

Junos® OS for EX Series Ethernet Switches

Port Security Feature Guide for EX2300, EX3400, and EX4300 Switches

Release

14.1x53 (THIS IS A !5.1 page, with BAD release: saving
into 13.2X51)



Modified: 2017-01-05

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
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[®] OS for EX Series Ethernet Switches Port Security Feature Guide for EX2300, EX3400, and EX4300 Switches
Release 14.1x53 (THIS IS A !5.1 page, with BAD release: saving into 13.2X51)
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 | xiii |
| | Documentation and Release Notes | xiii |
| | Supported Platforms | xiii |
| | Using the Examples in This Manual | xiii |
| | Merging a Full Example | xiv |
| | Merging a Snippet | xiv |
| | Documentation Conventions | xv |
| | Documentation Feedback | xvii |
| | Requesting Technical Support | xvii |
| | Self-Help Online Tools and Resources | xvii |
| | Opening a Case with JTAC | xviii |
| Chapter 1 | Port Security Overview | 19 |
| | Understanding Port Security Features to Protect the Access Ports on Your Device | |
| | Against the Loss of Information and Productivity | 19 |
| | Configuring Port Security Features | 21 |
| | Configuring Port Security (J-Web Procedure) | 23 |
| Chapter 2 | Configuring IP Source Guard to Prevent IP Spoofing Attacks | 29 |
| | Understanding IP Source Guard for Port Security on EX Series Switches | 29 |
| | IP Address Spoofing | 29 |
| | How IP Source Guard Works | 29 |
| | IPv6 Source Guard | 30 |
| | The DHCP Snooping Table | 30 |
| | Typical Uses of Other Junos OS Features with IP Source Guard | 31 |
| | Configuring IP Source Guard (CLI Procedure) | 32 |
| | Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect | |
| | the Switch from IP Spoofing and ARP Spoofing | 33 |
| | Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to | |
| | Protect a Switch from IPv6 Address Spoofing | 38 |
| Chapter 3 | Configuring Dynamic ARP Inspection to Prevent ARP Spoofing | |
| | Attacks | 45 |
| | Understanding Dynamic ARP Inspection for Protecting Switching Devices Against | |
| | ARP Spoofing | 45 |
| | Address Resolution Protocol | 45 |
| | ARP Spoofing | 46 |
| | Dynamic ARP Inspection | 46 |
| | Prioritizing Inspected Packets | 47 |
| | Enabling Dynamic ARP Inspection (CLI Procedure) | 48 |
| | Enabling Dynamic ARP Inspection (J-Web Procedure) | 49 |

| | | |
|------------------|--|-----------|
| Chapter 4 | Configuring DHCP Snooping to Filter DHCP Messages from Untrusted Hosts | 51 |
| | Understanding DHCP Snooping for Port Security | 51 |
| | DHCP Snooping Basics | 52 |
| | Enabling DHCP Snooping | 53 |
| | DHCP Snooping Process | 53 |
| | DHCPv6 Snooping | 54 |
| | Rapid Commit for DHCPv6 | 55 |
| | DHCP Server Access | 55 |
| | Switch, DHCP Clients, and the DHCP Server Are All on the Same VLAN | 55 |
| | Switch Acts as the DHCP Server | 56 |
| | Switch Acts as a Relay Agent | 57 |
| | Static IP Address Additions to the DHCP Snooping Database | 58 |
| | Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks | 59 |
| | DHCP Option 82 Overview | 59 |
| | Suboption Components of Option 82 | 60 |
| | Switching Device Configurations That Support Option 82 | 61 |
| | Switching Device, DHCP Clients, and the DHCP Server Are on the Same VLAN or Bridge Domain | 61 |
| | Switching Device Acts as a Relay Agent | 61 |
| | DHCPv6 Options | 62 |
| | Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) | 63 |
| | Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure) | 65 |
| | Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure) | 66 |
| Chapter 5 | Enabling IPv6 ND Inspection and RA Guard to Prevent IPv6 Spoofing Attacks | 69 |
| | Understanding IPv6 Neighbor Discovery Inspection | 69 |
| | Enabling IPv6 Neighbor Discovery Inspection | 71 |
| | Understanding IPv6 Router Advertisement Guard | 71 |
| | Configuring Stateless IPv6 Router Advertisement Guard on Switches | 74 |
| | Configuring a Discard Policy for RA Guard | 75 |
| | Configuring an Accept Policy for RA Guard | 76 |
| | Enabling Stateless RA Guard on an Interface | 78 |
| | Enabling Stateless RA Guard on a VLAN | 78 |
| | Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard | 79 |
| | Configuring Stateful IPv6 Router Advertisement Guard on Switches | 79 |
| | Configuring a Discard Policy for RA Guard | 80 |
| | Configuring an Accept Policy for RA Guard | 81 |
| | Enabling Stateful RA Guard on an Interface | 83 |
| | Enabling Stateful RA Guard on a VLAN | 83 |
| | Configuring the Learning State on an Interface | 83 |
| | Configuring the Forwarding State on an Interface | 85 |

| | | |
|------------------|--|------------|
| | Configuring the Blocking State on an Interface | 85 |
| | Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard | 85 |
| Chapter 6 | Configuring MAC Limiting, MAC Move Limiting and Persistent MAC Learning to Prevent DHCP Starvation Attacks | 87 |
| | Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches | 88 |
| | MAC Limiting | 88 |
| | MAC Move Limiting | 89 |
| | Actions for MAC Limiting and MAC Move Limiting | 89 |
| | Configuring MAC Limiting (CLI Procedure) | 91 |
| | Limiting the Number of MAC Addresses Learned by an Interface | 91 |
| | Limiting the Number of MAC Addresses Learned by a VLAN | 91 |
| | Configuring MAC Limiting (J-Web Procedure) | 92 |
| | Configuring MAC Move Limiting (CLI Procedure) | 94 |
| | Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) | 96 |
| | Understanding Persistent MAC Learning (Sticky MAC) | 97 |
| | Configuring Persistent MAC Learning (CLI Procedure) | 98 |
| Chapter 7 | Configuring MACSec to Provide Point-to-Point Security on Ethernet Links | 101 |
| | Understanding Media Access Control Security (MACsec) | 101 |
| | How MACsec Works | 102 |
| | Understanding Connectivity Associations and Secure Channels | 102 |
| | Understanding MACsec Security Modes | 103 |
| | Understanding Static Connectivity Association Key Security Mode (Recommended Security Mode for Switch-to-Switch Links) | 103 |
| | Understanding Dynamic Secure Association Key Security Mode (Switch-to-Host Links) | 104 |
| | Understanding Static Secure Association Key Security Mode (Supported for Switch-to-Switch Links) | 104 |
| | Understanding the Requirements to Enable MACsec on a Switch-to-Host Link | 105 |
| | MACsec Hardware and Software Support Summary | 105 |
| | Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches | 106 |
| | Understanding MACsec Software Requirements for EX Series and QFX Series Switches | 107 |
| | Understanding the MACsec Feature License Requirement | 108 |
| | MACsec Limitations | 108 |
| | Configuring Media Access Control Security (MACsec) | 109 |
| | Acquiring and Downloading the Junos OS Software | 109 |
| | Acquiring and Downloading the MACsec Feature License | 111 |
| | Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only) | 111 |
| | Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links) | 113 |

| | | |
|------------------|---|------------|
| | Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link | 117 |
| | Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link | 121 |
| Chapter 8 | Enabling Trusted DHCP Servers to Protect Against Rogue DHCP Servers | 127 |
| | Understanding Trusted DHCP Servers for Port Security | 127 |
| | Enabling a Trusted DHCP Server (CLI Procedure) | 128 |
| | Enabling a Trusted DHCP Server (J-Web Procedure) | 128 |
| Chapter 9 | Configuration Statements | 131 |
| | [edit vlans] Configuration Statement Hierarchy on EX Series Switches | 133 |
| | Supported Statements in the [edit vlans] Hierarchy Level | 134 |
| | Unsupported Statements in the [edit vlans] Hierarchy Level | 136 |
| | accept | 137 |
| | access-security | 138 |
| | arp-inspection | 139 |
| | cak | 140 |
| | circuit-id | 141 |
| | ckn | 142 |
| | connectivity-association | 143 |
| | connectivity-association (MACsec Interfaces) | 144 |
| | dhcp-security | 145 |
| | dhcp-service | 148 |
| | dhcp-snooping-file | 149 |
| | dhcpv6-options | 150 |
| | dhcpv6-snooping-file | 151 |
| | direction | 152 |
| | discard | 153 |
| | encryption (MACsec) | 154 |
| | exclude-protocol | 155 |
| | group (DHCP Security) | 156 |
| | host-name | 157 |
| | id | 158 |
| | interface (DHCP Security) | 159 |
| | interface (RA Guard) | 160 |
| | interface-mac-limit | 161 |
| | interfaces (MACsec) | 163 |
| | ip-source-guard | 164 |
| | ipv6-prefix-list | 166 |
| | ipv6-source-address-list | 167 |
| | ipv6-source-guard | 168 |
| | key (MACsec) | 169 |
| | key-server-priority (MACsec) | 170 |
| | mac | 171 |
| | mac-address (MACsec) | 172 |
| | mac-list | 173 |
| | mac-move-limit | 174 |
| | macsec | 176 |

| | |
|--|-----|
| mark-interface (RA Guard) | 177 |
| match-list | 178 |
| match-option | 180 |
| mka | 181 |
| must-secure | 182 |
| neighbor-discovery-inspection | 183 |
| no-dhcpv6-options | 184 |
| no-dhcpv6-snooping | 184 |
| no-encryption (MACsec) | 185 |
| no-option16 | 186 |
| no-option18 | 186 |
| no-option37 | 187 |
| offset | 188 |
| option-16 (DHCPv6 Snooping) | 189 |
| option-18 (DHCPv6 Snooping) | 190 |
| option-37 (DHCPv6 Snooping) | 191 |
| overrides (DHCP Security) | 192 |
| packet-action | 193 |
| persistent-learning | 195 |
| policy | 196 |
| port-id | 197 |
| pre-shared-key | 198 |
| prefix (DHCPv6 Options) | 199 |
| prefix-list | 200 |
| recovery-timeout | 201 |
| remote-id | 203 |
| replay-protect | 204 |
| replay-window-size | 205 |
| router-advertisement-guard | 206 |
| routing-instance-name (circuit-id) | 208 |
| secure-channel | 209 |
| security-association | 210 |
| security-mode | 211 |
| source-mac-address-list | 212 |
| stateful | 213 |
| stateless | 214 |
| static-ip | 215 |
| static-ipv6 | 216 |
| traceoptions (DHCP) | 217 |
| transmit-interval (MACsec) | 219 |
| trusted | 220 |
| untrusted | 220 |
| use-interface-description | 221 |
| use-interface-index | 222 |
| use-interface-name | 223 |
| use-string | 224 |
| use-vlan-id | 226 |
| vendor-id | 228 |
| vlan (RA Guard) | 229 |

Chapter 10

| | |
|--|------------|
| write-interval | 230 |
| Operational Commands | 231 |
| clear access-security router-advertisement statistics | 232 |
| clear arp | 233 |
| clear dhcp-security binding | 235 |
| clear dhcp-security ipv6 binding | 236 |
| clear dot1x | 237 |
| clear ethernet-switching recovery-timeout | 239 |
| clear security mka statistics | 240 |
| request access-security router-advertisement-guard-forward | 241 |
| request access-security router-advertisement-guard-block | 242 |
| request access-security router-advertisement-guard-learn interface | 243 |
| show access-security router-advertisement statistics | 244 |
| show access-security router-advertisement state | 245 |
| show dhcp-security arp inspection statistics | 247 |
| show dhcp-security binding | 249 |
| show dhcp-security binding ip-source-guard | 252 |
| show dhcp-security ipv6 binding | 254 |
| show dhcp-security ipv6 statistics | 256 |
| show dhcp-security neighbor-discovery-inspection statistics | 259 |
| show security macsec connections | 261 |
| show security macsec statistics | 263 |
| show security mka sessions | 267 |
| show security mka statistics | 269 |

List of Figures

| | | |
|------------------|--|-----------|
| Chapter 2 | Configuring IP Source Guard to Prevent IP Spoofing Attacks | 29 |
| | Figure 1: Network Topology for Basic Port Security | 35 |
| | Figure 2: Network Topology for Basic Port Security | 40 |
| Chapter 4 | Configuring DHCP Snooping to Filter DHCP Messages from Untrusted Hosts | 51 |
| | Figure 3: DHCP Server Connected Directly to a Switch | 56 |
| | Figure 4: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port | 56 |
| | Figure 5: Switch Is the DHCP Server | 57 |
| | Figure 6: Switch Acting as a Relay Agent Through a Router to the DHCP Server | 58 |
| | Figure 7: DHCP Clients, Switching Device, and the DHCP Server Are All on the Same VLAN or Bridge Domain | 61 |
| | Figure 8: Switching Device Acting as an Extended Relay Server | 62 |
| Chapter 5 | Enabling IPv6 ND Inspection and RA Guard to Prevent IPv6 Spoofing Attacks | 69 |
| | Figure 9: Stateful RA Guard State Transitions | 73 |

List of Tables

| | | |
|-------------------|---|-------------|
| | About the Documentation | xiii |
| | Table 1: Notice Icons | xv |
| | Table 2: Text and Syntax Conventions | xv |
| Chapter 1 | Port Security Overview | 19 |
| | Table 3: Port Security Settings on VLANs | 24 |
| | Table 4: Port Security on Interfaces | 26 |
| Chapter 2 | Configuring IP Source Guard to Prevent IP Spoofing Attacks | 29 |
| | Table 5: Components of the Port Security Topology | 35 |
| | Table 6: Components of the Port Security Topology | 40 |
| Chapter 4 | Configuring DHCP Snooping to Filter DHCP Messages from Untrusted Hosts | 51 |
| | Table 7: DHCPv6 Messages and DHCPv4 Equivalent Messages | 54 |
| Chapter 5 | Enabling IPv6 ND Inspection and RA Guard to Prevent IPv6 Spoofing Attacks | 69 |
| | Table 8: IPv6 RA guard states | 73 |
| Chapter 7 | Configuring MACSec to Provide Point-to-Point Security on Ethernet Links | 101 |
| | Table 9: MACsec Hardware and Software Support Summary for EX Series and QFX Series Switches | 106 |
| Chapter 9 | Configuration Statements | 131 |
| | Table 10: Unsupported [edit vlans] Configuration Statements on EX Series Switches | 136 |
| Chapter 10 | Operational Commands | 231 |
| | Table 11: show access-security router-advertisement statistics Output Fields | 244 |
| | Table 12: show access-security router-advertisement state Output Fields | 245 |
| | Table 13: show dhcp-security arp inspection statistics Output Fields | 247 |
| | Table 14: show dhcp-security binding Output Fields | 249 |
| | Table 15: show dhcp-security binding ip-source-guard Output Fields | 252 |
| | Table 16: show dhcp-security ipv6 binding Output Fields | 254 |
| | Table 17: show dhcp-security ipv6 statistics Output Fields | 257 |
| | Table 18: show dhcp-security neighbor-discovery-inspection statistics Output Fields | 259 |
| | Table 19: show security macsec connections Output Fields | 261 |
| | Table 20: show security macsec statistics Output Fields | 263 |
| | Table 21: show security mka sessions Output Fields | 267 |
| | Table 22: show security mka statistics Output Fields | 269 |

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

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:

- EX 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 xv 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 xv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

| Convention | Description | Examples |
|----------------------------|--------------------------------|--|
| Bold text like this | Represents text that you type. | To enter configuration mode, type the configure command: user@host> configure |

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

| Convention | Description | Examples |
|--------------------------------|---|--|
| Fixed-width text like this | Represents output that appears on the terminal screen. | user@host> show chassis alarms No alarms currently active |
| <i>Italic text like this</i> | <ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles. | <ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i> |
| <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@# set system domain-name <i>domain-name</i> |
| 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">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE. |
| < > (angle brackets) | Encloses optional keywords or variables. | stub <default-metric <i>metric</i>>; |
| (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. | broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>) |
| # (pound sign) | Indicates a comment specified on the same line as the configuration statement to which it applies. | rsvp { # Required for dynamic MPLS only |
| [] (square brackets) | Encloses a variable for which you can substitute one or more values. | community name members [<i>community-ids</i>] |
| Indentation and braces ({ }) | Identifies a level in the configuration hierarchy. | [edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } } |
| ;(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">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel. |

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

| Convention | Description | Examples |
|------------------------------|---|--|
| > (bold right angle bracket) | Separates levels in a hierarchy of menu selections. | In the configuration editor hierarchy, select Protocols>Ospf . |

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. 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>.

CHAPTER 1

Port Security Overview

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring Port Security Features on page 21](#)
- [Configuring Port Security \(J-Web Procedure\) on page 23](#)

Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your device against the loss of information and productivity that such attacks can cause.

The Juniper Networks Junos operating system (Junos OS) provides features to help secure ports on a device. Ports can be categorized as either trusted or untrusted. You apply policies appropriate to each category to protect ports against various types of attacks.

Basic port security features are enabled in the device's default configuration. You can configure additional features with minimal configuration steps.

Depending on the particular feature, you can configure the feature either on VLANs or bridge domain interfaces.

Port security features supported on switching devices are:

- DHCP snooping—Filters and blocks ingress Dynamic Host Configuration Protocol (DHCP) server messages on untrusted ports, and builds and maintains a database of DHCP lease information, which is called the DHCP snooping database.



NOTE: DHCP snooping is not enabled in the default configuration of the switching device. DHCP snooping is enabled on a VLAN or bridge domain. The details of enabling DHCP snooping depend on the particular device.

- DHCPv6 snooping—DHCP snooping for IPv6.
- DHCP option 82—Also known as the DHCP Relay Agent Information option. This DHCPv4 feature helps protect the switching device against attacks such as spoofing

of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

- DHCPv6 option 37—Option 37 is the remote ID option for DHCPv6 and is used to insert information about the network location of the remote host into DHCPv6 packets. You enable option 37 on a VLAN.



NOTE: DHCPv6 snooping with option 37 is not supported on the MX Series.

- DHCPv6 option 18—Option 18 is the circuit ID option for DHCPv6 and is used to insert information about the client port into DHCPv6 packets. This option includes other details that can be optionally configured, such as the prefix and the interface description.
- DHCPv6 option 16—Option 16 is the vendor ID option for DHCPv6 and is used to insert information about the vendor of the client hardware into DHCPv6 packets.
- Dynamic ARP inspection (DAI)—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable DAI on a VLAN.
- IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable neighbor discovery inspection on a VLAN.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the DHCP snooping database. If the packet cannot be validated, it is discarded. You enable IP source guard on a VLAN or bridge domain.



NOTE: IP source guard is not supported on the QFX Series.

- IPv6 source guard—IP source guard for IPv6.



NOTE: IPv6 source guard is not supported on the QFX Series.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You can enable MAC limiting on an interface.
- MAC move limiting—Tracks MAC movement and detects MAC spoofing on access ports. You enable this feature on a VLAN or bridge domain.

- Persistent MAC learning—Also known as sticky MAC. Persistent MAC learning enables interfaces to retain dynamically learned MAC addresses across switch reboots. You enable this feature on an interface.
- Trusted DHCP server—Configuring the DHCP server on a trusted port protects against rogue DHCP servers sending leases. You enable this feature on an interface (port). By default, access ports are untrusted, and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect an Ethernet switch to other switches or to routers.)

Related Documentation

- *Security Features for EX Series Switches Overview*
- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*
- [Understanding DHCP Snooping for Port Security on page 51](#)
- [Understanding IPv6 Neighbor Discovery Inspection on page 69](#)
- [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 45](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches on page 29](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 88](#)
- [Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59](#)

Configuring Port Security Features



NOTE: The features described are supported on EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Security (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

Ethernet LANs are vulnerable to attacks such as address spoofing and Layer 2 denial of service (DoS) on network devices. DHCP port security features help protect the access ports on the switch against the losses of information and productivity that can result from such attacks.

The following port security features are supported for DHCPv4:

- DHCP snooping
- Dynamic ARP inspection (DAI)
- IP source guard
- DHCP option 82

The following port security features are supported for DHCPv6:

- DHCPv6 snooping
- IPv6 Neighbor discovery inspection
- IPv6 source guard
- DHCPv6 option 37, option 18 and option 16

DHCP snooping and DHCPv6 snooping are disabled in the default configuration. There is no explicit configuration required for enabling DHCP snooping or DHCPv6 snooping. If you configure any other port security features for a VLAN at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level, then DHCP snooping and DHCPv6 snooping are automatically enabled on that VLAN.



NOTE: You can enable DHCP snooping or DHCPv6 snooping on a VLAN by configuring any CLI statement under the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level. For example, by configuring **no-option82** at the **[edit vlans *vlan-name* forwarding-options dhcp-security group *group-name* overrides]** hierarchy level, you can enable snooping without configuring any of the port security features listed above.

DAI, IPv6 neighbor discovery inspection, IP source guard, IPv6 source guard, DHCP option 82 and DHCPv6 options are configured per VLAN. You must configure a VLAN before configuring these DHCP port security features. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

The DHCP port security features that you specify for the VLAN apply to all the interfaces included within that VLAN. However, you can assign different attributes to an access interface or a group of access interfaces within the VLAN. The access interface or interfaces must first be configured as a group using the **group** statement at the **[edit vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level. A group must have at least one interface.



NOTE:

- Configuring a group of access interfaces on a VLAN automatically enables DHCP snooping for that VLAN.

Attributes that can be specified for access interfaces using the **group** statement are:

- Specifying that the interface have a static IP-MAC address (**static-ip** or **static-ipv6**)
- Specifying an access interface to act as a trusted interface to a DHCP server (**trusted**)
- Specifying an interface not to transmit DHCP option 82 (**no-option82**) or DHCPv6 options (**no-option37**)



NOTE: Trunk interfaces are trusted by default. However, on an EX9200 switch, you can override this default behavior and set a trunk interface as **untrusted**.

For additional details, see:

- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
- [Enabling IPv6 Neighbor Discovery Inspection on page 71](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)

You can override the general port security settings for the VLAN by configuring a group of access interfaces within that VLAN. For details, see:

- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 128](#)

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices](#)

Configuring Port Security (J-Web Procedure)



NOTE: This topic applies only to the J-Web Application package.

To configure port security on an EX Series switch using the J-Web interface:

1. Select **Configure > Security > Port Security**.

The VLAN List table lists all the VLAN names, VLAN identifiers, port members, and port security VLAN features.

The Interface List table lists all the ports and indicates whether security features have been enabled on the ports.



NOTE: After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Edit**—Click this option to modify the security features for the selected port or VLAN.

Enter information as specified in [Table 3 on page 24](#) to modify port security settings on VLANs.

Enter information as specified in [Table 4 on page 26](#) to modify port security settings on interfaces.

- **Activate/Deactivate**—Click this option to enable or disable security on the switch.



NOTE: This option is not supported on EX4300 switches.

- **Delete**—Click this option to delete the security features of the selected port or VLAN.



NOTE: This option is supported only on EX4300 switches.

Table 3: Port Security Settings on VLANs

| Field | Function | Your Action |
|--|---|--|
| General tab | | |
| Enable DHCP Snooping on VLAN NOTE: On EX4300 switches, DHCP snooping is enabled implicitly for all VLANs if you configure dhcp-security on one or more VLANs. | Allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. Builds and maintains a database of valid IP addresses/MAC address bindings. (By default, access ports are untrusted and trunk ports are trusted.) | Select to enable DHCP snooping on a specified VLAN or all VLANs. TIP: For private VLANs (P-VLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from P-VLAN trunk ports are not snooped. |
| Enable ARP Inspection on VLAN | Uses information in the DHCP snooping database to validate ARP packets on the LAN and protect against ARP cache poisoning. | Select to enable ARP inspection on a specified VLAN or all VLANs. (Configure any port on which you do not want ARP inspection to occur as a trusted DHCP server port.) |
| MAC movement | Number of MAC movements allowed on the given VLAN. | Enter a number. The default is unlimited. |

Table 3: Port Security Settings on VLANs (*continued*)

| Field | Function | Your Action |
|--|---|---|
| MAC movement action | Specifies the action to be taken if the MAC movement limit is exceeded. | <p>Select one of the following options:</p> <ul style="list-style-type: none"> log—Generate a system log entry, an SNMP trap, or an alarm. drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm (default). shutdown—Shut down the VLAN and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a disable timeout value. See <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i>. none—Take no action. <p>EX4300 switches have an additional option:</p> <ul style="list-style-type: none"> drop-and-log—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. |
| DHCP Groups | | |
| Group Name <i>NOTE:</i> This option is supported only on EX4300 switches. | Specifies the DHCP name of the group. | Enter a name. |
| Trusted <i>NOTE:</i> This option is supported only on EX4300 switches. | Specifies trusting DHCP packets on the selected interface. By default, trunk ports are dhcp-trusted . | To enable this option, select the check box. |
| No Option-82 <i>NOTE:</i> This option is supported only on EX4300 switches. | Enable or disable the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. | To enable this option, select the check box. |
| Interfaces <i>NOTE:</i> This option is supported only on EX4300 switches. | Specifies the DHCP interface. | Select the required interface. |
| Ports | | |
| Interface <i>NOTE:</i> This option is supported only on EX4300 switches. | Name of the interface. | Click the Edit button of the selected interface, to configure the MAC limit and the MAC limit action. |
| MAC Limit <i>NOTE:</i> This option is supported only on EX4300 switches. | Maximum number of MAC addresses learned on the interface. | Enter a number. The default is unlimited. |

Table 3: Port Security Settings on VLANs (*continued*)

| | | |
|------------------|---|--|
| MAC Limit Action | Specifies the action to be taken if the MAC move limit is exceeded. | Action to be taken when MAC limit is reached. The options are: <ul style="list-style-type: none"> • drop—Drop the packet and do not learn. Default is forward. • drop-and-log—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. • log—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry. • none—Forward the packet. • shutdown—Disable the interface and generate an alarm, an SNMP trap, or a system log entry. |
|------------------|---|--|

NOTE: This option is supported only on EX4300 switches.

Table 4: Port Security on Interfaces

| Field | Function | Your Action |
|------------------|--|---|
| Trust DHCP | Specifies trusting DHCP packets on the selected interface. By default, trunk ports are dhcp-trusted . | Select to enable DHCP trust. |
| | NOTE: This option is not supported on EX4300 switches. | |
| MAC Limit | Specifies the number of MAC addresses that can be learned on a single Layer 2 access port. This option is not valid for trunk ports. | Enter a number. |
| | NOTE: Trunk ports are supported only on EX4300 switches. | |
| MAC Limit Action | Specifies the action to be taken if the MAC limit is exceeded. This option is not valid for trunk ports. | <p>Select one of the following:</p> <ul style="list-style-type: none"> • log—Generate a system log entry, an SNMP trap, or an alarm. • drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm. (Default) • shutdown—Shut down the interface and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a disable timeout value. See <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i> • none—Take no action. <p>EX4300 switches have an additional option:</p> <ul style="list-style-type: none"> • drop-and-log—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. |

Table 4: Port Security on Interfaces (*continued*)

| Field | Function | Your Action |
|------------------|---|---|
| Allowed MAC List | Specifies the MAC addresses that are allowed for the interface. | To add a MAC address: <ol style="list-style-type: none">1. Click Add.2. Enter the MAC address.3. Click OK. |

- Related Documentation**
- [Configuring Port Security \(CLI Procedure\)](#)
 - [Example: Configuring Basic Port Security Features](#)
 - [Monitoring Port Security](#)
 - [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)

CHAPTER 2

Configuring IP Source Guard to Prevent IP Spoofing Attacks

- [Understanding IP Source Guard for Port Security on EX Series Switches on page 29](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)
- [Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38](#)

Understanding IP Source Guard for Port Security on EX Series Switches

Ethernet LAN switches are vulnerable to attacks that involve spoofing (forging) of source IP addresses or source MAC addresses. You can use the IP source guard access port security feature on Juniper Networks EX Series Ethernet Switches to mitigate the effects of these attacks.

- [IP Address Spoofing on page 29](#)
- [How IP Source Guard Works on page 29](#)
- [IPv6 Source Guard on page 30](#)
- [The DHCP Snooping Table on page 30](#)
- [Typical Uses of Other Junos OS Features with IP Source Guard on page 31](#)

IP Address Spoofing

Hosts on access interfaces can spoof source IP addresses and source MAC addresses by flooding the switch with packets containing invalid addresses. Such attacks combined with other techniques such as TCP SYN flood attacks can cause denial-of-service (DoS) attacks. With source IP address or source MAC address spoofing, the system administrator cannot identify the source of the attack. The attacker can spoof addresses on the same subnet or on a different subnet.

How IP Source Guard Works

IP source guard examines each packet sent from a host attached to an untrusted access interface on the switch. The IP address, MAC address, VLAN and interface associated with the host is checked against entries stored in the DHCP snooping database. If the

packet header does not match a valid entry in the DHCP snooping database, the switch does not forward the packet—that is, the packet is discarded.



NOTE:

- If your switch uses Junos OS for EX Series with support for the Enhanced Layer 2 Software (ELS) configuration style, DHCP snooping is enabled automatically when you enable IP source guard on a VLAN. See [“Configuring IP Source Guard \(CLI Procedure\)”](#) on page 32.
- If your switch uses Junos OS for EX Series without support the Enhanced Layer 2 Software (ELS) configuration style and you enable IP source guard on a VLAN, you must also explicitly enable DHCP snooping on that VLAN. Otherwise, the default value of no DHCP snooping applies to the VLAN. See [Configuring IP Source Guard \(CLI Procedure\)](#).

IP source guard examines packets sent from untrusted access interfaces on those VLANs. By default, on EX Series switches, access interfaces are untrusted and trunk interfaces are trusted. IP source guard does not examine packets that have been sent to the switch by devices connected to trusted interfaces so that a DHCP server can be connected to that interface to provide dynamic IP addresses.



NOTE: On an EX9200 switch, you can set a trunk interface as untrusted so that it supports IP source guard.

IPv6 Source Guard

IPv6 source guard is available on switches that support DHCPv6 snooping. To determine whether your switch supports DHCPv6 snooping, see the *EX Series Switch Software Features Overview*.

The DHCP Snooping Table

IP source guard obtains information about IP address to MAC address bindings (IP-MAC binding) from the DHCP snooping table, also known as the DHCP binding table. The DHCP snooping table is populated either through dynamic DHCP snooping or through configuration of specific static IP address to MAC address bindings. For more information about the DHCP snooping table, see [“Understanding DHCP Snooping for Port Security”](#) on page 51.

To display the DHCP snooping table, issue the operational mode command that appears in the switch CLI.

For DHCP snooping:

- (For non-ELS switches) **show ip-source-guard**
- (ELS switches only) **show dhcp-security binding**

For DHCPv6 snooping:

- (For non-ELS switches) **show dhcpv6 snooping binding**
- (ELS switches only) **show dhcp-security ipv6 binding**

Typical Uses of Other Junos OS Features with IP Source Guard

You can configure IP source guard with various other features on the EX Series switch to provide access port security, including:

- VLAN tagging (used for voice VLANs)
- GRES (graceful Routing Engine switchover)
- Virtual Chassis configurations (See *EX Series Switch Software Features Overview* for list of models that support IP Source Guard.)
- Link aggregation groups (LAGs)
- 802.1X user authentication in single supplicant, single-secure supplicant, or multiple supplicant mode.



NOTE: While implementing 801.X user authentication in single-secure supplicant or multiple supplicant mode, use the following configuration guidelines:

- If the 802.1X interface is part of an untagged MAC-based VLAN and you want to enable IP source guard and DHCP snooping on that VLAN, you must enable IP source guard and DHCP snooping on all dynamic VLANs in which the interface has untagged membership. This also applies to IPv6 source guard and DHCPv6 snooping.
- If the 802.1X interface is part of a tagged MAC-based VLAN and you want to enable IP source guard and DHCP snooping on that VLAN, you must enable IP source guard and DHCP snooping on all dynamic VLANs in which the interface has tagged membership. This also applies to IPv6 source guard and DHCPv6 snooping.

Related Documentation

- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
- *Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN*
- *Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces*
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)

- [Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38](#)

Configuring IP Source Guard (CLI Procedure)



NOTE: This task uses Junos OS for MX Series routers and EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device runs software that does not support ELS, see *Configuring IP Source Guard (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.



NOTE: On EX9200 switches, IP source guard is not supported in an MC-LAG scenario.

You can use the IP source guard access port security feature on EX Series switches to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, then IP source guard ensures that the switch does not forward the packet—that is, the packet is discarded.

You configure the IP source guard feature on a specific VLAN. When you configure IP source guard on a VLAN, the switch automatically enables DHCP snooping on that VLAN.

IPv6 source guard is supported on switches with support for DHCPv6 snooping. On these switches, configuring IP source guard or IPv6 source guard on a VLAN automatically enables DHCP snooping and DHCPv6 snooping on that VLAN.

IP source guard and IPv6 source guard can be applied only to untrusted interfaces. Access interfaces are untrusted by default.

IP source guard and IPv6 source guard can be used together with 802.1X user authentication in single supplicant, single-secure supplicant, or multiple supplicant mode.

Before you can configure IP source guard or IPv6 source guard on a VLAN, you must configure the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To configure IP source guard on a specific VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set ip-source-guard
```

To configure IPv6 source guard on a specific VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set ipv6-source-guard
```

Related Documentation

- [Verifying That IP Source Guard Is Working Correctly](#)

- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)
- [Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches on page 29](#)

Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing



NOTE: This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.



NOTE: On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

This example describes how to enable IP source guard and Dynamic ARP Inspection (DAI) on a specified VLAN to protect the switch against spoofed IP/MAC addresses and ARP spoofing attacks. When you enable either IP source guard or DAI, the configuration automatically enables DHCP snooping for the same VLAN.

- [Requirements on page 33](#)
- [Overview and Topology on page 34](#)
- [Configuration on page 35](#)
- [Verification on page 36](#)

Requirements

This example uses the following hardware and software components:

- One EX4300 switch or EX9200 switch
- Junos OS Release 13.2X50-D10 or later for EX Series switches
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure IP source guard to prevent IP/MAC spoofing or DAI to mitigate ARP spoofing attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN to which you are adding DHCP security features. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

Overview and Topology

Ethernet LAN switches are vulnerable to attacks on security that involve spoofing (forging) of source MAC addresses or source IP addresses. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch against entries stored in the DHCP snooping database. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

Another type of security attack is ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP-spoofing is a way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switch sending traffic to the proper network device, it sends it to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that should have gone to another device. The result is that traffic from the switch is misdirected and cannot reach its proper destination.



NOTE: When dynamic ARP inspection (DAI) is enabled, the switch logs the number of invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. You can use these log messages to discover ARP spoofing on the network.


This example shows how to configure these important port security features on a switch that is connected to a DHCP server. The setup for this example includes the VLAN **employee-vlan** on the switch. [Figure 1 on page 35](#) illustrates the topology for this example.



NOTE:

The trunk interface connecting to the DHCP server interface is a trusted port by default.

Figure 1: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 5 on page 35](#).

Table 5: Components of the Port Security Topology

| Properties | Settings |
|-------------------------------------|--|
| Switchhardware | One EX4300 or EX9200 switch |
| VLAN name and ID | employee-vlan , tag 20 |
| VLAN subnets | 192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is the subnet's broadcast address |
| Interfaces in employee-vlan | ge-0/0/1 , ge-0/0/2 , ge-0/0/3 , ge-0/0/8 |
| Interface connecting to DHCP server | ge-0/0/8 |

In this example, the switch has already been configured as follows:

- All access ports are untrusted, which is the default setting.
- The trunk port (**ge-0/0/8**) is trusted, which is the default setting.
- The VLAN (**employee-vlan**) has been configured to include the specified interfaces.

