

JKL

j1

Description:	Enables the J1 variant (Japan) of the T1 framing. The no version disables the feature (default).
Syntax:	[no] j1
Mode(s):	Controller Configuration

key

Description:	From Radius Configuration mode, specifies the authentication or accounting server secret. The no version removes the secret. From Manual Key Configuration mode, configures a manual ISAKMP/IKE preshared key. There is no no version. To delete a key, use the no version of the ipsec key manual command.
Syntax:	To assign a RADIUS key: key <i>secret</i> no key <ul style="list-style-type: none"> <i>secret</i> – authentication or accounting server secret text string used by RADIUS to encrypt the client and server authenticator field during exchanges between the system and a RADIUS server. The system encrypts PPP PAP passwords using this text string. To assign an ISAKMP/IKE key: key <i>keyString</i> <ul style="list-style-type: none"> <i>keyString</i> – key value in ASCII format; up to 200 characters
Mode(s):	Radius Configuration, Manual Key Configuration

l2f checksum

Description:	Enables the generation of checksums for L2F data packets running over IP/UDP. The no version disables the generation of checksums for data packets running over IP/UDP. The default setting is disabled.
Syntax:	[no] l2f checksum
Mode(s):	Global Configuration

I2f destruct-timeout

- Description:** Specifies the maximum time for which the system maintains dynamic destinations, tunnels, and sessions that have terminated. If resources are low, the system will replace the terminated objects with new requests. The **no** version restores the default value, 600 seconds.
- Syntax:** I2f destruct-timeout *seconds*
no I2f destruct-timeout
- *seconds* – a time in the range 10–3600 seconds (1 hour)
- Mode(s):** Global Configuration

I2f drain

- Description:** Prevents the creation of new destinations, tunnels, and sessions for the system. This command works with the **I2f shutdown** command. Both commands affect the administrative state of L2F on the system. The **I2f drain** command sets the administrative state to drain, and the **I2f shutdown** command sets the administrative state to disabled. The **no** version allows the creation of new destinations, tunnels, and sessions for the system.
- Syntax:** [no] I2f drain
- Mode(s):** Global Configuration

I2f drain destination

- Description:** Prevents the creation of new tunnels and sessions at a destination. This command works with the **I2f shutdown destination** command. Both commands affect the status of the administrative state of L2F for the destination. The **I2f drain destination** command sets the administrative state to drain, and the **I2f shutdown destination** command sets the administrative state to disabled. The **no** version allows the creation of new tunnels and sessions at a destination.
- Syntax:** [no] I2f drain destination { *destinationName* |
[virtual-router *virtualRouterName*] ip *ipAddress* }
- *destinationName* – name the system assigns to the home gateway
 - *virtualRouterName* – name of the virtual router on which the destination exists
 - *ipAddress* – IP address of the home gateway
- Mode(s):** Global Configuration

l2f drain tunnel

Description: Prevents the assignment of new sessions to a tunnel. This command works with the **l2f shutdown tunnel** command. Both commands affect the status of the administrative state of L2F for the tunnel. The **l2f drain tunnel** command sets the administrative state to drain, and the **l2f shutdown tunnel** command sets the administrative state to disabled. The **no** version allows the assignment of new sessions to a tunnel.

Syntax: [no] l2f drain tunnel { *destinationName* | [virtual-router *vrName*]
ip *ipAddress tunnelName* }

- *destinationName* – name the system assigns to the home gateway
- *vrName* – name of the virtual router on which the tunnel exists
- *ipAddress* – IP address of the home gateway
- *tunnelName* – name of the tunnel

Mode(s): Global Configuration

l2f ignore-receive-data-sequencing

Description: Suppresses sequence number checking for data packets received on all L2F tunnels in the system. This setting affects only packets received on a tunnel, not packets sent on a tunnel. The L2F NAS still inserts sequence numbers into data packets if the NAS receives packets from the home gateway that contain sequence numbers. The **no** version, which is the default, causes the system to check sequence numbers in data packets that it receives on L2F tunnels.



