

Chapter 9

Summary of Protocol-Independent Routing Properties Configuration Statements

This chapter explains each of the protocol-independent routing configuration statements. The statements are organized alphabetically.

active

Syntax	(active passive);
Hierarchy Level	[edit routing-options static (defaults route)] [edit routing-options rib <i>routing-table-name</i> static (defaults route)]
Description	<p>Configure whether static routes are removed from the routing and forwarding tables when they become inactive. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with reject next hops when they are inactive.</p> <p>active—Remove a route from the routing and forwarding tables when it becomes inactive.</p> <p>passive—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.</p>
Default	active
Usage Guidelines	See “Specify Whether Inactive Routes Are Removed from the Routing or Forwarding Table” on page 81, “Configure Static Routes” on page 75.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

as-path

Syntax `as-path <as-path> <origin (egp | igp | incomplete)> <atomic-aggregate>
<aggregator as-number in-address>;`

Hierarchy Level [edit routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]

Description Associate Border Gateway Protocol (BGP) autonomous system (AS) path information with a static route.

Options *aggregator*—(Optional) Attach the BGP aggregator path attribute to the aggregate route. When using this option, you must specify the last AS number that formed the aggregate route (encoded as two octets), followed by the IPv6 address of the BGP system that formed the aggregate route.

as-number—(Optional) The last AS number that formed the aggregate route.

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for *as-path*; you must use a full, valid AS path.

atomic-aggregate—(Optional) Attach the BGP atomic-aggregate path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

in-address—(Optional) IPv6 address of the BGP system that formed the aggregate route.

origin *egp*—BGP origin attribute that indicates that the path information originated in another AS.

origin *igp*—BGP origin attribute that indicates that the path information originated within the local AS.

origin *incomplete*—BGP origin attribute that indicates that the path information was learned by some other means.

Usage Guidelines See “Configure Static Routes” on page 75.

Required Privilege Level *routing*—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

autonomous-system

Statement	autonomous-system <i>autonomous-system</i> <loops <i>number</i> >;
Hierarchy Level	[edit routing-options]
Description	Specify the router's AS number.
Options	<p><i>autonomous-system</i>—AS number. Use a number assigned to you by the Network Information Center (NIC). Range: 1 through 65,535</p> <p><i>loops number</i>—(Optional) How many times this AS number can appear in an AS path. Range: 1 through 65,535 Default: 1 (AS number can appear once)</p>
Usage Guidelines	See "Configure the AS Number" on page 87.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

color

See preference on page 112

community

Syntax community ([*community-ids*] | no-advertise | no-export | no-export-subconfed | none);

Hierarchy Level [edit routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]

Description Associate BGP community information with a static route.

Options *community-ids*—One or more community identifiers. The *community-ids* format varies according to the type of attribute that you use.

The BGP community attribute format is *as-number :community-value*:

as-number—AS number of the community member. It can be a value from 0 through 65,535.

community-value—Identifier of the community member. It can be a number from 0 through 65,535.

For specifying the BGP community attribute only, you also can specify *community-ids* as one of the following well-known community names defined in RFC 1997:

no-advertise—Routes containing this community name are not advertised to other BGP peers.

no-export—Routes containing this community name are not advertised outside a BGP confederation boundary.

no-export-subconfed—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

none—Explicitly exclude BGP community information with a static route. Include this option when configuring an individual route in the route portion to override a community option specified in the defaults portion.

The BGP extended communities attribute format is *type:administrator:assigned-number*:

type is the type of extended community and can be either a target or origin community. The target community identifies the destination to which the route is going. The origin community identifies where the route originated.

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

Usage Guidelines See "Specify Community Information" on page 79.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

confederation

Statement	<code>confederation confederation-autonomous-system members [autonomous-systems];</code>
Hierarchy Level	[edit routing-options]
Description	Specify the router's confederation AS number.
Options	<p><i>autonomous-system</i>—AS numbers of the confederation members. Range: 1 through 65,535</p> <p><i>confederation-autonomous-system</i>—Confederation AS number. Use one of the numbers assigned to you by the NIC. Range: 1 through 65,535</p> <p><i>members</i>—Confederation members.</p>
Usage Guidelines	See "Configure AS Confederation Members" on page 88.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

export-rib

Syntax	<code>export-rib routing-table-name;</code>
Hierarchy Level	[edit routing-options rib-group <i>group-name</i>]
Description	Name of the routing table from which the JUNOS software should export routing information.
Options	<i>routing-table-name</i> —Routing table group name.
Usage Guidelines	See "Create Routing Table Groups" on page 85.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

graceful-restart

Statement graceful-restart {
 disable;
 path-selection-defer-time-limit *time-limit*;
 }

Hierarchy Level [edit routing-options]

Description Configure graceful restart.