Configuration

To configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping) to protect the switch against IP spoofing and ARP attacks:

| | |
|--------------------------------|---|
| CLI Quick Configuration | <p>To quickly configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping), copy the following commands and paste them into the switch terminal window:</p> <pre>[edit] set vlans employee-vlan forwarding-options dhcp-security ip-source-guard set vlans employee-vlan forwarding-options dhcp-security arp-inspection</pre> |
| Step-by-Step Procedure | <p>Configure IP source guard and DAI (and thereby, also automatically configure DHCP snooping) on the VLAN:</p> <ol style="list-style-type: none">1. Configure IP source guard on the VLAN: <pre>[edit vlans employee-vlan forwarding-options dhcp-security] user@switch# set ip-source-guard</pre>2. Enable DAI on the VLAN: <pre>[edit vlans employee-vlan forwarding-options dhcp-security] user@switch# set arp-inspection</pre> |
| Results | <p>Check the results of the configuration:</p> <pre>user@switch> show vlans employee-vlan forwarding-options employee-vlan { forwarding-options { dhcp-security { arp-inspection; ip-source-guard; } } }</pre> |

Verification

Confirm that the configuration is working properly.

- [Verifying That DHCP Snooping Is Working Correctly on the Switch on page 36](#)
- [Verifying That IP Source Guard is Working on the VLAN on page 37](#)
- [Verifying That DAI Is Working Correctly on the Switch on page 37](#)

Verifying That DHCP Snooping Is Working Correctly on the Switch

| | |
|----------------|---|
| Purpose | Verify that DHCP snooping is working on the switch. |
|----------------|---|

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the port on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

user@switch> `show dhcp-security binding`

| IP Address | MAC Address | Vlan | Expires | State | Interface |
|------------|-------------------|---------------|---------|-------|------------|
| 192.0.2.17 | 00:05:85:3A:82:77 | employee-vlan | 86265 | BOUND | ge-0/0/1.0 |
| 192.0.2.18 | 00:05:85:3A:82:79 | employee-vlan | 86265 | BOUND | ge-0/0/1.0 |
| 192.0.2.19 | 00:05:85:3A:82:80 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.20 | 00:05:85:3A:82:81 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.21 | 00:05:85:3A:82:83 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.22 | 00:05:85:27:32:88 | employee-vlan | 86254 | BOUND | ge-0/0/3.0 |

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for the assigned IP address, the device's MAC address, the VLAN name, and the time, in seconds, remaining before the lease expires.

Verifying That IP Source Guard is Working on the VLAN

Purpose Verify that IP source guard is enabled and working on the VLAN.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch. View the IP source guard information for the data VLAN.

user@switch> `show dhcp-security binding ip-source-guard`

| IP Address | MAC Address | Vlan | Expires | State | Interface |
|------------|-------------------|---------------|---------|-------|------------|
| 192.0.2.17 | 00:05:85:3A:82:77 | employee-vlan | 86265 | BOUND | ge-0/0/1.0 |
| 192.0.2.18 | 00:05:85:3A:82:79 | employee-vlan | 86265 | BOUND | ge-0/0/1.0 |
| 192.0.2.19 | 00:05:85:3A:82:80 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.20 | 00:05:85:3A:82:81 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.21 | 00:05:85:3A:82:83 | employee-vlan | 86287 | BOUND | ge-0/0/2.0 |
| 192.0.2.22 | 00:05:85:27:32:88 | employee-vlan | 86254 | BOUND | ge-0/0/3.0 |

Meaning The IP source guard database table contains the VLANs enabled for IP source guard.

Verifying That DAI Is Working Correctly on the Switch

Purpose Verify that DAI is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show dhcp-security arp inspection statistics
```

ARP inspection statistics:

| Interface | Packets received | ARP inspection pass | ARP inspection failed |
|------------|------------------|---------------------|-----------------------|
| ge-0/0/1.0 | 7 | 5 | 2 |
| ge-0/0/2.0 | 10 | 10 | 0 |
| ge-0/0/3.0 | 12 | 12 | 0 |

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
 - [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
 - [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 49](#)

Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing



NOTE: This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

This example describes how to enable IPv6 source guard and neighbor discovery inspection on a specified VLAN to protect an EX Series switch against IPv6 address spoofing attacks. When you enable either IPv6 source guard or neighbor discovery inspection, DHCPv6 snooping is automatically enabled on the same VLAN.

- [Requirements on page 38](#)
- [Overview and Topology on page 39](#)
- [Configuration on page 41](#)
- [Verification on page 41](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch that supports the Enhanced Layer 2 Software configuration style.
- Junos OS Release 13.2X51-D20 or later for EX Series switches

- A DHCPv6 server to provide IPv6 addresses to network devices on the switch

Before you configure IPv6 source guard and neighbor discovery inspection to prevent IPv6 address spoofing attacks, be sure you have:

- Connected the DHCPv6 server to the switch.
- Configured the VLAN to which you are adding DHCPv6 security features. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

Overview and Topology

Ethernet LAN switches are vulnerable to attacks on security that involve spoofing (forging) of source MAC addresses or source IPv6 addresses. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. For more information on IPv6 address spoofing attacks, see [“Understanding IPv6 Neighbor Discovery Inspection” on page 69](#).

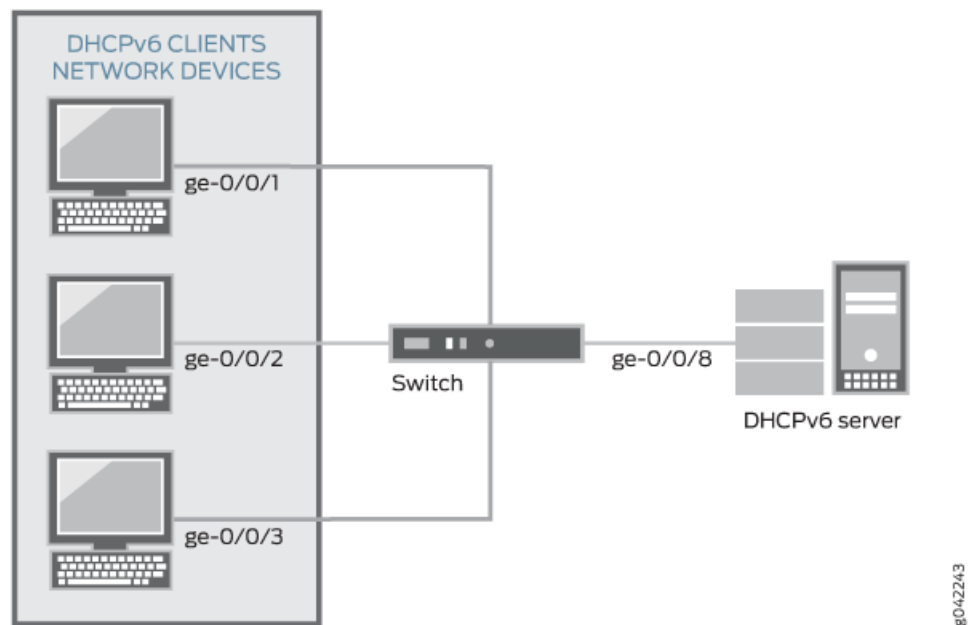
By using the DHCPv6 snooping table, also known as the binding table, IPv6 source guard and neighbor discovery inspection mitigate the risk of IPv6 spoofing attacks. The DHCPv6 snooping table contains the IP address, MAC address, VLAN and interface ID for each host associated with the VLAN. When a packet is sent from a host attached to an untrusted access interface on the switch, IPv6 source guard checks it against the entries in the DHCPv6 snooping table. If there is no match in the table, the switch does not forward the packet—that is, the packet is discarded. Neighbor discovery inspection verifies neighbor discovery messages sent between IPv6 nodes on the same network link against the DHCPv6 snooping table, and also discards the packet if no match is found.

This example shows how to configure these important port security features on a switch that is connected to a DHCPv6 server. The setup for this example includes the VLAN sales on the switch. [Figure 1 on page 35](#) illustrates the topology for this example.



NOTE: The trunk interface connecting to the DHCPv6 server interface is a trusted port by default.

Figure 2: Network Topology for Basic Port Security



The components of the topology for this example are shown in [Table 5 on page 35](#).

Table 6: Components of the Port Security Topology

| Properties | Settings |
|---------------------------------------|--|
| Switch hardware | One EX Series switch that supports the Enhanced Layer 2 Software configuration style. |
| VLAN name and ID | sales, tag 20 |
| VLAN subnets | 192.0.2.16/28 192.0.2.17 through 192.0.2.30 192.0.2.31 is the subnet's broadcast address |
| Interfaces in sales | ge-0/0/1, ge-0/0/2, ge-0/0/3, ge-0/0/8 |
| Interface connecting to DHCPv6 server | ge-0/0/8 |

In this example, the switch has already been configured as follows:

- All access ports are untrusted, which is the default setting.
- The trunk port (ge-0/0/8) is trusted, which is the default setting.
- The VLAN (sales) has been configured to include the specified interfaces.

Configuration

CLI Quick Configuration To quickly configure IPv6 source guard and neighbor discovery inspection (and thereby, also automatically configure DHCPv6 snooping), copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans sales forwarding-options dhcp-security ipv6-source-guard
set vlans sales forwarding-options dhcp-security neighbor-discovery-inspection
```

Step-by-Step Procedure Configure IPv6 source guard and neighbor discovery inspection (and thereby, also automatically configure DHCPv6 snooping) on the VLAN:

1. Configure IPv6 source guard on the VLAN:

```
[edit vlans sales forwarding-options dhcp-security]
user@switch# set ipv6-source-guard
```

2. Enable neighbor discovery inspection on the VLAN:

```
[edit vlans sales forwarding-options dhcp-security]
user@switch# set neighbor-discovery-inspection
```

Results Check the results of the configuration:

```
user@switch> show vlans sales forwarding-options
dhcp-security {
  neighbor-discovery-inspection;
  ipv6-source-guard;
}
```

Verification

Confirm that the configuration is working properly.

- [Verifying That DHCPv6 Snooping Is Working Correctly on the Switch on page 41](#)
- [Verifying That Neighbor Discovery Inspection Is Working Correctly on the Switch on page 42](#)

Verifying That DHCPv6 Snooping Is Working Correctly on the Switch

Purpose Verify that DHCPv6 snooping is working on the switch.

Action Send DHCPv6 requests from network devices (in this example, these are DHCPv6 clients) connected to the switch.

Display the DHCPv6 snooping information when the port on which the DHCPv6 server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IPv6 addresses and leases:

```
user@switch> show dhcp-security ipv6 binding
```

| IPv6 address | MAC address | Vlan | Expires | State | Interface |
|-----------------------|-------------------|--------|---------|-------|------------|
| 2001:db8:fe10:: | 00:10:94:00:55:0b | vlan20 | 3456 | BOUND | ge-0/0/1.0 |
| fe80::210:94ff:fe00:1 | 00:10:94:00:55:0b | vlan20 | 3456 | BOUND | ge-0/0/1.0 |
| 2001:db8:fe12:: | 00:10:94:00:00:34 | vlan20 | 3456 | BOUND | ge-0/0/2.0 |
| fe80::210:94ff:fe00:2 | 00:10:94:00:00:34 | vlan20 | 3456 | BOUND | ge-0/0/2.0 |
| 2001:db8:fe14:: | 00:10:94:00:00:55 | vlan20 | 3456 | BOUND | ge-0/0/3.0 |
| fe80::210:94ff:fe00:3 | 00:10:94:00:00:55 | vlan20 | 3456 | BOUND | ge-0/0/3.0 |

Meaning The output shows the assigned IPv6 addresses, the MAC address, the VLAN name, and the time, in seconds, remaining before the lease expires. Because IPv6 hosts usually have more than one IPv6 address assigned to each of their IPv6-enabled network interfaces, there are two entries added for each client: one with the link-local IPv6 address, which is used by the client for DHCP transactions, and another with the IPv6 address assigned by the server. The link-local address always has the prefix **fe80::/10**.

Verifying That Neighbor Discovery Inspection Is Working Correctly on the Switch

Purpose Verify that neighbor discovery inspection is working on the switch.

Action Send neighbor discovery packets from network devices connected to the switch.

Display the neighbor discovery information:

```
user@switch> show dhcp-security neighbor-discovery-inspection statistics
```

ND inspection statistics:

| Interface | ND Packets received | ND inspection pass | ND inspection failed |
|------------|---------------------|--------------------|----------------------|
| ge-0/0/1.0 | 7 | 5 | 2 |
| ge-0/0/2.0 | 10 | 10 | 0 |
| ge-0/0/3.0 | 12 | 12 | 0 |

Meaning The sample output shows the number of neighbor discovery packets received and inspected per interface, with a list of the number of packets that passed and the number of packets that failed the inspection on each interface. The switch compares the neighbor discovery requests and replies against the entries in the DHCPv6 snooping database. If a MAC address or IPv6 address in the neighbor discovery packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
 - [Enabling IPv6 Neighbor Discovery Inspection on page 71](#)
 - [Configuring Port Security Features on page 21](#)

CHAPTER 3

Configuring Dynamic ARP Inspection to Prevent ARP Spoofing Attacks

- [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 45](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 49](#)

Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing

Dynamic ARP inspection (DAI) protects switching devices against ARP spoofing.

DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address with entries in the database. If the media access control (MAC) address or IP address in the ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are sent to the Routing Engine and are rate-limited to protect the switching device from CPU overload.

- [Address Resolution Protocol on page 45](#)
- [ARP Spoofing on page 46](#)
- [Dynamic ARP Inspection on page 46](#)
- [Prioritizing Inspected Packets on page 47](#)

Address Resolution Protocol

Sending IP packets on a multi-access network requires mapping an IP address to an Ethernet MAC address.

Ethernet LANs use ARP to map MAC addresses to IP addresses.

The switching device maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

ARP Spoofing

ARP spoofing is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switching device sending traffic to the proper network device, it sends the traffic to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that must have gone to another device. The result is that traffic from the switching device is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an attacker can poison the ARP cache of a network device by sending an ARP response to the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, the switches examine ARP responses through DAI.

Dynamic ARP Inspection

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid. ARP probe packets are not subjected to dynamic ARP inspection. The switch always forwards such packets.

Junos OS for EX Series switches and the QFX Series uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, and therefore ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

If you set an interface to be a DHCP trusted port, it is also trusted for ARP packets.

**NOTE:**

- If your switching device is an EX Series switch and uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see [“Enabling a Trusted DHCP Server \(CLI Procedure\)” on page 128](#) for information about configuring an access interface to be a DHCP trusted port.
- If your switching device is an EX Series switch and is *not* using Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.

For packets directed to the switching device to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Packet Forwarding Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

Prioritizing Inspected Packets



NOTE: Prioritizing inspected packets is not supported on the QFX Series and the EX4600 switch.

You can use class-of-service (CoS) forwarding classes and queues to prioritize DAI packets for a specified VLAN. This type of configuration places inspected packets for that VLAN in the egress queue, that you specify, ensuring that the security procedure does not interfere with the transmission of high-priority traffic.

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
- *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)

- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 49](#)

Enabling Dynamic ARP Inspection (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#). For ELS details, see [Getting Started with Enhanced Layer 2 Software](#).



NOTE: On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

Dynamic ARP inspection (DAI) protects switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

Before you can enable DAI on a VLAN, you must configure the VLAN. See [Configuring VLANs for EX Series Switches \(CLI Procedure\)](#).

To enable DAI on a VLAN by using the CLI:

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set arp-inspection
```

Related Documentation

- [Enabling Dynamic ARP Inspection \(J-Web Procedure\) on page 49](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)
- [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 45](#)

Enabling Dynamic ARP Inspection (J-Web Procedure)

Dynamic ARP inspection (DAI) protects EX Series switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

To enable DAI on one or more VLANs by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more VLANs from the VLAN list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Enable ARP Inspection on VLAN** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.



NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

Related Documentation

- *Enabling Dynamic ARP Inspection (CLI Procedure)*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
- *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
- *Verifying That DAI Is Working Correctly*
- *Monitoring Port Security*
- [Understanding Dynamic ARP Inspection for Protecting Switching Devices Against ARP Spoofing on page 45](#)

CHAPTER 4

Configuring DHCP Snooping to Filter DHCP Messages from Untrusted Hosts

- [Understanding DHCP Snooping for Port Security on page 51](#)
- [Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)
- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)
- [Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance \(CLI Procedure\) on page 66](#)

Understanding DHCP Snooping for Port Security



NOTE: This topic includes information about enabling Dynamic Host Configuration Protocol (DHCP) snooping when using Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs Junos OS software that does not support ELS, see *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

DHCP snooping enables the switching device, which can be either a switch or a router, to monitor DHCP messages received from untrusted devices connected to the switching device. When DHCP snooping is enabled on a VLAN, the system examines DHCP messages sent from untrusted hosts associated with the VLAN and extracts their IP addresses and lease information. This information is used to build and maintain the DHCP snooping database. Only hosts that can be verified using this database are allowed access to the network.

- [DHCP Snooping Basics on page 52](#)
- [Enabling DHCP Snooping on page 53](#)
- [DHCP Snooping Process on page 53](#)

- [DHCPv6 Snooping on page 54](#)
- [Rapid Commit for DHCPv6 on page 55](#)
- [DHCP Server Access on page 55](#)
- [Static IP Address Additions to the DHCP Snooping Database on page 58](#)

DHCP Snooping Basics

DHCP allocates IP addresses dynamically, *leasing* addresses to devices so that the addresses can be reused when they are no longer needed by the devices to which they were assigned. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port).

By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping.

When DHCP snooping is enabled, the lease information from the server is used to create the DHCP snooping table, also known as the DHCP binding table. The table shows current IP-MAC address bindings, as well as lease time, type of binding, names of associated VLANs and interfaces.

Entries in the DHCP snooping table are updated in the following events:

- When a network device releases an IP address (sends a DHCPRELEASE message). In this event, the associated mapping entry is deleted from the database.
- When you move a network device from one VLAN to another. In this event, typically the device needs to acquire a new IP address. Therefore, its entry in the database, including the VLAN name, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires. In this event, the associated entry is deleted from the database.
- When the network device renews its lease by sending a unicast DHCPREQUEST message and receiving a positive response from the DHCP server. In this event, the lease time is updated in the database.
- If the network device cannot reach the DHCP server that originally granted the lease, it sends a broadcast DHCPREQUEST message and rebinds to the DHCP server that responds. In this event, the client receives a new IP address and the binding is updated in the DHCP snooping table.
- If a network device with a fixed IP allocation from the DHCP server is replaced by a new device with a different MAC address. In this event, the new IP-MAC address binding is stored until the server sends a DHCPACK message, and then the entry in the DHCP snooping table is updated with the new address binding.



TIP: By default, the IP-MAC bindings are lost when the switch is rebooted, and the DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

You can configure the switch to snoop DHCP server responses only from specific VLANs. Doing this prevents spoofing of DHCP server messages.

Enabling DHCP Snooping

DHCP snooping is not enabled in the default switch configuration. DHCP snooping is enabled automatically by Junos OS when you configure any port security features at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level. You enable DHCP snooping per VLAN, not per interface (port). For additional information about enabling DHCP snooping, see [“Configuring Port Security Features” on page 21](#).

DHCP Snooping Process

The DHCP snooping process consists of the following steps:



NOTE: When DHCP snooping is enabled for a VLAN, all DHCP packets sent from network devices in that VLAN are subjected to DHCP snooping. The final IP-MAC binding occurs when the DHCP server sends a DHCPACK packet to the DHCP client.

1. The network device sends a DHCPDISCOVER packet to request an IP address.
2. The switch forwards the packet to the DHCP server.
3. The server sends a DHCPOFFER packet to offer an address. If the DHCPOFFER packet is from a trusted interface, the switch forwards the packet to the network device.
4. The network device sends a DHCPREQUEST packet to accept the IP address. The switch adds an IP-MAC placeholder binding to the DHCP snooping table. The entry is considered a placeholder until a DHCPACK packet is received from the server. Until then, the IP address could still be assigned to some other host.



NOTE: If the client entry already exists in the DHCP snooping table with the same IP address, but a different MAC address, the switch adds an IP-MAC placeholder binding to the DHCP snooping table using the current information from the DHCPREQUEST message. This entry is updated to the DHCP snooping database after a DHCPACK of the requested IP-address is received from the server.

5. The server sends a DHCPACK packet to assign the IP address or a DHCPNAK packet to deny the address request.
6. The switch updates the DHCP database according to the type of packet received:
 - If the switch receives a DHCPACK packet, it updates lease information for the IP-MAC address bindings in its database.
 - If the switch receives a DHCPNACK packet, it deletes the placeholder.



NOTE: The DHCP database is updated only after the DHCPREQUEST packet is sent.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the [Junos OS System Basics Configuration Guide](#).

DHCPv6 Snooping

DHCP snooping is also supported for IPv6 packets. The process for DHCPv6 snooping is similar to that for DHCP snooping, but uses different names for the messages exchanged between the client and server to assign IPv6 addresses. [Table 7 on page 54](#) shows DHCPv6 messages and their DHCPv4 equivalents.

Table 7: DHCPv6 Messages and DHCPv4 Equivalent Messages

| Sent by | DHCPv6 Messages | DHCPv4 Equivalent Messages |
|---------|-------------------------|----------------------------|
| Client | SOLICIT | DHCPDISCOVER |
| Server | ADVERTISE | DHCPOFFER |
| Client | REQUEST, RENEW, REBIND | DHCPREQUEST |
| Server | REPLY | DHCPACK/DHCPNAK |
| Client | RELEASE | DHCPRELEASE |
| Client | INFORMATION-REQUEST | DHCPINFORM |
| Client | DECLINE | DHCPDECLINE |
| Client | CONFIRM | none |
| Server | RECONFIGURE | DHCPFORCERENEW |
| Client | RELAY-FORW, RELAY-REPLY | none |

Rapid Commit for DHCPv6

The DHCPv6 Rapid Commit option can shorten the exchange of messages between the client and server. When supported by the server and set by the client, this option shortens the exchange from a four-way relay to a two-message handshake. For more information about enabling the Rapid Commit option, see *Enabling DHCPv6 Rapid Commit Support*.

When the Rapid Commit option is enabled, the exchange of messages is as follows:

1. The DHCPv6 client sends out a SOLICIT message that contains a request that rapid assignment of address, prefix, and other configuration parameters be preferred.
2. If the DHCPv6 server supports rapid assignment, it responds with a REPLY message, which contains the assigned IPv6 address and prefix and other configuration parameters.

DHCP Server Access

A switch's access to the DHCP server can be configured in three ways:

- [Switch, DHCP Clients, and the DHCP Server Are All on the Same VLAN on page 55](#)
- [Switch Acts as the DHCP Server on page 56](#)
- [Switch Acts as a Relay Agent on page 57](#)

Switch, DHCP Clients, and the DHCP Server Are All on the Same VLAN

When the switch, DHCP clients, and DHCP server are *all members of the same VLAN*, the DHCP server can be connected to the switch in one of two ways:



NOTE: To enable DHCP snooping on the VLAN, configure the `dhcp-security` statement at the `[edit vlans vlan-name forwarding-options]` hierarchy.

- (See [Figure 3 on page 56](#).) The server is directly connected to the same switch as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port.
- (See [Figure 4 on page 56](#).) The server is connected to an intermediary switch (Switch 2) that is connected through a trunk port to the switch (Switch 1) that the DHCP clients are connected to. Switch 2 is being used as a transit switch. The VLAN is enabled for DHCP snooping to protect the untrusted access ports of Switch 1. The trunk port is configured by default as a trusted port. In [Figure 4 on page 56](#), ge-0/0/11 is a trusted trunk port.

Figure 3: DHCP Server Connected Directly to a Switch






Figure 4: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port



Switch Acts as the DHCP Server

You can configure DHCP local server options on the switch, which enables the switch to function as an extended DHCP local server. In [Figure 5 on page 57](#), the DHCP clients are connected to the extended DHCP local server through untrusted access ports.

Figure 5: Switch Is the DHCP Server




Switch Acts as a Relay Agent

The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface (on a switch or a router). The Layer 3 interfaces on the switch are configured as routed VLAN interfaces (RVIs)—also called integrated routing and bridging (IRB) interfaces. The trunk interfaces are trusted by default.

The switch can act as a relay agent in these two scenarios:

- The DHCP server and clients are in different VLANs.
- The switch is connected to a router that is, in turn, connected to the DHCP server. See [Figure 6 on page 58](#).

Figure 6: Switch Acting as a Relay Agent Through a Router to the DHCP Server



Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add a static IP address, you provide the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. You do not assign a lease time to the entry. The statically configured entry never expires.

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring Port Security Features on page 21](#)
- [Understanding Trusted DHCP Servers for Port Security on page 127](#)
- [Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59](#)
- [Understanding DHCP Services for Switches](#)
- [DHCP/BOOTP Relay for Switches Overview](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 128](#)
- [Enabling DHCP Snooping \(J-Web Procedure\)](#)
- [Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance \(CLI Procedure\) on page 66](#)

Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect Juniper Networks EX Series Ethernet Switches and MX Series 3D Universal Edge Routers against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on an Ethernet LAN switching device send requests for IP addresses to access the Internet. The switching device forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to penetrate the network by address spoofing.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Junos OS implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Overview on page 59](#)
- [Suboption Components of Option 82 on page 60](#)
- [Switching Device Configurations That Support Option 82 on page 61](#)
- [DHCPv6 Options on page 62](#)

DHCP Option 82 Overview

If DHCP option 82 is enabled on a VLAN or bridge domain, then when a network device—a DHCP client—that is connected to the VLAN or bridge domain on an untrusted interface sends a DHCP request, the switching device inserts information about the client's network location into the packet header of that request. The switching device then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “[Suboption Components of Option 82](#)” on page 60 for more information about option 82.



NOTE: On EX4300 switches, DHCP option 82 information is added to DHCP packets received on trusted interfaces as well as untrusted interfaces.

If option 82 is enabled on a VLAN or bridge domain, the following sequence of events occurs when a DHCP client sends a DHCP request:

1. The switching device receives the request and inserts the option 82 information in the packet header.
2. The switching device forwards (or relays) the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response to the switching device. It does not alter the option 82 information.

4. The switching device strips the option 82 information from the response packet.
5. The switching device forwards the response packet to the client.

To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If the DHCP server is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information for setting parameters and it does not echo the information in its response message.



NOTE: If your switching device is an EX Series switch and uses Junos OS with Enhanced Layer 2 Software (ELS) configuration style, you can enable DHCP option 82 only for a specific VLAN. See [“Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)”](#) on page 63.

If your switching device is an EX Series switch and does *not* use Junos OS with Enhanced Layer 2 Software (ELS) configuration style, you can enable DHCP option 82 either for a specific VLAN or for all VLANs. See [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#).

Suboption Components of Option 82

Option 82 as implemented on a switching device comprises the suboptions circuit ID, remote ID, and vendor ID. These suboptions are fields in the packet header:

- circuit ID—Identifies the circuit (interface or VLAN) on the switching device on which the request was received. The circuit ID contains the interface name and VLAN name, with the two elements separated by a colon—for example, ge-0/0/10:vlan1, where ge-0/0/10 is the interface name and vlan1 is the VLAN name. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, ge-0/0/10.

Use the prefix option to add an optional prefix to the circuit ID. If you enable the prefix option, the hostname for the switching device is used as the prefix; for example, device1:ge-0/0/10:vlan1, where device1 is the hostname.

You can also specify that the interface description be used rather than the interface name or that the VLAN ID be used rather than the VLAN name.

- remote ID—Identifies the remote host. See [remote-id](#) for details.
- vendor ID—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value Juniper is used. To specify a value, you type a character string.

Switching Device Configurations That Support Option 82

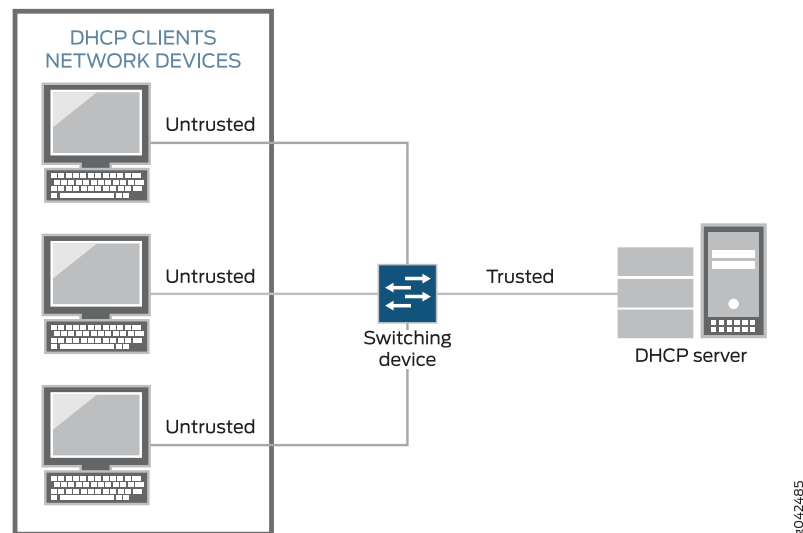
Switching device configurations that support option 82 are:

- [Switching Device, DHCP Clients, and the DHCP Server Are on the Same VLAN or Bridge Domain on page 61](#)
- [Switching Device Acts as a Relay Agent on page 61](#)

Switching Device, DHCP Clients, and the DHCP Server Are on the Same VLAN or Bridge Domain

If the switching device, the DHCP clients, and the DHCP server are all on the same VLAN or bridge domain, the switching device forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. See [Figure 7 on page 61](#).


Figure 7: DHCP Clients, Switching Device, and the DHCP Server Are All on the Same VLAN or Bridge Domain



Switching Device Acts as a Relay Agent

The switching device functions as a relay agent (extended relay server) when the DHCP clients or the DHCP server is connected to the switching device through a Layer 3 interface. On the switching device, these interfaces are configured as routed VLAN interfaces (RVIs). [Figure 8 on page 62](#) illustrates a scenario for the switching device acting as an extended relay server; in this instance, the switching device relays requests to the server.

Figure 8: Switching Device Acting as an Extended Relay Server



DHCPv6 Options



NOTE: MX Series routers do not support DHCPv6.

DHCPv6 provides several options that can be used to insert information into the DHCPv6 request packets that are relayed to a server from a client. These options are equivalent to the sub-options of DHCP option 82.

- Option 37—Identifies the remote host. Option 37 is equivalent to the **remote-id** sub-option of DHCP option 82.
- Option 18—Identifies the interface on which the DHCP request packet was received from the client. Option 18 is equivalent to the **circuit-id** sub-option of DHCP option 82.
- Option 16—Identifies the vendor of the hardware on which the client is hosted. Option 16 is equivalent to the **vendor-id** sub-option of DHCP option 82.

DHCPv6 options are not enabled automatically when DHCPv6 snooping is enabled on a VLAN. They must be configured using the **dhcpv6-options** statement.

Related Documentation

- [Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks \(CLI Procedure\)](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)

Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's responses to the clients. This topic describes this configuration.
- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as integrated routing and bridging (IRB) interfaces. The switch relays the clients' requests to the server and then forwards the server's responses to the clients. This configuration is described in *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*.

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.



NOTE: Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configure a VLAN on the switch and associate the interfaces on which the clients and the server connect to the switch with that VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*

To configure DHCP option 82:

1. Specify DHCP option 82 for the VLAN that you configured.

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set option-82
```



NOTE: If you want to enable DHCP option 82 on all VLANs, you must configure it separately for each specific VLAN.

The remaining steps are optional.

2. Configure the prefix for the circuit ID suboption to include the switch's hostname or the routing instance name for the VLAN:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id prefix (host-name | routing-instance-name)
```

3. Specify that the circuit ID suboption value contains the interface description rather than the interface name (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id use-interface-description
```



NOTE: When you use the interface description rather than the interface name, the interface description has to be specified under interface unit ("set interfaces ge-0/0/0 unit 0 description "client"). If you do not do this, then the interface name is used.

4. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set circuit-id use-vlan-id
```

5. Specify that the remote ID suboption is included in the DHCP option 82 information:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id
```



NOTE: If you do not specify a keyword after *remote-id*, the default value for the *remote-id* suboption is the interface name.

6. Specify that the remote ID suboption is the hostname of the switch:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id host-name
```

7. Specify that the remote ID suboption value contains the interface description:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id use-interface-description
```

8. Specify that the remote ID suboption value contains a character string:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]  
user@switch# set remote-id use-string mystring
```


9. Configure a vendor ID suboption:

- To use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]
user@switch# set vendor-id
```

- To configure that the vendor ID suboption value contains a character string value that you specify rather than **Juniper** (the default):

```
[edit vlans vlan-name forwarding-options dhcp-security option-82]
user@switch# set vendor-id use-string mystring
```

Related Documentation

- [Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59](#)
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database. These bindings are labeled *static* in the database, while those bindings that have been added through the process of DHCP snooping are labeled *dynamic*. Static IPv6 address assignment is also available for DHCPv6.

Before you can perform this procedure, you must configure the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To configure a static IP address to MAC address (IP-MAC) binding in the DHCP snooping database, you must first create a group of access interfaces under the **[edit vlans vlan-name forwarding-options dhcp-security]** hierarchy. Creating this group automatically enables DHCP snooping, which is a prerequisite for creating the DHCP snooping database. You can then configure a specific interface within the group to have a static IP address that is bound to a fixed MAC address. If you want to have multiple static IP addresses, configure additional interfaces within the same group.



NOTE: On switches that support DHCPv6, creating the group of interfaces will automatically enable both DHCP and DHCPv6 snooping.

To configure a static IP-MAC address binding in the DHCP snooping database:

- [edit vlans *vlan-name* forwarding-options dhcp-security]
user@switch# set group *group-name* interface *interface-name* static-ip *ip-address* mac *mac-address*

To configure a static IPv6-MAC address binding in the DHCPv6 snooping database:

- [edit vlans *vlan-name* forwarding-options dhcp-security]
user@switch# set group *group-name* interface *interface-name* static-ipv6 *ip-address* mac *mac-address*

**Related
Documentation**

- [show dhcp-security binding on page 249](#)
- *Verifying That DHCP Snooping Is Working Correctly*
- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*

Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)



NOTE: This task uses Junos OS for MX Series routers and EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Making IP-MAC Bindings in the DHCP Snooping Database Persistent](#). For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

By default, IP-MAC address bindings in the DHCP snooping database do not persist through device reboots. You can improve network performance by configuring the IP-MAC address bindings in the DHCP snooping database to persist through reboots so that the table does not need to be rebuilt after rebooting. Do this by configuring a storage location for the DHCP snooping database file, where you must specify how frequently the device writes the database entries into the DHCP snooping database file.



NOTE: You can also configure persistent bindings for IPv6 addresses and MAC addresses on devices that support DHCPv6 snooping.

DHCPv6 is not supported on the MX Series routers.

The DHCP snooping database of IP-MAC bindings is created when you enable any of the port security features for a specific VLAN or bridge domain in either of the following hierarchy levels:

- [edit vlans *vlan-name* forwarding-options dhcp-security]
- [edit bridge-domains *bridge-domain-name* forwarding-options dhcp-security]

On devices that support DHCPv6, enabling any port security features will automatically enable DHCPv6 snooping. DHCP snooping and DHCPv6 snooping are not enabled by default.