Note: *If you are using IP reassembly, we recommend that you set up the system to ignore sequence numbers in received data packets. Because IP reassembly may reorder L2F packets, out-of-order packets may be dropped if sequence numbers are being used on L2F data packets.*

Syntax: [no] l2f ignore-receive-data-sequencing

Mode(s): Global Configuration

I2f shutdown

- Description:** Closes all destinations, tunnels, and sessions and prevents the creation of new destinations, tunnels, and sessions for the system. This command works with the **I2f drain** command. Both commands affect the administrative state of L2F on the system. The **I2f shutdown** command sets the administrative state to disabled, and the **I2f drain** command sets the administrative state to drain. The **no** version allows the creation of new destinations, tunnels, and sessions for the system.
- Syntax:** [no] I2f shutdown
- Mode(s):** Global Configuration

I2f shutdown destination

- Description:** Closes all tunnels and sessions at a destination, and prevents the creation of new tunnels and sessions at that destination. This command works with the **I2f drain destination** command. Both commands affect the status of the administrative state of L2TP on the destination. The **I2f shutdown destination** command sets the administrative state to disabled, and the **I2f drain destination** command sets the administrative state to drain. The **no** version enables the creation of new tunnels and sessions at that destination.
- Syntax:** [no] I2f shutdown destination { *destinationName* | [virtual-router *vrName*] ip *ipAddress* }
- *destinationName* – name the system assigns to the home gateway
 - *vrName* – name of the virtual router on which the destination exists
 - *ipAddress* – IP address of the home gateway
- Mode(s):** Global Configuration

I2f shutdown session

- Description:** Closes a specific session. The **no** version has no effect because all L2F sessions are dynamic and cannot be restarted after they have been shut down.
- Syntax:** [no] I2f shutdown session { *destinationName* | [virtual-router *vrName*] ip *ipAddress* *sessionName* }
- *destinationName* – name that the system assigns to the LNS
 - *vrName* – name of the virtual router on which the destination exists
 - *ipAddress* – IP address of the home gateway
 - *sessionName* – name of the session
- Mode(s):** Global Configuration

I2f shutdown tunnel

Description: Closes all sessions in a tunnel, and prevents the creation of new sessions in that tunnel. This command works with the **I2f drain tunnel** command. Both commands affect the status of the administrative state of L2F on the tunnel. The **I2f shutdown tunnel** command sets the administrative state to disabled, and the **I2f drain tunnel** command sets the administrative state to drain. The **no** version enables the creation of new sessions in that tunnel.

Syntax: [no] I2f shutdown tunnel { *destinationName* |
[virtual-router *vrName*] ip *ipAddress tunnelName* }

- *destinationName* – name the system assigns to the home gateway
- *virtualRouterName* – name of the virtual router on which the tunnel exists
- *ipAddress* – IP address of the home gateway
- *tunnelName* – name of the tunnel

Mode(s): Global Configuration

I2tp checksum

Description: Enables the generation of checksums for data packets running over IP/UDP. The **no** version disables the generation of checksums for data packets running over IP/UDP. The default setting is disabled.

Syntax: [no] I2tp checksum

Mode(s): Global Configuration

I2tp destination profile

Description: Defines the location of the LAC(s) by virtual router and IP address. Accesses the L2TP Destination Profile Configuration mode. The **no** version removes the L2TP destination profile.

Syntax: I2tp destination profile { *profileName* [[virtual-router *vrName*]
ip address *ipAddress*] | [virtual-router *vrName*] ip address *ipAddress* }

no I2tp destination profile { *profileName* |
[virtual-router *vrName*] ip address *ipAddress* }

- *profileName* – name of the L2TP destination profile
- *vrName* – name of the virtual router to be used to reach the destination (that is, the LAC). If you do not specify a virtual router, the current virtual router context is used.
- *ipAddress* – IP address to be used to reach the destination



Note: To manage an existing destination profile, use this version of the command: **I2tp destination profile** *profileName*. This version assumes that the L2TP destination profile already exists.

Mode(s): Global Configuration

I2tp destruct-timeout

Description: Specifies the maximum time for which the system maintains dynamic destinations, tunnels, and sessions that have terminated. If resources are low, the system will replace the terminated objects with new requests. The **no** version restores the default value, 600 seconds.

Syntax: I2tp destruct-timeout *seconds*

no I2tp destruct-timeout

- *seconds* – time in the range 10–3600 seconds (1 hour)

Mode(s): Global Configuration

I2tp drain

Description: Prevents the creation of new destinations, tunnels, and sessions for the system. This command works in conjunction with the **I2tp shutdown** command. Both commands affect the status of the administrative state of L2TP on the system; the **I2tp drain** command sets the administrative state to drain. The **no** version allows the creation of new destinations, tunnels, and sessions for the system.

Syntax: [no] I2tp drain

Mode(s): Global Configuration

I2tp drain destination

Description: Prevents the creation of new tunnels and sessions at a destination. This command works in conjunction with the **I2tp shutdown destination** command. Both commands affect the status of the administrative state of L2TP for the destination; the **I2tp drain destination** command sets the administrative state to drain. The **no** version allows the creation of new tunnels and sessions at a destination.