Option disable—Name of one or more policies.

time-limit—Grace period for graceful restart, in seconds.

Range: 1 through 600

Usage Guidelines See “Configure Graceful Restart” on page 88.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

import-policy

Syntax import-policy [*policy-names*];

Hierarchy Level [edit routing-options rib-group *group-name*]

Description Apply one or more policies to routes being imported into the routing table group. The import-policy statement compliments the import-rib statement, and cannot be used without first specifying the routing table from where the routes are being imported.

Options *policy-names*—Name of one or more policies.

Usage Guidelines See “Create Routing Table Groups” on page 85.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

import-rib

Syntax	import-rib [<i>routing-table-names</i>];
Hierarchy Level	[edit routing-options rib-group <i>group-name</i>]
Description	Name of the routing table into which the JUNOS software should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv6 import routing tables, the primary routing table must be inet6.0.
Options	<i>routing-table-name</i> —Name of one or more routing tables.
Usage Guidelines	See “Create Routing Table Groups” on page 85.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

install

Syntax	(install no-install);
Hierarchy Level	[edit routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)]
Description	Configure whether the JUNOS software installs all static routes into the forwarding table even if they do not have the lowest preference values. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols. <p>install—Have the JUNOS software install all static routes into the forwarding table even if they do not have the lowest preference values. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.</p> <p>no-install—Do not install the route into the forwarding table even if it is the route with the lowest preference.</p>
Default	install
Usage Guidelines	See “Configure Static Routes” on page 75.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

• interface-routes

• **Syntax** interface-routes {
• rib-group *group-name*;
• }

• **Hierarchy Level** [edit routing-options]

• **Description** Associate a routing table group with the router's interfaces and specify the routing table groups into which interface routes are imported.

• **Options** The statement is explained separately.

• **Usage Guidelines** See "Configure How Interface Routes Are Imported into Routing Tables" on page 86.

• **Required Privilege Level** routing—To view this statement in the configuration.
• routing-control—To add this statement to the configuration.

martians

Syntax	martians { <i>destination-prefix match-type</i> <allow>; }
Hierarchy Level	[edit routing-options], [edit routing-options rib <i>routing-table-name</i>]
Description	Configure martian addresses.
Options	allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. <i>destination-prefix</i> —Destination route you are configuring: <i>destination-prefix/prefix-length</i> — <i>destination-prefix</i> is the network portion of the IPv6 address, and <i>prefix-length</i> is the destination prefix length. <i>match-type</i> —Criteria that the destination must match: exact—Exactly match the route's mask length. longer—The route's mask length is greater than the specified mask length. orlonger—The route's mask length is equal to or greater than the specified mask length. through <i>destination-prefix</i> —The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix. upto <i>prefix-length</i> —The route's mask length falls between the two destination prefix lengths, inclusive. Default: exact
Usage Guidelines	See "Configure Martian Addresses" on page 84.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

maximum-routes

Statement	maximum-routes <i>route-limit</i> <(log-only threshold <i>value</i>)>;
Hierarchy Level	[edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Description	Configures the maximum number of routes to install into a routing table.
Options	<i>route-limit</i> —Maximum number of routes. If this limit is reached, a warning is triggered and any additional routes are rejected. Range: 1 through 4,294,967,295
	log-only—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.
	threshold <i>value</i> —(Optional) Threshold value for the advisory limit that triggers a warning. Range: 1 through 100
Usage Guidelines	See “Configure a Route Limit for Routing Tables” on page 86.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

metric

Syntax	metric <i>metric</i> <type <i>type</i> >;
Hierarchy Level	[edit routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)]
Description	Metric value for a static route.
Options	<i>metric</i> —Metric value. Range: 1 through 65,535
	type <i>type</i> —(Optional) Type of route. Range: 1 through 16
Usage Guidelines	See “Specify the Route Metric” on page 79.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-install

See install on page 107

no-readvertise

See readvertise on page 113

no-resolve

See resolve on page 114

no-retain

See retain on page 114

options

Syntax options {
 syslog (level *level* | upto level);
 }

Hierarchy Level [edit routing-options]

Description Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.

Options *level*—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:

 emergency—Panic or other conditions that cause the system to become unusable.

 alert—Conditions that should be corrected immediately, such as a corrupted system database.

 critical—Critical conditions, such as hard drive errors.

 error—Standard error conditions.

 warning—System warning messages.

 notice—Conditions that are not error conditions, but that might warrant special handling.

 info—Informational messages.

 debug—Software debugging messages.