To configure a *local* storage location for the DHCP snooping database file:

- For DHCP snooping:

```
[edit system processes]
user@device# set dhcp-service dhcp-snooping-file local-pathname write-interval seconds
For example:
```

```
[edit system processes]
user@device# set dhcp-service dhcp-snooping-file /var/tmp/test.log write-interval 60
```

- For DHCPv6 snooping:

```
[edit system processes]
user@device# set dhcp-service dhcpv6-snooping-file local-pathname write-interval seconds
For example:
```

```
[edit system processes]
user@device# set dhcp-service dhcpv6-snooping-file /var/tmp/test.log write-interval 60
```

To configure a *remote* storage location for IP-MAC bindings, use **tftp://ip-address** or **ftp://hostname/path** as the remote URL, or the local pathname for the storage location of the DHCP or DHCPv6 snooping database file:

- For DHCP snooping:

```
[edit system processes]
user@device# set dhcp-service dhcp-snooping-file remote_url write-interval seconds
For example:
```

```
[edit system processes]
user@device# set dhcp-service dhcp-snooping-file tftp://@14.1.2.1 write-interval 60
```

- For DHCPv6 snooping:

```
[edit system processes]
user@device# set dhcp-service dhcpv6-snooping-file remote_url write-interval seconds
For example:
```

```
[edit system processes]
user@device# set dhcp-service dhcpv6-snooping-file tftp://@14.1.2.1 write-interval 60
```

**Related
Documentation**

- *Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices*

CHAPTER 5

Enabling IPv6 ND Inspection and RA Guard to Prevent IPv6 Spoofing Attacks

- [Understanding IPv6 Neighbor Discovery Inspection on page 69](#)
- [Enabling IPv6 Neighbor Discovery Inspection on page 71](#)
- [Understanding IPv6 Router Advertisement Guard on page 71](#)
- [Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74](#)
- [Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79](#)

Understanding IPv6 Neighbor Discovery Inspection

IPv6 nodes (hosts and routers) use Neighbor Discovery Protocol (NDP) to discover the presence and link-layer addresses of other nodes residing on the same link. Hosts use NDP to find neighboring routers that are willing to forward packets on their behalf, while routers use it to advertise their presence. Nodes also use NDP to maintain reachability information about the paths to active neighbors. When a router or the path to a router fails, a host can search for alternate paths.

The NDP process is based on the exchange of neighbor solicitation and advertisement messages. NDP messages are unsecured, which makes NDP susceptible to attacks that involve the spoofing (or forging) of link-layer addresses. An attacking node can cause packets for legitimate nodes to be sent to some other link-layer address by either sending a neighbor solicitation message with a spoofed source MAC address, or by sending a neighbor advertisement address with a spoofed target MAC address. The spoofed MAC address is then associated with a legitimate network IPv6 address by the other nodes.

IPv6 neighbor discovery inspection is based on DHCPv6 snooping; it mitigates NDP security vulnerabilities by inspecting neighbor discovery messages and verifying them against the DHCPv6 snooping table. The DHCPv6 snooping table, which is built by snooping DHCPv6 message exchanges, includes the IPv6 address, MAC address, VLAN and interface for each host associated with the VLAN. When a neighbor discovery message is received on an untrusted interface, neighbor discovery inspection discards the packet unless the source IPv6 and MAC addresses, VLAN, and interface can be matched to an entry in the DHCPv6 snooping table. Entries can be added to the DHCPv6 snooping table by configuring the `static-ipv6` CLI statement.



NOTE: Neighbor discovery messages are always allowed on trusted interfaces.

Neighbor discovery inspection verifies five different ICMPv6 message types: Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement, and Redirect. By discarding message packets that can not be verified against the DHCPv6 snooping table, neighbor discovery inspection can prevent the following types of attacks:

- Cache poisoning attacks—Neighbor discovery cache poisoning is the IPv6 equivalent of ARP spoofing, in which an attacker uses a forged address to send an unsolicited advertisement to other hosts on the network, for associating its own MAC address with a legitimate network IP address. These bindings between IPv6 addresses and MAC addresses are stored by each node in its neighbor cache. Once the caches are updated with the malicious bindings, the attacker can initiate a man-in-the-middle attack, intercepting traffic that was intended for a legitimate host.
- Routing denial-of-service (DoS) attacks—An attacker could cause a host to disable its first-hop router by spoofing the address of a router and sending a neighbor advertisement message with the *router* flag cleared. The victim host assumes that the device that used to be its first-hop router is no longer a router.
- Redirect attacks—Routers use ICMPv6 redirect requests to inform a host of a more efficient route to a destination. Hosts can be redirected to a better first-hop router, but can also be informed by a Router Redirect message that the destination is in fact a neighbor. An attacker using this provision can achieve an effect similar to cache poisoning and intercept all traffic from the victim host. Neighbor discovery inspection checks that Router Redirect messages are sent only by trusted routers.

**Related
Documentation**

- [IPv6 Neighbor Discovery Protocol Overview](#)
- [Enabling IPv6 Neighbor Discovery Inspection on page 71](#)
- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring Port Security \(CLI Procedure\)](#)
- [Understanding DHCP Snooping for Port Security on page 51](#)

Enabling IPv6 Neighbor Discovery Inspection



NOTE: This procedure uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch uses software that does not support ELS, see *Configuring Port Security (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

IPv6 neighbor discovery inspection protects switches against IPv6 address spoofing. Neighbor discovery inspection validates IPv6 packets carrying neighbor discovery messages against the DHCPv6 binding table. The source IP address, source MAC address, VLAN and interface ID of each packet are checked against the table, and if a valid match is not found, the packet is dropped.

Before you can enable neighbor discovery inspection on a VLAN, you must configure the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To enable neighbor discovery inspection on a VLAN:

```
[edit vlans vlan-name forwarding-options dhcp-security]  
user@switch# set neighbor-discovery-inspection
```



NOTE: DHCPv6 snooping is enabled automatically when neighbor discovery inspection is configured. There is no explicit configuration required for DHCPv6 snooping.

Related Documentation

- [Understanding IPv6 Neighbor Discovery Inspection on page 69](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33](#)
- [Understanding DHCP Snooping for Port Security on page 51](#)

Understanding IPv6 Router Advertisement Guard

In an IPv6 deployment, routers periodically multicast Router Advertisement (RA) messages to announce their availability and convey information to neighboring nodes that enable them to be automatically configured on the network. RA messages are used by Neighbor Discovery Protocol (NDP) to detect neighbors, advertise IPv6 prefixes, assist in address provisioning, and share link parameters such as maximum transmission unit (MTU), hop limit, advertisement intervals, and lifetime. Hosts listen for RA messages for IPv6 address autoconfiguration and discovery of link-local addresses of the neighboring routers, and can also send a Router Solicitation (RS) message to request immediate advertisements.

RA messages are unsecured, which makes them susceptible to attacks on the network that involve the spoofing (or forging) of link-layer addresses. Also, unintended

misconfiguration by users or administrators might lead to the presence of unwanted, or rogue, RA messages, which can cause operational problems for neighboring hosts. You can configure IPv6 Router Advertisement (RA) guard to protect your network against rogue RA messages generated by unauthorized or improperly configured routers connecting to the network segment.

RA guard works by validating RA messages on the basis of whether they meet certain criteria, configured on the switch using policies. RA guard inspects RA messages and compares the information contained in the message attributes to the configured policy. Depending on the policy, RA guard either drops or forwards the RA messages that match the conditions.

The following information contained in RA message attributes can be used by RA guard to validate the source of the RA message:

- Source MAC address
- Source IPv6 address
- Source IPv6 address prefix
- Hop-count limit
- Router preference priority
- *Managed* configuration flag
- *Other* configuration flag

You can configure RA guard to operate in either stateless or stateful mode. In stateless mode, in the default state, an RA message that is received on an interface is examined and filtered on the basis of whether it matches the conditions configured in the policy attached to that interface. If the content of the RA message is validated, it forwards the RA message to its destination; otherwise, the RA message is dropped. The state of an interface operating in stateless mode can be changed by configuration. If the interface is configured as *trusted*, all RA messages are forwarded without being validated against the policy. If the interface is configured as *blocked*, all RA messages are dropped without being validated against the policy.

In stateful mode, an interface can dynamically transition from one state to another based on information gathered during a learning period. During this period, known as the *learning* state, ingress RA messages are validated against a policy to determine which interfaces are attached to links with valid IPv6 routers. At the end of the learning period, interfaces attached to legitimate senders of RA messages transition dynamically to the *forwarding* state, in which RA messages are forwarded if they can be validated against a policy. Interfaces that do not receive valid RA messages during the learning period transition dynamically to the *blocked* state, in which all ingress RA messages are dropped.

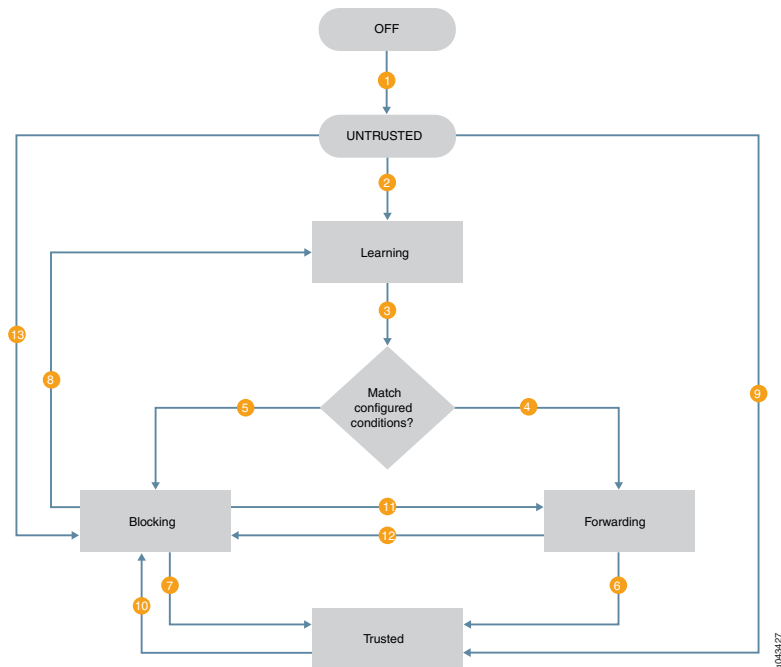
[Table 8 on page 73](#) summarizes the states of IPv6 RA guard for both stateless and stateful mode.

Table 8: IPv6 RA guard states

| State | Description | Mode |
|------------|--|--------------------|
| Off | The interface operates as if RA guard is not available. | Stateless/stateful |
| Untrusted | The interface forwards ingress RA messages if received RA messages are validated against the configured policy rules; otherwise, it drops the RA message. Untrusted state is the default state of an interface enabled for RA guard. | Stateless/stateful |
| Blocked | The interface blocks ingress RA messages. | Stateless/stateful |
| Forwarding | The interface forwards ingress RA messages if received RA messages are validated against the configured policy rules; otherwise, it drops the RA messages. | Stateful |
| Learning | The switch actively acquires information about the IPv6 routing device connected to the interface. The learning process takes place over a predefined period of time. | Stateful |
| Trusted | The interface forwards all RA messages directly, without validating them against the policy. | Stateless/stateful |

Figure 9 on page 73 illustrates the transition of states when stateful RA guard is enabled. The numbers shown on the illustrations are described in the text that follows; these are not sequential steps.

Figure 9: Stateful RA Guard State Transitions



1. When RA guard is enabled on an interface it moves to the *untrusted* state from the *off* state. The *untrusted* state is the default state of an interface that is enabled for RA guard.
2. When the command requesting the learning state is issued, the interface is moved from the *off* state to the *learning* state.
3. RA messages received during the learning state are compared to the configured policy.
4. If RA messages are validated against the configured policy, the interface moves to *forwarding* state.
5. If RA messages are not validated against the configured policy, the interface moves to *blocked* state.
6. If **mark-interface trust** is configured on the validated interface, then it moves from *forwarding* state to *trusted* state.
7. If **mark-interface trust** is configured on the blocked interface, then it moves from *blocked* state to *trusted* state.
8. If learning is requested on a blocked interface, then the interface moves from the *blocked* state to the *learning* state.
9. If an interface in the default *untrusted* state is configured as **mark-interface trust**, it moves directly to the *trusted* state. In this case a policy can not be applied on that interface.
10. If the **mark-interface trust** configuration is deleted, and no valid RAs are received on the interface, then the interface moves to the *blocked* state.
11. If the command requesting the forwarding state is issued, then the interface moves directly from *blocked* to *forwarding* state.
12. If the command requesting the blocking state is issued, then the interface moves directly from *forwarding* to *blocked*.
13. If an interface in the default *untrusted* state is configured as **mark-interface block**, it moves directly to the *blocked* state. In this case a policy can not be applied on that interface.

**Related
Documentation**

- [IPv6 Neighbor Discovery Protocol Overview](#)
- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring Port Security \(CLI Procedure\)](#)

Configuring Stateless IPv6 Router Advertisement Guard on Switches

Stateless IPv6 Router Advertisement (RA) guard enables the switch to examine incoming RA messages and filter them based on a predefined set of criteria. If the switch validates the content of the RA message, it forwards the RA message to its destination; otherwise, the RA message is dropped.

Before you can enable IPv6 RA guard, you must configure a policy with the criteria to be used for validating RA messages received on an interface. You can configure the policy to either accept or discard RA messages on the basis of whether they meet the criteria. The criteria are compared to information included in the RA messages. If the criteria for the policy includes source addresses or address prefixes, you must configure a list of the addresses before configuring the policy.

- [Configuring a Discard Policy for RA Guard on page 75](#)
- [Configuring an Accept Policy for RA Guard on page 76](#)
- [Enabling Stateless RA Guard on an Interface on page 78](#)
- [Enabling Stateless RA Guard on a VLAN on page 78](#)
- [Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard on page 79](#)

Configuring a Discard Policy for RA Guard

You can configure a discard policy to drop RA messages from predefined sources. You must first configure a list or lists of the source addresses or address prefixes, and then associate them with a policy. The following lists can be associated with discard policy:

- **ipv6-source-address-list**
- **source-mac-address-list**
- **ipv6-prefix-list**



NOTE: You can include more than one type of list in a discard policy. If the information contained in a received RA message matches any one of the list parameters, then that RA message is discarded.

To configure a discard policy for RA guard:

1. Define one or more lists of disallowed source addresses or address prefixes that RA guard will use to filter incoming RA messages. Add one address or address prefix per line in the configuration.
 - To define a list of IPv6 source addresses:

```
[edit]
user@switch# set policy-options prefix-list address-list-name ipv6-address
```
 - To define a list of IPv6 address prefixes:

```
[edit]
user@switch# set policy-options prefix-list prefix-list-name ipv6-prefix
```
 - To define a list of MAC source addresses:

```
[edit]
user@switch# set policy-options mac-list address-list-name mac-address
```
2. Configure the policy name:

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard policy  
policy-name
```

3. Specify the discard action:

```
[edit forwarding-options access-security router-advertisement-guard policy  
policy-name]  
user@switch# set discard
```

4. Associate the policy with the list or lists defined in 1. For example, to discard RA messages that match a source MAC address in the list:

```
[edit forwarding-options access-security router-advertisement-guard policy  
policy-name discard]  
user@switch# set source-mac-address-list address-list-name
```

Configuring an Accept Policy for RA Guard

You can configure an accept policy to forward RA messages on the basis of certain criteria. You can configure either match lists of source address or address prefixes as the criteria, or you can other match conditions, such as hop limit, configuration flags, or router preference as the criteria.

The following lists can be associated with an accept policy by using the **match-list** option:

- **ipv6-source-address-list**
- **source-mac-address-list**
- **ipv6-prefix-list**



NOTE: You can associate more than one type of match list with an accept policy. If the **match-all** suboption is configured, then a received RA message must match all configured match lists in order to be forwarded; otherwise, it is discarded. If the **match-any** option is configured, then a received RA message must match any one of the configured match lists in order to be forwarded; if it does not match any of the configured lists, then it is discarded.

The following match conditions can be configured using the **match-option** option:

- **hop-limit**—Configure the RA guard policy to verify the minimum or maximum hop count for an incoming RA message.
- **managed-config-flag**—Configure the RA guard policy to verify that the managed address configuration flag of an incoming RA message is set.
- **other-config-flag**—Configure the RA guard policy to verify that the other configuration flag of an incoming RA message is set.
- **router-preference-maximum**—Configure the RA guard policy to verify that the default router preference parameter value of an incoming RA message is lower than or equal to a specified limit.



NOTE: The **match-list** and **match-option** options are used only in accept policies, not in discard policies.

To configure an accept policy for RA guard by using the **match-list** option:

1. Define one or more lists of authorized source addresses or address prefixes that RA guard will use to filter incoming RA messages. Add one address or address prefix per line in the configuration.

- To define a list of IPv6 source addresses:

```
[edit]
user@switch# set policy-options prefix-list address-list-name ipv6-address
```

- To define a list of IPv6 address prefixes:

```
[edit]
user@switch# set policy-options prefix-list prefix-list-name ipv6-prefix
```

- To define a list of MAC source addresses:

```
[edit]
user@switch# set policy-options mac-list address-list-name mac-address
```

2. Specify the policy name:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard policy
policy-name
```

3. Specify the accept action:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name]
user@switch# set accept
```

4. Specify whether RA guard must meet the criteria in all lists or in any of the lists configured in 1:

- To match on all lists:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
user@switch# set match-list match-all
```

- To match on any of the lists:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
user@switch# set match-list match-any
```

5. Associate the accept policy with the list or lists configured in Step 1. For example:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept match-list match-all]
user@switch# set source-mac-address-list address-list-name
```

To configure an accept policy for RA guard using the **match-option** option:

1. Specify the policy name:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard policy
policy-name
```

2. Specify the accept action:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name]
```

```
user@switch# set accept
```

3. Specify the match conditions by using the **match-option** option. For example, to specify a match on the maximum number of hops:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
```

```
user@switch# set match-option hop-limit maximum value
```

Enabling Stateless RA Guard on an Interface

You can enable stateless RA guard on an interface. You must first configure a policy, which is applied to incoming RA messages on the interface or interfaces. To configure an accept policy, see [“Configuring an Accept Policy for RA Guard” on page 76](#). To configure a discard policy, see [“Configuring a Discard Policy for RA Guard” on page 75](#). After you apply a policy to an interface, you must also enable RA guard on the corresponding VLAN; otherwise, the policy applied to the interface does not have any impact on received RA packets.

To enable stateless RA guard on an interface:

1. Apply a policy to an interface:

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name policy policy-name
```

2. Configure the **stateless** option on the interface:

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name policy policy-name stateless
```

3. Enable stateless RA guard on the corresponding VLAN:

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name
```

Enabling Stateless RA Guard on a VLAN

You can enable stateless RA guard on a VLAN. You must first configure a policy, which is used to validate incoming RA messages in the learning state. To configure an accept policy, see [“Configuring an Accept Policy for RA Guard” on page 76](#). To configure a discard policy, see [“Configuring a Discard Policy for RA Guard” on page 75](#).

To enable stateless RA guard on a VLAN:

1. Apply a policy to a VLAN.

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name policy policy-name
```

2. Configure the **stateless** option on the VLAN:

```
[edit]
```

```
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name stateless
```

Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard

You can configure an interface as trusted or blocked to bypass inspection of RA messages by RA guard. When an RA message is received on a trusted or blocked interface, it is not subject to validation against the configured policy. A trusted interface forwards all RA messages. A blocked interface discards all RA messages.

- To configure an interface as trusted:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name mark-interface trusted
```

- To configure an interface as blocked:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name mark-interface block
```

Related Documentation

- [Understanding IPv6 Router Advertisement Guard on page 71](#)
- [Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79](#)

Configuring Stateful IPv6 Router Advertisement Guard on Switches

Stateful IPv6 Router Advertisement (RA) guard enables a switch to learn about the sources of RA messages for a certain period of time. During this period, during which the switch is known to be in the learning state, the information contained in received RA message attributes is stored and compared to the policy. At the end of the learning period, the switch has a record of which interfaces are attached to links with valid IPv6 routers. If there is no valid IPv6 router attached to an interface, the switch dynamically transitions the interface from the learning state into the blocking state. Subsequent RA messages received after the transition to blocking state are dropped. If there is a valid IPv6 router attached to the interface, the interface transitions to the forwarding state. In the forwarding state, RA messages that can be validated against the configured policy are forwarded.

You can override the dynamic state transitions by statically configuring the forwarding or blocking states on an interface. When you statically configure the state on an interface, the state can be changed only through configuration. For example, if you configure the forwarding state on an interface, the interface remains in the forwarding state until you configure a different state on that interface.

Before you can enable IPv6 RA guard on an interface or a VLAN, you must configure a policy. Stateful RA guard uses the policy to determine whether the RA messages received on an interface are from valid senders. You can configure the policy to either accept or discard RA messages that meet the predefined criteria. If the criteria for the policy includes source addresses or address prefixes, you must configure a list of the addresses before configuring the policy.

- [Configuring a Discard Policy for RA Guard on page 80](#)
- [Configuring an Accept Policy for RA Guard on page 81](#)

- [Enabling Stateful RA Guard on an Interface on page 83](#)
- [Enabling Stateful RA Guard on a VLAN on page 83](#)
- [Configuring the Learning State on an Interface on page 83](#)
- [Configuring the Forwarding State on an Interface on page 85](#)
- [Configuring the Blocking State on an Interface on page 85](#)
- [Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard on page 85](#)

Configuring a Discard Policy for RA Guard

You can configure a discard policy to drop RA messages from predefined sources. You must first configure a list or lists of the source addresses or address prefixes, and then associate them with a policy. The following lists can be associated with discard policy:

- **ipv6-source-address-list**
- **source-mac-address-list**
- **ipv6-prefix-list**



NOTE: You can include more than one type of list in a discard policy. If the information contained in a received RA message matches any one of the list parameters, then that RA message is discarded.

To configure a discard policy for RA guard:

1. Define one or more lists of disallowed source addresses or address prefixes that RA guard will use to filter incoming RA messages. Add one address or address prefix per line in the configuration.
 - To define a list of IPv6 source addresses:

```
[edit]  
user@switch# set policy-options prefix-list address-list-name ipv6-address
```
 - To define a list of IPv6 address prefixes:

```
[edit]  
user@switch# set policy-options prefix-list prefix-list-name ipv6-prefix
```
 - To define a list of MAC source addresses:

```
[edit]  
user@switch# set policy-options mac-list address-list-name mac-address
```
2. Configure the policy name:

```
[edit]  
user@switch# set forwarding-options access-security router-advertisement-guard policy  
policy-name
```
3. Specify the discard action:

```
[edit forwarding-options access-security router-advertisement-guard policy  
policy-name]
```



```
user@switch# set discard
```

4. Associate the policy with the list or lists defined in 1. For example, to discard RA messages that match a source MAC address in the list:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name discard]
user@switch# set source-mac-address-list address-list-name
```

Configuring an Accept Policy for RA Guard

You can configure an accept policy to forward RA messages on the basis of certain criteria. You can configure either match lists of source address or address prefixes as the criteria, or you can other match conditions, such as hop limit, configuration flags, or router preference as the criteria.

The following lists can be associated with an accept policy by using the **match-list** option:

- **ipv6-source-address-list**
- **source-mac-address-list**
- **ipv6-prefix-list**



NOTE: You can associate more than one type of match list with an accept policy. If the **match-all** suboption is configured, then a received RA message must match all configured match lists in order to be forwarded; otherwise, it is discarded. If the **match-any** option is configured, then a received RA message must match any one of the configured match lists in order to be forwarded; if it does not match any of the configured lists, then it is discarded.

The following match conditions can be configured using the **match-option** option:

- **hop-limit**—Configure the RA guard policy to verify the minimum or maximum hop count for an incoming RA message.
- **managed-config-flag**—Configure the RA guard policy to verify that the managed address configuration flag of an incoming RA message is set.
- **other-config-flag**—Configure the RA guard policy to verify that the other configuration flag of an incoming RA message is set.
- **router-preference-maximum**—Configure the RA guard policy to verify that the default router preference parameter value of an incoming RA message is lower than or equal to a specified limit.



NOTE: The **match-list** and **match-option** options are used only in accept policies, not in discard policies.

To configure an accept policy for RA guard by using the **match-list** option:

1. Define one or more lists of authorized source addresses or address prefixes that RA guard will use to filter incoming RA messages. Add one address or address prefix per line in the configuration.

- To define a list of IPv6 source addresses:

```
[edit]
user@switch# set policy-options prefix-list address-list-name ipv6-address
```

- To define a list of IPv6 address prefixes:

```
[edit]
user@switch# set policy-options prefix-list prefix-list-name ipv6-prefix
```

- To define a list of MAC source addresses:

```
[edit]
user@switch# set policy-options mac-list address-list-name mac-address
```

2. Specify the policy name:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard policy
policy-name
```

3. Specify the accept action:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name]
user@switch# set accept
```

4. Specify whether RA guard must meet the criteria in all lists or in any of the lists configured in 1:

- To match on all lists:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
user@switch# set match-list match-all
```

- To match on any of the lists:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
user@switch# set match-list match-any
```

5. Associate the accept policy with the list or lists configured in Step 1. For example:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept match-list match-all]
user@switch# set source-mac-address-list address-list-name
```

To configure an accept policy for RA guard using the **match-option** option:

1. Specify the policy name:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard policy
policy-name
```

2. Specify the accept action:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name]
user@switch# set accept
```

3. Specify the match conditions by using the **match-option** option. For example, to specify a match on the maximum number of hops:

```
[edit forwarding-options access-security router-advertisement-guard policy
policy-name accept]
user@switch# set match-option hop-limit maximum value
```

Enabling Stateful RA Guard on an Interface

You can enable stateful RA guard on an interface. You must first configure a policy, which is used to validate incoming RA messages during the learning period. To configure an accept policy, see [“Configuring an Accept Policy for RA Guard” on page 76](#). To configure a discard policy, see [“Configuring a Discard Policy for RA Guard” on page 75](#). After you apply an RA guard policy to an interface, you must enable RA guard on the corresponding VLAN.

To enable stateless RA guard on an interface:

1. Apply a policy to an interface.

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name policy policy-name
```

2. Configure the **stateful** option on the interface:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name stateful
```

3. Enable stateful RA guard on the corresponding VLAN:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name stateful
```

Enabling Stateful RA Guard on a VLAN

You can enable stateful RA guard on a VLAN. You must first configure a policy, which is used to validate incoming RA messages during the learning state. To configure an accept policy, see [“Configuring an Accept Policy for RA Guard” on page 76](#). To configure a discard policy, see [“Configuring a Discard Policy for RA Guard” on page 75](#).

To enable stateful RA guard on a VLAN:

1. Apply a policy to a VLAN.

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name policy policy-name
```

2. Configure the **stateful** option on the VLAN:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard vlan
vlan-name stateful
```

Configuring the Learning State on an Interface

When stateful RA guard is first enabled, the default state is *off*. An interface in the off state operates as if RA guard is not available. To transition an interface to the learning

state, you must request learning on the interface. An interface in the learning state actively acquires information from the RA messages that it receives.

To configure stateful RA guard learning on an interface:

1. Request learning on the interface.

```
[edit]
user@switch# request access-security router-advertisement-guard-learn interface
interface-name
```

2. Configure the learning period in seconds.

```
[edit]
user@switch# request access-security router-advertisement-guard-learn interface
interface-name duration seconds
```

3. Configure the action to take on ingress RA messages received during the learning period. To forward RA messages received during the learning period, configure forwarding on the interface.

- To forward RA messages during the learning period:

```
[edit]
user@switch# request access-security router-advertisement-guard-learn interface
interface-name duration seconds forward
```

- To block RA messages during the learning period:

```
[edit]
user@switch# request access-security router-advertisement-guard-learn interface
interface-name duration seconds block
```

Configuring the Forwarding State on an Interface

An interface in the forwarding state accepts ingress RA messages that can be validated against the configured policy and forwards them to their destination. An interface can dynamically transition to the forwarding state directly from the learning state, or the forwarding state can be statically configured on the interface.

- To configure the forwarding state on an interface:

```
[edit]
user@switch# request access-security router-advertisement-guard-forward interface
interface-name
```

Configuring the Blocking State on an Interface

An interface in the blocking state blocks ingress RA messages. An interface can dynamically transition to the blocking state directly from the learning state, or the blocking state can be statically configured on the interface. An interface that has been statically configured to be in the blocking state will remain in the blocking state until another state is configured on that interface.

- To configure the blocking state on an interface:

```
[edit]
user@switch# request access-security router-advertisement-guard-block interface
interface-name
```

Configuring an Interface as Trusted or Blocked to Bypass Inspection by RA Guard

You can configure an interface as trusted or blocked to bypass inspection of RA messages by RA guard. When an RA message is received on a trusted or blocked interface, it is not subject to validation against the configured policy. A trusted interface forwards all RA messages. A blocked interface discards all RA messages.

- To configure an interface as trusted:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name mark-interface trusted
```

- To configure an interface as blocked:

```
[edit]
user@switch# set forwarding-options access-security router-advertisement-guard interface
interface-name mark-interface block
```

Related Documentation

- [Understanding IPv6 Router Advertisement Guard on page 71](#)
- [Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74](#)

CHAPTER 6

Configuring MAC Limiting, MAC Move Limiting and Persistent MAC Learning to Prevent DHCP Starvation Attacks

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 88](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 91](#)
- [Configuring MAC Limiting \(J-Web Procedure\) on page 92](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 94](#)
- [Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 96](#)
- [Understanding Persistent MAC Learning \(Sticky MAC\) on page 97](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 98](#)

Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches

MAC limiting enhances port security by limiting the number of MAC addresses that can be learned within a VLAN. Limiting the number of MAC addresses protects the switch from flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). Flooding occurs when the number of new MAC addresses that are learned causes the Ethernet switching table to overflow, and previously learned MAC addresses are flushed from the table. The switch then reverts to flooding the previously-learned MAC addresses, which can impact performance and introduce security vulnerabilities.

MAC move limiting provides additional security by controlling the number of MAC address moves that are allowed in a VLAN within one second. A MAC address move occurs when the switch receives a packet with a source MAC address that has already been learned by the switch, but on a different interface. The Ethernet switching table is then updated to reflect the association of the MAC address with the new interface. Because the Ethernet switching table must be updated for each MAC address move, frequent move events can lead to exhaustion of the switch's processing resources. This might occur as the result of a MAC spoofing attack or a loop in the network.

- [MAC Limiting on page 88](#)
- [MAC Move Limiting on page 89](#)
- [Actions for MAC Limiting and MAC Move Limiting on page 89](#)

MAC Limiting

With MAC limiting, you limit the MAC addresses that can be learned on Layer 2 access interfaces by either limiting the number of MAC addresses or by specifying allowed MAC addresses:

- Limiting the number of MAC addresses—You configure the maximum number of MAC addresses that can be dynamically learned (added to the Ethernet switching table) per interface. You can specify that incoming packets with new MAC addresses be ignored, dropped, or logged when the limit is exceeded. You can also specify that the interface be shut down or temporarily disabled.



NOTE: Static MAC addresses do not count toward the limit you specify for dynamic MAC addresses.

- Specifying allowed MAC addresses—You configure the allowed MAC addresses for an interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. An allowed MAC address is bound to a VLAN so that the address is not registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

MAC limiting is configured on Layer 2 interfaces. You can specify the maximum number of dynamic MAC addresses that can be learned on a single interface, all interfaces, or a

specific interface on the basis of its membership within a VLAN (VLAN membership MAC limit).

When you are configuring the maximum MAC limit for an interface, you can choose the action that occurs on incoming packets when the MAC limit is exceeded. You can specify that incoming packets be ignored, dropped, or logged when the limit is exceeded. You can also specify that the interface be shut down or temporarily disabled.

MAC limiting is not enabled by default. For additional information about configuring MAC limit for an interface on a device that supports ELS, see [“Configuring MAC Limiting \(CLI Procedure\)” on page 91](#). For additional information about configuring MAC limit for an interface on a device that does not support Enhanced Layer 2 Software (ELS), see *Configuring MAC Limiting (CLI Procedure)*.

See *Getting Started with Enhanced Layer 2 Software* for additional information on ELS.

MAC Move Limiting

With MAC move limiting, you limit the number of times a MAC address can move to a new interface within one second. When MAC move limiting is configured, MAC address movements are tracked by the switch. The first time a MAC address moves is always considered a good move and will not count toward the configured MAC move limit. Monitoring of MAC address moves comes into effect after the first move, even if the MAC move limit is configured as 1.

You configure MAC move limiting on a per-VLAN basis. Although you enable this feature on VLANs, the MAC move limit applies to the number of movements for each individual MAC address rather than the total number of MAC address moves in the VLAN. For example, if the MAC move limit is set to 1, the switch allows an unlimited number of MAC address movements within the VLAN as long as the same MAC address does not move more than once within a second.

You can configure an action to be taken if the MAC address move limit is exceeded. You can specify that incoming packets be ignored, dropped, or logged when the limit is exceeded. You can also specify that the interface be shut down or temporarily disabled.

MAC move limiting is not enabled by default. For additional information about configuring MAC move limiting on a device that does not support ELS, see *Configuring MAC Move Limiting (CLI Procedure)*. For additional information about configuring MAC move limiting on a device that supports ELS, see [“Configuring MAC Move Limiting \(CLI Procedure\)” on page 94](#).

Actions for MAC Limiting and MAC Move Limiting

You can choose to have one of the following actions performed when the MAC limit or the MAC move limit is exceeded:



NOTE: There is no default action.

- **drop**—Drop the packet, but do not generate an alarm.
- **drop-and-log**—Drop the packet and generate an alarm, an SNMP trap, or system log entry.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—Forward packets with new source MAC addresses, and learn the new source MAC address.
- **shutdown**—Disable the interface in the VLAN and generate an alarm, an SNMP trap, or a system log entry.
- **vlan-member-shutdown**—(EX9200 only) Block an interface on the basis of its membership in a specific VLAN and generate an alarm, an SNMP trap, or a system log entry.

In the event of shutdown, you can configure the switch to automatically restore the disabled interfaces to service after a specified period of time. To configure autorecovery on a device that supports ELS, see [“Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)” on page 96](#). To configure autorecovery on a device that does not support ELS, see [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#).



NOTE: If you do not configure the switch for autorecovery from the disabled condition, you can bring up the disabled interfaces by running one of the following commands:

- (For devices that support ELS)—[clear ethernet-switching recovery-timeout](#)
- (For devices that do not support ELS)—[clear ethernet-switching port-error](#)

Related Documentation

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 91](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 94](#)
- [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
- [Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 96](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\)](#)
- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\)](#)

Configuring MAC Limiting (CLI Procedure)



NOTE: This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Limiting (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.



NOTE: On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

The different ways of setting a MAC limit are described in the following sections:

- [Limiting the Number of MAC Addresses Learned by an Interface on page 91](#)
- [Limiting the Number of MAC Addresses Learned by a VLAN on page 91](#)

Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action
action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

1. Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-size limit packet-action
action
```