Syntax: [no] I2tp drain destination { *destinationName* |
[virtual-router *vrName*] ip *address* }

- *destinationName* – name the system assigns to the LNS
- *vrName* – name of the virtual router on which the destination exists
- *address* – IP address of the LNS

Mode(s): Global Configuration

I2tp drain tunnel

Description: Prevents the assignment of new sessions to a tunnel. This command works in conjunction with the **I2tp shutdown tunnel** command. Both commands affect the status of the administrative state of L2TP for the tunnel; the **I2tp drain tunnel** command sets the administrative state to drain. The **no** version allows the assignment of new sessions to a tunnel.

Syntax: [no] I2tp drain tunnel { *destinationName* |
[virtual-router *vrName*] ip *address tunnelName* }

- *destinationName* – name the system assigns to the LNS
- *vrName* – name of the virtual router on which the tunnel exists
- *address* – IP address of the LNS
- *tunnelName* – name of the tunnel

Mode(s): Global Configuration

l2tp ignore-receive-data-sequencing

Description: Suppresses sequence number checking for data packets received on all L2TP tunnels in the system. This setting affects only packets received on a tunnel, not packets sent on a tunnel. The L2TP LAC still inserts sequence numbers into data packets if the LAC receives packets from the LNS that contain sequence numbers. The **no** version, which is the default, causes the system to check the sequence numbers in data packets that it receives on L2TP tunnels.



Note: *If you are using IP reassembly, we recommend that you set up the system to ignore sequence numbers in received data packets. Because IP reassembly may reorder L2TP packets, out-of-order packets may be dropped if sequence numbers are being used on L2TP data packets.*

Syntax: [no] l2tp ignore-receive-data-sequencing

Mode(s): Global Configuration

l2tp retransmission

Description: Sets the number of retransmission retries. The **no** version resets the number of retransmissions to the default value, 5.

Syntax: l2tp retransmission *retries*
no l2tp retransmission

- *retries* – in the range 2–7

Mode(s): Global Configuration

l2tp shutdown

Description: Closes all destinations, tunnels, and sessions and prevents the creation of new destinations, tunnels, and sessions for the system. This command works in conjunction with the **l2tp drain** command. Both commands affect the status of the administrative state of L2TP on the system; the **l2tp shutdown** command sets the administrative state to disabled. The **no** version allows the creation of new destinations, tunnels, and sessions for the system.

Syntax: [no] l2tp shutdown

Mode(s): Global Configuration

I2tp shutdown destination

Description: Closes all tunnels and sessions at a destination, and prevents the creation of new tunnels and sessions at that destination. This command works in conjunction with the **I2tp drain destination** command. Both commands affect the status of the administrative state of L2TP on the destination; the **I2tp shutdown destination** command sets the administrative state to disabled. The **no** version enables the creation of new tunnels and sessions at that destination.

Syntax: [no] I2tp shutdown destination { *destinationName* |
[virtual-router *vrName*] ip *ipAddress* }

- *destinationName* – name the system assigns to the LNS
- *vrName* – name of the virtual router on which the destination exists
- *ipAddress* – IP address of the LNS

Mode(s): Global Configuration

I2tp shutdown session

Description: Closes a specific session. The **no** version has no effect because all L2TP sessions are dynamic and cannot be restarted after they have been shut down.

Syntax: [no] I2tp shutdown session { *destinationName* |
[virtual-router *vrName*] ip *ipAddress* *sessionName* }

- *destinationName* – name that the system assigns to the LNS
- *vrName* – name of the virtual router on which the destination exists
- *ipAddress* – IP address of the LNS
- *sessionName* – name of the session

Mode(s): Global Configuration

I2tp shutdown tunnel

Description: Closes all sessions in a tunnel, and prevents the creation of new sessions in that tunnel. This command works in conjunction with the **I2tp drain tunnel** command. Both commands affect the status of the administrative state of L2TP on the tunnel; the **I2tp shutdown tunnel** command sets the administrative state to disabled. The **no** version enables the creation of new sessions in that tunnel.

Syntax: [no] I2tp shutdown tunnel { *destinationName* |
[virtual-router *vrName*] ip *address tunnelName* }

- *destinationName* – name the system assigns to the LNS
- *vrName* – name of the virtual router on which the tunnel exists
- *address* – IP address of the LNS
- *tunnelName* – name of the tunnel

Mode(s): Global Configuration

I2tp tunnel idle-timeout

Description: Configures the tunnel idle-timeout value and creates persistent tunnels by setting the value to 0. There is no **no** version.

Syntax: I2tp tunnel idle-timeout [*timerValue*]

- *timerValue* – range is 0–86400 seconds

Mode(s): Global Configuration

I2tp tunnel-switching

Description: Enables tunnel switching chassis-wide. The **no** version disables tunnel switching. Disabled is the default.