Default: info

 level—Log messages at a particular level.

 upto—Log all messages up to a particular level.

Usage Guidelines See “Configure Logging for the Routing Protocol Process” on page 89.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

passive

See active on page 101

preference

Syntax (preference | preference2 | color | color2) *preference* <type *type*>;

Hierarchy Level [edit routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]

Description Preference value for a static route. You also can specify a secondary preference value (preference2) as well as colors, which are even finer-grained preference values (color and color2).

Options *preference*—Preference value. A lower number indicates a more preferred route.
Range: 1 through 255
Default: 5 (for static routes)

type—(Optional) Type of route.
Range: 1 through 16

Usage Guidelines See “Specify the Route Preference” on page 79.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

qualified-next-hop

Syntax qualified-next-hop *address* {
metric *metric*;
preference *preference*;
}

Hierarchy Level [edit routing-options rib inet6.0 static route *destination-prefix*]

Description Configure an independent metric or preference on a static route.

Options *address*—IPv6 address for the next hop.

metric—Metric value.
Range: 1 through 65,535

preference—Preference value. A lower number indicates a more preferred route.
Range: 1 through 255
Default: 5

Usage Guidelines See “Specify an Independent Preference for a Static Route” on page 77.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

readvertise

Syntax	(readvertise no-readvertise);
Hierarchy Level	[edit routing-options static (defaults route)], [edit routing-options rib <i>routing-table-name</i> static (defaults route)]
Description	Configure whether static routes are eligible to be readvertised by routing protocols: readvertise—Readvertise static routes. no-readvertise—Mark a static route as being ineligible for readvertisement; include the no-readvertise option when configuring the route.
Default	readvertise
Usage Guidelines	See “Configure Static Routes” on page 75.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

resolution

Statement	resolution { tracefilter [<i>policy-names</i>]; traceoptions { file <i>name</i> <replace> <size <i>size</i> > <files <i>number</i> > <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <flag-modifier> <disable>; } }
Hierarchy Level	[edit routing-options]
Description	Route next-hop resolution options.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

resolve

Syntax (resolve | no-resolve);

Hierarchy Level [edit routing-options rib *routing-table-name* static (defaults | route)]

Description Configure statically configured routes to be resolved to a next hop that is not directly connected.

resolve—Resolve a static route to an IP prefix that is not a directly connected next hop. In this case, the route is resolved through the inet.0 and inet.3 routing tables.

no-resolve—Prevent a static route from being resolved to an IP prefix that is not a directly connected next hop.

Default no-resolve

Usage Guidelines See “Specify When the Route Can Be Resolved to a Prefix Not Directly Connected” on page 82.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

retain

Syntax (retain | no-retain);

Hierarchy Level [edit routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]

Description Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally.

retain—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

no-retain—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally.

Default no-retain

Usage Guidelines See “Specify Whether the Route Is Permanently Installed in the Forwarding Table” on page 81.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

rib

Syntax	<pre> rib <i>routing-table-name</i> { static { defaults { <i>static-options</i>; } rib-group <i>group-name</i>; route <i>destination-prefix</i> { <i>next-hop</i>; <i>static-options</i>; } } } </pre>
Hierarchy Level	[edit routing-options]
Description	<p>Create a routing table.</p> <p>Explicitly creating a routing table with the <i>routing-table-name</i> statement is optional if you are not adding any static routes to the routing table and if you also are creating a routing table group. Simply including the rib-groups statement to declare that a routing table is part of a routing table group is sufficient to create it.</p>
Default	If you do not specify a routing table name with the <i>routing-table-name</i> statement, the software uses the default IPv4 routing tables, which are inet.0 for unicast routes and inet.1 for the multicast cache.
Options	<p><i>routing-table-name</i>—Name of the routing table, in the following format:</p> <p style="text-align: center;"><i>protocol</i>[.<i>identifier</i>]</p> <p><i>protocol</i> is the protocol family. It can be inet6 for the IPv6 family or inet for the IPv4 family.</p> <p><i>identifier</i> is a positive integer that specifies the instance of the routing table.</p> <p>Default: inet.0</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Create IPv6 Routing Tables” on page 74.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

rib-group

Syntax `rib-group group-name;`

Hierarchy Level [edit routing-options interface-routes]

Description Configure the routing table groups into which interface routes are imported.

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.