2. Set the maximum number of MAC addresses that can be learned by one or all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action
action
```



NOTE: If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the **drop** and **drop-and-log** options are supported.

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

**Related
Documentation**

- [Understanding Bridging and VLANs on EX Series Switches](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 98](#)

Configuring MAC Limiting (J-Web Procedure)

MAC limiting protects against flooding of the Ethernet switching table on an EX Series switch. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- Maximum number of dynamic MAC addresses allowed per interface—If the limit is exceeded, incoming packets with new MAC addresses are dropped.
- Specific “allowed” MAC addresses for the access interface—Any MAC address that is not in the list of configured addresses is not learned.

You configure MAC limiting for each interface, not for each VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface or on all Layer 2 access interfaces. The default action that the switch will take if that maximum number is exceeded is **drop**—drop the packet and generate an alarm, an SNMP trap, or a system log entry.

To enable MAC limiting on one or more interfaces using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the **Interface List**.

3. Click the **Edit** button. If a message appears asking whether you want to enable port security, click **Yes**.
4. To set a dynamic MAC limit:
 1. Type a limit value in the **MAC Limit** box.
 2. Select an action from the **MAC Limit Action** box (optional). The switch takes this action when the MAC limit is exceeded. If you do not select an action, the switch applies the default action, **drop**.
 - Log—Generate a system log entry.
 - Drop—Drop the packets and generate a system log entry. (Default)
 - Shutdown—Shut down the VLAN and generate a system log entry. You can mitigate the effect of this option by configuring the switch for autorecovery from the disabled state and specifying a **disable timeout** value. See *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*. If you have not configured autorecovery from the disabled state, you can bring up the interfaces by running the **clear ethernet-switching port-error** command.
 - None—No action to be taken.
5. To add allowed MAC addresses:
 1. Click **Add**.
 2. Type the allowed MAC address and click **OK**.

Repeat this step to add more allowed MAC addresses.
6. Click **OK** when you have finished setting MAC limits.
7. Click **OK** after the configuration has been successfully delivered.



NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), a message asking whether you want to enable port security appears.

Related Documentation

- *Configuring MAC Limiting (CLI Procedure)*
- *Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks*
- *Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks*
- *Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks*
- *Verifying That MAC Limiting Is Working Correctly*

- [Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces \(CLI Procedure\)](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 88](#)

Configuring MAC Move Limiting (CLI Procedure)



NOTE: This topic uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring MAC Move Limiting \(CLI Procedure\)](#). For ELS details, see [Getting Started with Enhanced Layer 2 Software](#).

When MAC move limiting is configured, the switch tracks MAC address movements on access and trunk interfaces. A MAC address move occurs when the switch receives a packet with a source MAC address that has already been learned by the switch, but on a different interface. If a MAC address changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged or ignored, or the interface is shut down, as specified in the configuration.

MAC move limiting is not configured by default.

You can choose to have one of the following actions performed when the MAC move limit is exceeded:



NOTE: There is no default action.

- **drop**—(EX2300, EX3400 and EX4300) Drop the packet, but do not generate an alarm.
- **drop-and-log**—(EX2300, EX3400 and EX4300 only) Drop the packet and generate an alarm, an SNMP trap, or system log entry.
- **log**—(EX4300 and EX9200) Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—(EX4300 and EX9200) Forward packets with new source MAC addresses, and learn the new source MAC address.
- **shutdown**—Disable the interface in the VLAN and generate an alarm, an SNMP trap, or a system log entry. If you configure an interface with the [recovery-timeout](#) statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you do not configure the switch for autorecovery from the disabled condition, you can bring up the disabled interfaces by running the [clear ethernet-switching recovery-timeout](#) command.
- **vlan-member-shutdown**—(EX9200 only) Block an interface on the basis of its membership in a specific VLAN and generate an alarm, an SNMP trap, or a system log

entry. If you configure an interface with the **recovery-timeout** statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you do not configure **recovery-timeout**, then the interface remains blocked for 180 seconds, after which it is automatically restored. You can recover all of the blocked interfaces by running the **clear ethernet-switching recovery-timeout** command, or recover a specific interface by using the **set ethernet-switching recovery-timeout interface interface-name vlan vlan-name** command.

To configure a MAC move limit for MAC addresses within a specific VLAN:

- To limit the number of MAC address movements that can be made by an individual MAC address within the specified VLAN:

```
[edit edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit
```

- To limit the number of MAC address movements that can be made by an individual MAC address and to specify the action to be taken when the limit is reached:

```
[edit edit vlans vlan-name switch-options]
user@switch# set mac-move-limit limit packet-action action
```

The switch performs the specified action if it tracks that an individual MAC address within the specified VLAN has moved more than the specified number of times within one second.

- (EX9200 switches only) If an action has been configured, you can determine the priority for an interface involved in the MAC move to be selected for the action:

```
[edit edit vlans vlan-name switch-options]
user@switch# set mac-move-limit interface interface-name action-priority value
```

The interface with the lowest value configured for **action-priority** has the highest priority.



NOTE: You can use the action priority to decrease the likelihood of blocking a trusted interface. The trusted interface should have the lowest priority if the configured action is shutdown or vlan-member-shutdown. To assign a low priority, configure a high value for action-priority.

Related Documentation

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches on page 88](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 91](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 98](#)

Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)



NOTE: This example uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switching device is an EX Series switch and runs software that does not support ELS, see *Understanding Storm Control on EX Series Switches*. If your switching device is an EX Series switch and runs software that does support ELS, see *Getting Started with Enhanced Layer 2 Software*.

An Ethernet switching access interface on a switching device might shut down or be disabled as a result of one of the following port-security or storm-control configurations:

- MAC limiting—(Not supported on MX Series routers) The **mac-limit** statement is configured with the **action-shutdown** statement.
- MAC move limiting—(Not supported on MX Series routers) The **mac-move-limit** statement is configured with the **action-shutdown** statement.
- Storm control—The **storm-control** statement is configured with the **action-shutdown** statement.

You can configure the switching device to automatically restore the disabled interfaces to service after a specified period of time. The specified time configured in the **recovery-timeout** statement applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.



NOTE: To enable autorecovery, specify the recovery timeout value for the interfaces to recover automatically. There is no default recovery timeout. If you do not specify a timeout value, you need to use the [clear ethernet-switching recovery-timeout](#) command for EX Series switches and the [clear bridge recovery-timeout](#) command for MX Series routers to clear the errors and restore the interfaces to service.

To specify the recovery timeout period for the interface:

- Set the **recovery-timeout** statement.

For EX Series switches:

```
[edit interfaces interface-name family unit 0 ethernet-switching]  
user@switch# set recovery-timeout seconds
```

For MX Series routers:

```
[edit interfaces interface-name family unit 0 bridge]  
user@switch# set recovery-timeout seconds
```

**Related
Documentation**

- [Configuring MAC Limiting \(CLI Procedure\) on page 91](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 94](#)
- [Configuring or Disabling Storm Control \(CLI Procedure\)](#)

Understanding Persistent MAC Learning (Sticky MAC)

Persistent MAC learning, also known as sticky MAC, is a port security feature that enables an interface to retain dynamically learned MAC addresses when the switch is restarted or if the interface goes down and is brought back online.

Persistent MAC address learning is disabled by default. You can enable persistent MAC address learning in conjunction with MAC limiting to restrict the number of persistent MAC addresses. You enable this feature on interfaces.

Configure persistent MAC learning on an interface to:

- Prevent traffic losses for trusted workstations and servers because the interface does not have to relearn the addresses from ingress traffic after a restart.
- Protect the switch against security attacks. Use persistent MAC learning in combination with MAC limiting to protect against attacks, such as Layer 2 denial-of-service (DoS) attacks, overflow attacks on the Ethernet switching table, and DHCP starvation attacks, by limiting the MAC addresses allowed while still allowing the interface to dynamically learn a specified number of MAC addresses. The interface is secured because after the limit has been reached, additional devices cannot connect to the port.

By configuring persistent MAC learning along with MAC limiting, you enable interfaces to learn MAC addresses of trusted workstations and servers from the time when you connect the interface to your network until the limit for MAC addresses is reached, and ensure that after this limit is reached, new devices will not be allowed to connect to the interface even if the switch restarts. As an alternative to using persistent MAC learning with MAC limiting, you can statically configure each MAC address on each port or allow the port to continuously learn new MAC addresses after restarts or interface-down events. Allowing the port to continuously learn MAC addresses represents a security risk.



NOTE: While a switch is restarting or an interface is coming back up, there might be a short delay before the interface can learn more MAC addresses. This delay occurs while the system re-enters previously learned persistent MAC addresses into the forwarding database for the interface.



TIP: If you move a device within your network that has a persistent MAC address entry on the switch, use the `clear ethernet-switching table persistent-mac` command to clear the persistent MAC address entry from the interface. If you move the device and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address of the device and the device will not be able to connect.

If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect. However, if you do not clear the persistent MAC address on the original port, then when the port restarts, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the persistent MAC address is removed from the new port and the device loses connectivity.

Consider the following configuration guidelines when configuring persistent MAC learning:

- Interfaces must be configured in access mode (use the **port-mode** configuration statement or, for switches operating on the Enhanced Layer 2 Software (ELS) configuration style, the **interface-mode** configuration statement).
- You cannot enable persistent MAC learning on an interface on which 802.1x authentication is configured.
- You cannot enable persistent MAC learning on an interface that is part of a redundant trunk group.
- You cannot enable persistent MAC learning on an interface on which **no-mac-learning** is enabled.

**Related
Documentation**

- [Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\)](#)
- [Configuring Persistent MAC Learning \(CLI Procedure\) on page 98 \(ELS\)](#)

Configuring Persistent MAC Learning (CLI Procedure)



NOTE: This topic uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Persistent MAC Learning (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

Persistent MAC address learning is disabled by default. You can enable it to:

- Help prevent traffic losses for trusted workstations and servers because, if persistent MAC address learning is enabled on an interface, the interface does not have to relearn the addresses from ingress traffic after a restart.
- Protect the switch against security attacks—Use persistent MAC learning in combination with MAC limiting to protect against attacks while still obviating the need to statically configure MAC addresses. When the initial learning of MAC addresses up to the number specified by the MAC limit is done, new addresses are not allowed even after a restart. The port is secured because after the limit has been reached, additional devices cannot connect to the interface.

To configure persistent MAC learning on an interface and limit the number of allowed MAC addresses:

1. Enable persistent MAC learning on an interface:

```
[edit switch-options]
user@switch# set interface interface-name persistent-learning
```

2. Configure the MAC limit on an interface, and specify the action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

Values for *action* are:

drop—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.

drop-and-log—(EX Series switches only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

log—(EX Series switches only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.

none—(EX Series switches only) Forward packets with new source MAC addresses, and learn the new source MAC address.

shutdown—(EX Series switches only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.



TIP: If you move a device within your network that has a persistent MAC address entry on the switch, use the `clear ethernet-switching table persistent-mac` command to clear the persistent MAC address entry from the interface. If you move the device and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address of the device and the device will not be able to connect.

If the original port is down when you move the device, then the new port will learn the MAC address and the device can connect. However, if you do not clear the persistent MAC address on the original port, then when the port restarts, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the persistent MAC address is removed from the new port and the device loses connectivity.

Related Documentation

- [Configuring Persistent MAC Learning \(CLI Procedure\)](#)
- [Configuring MAC Move Limiting \(CLI Procedure\) on page 94](#)
- [Understanding Persistent MAC Learning \(Sticky MAC\) on page 97](#)

CHAPTER 7

Configuring MACSec to Provide Point-to-Point Security on Ethernet Links

- [Understanding Media Access Control Security \(MACsec\) on page 101](#)
- [Configuring Media Access Control Security \(MACsec\) on page 109](#)

Understanding Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for all traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks.

MACsec allows you to secure an Ethernet link for almost all traffic, including frames from the Link Layer Discovery Protocol (LLDP), Link Aggregation Control Protocol (LACP), Dynamic Host Configuration Protocol (DHCP), Address Resolution Protocol (ARP), and other protocols that are not typically secured on an Ethernet link because of limitations with other security solutions. MACsec can be used in combination with other security protocols such as IP Security (IPsec) and Secure Sockets Layer (SSL) to provide end-to-end network security.

MACsec is standardized in IEEE 802.1AE. The IEEE 802.1AE standard can be seen on the IEEE organization website at [IEEE 802.1: BRIDGING & MANAGEMENT](#).

This topic contains the following sections:

- [How MACsec Works on page 102](#)
- [Understanding Connectivity Associations and Secure Channels on page 102](#)
- [Understanding MACsec Security Modes on page 103](#)
- [Understanding the Requirements to Enable MACsec on a Switch-to-Host Link on page 105](#)
- [MACsec Hardware and Software Support Summary on page 105](#)
- [Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches on page 106](#)
- [Understanding MACsec Software Requirements for EX Series and QFX Series Switches on page 107](#)

- [Understanding the MACsec Feature License Requirement on page 108](#)
- [MACsec Limitations on page 108](#)

How MACsec Works

MACsec provides industry-standard security through the use of secured point-to-point Ethernet links. The point-to-point links are secured after matching security keys—a user-configured pre-shared key when you enable MACsec using static connectivity association key (CAK) security mode, a user-configured static secure association key when you enable MACsec using static secure association key (SAK) security mode, or a dynamic key included as part of the AAA handshake with the RADIUS server when you enable MACsec using dynamic security mode—are exchanged and verified between the interfaces at each end of the point-to-point Ethernet link. Other user-configurable parameters, such as MAC address or port, must also match on the interfaces on each side of the link to enable MACsec. See [“Configuring Media Access Control Security \(MACsec\)” on page 109](#).

Once MACsec is enabled on a point-to-point Ethernet link, all traffic traversing the link is MACsec-secured through the use of data integrity checks and, if configured, encryption.

The data integrity checks verify the integrity of the data. MACsec appends an 8-byte header and a 16-byte tail to all Ethernet frames traversing the MACsec-secured point-to-point Ethernet link, and the header and tail are checked by the receiving interface to ensure that the data was not compromised while traversing the link. If the data integrity check detects anything irregular about the traffic, the traffic is dropped.

MACsec can also be used to encrypt all traffic on the Ethernet link. The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anybody monitoring traffic on the link. MACsec encryption is optional and user-configurable; you can enable MACsec to ensure the data integrity checks are performed while still sending unencrypted data “in the clear” over the MACsec-secured link, if desired.

MACsec is configured on point-to-point Ethernet links between MACsec-capable interfaces. If you want to enable MACsec on multiple Ethernet links, you must configure MACsec individually on each point-to-point Ethernet link.

Understanding Connectivity Associations and Secure Channels

MACsec is configured in connectivity associations. MACsec is enabled when a connectivity association is assigned to an interface.

When you are configuring MACsec using static secure association key (SAK) security mode, you must configure secure channels within a connectivity association. The secure channels are responsible for transmitting and receiving data on the MACsec-enabled link, and also responsible for transmitting SAKs across the link to enable and maintain MACsec. A single secure channel is uni-directional—it can only be used to apply MACsec to inbound or outbound traffic. A typical connectivity association when MACsec is enabled using SAK security mode contains two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic.

When you enable MACsec using static CAK or dynamic security mode, you have to create and configure a connectivity association. Two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic—are automatically created. The automatically-created secure channels do not have any user-configurable parameters; all configuration is done in the connectivity association outside of the secure channels.

Understanding MACsec Security Modes

Understanding Static Connectivity Association Key Security Mode (Recommended Security Mode for Switch-to-Switch Links)

When you enable MACsec using static connectivity association key (CAK) security mode, two security keys—a connectivity association key (CAK) that secures control plane traffic and a randomly-generated secure association key (SAK) that secures data plane traffic—are used to secure the point-to-point Ethernet link. Both keys are regularly exchanged between both devices on each end of the point-to-point Ethernet link to ensure link security.

You initially establish a MACsec-secured link using a pre-shared key when you are using static CAK security mode to enable MACsec. A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

Once matching pre-shared keys are successfully exchanged, the MACsec Key Agreement (MKA) protocol is enabled. The MKA protocol is responsible for maintaining MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

We recommend enabling MACsec on switch-to-switch links using static CAK security mode. Static CAK security mode ensures security by frequently refreshing to a new random security key and by only sharing the security key between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode.

See [“Configuring Media Access Control Security \(MACsec\)” on page 109](#) for step-by-step instructions on enabling MACsec using static CAK security mode.

Understanding Dynamic Secure Association Key Security Mode (Switch-to-Host Links)

Dynamic secure association key security mode is used to enable MACsec on a switch-to-host link.

To enable MACsec on a link connecting an endpoint device—such as a server, phone, or personal computer—to a switch, the endpoint device must support MACsec and must be running software that allows it to enable a MACsec-secured connection. When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

A secure association using dynamic secure association security mode must be configured on the switch's Ethernet interface that connects to the host in order for the switch to create a MACsec-secured connection after receiving the MKA keys from the RADIUS server.

The RADIUS server must be using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) in order to support MACsec. The RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used to support MACsec. In order to enable MACsec on a switch to secure a connection to a host, you must be using 802.1X authentication on the RADIUS server. MACsec must be configured into dynamic mode. MACsec is still enabled using connectivity associations when enabled on a switch-to-host link, as it is on a switch-to-switch link.

Understanding Static Secure Association Key Security Mode (Supported for Switch-to-Switch Links)

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured SAKs is used to secure data traffic on the point-to-point Ethernet link. All SAK names and values are configured by the user; there is no key server or other tool that creates SAKs. Security is maintained on the point-to-point Ethernet link by periodically rotating between the two security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure SAKs within secure channels when you enable MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two manually-configured SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

We recommend enabling MACsec using static CAK security mode. You should only use static SAK security mode if you have a compelling reason to use it instead of static CAK security mode.

See [“Configuring Media Access Control Security \(MACsec\)” on page 109](#) for step-by-step instructions on enabling MACsec using SAKs.

Understanding the Requirements to Enable MACsec on a Switch-to-Host Link

When configuring MACsec on a switch-to-host link, the MACsec Key Agreement (MKA) keys, which are included as part of 802.1X authentication, are retrieved from a RADIUS server as part of the AAA handshake. A master key is passed from the RADIUS server to the switch and from the RADIUS server to the host in independent authentication transactions. The master key is then passed between the switch and the host to create a MACsec-secured connection.

The following requirements must be met in order to enable MACsec on a link connecting a host device to a switch.

The host device:

- must support MACsec and must be running software that allows it to enable a MACsec-secured connection with the switch.

The switch:

- must be an EX4200, EX4300, or EX4550 switch running Junos OS Release 14.1X53-D10 or later, or an EX9200 switch running Junos OS Release 15.1R1 or later.
- must be configured into dynamic secure association key security mode.
- must be using 802.1X authentication to communicate with the RADIUS server.

The RADIUS server:

- must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.



NOTE: RADIUS servers that support other widely-used authentication frameworks, such as password-only or md5, cannot be used to support MACsec.

- must be using 802.1X authentication.
- can be multiple hops from the switch and the host device.

MACsec Hardware and Software Support Summary

[Table 9 on page 106](#) summarizes MACsec hardware and software support for EX Series and QFX Series switches.

MACsec hardware and software support is discussed in greater detail in the remaining sections.

Table 9: MACsec Hardware and Software Support Summary for EX Series and QFX Series Switches

| Switch | MACsec-capable Interfaces | Switch-to-Switch Support Introduction | Switch-to-Host Support Introduction | Required Software Package |
|---------|--|---------------------------------------|-------------------------------------|---|
| EX4200 | All uplink port connections on the SFP+ MACsec uplink module. | 13.2X50-D15 | 14.1X53-D10 | controlled |
| EX4300 | All access and uplink ports. | 13.2X50-D15 | 14.1X53-D10 | controlled |
| EX4550 | All EX4550 optical interfaces that use the LC connection type. | 13.2X50-D15 | 14.1X53-D10 | controlled |
| EX4600 | All twenty-four fixed 1GbE SFP/10GbE SFP+ interfaces and all interfaces that support the copper Gigabit Interface Converter (GBIC). All eight SFP+ interfaces on the EX4600-EM-8F expansion module. | 14.1X53-D15 | Not supported | controlled |
| EX9200 | All forty SFP interfaces on the EX9200-40F-M. | 15.1R1 | 15.1R1 | Junos image <i>NOTE:</i> MACsec is available on the Junos OS image in EX9200 switches only. MACsec is not available on the limited Junos OS image package. |
| QFX5100 | All eight SFP+ interfaces on the EX4600-EM-8F expansion module. | 14.1X53-D15 | Not supported | controlled |

Understanding MACsec Hardware Requirements for EX Series and QFX Series Switches

MACsec is currently supported on the following EX Series and QFX Series switch interfaces:

- The uplink port connections on the SFP+ MACsec uplink module that can be installed on EX4200 series switches.
- All access and uplink ports on EX4300 switches.
- All EX4550 optical interfaces that use the LC connection type. See *Pluggable Transceivers Supported on EX4550 Switches*.

- All twenty-four fixed 1GbE SFP/10GbE SFP+ interfaces on an EX4600 switch and all interfaces that support the copper Gigabit Interface Converter (GBIC).
- All eight SFP+ interfaces on the EX4600-EM-8F expansion module, when installed in an EX4600 or QFX5100-24Q switch.



NOTE: MACsec is not supported on EX4600 or QFX5100-24Q switches in Junos OS Release 15.1.

See [Feature Explorer](#) for a full listing of Junos OS releases that support MACsec.

- All forty SFP interfaces on the EX9200-40F-M line card, when the line card is installed in an EX9200 series switch.

MACsec can be configured on supported switch interfaces when those switches are configured in a Virtual Chassis or Virtual Chassis Fabric (VCF), including when MACsec-supported interfaces are on member switches in a mixed Virtual Chassis or VCF that includes switch interfaces that do not support MACsec. MACsec, however, cannot be enabled on Virtual Chassis ports (VCPs) to secure traffic travelling between member switches in a Virtual Chassis or VCF.

Understanding MACsec Software Requirements for EX Series and QFX Series Switches

See [Feature Explorer](#) for a full listing of Junos OS releases and platforms that support MACsec.

MACsec was initially released on EX4200, EX4300, and EX4550 switches in Junos OS Release 13.2X50-D15.

MACsec support for dynamic security mode, which allows MACsec to be configured on switch-to-host links, for EX4200, EX4300, and EX4550 switches was introduced in Junos OS Release 14.1X53-D10.

The switches on each end of a MACsec-secured switch-to-switch link must either both be using Junos OS Release 14.1X53-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

MACsec support for EX4600 switches and QFX5100-24Q switches was introduced in Junos OS Release 14.1X53-D15. The EX4600 and QFX5100-24Q switches supports MACsec on switch-to-switch links only.



NOTE: MACsec is not supported on EX4600 or QFX5100-24Q switches in Junos OS Release 15.1.

See [Feature Explorer](#) for a full listing of Junos OS releases and platforms that support MACsec.

MACsec support for EX9200 switches for both switch-to-switch links and for switch-to-host links was introduced in Junos OS Release 15.1R1.

You must download the controlled version of your Junos OS software to enable MACsec on EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches. MACsec software support is not available in the domestic version of your Junos OS software on these platforms.

You must download the standard Junos image to enable MACsec on EX9200 switches. MACsec is not supported on the limited image.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches includes all features and functionality available in the domestic version of Junos OS, while also supporting MACsec. The domestic version of Junos OS software is shipped on all switches that support MACsec, so you must download and install a controlled version of Junos OS software for your switch before you can enable MACsec.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The standard version of Junos OS software for EX9200 switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of this Junos OS software is strictly controlled under United States export laws. The export, import, and use of this Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring this version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

Understanding the MACsec Feature License Requirement

A feature license is required to configure MACsec on a switch.

To purchase a feature license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

The MACsec feature license is an independent feature license; the feature licenses that must be purchased to enable other groups of features on your switches cannot be purchased to enable MACsec.

MACsec Limitations

All types of Spanning Tree Protocol frames cannot currently be encrypted using MACsec.

- Related Documentation**
- [Configuring Media Access Control Security \(MACsec\) on page 109](#)

Configuring Media Access Control Security (MACsec)

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly-connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks. MACsec is standardized in IEEE 802.1AE.

You can configure MACsec to secure point-to-point Ethernet links connecting EX Series or QFX Series switches, or on Ethernet links connecting a switch to a host device such as a PC, phone, or server. Each point-to-point Ethernet link that you want to secure using MACsec must be configured independently. You can enable MACsec on switch-to-switch links using static secure association key (SAK) security mode or static connectivity association key (CAK) security mode. Both processes are provided in this document.



BEST PRACTICE: We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

The configuration steps for both processes are provided in this document.

- [Acquiring and Downloading the Junos OS Software on page 109](#)
- [Acquiring and Downloading the MACsec Feature License on page 111](#)
- [Configuring the PIC Mode of the MACsec-capable Interfaces \(EX4200 switches only\) on page 111](#)
- [Configuring MACsec Using Static Connectivity Association Key Security Mode \(Recommended for Enabling MACsec on Switch-to-Switch Links\) on page 113](#)
- [Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link on page 117](#)
- [Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link on page 121](#)

Acquiring and Downloading the Junos OS Software

MACsec was initially released on EX4200, EX4300, and EX4550 switches in Junos OS Release 13.2X50-D15. MACsec was released on EX4600 and QFX5100-24Q switches in Junos OS Release 14.1X53-D15, and on EX9200 series switches in Junos OS Release 15.1R1. The switches on each end of a MACsec-secured switch-to-switch link must either both

be using Junos OS Release 14.1X51-D10 or later, or must both be using an earlier version of Junos, in order to establish a MACsec-secured connection when using static CAK security mode.

You must download the controlled version of your Junos OS software to enable MACsec on EX4200, EX4300, EX4550, EX4600, and QFX5100-24Q switches.

You must download the standard version of your Junos OS software to enable MACsec on EX9200 switches. MACsec is not supported in the limited version of Junos OS on EX9200 switches.

See [“Understanding Media Access Control Security \(MACsec\)” on page 101](#) for additional information on the versions of Junos OS software that are required for MACsec.

You can identify whether a software package is the controlled or standard version of Junos OS by viewing the package name. A software package for a controlled version of Junos OS is named using the following format:

package-name-m.nZx.y-controlled-signed.tgz

A software package for a standard version of Junos OS on an EX9200 switch is named using the following format:

package-name-m.nZx.y-.tgz

If you are unsure which version of Junos OS is running on your switch, enter the **show version** command. If the “JUNOS Crypto Software Suite” description appears in the output, you are running the controlled version of Junos OS.

The controlled version of Junos OS software for EX4200, EX4300, EX4550, EX4600, or QFX5100-24Q switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of the controlled version of Junos OS software is strictly controlled under United States export laws. The export, import, and use of the controlled version of Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring the controlled version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The standard version of Junos OS software for EX9200 switches contains encryption and is, therefore, not available to customers in all geographies. The export and re-export of this Junos OS software is strictly controlled under United States export laws. The export, import, and use of this Junos OS software is also subject to controls imposed under the laws of other countries. If you have questions about acquiring this version of your Junos OS software, contact Juniper Networks Trade Compliance group at compliance_helpdesk@juniper.net.

The process for installing the controlled or standard version of Junos OS software onto your switch is identical to installing any other version of Junos OS software. You must enter the **request system software add** statement to download the Junos OS image, and the **request system reboot** statement to reboot the switch to complete the upgrade procedure. See *Downloading Software Packages from Juniper Networks, Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*, and *Installing Software*

on an EX Series Switch with Redundant Routing Engines (CLI Procedure) for detailed information about acquiring and installing Junos OS software images for your switches.

Acquiring and Downloading the MACsec Feature License

A feature license is required to configure MACsec on an EX Series or a QFX Series switch.

The MACsec feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on EX Series or QFX Series switches cannot be purchased to enable MACsec.

To purchase a software license for MACsec, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with a feature license file and a license key. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show chassis hardware** command.

For a Virtual Chassis deployment, two MACsec license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To add one or more new MACsec license keys on the switch, follow this procedure:

1. Add the license key or keys:
 - To add one or more license keys from a file or URL, specify the filename of the file or the URL where the key is located:

```
user@switch> request system license add filename |url
```

- To add a license key from the terminal:

```
user@switch> request system license add terminal
```

2. When prompted, enter the license key, separating multiple license keys with a blank line.

If the license key you enter is invalid, an error appears in the CLI output when you press Ctrl+d to exit the license entry mode.

A MACsec feature license is installed and maintained like any other switch license. See *Managing Licenses for the EX Series Switch (CLI Procedure)* or *Adding New Licenses (CLI Procedure)* for more detailed information on configuring and managing your MACsec software license.

Configuring the PIC Mode of the MACsec-capable Interfaces (EX4200 switches only)

To configure MACsec on an EX4200 switch, you must install the SFP+ MACsec uplink module. The interfaces on the SFP+ MACsec uplink module are the only MACsec-capable interfaces available for EX4200 switches. All four ports on the uplink module are MACsec-capable.

The SFP+ MACsec uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four

ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

The PIC mode is set to **10g**, by default. You only need to perform this procedure if you want to operate your uplink in 1-gigabit mode, or if you previously set the uplink module to 1-gigabit mode and would like to return it to 10-gigabit mode.

To configure the PIC mode:

```
[edit chassis]
```

```
user@switch# set fpc fpc-slot-number pic 1 sfpplus pic-mode (1g | 10g)
```

where *fpc-slot-number* is the FPC slot number, *pic-slot-number* is the PIC slot number, and the **[1g | 10g]** option configures the MACsec capability of the four SFP+ ports on the MACsec uplink module.

The *fpc-slot-number* is always 0 on standalone EX4200 switches, and is the member ID of the member switch in an EX4200 Virtual Chassis.

The PIC slot number is always 1 for the uplink module port slot on an EX4200 switch, so **pic 1** is always the specified PIC slot number.

The PIC mode is set to **10g** by default. When the PIC mode is set to **10g**, uplink ports 0 and 2 on the MACsec uplink module support MACsec at 10-Gbps speeds. Ports 1 and 3 cannot be used to send any traffic.

When the PIC mode is set to **1g**, all four SFP+ ports on the MACsec uplink module support MACsec at 1-Gbps speeds.

Configuring MACsec Using Static Connectivity Association Key Security Mode (Recommended for Enabling MACsec on Switch-to-Switch Links)

You can enable MACsec using static connectivity association key (CAK) security mode or static secure association keys (SAK) security mode on a point-to-point Ethernet link connecting switches. This procedure shows you how to configure MACsec using static CAK security mode.



BEST PRACTICE: We recommend enabling MACsec using static CAK security mode on switch-to-switch links. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available for MACsec-secured switch-to-switch connections that are enabled using static CAK security mode.

When you enable MACsec using static CAK security mode, a pre-shared key is exchanged between the switches on each end of the point-to-point Ethernet link. The pre-shared key includes a connectivity association name (CKN) and a connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

After the pre-shared keys are exchanged and verified, the MACsec Key Agreement (MKA) protocol, which enables and maintains MACsec on the link, is enabled. The MKA is responsible for selecting one of the two switches on the point-to-point link as the key server. The key server then creates a randomized security key that is shared only with the other device over the MACsec-secured link. The randomized security key enables and maintains MACsec on the point-to-point link. The key server will continue to periodically create and share a randomly-created security key over the point-to-point link for as long as MACsec is enabled.

You enable MACsec using static CAK security mode by configuring a connectivity association on both ends of the link. All configuration is done within the connectivity association but outside of the secure channel. Two secure channels—one for inbound traffic and one for outbound traffic—are automatically created when using static CAK security mode. The automatically-created secure channels do not have any user-configurable parameters that cannot already be configured in the connectivity association.

To configure MACsec using static CAK security mode to secure a switch-to-switch Ethernet link:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-cak** for the connectivity association:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name security-mode
static-cak
```

For instance, to configure the MACsec security mode to **static-cak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-cak
```

3. Create the pre-shared key by configuring the connectivity association key name (CKN) and connectivity association key (CAK):

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name pre-shared-key
ckn hexadecimal-number
user@switch# set connectivity-association connectivity-association-name pre-shared-key
cak hexadecimal-number
```

A pre-shared key is exchanged between directly-connected links to establish a MACsec-secure link. The pre-shared-key includes the CKN and the CAK. The CKN is a 64-digit hexadecimal number and the CAK is a 32-digit hexadecimal number. The CKN and the CAK must match on both ends of a link to create a MACsec-secured link.



NOTE: To maximize security, we recommend configuring all 64 digits of a CKN and all 32 digits of a CAK.

If you do not configure all 64 digits of a CKN or all 32 digits of a CAK, however, all remaining digits will be auto-configured to 0.

After the pre-shared keys are successfully exchanged and verified by both ends of the link, the MACsec Key Agreement (MKA) protocol is enabled and manages the secure link. The MKA protocol then elects one of the two directly-connected switches as the key server. The key server then shares a random security with the other device over the MACsec-secure point-to-point link. The key server will continue to periodically create and share a random security key with the other device over the MACsec-secured point-to-point link as long as MACsec is enabled.

To configure a CKN of **37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311** and CAK of **228ef255aa23ff6729ee664acb66e91f** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 pre-shared-key ckn
37c9c2c45ddd012aa5bc8ef284aa23ff6729ee2e4acb66e91fe34ba2cd9fe311
user@switch# set connectivity-association ca1 pre-shared-key cak
228ef255aa23ff6729ee664acb66e91f
```



NOTE: MACsec is not enabled until a connectivity association is attached to an interface. See the final step of this procedure to attach a connectivity association to an interface.

4. (Required on switches when connecting to EX4300 switches only) Enable SCI tagging:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set include-sci
```

You must enable SCI tagging on a switch that is enabling MACsec on an Ethernet link connecting to an EX4300 switch.

SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. This option is, therefore, not available on EX4300 switches.

You should only use this option when enabling MACsec on a link to an EX4300 switch. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16.

If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower *interval* increases bandwidth overhead on the link; a higher *interval* optimizes MKA protocol communication.

The default *interval* is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link when MACsec using static CAK security mode is enabled.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca1* is attached to an interface:

```
[edit security macsec connectivity-association ca1]
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using static CAK security mode, by default.

When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named *ca1*:

```
[edit security macsec connectivity-association ca1]
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an **offset** is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link.

When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance,

if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association **ca1**:

```
[edit security macsec connectivity-association ca1]
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association **ca1** to interface **xe-0/0/1**:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static CAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and contains pre-shared keys that match on both ends of the link.

Configuring MACsec on the Switch Using Dynamic Secure Association Key Security Mode to Secure a Switch-to-Host Link

Before you begin to enable MACsec on a switch-to-host link:

- Confirm that MACsec on switch-to-host links is supported on your switch. See [“Understanding Media Access Control Security \(MACsec\)” on page 101](#).
- Configure a RADIUS server. The RADIUS server:
 - must be configured as the user database for 802.1X authentication.

- must be using the Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication framework.
- must have connectivity to the switch and to the host. The RADIUS server can be multiple hops from the switch or the host.

See Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch.

- Enable MACsec on the host device.

The procedures for enabling MACsec on the host device varies by host device, and is beyond the scope of this document.

To configure MACsec using dynamic security mode to secure a switch-to-host Ethernet link:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca-dynamic1`, enter:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1
```

2. Configure the MACsec security mode as dynamic for the connectivity association:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name security-mode dynamic
```

For instance, to configure the MACsec security mode to dynamic on connectivity association `ca-dynamic1`:

```
[edit security macsec]
```

```
user@switch# set connectivity-association ca-dynamic1 security-mode dynamic
```

3. (Optional) Configure the **must-secure** option:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name mka must-secure
```

When the **must-secure** option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.

When the **must-secure** option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.

The **must-secure** option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the **must-secure** option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.

The **must-secure** option is disabled, by default.

4. (Required only if the host device requires SCI tagging) Enable SCI tagging:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set include-sci
```

You should only use this option when connecting a switch to a host that requires SCI tags. SCI tags are eight octets long, so appending an SCI tag to all traffic on the link adds a significant amount of unneeded overhead.

5. (Optional) Set the MKA key server priority:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka key-server-priority priority-number
```

Specifies the key server priority used by the MKA protocol to select the key server. The switch with the lower *priority-number* is selected as the key server.

The default *priority-number* is 16. If the **key-server-priority** is identical on both sides of the point-to-point link, the MKA protocol selects the interface with the lower MAC address as the key server. Therefore, if this statement is not configured in the connectivity associations at each end of a MACsec-secured point-to-point link, the interface with the lower MAC address becomes the key server.

To change the key server priority to 0 to increase the likelihood that the current device is selected as the key server when MACsec is enabled on the interface using connectivity association *ca1*:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka key-server-priority 0
```

To change the key server priority to 255 to decrease the likelihood that the current device is selected as the key server in connectivity association *ca-dynamic1*:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka key-server-priority 255
```

6. (Optional) Set the MKA transmit interval:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set mka transmit-interval interval
```

The MKA transmit interval setting sets the frequency for how often the MKA protocol data unit (PDU) is sent to the directly connected device to maintain MACsec connectivity on the link. A lower interval increases bandwidth overhead on the link; a higher interval optimizes MKA protocol communication.

The default interval is 2000ms. We recommend increasing the interval to 6000 ms in high-traffic load environments. The transmit interval settings must be identical on both ends of the link.

For instance, if you wanted to increase the MKA transmit interval to 6000 milliseconds when connectivity association *ca-dynamic1* is attached to an interface:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set mka transmit-interval 6000
```

7. (Optional) Disable MACsec encryption:

```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set no-encryption
```

Encryption is enabled for all traffic entering or leaving the interface when MACsec is enabled using dynamic security mode, by default. When encryption is disabled, traffic is forwarded across the Ethernet link in clear text. You are able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec

header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic sent or received on the link has not been tampered with and does not represent a security threat.

