memory limit for that script type, it exits gracefully.</p>
<div>  <p><b>NOTE:</b> For op scripts, the <code>max-datasize</code> statement is only enforced for op scripts that are local to the device. If you execute an op script from a remote location using the <code>op url</code> command, Junos OS uses the default memory allocation settings.</p> </div>	
<b>Default</b>	<p>If you do not include the <code>max-datasize</code> statement, the system allocates half of the total available memory of the system up to a maximum value of 128 MB for the data segment portion of the executed script.</p>
<b>Options</b>	<p><b>size</b>—Maximum amount of memory allocated for the data segment during execution of a script of the given type. If you do not specify a unit of measure, the default is bytes.</p> <p><b>Syntax:</b> <i>size</i> to specify bytes, <i>sizek</i> to specify KB, <i>sizem</i> to specify MB, or <i>sizeg</i> to specify GB</p> <p><b>Range:</b> 23,068,672 (22 MB) through 1,073,741,824 (1 GB)</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">max-policies on page 786</a></li> <li>• <a href="#">Understanding Limits on Executed Event Policies and Memory Allocation for Scripts on page 693</a></li> </ul>

- [Example: Configuring Limits on Executed Event Policies and Memory Allocation for Scripts on page 695](#)

## translation

---

<b>Syntax</b>	<code>translation {     max-datasize; }</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 16.1.
<b>Description</b>	Configure options for translation scripts.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Storing and Enabling Scripts on page 661</a></li></ul>





## PART 13

# Operational Commands

- [Operational Commands Common to All Scripts on page 843](#)
- [Op Script Operational Commands on page 853](#)
- [Event Script Operational Commands on page 859](#)



## CHAPTER 53

# Operational Commands Common to All Scripts

- file checksum md5
- file checksum sha-256
- file checksum sha1
- request system scripts convert
- request system scripts refresh-from
- request system scripts synchronize

## file checksum md5

---

<b>Syntax</b>	<code>file checksum md5 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
<b>Description</b>	Calculate the Message Digest 5 (MD5) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the MD5 checksum.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li><li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li><li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li><li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li><li>• <a href="#">file checksum sha-256 on page 845</a></li><li>• <a href="#">file checksum sha1 on page 846</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum md5 on page 844</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = $ABC123
```

## file checksum sha-256

<b>Syntax</b>	<code>file checksum sha-256 &lt;pathname&gt; filename</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
<b>Description</b>	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
<b>Options</b>	<p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the SHA-256 checksum.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li> <li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li> <li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li> <li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li> <li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li> <li>• <a href="#">file checksum md5 on page 844</a></li> <li>• <a href="#">file checksum sha1 on page 846</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha-256 on page 845</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =$ABC123
```

## file checksum sha1

---

<b>Syntax</b>	<code>file checksum sha1 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.
<b>Description</b>	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the SHA-1 checksum.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Checksum Hashes for a Commit Script on page 264</a></li><li>• <a href="#">Configuring Checksum Hashes for an Event Script on page 628</a></li><li>• <a href="#">Configuring Checksum Hashes for an Op Script on page 460</a></li><li>• <a href="#">Configuring Checksum Hashes for an SNMP Script on page 647</a></li><li>• <a href="#">Executing an Op Script from a Remote Site on page 462</a></li><li>• <a href="#">file checksum md5 on page 844</a></li><li>• <a href="#">file checksum sha-256 on page 845</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha1 on page 846</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax  
  
SHA1 (/var/db/scripts/commitscript.slax) = $ABC123
```

## request system scripts convert

<b>Syntax</b>	<code>request system scripts convert (slax-to-xslt   xslt-to-slax) source <i>source/filename</i> destination <i>destination/&lt;filename&gt;</i> &lt;partial&gt; &lt;version (1.0   1.1)&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>partial</b> and <b>version</b> options added in Junos OS Release 12.2.
<b>Description</b>	Convert an Extensible Stylesheet Language Transformations (XSLT) script to Stylesheet Language Alternative Syntax (SLAX), or convert a SLAX script to XSLT.
<b>Options</b>	<p><b>destination <i>destination/&lt;filename&gt;</i></b>—Path and filename for the converted file.</p> <p>Optionally, you can specify a filename for the converted file. If you do not specify a filename, the software assigns one automatically. The default destination filename is <b>SLAX-Conversion-Temp</b> or <b>slax-temp</b> depending on the Junos OS release, with a randomly generated, five-character, alpha-numeric extension. For example, the software converts a source file called <b>test.xml</b> to <b>slax-temp.kWwQk</b>. The software converts a source file called <b>test1.slax</b> to <b>slax-temp.zN61h</b>.</p> <p><b>partial</b>—(Optional) Convert partial script input.</p> <p><b>slax-to-xslt</b>—Convert a SLAX script to XSLT.</p> <p><b>source <i>source/filename</i></b>—Path and filename of the file to convert.</p> <p><b>version</b>—(Optional) SLAX version number that will be included in the version statement of the generated script for XSLT-to-SLAX conversions. Acceptable values are 1.0 and 1.1.</p> <p><b>Default:</b> 1.1</p> <p><b>xslt-to-slax</b>—Convert an XSLT script to SLAX.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Converting Scripts Between SLAX and XSLT on page 711</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system scripts convert slax-to-xslt on page 847</a> <a href="#">request system scripts convert xslt-to-slax on page 848</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system scripts convert slax-to-xslt

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op
conversion complete
```