Syntax: [no] I2tp tunnel-switching

Mode(s): Global Configuration

lease

Description: Specifies the time period for which the supplied IP address is valid. The **no** version restores the default lease time, one day.

Syntax: lease { *days* [*hours* [*minutes* [*seconds*]]] | infinite }

no lease

- *days* – number of days for which the IP address is valid in the range 0–32768
- *hours* – number of hours for which the IP address is valid in the range 0–24
- *minutes* – number of minutes for which the IP address is valid in the range 0–60
- *seconds* – number of seconds for which the IP address is valid in the range 0–60
- infinite – assigns a lease that does not expire

Mode(s): Pool Configuration

license b-ras

Description: Specifies the B-RAS license provided by your sales representative or Juniper Networks Customer Service. Depending on the license purchased, the system supports up to 2,000, 4,000, 8,000, 16,000, or 20,000 authenticated PPP or SDX (formerly SSC) sessions. The **no** version disables the license.

Syntax: license b-ras *licenseKey*

no license b-ras

- *licenseKey* – unique string of alphanumeric characters up to 15 characters long that we provide to you

Mode(s): Global Configuration

lifetime

Description: Associates a lifetime with IKE SAs established using this IKE policy. The **no** version restores the lifetime to its default, 28800 seconds (8 hours).

Syntax: lifetime *seconds*

no lifetime

- *seconds* – number of seconds an SA lives before expiring; range is 7200 to 864000 (10 days)

Mode(s): ISAKMP Policy Configuration

limits

Description: Sets memory limits for BGP internal tables maintained by BGP software. If you set a particular memory limit to a value lower than the current value and the system uses the memory up to the previous limit, then memory allocations will start to fail when the new value takes effect. The **no** version restores the default values; entering an optional value in the **no** version has the same effect as entering no optional values.

Syntax:

```
limits { { nlr | received-route } receivedRouteLimit |
{ nlri | destination } destinationLimit | path-attribute pathAttributeLimit |
vrf vrfLimit | address-family addressFamilyLimit | peer peerLimit |
peer-address-family peerAddressFamilyLimit |
peer-group peer-groupLimit | peer-group-address-family
peer-groupAddressFamilyLimit | dampening dampeningLimit |
network-route networkRouteLimit | aggregated-route aggregatedRouteLimit |
redistributed-route redistributedRouteLimit |
auto-summary-route autoSummaryRouteLimit | next-hop nextHopLimit |
route-flap-history routeFlapHistoryLimit | rib-out ribOutLimit |
group-rib-out groupRibOutLimit | send-queue-entry sendQueueEntryLimit |
route-target-entry routeTargetEntryLimit }

no limits { { nlr | received-route } [ receivedRouteLimit ] |
{ nlri | destination } [ destinationLimit ] | path-attribute [ pathAttributeLimit ] |
vrf [ vrfLimit ] | address-family [ address-family-Limit ] | peer [ peerLimit ] |
peer-address-family [ peerAddressFamilyLimit ] |
peer-group [ peer-groupLimit ] |
peer-group-address-family [ peer-groupAddressFamilyLimit ] |
dampening [ dampeningLimit ] | network-route [ networkRouteLimit ] |
aggregated-route [ aggregatedRouteLimit ] |
redistributed-route [ redistributedRouteLimit ] |
auto-summary-route [ autoSummaryRouteLimit ] | next-hop [ nextHopLimit ] |
route-flap-history [ routeFlapHistoryLimit ] | rib-out [ ribOutLimit ] |
group-rib-out [ groupRibOutLimit ] |
send-queue-entry [ sendQueueEntryLimit ] |
route-target-entry [ routeTargetEntryLimit ] }
```



Note: The **nlre** and **received-route** keywords have the same purpose in this command, to set a limit on the received routes table. Similarly, the **nlri** and **destination** keywords have the same purpose, to set a limit on the BGP destination table. The **nlre** and **nlri** keywords are maintained for compatibility with previous software releases.