8. (Optional) Set an offset for all packets traversing the link:

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set offset (0 | 30 | 50)
```

For instance, if you wanted to set the offset to 30 in the connectivity association named ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set offset 30
```

The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an offset is not set.

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

9. (Optional) Enable replay protection.

```
[edit security macsec connectivity-association connectivity-association-name]  
user@switch# set replay-protect replay-window-size number-of-packets
```

When MACsec is enabled on a link, an ID number is assigned to each packet on the MACsec-secured link. When replay protection is enabled, the receiving interface checks the ID number of all packets that have traversed the MACsec-secured link. If a packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.

Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.

Replay protection should not be enabled in cases where packets are expected to arrive out of order.

You can require that all packets arrive in order by setting the replay window size to 0.

To enable replay protection with a window size of five on connectivity association ca-dynamic1:

```
[edit security macsec connectivity-association ca-dynamic1]  
user@switch# set replay-protect replay-window-size 5
```

10. (Optional) Exclude a protocol from MACsec:


```
[edit security macsec connectivity-association connectivity-association-name]
user@switch# set exclude-protocol protocol-name
```

For instance, if you did not want Link Level Discovery Protocol (LLDP) to be secured using MACsec:

```
[edit security macsec connectivity-association ca-dynamic1]
user@switch# set exclude-protocol lldp
```

When this option is enabled, MACsec is disabled for all packets of the specified protocol—in this case, LLDP—that are sent or received on the link.

11. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface. For instance, to assign connectivity association `ca-dynamic1` to interface `xe-0/0/1`:

```
[edit security macsec]
user@switch# set interfaces xe-0/1/0 connectivity-association ca-dynamic1
```

Configuring MACsec Using Static Secure Association Key Security Mode to Secure a Switch-to-Switch Link

When you enable MACsec using static secure association key (SAK) security mode, one of up to two manually configured security keys is used to secure the point-to-point Ethernet link between the switches. All security key names and values are configured by the user; there is no key server or other tool that creates security keys. Security is maintained on the point-to-point Ethernet link by periodically rotating the security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

You configure static SAKs within secure channels when you are enabling MACsec using static SAK security mode. You configure secure channels within connectivity associations. A typical connectivity association for MACsec using static SAK security mode contains two secure channels—one for inbound traffic and one for outbound traffic—that have each been configured with two static SAKs. You must attach the connectivity association with the secure channel configurations to an interface to enable MACsec using static SAK security mode.

To configure MACsec on a switch-to-switch Ethernet link using static SAK security mode:

1. Create a connectivity association. You can skip this step if you are configuring an existing connectivity association.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name
```

For instance, to create a connectivity association named `ca1`, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1
```

2. Configure the MACsec security mode as **static-sak** for the connectivity association:

```
[edit security macsec]
```

```
user@switch# set connectivity-association connectivity-association-name security-mode static-sak
```

For instance, to configure the MACsec security mode to **static-sak** on connectivity association **ca1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 security-mode static-sak
```

3. Create a secure channel within the connectivity association. You can skip this step if you are configuring an existing secure channel.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name
```

For instance, to create secure channel **sc1** in connectivity association **ca1**, enter:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1
```

4. Define the security associations and the static SAKs for the secure channel:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name security-association number key key-string
```

where the **security-association number** is a number between 0 and 3, and the **key-string** is a 32-digit key defined statically by the network administrator.

The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec.

A secure channel must have at least two security associations with unique key strings. MACsec uses a security associations to establish a secure communications link, and periodically rotates to a new security association to keep the link secure. MACsec, therefore, must have at least one backup security association and key at all times.

To create one secure channel with two security associations and keys, for example:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 0 key d183c4002fa6fe3d2d9a852c20ab8412
user@switch# set connectivity-association ca1 secure-channel sc1 security-association 1 key b976c7494ab6fe2f2d4c432a90fd90a8
```

5. Specify whether the secure channel should be applied to traffic entering or leaving the switch:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel secure-channel-name direction [inbound | outbound]
```

where **inbound** applies the secure channel to traffic entering the switch, and **outbound** applies the secure channel to traffic leaving the switch.



NOTE: A secure channel can only be applied to traffic entering (inbound) or leaving (outbound) an interface on the switch.

If you need to configure MACsec using SAKs on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is assigned to an interface later in this process.

For instance, to configure secure channel **sc1** to apply MACsec to incoming traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 direction inbound
```

To configure secure channel **sc2** to apply MACsec to outgoing traffic:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc2 direction outbound
```

6. Specify a MAC address:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id mac-address mac-address
```

If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the **mac-address**.

If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the **mac-address**.

The **mac-address** variables must match on the sending and receiving secure channel on each side of a link to enable MACsec using static SAK security mode.



NOTE: You can see the MAC address of an interface in the **show interfaces** output.

To configure MACsec to accept frames from MAC address **12:34:56:ab:cd:ef** on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id mac-address
12:34:56:ab:cd:ef
```

7. Specify a port:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name id port-id port-id-number
```

The **port-id-number** variables must match on a sending and receiving secure channel on each side of a link to enable MACsec.



NOTE: The only requirement for port numbers in this implementation of MACsec is that they match on the sending and receiving ends of an Ethernet link. When the port numbers match, MACsec is enabled for all traffic on the connection.

To specify port ID 4 on secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 id port-id 4
```

8. (Optional) Enable encryption:

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name encryption
```

You can enable MACsec without enabling encryption. If a secure channel is configured on an interface without encryption, traffic is forwarded across the Ethernet link in clear text, and you will be able to view unencrypted data in the Ethernet frame traversing the link when you are monitoring it. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic on the link does not represent a security threat.

Encryption is disabled by default when you are enabling MACsec using static SAK security mode. To ensure all traffic traversing secure-channel **sc1** is encrypted:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 encryption
```

9. (Optional) Set an offset to send the first 30 or 50 octets in unencrypted plain text when encryption is enabled.

```
[edit security macsec]
user@switch# set connectivity-association connectivity-association-name secure-channel
secure-channel-name offset [0 | 30 | 50]
```

When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic. When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.

You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.

The default offset is 0, so all traffic on the link is encrypted when the **encryption** option is enabled and an **offset** is not set.

To change the offset to 30 for secure channel **sc1**:

```
[edit security macsec]
user@switch# set connectivity-association ca1 secure-channel sc1 offset 30
```

10. Assign the connectivity association to an interface:

```
[edit security macsec]
user@switch# set interfaces interface-names connectivity-association
connectivity-association-name
```

Assigning the connectivity association to an interface is the final configuration step to enabling MACsec on an interface.

For instance, to assign connectivity association ca1 to interface xe-0/0/1:

```
[edit security macsec]  
user@switch# set interfaces xe-0/1/0 connectivity-association ca1
```

MACsec using static SAK security mode is not enabled until a connectivity association on the opposite end of the link is also configured, and the configuration match on both ends of the link.

**Related
Documentation**

- [Understanding Media Access Control Security \(MACsec\) on page 101](#)

CHAPTER 8

Enabling Trusted DHCP Servers to Protect Against Rogue DHCP Servers

- [Understanding Trusted DHCP Servers for Port Security on page 127](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 128](#)
- [Enabling a Trusted DHCP Server \(J-Web Procedure\) on page 128](#)

Understanding Trusted DHCP Servers for Port Security

Any interface on the switching device that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

Related Documentation

- *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 128](#)

Enabling a Trusted DHCP Server (CLI Procedure)



NOTE: This example uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Enabling a Trusted DHCP Server (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can configure any interface on a switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

By default, all access interfaces are untrusted, and all trunk interfaces are trusted. However, you can override the default setting for access interfaces by configuring a group of access interfaces within a VLAN, specifying an interface to belong to that group, and then configuring the group as trusted.

Before you can configure a trusted DHCP server, you must configure a VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To configure an untrusted access interface as a trusted interface for a DHCP server by using the CLI :

1. Configure a group within a VLAN with a specific access interface:

```
[edit vlans vlan-name forwarding-options dhcp-security ]
user@switch# set group group-name interface interface-name
```

2. Configure that group as **trusted** to make the specified interface contained within the group a trusted interface:

```
[edit vlans vlan-name forwarding-options dhcp-security group group-name]
user@switch# set overrides trusted
```

Related Documentation

- [Enabling a Trusted DHCP Server \(J-Web Procedure\) on page 128](#)
- [Understanding Trusted DHCP Servers for Port Security on page 127](#)

Enabling a Trusted DHCP Server (J-Web Procedure)

You can configure any interface on the EX Series switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

You configure a trusted DHCP server on an interface, not on a VLAN. By default, all access interfaces are untrusted and all trunk interfaces are trusted.

To enable a trusted DHCP server on one or more interfaces by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the Port list.

3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Trust DHCP** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.



NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

**Related
Documentation**

- *Enabling a Trusted DHCP Server (CLI Procedure)*
- *Example: Configuring Basic Port Security Features*
- *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
- *Verifying That a Trusted DHCP Server Is Working Correctly*
- *Monitoring Port Security*
- [Understanding Trusted DHCP Servers for Port Security on page 127](#)

CHAPTER 9

Configuration Statements

- [\[edit vlans\] Configuration Statement Hierarchy on EX Series Switches on page 133](#)
- [accept on page 137](#)
- [access-security on page 138](#)
- [arp-inspection on page 139](#)
- [cak on page 140](#)
- [circuit-id on page 141](#)
- [ckn on page 142](#)
- [connectivity-association on page 143](#)
- [connectivity-association \(MACsec Interfaces\) on page 144](#)
- [dhcp-security on page 145](#)
- [dhcp-service on page 148](#)
- [dhcp-snooping-file on page 149](#)
- [dhcpv6-options on page 150](#)
- [dhcpv6-snooping-file on page 151](#)
- [direction on page 152](#)
- [discard on page 153](#)
- [encryption \(MACsec\) on page 154](#)
- [exclude-protocol on page 155](#)
- [group \(DHCP Security\) on page 156](#)
- [host-name on page 157](#)
- [id on page 158](#)
- [interface \(DHCP Security\) on page 159](#)
- [interface \(RA Guard\) on page 160](#)
- [interface-mac-limit on page 161](#)
- [interfaces \(MACsec\) on page 163](#)
- [ip-source-guard on page 164](#)
- [ipv6-prefix-list on page 166](#)
- [ipv6-source-address-list on page 167](#)

- [ipv6-source-guard](#) on page 168
- [key \(MACsec\)](#) on page 169
- [key-server-priority \(MACsec\)](#) on page 170
- [mac](#) on page 171
- [mac-address \(MACsec\)](#) on page 172
- [mac-list](#) on page 173
- [mac-move-limit](#) on page 174
- [macsec](#) on page 176
- [mark-interface \(RA Guard\)](#) on page 177
- [match-list](#) on page 178
- [match-option](#) on page 180
- [mka](#) on page 181
- [must-secure](#) on page 182
- [neighbor-discovery-inspection](#) on page 183
- [no-dhcpv6-options](#) on page 184
- [no-dhcpv6-snooping](#) on page 184
- [no-encryption \(MACsec\)](#) on page 185
- [no-option16](#) on page 186
- [no-option18](#) on page 186
- [no-option37](#) on page 187
- [offset](#) on page 188
- [option-16 \(DHCPv6 Snooping\)](#) on page 189
- [option-18 \(DHCPv6 Snooping\)](#) on page 190
- [option-37 \(DHCPv6 Snooping\)](#) on page 191
- [overrides \(DHCP Security\)](#) on page 192
- [packet-action](#) on page 193
- [persistent-learning](#) on page 195
- [policy](#) on page 196
- [port-id](#) on page 197
- [pre-shared-key](#) on page 198
- [prefix \(DHCPv6 Options\)](#) on page 199
- [prefix-list](#) on page 200
- [recovery-timeout](#) on page 201
- [remote-id](#) on page 203
- [replay-protect](#) on page 204
- [replay-window-size](#) on page 205
- [router-advertisement-guard](#) on page 206

- [routing-instance-name \(circuit-id\) on page 208](#)
- [secure-channel on page 209](#)
- [security-association on page 210](#)
- [security-mode on page 211](#)
- [source-mac-address-list on page 212](#)
- [stateful on page 213](#)
- [stateless on page 214](#)
- [static-ip on page 215](#)
- [static-ipv6 on page 216](#)
- [traceoptions \(DHCP\) on page 217](#)
- [transmit-interval \(MACsec\) on page 219](#)
- [trusted on page 220](#)
- [untrusted on page 220](#)
- [use-interface-description on page 221](#)
- [use-interface-index on page 222](#)
- [use-interface-name on page 223](#)
- [use-string on page 224](#)
- [use-vlan-id on page 226](#)
- [vendor-id on page 228](#)
- [vlan \(RA Guard\) on page 229](#)
- [write-interval on page 230](#)

[edit vlans] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level on page 134](#)
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level on page 136](#)

Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the [edit vlans] configuration statements supported on one or more of the EX Series switches:

```
vlans {  
  vlan-name {  
    description text-description;  
    domain-type bridge;  
    forwarding-options {  
      dhcp-security {  
        arp-inspection;  
        group group-name {  
          interface interface-name {  
            static-ip ip-address {  
              mac mac-address;  
            }  
          }  
        }  
        overrides {  
          no-option82;  
          trusted;  
        }  
      }  
    }  
    ip-source-guard;  
    no-dhcp-snooping;  
    option-82 {  
      circuit-id {  
        prefix {  
          host-name;  
          logical-system-name;  
          routing-instance-name;  
        }  
        use-interface-description (device | logical);  
        use-vlan-id;  
      }  
      remote-id {  
        host-name;  
        use-interface-description (device | logical);  
        use-string string;  
      }  
      vendor-id {  
        use-string string;  
      }  
    }  
  }  
  filter {  
    input filter-name;  
    output filter-name;  
  }  
  flood {  
    input filter-name;  
  }  
}  
l3-interface irb.logical-unit-number;  
multicast-snooping-options {
```

```

flood-groups [group-names];
forwarding-cache {
    threshold {
        reuse threshold;
        suppress threshold;
    }
}
graceful-restart {
    disable;
    restart-duration duration;
}
host-outbound-traffic {
    dot1p bits;
    forwarding-class forwarding-class;
}
multichassis-lag-replicate-state;
nexthop-hold-time time;
options {
    syslog {
        level level;
        mark interval;
        upto level;
    }
}
traceoptions {
    file filename {
        files number;
        no-world-readable;
        size file-size;
        world-readable;
    }
    flag flag {
        disable;
    }
}
}
switch-options {
    interface interface-name {
        interface-mac-limit limit {
            packet-action action;
        }
        static-mac mac-address;
    }
    interface-mac-limit limit {
        packet-action action;
    }
    mac-move-limit limit {
        packet-action action;
    }
    mac-table-size limit {
        packet-action drop;
    }
    no-mac-learning;
}
vlan-id number;
vlan-id-list [vlan-id | vlan-id-vlan-id];

```

```
}  
}
```

Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the [edit vlans] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 10: Unsupported [edit vlans] Configuration Statements on EX Series Switches

| Statement | Hierarchy Level |
|--|-----------------|
| NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies. | |
| mcae-mac-synchronize | [edit vlans] |
| no-irb-layer-2-copy | [edit vlans] |

Related Documentation

- *Example: Connecting Access Switches to a Distribution Switch*

accept

| | |
|---------------------------------|---|
| Syntax | <pre> accept { match-list { (match-all match-any); ipv6-prefix-list <i>prefix-list-name</i>; ipv6-source-address-list <i>address-list-name</i>; source-mac-address-list <i>address-list-name</i>; } match-option { hop-limit { (maximum minimum) <i>value</i>; } managed-config-flag; other-config-flag; router-preference (high low medium); } } </pre> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i>] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure the accept policy for an IPv6 Router Advertisement (RA) guard policy. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>When RA guard is enabled, the switch compares the information contained in the attributes of RA messages to the information contained in the policy. If RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped.</p> <p>The criteria are configured either as one or more lists of source address or address prefixes, which are associated with the accept policy by using the match-list statement, or match condition parameters, which are associated with the accept policy by using the match-option statement.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

access-security

```
Syntax  access-security {
        router-advertisement-guard {
            interface interface-name {
                mark-interface (trusted | block);
                policy policy-name;
                (stateful | stateless);
            }
            vlan vlan-name {
                policy policy-name;
                (stateful | stateless);
            }
            policy policy-name {
                accept {
                    match-list {
                        (match-all | match-any);
                        ipv6-prefix-list prefix-list-name;
                        ipv6-source-address-list address-list-name;
                        source-mac-address-list address-list-name;
                    }
                    match-option {
                        hop-limit {
                            (maximum | minimum) value;
                        }
                        managed-config-flag;
                        other-config-flag;
                        router-preference (high | low | medium);
                    }
                }
                discard {
                    ipv6-prefix-list prefix-list-name;
                    ipv6-source-address-list address-list-name;
                    source-mac-address-list address-list-name;
                }
            }
        }
    }
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.
Statement introduced in Junos OS Release 16.1 for EX Series switches.

Description Configure IPv6 access security options.



The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74](#)
- [Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79](#)

arp-inspection


| | |
|---|---|
| Syntax | <pre>arp-inspection { forwarding-class <i>class-name</i>; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: <ul style="list-style-type: none"> [edit vlans <i>vlan-name</i> forwarding-options dhcp-security], [edit forwarding-options dhcp-relay] For platforms without ELS: <ul style="list-style-type: none"> [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)], [edit forwarding-options dhcp-relay] |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> |
| Description | <p>Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.</p> <p>When DAI is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses. ARP probe packets, which have the sender IP address 0.0.0.0, are validated by DAI.</p> |
| <div>  <p>NOTE: If you configure DAI at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level:</p> <ul style="list-style-type: none"> DAI can be configured only for a specific VLAN, not for a list or a range of VLAN IDs. DHCP snooping is automatically enabled on the specified VLAN. The forwarding-class statement is not available at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level. <p>See “Enabling Dynamic ARP Inspection (CLI Procedure)” on page 48 for more information about this configuration.</p> </div> | |
| <div>  <p>NOTE: On EX9200 switches, DAI is not supported in an MC-LAG scenario.</p> <p>The remaining statement is explained separately.</p> </div> | |
| Default | Disabled. |

| | |
|---------------------------------|--|
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch</i>• <i>Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks</i>• Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33• <i>Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic</i>• <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i>• Enabling Dynamic ARP Inspection (J-Web Procedure) on page 49 |

cak

| | |
|---------------------------------|---|
| Syntax | <code>ckn hexadecimal-number;</code> |
| Hierarchy Level | [edit security macsec connectivity-association connectivity-association-name pre-shared-key] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the connectivity association key (CAK) for a pre-shared key.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a CAK. A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p> |
| Default | No CAK exists, by default. |
| Options | <p>hexadecimal-number—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p> |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

circuit-id

| | |
|---|---|
| Syntax | <pre> circuit-id { prefix { host-name; logical-system-name; routing-instance-name; } use-interface-description (device logical); use-vlan-id; } </pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with Enhanced Layer 2 Software (ELS): [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82], [edit forwarding-options helpers bootp dhcp-option82] , [edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82] For MX Series platforms: [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82] |
| Release Information | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure the circuit-id suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.</p> <p>The remaining statements are explained separately.</p> |
| Default | <p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i>.</p> |
| <div>  <p>NOTE: When you configure circuit-id, remote-id is also enabled, even if you do not explicitly configure remote-id .</p> </div> | |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |

- Related Documentation**
- *Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)*
 - *Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server*
 - *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
 - *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
 - *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
 - [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\)](#) on page 63
 - RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

ckn

| | |
|---------------------------------|---|
| Syntax | <code>ckn hexadecimal-number;</code> |
| Hierarchy Level | [edit security macsec connectivity-association connectivity-association-name pre-shared-key] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the connectivity association key name (CKN) for a pre-shared key.</p> <p>A pre-shared key includes a CKN and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using dynamic security keys. The MACsec Key Agreement (MKA) protocol is enabled once the pre-shared keys are successfully exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link</p> |
| Default | No CKN exists, by default. |
| Options | <p>hexadecimal-number—The key name, in hexadecimal format.</p> <p>The key name is 32 hexadecimal characters in length. If you enter a key name that is less than 32 characters long, the remaining characters are set to 0.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

connectivity-association

| | |
|---------------------------------|---|
| Syntax | <pre> connectivity-association <i>connectivity-association-name</i> { <i>exclude-protocol</i> <i>protocol-name</i>; include-sci; mka { must-secure; key-server-priority <i>priority-number</i>; transmit-interval <i>interval</i>; } no-encryption; offset (0 30 50); pre-shared-key { cak <i>hexadecimal-number</i>; ckn <i>hexadecimal-number</i>; } replay-protect { replay-window-size <i>number-of-packets</i>; } secure-channel <i>secure-channel-name</i> { direction (inbound outbound); encryption (MACsec); id { mac-address <i>mac-address</i>; port-id <i>port-id-number</i>; } offset (0 30 50); security-association <i>security-association-number</i> { key <i>key-string</i>; } } security-mode <i>security-mode</i>; } </pre> |
| Hierarchy Level | [edit security macsec] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Create or configure a MACsec connectivity association.</p> <p>A connectivity association is not applying MACsec to traffic until it is associated with an interface. MACsec connectivity associations are associated with interfaces using the interfaces statement in the [edit security macsec] hierarchy.</p> |
| Default | No connectivity associations are present, by default. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |

Related Documentation • [Configuring Media Access Control Security \(MACsec\) on page 109](#)

connectivity-association (MACsec Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>connectivity-association <i>connectivity-association-name</i>;</code> |
| Hierarchy Level | [edit security macsec interfaces <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Applies a connectivity association to an interface, which enables Media Access Control Security (MACsec) on that interface. |
| Default | No connectivity associations are associated with any interfaces. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | • Configuring Media Access Control Security (MACsec) on page 109 |

dhcp-security

```
Syntax  dhcp-security {
        arp-inspection;
        dhcpv6-options {
            option-16 {
                use-string string;
            }
            option-18 {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                    vlan-id;
                    vlan-name;
                }
                use-interface-mac;
                use-interface-index (device | logical);
                use-interface-description (device | logical);
                use-interface-name (device | logical);
                use-string string;
            }
            option-37 {
                prefix {
                    host-name;
                    logical-system-name;
                    routing-instance-name;
                    vlan-id;
                    vlan-name;
                }
                use-interface-mac;
                use-interface-index (device | logical);
                use-interface-description (device | logical);
                use-interface-name (device | logical);
                use-string string;
            }
        }
        group group-name {
            interface interface-name {
                static-ip ip-address {
                    mac mac-address;
                }
                static-ipv6 ip-address {
                    mac mac-address;
                }
            }
        }
        overrides {
            no-dhcpv6-options;
            no-option16;
            no-option18;
            no-option37;
            no-option82;
            trusted;
            untrusted;
        }
    }
```

```
    }  
  }  
  ip-source-guard;  
  ipv6-source-guard;  
  neighbor-discovery-inspection;  
  no-dhcp-snooping;  
  no-dhcpv6-snooping;  
  option-82 {  
    circuit-id {  
      prefix {  
        host-name;  
        logical-system-name;  
        routing-instance-name;  
      }  
      use-interface-description (device | logical);  
      use-vlan-id;  
    }  
    remote-id {  
      host-name hostname;  
      use-interface-description (device | logical);  
      mac (Option 82);  
      use-string string;  
    }  
    vendor-id {  
      use-string string;  
    }  
  }  
}
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Support for **static-ipv6**, **neighbor-discovery-inspection**, **ipv6-source-guard**, **no-dhcpv6-snooping**, and **no-option37** introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support for **dhcpv6-options**, **option-16**, **option-18**, **option-37**, **no-dhcpv6-options**, **no-option16**, **no-option18**, and **no-option37** introduced in Junos OS Release 14.2 for EX Series switches.

Description Configure port security features on the switch. DHCP snooping is enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard
- DHCP option 82
- Static IP

For switches that support DHCPv6, both DHCP snooping and DHCPv6 snooping are enabled automatically if you configure any of the afore-mentioned features or any of the following IPv6 features:

- IPv6 neighbor discovery inspection
- IPv6 source guard
- Static IPv6



NOTE: On EX9200 switches, DHCP Snooping, DHCPv6 Snooping and Port Security features are not supported in MC-LAG scenario.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
- [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)
- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)

dhcp-service

| | |
|---------------------------------|---|
| Syntax | <pre>dhcp-service { dhcp-snooping-file (local_pathname remote_URL); write-interval interval; }</pre> |
| Hierarchy Level | [edit system processes] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for MX Series routers. |
| Description | <p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure) on page 66 |

dhcp-snooping-file

| | |
|---------------------------------|--|
| Syntax | <code>dhcp-snooping-file (<i>local_pathname</i> <i>remote_URL</i>); <i>write-interval</i> <i>seconds</i>; }</code> |
| Hierarchy Level | [edit system processes dhcp-service] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 14.1 for MX Series routers. |
| Description | <p>Ensure that IP-MAC bindings persist through the device reboots by specifying a local pathname or a remote URL for the storage location of the DHCP snooping database file. You <i>must</i> specify how frequently the device writes the database entries into the DHCP snooping database file.</p> <p>The remaining statement is explained separately.</p> |
| Default | The IP-MAC bindings in the DHCP snooping database file are not persistent by default. If the device is rebooted, the bindings are lost, and the table must be rebuilt on reboot. |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure) on page 66 • Configuring Static IP Addresses for DHCP Bindings on Access Ports for MX Series Routers (CLI Procedure) • Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

dhcpv6-options

```
Syntax  dhcpv6-options {
        option-16 {
            use-string string;
        }
        option-18 {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
        option-37 {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
                vlan-id;
                vlan-name;
            }
            use-interface-mac;
            use-interface-index (device | logical);
            use-interface-description (device | logical);
            use-interface-name (device | logical);
            use-string string;
        }
    }
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options [dhcp-security](#)]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Configure optional information to be included in DHCPv6 packets during the DHCPv6 snooping process.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [no-dhcpv6-options on page 184](#)
- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)

- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\)](#) on page 65

dhcpv6-snooping-file

| | |
|---------------------------------|---|
| Syntax | <pre>dhcpv6-snooping-file (<i>local_pathname</i> <i>remote_URL</i>); location <i>local_pathname</i> <i>remote_URL</i>; timeout <i>seconds</i>; write-interval <i>seconds</i>; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> • For platforms with Enhanced Layer 2 Software (ELS): [edit system processes dhcp-service]; • For platforms without ELS: [edit ethernet-switching-options secure-access-port] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support at the [edit ethernet-switching-options secure-access-port] hierarchy level introduced in Junos OS 14.1X53-D10 for EX Series switches. |
| Description | <p>Ensure that IP-MAC bindings persist through switch reboots by specifying a local pathname or a remote URL for the storage location of the DHCPv6 snooping database file.</p> <p>The remaining statements are explained separately.</p> |
| Default | The IP-MAC bindings in the DHCPv6 snooping database are not persistent. If the switch is rebooted, the bindings are lost. |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure) on page 66 • Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

direction

| | |
|---------------------------------|--|
| Syntax | direction (inbound outbound); |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Configure whether the secure channel applies MACsec security to traffic entering or leaving an interface.</p> <p>If you need to apply MACsec on traffic entering and leaving an interface, you need to create one secure channel to apply MACsec on incoming traffic and another secure channel to apply MACsec on outgoing traffic within the same connectivity association. When you associate the connectivity association with an interface, MACsec is applied on traffic entering and leaving that interface.</p> <p>You only use this configuration option when you are configuring MACsec using static secure association keys (SAK) security mode. When you are configuring MACsec using static connectivity association keys (CAK) security mode, two secure channels that are not user-configurable—one inbound secure channel and one outbound secure channel—are automatically created within the connectivity association.</p> |
| Default | <p>This statement does not have a default value.</p> <p>If you have configured a secure channel to enable MACsec using static SAK security mode, you must specify whether the secure channel applies MACsec to traffic entering or leaving an interface. A candidate configuration that contains a secure channel that has not configured a direction cannot be committed.</p> |
| Options | <p>inbound—Enable MACsec security on traffic entering the interface that has applied the secure channel.</p> <p>outbound—Enable MACsec security on traffic leaving the interface that has applied the secure channel.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

discard

| | |
|---------------------------------|--|
| Syntax | <pre>discard { ipv6-prefix-list <i>prefix-list-name</i>; ipv6-source-address-list <i>address-list-name</i>; source-mac-address-list <i>address-list-name</i>; }</pre> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i>] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure a discard policy for an IPv6 Router Advertisement (RA) guard policy. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>When RA guard is enabled, the switch compares the information contained in the attributes of RA messages to the criteria configured in the policy. When RA guard is enabled by using a discard policy, any RA messages that match the conditions defined in the policy are dropped, and RA messages that do not match the conditions are forwarded.</p> <p>The criteria are configured as one or more lists of IPv6 addresses, MAC addresses, or IPv6 address prefixes associated with the policy. RA guard compares the source address or address prefix of incoming RA messages with the configured lists. You configure the lists at the [edit policy-options] hierarchy level, by using the prefix-list option for an IPv6 address or address prefix list, and the mac-list option for a MAC address list.</p> <p>If more than one list is associated with a discard policy, then an incoming RA message that meets the criteria in any of the lists is discarded.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

encryption (MACsec)

| | |
|---------------------------------|--|
| Syntax | encryption; |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Enable MACsec encryption within a secure channel.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association with a secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the frame, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic has not been tampered with and does not represent a security threat.</p> <p>Traffic traversing a MAC-enabled point-to-point Ethernet link traverses the link at the same speed regardless of whether encryption is enabled or disabled. You cannot increase the speed of traffic traversing a MACsec-enabled Ethernet link by disabling encryption.</p> <p>This command is used to enable encryption when MACsec is configured using secure association key (SAK) security mode only. When MACsec is configuring using static connectivity association key (CAK) security mode, the encryption setting is configured outside of the secure channel using the no-encryption configuration statement.</p> |
| Default | MACsec encryption is disabled when MACsec is configured using static SAK security mode, by default. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

exclude-protocol

| | |
|---------------------------------|---|
| Syntax | <code>exclude-protocol <i>protocol-name</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies protocols whose packets are not secured using Media Access Control Security (MACsec) when MACsec is enabled on a link using static connectivity association key (CAK) security mode.</p> <p>When this option is enabled in a connectivity association that is attached to an interface, MACsec is not enabled for all packets of the specified protocols that are sent and received on the link.</p> |
| Default | <p>Disabled.</p> <p>All packets are secured on a link when MACsec is enabled, with the exception of all types of Spanning Tree Protocol (STP) packets.</p> |
| Options | <p><i>protocol-name</i>—Specifies the name of the protocol that should not be MACsec-secured. Options include:</p> <ul style="list-style-type: none"> • cdp—Cisco Discovery Protocol. • lACP—Link Aggregation Control Protocol. • lldp—Link Level Discovery Protocol. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

group (DHCP Security)

```
Syntax  group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
            static-ipv6 ip-address {
                mac mac-address;
            }
        }
        overrides {
            no-dhcpv6-options;
            no-option16;
            no-option18;
            no-option37;
            no-option82;
            trusted;
            untrusted;
        }
    }
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options **dhcp-security**]

Release Information Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. Support for the **static-ipv6** and **no-option37** statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

Description Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN. A group must contain at least one interface.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)
- [Enabling a Trusted DHCP Server \(CLI Procedure\) on page 128](#)
- [Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices](#)

host-name

| | |
|------------------------------------|--|
| Syntax | host-name <i>host-name</i> ; |
| Hierarchy Level (EX Series) | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82 remote-id] |
| Hierarchy Level (MX Series) | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security remote-id option-82] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D10. Statement introduced in Junos OS Release 14.1 for the MX Series. |
| Description | Use the hostname of the switching device as the remote-id suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 63 • Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure) • RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046 |

id

| | |
|---------------------------------|--|
| Syntax | <pre>id { mac-address <i>mac-address</i>; port-id <i>port-id-number</i>; }</pre> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Specify a MAC address and a port that traffic on the link must be from to be accepted by the interface when MACsec is enabled using static secure association key (SAK) security mode. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

interface (DHCP Security)

| | |
|---------------------------------|--|
| Syntax | <pre> interface <i>interface-name</i> { static-ip <i>ip-address</i> { mac <i>mac-address</i>; } static-ipv6 <i>ip-address</i> { mac <i>mac-address</i>; } } </pre> |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i>] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Support for the static-ipv6 statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p> |
| Description | <p>Configure an interface for a static IPv4 or IPv6 address to MAC address binding (IP-MAC binding) or configure an interface to belong to a group within the VLAN that has DHCP security attributes that are different from the attributes of other interfaces in the VLAN.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure) on page 65 • Enabling a Trusted DHCP Server (CLI Procedure) on page 128 • Configuring Port Security Features on page 21 |

interface (RA Guard)

Syntax `interface (interface-name) {
 mark-interface (trusted | block);
 policy policy-name;
 (stateful | stateless);
 }`

Hierarchy Level [edit forwarding-options [access-security router-advertisement-guard](#)]

Release Information Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.
Statement introduced in Junos OS Release 16.1 for EX Series switches.

Description Configure IPv6 Router Advertisement (RA) guard on an interface. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.

Before you can configure RA guard on an interface, you must first configure a policy at the [edit forwarding-options [access-security router-advertisement-guard](#)] hierarchy level. The policy is then applied to an interface at the [edit forwarding-options [access-security router-advertisement-guard interface *interface-name*](#)] hierarchy level.



NOTE: If you apply an RA guard policy on an interface, you must enable RA guard on the VLAN that is associated with that interface by using the `vlan` statement at the [edit forwarding-options [access-security router-advertisement-guard](#)] hierarchy level.

The remaining statements are explained separately.

Options *interface-name*—Configure RA guard parameters on the specified interface.

Required Privilege Level `interface`—To view this statement in the configuration.
 `interface-control`—To add this statement to the configuration.

Related Documentation

- [Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74](#)
- [Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79](#)

interface-mac-limit

| | |
|----------------------------|--|
| Syntax | <pre>interface-mac-limit <i>limit</i> { packet-action drop; }</pre> |
| Hierarchy Level | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options], [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> switch-options], [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> |
| Description | <p>Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p> |



NOTE: For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at configuration` statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers or switches by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

Default For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

Options *limit*—Maximum number of MAC addresses learned from an interface.

