



Global Interface Commands

This module describes the global command line interface (CLI) commands for configuring interfaces on the Cisco 8000 Series Routers.

To use commands of this module, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using any command, contact your AAA administrator for assistance.

- [bandwidth \(global\), on page 2](#)
- [clear interface, on page 3](#)
- [dampening, on page 4](#)
- [encapsulation dot1ad dot1q, on page 6](#)
- [encapsulation dot1q, on page 7](#)
- [interface \(global\), on page 8](#)
- [lACP system , on page 10](#)
- [mtu, on page 11](#)
- [rewrite ingress tag, on page 14](#)
- [show im dampening, on page 16](#)
- [show interfaces, on page 19](#)

bandwidth (global)

To configure the bandwidth of an interface, use the **bandwidth** command in interface configuration mode.

bandwidth *rate*

Syntax Description	<i>rate</i> Amount of bandwidth to be allocated on the interface, in Kilobits per second (kbps). Range is from 0 through 4294967295.
---------------------------	--

Command Default	The default bandwidth depends on the interface type.
------------------------	--

Command Modes	Interface configuration
----------------------	-------------------------

Command History	Release	Modification
	Release 7.0.12	This command was introduced.

Usage Guidelines



Note	To obtain the default bandwidth for a specific interface, use the show interfaces command after you first bring up the interface. The default interface bandwidth is displayed in the show interfaces command output.
-------------	---

Task ID	Task ID	Operations
	interface	execute
	basic-services	read, write

Examples

This example shows how to configure the bandwidth on a Ten Gigabit Ethernet interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router# interface TenGigE 0/4/1/0
RP/0/RP0/CPU0:router# bandwidth 4000000
```

clear interface

To clear interface statistics or packet counters, use the **clear interface** command in XR EXEC mode .

clear interface *type interface-path-id*

Syntax Description

type Interface type. For more information, use the question mark (?) online help function.

interface-path-id Physical interface or virtual interface.

Note Use the **show interfaces** command to see a list of all interfaces currently configured on the router.

For more information about the syntax for the router, use the question mark (?) online help function.

Command Default

No default behavior or values

Command Modes

EXEC

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Task ID

Task ID	Operations
interface	execute
basic-services	read, write

Examples

This example shows how to use the **clear interface** command to clear the loopback interface 2:

```
RP/0/RP0/CPU0:router# clear interface loopback 2
```

dampening

To limit propagation of transient or frequently changing interface states on Interface Manager (IM) clients, turn on event dampening by using the **dampening** command in interface configuration mode. To turn dampening off, use the **no** form of this command.

dampening [*half-life* [*reuse suppress max-suppress-time*]]

Syntax Description

<i>half-life</i>	(Optional) Time (in minutes) after which a penalty is decreased. Once the interface has been assigned a penalty, the penalty is decreased by half after the half-life period. The process of reducing the penalty happens every 5 seconds. The range of the half-life period is 1 to 45 minutes. The default is 1 minute.
<i>reuse</i>	(Optional) Penalty value below which a stable interface is unsuppressed. Range is from 1 through 20000. Default value is 750.
<i>suppress</i>	(Optional) Limit at which an interface is suppressed when its penalty exceeds that limit. Range is from 1 through 20000, and must be greater than the reuse threshold. The default value is 2000.
<i>max-suppress-time</i>	(Optional) Maximum time (in minutes) that an interface can be suppressed. This value effectively acts as a ceiling that the penalty value cannot exceed. Default value is four times the half-life period.

Command Default

Dampening is turned off by default. When you use the **dampening** command, the following default values are enabled for any optional parameters that you do not enter:

- *half-life*: 1 minute
- *reuse*: 750
- *suppress*: 2000
- *max-suppress-time*: Four times the half-life

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Event dampening suppresses a constantly unstable interface until it remains stable for a period of time. Enabling dampening on an interface that already has dampening configured has the effect of resetting the penalty associated with that interface to zero. The reuse threshold must always be less than the suppress threshold.

Consider the following guidelines when configuring event dampening:

- Configuring dampening on both a subinterface and its parent is usually unnecessary because their states are almost always the same and dampening would be triggered at the same time on each interface.
- If all subinterfaces require dampening, then apply dampening to the main interface only. Applying configuration to large numbers of subinterfaces requires an abundance of memory and increases the time required to process the configuration during boot and failover.

- When dampening is enabled, an interface has a penalty value associated with it. The value starts at 0 and is increased by 1000 whenever the underlying state of the interface changes from up to down.
- The penalty value decreases exponentially while the interface state is stable. If the penalty value exceeds a configured suppress threshold, then the state of the interface is suppressed and IM will not notify upper layers of further state transitions. The suppressed state remains until the penalty value decreases past a configured reuse threshold.

Task ID	Task ID	Operations
	interface	read, write

Examples

This example shows how to enable dampening with default values on an interface:

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 0/4/0/0
RP/0/RP0/CPU0:router(config-if)# dampening
```

Related Commands

Command	Description
show im dampening, on page 16	Displays the state of all interfaces on which dampening has been configured.

encapsulation dot1ad dot1q

To define the matching criteria to be used in order to map single-tagged 802.1ad frames ingress on an interface to the appropriate service instance, use the **encapsulation dot1ad dot1q** command in sub-interface configuration mode. To remove the configuration, use the **no** form of this command.

encapsulation dot1ad *vlan-id* **dot1q** *vlan-id*

Syntax Description

dot1ad Indicates that the IEEE 802.1ad provider bridges encapsulation type is used for the outer tag.

dot1q Indicates that the IEEE 802.1q standard encapsulation type is used for the inner tag.

vlan-id VLAN ID, can be given as single ID.