#### request system scripts convert xslt-to-slax

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/commit/script1.xml  
destination /var/db/scripts/commit partial version 1.0  
conversion complete
```



## request system scripts refresh-from

<b>Syntax</b>	<code>request system scripts refresh-from (commit   event   op   snmp) file <i>filename</i> url <i>url</i> &lt;sync&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1. Option <b>sync</b> introduced in Junos OS Release 13.2.
<b>Description</b>	<p>Refresh a script from a remote source at the specified URL. The system uses the script type to determine the directory on the device where the script resides. If you include the <b>sync</b> option, the system refreshes the script on the requesting Routing Engine and then synchronizes the script on the other Routing Engine. If the <b>load-scripts-from-flash</b> statement is present in the configuration for that Routing Engine, the system refreshes the script in flash memory. Otherwise, the system refreshes the script on the hard disk.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-script-refresh-from&gt;   &lt;type&gt;(commit   op   event   snmp)&lt;/type&gt;   &lt;file&gt;<i>filename</i>&lt;/file&gt;   &lt;url&gt;<i>url</i>&lt;/url&gt;   &lt;sync/&gt; &lt;/request-script-refresh-from&gt;</pre>
<b>Options</b>	<p><b>(commit   event   op   snmp)</b>—Specify that the script resides in the commit, event, op, or snmp script subdirectory under the scripts directory.</p> <p><b>file <i>filename</i></b>—Filename of the script to refresh.</p> <p><b>sync</b>—(Optional) Synchronize the refreshed script from the requesting Routing Engine to the other Routing Engine.</p> <p><b>url <i>url</i></b>—URL of the file that will replace the script provided as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Refreshing a Script from an Alternate Location on page 673</a></li> <li>• <a href="#">Example: Refreshing a Script from an Alternate Source on page 675</a></li> <li>• <a href="#">Example: Configuring and Refreshing from the Master Source for a Script on page 671</a></li> <li>• <a href="#">Synchronizing Scripts Between Routing Engines on page 702</a></li> <li>• <a href="#">request system scripts synchronize on page 851</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">request system scripts refresh-from on page 850</a></p> <p><a href="#">request system scripts refresh-from sync on page 850</a></p>

## Sample Output

### request system scripts refresh-from

The following command refreshes the op script **config.slax** from the remote source specified by the URL.

```
user@host> request system scripts refresh-from op file config.slax url
http://host1.juniper.net/config.slax
config.slax                                100%    9    0.0KB/s   00:00
```

The following command refreshes the op script **config1.slax** from the remote source specified by the URL. However, in this case, the script does not exist at the remote site, and the device generates an error message.

```
user@host> request system scripts refresh-from op file config1.slax url
http://host1.juniper.net/config1.slax
refreshing 'config1.slax' from 'http://host1.juniper.net/config1.slax'
fetch-secure: http://host1.juniper.net/config1.slax: Not Found
error: file-fetch failed
error: communication error: rpc failed (file-fetch)
error: error retrieving file http://host1.juniper.net/config1.slax
```

## Sample Output

### request system scripts refresh-from sync

The following command refreshes the op script **config.slax** on the requesting Routing Engine from the remote source specified by the URL. The system then synchronizes the updated script to the other Routing Engine.