- *receivedRouteLimit* – maximum number of received routes stored by BGP in the range 0–2147483648; the default is 5,000,000
- *destinationLimit* – maximum number of BGP destinations stored by BGP in the range 0–2147483648; the default is 5,000,000
- *pathAttributeLimit* – maximum number of path attributes stored by BGP in the range 0–2147483648; the default is 5,000,000

- *vrfLimit* – maximum number of VRFs stored by BGP in the range 0–2147483648; the default is 5,000,000
- *addressFamilyLimit* – maximum number of address families stored by BGP in the range 0–2147483648; the default is 5,000,000
- *peerLimit* – maximum number of peers stored by BGP in the range 0–2147483648; the default is 5,000,000
- *peerAddressFamilyLimit* – maximum number of peers per address family stored by BGP in the range 0–2147483648; the default is 5,000,000
- *peer-groupLimit* – maximum number of peer groups stored by BGP in the range 0–2147483648; the default is 5,000,000
- *peer-groupAddressFamilyLimit* – maximum number of peer-groups per address family stored by BGP in the range 0–2147483648; the default is 5,000,000
- *dampeningLimit* – maximum number of dampening parameter blocks stored by BGP in the range 0–2147483648; the default is 5,000,000. BGP creates a dampening parameter block for each unique set of dampening parameters—such as suppress threshold, reuse threshold, and so on—used by BGP. For example, if you have a route map that sets the dampening parameters to one set of values for some routes and to another set of values for the remaining routes, BGP uses and stores two dampening parameter blocks, one for each set.
- *networkRouteLimit* – maximum number of network routes stored by BGP in the range 0–2147483648; the default is 5,000,000
- *aggregatedRouteLimit* – maximum number of aggregated routes stored by BGP in the range 0–2147483648; the default is 5,000,000
- *redistributedRouteLimit* – maximum number of redistributed routes stored by BGP in the range 0–2147483648; the default is 5,000,000
- *autoSummaryRouteLimit* – maximum number of automatically summarized routes stored by BGP in the range 0–2147483648; the default is 5,000,000
- *nextHopLimit* – maximum number of next hops stored by BGP in the range 0–2147483648; the default is 5,000,000
- *routeFlapHistoryLimit* – maximum number of route-flap histories stored by BGP in the range 0–2147483648; the default is 5,000,000
- *ribOutLimit* – maximum number of RIB-Out routes stored by BGP for individual peers in the range 0–2147483648; the default is 1,000,000
- *groupRibOutLimit* – maximum number of RIB-Out routes stored by BGP for peer groups in the range 0–2147483648; the default is 1,000,000
- *sendQueueEntryLimit* – maximum number of send queue entries stored by BGP in the range 0–2147483648; the default is 5,000,000
- *routeTargetEntryLimit* – maximum number of combined import and export route-target entries in the range 0–2147483648; the default is 5,000,000

Mode(s): Router Configuration

line

Description: Opens virtual terminal lines or the console line and allows you to configure the lines. By default five vty lines (0–4) are open.

The **no** version removes a vty line or a range of lines from your configuration; users will not be able to run Telnet, SSH, or FTP to lines that you remove. When you remove a vty line, the system removes all lines above that line. For example, **no line vty 6** causes the system to remove lines 6 through 19. You cannot remove lines 0 through 4.



Note: Once lines are open, login is enabled by default. Before users can access the lines, you must configure a password, disable login using the **no login** command, or configure AAA authentication on the line.

Syntax: line { console *lineNumber* | vty *lineRangeStart* [*lineRangeEnd*] }

no line vty *lineNumber*

- console – specifies the console line
- vty – specifies vty lines
- *lineNumber* – number of a single line; 0 for the console line
- *lineRangeStart* – start of the vty line range; a number from 0–19;
- *lineRangeEnd* – end of the vty line range; a number from 0–19

Mode(s): Global Configuration

lineCoding

Description: Specifies the type of line coding used by a CE1 or CT1 interface. The **no** version restores the default—hdb3 for CE1 interfaces and b8zs for CT1 interfaces.

Syntax: lineCoding *linecodingType*

no lineCoding

- *linecodingType* – one of the following:
 - › ami – alternate mark inversion
 - › b8zs – bipolar with eight-zero substitution; CT1 default
 - › hdb3 – high-density bipolar 3; CE1 default

Mode(s): Controller Configuration

link

- Description:** Links the pool currently being configured to another DHCP local address pool. The linked pool acts as a backup pool. The **no** version removes the link.
- Syntax:** link *poolName*
no link
- *poolName* – name of pool to which you want to link the pool currently being configured
- Mode(s):** Pool Configuration

list

- Description:** Lists the currently configured MPLS explicit path (optionally starting at a particular index). There is no **no** version.
- Syntax:** list [*index*]
- *index* – number of a node in an ordered set of abstract nodes, a value ranging from 1–255; set with the **index** command
- Mode(s):** Explicit Path Configuration

load-interval

- Description:** Sets the time interval at which the system calculates bit rates and packet rates for an interface. The **no** version restores the default time interval, which is 300 seconds. This command is not available for the Ethernet interface on the SRP module.
- Syntax:** load-interval *timeInterval*
no load-interval
- *timeInterval* – a multiple of 30 seconds in the range 30–300
- Mode(s):** Interface Configuration

local host

- Description:** Configures an L2TP local hostname to be used with a remote host. The **no** version removes the local hostname from use with a remote host.
- Syntax:** local host *hostname*
no local host
- *hostname* – an L2TP local hostname; can be up to 64 characters in length (no spaces)
- Mode(s):** L2TP Destination Profile Host Configuration

local ip address

Description: Configures a local IP address for use with a remote host. The **no** version removes the local IP address from use with a remote host.