Range: 1 through 524287 MAC addresses per interface

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Understanding Layer 2 Learning and Forwarding for Bridge Domains*
- *Layer 2 Learning and Forwarding for VLANs Overview*
- *Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

interfaces (MACsec)

| | |
|---------------------------------|--|
| Syntax | <pre>interfaces <i>interface-name</i> { connectivity-association <i>connectivity-association-name</i>; }</pre> |
| Hierarchy Level | [edit security macsec] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Applies the specified connectivity association to the specified interface to enable MACsec.</p> <p>One connectivity association can be applied to multiple interfaces.</p> <p>You must always use this statement to apply a connectivity association to an interface to enable MACsec. You must complete this configuration step regardless of whether MACsec is enabled using static connectivity association key (CAK) security mode or static secure association key (SAK) security mode.</p> <p>If you are enabling MACsec using static SAK security mode and need to configure MACsec on inbound and outbound traffic on the same interface, you must configure a connectivity association with one secure channel for inbound traffic and a second secure channel for outbound traffic. The connectivity association is then applied to the interface using this statement to enable MACsec for traffic entering and leaving the interface.</p> |
| Default | Interfaces are not associated with any connectivity associations, by default. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

ip-source-guard

| | |
|----------------------------|---|
| Syntax | <code>ip-source-guard;</code> |
| Hierarchy Level | <ul style="list-style-type: none">For platforms with ELS: [edit vlans <i>vlan-name</i> forwarding-options dhcp-security]For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] |
| Release Information | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> |
| Description | <p>Perform IP source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN or VLAN range. Forward packets with valid addresses and drop those with invalid addresses.</p> <ul style="list-style-type: none">ip-source-guard—Enable IP source guard checking.no-ip-source-guard—(Not available in [edit vlans <i>vlan-name</i> forwarding-options dhcp-security]) Disable IP source guard checking. <p>If you configure IP source guard at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level:</p> <ul style="list-style-type: none">IP source guard can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.DHCP snooping is automatically enabled. <p>See “Configuring IP Source Guard (CLI Procedure)” on page 32 for more information about this configuration.</p> <p>If you configure IP source guard at the [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] hierarchy level:</p> <ul style="list-style-type: none">You must enable DHCP snooping on all VLANs if you configure IP source guard on all VLANs.You must enable DHCP snooping for the specific VLAN if you configure IP source guard on that specific VLAN. Otherwise, the default behavior of no DHCP snooping applies to that VLAN. <p>See <i>Enabling DHCP Snooping (CLI Procedure)</i> for more information about this configuration.</p> |



NOTE: On EX9200 switches, IP source guard is not supported in an MC-LAG scenario.

| | |
|---------------------------------|--|
| Default | Disabled. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN</i>• <i>Example: Configuring IP Source Guard with Other EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces</i>• Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33• Configuring IP Source Guard (CLI Procedure)• Configuring IP Source Guard (CLI Procedure) on page 32 |

ipv6-prefix-list

| | |
|---------------------------------|--|
| Syntax | <code>ipv6-prefix-list <i>prefix-list-name</i>;</code> |
| Hierarchy Level | <code>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> discard]</code> <code>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> accept match-list]</code> |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure a list of IPv6 address prefixes for an IPv6 Router Advertisement (RA) guard policy. The policy is used to validate the source IPv6 address prefix of an incoming RA message against the IPv6 address prefixes in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>You can use a list of IPv6 address prefixes for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA guard policy, you must configure the list name at the <code>[edit policy-options prefix-list]</code> hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.</p> |
| Options | <i>prefix-list-name</i> —Configure a list of IPv6 address prefixes for an RA guard policy. The policy is used to validate the source of an incoming RA message by comparing the IPv6 address prefix of the RA message to the IPv6 address prefixes contained in the list. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74• Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

ipv6-source-address-list

| | |
|---------------------------------|--|
| Syntax | <code>ipv6-source-address-list <i>address-list-name</i>;</code> |
| Hierarchy Level | <p>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> discard]</p> <p>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> accept match-list]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.</p> <p>Statement introduced in Junos OS Release 16.1 for EX Series switches.</p> |
| Description | <p>Configure a list of IPv6 addresses for an IPv6 Router Advertisement (RA) guard policy to validate the source IPv6 address of an incoming RA message against the IPv6 addresses in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>You can use a list of IPv6 addresses for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA policy, you must configure the list name at the [edit policy-options prefix-list] hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.</p> |
| Options | <p><i>address-list-name</i>—Configure a list of IPv6 addresses to use in an RA guard policy. The policy is used to validate the source of an incoming RA message by comparing the IPv6 address of the RA message to the IPv6 addresses contained in the list.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

ipv6-source-guard

| | |
|--|--|
| Syntax | ipv6-source-guard; |
| Hierarchy Level | <ul style="list-style-type: none">For platforms with Enhanced Layer 2 Software (ELS): [edit vlans <i>vlan-name</i> forwarding-options dhcp-security];For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support at the [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Perform IPv6 source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN. Forward packets with valid addresses and drop those with invalid addresses. |
| <div> NOTE: If you configure the <code>ipv6-source-guard</code> statement at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.</div> <div>If you configure the <code>ipv6-source-guard</code> statement at the [edit ethernet-switching-options secure-access-port vlan <i>vlan-name</i>] hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN.</div> | |
| Default | Disabled. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38Configuring IP Source Guard (CLI Procedure) on page 32 |

key (MACsec)

| | |
|---------------------------------|--|
| Syntax | <code>key key-string;</code> |
| Hierarchy Level | [edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name security-association security-association-number] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the static security key to exchange to enable MACsec using static secure association key (SAK) security mode.</p> <p>The key string is a 32-digit hexadecimal number. The key string and the security association must match on both sides of an Ethernet connection to secure traffic using MACsec when enabling MACsec using SAK security mode.</p> <p>You must configure at least two security associations with unique security association numbers and key strings to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p> |
| Default | This statement does not have a default value. |
| Options | key-string —Specifies the key to exchange with the other end of the link on the secure channel. The <i>key-string</i> is a 32-digit hexadecimal string that is created by the user. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

key-server-priority (MACsec)

| | |
|---------------------------------|--|
| Syntax | <code>key-server-priority <i>priority-number</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> mka] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the key server priority used by the MACsec Key Agreement (MKA) protocol to select the key server when MACsec is enabled using static connectivity association key (CAK) security mode.</p> <p>The switch with the lower <i>priority-number</i> is selected as the key server.</p> <p>If the <i>priority-number</i> is identical on both sides of a point-to-point link, the MKA protocol selects the device with the lower MAC address as the key server.</p> |
| Default | The default key server priority number is 16. |
| Options | <p><i>priority-number</i>—Specifies the MKA server election priority number.</p> <p>The <i>priority-number</i> can be any number between 0 and 255. The lower the number, the higher the priority.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

mac

| | |
|---------------------------------|---|
| Syntax | <code>mac mac-address;</code> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with Enhanced Layer 2 Software (ELS): <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i> static-ip <i>ip-address</i>]</code> For platforms without ELS: <code>[edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>) static-ip <i>ip-address</i> vlan <i>vlan-name</i>]</code> For MX Series platforms: <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i> static-ip <i>ip-address</i>]</code> |
| Release Information | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 14.1 for the MX Series.</p> |
| Description | Configure the media access control (MAC) address or hardware address of the device connected to the specified interface. |
| Options | <i>mac-address</i> —Value (in hexadecimal format) of the address assigned to this device. |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)</i> <i>Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure)</i> on page 65 <i>Configuring Static IP Addresses for DHCP Bindings on Access Ports for MX Series Routers (CLI Procedure)</i> |

mac-address (MACsec)

| | |
|---------------------------------|---|
| Syntax | <code>mac-address <i>mac-address</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i> id] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specify a MAC address to enable MACsec using static secure association key (SAK) security mode. The mac-address variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.</p> <p>If you are configuring a MAC address on a secure channel in the outbound direction, you should specify the MAC address of the interface as the mac-address.</p> <p>If you are configuring a MAC address on a secure channel in the inbound direction, you should specify the MAC address of the interface at the other end of the link as the mac-address.</p> <p>You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.</p> |
| Default | No MAC address is specified in the secure channel, by default. |
| Options | mac-address —The MAC address, in six groups of two hexadecimal digits. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

mac-list

| | |
|---------------------------------|--|
| Syntax | <code>mac-list <i>name</i> { <i>mac-addresses</i>; }</code> |
| Hierarchy Level | [edit policy-options] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | Define a list of MAC addresses for use in an IPv6 Router Advertisement (RA) guard policy. |
| Options | <i>mac-addresses</i> —List of MAC addresses, one MAC address per line in the configuration. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Understanding Prefix Lists for Use in Routing Policy Match Conditions</i> |

mac-move-limit

| | |
|----------------------------|---|
| Syntax | <pre>mac-move-limit { limit; <action action packet-action action>; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: [edit vlans <i>vlan-name</i> switch-options] For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> switch-options] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> |
| Description | Specify the number of times a MAC address can move to a new interface (port) in one second and the action to be taken by the switch if the MAC address move limit is exceeded. |
| Default | If you do not specify mac-move-limit , the default MAC address move limit is unlimited. |
| Options | <p>limit <i>limit</i>—Maximum number of moves to a new interface per second.</p> <ul style="list-style-type: none"> action <i>action</i>—(Optional) (Available <i>only</i> under the hierarchy level [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) mac-move-limit]) Action to take when the MAC address move limit is reached: <ul style="list-style-type: none"> drop—Drop the packet and generate a system log entry. This is the default. log—Do not drop the packet but generate a system log entry. none—No action. shutdown—Logically disable the interface and generate a system log entry. If you have configured the switch with the port-error-disable statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the clear ethernet-switching port-error command. packet-action <i>action</i>—(Optional) (Available <i>only</i> under the hierarchy level, [edit vlans <i>vlan-name</i> switch-options mac-move-limit]) Action to take when the MAC address move limit is reached: |



NOTE: There is no default action.

- drop**—Drop the packet and do not generate an alarm.

- **drop and log**—Drop the packet and generate an alarm, an SNMP trap, or system log entry.
- **log**— Do not drop the packet, but generate an alarm, an SNMP trap, or a system log entry.
- **none**—No action.
- **shutdown**—Logically disable the interface and generate an alarm or an SNMP trap. If you have configured the interface with the [recovery-timeout](#) statement, the disabled interface recovers automatically upon expiration of the specified timeout. If you have not configured the interface for a recovery timeout, you can bring up the disabled interface by running the operational command [clear ethernet-switching recovery-timeout](#).

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Configuring Basic Port Security Features*
 - *Configuring MAC Move Limiting (CLI Procedure)*
 - [Configuring MAC Move Limiting \(CLI Procedure\) on page 94 \(ELS\)](#)
 - [Configuring Persistent MAC Learning \(CLI Procedure\) on page 98](#)
 - *Configuring MAC Move Limiting (J-Web Procedure)*
 - *Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)*
 - [Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\) on page 96](#)

macsec

```
Syntax  macsec {
        connectivity-association connectivity-association-name {
            exclude-protocol protocol-name;
            include-sci;
            mka {
                must-secure;
                key-server-priority priority-number;
                transmit-interval interval;
            }
            no-encryption;
            offset (0|30|50);
            pre-shared-key {
                cak hexadecimal-number;
                ckn hexadecimal-number;
            }
            replay-protect {
                replay-window-size number-of-packets;
            }
            secure-channel secure-channel-name {
                direction (inbound | outbound);
                encryption (MACsec);
                id {
                    mac-address mac-address;
                    port-id port-id-number;
                }
                offset (0|30|50);
                security-association security-association-number {
                    key key-string;
                }
            }
            security-mode security-mode;
        }
        interfaces interface-name {
            connectivity-association connectivity-association-name;
        }
    }
```

Hierarchy Level [edit security]

Release Information Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Configure Media Access Control Security (MACsec)..

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- [Configuring Media Access Control Security \(MACsec\) on page 109](#)

mark-interface (RA Guard)

| | |
|---------------------------------|--|
| Syntax | <code>mark-interface (trusted block);</code> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard interface interface-name] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure an interface as blocked or trusted for IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. RA guard performs checks on incoming RA messages to make sure that they are sent from legitimate routers. If the sender of the RA message cannot be validated, the RA message is dropped.</p> <p>You can configure the mark-interface statement on an interface to bypass RA guard policy checks on that interface. If an interface is configured as either a trusted interface or a blocked interface, RA messages received on the interface are not subject to inspection by RA guard, even if the interface or VLAN is enabled for RA guard. If the interface is trusted, it forwards all RA messages. If the interface is blocked, it drops all RA messages.</p> |
| Options | <p>block—Configure an interface as blocked for bypassing inspection of RA messages received on that interface by RA guard. When you configure an interface as blocked, all RA messages received on the interface are dropped.</p> <p>trusted—Configure an interface as trusted for bypassing inspection of RA messages received on that interface by RA guard. When you configure an interface as trusted, all RA messages received on the interface are forwarded.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

match-list

| | |
|--------------------------|---|
| Syntax | <pre>match-list { (match-any match-all); ipv6-prefix-list <i>prefix-list-name</i>; ipv6-source-address-list <i>address-list-name</i>; source-mac-address-list <i>address-list-name</i>; }</pre> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard <i>policy</i> <i>policy-name</i> accept] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure one or more lists of IPv6 addresses, MAC addresses, or IPv6 address prefixes to be associated with an IPv6 Router Advertisement (RA) guard <i>accept</i> policy.</p> <p>RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>You can configure match lists in either an accept policy or a discard policy. When RA guard is enabled by using an accept policy, RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.</p> <p>You can associate match lists or match conditions (see match-option) with an accept policy. You can configure match lists that be associated with an accept policy by using the match-list statement. The lists configured by using the match-list statement can contain IPv6 addresses, MAC addresses, or IPv6 address prefixes. RA guard examines the source address or address prefix. You configure the lists at the [edit policy-options] hierarchy level by using the prefix-list option for an IPv6 address or address prefix list, and mac-list for a MAC address list.</p> |
| Options | <p>match-all—Configure the RA guard policy so that a received RA message is accepted only if it matches criteria in all of the lists configured under match-list; otherwise, the message is discarded.</p> <p>match-any—Configure the RA guard policy so that a received RA message is accepted if it matches criteria in any of the lists configured under match-list; otherwise, the message is discarded.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |

- Related Documentation**
- [Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74](#)
 - [Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79](#)

match-option

| | |
|----------------------------|--|
| Syntax | <pre>match-option { hop-limit { (maximum minimum) value; } managed-config-flag; other-config-flag; router-preference maximum (high low medium); }</pre> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> accept] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure one or more parameters such as hop-count limit, managed configuration flag, other configuration flag, or router preference priority as the match condition to be associated with an IPv6 Router Advertisement (RA) guard <i>accept</i> policy.</p> <p>RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>You can associate match lists (see match-list) or match conditions with an accept policy. You can configure match conditions by using the match-option statement in an RA guard accept policy. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped.</p> |
| Options | <p>hop-limit—Configure the RA guard policy to verify the minimum or maximum hop count for an incoming RA message. Use maximum to set a maximum hop count, or minimum to set a minimum hop count.</p> <p>managed-config-flag—Configure the RA guard policy to verify that the managed address configuration flag of an incoming RA message is set. When the managed address configuration flag is set, it indicates that addresses are available for allocation by Dynamic Host Configuration Protocol version 6 (DHCPv6).</p> <p>other-config-flag—Configure the RA guard policy to verify that the other configuration flag of an incoming RA message is set. When this flag is set, it indicates that other configuration information is available through DHCPv6. Examples of such information are DNS-related information or information on other servers within the network.</p> <p>router-preference-maximum—Configure the RA guard policy to verify that the default router preference parameter value of an incoming RA message is lower than or equal to a specified limit. The default router preference value improves the ability of IPv6 hosts to select a default router to reach a remote destination when the host has</p> |

multiple routers on its default router list. Use **high**, **medium**, or **low** to set the maximum preference.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |


mka

| | |
|---------------------------------|--|
| Syntax | <pre>mka { must-secure; key-server-priority <i>priority-number</i>; transmit-interval <i>interval</i>; }</pre> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15. |
| Description | Specify parameters for the MACsec Key Agreement (MKA) protocol. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

must-secure

| | |
|---------------------------------|---|
| Syntax | must-secure; |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> mka] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10. |
| Description | <p>Specifies that all traffic travelling on the MACsec-secured link must be MACsec-secured to be forwarded onward.</p> <p>When the must-secure option is enabled, all traffic that is not MACsec-secured that is received on the interface is dropped.</p> <p>When the must-secure option is disabled, all traffic from devices that support MACsec is MACsec-secured while traffic received from devices that do not support MACsec is forwarded through the network.</p> <p>The must-secure option is particularly useful in scenarios where multiple devices, such as a phone and a PC, are accessing the network through the same Ethernet interface. If one of the devices supports MACsec while the other device does not support MACsec, the device that doesn't support MACsec can continue to send and receive traffic over the network—provided the must-secure option is disabled—while traffic to and from the device that supports MACsec is MACsec-secured. In this scenario, traffic to the device that is not MACsec-secured must be VLAN-tagged.</p> |
| Default | The must-secure option is disabled. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

neighbor-discovery-inspection

| | |
|---|---|
| Syntax | neighbor-discovery-inspection; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security]; [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support at the [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | <p>Perform dynamic IPv6 neighbor discovery inspection on the specified VLAN.</p> <p>When neighbor discovery inspection is configured, the switch inspects IPv6 packets with neighbor discovery messages and validates them against the DHCPv6 binding table. The source IP address and source MAC address of each packet are checked against the table, and if a valid match is not found, the packet is dropped.</p> |
| <div>  <p>NOTE: If you configure the neighbor-discovery-inspection statement at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level, DHCPv6 snooping is automatically enabled for the specified VLAN.</p> <p>See “Enabling IPv6 Neighbor Discovery Inspection” on page 71 for more information about this configuration.</p> <p>If you configure the neighbor-discovery-inspection statement at the [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>)] hierarchy level, you must also enable DHCPv6 snooping for the specified VLAN or VLANs.</p> </div> | |
| Default | Disabled. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Port Security Features on page 21 • Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38 • Enabling IPv6 Neighbor Discovery Inspection on page 71 |

no-dhcpv6-options

| | |
|---------------------------------|--|
| Syntax | no-dhcpv6-options; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> overrides] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure a specific group of one or more access interfaces within the VLAN not to add any DHCPv6 options, even if the VLAN is configured to perform DHCPv6 snooping. DHCPv6 options include option 16, option 18, and option 37. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• dhcpv6-options on page 150• Understanding DHCP Snooping for Port Security on page 51• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59 |

no-dhcpv6-snooping

| | |
|---------------------------------|---|
| Syntax | no-dhcpv6-snooping; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Disable DHCPv6 snooping for the specified VLAN. |
| Default | DHCPv6 snooping is not enabled by default. There is no configuration statement that explicitly enables DHCPv6 snooping. DHCPv6 snooping is enabled automatically by Junos OS if any port security feature, such as IPv6 neighbor discovery inspection or IPv6 source guard, is configured at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

no-encryption (MACsec)

| | |
|---------------------------------|---|
| Syntax | no-encryption; |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Disables MACsec encryption for a connectivity association that is configured to enable MACsec using static connectivity association key (CAK) or dynamic security mode.</p> <p>You can enable MACsec without enabling encryption. If a connectivity association that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link. The MACsec header is still applied to the packet, however, and all MACsec data integrity checks are run on both ends of the link to ensure the traffic does not represent a security threat.</p> <p>This command is used to disable encryption when MACsec is configured using static CAK or dynamic security mode only. When MACsec is configuring using static secure association key (SAK) security mode, the encryption setting is managed in the secure channel using the encryption configuration statement.</p> |
| Default | MACsec encryption is enabled if MACsec is enabled using static CAK or dynamic security mode. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

no-option16

| | |
|---------------------------------|---|
| Syntax | no-option16; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group group-name overrides] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure a specific group of one or more access interfaces within the VLAN not to transmit DHCPv6 option 16 information, even if the VLAN is configured to perform DHCPv6 snooping. Option 16 information that has already been added by a DHCPv6 client will be forwarded as is. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• option-16 on page 189• Understanding DHCP Snooping for Port Security on page 51• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59 |


no-option18

| | |
|---------------------------------|---|
| Syntax | no-option18; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group group-name overrides] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure a specific group of one or more access interfaces within the VLAN <i>not</i> to transmit DHCP option 18 information, even if the VLAN is configured to perform DHCPv6 snooping. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• option-18 on page 190• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59• Understanding DHCP Snooping for Port Security on page 51 |

no-option37

| | |
|---------------------------------|---|
| Syntax | no-option37; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> overrides] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Configure a specific group of one or more access interfaces within the VLAN <i>not</i> to transmit DHCP option 37 information, even if the VLAN is configured to perform DHCPv6 snooping. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• option-37 on page 191• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

offset

| | |
|----------------------------|---|
| Syntax | offset (0 30 50); |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the number of octets in an Ethernet frame that are sent in unencrypted plain-text when encryption is enabled for MACsec.</p> <p>Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.</p> <p>You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.</p> <p>You configure the offset in the [edit security macsec connectivity-association <i>connectivity-association-name</i>] hierarchy when you are enabling MACsec using static connectivity association key (CAK) or dynamic security mode.</p> <p>You configure the offset in the [edit security macsec connectivity-association <i>connectivity-association-name</i> secure-channel <i>secure-channel-name</i>] hierarchy when you are enabling MACsec using static secure association key (SAK) security mode.</p> |
| Default | 0 |
| Options | <p>0—Specifies that no octets are unencrypted. When you set the offset to 0, all traffic on the interface where the connectivity association or secure channel is applied is encrypted.</p> <p>30—Specifies that the first 30 octets of each Ethernet frame are unencrypted.</p> |
| | <p> NOTE: In IPv4 traffic, setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 30, therefore, is typically used when a feature needs this information to perform a task on IPv4 traffic.</p> |
| | <p>50—Specified that the first 50 octets of each Ethernet frame are unencrypted.</p> |



NOTE: In IPv6 traffic, setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the rest of the traffic. An offset of 50, therefore, is typically used when a feature needs this information to perform a task on IPv6 traffic.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation • [Configuring Media Access Control Security \(MACsec\) on page 109](#)

option-16 (DHCPv6 Snooping)

Syntax `option-16 {
 use-string string;
}`

Hierarchy Level [edit vlans *vlan-name* forwarding-options [dhcp-security](#) [dhcpv6-options](#)]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Configure the DHCPv6 Vendor ID option (option 16) to be included in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 16 provides the server with information about the vendor that manufactured the hardware on which the DHCP client is running. When configured, the switch will overwrite any existing option 16 information sent by clients in the DHCPv6 packets.

Option 16 is the DHCPv6 equivalent of the [vendor-id](#) sub-option of DHCP option 82.

Options `use-string string`—Define a custom string to be used as the DHCPv6 vendor identifier.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [Enabling Dynamic ARP Inspection \(CLI Procedure\) on page 48](#)
• [Configuring IP Source Guard \(CLI Procedure\) on page 32](#)
• [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)
• [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)

option-18 (DHCPv6 Snooping)

Syntax

```
option-18 {
  prefix {
    host-name;
    logical-system-name;
    routing-instance-name;
    vlan-id;
    vlan-name;
  }
  use-interface-index (device | logical);
  use-interface-description (device | logical);
  use-interface-mac;
  use-interface-name (device | logical);
  use-string string;
}
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options **dhcp-security** **dhcpv6-options**]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Configure the DHCPv6 Relay Agent Interface ID option (option 18) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 18 provides information about the port on which the request was received, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 18 is configured, a unique interface ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the interface ID or to change the type of information used to identify the interface. The default fields included in option 18 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 18 is the DHCPv6 equivalent of the **circuit-id** sub-option of DHCP option 82.



NOTE: DHCPv6 packets that already contain option 18 information when received from a client are dropped by the switch.

Options **use-interface-mac**—Use the MAC address of the interface in the DHCPv6 interface ID.

use-string *string*—Use a custom string in the DHCPv6 interface ID.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [no-option18 on page 186](#)
- [no-dhcpv6-options on page 184](#)

option-37 (DHCPv6 Snooping)

```
Syntax  option-37 {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
            vlan-id;
            vlan-name;
        }
        use-interface-index (device | logical);
        use-interface-description (device | logical);
        use-interface-mac;
        use-interface-name (device | logical);
        use-string string;
    }
```

Hierarchy Level [edit vlans *vlan-name* forwarding-options [dhcp-security dhcpv6-options](#)]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Configure the DHCPv6 Relay Agent Remote ID option (option 37) to insert information in DHCPv6 requests from clients before forwarding them to a DHCPv6 server. Option 37 provides information about the remote host, which the server can use to assign IP addresses, prefixes, and other configuration parameters for the client.

When option 37 is configured, a unique remote ID is inserted into the DHCPv6 packet headers. Suboptions can be configured to include a prefix with the remote ID or to change the interface portion of the ID. The default fields included in option 37 are the Juniper Enterprise ID, VLAN, and MAC address of the interface.

Option 37 is the DHCPv6 equivalent of the [remote-id](#) sub-option of DHCP option 82.



NOTE: DHCPv6 packets that already contain option 37 information when received from a client are dropped by the switch.

Options **use-interface-mac**—Use the MAC address of the interface in the DHCPv6 remote ID.

use-string *string*—Use a custom string in the DHCPv6 remote ID.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [no-option37 on page 187](#)
- [no-dhcpv6-options on page 184](#)

- [Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server \(CLI Procedure\) on page 63](#)
- [Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports \(CLI Procedure\) on page 65](#)

overrides (DHCP Security)

| | |
|---------------------------------|---|
| Syntax | <pre>overrides { no-dhcpv6-options; no-option16; no-option18; no-option37; no-option82; trusted; untrusted; }</pre> |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i>] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Support for the no-option37 option introduced in Junos OS Release 13.2X51-D20 for EX Series switches.</p> <p>Support for the no-dhcpv6-options, no-option16 and no-option18 options introduced in Junos OS Release 14.2 for EX Series switches.</p> |
| Description | Modify selected DHCP attributes for a group of interfaces that is configured within a specified VLAN. |
| Options | <p>trusted—The interface specified in this group is trusted. DHCP snooping and DHCPv6 snooping do not apply to the trusted interface. Likewise, DAI, IP source guard, IPv6 source guard, and IPv6 neighbor discovery inspection—even if they are enabled for the VLAN—do not apply to the interface that is configured with the overrides and the trusted options. Access interfaces are untrusted by default.</p> <p>untrusted—(Only for EX9200) The interface specified in this group is untrusted. Trunk interface are trusted by default. Access interfaces are untrusted by default.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Enabling a Trusted DHCP Server (CLI Procedure) on page 128• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59 |

packet-action

Syntax `packet-action action;`

Hierarchy Level [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],
 [edit protocols l2-learning global-mac-limit *limit*],
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (VPLS)],
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (VPLS)],
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit switch-options **interface-mac-limit** *limit*],
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit switch-options **interface-mac-limit** *limit*],
 [edit switch-options mac-table-size *limit*],
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],
 [edit vlans *vlan-name* switch-options mac-table-size *limit*]
 [edit [vlans on page 133](#) *vlan-name* switch-options **interface-mac-limit** *limit*],
 [edit [vlans on page 133](#) *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit [vlans on page 133](#) *vlan-name* switch-options mac-table-size *limit*]

Release Information Statement introduced in Junos OS Release 8.4.
 Support for the **switch-options** statement added in Junos OS Release 9.2.
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy

supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

Description Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

Default



NOTE: On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

Options

- drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.
- drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.
- shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

| | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring EVPN Routing Instances</i> • <i>Configuring EVPN Routing Instances on EX9200 Switches</i> • Configuring MAC Limiting (CLI Procedure) on page 91 • Configuring Persistent MAC Learning (CLI Procedure) on page 98 • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains</i> • <i>Layer 2 Learning and Forwarding for VLANs Overview</i> • <i>Understanding Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i> • <i>Layer 2 Learning and Forwarding for VLANs Overview</i> • <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i> |
|------------------------------|---|

persistent-learning

| | |
|---------------------------------|--|
| Syntax | <code>persistent-learning;</code> |
| Hierarchy Level | <code>[edit switch-options interface <i>interface-name</i>]</code> |
| Release Information | Hierarchy level <code>[edit switch-options interface interface-name]</code> introduced in Junos OS Release 13.2X50-D10 |
| Description | Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC. |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring Basic Port Security Features</i> • Configuring Persistent MAC Learning (CLI Procedure) on page 98 |

policy

Syntax

```

policy policy-name {
  discard {
    ipv6-source-address-list address-list-name;
    source-mac-address-list address-list-name;
    ipv6-prefix-list prefix-list-name;
  }
  accept {
    match-list {
      (match-any | match-all);
      ipv6-source-address-list address-list-name;
      source-mac-address-list address-list-name;
      ipv6-prefix-list prefix-list-name;
    }
    match-option {
      hop-limit {
        (maximum | minimum) value;
      }
      managed-config-flag;
      other-config-flag;
      router-preference (high | low | medium);
    }
  }
}

```

Hierarchy Level

```

[edit forwarding-options access-security router-advertisement-guard]
[edit forwarding-options access-security router-advertisement-guard interface interface-name]
[edit forwarding-options access-security router-advertisement-guard vlan vlan-name]

```

Release Information

Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.
Statement introduced in Junos OS Release 16.1 for EX Series switches.

Description

Configure the policy for an IPv6 Router Advertisement (RA) guard. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages based on whether they match the conditions defined in the policy.

RA guard compares the information contained in attributes of RA messages to the information contained in the policy. You must configure the policy before you can enable RA guard. You can configure either an accept policy or a discard policy and enable it on an interface or on a VLAN. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions defined in the policy are dropped, and RA messages that do not match the conditions are forwarded.

The remaining statements are explained separately.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

port-id

| | |
|---------------------------------|--|
| Syntax | <code>port-id <i>port-id-number</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association connectivity-association-name secure-channel secure-channel-name <i>id</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specify a port ID in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.</p> <p>Once the port numbers match, MACsec is enabled for all traffic on the connection.</p> <p>You only use this configuration option when you are configuring MACsec using static SAK security mode. This option does not need to be specified when you are enabling MACsec using static connectivity association key (CAK) security mode.</p> |
| Default | No port ID is specified. |
| Options | <i>port-id-number</i> —The port ID number. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

pre-shared-key

| | |
|---------------------------------|---|
| Syntax | <pre>pre-shared-key { cak hexadecimal-number; ckn hexadecimal-number; }</pre> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the pre-shared key used to enable MACsec using static connectivity association key (CAK) security mode.</p> <p>A pre-shared key includes a connectivity association key name (CKN) and a connectivity association key (CAK). A pre-shared key is exchanged between two devices at each end of a point-to-point link to enable MACsec using static CAK security mode. The MACsec Key Agreement (MKA) protocol is enabled after the pre-shared keys are successfully verified and exchanged. The pre-shared key—the CKN and CAK—must match on both ends of a link.</p> |
| Default | No pre-shared keys exist, by default. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |


prefix (DHCPv6 Options)

| | |
|---------------------------------|--|
| Syntax | <pre> prefix { host-name; logical-system-name; routing-instance-name; vlan-id; vlan-name; } </pre> |
| Hierarchy Level | [edit vlans forwarding-options dhcp-security dhcpv6-options option-18] [edit vlans forwarding-options dhcp-security dhcpv6-options option-37] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure a prefix for DHCPv6 option 18 (Interface ID) or option 37 (Remote ID). When configured, the prefix is inserted into DHCPv6 packets during the DHCPv6 snooping process. |
| Default | If the prefix statement is not explicitly specified, no prefix is inserted in DHCPv6 packets. |
| Options | <p>host-name—Add the host name of the switch to DHCPv6 options.</p> <p>logical-system-name—Add the logical system name to the DHCPv6 options.</p> <p>routing-instance-name—Add the routing instance name to the DHCPv6 options.</p> <p>vlan-id—Add the VLAN ID to the DHCPv6 options.</p> <p>vlan-name—Add the VLAN name to the DHCPv6 options.</p> |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • option-37 (DHCPv6 Snooping) on page 191 • option-18 (DHCPv6 Snooping) on page 190 |

prefix-list

| | |
|---------------------------------|---|
| Syntax | <pre>prefix-list name { ip-addresses; apply-path path; }</pre> |
| Hierarchy Level | [edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for configuration in the dynamic database introduced in Junos OS Release 9.5. Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches. Support for the vpls protocol family introduced in Junos OS Release 10.2. |
| Description | <p>Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.</p> <p>You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.</p> |
| Options | <p>name—Name that identifies the list of IPv4 or IPv6 address prefixes.</p> <p>ip-addresses—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Understanding Prefix Lists for Use in Routing Policy Match Conditions</i>• <i>Configuring Routing Policies and Policy Objects in the Dynamic Database</i>• <i>dynamic-db</i>• "Firewall Filter Match Conditions Based on Address Fields" in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>• "Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List" in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i> |

recovery-timeout

| | |
|---|---|
| Syntax | <code>recovery-timeout seconds;</code> |
| Hierarchy Level (EX Series and QFX Series) | [edit interfaces <i>interface-name</i> unit 0 family ethernet-switching] |
| Hierarchy Level (MX Series) | [edit interfaces <i>interface-name</i> unit 0 family bridge] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for the MX Series routers. |
| Description | <p>Configure an interface to be temporarily disabled when MAC limiting, MAC move limiting, or rate-limiting is in effect with the action shutdown. This enables the affected interface to recover automatically from the error condition after the specified period of time:</p> <ul style="list-style-type: none"> • If you configure MAC limiting with the shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the MAC address limit is reached. The interface will recover automatically after the number of seconds specified. • If you enable MAC move limiting with the shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified. • If you enable MAC move limiting with the vlan-member-shutdown option and you enable recovery-timeout, the interface is temporarily disabled when the maximum number of MAC address moves is reached. The interface will recover automatically after the number of seconds specified. If the recovery timeout is not configured, the interface will recover automatically after 180 seconds. • If you enable storm control with the action-shutdown option and you enable recovery-timeout, the interface is temporarily disabled when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic can include broadcast, unknown unicast, and multicast traffic. |
| | <p> NOTE: The recovery-timeout configuration does not apply to preexisting error conditions. It impacts only error conditions that are detected after the recovery-timeout statement is configured and committed. To clear a preexisting error condition and restore the interface to service, use the operational mode commands clear ethernet-switching recovery-timeout for EX Series and QFX Series and clear bridge recovery-timeout for MX Series routers.</p> |
| Default | The interface does not automatically recover from an error condition. |



NOTE: On EX9200 switches, if a MAC move limit is configured with the action `vlan-member-shutdown`, the interface automatically recovers from the disabled condition after 180 seconds by default.

| | |
|---------------------------------|--|
| Options | <i>seconds</i> — Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery. Range: 10 through 3600 |
| Required Privilege Level | system—To view this statement in the configuration. system—control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>action-shutdown</i>• Configuring MAC Limiting (CLI Procedure) on page 91• Configuring MAC Move Limiting (CLI Procedure) on page 94• Configuring or Disabling Storm Control (CLI Procedure) |

remote-id

| | |
|----------------------------|---|
| Syntax | <pre>remote-id { host-name host-name; mac (Option 82); prefix (hostname mac none); use-interface-description (logical device); use-string string; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with Enhanced Level 2 Software (ELS): [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] For platforms without ELS: [edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82], [edit forwarding-options helpers bootp dhcp-option82], [edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82] |
| Release Information | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> |
| Description | <p>Insert the remote-id suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.</p> <p>The remaining statements are explained separately, and their availability depends on the hierarchy level at which the remote-id suboption is specified, as follows:</p> <ul style="list-style-type: none"> The statement prefix, is <i>not</i> supported at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] hierarchy level. The statement host-name is supported <i>only</i> at the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82] hierarchy level. |
| Default | <p>If the remote-id statement is not explicitly set, no remote ID value is inserted in the DHCP request packet header.</p> <p>If the remote-id statement is explicitly set, but is not qualified by a keyword, the following are true:</p> <ul style="list-style-type: none"> At the [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] hierarchy level, the default keyword value is <i>interface-name</i>. At all other hierarchy levels, the default value of the remote-id keyword is the MAC address of the switch. |



NOTE: When you configure `remote-id`, `circuit-id` is also enabled, even if you do not explicitly configure `circuit-id`.

| | |
|---------------------------------|--|
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i>• <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> on page 63• <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i>• RFC 3046, <i>DHCP Relay Agent Information Option</i>, at http://tools.ietf.org/html/rfc3046 |

replay-protect

| | |
|---------------------------------|---|
| Syntax | <pre>replay-protect { replay-window-size number-of-packets; }</pre> |
| Hierarchy Level | [edit security <code>macsec connectivity-association connectivity-association-name</code>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Enable replay protection for MACsec. A replay window size specified using the <code>replay-window-size number-of-packets</code> statement must be specified to enable replay protection. |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Media Access Control Security (MACsec)</i> on page 109 |

replay-window-size

| | |
|---------------------------------|--|
| Syntax | <code>replay-window-size <i>number-of-packets</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i> replay-protect] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the size of the replay protection window.</p> <p>This statement has to be configured to enable replay protection.</p> <p>When MACsec is enabled on an Ethernet link, an ID number is assigned to each packet entering the link. The ID number of the packet is checked by the receiving interface after the packet has traversed the MACsec-enabled link.</p> <p>When replay protection is enabled, the sequence of the ID number of received packets are checked. If the packet arrives out of sequence and the difference between the packet numbers exceeds the replay protection window size, the packet is dropped by the receiving interface. For instance, if the replay protection window size is set to five and a packet assigned the ID of 1006 arrives on the receiving link immediately after the packet assigned the ID of 1000, the packet that is assigned the ID of 1006 is dropped because it falls outside the parameters of the replay protection window.</p> <p>Replay protection is especially useful for fighting man-in-the-middle attacks. A packet that is replayed by a man-in-the-middle attacker on the Ethernet link will arrive on the receiving link out of sequence, so replay protection helps ensure the replayed packet is dropped instead of forwarded through the network.</p> <p>Replay protection should not be enabled in cases where packets are expected to arrive out of order.</p> |
| Default | Replay protection is disabled. |
| Options | <p><i>number-of-packets</i>—Specifies the size of the replay protection window, in packets.</p> <p>When this variable is set to 0, all packets that arrive out-of-order are dropped.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Media Access Control Security (MACsec) on page 109 |

router-advertisement-guard