```
user@host> request system scripts refresh-from op file config.slax url
http://host1.juniper.net/config.slax sync
config.slax                                100%    9    0.0KB/s   00:00
```

## request system scripts synchronize

<b>Syntax</b>	request system scripts synchronize (all   commit   event   lib   op   snmp) <file <i>filename</i> > <newer-than YYYY-MM-DD.HH:MM:SS>
<b>Release Information</b>	Command introduced in Junos OS Release 13.2.
<b>Description</b>	<p>This command is for devices with multiple Routing Engines only. Synchronize scripts from the requesting Routing Engine to the other Routing Engine.</p> <p>Specify <b>all</b> to synchronize all scripts present on the requesting Routing Engine to the responding Routing Engine. Specify <b>commit</b>, <b>event</b>, <b>lib</b>, <b>op</b>, or <b>snmp</b> to synchronize all scripts of the given type to the other Routing Engine. Include the <b>file</b> option or the <b>newer-than</b> option to narrow the scope to only synchronize scripts with the specified filename or date criteria.</p> <p>If the <b>load-scripts-from-flash</b> statement is configured for the requesting Routing Engine, the device synchronizes the scripts that are in flash memory. Otherwise, the device synchronizes the scripts that are on the hard disk. If the <b>load-scripts-from-flash</b> statement is configured for the responding Routing Engine, the scripts are synchronized to flash memory. Otherwise, the scripts are synchronized to the hard disk. The device synchronizes a script regardless of whether it is enabled in the configuration.</p> <p>The Junos XML management protocol equivalent for this operational mode command is:</p> <pre>&lt;request-scripts-synchronize&gt;   &lt;all/&gt;   &lt;commit/&gt;   &lt;event/&gt;   &lt;lib/&gt;   &lt;op/&gt;   &lt;file&gt;<i>filename</i>&lt;/file&gt;   &lt;newer-than&gt;YYYY-MM-DD.HH:MM:SS&lt;/newer-than&gt; &lt;/request-scripts-synchronize&gt;</pre>
<b>Options</b>	<p>(all   commit   event   lib   op   snmp)—Synchronize all scripts or all scripts of a given type that are present on the requesting Routing Engine to the responding Routing Engine.</p> <p><b>file <i>filename</i></b>—(Optional) Synchronize an individual script with the specified filename from the requesting Routing Engine to the responding Routing Engine. This option can only be used with a specific script type.</p> <p><b>newer-than YYYY-MM-DD.HH:MM:SS</b>—(Optional) Synchronize scripts that are more current than the specified date and time from the requesting Routing Engine to the responding Routing Engine.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Script Synchronization Between Routing Engines on page 701</a></li> <li>• <a href="#">Synchronizing Scripts Between Routing Engines on page 702</a></li> <li>• <a href="#">synchronize on page 722</a></li> </ul>

- [request system scripts refresh-from on page 849](#)

**List of Sample Output**    [request system scripts synchronize all on page 852](#)  
[request system scripts synchronize commit file on page 852](#)  
[request system scripts synchronize event file on page 852](#)  
[request system scripts synchronize op newer-than on page 852](#)

## Sample Output

### [request system scripts synchronize all](#)

The following command synchronizes all scripts from the requesting Routing Engine to the responding Routing Engine.

```
user@host> request system scripts synchronize all
user@host>
```

### [request system scripts synchronize commit file](#)

The following command synchronizes the commit script **vpn-commit.slax** from the requesting Routing Engine to the responding Routing Engine. In this case, the script does not exist on the requesting Routing Engine, and the device issues an error message.

```
user@host> request system scripts synchronize commit file vpn-commit.slax
error: Invalid directory: No such file or directory
warning: No script will be pushed to other RE
```

### [request system scripts synchronize event file](#)

The following command synchronizes the event script **ospf-neighbor.slax** from the requesting Routing Engine to the responding Routing Engine. In this case, the responding Routing Engine does not have the resources to perform the operation, and the device issues an error message.

```
user@host> request system scripts synchronize event file ospf-neighbor.slax
error: Unable to copy scripts to re: re1
```

### [request system scripts synchronize op newer-than](#)

The following command synchronizes any op scripts that have been modified after the specified date. In this case, no op scripts meet this requirement.