Usage Guidelines See “Configure How Interface Routes Are Imported into Routing Tables” on page 86, and “Create Routing Table Groups” on page 85.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

rib-groups

Syntax

```
rib-groups {
    group-name {
        import-policy [ policy-names ];
        import-rib [ group-names ];
        export-rib group-name;
    }
}
```

Hierarchy Level [edit routing-options]

Description Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

Each routing table group must contain one or more routing tables that the JUNOS software uses when importing routes (specified in the `import-rib` statement) and optionally can contain one routing table group that the JUNOS software uses when exporting routes to the routing protocols (specified in the `export-rib` statement).

Options *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately in this chapter.

Usage Guidelines See “Create Routing Table Groups” on page 85.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

router-id

Statement	router-id <i>address</i> ;
Hierarchy Level	[edit routing-options]
Description	Specify the router's IPv6 address.
Options	<i>address</i> —IPv6 address of the router. Default: Address of the first interface encountered by the JUNOS software.
Usage Guidelines	See “Configure the Router Identifier” on page 88.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

routing-options

Statement	routing-options { ... }
Hierarchy Level	[edit]
Description	Configure protocol-independent routing properties.
Usage Guidelines	See “Configure Protocol-Independent Routing Properties” on page 73.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

static

```

Syntax  static {
           defaults {
             static-options;
           }
           rib-group group-name;
           route destination-prefix {
             next-hop;
             qualified-next-hop address {
               metric metric;
               preference preference;
             }
             (active | passive);
             as-path < as-path> < origin (egp | igp | incomplete)> < atomic-aggregate>
               < aggregator as-number in-address>;
             community [ community-ids ];
             (install | no-install);
             (metric | metric2 | metric3 | metric4) value < type type> ;
             (preference | preference2 | color | color2) preference < type type> ;
             (readvertise | no-readvertise);
             (resolve | no-resolve);
             (no-retain | retain);
           }
         }

```

Hierarchy Level [edit routing-options rib inet6.0]

Description Configure static routes to be installed in the routing table. You can specify any number of routes within a single static statement, and you can specify any number of static statements in the configuration.

Options *destination-prefix*—Destination of the generated route, specified as:

destination-prefix/prefix-length—*destination-prefix* is the network portion of the IPv6 address, and *prefix-length* is the destination prefix length.

next-hop—How to reach the destination, specified as:

next-hop address—Reach the next-hop router by specifying an IPv6 address or an interface name.

next-table routing-table-name—Name of the next routing table to the destination.


reject—Do not forward packets addressed to this destination. Instead, drop the packets, send Internet Control Message Protocol version 6 (ICMPv6) unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

discard—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMPv6 unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

receive—Install a receive route for this destination into the routing table.

qualified-next-hop—An independent metric or preference for a static route.

address—Reach the next-hop router by specifying an IPv6 address.

 <p>Note</p>	<p>The qualified-next-hop statement is mutually exclusive with all other types of next hops, except for next-hop address. Therefore, next-hop reject, next-hop discard, and next-hop receive cannot be configured with qualified-next-hop for the same destination.</p>
--	---

The remaining statements are explained separately.

Usage Guidelines See “Configure Static Routes” on page 75.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

tracefilter

Syntax tracefilter [*policy-names*];

Hierarchy Level [edit routing-options resolution]

Description Define the routing policy used to filter trace output.

Options *policy-name*—Name of the routing policy used to filter trace output.

Required Privilege Level routing and trace—To view this statement in the configuration.
 routing-control and trace-control—To add this statement to the configuration.

traceoptions

Syntax traceoptions {
 file *name* <replace> <size *size*> <files *number*> <no-stamp>
 <(world-readable | no-world-readable)>;
 flag *flag* <*flag-modifier*> <disable>;
 }

Hierarchy Level [edit routing-options],
 [edit routing-options resolution]

Description Define tracing operations that track all routing protocol functionality in the router.

To specify more than one tracing operation, include multiple flag statements.

Default If you do not include this statement, no global tracing operations are performed.

Options disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *name*—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. We recommend that you place global routing protocol tracing output in the file routing-log.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the size option.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. Global routing protocol tracing options include the following:

all—All tracing operations

config-internal—Configuration internals

event—Event processing (resolution statement only)

flash—Flash processing (resolution statement only)

general—All normal operations and routing table changes (a combination of the normal and route trace operations)

indirect—Indirect next-hop add or change or delete (resolution statement only)

kernel—Kernel communication (resolution statement only)

normal—All normal operations

parse—Configuration parsing

policy—Routing policy operations and actions

regex-parse—Regular-expression parsing

route—Routing table changes

state—State transitions

task—Interface transactions and processing

timer—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

detail—Detailed trace information

receive—Packets being received

send—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you do not include this option, timestamp information is placed at the beginning of each line of the tracing output.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the files option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

Usage Guidelines See “Trace Global Routing Protocol Operations” on page 90.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