```
Syntax router-advertisement-guard {
    interface interface-name {
        policy policy-name;
        mark-interface (trusted | block);
        (stateful | stateless);
    }
    vlan vlan-name {
        policy policy-name;
        (stateful | stateless);
    }
    policy policy-name {
        discard {
            ipv6-source-address-list address-list-name;
            source-mac-address-list address-list-name;
            ipv6-prefix-list prefix-list-name;
        }
        accept {
            match-list {
                (match-any | match-all);
                ipv6-source-address-list address-list-name;
                source-mac-address-list address-list-name;
                ipv6-prefix-list prefix-list-name;
            }
            match-option {
                hop-limit {
                    (maximum | minimum) value;
                }
                managed-config-flag;
                other-config-flag;
                router-preference (high | low | medium);
            }
        }
    }
}
```

Hierarchy Level [edit forwarding-options [access-security](#)]

Release Information Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.
Statement introduced in Junos OS Release 16.1 for EX Series switches.

Description Configure IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy. The policy can be either an accept policy or a discard policy. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.

You can enable RA guard on an interface or on a VLAN. You must first configure a policy at the `[edit forwarding-options access-security router-advertisement-guard]` hierarchy level. The policy is then applied to an interface at the `[edit forwarding-options access-security router-advertisement-guard interface interface-name]` hierarchy level, or to a VLAN at the `[edit forwarding-options access-security router-advertisement-guard vlan vlan-name]` hierarchy level.



NOTE: If you apply an RA guard policy on an interface, you must enable RA guard on the VLAN that is associated with that interface using the `vlan` statement at the `[edit forwarding-options access-security router-advertisement-guard]` hierarchy level.

You can configure RA guard to be stateless or stateful. Stateless RA guard enables a switch to examine incoming RA messages and filter each message on the basis of whether it matches the conditions configured in the policy. For example, an interface can be statically configured to forward RA messages only from predefined sources. Stateful RA guard enables a switch to learn about legitimate senders of RA messages and store this information, which is used to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages from legitimate senders dynamically transitions to the forwarding state, in which RA messages from valid senders are forwarded to their destination.

The remaining statements are explained separately.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

routing-instance-name (circuit-id)

| | |
|---------------------------------|---|
| Syntax | routing-instance--name; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82 circuit-id prefix] |
| Release Information | Statement introduced in Junos OS Release 13.2 for EX Series switches. |
| Description | Specify that the routing instance name used by the VLAN is included with the circuit ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 63• Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59 |

secure-channel

| | |
|---------------------------------|--|
| Syntax | <pre> secure-channel <i>secure-channel-name</i> { direction (inbound outbound); encryption (MACsec); id { mac-address <i>mac-address</i>; port-id <i>port-id-number</i>; } offset (0 30 50); security-association <i>security-association-number</i> { key <i>key-string</i>; } } </pre> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Create and configure a secure channel to enable and configure MACsec when MACsec is enabled using static secure association key (SAK) security mode.</p> <p>You do not need to use this option to enable MACsec using static connectivity association key (CAK) security mode. All configuration for MACsec using static CAK security mode is done inside of the connectivity association but outside of the secure channel. When MACsec is enabled using static CAK security mode, an inbound and an outbound secure channel—neither of which is user-configurable—is automatically created within the connectivity association.</p> |
| Options | The remaining statements are explained separately. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

security-association

| | |
|---------------------------------|--|
| Syntax | <code>security-association <i>security-association-number</i> { <i>key</i> <i>key-string</i>; }</code> |
| Hierarchy Level | [edit security <i>macsec</i> <i>connectivity-association</i> <i>connectivity-association-name</i> <i>secure-channel</i> <i>secure-channel-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the number of one of the security associations in the secure channel when MACsec is enabled using static secure association key (SAK) security mode. Because SAKs are created by the key server when MACsec is enabled using static connectivity association key (CAK) security mode, the security-association statement is not used when enabling MACsec using static CAK security mode.</p> <p>You must configure at least two security associations to enable MACsec using static SAK security mode. MACsec initially establishes a secure connection when a security association number and key match on both ends of an Ethernet link. After a certain number of Ethernet frames are securely transmitted across the Ethernet link, MACsec automatically rotates to a new security association with a new security association number and key to maintain the secured Ethernet link. This rotation continues each time a certain number of Ethernet frames are securely transmitted across the secured Ethernet link, so you must always configure MACsec to have at least two security associations.</p> |
| Default | No security keys are configured, by default. |
| Options | <p><i>security-association-number</i>—Specifies the security association number and creates the SAK.</p> <p>The security association number is a whole number between 0 and 3. You can configure two security associations in a secure channel when enabling MACsec using static security keys.</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Media Access Control Security (MACsec) on page 109 |

security-mode

| | |
|---------------------------------|---|
| Syntax | <code>security-mode <i>security-mode</i>;</code> |
| Hierarchy Level | [edit security macsec connectivity-association <i>connectivity-association-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15. The dynamic security mode option was introduced in Junos OS Release 14.1X53-D10. |
| Description | Configure the MACsec security mode for the connectivity association. We recommend enabling MACsec on switch-to-switch Ethernet links using static connectivity association key (CAK) security mode. Static CAK security mode ensures security by frequently refreshing to a new random secure association key (SAK) and by only sharing the SAK between the two devices on the MACsec-secured point-to-point link. Additionally, some optional MACsec features—replay protection, SCI tagging, and the ability to exclude traffic from MACsec—are only available when you enable MACsec using static CAK security mode. |
| Options | <p>security-mode—Specifies the MACsec security mode. Options include:</p> <ul style="list-style-type: none"> • dynamic—Dynamic mode. Dynamic security mode is used to enable MACsec on switch-to-host Ethernet links. In dynamic mode, a master key is retrieved from a RADIUS server by a switch and a host as part of the AAA handshake in separate transactions. The MKA protocol is enabled when the master key is exchanged between the switch and the host. • static-cak—Static connectivity association key (CAK) mode. Static CAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In static-cak mode, the switch at one end of the point-to-point link acts as the key server and regularly transmits a randomized key using a process that does not transmit any traffic outside of the MACsec-secured point-to-point link. • static-sak—Static secure association key (SAK) mode. Static SAK security mode is used to enable MACsec on switch-to-switch Ethernet links. In static-sak mode, one of two user-configured security keys is used to secure the point-to-point link. The two security keys are regularly rotated. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Media Access Control Security (MACsec) on page 109 |

source-mac-address-list

| | |
|---------------------------------|--|
| Syntax | <code>source-mac-address-list <i>address-list-name</i>;</code> |
| Hierarchy Level | <code>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> discard]</code> <code>[edit forwarding-options access-security router-advertisement-guard policy <i>policy-name</i> accept match-list]</code> |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure a list of MAC addresses for an IPv6 Router Advertisement (RA) guard policy to validate the source MAC address of an incoming RA message against the MAC addresses in this list. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in the policy.</p> <p>You can use a list of MAC address for validating RA messages as part of either an accept policy or a discard policy. Before you can include a list in an RA policy, you must configure the list name at the <code>[edit policy-options mac-list]</code> hierarchy level. When RA guard is enabled by using an accept policy, any RA messages that match the conditions defined in the policy are forwarded, and RA messages that do not match the conditions are dropped. When RA guard is enabled by using a discard policy, any RA messages that match the conditions are dropped, and RA messages that do not match the conditions are forwarded.</p> |
| Options | <i>address-list-name</i> —Configure the RA guard policy to match the MAC source address of an incoming RA message to a MAC address contained in the list. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74• Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |


stateful

| | |
|---------------------------------|--|
| Syntax | (stateful stateless); |
| Hierarchy Level | <p>[edit forwarding-options access-security router-advertisement-guard interface (<i>interface-name</i> <i>interface-range-name</i>)]</p> <p>[edit forwarding-options access-security router-advertisement-guard vlan <i>vlan-name</i>]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches.</p> <p>Statement introduced in Junos OS Release 16.1 for EX Series switches.</p> |
| Description | <p>Configure stateful IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. RA guard performs checks on incoming RA messages to make sure that they are sent from legitimate routers. If the sender of the RA message cannot be validated, the RA message is dropped.</p> <p>Stateful RA guard enables the switch to learn about the sources of RA messages for a certain period of time. During this period, when the switch is known to be in the learning state, the information contained in attributes of received RA messages is stored and compared to the policy. When the learning period ends, the switch has a record of which interfaces are attached to links with valid IPv6 routers. If there is no valid IPv6 router attached to the interface, the switch dynamically transitions the interface from the learning state into the blocking state. Subsequent RA messages received after the transition to blocking state are dropped. If there is a valid IPv6 router attached to the interface, the interface transitions to the forwarding state, and subsequent RA messages that can be validated against the configured policy are forwarded.</p> <p>You can enable stateful RA guard on an interface or on a VLAN. When you enable stateful RA guard, the initial state is Off. You initiate the learning state by issuing the request access-security router-advertisement-guard-learn command.</p> |
| Default | RA guard is stateless by default. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

stateless

| | |
|---------------------------------|---|
| Syntax | (stateful stateless); |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard interface (<i>interface-name</i> <i>interface-range-name</i>)] [edit forwarding-options access-security router-advertisement-guard vlan <i>vlan-name</i>] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure stateless IPv6 Router Advertisement (RA) guard. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.</p> <p>You can configure RA guard to be stateless or stateful. If stateless RA guard is enabled, the switch examines incoming RA messages and filters each message on the basis of whether it matches the conditions configured in the policy. After the switch has validated the content of the RA message, it forwards the RA message to its destination; otherwise, the RA message is dropped. For example, an interface can be statically configured to forward RA messages only from predefined sources.</p> <p>You can enable stateless RA guard on an interface or on a VLAN.</p> |
| Default | RA guard is stateless by default. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74• Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

static-ip

| | |
|---|--|
| Syntax | <pre>static-ip <i>ip-addresses</i> { vlan <i>vlan-name</i>; mac <i>mac-address</i>; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>] For platforms without ELS: [edit ethernet-switching-options secure-access-port interface (all <i>interface-name</i>)] |
| Release Information | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> |
| Description | Configure a static IP address to MAC address (IP-MAC) binding to be added to the DHCP snooping database. |
| <div style="display: flex; align-items: center;">  <div> <p>NOTE: The VLAN is specified at the higher hierarchy level when static-ip is configured at [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>].</p> </div> </div> | |
| Options | <p>ip-address—Static IP address assigned to a device connected on the specified interface.</p> <p>mac mac-address—Static MAC address assigned to a device connected on the specified interface.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure) Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure) on page 65 |

static-ipv6

| | |
|---------------------------------|--|
| Syntax | <code>static-ipv6 ip-address { mac mac-address; }</code> |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> interface <i>interface-name</i>]; [edit ethernet-switching-options secure-access-port interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches. Support at the [edit ethernet-switching-options secure-access-port interface <i>interface-name</i>] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure a static IP-MAC binding to be added to the DHCPv6 snooping database. |
| Options | <i>ip-address</i> —Static IPv6 address assigned to a device connected on the specified interface. <i>mac mac-address</i> —Static MAC address assigned to a device connected on the specified interface. The remaining statements are explained separately. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Static IP Addresses in the DHCP Snooping Database for Access Ports (CLI Procedure) on page 65 |

traceoptions (DHCP)

| | |
|----------------------------|---|
| Syntax | <pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; level (all error info notice verbose warning); no-remote-trace; } </pre> |
| Hierarchy Level | [edit system processes dhcp-service] |
| Release Information | <p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> |
| Description | <p>Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.</p> <p>This statement replaces the deprecated traceoptions statements at the [edit forwarding-options dhcp-relay] and [edit system services dhcp-local-server] hierarchy levels.</p> |
| Options | <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements:</p> <ul style="list-style-type: none"> • all—Trace all events. • auth—Trace authentication events. • database—Trace database events. • fwd—Trace firewall process events. • general—Trace miscellaneous events. • ha—Trace high availability-related events. • interface—Trace interface operations. • io—Trace I/O operations. • liveness-detection—Trace liveness detection operations. • packet—Trace packet and option decoding operations. • performance—Trace performance measurement operations. |

- **profile**—Trace profile operations.
- **rpd**—Trace routing protocol process events.
- **rtsock**—Trace routing socket operations.
- **security-persistence**—Trace security persistence events.
- **session-db**—Trace session database events.
- **state**—Trace changes in state.
- **statistics**—Trace baseline statistics.
- **ui**—Trace user interface operations.

level—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:

- **all**—Match messages of all levels.
- **error**—Match error messages.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.
- **warning**—Match warning messages.

Default: error

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access, allowing only the user **root** and users who have the Junos OS **maintenance** permission to access the trace files.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (***maximum-file-sizek***), megabytes (***maximum-file-sizem***), or gigabytes (***maximum-file-sizeg***). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Range: 10,240 through 1,073,741,824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

| | |
|---------------------------------|---|
| Required Privilege Level | trace—To view this statement in the configuration. trace-control—To add this statement to the configuration. |
|---------------------------------|---|

Related Documentation • [Tracing Extended DHCP Operations](#)

transmit-interval (MACsec)

| | |
|---------------------------------|---|
| Syntax | <code>transmit-interval <i>interval</i>;</code> |
| Hierarchy Level | [edit security <code>macsec connectivity-association connectivity-association-name mka</code>] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | <p>Specifies the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs).</p> <p>The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower <i>interval</i> increases bandwidth overhead on the link; a higher <i>interval</i> optimizes the MKA protocol data unit exchange process.</p> <p>The transmit interval settings must be identical on both ends of the link when MACsec using static connectivity association key (CAK) security mode is enabled.</p> <p>We recommend increasing the interval to 6000 ms in high-traffic load environments.</p> |
| Default | The default transmit interval is 2000 milliseconds. |
| Options | <i>interval</i> —Specifies the transmit interval, in milliseconds. |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | • Configuring Media Access Control Security (MACsec) on page 109 |

trusted

| | |
|---------------------------------|--|
| Syntax | trusted; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> overrides] |
| Release Information | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. |
| Description | Allow DHCP responses from the specified interface. The interface is not subject to DHCP snooping, even if the VLAN is enabled for DHCP snooping. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling a Trusted DHCP Server (CLI Procedure) on page 128• Understanding Trusted DHCP Servers for Port Security on page 127• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

untrusted

| | |
|---------------------------------|--|
| Syntax | untrusted; |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security group <i>group-name</i> overrides] |
| Release Information | Statement introduced in Junos OS Release 13.2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX series. |
| Description | Override the default behavior of a trunk interface from trusted to untrusted. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Enabling a Trusted DHCP Server (CLI Procedure) on page 128• Understanding Trusted DHCP Servers for Port Security on page 127• Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices |

use-interface-description

| | |
|---|---|
| Syntax | use-interface-description (device logical); |
| For Platforms with Enhanced Layer 2 Software (ELS) | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82 circuit-id] |
| For Platforms Without ELS | <p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 circuit-id],</p> <p>[edit forwarding-options helpers bootp dhcp-option82 circuit-id],</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 circuit-id],</p> <p>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 remote-id],</p> <p>[edit forwarding-options helpers bootp dhcp-option82 remote-id],</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 remote-id]</p> |
| For MX Series Platforms | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 circuit-id] |
| Release Information | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Hierarchy level [edit bridge-domains <i>bridge domain name</i> forwarding-options dhcp-security] introduced in Junos OS Release 14.1 for the MX Series.</p> |
| Description | <p>Use the interface description rather than the interface name (which is the default value) in the circuit ID or remote ID value in the DHCP option 82 information.</p> <p>The textual description is configured using the description statement at the [edit interfaces <i>interface-name</i>] hierarchy level. If you specify that the textual description be used and no description is configured for the interface, DHCP relay defaults to using the Layer 2 interface name. When you use the interface description rather than the interface name, the interface description has to be specified under interface unit ("set interfaces ge-0/0/0 unit 0 description "client"). If you do not do this, then the interface name is used.</p> |
| Options | <p>device—Use the device interface description. Only available for MX Series platform configuration.</p> <p>logical—Use the logical interface description. Only available for MX Series platform configuration.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server |

- *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 63*
- *Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)*
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

use-interface-index

| | |
|---------------------------------|--|
| Syntax | use-interface-index (logical device); |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options option-18], [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options option-37], |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Use the index number of the interface instead of the interface name in the DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID). These options are used by a relay agent to insert information in DHCPv6 requests before the relay agent forwards them to a DHCPv6 server. |
| Options | logical —Use the textual description that is configured for the logical interface. device —Use the textual description that is configured for the device interface. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Including a Textual Description in DHCP Options</i>• <i>Using DHCP Relay Agent Option 82 Information</i>• <i>Configuring DHCPv6 Relay Agent Options</i> |

use-interface-name




| | |
|---------------------------------|--|
| Syntax | use-interface-name (logical device); |
| Hierarchy Level | [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options option-18], [edit vlans <i>vlan-name</i> forwarding-options dhcp-security dhcpv6-options option-37], |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches. |
| Description | Configure DHCPv6 option 18 (Relay Agent Interface-ID) or option 37 (Relay Agent Remote-ID) to use the interface name to identify the port identity of the DHCP client to the DHCP server. |
| Options | logical —Use the name that is configured for the logical interface. device —Use the name that is configured for the device interface. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Including a Textual Description in DHCP Options</i>• <i>Using DHCP Relay Agent Option 82 Information</i>• <i>Configuring DHCPv6 Relay Agent Options</i> |

use-string

| | |
|---|---|
| Syntax | <code>use-string <i>string</i>;</code> |
| For Platforms with Enhanced Layer 2 Software (ELS) | <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82 remote-id]</code> |
| For Platforms Without ELS | <code>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82 remote-id],</code> <code>[edit forwarding-options helpers bootp dhcp-option82 remote-id],</code> <code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82 remote-id]</code> |
| For MX Series Platforms | <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 circuit-id]</code> |
| Release Information | Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 14.1 for the MX Series. |
| Description | Use a string rather than the MAC address of the host system (the default) in the remote ID value in the DHCP option 82 information. |
| Options | string —Character string used as the remote ID value. Range: 1–255 characters |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> <i>Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)</i> Understanding DHCP Option 82 for Protecting Switching Devices Against Attacks on page 59 <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i> <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i> <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> <i>Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)</i> |

- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 63
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

use-vlan-id

| | |
|---|--|
| Syntax | <code>use-vlan-id;</code> |
| For Platforms with Enhanced Layer 2 Software (ELS) | <p>[edit forwarding-options helpers bootp dhcp-option82-circuit-id]</p> <p>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82-circuit-id]</p> |
| For MX Series Platforms | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82 <i>circuit-id</i>] |
| Release Information | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.)</p> <p>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 14.1 for the MX Series.</p> |
| <div>  <p>NOTE: The EX Series switches that support the <code>use-vlan-id</code> statement are the EX4300, EX4600, and EX9200 switches.</p> </div> | |
| Description | Use the VLAN ID rather than the VLAN name (the default) in the circuit ID or remote ID value in the DHCP option 82 information. |
| <div>  <p>NOTE: The <code>use-vlan-id</code> statement is mutually exclusive with the <code>use-interface-description</code> and <code>no-vlan-interface-name</code> statements.</p> </div> <p>The <code>use-vlan-id</code> statement only applies to interfaces in a bridge domain. The format of the Agent Circuit ID or Agent Remote ID information for Fast Ethernet or Gigabit Ethernet interfaces is as follows:</p> <pre>(fe ge)-fpc/pic/port.subunit:svlan_id-vlan_id</pre> | |
| <div>  <p>NOTE: The <i>subunit</i> is required and used to differentiate the interface for remote systems, and <i>svlan_id-vlan_id</i> represents the VLANs associated with the bridge domain.</p> </div> | |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |

**Related
Documentation**

- *Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server*
- *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *[edit forwarding-options] Configuration Statement Hierarchy on EX Series Switches*
- RFC 3046, DHCP Relay Agent Information Option, at <http://tools.ietf.org/html/rfc3046>.

vendor-id

| | |
|---|---|
| Syntax | <code>vendor-id <string>;</code> |
| For Platforms with Enhanced Layer 2 Software (ELS) | <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security option-82]</code> |
| For Platforms Without ELS | <code>[edit ethernet-switching-options secure-access-port vlan (all <i>vlan-name</i>) dhcp-option82],</code> <code>[edit forwarding-options helpers bootp dhcp-option82],</code> <code>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</code> |
| For MX Series Platforms | <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security option-82]</code> |
| Release Information | Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Hierarchy level <code>[edit vlans <i>vlan-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 13.2X50-D10. (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Hierarchy level <code>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]</code> introduced in Junos OS Release 14.1 for the MX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server. |
| Default | If vendor-id is not explicitly configured for DHCP option 82, then no vendor ID is set. |
| Options | string —(Optional) A single string that designates the vendor ID. Range: 1–255 characters Default: If you specify vendor-id with no string value, then the default vendor ID Juniper Networks is configured. |
| Required Privilege Level | system —To view this statement in the configuration. system-control —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> <i>Configuring DHCP Option 82 to Help Protect the Switching Devices Against Attacks (CLI Procedure)</i> <i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i> <i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i> Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 63 |

vlan (RA Guard)

| | |
|---------------------------------|---|
| Syntax | <pre> vlan <i>vlan-name</i> { <i>policy</i> <i>policy-name</i>; (stateful stateless); } </pre> |
| Hierarchy Level | [edit forwarding-options access-security router-advertisement-guard] |
| Release Information | Statement introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Statement introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Configure IPv6 Router Advertisement (RA) guard on a VLAN. In an IPv6 deployment, RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. An RA guard policy is used to validate incoming RA messages on the basis of whether they match the conditions defined in a policy.</p> <p>Before you can configure RA guard on a VLAN, you must first configure a policy at the [edit forwarding-options access-security router-advertisement-guard] hierarchy level. The policy is then applied to the VLAN at the [edit forwarding-options access-security router-advertisement-guard vlan <i>vlan-name</i>] hierarchy level.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateless IPv6 Router Advertisement Guard on Switches on page 74 • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |

write-interval

| | |
|---|--|
| Syntax | <code>write-interval seconds;</code> |
| For Platforms with Enhanced Layer 2 Software (ELS) | [edit system processes dhcp-service dhcp-snooping-file], [edit system processes dhcp-service dhcpv6-snooping-file] |
| For Platforms Without ELS | [edit ethernet-switching-options secure-access-port dhcp-snooping-file], [edit ethernet-switching-options secure-access-port dhcpv6-snooping-file] |
| For MX Series Platforms | [edit system processes dhcp-service dhcp-snooping-file] |
| Release Information | <p>Statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Support at the [edit system processes dhcp-service dhcp-snooping-file] hierarchy level introduced in Junos OS Release 13.2X50-D10.</p> <p>Support at the [edit system processes dhcp-service dhcpv6-snooping-file] hierarchy level introduced in Junos OS Release 13.2X51-D20.</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p> <p>Support at the [edit ethernet-switching-options secure-access-port dhcpv6-snooping-file] hierarchy level introduced in Junos OS Release 14.1X53-D10 for EX Series switches.</p> |
| Description | <p>Specify how frequently the device writes the database entries from memory into the DHCP snooping database file.</p> <ul style="list-style-type: none"> If you are configuring write-interval at the [edit ethernet-switching-options secure-access-port dhcp-snooping-file] or the [edit ethernet-switching-options secure-access-port dhcpv6-snooping-file] hierarchy level, see <i>Making IP-MAC Bindings in the DHCP Snooping Database Persistent (CLI Procedure)</i>. If you are configuring write-interval at the [edit system processes dhcp-service dhcp-snooping-file] or the [edit system processes dhcp-service dhcpv6-snooping-file] hierarchy level, see "Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)" on page 66. |
| Options | <p>seconds—Value in seconds.</p> <p>Range: 60 through 86,400 seconds.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Understanding DHCP Snooping for Monitoring DHCP Messages Received from Untrusted Devices</i> |

CHAPTER 10

Operational Commands

- clear access-security router-advertisement statistics
- clear arp
- clear dhcp-security binding
- clear dhcp-security ipv6 binding
- clear dot1x
- clear ethernet-switching recovery-timeout
- clear security mka statistics
- request access-security router-advertisement-guard-forward
- request access-security router-advertisement-guard-block
- request access-security router-advertisement-guard-learn interface
- show access-security router-advertisement statistics
- show access-security router-advertisement state
- show dhcp-security arp inspection statistics
- show dhcp-security binding
- show dhcp-security binding ip-source-guard
- show dhcp-security ipv6 binding
- show dhcp-security ipv6 statistics
- show dhcp-security neighbor-discovery-inspection statistics
- show security macsec connections
- show security macsec statistics
- show security mka sessions
- show security mka statistics

clear access-security router-advertisement statistics

| | |
|---------------------------------|---|
| Syntax | clear access-security router-advertisement statistics (fail success) (all interface <i>interface-name</i> vlan <i>vlan-name</i>) |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | Clear the IPv6 Router Advertisement (RA) guard entries for received RA messages. If RA guard is enabled on a switch, the switch examines incoming RA messages and filters them on the basis of a predefined set of criteria. If the switch validates the sender of the RA message as a legitimate IPv6 router, it forwards the RA message to its destination; otherwise, the RA message is dropped. |
| Options | all —Clear the RA guard entries on all VLANs. fail —Clear RA guard entries for RA messages that were discarded. interface <i>interface-name</i> —Clear the RA guard entries for the specified interface. success —Clear the RA guard entries for RA messages that were accepted. vlan <i>vlan-name</i> —Clear the RA guard entries for the specified VLAN. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show access-security router-advertisement statistics on page 244 |
| Output Fields | This command generates no output. |

clear arp

| | |
|---------------------------------|---|
| Syntax | <pre>clear arp <all> <hostname <i>hostname</i>> <interface <i>interface-name</i>> <logical-system <i>logical-system-name</i>> <vpn <i>vpn</i>></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 14.1 for the MX Series.</p> <p>all option introduced in Junos OS Release 14.2.</p> |
| Description | <p>Remove entries from the Address Resolution Protocol (ARP) table for the current CLI view. To clear entries for a specific logical system, you must first enter the set cli logical-system <i>logical-system-name</i> command, and then issue the clear arp command.</p> |
| Options | <p>none all— (Optional) Clear all entries from the ARP table.</p> <p>Both clear arp and clear arp all function identically.</p> <p>hostname <i>hostname</i>—(Optional) Clear only the specified host entry from the ARP table.</p> <p>interface <i>interface-name</i>—(Optional) Clear entries only for the specified interface from the ARP table.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Clear entries for only the specified logical system from the ARP table (only available in main router context).</p> <p>vpn <i>vpn</i>—(Optional) Clear entries from the ARP table for the specified virtual private network (VPN).</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • set cli logical-system • show arp • show dhcp-security arp inspection statistics on page 247 • Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19 |
| List of Sample Output | <p>clear arp on page 234</p> <p>clear arp all on page 234</p> <p>clear arp logical-system ls1 on page 234</p> |
| Output Fields | When you enter this command, you are provided feedback on the status of your request. |

Sample Output

clear arp

```
user@host> clear arp
192.168.71.254  deleted
192.168.65.46   deleted
192.168.64.10   deleted
10.0.12.14      deleted
10.0.17.14      deleted
```

clear arp all

```
user@host> clear arp all
192.168.71.254  deleted
192.168.65.46   deleted
192.168.64.10   deleted
10.0.12.14      deleted
10.0.17.14      deleted
```

clear arp logical-system ls1

```
user@host> clear arp logical-system ls1
192.168.71.254  deleted
192.168.65.46   deleted
192.168.64.10   deleted
10.0.12.14      deleted
10.0.17.14      deleted
```

clear dhcp-security binding

| | |
|---------------------------------|--|
| Syntax | clear dhcp-security binding <interface <i>interface-name</i> > <ip-address <i>ip-address</i> > <statistics> <vlan <i>vlan-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Command introduced in Junos OS Release 14.1 for the MX Series. |
| Description | Clear the DHCP snooping database information. |
| Options | <p>interface <i>interface-name</i>—(Optional) Clear DHCP snooping database information for the specified interface.</p> <p>ip-address <i>ip-address</i>—(Optional) Clear DHCP snooping database information for the specified IP address.</p> <p>statistics—(Optional) Clear all DHCP snooping database statistics.</p> <p>vlan <i>vlan-name</i>—(Optional) Clear DHCP snooping database information for the specified VLAN.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security binding on page 249 • Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33 • Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks • Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19 |

clear dhcp-security ipv6 binding

| | |
|---------------------------------|---|
| Syntax | clear dhcp-security ipv6 binding <all> <interface <i>interface-name</i> > <ipv6-address <i>ipv6-address</i> > <vlan <i>vlan-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Clear the DHCPv6 snooping database information. |
| Options | <p>all—(Optional) Clear all DHCPv6 snooping database statistics.</p> <p>interface <i>interface-name</i>—(Optional) Clear DHCPv6 snooping database information for the specified interface.</p> <p>ipv6-address <i>ipv6-address</i>—(Optional) Clear DHCPv6 snooping database information for the specified IPv6 address.</p> <p>vlan <i>vlan-name</i>—(Optional) Clear DHCPv6 snooping database information for the specified VLAN.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show dhcp-security ipv6 binding on page 254• Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38 |
| List of Sample Output | clear dhcp-security ipv6 binding on page 236 |
| Output Fields | This command produces no output. |

Sample Output

clear dhcp-security ipv6 binding

```
user@switch> clear dhcp-security ipv6 binding
```

clear dot1x

Syntax `clear dot1x (firewall <counter-name> | interface <[interface-name]> | mac-address [mac-addresses] | statistics <interface interface-name>)`

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
firewall option added in Junos OS Release 9.5 for EX Series switches.
 Command introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description Reset the authentication state of an interface or delete 802.1X statistics from the switch. When you reset an interface using the **interface** or **mac-address** options, reauthentication on the interface is also triggered. The switch sends out a multicast message on the interface to restart the authentication of all connected supplicants. If a MAC address is reset, then the switch sends out a unicast message to that specific MAC address to restart authentication.

If a supplicant is sending traffic when the **clear dot1x interface** command is issued, the authenticator immediately initiates reauthentication. This process happens quickly, and it might seem that reauthentication did not occur. To verify that reauthentication has happened, issue the **show dot1x interface detail** command. The values for **Reauthentication due** and **Reauthentication interval** will be about the same.



CAUTION: When you clear the learned MAC addresses from an interface using the **clear dot1x interface** command, all MAC addresses are cleared, including those in static MAC bypass list.

If you have enabled Media Access Control Security (MACsec) using static secure association key (SAK) security mode on an EX Series switch, the SAKs are rotated when the **clear dot1x** command is entered. The **clear dot1x** command has no impact on MACsec when MACsec is enabled using static connectivity association keys (CAK) or any other security mode.

Options **firewall <counter-name>**—Clear 802.1X firewall counter statistics. If the *counter-name* option is specified, clear 802.1X firewall statistics for that counter.

interface <[interface-name]>—Reset the authentication state of all the supplicants (also, clears all the authentication bypassed clients) connected to the specified interface (when the interface is an authenticator) or reset the authentication state for the interface itself (when the interface is a supplicant).

mac-address [mac-addresses]—Reset the authentication state of the specified MAC addresses.

statistics <interface interface-name>—Clear 802.1X statistics on all 802.1X-enabled interfaces. If the **interface** option is specified, clear 802.1X firewall statistics for that interface or interfaces.

Required Privilege Level view

Related Documentation

- *show dot1x*
- *Example: Setting Up 802.1X for Single-Suppliant or Multiple-Suppliant Configurations on an EX Series Switch*
- *Filtering 802.1X Suplicants by Using RADIUS Server Attributes*

List of Sample Output

- [clear dot1x firewall on page 238](#)
- [clear dot1x interface \(Specific Interfaces\) on page 238](#)
- [clear dot1x mac-address \(Specific MAC Address\) on page 238](#)
- [clear dot1x statistics interface \(Specific Interface\) on page 238](#)

Sample Output

clear dot1x firewall

```
user@switch> clear dot1x firewall c1
```

clear dot1x interface (Specific Interfaces)

```
user@switch> clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0
```

clear dot1x mac-address (Specific MAC Address)

```
user@switch> clear dot1x mac-address 00:04:ae:cd:23:5f
```

clear dot1x statistics interface (Specific Interface)

```
user@switch> clear dot1x statistics interface ge-1/0/1
```

clear ethernet-switching recovery-timeout

| | |
|---------------------------------|---|
| Syntax | clear ethernet-switching recovery-timeout |
| Release Information | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches. |
| Description | Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch, and restore the interfaces to service. |
| Options | interface <i>interface-name</i> vlan <i>vlan-name</i> —(EX9200 switches) Unblock an interface on the basis of its membership in the specified VLAN. This option can be used to restore an interface that is blocked because of a vlan-member-shutdown action. |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• Configuring Autorecovery from the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 96 |
| Output Fields | This command produces no output. |

clear security mka statistics

| | |
|---------------------------------|--|
| Syntax | clear security mka statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Command introduced in Junos OS Release 14.1X53-D15 for QFX Series switches. |
| Description | <p>Clear—reset to zero (0)—all MACsec Key Agreement (MKA) protocol statistics.</p> <p>You are clearing the statistics that are viewed using the show security mka statistics when you enter this command.</p> |
| Options | <p>none—Clear all MKA counters for all interfaces on the switch.</p> <p>interface <i>interface-name</i>—(Optional) Clear MKA traffic counters for the specified interface only.</p> |
| Required Privilege Level | clear |
| Related Documentation | <ul style="list-style-type: none">• show security mka statistics on page 269• show security mka sessions on page 267• Understanding Media Access Control Security (MACsec) on page 101 |

Sample Output

clear security mka statistics