```
user@host> request system scripts synchronize op newer-than 2012-05-30
warning: No script will be pushed to other RE
```

## CHAPTER 54

# Op Script Operational Commands

- `op`
- `op invoke-debugger cli`

## op

---

**Syntax**    `op filename`  
              `<detail>`  
              `<argument-name argument-value>`  
              `<key (md5 | sha-256 | sha1) key-value`  
              `<url url>`

**Release Information**    Command introduced in Junos OS Release 7.6.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              **key** option introduced in Junos OS Release 10.0.  
                              **url** option introduced in Junos OS Release 10.0.

**Description**    Execute an op script configured under the **[edit system scripts op file]** hierarchy. Op scripts can be stored in one of the following locations:

- On the local device in the `/var/db/scripts/op` directory on the hard disk
- On the local device in the `/config/scripts/op` directory on the flash drive.

To execute op scripts from the flash drive, you must configure the **load-scripts-from-flash** statement at the **[edit system scripts]** hierarchy level.

- At a remote URL

**Options**    **detail**—(Optional) Display detailed output.

**argument-name argument-value**—(Optional) Specify one or more arguments to the script.  
For each argument you include on the command line, you must specify a corresponding value for the argument.

**key (md5 | sha-256 | sha1) key-value**—(Optional) With the `<url>` option, specify a checksum hash to verify the integrity of the script. You can include the `<key>` option if the **checksum** statement is included at the **[edit system scripts op file filename]** hierarchy level.

**url url**—(Optional) Specify the URL for a remote script.



**NOTE:** Statements configured under the **[edit system scripts op]** hierarchy level are only enforced for op scripts that are local to the device. If you configure memory allocation, script dampening, traceoptions, or other op script-specific statements within that hierarchy, Junos OS does not apply the configuration when you execute a remote script using the **op url** command.

---

**Additional Information**    For more information about op scripts, see the *Automation Scripting Feature Guide*.

**Required Privilege Level** maintenance

**Related Documentation**

- [Executing an Op Script on page 461](#)
- [Executing an Op Script from a Remote Site on page 462](#)
- [checksum on page 726](#)
- [file checksum md5 on page 844](#)
- [file checksum sha-256 on page 845](#)
- [file checksum sha1 on page 846](#)

**List of Sample Output** [op on page 855](#)  
[op url on page 855](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**op**

```
user@host> op script1 interface ge-0/2/0.0 protocol inet
```

**op url**

```
user@host> op url https://www.example.net/opscrip/2009-04-01.01.slax key md5  
abcd1234dcba4321cdef5678fedc8765 interface ge-0/2/0.0 protocol inet
```

## op invoke-debugger cli

---

<b>Syntax</b>	<b>op invoke-debugger cli</b> <i>filename</i> <detail> < <i>argument-name argument-value</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 13.1 for the SLAX debugger. Support for the Python debugger added in Junos OS Release 16.1.
<b>Description</b>	<p>Invoke an interactive debugger to trace the execution of an op script that is enabled in the configuration.</p> <p>When you execute this command for a SLAX script, the device invokes the SLAX debugger and displays the (sdb) prompt. When you execute this command for a Python script, the device invokes the <b>pdb</b> interactive debugger for Python programs and displays the (Pdb) prompt. For more information about the Python debugger, see the official Python documentation.</p>
<b>Options</b>	<p><b>detail</b>—Display detailed output.</p> <p><b><i>argument-name argument-value</i></b>—(Optional) Specify one or more arguments for the script. For each argument you include on the command line, you must specify a corresponding value for the argument.</p> <p><b><i>filename</i></b>—Script filename</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SLAX Debugger, Profiler, and callflow on page 158</a></li></ul>
<b>List of Sample Output</b>	<a href="#">op invoke-debugger cli (SLAX script) on page 856</a> <a href="#">op invoke-debugger cli (Python script) on page 856</a> <a href="#">op invoke-debugger cli script argument1 value1 on page 857</a>
<b>Output Fields</b>	When you enter the <b>op invoke-debugger cli</b> command, depending on the script type, the CLI displays either the (sdb) or (Pdb) prompt where you can enter debugging commands.

## Sample Output

### op invoke-debugger cli (SLAX script)

```
user@host> op invoke-debugger cli script1.slax
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```

### op invoke-debugger cli (Python script)

```
user@host> op invoke-debugger cli script1.py
(Pdb)
```



**op invoke-debugger cli script argument1 value1**

```
user@host> op invoke-debugger cli script2.slax interface ge-0/2/0.0
sdb: The SLAX Debugger (version )
Type 'help' for help
(sdb)
```



## CHAPTER 55

# Event Script Operational Commands

- request system scripts event-scripts

## request system scripts event-scripts

---

<b>Syntax</b>	<b>request system scripts event-scripts</b> <i>&lt;reload&gt;</i>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Reload all active event scripts into the eventd process' memory. This enables you to replace or update an existing event script on the device without changing the device's configuration or disrupting operations.
<b>Options</b>	<b>reload</b> —Reload all active event-scripts.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Replacing an Event Script on page 628</a></li></ul>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### request system scripts event-scripts reload

```
user@host> request system scripts event-scripts reload
Event scripts loaded
```