Syntax: local ip address *ipAddress*
no local ip address

- *ipAddress* – an IP address

Mode(s): L2TP Destination Profile Host Configuration

log

Description: Configures logging settings. The **no** version negates the command. The **suspend** version suspends the policy rule.

Syntax: [no] [suspend] log [classifier-group] [precedence]

- classifier-group – specifies the classifier control list for use in this log policy
- precedence – specifies the precedence of the policy rule

Mode(s): Policy Configuration

log-adjacency-changes

Description: Generates a log message when a NLSP adjacency changes state (up or down). The **no** version disables this function. This command manipulates the same log as the Global Configuration **log** commands.

Syntax: log-adjacency-changes [severity { *severityValue* | *severityNumber* }]
[verbosity *verbosityLevel*]

no log-adjacency-changes

- severity – minimum severity of the log messages for this category; described either by a descriptive term—*severityValue*—or by a corresponding number—*severityNumber*—in the range 0–7. The lower the number, the higher the priority:
 - › emergency *or* 0 – system unusable
 - › alert *or* 1 – immediate action needed
 - › critical *or* 2 – critical condition exists
 - › error *or* 3 – error condition
 - › warning *or* 4 – warning condition
 - › notice *or* 5 – normal but significant condition
 - › info *or* 6 – informational message
 - › debug *or* 7 – debug message
- verbosity – specifies the verbosity of this log category's messages
- *verbosityLevel* – specifies the verbosity of the log category's messages; can be any of the following:
 - › high – verbose
 - › low – terse
 - › medium – moderate detail

Mode(s): Router Configuration

log destination

Description: Configures the logging of system messages. You can direct messages to a destination, limit the messages logged based on severity level, or limit the event categories for which messages are logged. The **no** versions restore default settings or reverse the effect of previous commands that limited event categories.



Note: You can display traffic logs—such as *ipTraffic*, *icmpTraffic*, *tcpTraffic*, and *udpTraffic*—only via the **show log data** command or from the SRP module console. You cannot redirect traffic logs elsewhere, such as to a syslog or nonvolatile storage file, or to a Telnet session.

Syntax: To specify the destination and severity of messages logged:

```
log destination { console | nv-file | syslog ipAddress [ facility facilityId ] }  
{ severity { severityValue | severityNumber } | off }  
  
no log destination [ syslog [ ipAddress ] ]
```

To specify which event categories are logged to syslog:

```
log destination syslog ipAddress { include | exclude } category [ category ]*  
  
no log destination syslog ipAddress { include | exclude } [ category ]*
```

- console – configure or modify logging to the local console
- nv-file – configure or modify logging to the nonvolatile log file; the nv-file can accept only events at a severity level of critical or higher in importance
- syslog – configure or modify logging to a syslog server
- *ipAddress* – IP address of the syslog application on a remote host
- facility – specifies the syslog facility on the host
- *facilityId* – number in the range 0–7 that identifies the corresponding logging facility, local0–local7
- severity – minimum severity of the log messages displayed; described either by a descriptive term—*severityValue*—or by a corresponding number—*severityNumber*—in the range 0–7. The lower the number, the higher the priority:
 - › emergency or 0 – system unusable
 - › alert or 1 – immediate action needed
 - › critical or 2 – critical condition exists
 - › error or 3 – error condition
 - › warning or 4 – warning condition
 - › notice or 5 – normal but significant condition
 - › info or 6 – informational message

- › debug or 7 – debug message
- off – disable logging to this destination
- include – send only the specified event categories to the syslog server
- exclude – send all event categories except those specified to the syslog server

Issuing an **include** command after an **exclude** command (or vice versa) overrides the earlier command.

You can issue successive **include** commands or successive **exclude** commands. Successive commands expand the list of included or excluded categories.

- * – indicates that one or more parameters can be repeated multiple times in a list in the command line

Mode(s): Global Configuration

log destination syslog source

Description: Specifies a source interface type and location for events logged to a syslog server. Overrides the type and location of the actual source to enable server access behind firewalls. The **no** version restores the default state, which is to use the actual interface type and location of the source.