```
user@switch> clear security mka statistics
```


request access-security router-advertisement-guard-forward

| | |
|---------------------------------|---|
| Syntax | request access-security router-advertisement-guard-forward interface (<i>interface-name</i>) |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Initiate the forwarding state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages sent from legitimate sources will dynamically transition to the forwarding state. While the interface is in forwarding state, all RA messages received on that interface that can be validated against the configured policy are forwarded.</p> <p>You can override the dynamic state transitions by requesting the forwarding state on an interface. If you issue the request for the forwarding state on an interface, the interface will remain in forwarding state until either the learning or blocking state is requested on that interface.</p> |
| Options | interface <i>interface-name</i> —Initiate the forwarding state on the specified interface. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |
| Output Fields | This command produces no output. |

request access-security router-advertisement-guard-block

| | |
|---------------------------------|--|
| Syntax | request access-security router-advertisement-guard-block interface (<i>interface-name</i>) |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Initiate the blocking state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages that are not sent from valid IPv6 routers will dynamically transition to the blocking state. While the interface is in blocking state, all RA messages received on that interface are dropped.</p> <p>You can override the dynamic state transitions by requesting the blocking state on an interface. If you issue the request for the blocking state on an interface, the interface will remain in forwarding state until either the learning or forwarding state is requested on that interface.</p> |
| Options | interface <i>interface-name</i> —Initiate the blocking state on the specified interface. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |
| Output Fields | This command produces no output. |

request access-security router-advertisement-guard-learn interface

| | |
|---------------------------------|---|
| Syntax | request access-security router-advertisement-guard-learn interface (<i>interface-name</i>) duration <i>seconds</i> (forward block) |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | <p>Request the learning state on an interface or range of interfaces for stateful IPv6 Router Advertisement (RA) guard. Stateful RA guard learns about legitimate senders of RA messages and stores this information in order to validate senders of subsequent RA messages. For example, an interface that is in the learning state and receives RA messages sent from legitimate sources dynamically transitions to the forwarding state after the learning period ends. While the interface is in forwarding state, all RA messages received on that interface that can be validated against the configured policy are forwarded.</p> <p>Before you can request learning on an interface, you must enable RA guard at the [edit forwarding-options access-security router-advertisement-guard] hierarchy level and configure the stateful option. When you enable stateful RA guard, the default state is Off. An interface in the Off state operates as if RA guard is not available. The learning state can be initiated only by configuring the request access-security router-advertisement-guard-learn command.</p> <p>When you request the learning state, you must configure the duration of the learning period in seconds. This is the amount of time the interface will remain in the learning state before it transitions to another state. RA messages that are received during the learning period can be either forwarded or blocked. Configure the forward option to forward RA messages during the learning period, or configure the block option to block RA messages during the learning period.</p> |
| Options | <p>interface <i>interface-name</i>—Initiate the learning state on the specified interface.</p> <p>duration <i>seconds</i>—Configure the duration of the learning state in seconds. When the learning period ends, the state dynamically transitions to either the forwarding state or the blocking state.</p> <p>forward—Configure the interface to forward RA messages received during the learning period.</p> <p>block—Configure the interface to block RA messages received during the learning period.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Stateful IPv6 Router Advertisement Guard on Switches on page 79 |
| Output Fields | This command produces no output. |

show access-security router-advertisement statistics

| | |
|---------------------------------|--|
| Syntax | show access-security router-advertisement statistics <interface <i>interface-name</i>> |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | Display the IPv6 Router Advertisement (RA) guard entries for received RA messages. RA guard enables a switch to examine incoming RA messages and filter them on the basis of predefined set of criteria. Once the switch has validated that the sender of the RA message is a legitimate IPv6 router, it forwards the RA message to its destination; otherwise, the RA message is dropped. |
| Options | interface <i>interface-name</i> —(Optional) Display the RA guard entries for the specified interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> clear access-security router-advertisement statistics on page 232 |
| List of Sample Output | show access-security router-advertisement statistics on page 244 |
| Output Fields | Table 11 on page 244 lists the output fields for the show access-security router-advertisement statistics command. Output fields are listed in the approximate order in which they appear. |

Table 11: show access-security router-advertisement statistics Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| Interface | Interface on which the RA packet was received. | All levels |
| RA Packets | Total number of RA packets that were received. | All levels |
| RA inspection pass | Total number of RA packets that passed RA guard inspection. | All levels |
| RA inspection fail | Total number of RA packets that failed RA guard inspection. | All levels |

Sample Output

show access-security router-advertisement statistics

```

user@device> show access-security router-advertisement statistics
Interface    RA Packets received  RA inspection pass  RA inspection fail
ge-0/0/7.0   3                    2                    1
ge-0/0/15.0  8                    5                    3

```

show access-security router-advertisement state

| | |
|---------------------------------|--|
| Syntax | show access-security router-advertisement state <interface <i>interface-name</i>> |
| Release Information | Command introduced in Junos OS Release 15.1X53-D55 for EX3400 switches. Command introduced in Junos OS Release 16.1 for EX Series switches. |
| Description | Display the IPv6 Router Advertisement (RA) guard state information. Stateful RA guard enables the switch to learn about the sources of RA messages for a certain period of time. When the learning period ends, the switch has a record of which interfaces are attached to links with valid IPv6 routers. If there is no valid IPv6 router attached to the interface, the switch dynamically transitions the interface from the learning state into the blocking state. Subsequent RA messages received after the transition to blocking state are dropped. If there is a valid IPv6 router attached to the interface, the interface transitions to the forwarding state, and subsequent RA messages that can be validated against the configured policy are forwarded. |
| Options | interface <i>interface-name</i> —(Optional) Display the RA guard entries for the specified interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show access-security router-advertisement statistics on page 244 |
| List of Sample Output | show access-security router-advertisement state on page 246 |
| Output Fields | Table 11 on page 244 lists the output fields for the show access-security router-advertisement state command. Output fields are listed in the approximate order in which they appear. |

Table 12: show access-security router-advertisement state Output Fields

| Field Name | Field Description |
|------------------|--|
| Interface | Displays the interface on which stateful IPv6 RA guard is enabled. |
| State | Displays one of the following states: <ul style="list-style-type: none"> • OFF—The interface operates as if RA guard is not available. • BLOCKED—The interface blocks ingress RA messages. • FORWARDING—The interface forwards ingress RA messages that can be validated against the configured policy. • LEARNING—The switch is actively acquiring information about the IPv6 routing device connected to the interface. • TRUSTED—The interface forwards all ingress RA messages without performing policy checks. |

Sample Output

show access-security router-advertisement state

```
user@device> show access-security router-advertisement state
Interface      state
ge-0/0/0.0     LEARNING
ge-1/0/0.0     FORWARDING
ge-1/0/0.0     BLOCKED
```

show dhcp-security arp inspection statistics

| | |
|---------------------------------|--|
| Syntax | show dhcp-security arp inspection statistics |
| Release Information | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Command introduced in Junos OS Release 14.1 for the MX Series. |
| Description | Display Address Resolution Protocol (ARP) inspection statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security binding on page 249 • clear dhcp-security binding on page 235 • clear interfaces statistics • Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33 • Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks • Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19 |
| List of Sample Output | show dhcp-security arp inspection statistics on page 248 |
| Output Fields | <p>Table 13 on page 247 lists the output fields for the show dhcp-security arp inspection statistics command. Output fields are listed in the approximate order in which they appear.</p> <p>The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another.</p> |

Table 13: show dhcp-security arp inspection statistics Output Fields

| Field Name | Field Description | Level of Output |
|----------------------------|--|-----------------|
| Interface | Interface on which ARP inspection has been applied. | All levels |
| Packets received | Total number of packets that underwent ARP inspection. | All levels |
| ARP inspection pass | Total number of packets that passed ARP inspection. | All levels |
| ARP inspection fail | Total number of packets that failed ARP inspection. | All levels |

Sample Output

show dhcp-security arp inspection statistics

```
user@device> show dhcp-security arp inspection statistics
```

| Interface | Packets received | ARP inspection pass | ARP inspection fail |
|-------------|------------------|---------------------|---------------------|
| ge-0/0/30.0 | 7 | 7 | 0 |
| ge-0/0/4.0 | 3 | 3 | 0 |
| ge-0/0/6.0 | 72 | 4 | 68 |

show dhcp-security binding

| | |
|---------------------------------|---|
| Syntax | <pre>show dhcp-security binding <interface <i>interface-name</i>> <ip-address <i>ip-address</i>> <ip-source-guard <i>ip-sg-name</i>> <statistics> <vlan <i>vlan-name</i>></pre> |
| Release Information | <p>Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1 for the MX Series.</p> |
| Description | Display the DHCP snooping database information. |
| Options | <p>interface <i>interface-name</i>—(Optional) Display the DHCP snooping database information for an interface.</p> <p>ip-address <i>ip-address</i>—(Optional) Display the DHCP snooping database information for an IP address.</p> <p>vlan <i>vlan-name</i>—(Optional) Display the DHCP snooping database information for a VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security binding ip-source-guard on page 252 • clear dhcp-security binding on page 235 • Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33 • Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks • Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19 |
| List of Sample Output | <p>show dhcp-security binding on page 250</p> <p>show dhcp-security binding interface on page 250</p> <p>show dhcp-security binding ip-address on page 250</p> <p>show dhcp-security binding vlan on page 251</p> |
| Output Fields | <p>Table 14 on page 249 lists the output fields for the show dhcp-security binding command. Output fields are listed in the approximate order in which they appear.</p> |

Table 14: show dhcp-security binding Output Fields

| Field Name | Field Description | Level of Output |
|-------------|---|-----------------|
| IP Address | IP address of the network device; bound to the MAC address. | All levels |
| MAC address | MAC address of the network device; bound to the IP address. | All levels |

Table 14: show dhcp-security binding Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| VLAN | VLAN name of the network device whose MAC address is shown. | All levels |
| Expires | The time, in seconds, remaining before the lease of the IP address to the MAC address expires. This field is 0 for static entries. | All levels |
| State | Specifies whether the IP address is: <ul style="list-style-type: none"> • BOUND: Leased to the MAC address for a limited period of time. • STATIC: Attached to a fixed MAC address. | All levels |
| Interface | Interface address (port). | All levels |

Sample Output

show dhcp-security binding

```
user@device> show dhcp-security binding
```

| IP address | MAC address | Vlan | Expires | State | Interface |
|------------|-------------------|--------|---------|--------|------------|
| 10.1.1.10 | 00:10:00:20:00:01 | vlan20 | 0 | STATIC | ge-0/0/4.0 |
| 10.1.1.18 | 00:10:94:00:00:34 | vlan20 | 86287 | BOUND | ge-0/0/6.0 |
| 10.1.1.15 | 00:10:94:00:00:55 | vlan20 | 86265 | BOUND | ge-0/0/4.0 |
| 10.1.1.16 | 00:10:94:00:00:56 | vlan20 | 86265 | BOUND | ge-0/0/4.0 |
| 10.1.1.19 | 00:10:94:00:00:5b | vlan20 | 86287 | BOUND | ge-0/0/6.0 |
| 10.1.1.20 | 00:10:94:00:00:5c | vlan20 | 86287 | BOUND | ge-0/0/6.0 |
| 10.1.1.21 | 00:10:94:00:00:5d | vlan20 | 86287 | BOUND | ge-0/0/6.0 |
| 10.1.1.17 | 00:10:94:00:00:68 | vlan20 | 86265 | BOUND | ge-0/0/4.0 |

show dhcp-security binding interface

```
user@device> show dhcp-security binding interface ge-0/0/6
```

| IP address | MAC address | Vlan | Expires | State | Interface |
|------------|-------------------|--------|---------|-------|------------|
| 10.1.1.18 | 00:10:94:00:00:34 | vlan20 | 86282 | BOUND | ge-0/0/6.0 |
| 10.1.1.19 | 00:10:94:00:00:5b | vlan20 | 86282 | BOUND | ge-0/0/6.0 |
| 10.1.1.20 | 00:10:94:00:00:5c | vlan20 | 86282 | BOUND | ge-0/0/6.0 |
| 10.1.1.21 | 00:10:94:00:00:5d | vlan20 | 86282 | BOUND | ge-0/0/6.0 |

show dhcp-security binding ip-address

```
user@device> show dhcp-security binding ip-address
```

| IP address | MAC address | Vlan | Expires | State | Interface |
|------------|-------------------|--------|---------|-------|------------|
| 10.1.1.18 | 00:10:94:00:00:34 | vlan20 | 86282 | BOUND | ge-0/0/6.0 |

show dhcp-security binding vlan

```
user@device> show dhcp-security binding vlan vlan20
```

| IIP address | MAC address | Vlan | Expires | State | Interface |
|-------------|-------------------|--------|---------|-------|------------|
| 10.1.1.18 | 00:10:94:00:00:34 | vlan20 | 86282 | BOUND | ge-0/0/6.0 |

show dhcp-security binding ip-source-guard

| | |
|---------------------------------|---|
| Syntax | show dhcp-security binding ip-source-guard |
| Release Information | Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Command introduced in Junos OS Release 14.1 for the MX Series. |
| Description | Display IP source guard database table. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security binding on page 249 • clear dhcp-security binding on page 235 • Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing on page 33 • Example: Configuring IP Source Guard and Dynamic ARP Inspection on a Specified Bridge Domain to Protect the Devices Against Attacks • Understanding Port Security Features to Protect the Access Ports on Your Device Against the Loss of Information and Productivity on page 19 |
| List of Sample Output | show dhcp-security binding ip-source-guard on page 253 |
| Output Fields | <p>Table 15 on page 252 lists the output fields for the show dhcp-security binding ip-source-guard command. Output fields are listed in the approximate order in which they appear.</p> <p>The IP source guard database table shows the untrusted access interfaces in VLANs that have been enabled for IP source guard. The entries include the IP addresses and MAC addresses that are bound to one another.</p> |

Table 15: show dhcp-security binding ip-source-guard Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|---|-----------------|
| IP Address | IP address of the network device; bound to the MAC address. | All levels |
| MAC address | MAC address of the network device; bound to the IP address. | All levels |
| VLAN | VLAN name of the network device whose MAC address is shown. | All levels |
| Expires | The time, in seconds, remaining before the lease of the IP address to the MAC address expires. | All levels |
| State | Specifies whether the IP address is: <ul style="list-style-type: none"> • BOUND: Temporarily leased to the MAC address for a limited period of time. • STATIC: Attached to a fixed MAC address. | All levels |

Table 15: show dhcp-security binding ip-source-guard Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------|---------------------------|-----------------|
| Interface | Interface address (port). | All levels |

Sample Output

show dhcp-security binding ip-source-guard

```
user@device> show dhcp-security binding ip-source-guard
```

| IP address | MAC address | Vlan | Expires | State | Interface |
|------------|-------------------|--------|---------|--------|------------|
| 10.1.1.10 | 00:10:00:20:00:01 | vlan20 | 0 | STATIC | ge-0/0/4.0 |
| 10.1.1.18 | 00:10:94:00:00:34 | vlan20 | 86276 | BOUND | ge-0/0/6.0 |
| 10.1.1.15 | 00:10:94:00:00:55 | vlan20 | 86254 | BOUND | ge-0/0/4.0 |
| 10.1.1.16 | 00:10:94:00:00:56 | vlan20 | 86254 | BOUND | ge-0/0/4.0 |
| 10.1.1.19 | 00:10:94:00:00:5b | vlan20 | 86276 | BOUND | ge-0/0/6.0 |
| 10.1.1.20 | 00:10:94:00:00:5c | vlan20 | 86276 | BOUND | ge-0/0/6.0 |
| 10.1.1.21 | 00:10:94:00:00:5d | vlan20 | 86276 | BOUND | ge-0/0/6.0 |
| 10.1.1.17 | 00:10:94:00:00:68 | vlan20 | 86254 | BOUND | ge-0/0/4.0 |

show dhcp-security ipv6 binding

| | |
|---------------------------------|--|
| Syntax | show dhcp-security ipv6 binding <interface <i>interface-name</i> > <ipv6-address <i>ipv6-address</i> > <vlan <i>vlan-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Display bindings between IPv6 addresses and MAC addresses (IP-MAC bindings) along with other DHCP lease information, also known as the DHCPv6 binding table or DHCPv6 snooping database. |
| Options | <p>interface <i>interface-name</i>—(Optional) Display the DHCPv6 snooping table for the specified interface.</p> <p>ipv6-address <i>ipv6-address</i>—(Optional) Display the DHCPv6 snooping table for the specified IPv6 address.</p> <p>vlan <i>vlan-name</i>—(Optional) Display the DHCPv6 snooping table for a VLAN.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security ipv6 statistics on page 256 • clear dhcp-security ipv6 binding on page 236 • Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38 |
| List of Sample Output | show dhcp-security ipv6 binding on page 255 show dhcp-security ipv6 binding interface on page 255 |
| Output Fields | <p>Table 15 on page 252 lists the output fields for the show dhcp-security ipv6 binding command. Output fields are listed in the approximate order in which they appear.</p> <p>The DHCPv6 binding table shows the untrusted access interfaces in VLANs that have been enabled for DHCPv6 snooping. The entries include the IPv6 addresses and MAC addresses that are bound to one another.</p> |

Table 16: show dhcp-security ipv6 binding Output Fields

| Field Name | Field Description | Level of Output |
|--------------|---|-----------------|
| IPv6 address | IPv6 addresses of the network device; bound to the MAC address. There are two entries added for each client: one with the link-local IPv6 address, which is used by the client for DHCP transactions, and another with the IPv6 address assigned by the server. The link-local address always has the prefix fe80::/10. | All levels |
| MAC address | MAC address of the network device; bound to the IPv6 address. | All levels |

Table 16: show dhcp-security ipv6 binding Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------|---|-----------------|
| VLAN | VLAN name of the network device whose MAC address is shown. | All levels |
| Expires | The time, in seconds, remaining before the lease of the IPv6 address to the MAC address expires. This field is 0 for static entries. | All levels |
| State | Specifies whether the IPv6 address is: <ul style="list-style-type: none"> • BOUND: Temporarily leased to the MAC address for a limited period of time. • STATIC: Attached to a fixed MAC address. | All levels |
| Interface | Interface address (port). | All levels |

Sample Output

show dhcp-security ipv6 binding

```

user@switch> show dhcp-security ipv6 binding
IPv6 address      MAC address      Vlan    Expires  State  Interface
2001:db8:fe10::   00:10:94:00:55:0b v1an20   3456    BOUND  ge-0/0/1.0

fe80::210:94ff:fe00:1  00:10:94:00:55:0b v1an20   3456    BOUND  ge-0/0/1.0

2001:db8:fe12::   00:10:94:00:00:34 v1an20   3456    BOUND  ge-0/0/2.0

fe80::210:94ff:fe00:2  00:10:94:00:00:34 v1an20   3456    BOUND  ge-0/0/2.0

2001:db8:fe14::   00:10:94:00:00:55 v1an20   3456    BOUND  ge-0/0/3.0

fe80::210:94ff:fe00:3  00:10:94:00:00:55 v1an20   3456    BOUND  ge-0/0/3.0

```

Sample Output

show dhcp-security ipv6 binding interface

```

user@switch> show dhcp-security ipv6 binding interface ge-0/0/4.0
IPv6 address      MAC address      Vlan    Expires  State  Interface
2001:db8:fe16::   00:10:00:20:00:01 v1an20    0      STATIC  ge-0/0/4.0

fe80::210:94ff:fe00:4  00:10:00:20:00:01 v1an20    0      STATIC  ge-0/0/4.0

```

show dhcp-security ipv6 statistics

| | |
|---------------------------------|---|
| Syntax | show dhcp-security ipv6 statistics |
| Release Information | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Display DHCPv6 statistics. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• show dhcp-security ipv6 binding on page 254• show dhcp-security neighbor-discovery-inspection statistics on page 259 |
| List of Sample Output | show dhcp-security ipv6 statistics on page 258 |
| Output Fields | Table 17 on page 257 lists the output fields for the show dhcp-security ipv6 statistics command. Output fields are listed in the approximate order in which they appear. |

Table 17: show dhcp-security ipv6 statistics Output Fields

| Field Name | Field Description |
|-----------------|---|
| DHCPv6 messages | <p>Number of DHCPv6 messages exchanged.</p> <ul style="list-style-type: none"> • Total—Total number of DHCPv6 messages exchanged. • Solicit—Number of DHCPv6 messages of type Solicit. A client sends a Solicit message to locate servers. • Advertise—Number of DHCPv6 messages of type Advertise. A server sends an Advertise message, in response to a Solicit message, to indicate that it is available for DHCPv6 service. • Request—Number of DHCPv6 messages of type Request. A client sends a Request message to request configuration parameters from a server. • Reply—Number of DHCPv6 messages of type Reply. A server sends a Reply message in response to a Solicit, Request, Renew, Rebind, Confirm, Information Request, Release, or Decline message. • Confirm—Number of DHCPv6 messages of type Confirm. A client sends a Confirm message to any available server to determine whether the addresses it was assigned are still appropriate for the link to which the client is connected. • Decline—Number of DHCPv6 messages of type Decline. A client sends a Decline message to a server to indicate that one or more of the addresses assigned by the server are already in use on the link to which the client is connected. • Release—Number of DHCPv6 messages of type Release. A client sends a Release message to the server to indicate that the client will no longer use one or more of the assigned addresses. • Renew—Number of DHCPv6 messages of type Renew. A client sends a Renew message to the server to extend the lifetimes on the addresses assigned to the client by that server and to update other configuration parameters received by that server. • Rebind—Number of DHCPv6 messages of type Rebind. A client sends a Rebind message to any available server after receiving no reply to a Renew message. • Relay-forward—Number of DHCPv6 messages of type Relay-forward. A relay agent sends a Relay-forward message to relay messages to servers, either directly or through another relay agent. The received message is encapsulated in an option in the Relay-forward message. • Relay-reply—Number of DHCPv6 messages of type Relay-reply. A server sends a Relay-reply message to a relay agent containing a message that the relay agent delivers to a client. • Information-request—Number of DHCPv6 messages of type Information-request. A client sends an Information-request message to a server to request configuration parameters without the assignment of any IP addresses to the client. • Reconfigure—Number of DHCPv6 messages of type Reconfigure. A server sends a Reconfigure message to a client to inform the client that the server has new or updated configuration parameters, and that the client needs to initiate a Renew/Reply or Information-request/Reply transaction with the server in order to receive the updated information. |
| Packets dropped | <p>Number of packets not considered for DHCPv6 snooping because of errors.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by DHCPv6 snooping. • No configuration—Number of packets discarded because they did not have a valid configuration. • No VLAN—Number of packets discarded because they did not belong to a valid VLAN. • No interface—Number of packets discarded because they did not belong to a valid interface. • Request on trusted port—Number of packets discarded because a Request message was received on a trusted port. |

Sample Output

show dhcp-security ipv6 statistics

```
user@host> show dhcp-security ipv6 statistics
DHCPv6 messages:
  Total                32
  Solicit              1
  Advertise            1
  Request              3
  Reply                5
  Confirm              1
  Decline              2
  Release              9
  Renew                4
  Rebind               2
  Relay forward        1
  Relay reply          1
  Information request  1
  Reconfigure          2

Packets dropped:
  Total                0
  No configuration     0
  No VLAN              0
  No interface         0
  Request on trusted port 0
```

show dhcp-security neighbor-discovery-inspection statistics

| | |
|---------------------------------|---|
| Syntax | show dhcp-security neighbor-discovery-inspection statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X51-D20 for EX Series switches. |
| Description | Display IPv6 neighbor discovery inspection statistics to determine whether there is IPv6 address spoofing on the network. |
| Options | interface <i>interface-name</i> —(Optional) Display neighbor discovery inspection statistics for the specified interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show dhcp-security ipv6 binding on page 254 • Enabling IPv6 Neighbor Discovery Inspection on page 71 • Example: Configuring IPv6 Source Guard and Neighbor Discovery Inspection to Protect a Switch from IPv6 Address Spoofing on page 38 |
| List of Sample Output | show dhcp-security neighbor-discovery-inspection statistics on page 259 show dhcp-security neighbor-discovery-inspection statistics interface on page 260 |
| Output Fields | Table 13 on page 247 lists the output fields for the show dhcp-security neighbor-discovery-inspection statistics command. Output fields are listed in the approximate order in which they appear. |

Table 18: show dhcp-security neighbor-discovery-inspection statistics Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| Interface | Interface on which neighbor discovery inspection has been applied. | All levels |
| Packets received | Total number of packets that underwent neighbor discovery inspection. | All levels |
| ND inspection pass | Total number of packets that passed neighbor discovery inspection. | All levels |
| ND inspection fail | Total number of packets that failed neighbor discovery inspection. | All levels |

Sample Output

show dhcp-security neighbor-discovery-inspection statistics

```

user@switch> show dhcp-security neighbor-discovery-inspection statistics
Interface      ND Packets received  ND inspection pass  ND inspection failed
ge-0/0/1.0      7                    5                    2
ge-0/0/2.0     10                   10                   0
ge-0/0/3.0     12                   12                   0

```

Sample Output

show dhcp-security neighbor-discovery-inspection statistics interface

```
user@switch> show dhcp-security neighbor-discovery-inspection statistics interface ge-0/0/1.0
Interface      ND Packets received  ND inspection pass  ND inspection failed
ge-0/0/1.0          7                   5                   2
```

show security macsec connections

| | |
|---------------------------------|--|
| Syntax | show security macsec connections <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Display the status of the active MACsec connections on the switch. This command does not display output when MACsec is enabled using static secure association key (SAK) security mode. |
| Options | none —Display MACsec connection information for all interfaces on the switch. interface <i>interface-name</i> —(Optional) Display MACsec connection information for the specified interface only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • show security macsec statistics on page 263 |
| List of Sample Output | show security macsec connections on page 262 |
| Output Fields | Table 19 on page 261 lists the output fields for the show security macsec connections command. Output fields are listed in the approximate order in which they appear. |

Table 19: show security macsec connections Output Fields

| Field Name | Field Description |
|-----------------------------|--|
| Fields for Interface | |
| Interface name | Name of the interface. |
| CA name | <p>Name of the connectivity association.</p> <p>A connectivity association is named using the connectivity-association statement when you are enabling MACsec.</p> |
| Cipher suite | Name of the cipher suite used for encryption. |
| Encryption | <p>Encryption setting. Encryption is enabled when this output is on and disabled when this output is off.</p> <p>The encryption setting is set using the no-encryption statement in the connectivity association when using static connectivity association key (CAK) security mode and is set using the encryption statement in the secure channel when using static secure association key (SAK) or dynamic security mode.</p> |

Table 19: show security macsec connections Output Fields (*continued*)

| Field Name | Field Description |
|-------------------|--|
| Key server offset | <p>Offset setting.</p> <p>The offset is set using the offset statement when configuring the connectivity association when using static connectivity association key (CAK) or dynamic security mode or the secure channel when using static secure association key (SAK) security mode.</p> |
| Include SCI | <p>SCI tagging. The SCI tag is included on packets in a secure channel when this output is yes, and not included on packets in a secure channel when this output is no.</p> <p>You can enable SCI tagging using the include-sci statement in the connectivity association.</p> <p>NOTE: SCI tags are automatically appended to packets leaving a MACsec-enabled interface on an EX4300 switch. The include-sci option is, therefore, not available on EX4300 switches. The output for the Include SCI field is yes.</p> |
| Replay protect | <p>Replay protection setting. Replay protection is enabled when this output is on and disabled when this output is off.</p> <p>You can enable replay protection using the replay-protect statement in the connectivity association.</p> |
| Replay window | <p>Replay protection window setting. This output is set to 0 when replay protection is disabled, and is the size of the replay window, in number of packets, when replay protection is enabled.</p> <p>The size of the replay window is configured using the replay-window-size statement in the connectivity association.</p> |

Sample Output

show security macsec connections

```

user@host> show security macsec connections
Interface name: xe-0/1/0
  CA name: CA1
  Cipher suite: GCM-AES-128   Encryption: on
  Key server offset: 0        Include SCI: no
  Replay protect: off         Replay window: 0

```

show security macsec statistics

Syntax show security macsec statistics
<brief | detail>
<interface *interface-name*>

Release Information Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Description Display Media Access Control Security (MACsec) statistics.

This command does not display output when MACsec is enabled using static secure association key (SAK) security mode.

Options **none**—Display MACsec statistics in brief form for all interfaces on the switch.

brief | detail—(Optional) Display the specified level of output. Using the **brief** option is equivalent to entering the command with no options (the default). The **detail** option displays additional fields that are not visible in the **brief** output.



NOTE: The field names that only appear in this command output when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel.

interface *interface-name*—(Optional) Display MACsec statistics for the specified interface only.

Required Privilege Level view

Related Documentation • [show security macsec connections on page 261](#)

List of Sample Output [show security macsec statistics interface xe-0/1/0 detail on page 265](#)

Output Fields [Table 20 on page 263](#) lists the output fields for the **show security macsec statistics** command. Output fields are listed in the approximate order in which they appear.

The field names that appear in this command output only when you enter the **detail** option are mostly useful for debugging purposes by Juniper Networks support personnel. Those field names are, therefore, not included in this table.

Table 20: show security macsec statistics Output Fields

| Field Name | Field Description | Level of Output |
|-----------------------|------------------------|-----------------|
| Interface name | Name of the interface. | All levels |

Fields for Secure Channel transmitted

Table 20: show security macsec statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--|---|-----------------|
| Encrypted packets | <p>Total number of packets transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p> | All levels |
| Encrypted bytes | <p>Total number of bytes transmitted out of the interface in the secure channel that were secured and encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p> | All levels |
| Protected packets | <p>Total number of packets transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p> | All levels |
| Protected bytes | <p>Total number of bytes transmitted out of the interface in the secure channel that were secured but not encrypted using MACsec.</p> <p>Data packets are sent in the secure channel when MACsec is enabled, and are secured using a secure association key (SAK).</p> | All levels |
| Fields for Secure Association transmitted | | |
| Encrypted packets | <p>Total number of packets transmitted out of the interface in the connectivity association that were secured and encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p> | All levels |
| Protected packets | <p>Total number of packets transmitted out of the interface in the connectivity association that were secured but not encrypted using MACsec.</p> <p>The total includes the data packets transmitted in the secure channel and secured using a SAK and the control packets secured using a connectivity association key (CAK).</p> | All levels |
| Fields for Secure Channel received | | |
| Accepted packets | <p>The number of received packets that have been accepted by the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p> <p>This counter increments for traffic that is and is not encrypted using MACsec.</p> | All levels |

Table 20: show security macsec statistics Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---|---|-----------------|
| Validated bytes | <p>The number of bytes that have been validated by the MACsec integrity check and received on the secure channel on the interface. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>This counter does not increment when MACsec encryption is disabled.</p> | All levels |
| Decrypted bytes | <p>The number of bytes received in the secure channel on the interface that have been decrypted. The secure channel is used to send all data plane traffic on a MACsec-enabled link.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p> | All levels |
| Fields for Secure Association received | | |
| Accepted packets | <p>The number of received packets that have been accepted in the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check.</p> | All levels |
| Validated bytes | <p>The number of bytes that have been validated by the MACsec integrity check and received on the connectivity association on the interface. The counter includes all control and data plane traffic accepted on the interface.</p> <p>This counter does not increment when MACsec encryption is disabled.</p> | All levels |
| Decrypted bytes | <p>The number of bytes received in the connectivity association on the interface that have been decrypted. The counter includes all control and data plane traffic accepted on the interface.</p> <p>An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACsec.</p> | All levels |

Sample Output

show security macsec statistics interface xe-0/1/0 detail

```
user@host> show security macsec statistics interface xe-0/1/0 detail
```

```
Interface name: xe-0/1/0
Secure Channel transmitted
  Encrypted packets: 123858
  Encrypted bytes:   32190903
  Protected packets: 0
  Protected bytes:   0
Secure Association transmitted
```

```
Encrypted packets: 123858
Protected packets: 0
Secure Channel received
  Accepted packets: 123877
  Validated bytes: 0
  Decrypted bytes: 32196238
Secure Association received
  Accepted packets: 123877
  Validated bytes: 0
  Decrypted bytes: 32196238
Error and debug
Secure Channel transmitted packets
  Untagged: 0, Too long: 0
Secure Channel received packets
  Control: 0, Tagged miss: 3202804
  Untagged hit: 0, Untagged: 0
  No tag: 0, Bad tag: 0
  Unknown SCI: 0, No SCI: 0
  Control pass: 0, Control drop: 0
  Uncontrol pass: 123877, Uncontrol drop: 0
  Hit dropped: 0, Invalid accept: 0
  Late drop: 0, Delayed accept: 0
  Unchecked: 0, Not valid drop: 0
  Not using SA drop: 0, Unused SA accept: 0
```

show security mka sessions

| | |
|---------------------------------|---|
| Syntax | show security mka sessions <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Display MACsec Key Agreement (MKA) session information. |
| Options | <ul style="list-style-type: none"> interface <i>interface-name</i>—(Optional) Display the MKA session information for the specified interface only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show security mka statistics on page 269 show security macsec connections on page 261 show security macsec statistics on page 263 |
| List of Sample Output | show security mka sessions on page 268 |
| Output Fields | Table 21 on page 267 lists the output fields for the show security mka sessions command. Output fields are listed in the approximate order in which they appear. |

Table 21: show security mka sessions Output Fields

| Field Name | Field Description |
|-------------------|--|
| Interface name | Name of the interface. |
| Member identifier | Name of the member identifier. |
| CAK name | Name of the Connectivity Association Key (CAK). The CAK is configured using the cak keyword when configuring the pre-shared key. |
| Transmit interval | The transmit interval. |
| Outbound SCI | Name of the outbound secure channel identifier. |
| Message number | Number of the last data message. |
| Key number | Key number. |
| Key server | Key server status. The switch is the key server when this output is yes . The switch is not the key server when this output is no . |

Table 21: show security mka sessions Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|--|
| Key server priority | The key server priority. The key server priority can be set using the key-server-priority statement. |
| Latest SAK AN | Name of the latest secure association key (SAK) association number. |
| Latest SAK KI | Name of the latest secure association key (SAK) key identifier. |
| Fields for Peer list | |
| Member identifier | Name of the member identifier. |
| Hold time | Hold time, in seconds. |
| Message number | Number of the last data message |
| SCI | Name of the secure channel identifier. |
| Lowest acceptable PN | Number of the lowest acceptable packet number (PN). |

Sample Output

show security mka sessions

```
user@host> show security mka sessions
```

```
Interface name: xe-0/1/0
Member identifier: 0CCBEE42F8778300F8D0C1DC
CAK name: 1234567890
Transmit interval: 2000(ms)
Outbound SCI: 2C:6B:F5:9D:4B:1B/1
Message number: 1526465    Key number: 0
Key server: no            Key server priority: 15
Latest SAK AN: 0          Latest SAK KI: 4F18CE25228178FD15976E4C/1
Previous SAK AN: 0        Previous SAK KI: 000000000000000000000000/0
Peer list
1. Member identifier: 4F18CE25228178FD15976E4C (live)
   Message number: 1526484 Hold time: 14500 (ms)
   SCI: 2C:6B:F5:9D:3A:1B/1
   Lowest acceptable PN: 121198
```

show security mka statistics

| | |
|---------------------------------|---|
| Syntax | show security mka statistics <interface <i>interface-name</i> > |
| Release Information | Command introduced in Junos OS Release 13.2X50-D15 for EX Series switches. |
| Description | Display MACsec Key Agreement (MKA) protocol statistics. The output for this command does not include statistics for MACsec data traffic. For MACsec data traffic statistics, see show security macsec statistics . |
| Options | <ul style="list-style-type: none"> interface <i>interface-name</i>—(Optional) Display the MKA information for the specified interface only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> show security mka sessions on page 267 show security macsec statistics on page 263 show security macsec connections on page 261 |
| List of Sample Output | show security mka statistics on page 270 |
| Output Fields | Table 22 on page 269 lists the output fields for the show security mka statistics command. Output fields are listed in the approximate order in which they appear. |

Table 22: show security mka statistics Output Fields

| Field Name | Field Description |
|---------------------------------|---|
| Received packets | <p>Number of received MKA control packets.</p> <p>This counter increments for received MKA control packets only. This counter does not increment when data packets are received.</p> |
| Transmitted packets | <p>Number of transmitted MKA packets</p> <p>This counter increments for transmitted MKA control packets only. This counter does not increment when data packets are transmitted.</p> |
| Version mismatch packets | Number of version mismatch packets. |
| CAK mismatch packets | <p>Number of Connectivity Association Key (CAK) mismatch packets.</p> <p>This counter increments when the connectivity association key (CAK) and connectivity association key name (CKN), which are user-configured values that have to match to enable MACsec, do not match for an MKA control packet.</p> |

Table 22: show security mka statistics Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------------|---|
| ICV mismatch packets | Number of ICV mismatched packets. This counter increments when the connectivity association key (CAK) value does not match on both ends of a MACsec-secured Ethernet link. |
| Duplicate message identifier packets | Number of duplicate message identifier packets. |
| Duplicate message number packets | Number of duplicate message number packets. |
| Duplicate address packets | Number of duplicate source MAC address packets. |
| Invalid destination address packets | Number of invalid destination MAC address packets. |
| Formatting error packets | Number of formatting error packets. |
| Old Replayed message number packets | Number of old replayed message number packets. |

Sample Output

show security mka statistics

```
user@host> show security mka statistics
```

```

Received packets:          1525844
Transmitted packets:       1525841
Version mismatch packets:  0
CAK mismatch packets:      0
ICV mismatch packets:      0
Duplicate message identifier packets: 0
Duplicate message number packets: 0
Duplicate address packets:  0
Invalid destination address packets: 0
Formatting error packets:   0
Old Replayed message number packets: 0
```