Syntax: log destination syslog *ipAddress* source *interfaceType* *interfaceSpecifier*
no log destination syslog *ipAddress* source [*interfaceType* *interfaceSpecifier*]

- *ipAddress* – IP address of the syslog application
- *interfaceType* – type of interface; the source of the events logged; see *Interface Types and Specifiers* in *About This Guide*.
- *interfaceSpecifier* – particular interface; format varies according to interface type; see *Interface Types and Specifiers* in *About This Guide*.

Mode(s): Global Configuration

log engineering

Description: Enables engineering logs. The **no** version disables engineering logs.

Syntax: [no] log engineering

Mode(s): Global Configuration

log fields

- Description:** Selects optional fields to be added to all logs. The **no** version disables the optional log fields.
- Syntax:** log fields { timestamp | no-timestamp } { instance | no-instance }
{ calling-task | no-calling-task }
no log fields
- timestamp – include the timestamp in log messages
 - no-timestamp – do not include the timestamp in log messages
 - instance – include the event ID in log messages
 - no-instance – do not include the event ID in log messages
 - calling-task – include the logging task name in log messages
 - no-calling-task – do not include the logging task name in log messages
- Mode(s):** Global Configuration

log filters

- Description:** This command has only a **no** version. See the **no log filters** command for a complete description and syntax.

log here

- Description:** Enables the current terminal as a log console. The **no** version disables logs destined for a console from being displayed on the current terminal.
- Syntax:** [no] log here
- Mode(s):** User Exec, Privileged Exec, Global Configuration

log severity

Description: Sets the severity value for the selected category. The **no** version removes an override severity setting and returns a log to its default value or the system-wide setting.

Syntax: log severity { *severityValue* | off | *severityNumber* }
[*eventCategory* [*instanceTree*] | *eventCategory instanceTree* | *eventCategory*
]

no log severity [*severityValue* | off | *severityNumber*] [*eventCategory* [filters | *instanceTree*] | *eventCategory* { filters | *instanceTree* } | *eventCategory* | *]

- *severityValue* and *severityNumber* – minimum severity of the log messages displayed for the selected category; described either by a descriptive term—*severityValue*—or by a corresponding number—*severityNumber*—in the range 0–7. The lower the number, the higher the priority:
 - › emergency *or* 0 – system unusable
 - › alert *or* 1 – immediate action needed
 - › critical *or* 2 – critical condition exists
 - › error *or* 3 – error condition
 - › warning *or* 4 – warning condition
 - › notice *or* 5 – normal but significant condition
 - › info *or* 6 – informational message
 - › debug *or* 7 – debug message
- off – disables log messages for all event categories or for a specified event category
- *eventCategory* – log category; refer to the CLI online Help for available options
- filters – removes all log filters for the event category
- *instanceTree* – log-specific filter parameters; refer to the CLI online Help for available options
- * – resets all log severities, system wide and individual, to default settings

Mode(s): Global Configuration

log unlimit

- Description:** Removes the limit on the number of outstanding buffers for an event category. The **no** version returns the number of buffers to the default value.
- Syntax:** [no] log unlimit [*eventCategory*]
- *eventCategory* – log category; refer to the CLI online Help for available options
- Mode(s):** User Exec, Privileged Exec, Global Configuration

log verbosity

- Description:** Sets the verbosity level for a selected category. The **no** version returns the log verbosity to its default value, low.
- Syntax:** log verbosity *verbosityLevel* [*eventCategory*]
no log verbosity [*verbosityLevel*] [*eventCategory*]
- *verbosityLevel* – specifies the verbosity for the log category:
 - › low – terse (default)
 - › medium – moderate detail
 - › high – verbose
 - *eventCategory* – log category; refer to the CLI online Help for available options
- Mode(s):** Global Configuration

login

- Description:** Requires you to log in with a password. The **no** version removes the password requirement and allows connections without a password.
- Syntax:** [no] login
- Mode(s):** Line Configuration



Note: *If this command has been configured and no password has been configured, access to Telnet is refused.*

login authentication

- Description:** Applies an AAA authentication list to the vty sessions that you specified for AAA authentication. The **no** version removes all authentication methods, which means the system accepts Telnet sessions without challenge.
- Syntax:** login authentication *authListName*
no login authentication
- *authListName* – specifies an authentication list name of up to 32 characters
- Mode(s):** Line Configuration

logout subscribers

- Description:** Logs out the authenticated PPP users. If you do not specify a license, B-RAS configuration commands are disabled. There is no **no** version.
- Syntax:** logout subscribers { all | username *userName* | domain *domainName* | virtual-router *vrName* | port *interfaceLocation* }
- all – all PPP sessions
 - *userName* – active PPP session whose names match the username
 - *domainName* – active PPP session whose usernames have that domain name
 - *vrName* – active PPP session whose interfaces are bound to a specific virtual router
 - port – active PPP subscribers for the port
 - *interfaceLocation* – location of the port in slot/port format; format varies according to interface type; see *Interface Types and Specifiers* in *About This Guide*
- Mode(s):** Privileged Exec

loopback

Description: Specifies the loopback mode for a module controller or interface, or maps a loopback interface to a user domain name in Domain Map Configuration mode. The **no** version clears all loopback on the module or interface (the default), or deletes the mapping to the user domain name.

Syntax: **Module controllers** – The options available vary depending on the module being configured.

CE1 module:

```
loopback { local | network { payload | line } }
```

```
no loopback
```

CT1 module:

```
loopback { local | network { payload | line } | remote { line { fdl { ansi | bellcore }  
| inband } payload [ fdl ] [ ansi ] } }
```

```
no loopback [ remote ]
```

CT3, E3, or T3 module:

```
loopback { local | network | payload }
```

```
no loopback
```

cOCx/STMx SONET controller (SONET/SDH section layer), OCx/STMx line modules:

```
loopback { local | network }
```

```
no loopback
```

X.21/V.35 module:

```
[ no ] loopback
```

- **local** – loops the data back toward the router and sends an AIS out toward the network.
- **network payload** – loops the data toward the network after the framer has processed the data.
- **network line** – loops the data toward the network before the data reaches the framer.
- **remote line fdl ansi** – sends a repeating 16-bit ESF data link code word (00001110 11111111) to the remote end requesting that it enter into a network line loopback. Specify the **ansi** keyword to enable the remote line FDL ANSI bit loopback on the T1 line, according to the ANSI T1.403 specification.

- remote line fdl bellcore – sends a repeating 16-bit ESF data link code word (00010010 11111111) to the remote end requesting that it enter into a network line loopback. Specify the **bellcore** keyword to enable the remote line FDL Bellcore bit loopback on the T1 line, according to the Bellcore TR-TSY-000312 specification.
- remote line inband – sends a repeating 5-bit inband pattern (00001) to the remote end requesting that it enter into a network line loopback.
- remote payload [fdl] [ansi] – sends a repeating 16-bit ESF data link code word (00010100 11111111) to the remote end requesting that it enter into a network payload loopback. Enables the remote payload FDL ANSI bit loopback on the T1 line. You can specify **fdl and ansi**, but it is not necessary.
- remote – based on the last activate request sent to the remote end, sends the 16-bit ESF data link code word or inband pattern to deactivate the loopback at the remote end.
- network – loops the data toward the network before the data reaches the framer.
- payload – loops the data toward the network after the framer has processed the data.

Interfaces – The options available vary depending on the interface being configured.

ATM interface – This command cannot be used on a subinterface:

```
loopback { diagnostic | line }
```

```
no loopback
```

POS interface:

```
[ no ] loopback internal | line
```

- diagnostic – places the interface into internal loopback
- line – ATM interface: places the interface into external loopback; POS interface: connects the received network signal directly to the transmit network signal. When configured in line loopback mode, the system never receives data from the network.
- internal – connects the local transmitted signal to the local received signal

User Domain Name

```
loopback interface-number
```

```
no loopback
```

- *interface-number* – interface number in the range 1–32000

Mode(s): Controller Configuration, Interface Configuration, Domain Map Configuration

lsp-gen-interval

- Description:** Sets the minimum interval at which originated IS-IS link state packets are generated on a per LSP basis. The **no** version restores the default interval.
- Syntax:** lsp-gen-interval [level-1 | level-2] *seconds*
no lsp-gen-interval [level-1 | level-2]
- level-1 – sets interval for level 1 only
 - level-2 – sets interval for level 2 only
 - *seconds* – number in the range 0–120; the minimum interval in seconds; the default value is 5 seconds
- Mode(s):** Router Configuration

lsp-mtu

- Description:** Sets the maximum size of an IS-IS link state packet generated by the software. The **no** version restores the default MTU size of 1497 bytes.
- Syntax:** lsp-mtu *bytes*
no lsp-mtu
- *bytes* – number in the range 128–9180; the MTU size in bytes; the default value is 1497
- Mode(s):** Router Configuration

lsp-refresh-interval

- Description:** Sets the link state packet rate at which locally generated IS-IS link state packets are periodically transmitted. The **no** version restores the default refresh interval.
- Syntax:** lsp-refresh-interval *seconds*
no lsp-refresh-interval
- *seconds* – number in the range 1–65535; the refresh interval in seconds; the default value is 900
- Mode(s):** Router Configuration