Command Default

No matching criteria are defined.

Command Modes

Sub-interface configuration

Command History

Release	Modification
Release 7.2.12	This command was introduced.

Usage Guidelines

The outer VLAN tag is an 802.1ad VLAN tag, instead of an 802.1Q tag. An 802.1ad tag has an ethertype value of 0x88A8, instead of 0x8100 that 802.1Q uses.

Some of the fields in the 802.1ad VLAN header are interpreted differently per 802.1ad standard.

A **tunneling ethertype** command applied to the main interface does not apply to an 802.1ad sub-interface. An interface with encapsulation dot1ad causes the router to categorize the interface as an 802.1ad interface. This causes special processing for certain protocols and other features:

- MSTP uses the IEEE 802.1ad MAC STP address instead of the STP MAC address.
- Certain QoS functions may use the Drop Eligibility (DE) bit of the IEEE 802.1ad tag.

Examples

The following example shows how to map single-tagged 802.1ad ingress frames to a service instance:

```
Router# configure
Router# interface hundredGigE 0/0/0/1.10
Router(config-subif)# encapsulation dot1ad 100 dot1q 20
```

Related Commands

Command	Description
rewrite ingress tag, on page 14	Specifies the encapsulation adjustment that is to be performed on the frame ingress to the service instance.

encapsulation dot1q

To define the matching criteria to map 802.1Q frames ingress on an interface to the appropriate service instance, use the **encapsulation dot1q** command in the sub-interface configuration mode. To delete the matching criteria to map 802.1Q frames ingress on an interface to the appropriate service instance, use the **no** form of this command.

encapsulation dot1q *vlan-id*

Syntax Description

vlan-id VLAN ID, can be given as single ID.

Command Default

No matching criteria are defined.

Command Modes

Sub-interface configuration

Command History

Release	Modification
Release 7.2.12	This command was introduced.

Usage Guidelines

Only one encapsulation statement can be applied to a sub-interface. Encapsulation statements cannot be applied to main interfaces.

A single encapsulation dot1q statement specifies matching for frames with a single VLAN ID.

Examples

The following example shows how to map 802.1Q frames ingress on an interface to the appropriate service instance:

```
Router# configure
Router(config)#interface HundredGigE 0/0/0/24.10
Router(config-if)# encapsulation dot1q 10
```

The following example shows how to map 802.1Q frames ingress on an l2transport sub-interface:

```
Router# configure
Router(config)# interface HundredGigE 0/0/0/24.10 l2transport
Router(config-subif)# encapsulation dot1q 10
```

Related Commands

Command	Description
rewrite ingress tag, on page 14	Specifies the encapsulation adjustment that is to be performed on the frame ingress to the service instance.

interface (global)

To configure an interface or to create or configure a virtual interface, use the **interface** command in XR Config mode. To delete the interface configuration, use the **no** form of this command.

interface *type interface-path-id*

Syntax Description	<i>type</i>	Interface type. For more information, use the question mark (?) online help function.
	<i>interface-path-id</i>	Physical interface or virtual interface.
	Note	Use the show interfaces command to see a list of all interfaces currently configured on the router.
		For more information about the syntax for the router, use the question mark (?) online help function.

Command Default No interfaces are configured

Command Modes XR Config

Command History	Release	Modification
	Release 7.0.12	This command was introduced.

Usage Guidelines The **interface** command enters interface configuration mode to allow you to configure interfaces. If a virtual interface is configured, then the interface is created if it did not already exist.

The **no** form of this command applies only to virtual interfaces or to subinterfaces (that is, interfaces that have been created in global configuration mode).

Until Release 6.5.1, when you create an interface with some configurations, upon router or interface reload, interface configurations are lost. From Release 6.5.1, onwards, automatic shutdown config behavior is persistent and no shutdown configs are lost on interface or router reload.

Task ID	Task ID	Operations
		interface read, write

Examples

In the following example, the **interface** command is given for the card in location 0/2/0/1, and interface configuration mode is entered for that interface:

```
RP/0/RP0/CPU0:router(config)# interface POS 0/2/0/1
```


Related Commands

Command	Description
clear interface, on page 3	Clears interface statistics or packet counters.

lACP system

To set the default system parameters for the Link Aggregation Control Protocol (LACP) bundles, use the **lACP system** command in XR Config mode .

lACP system { **mac** | **priority** }

Syntax Description	mac	priority
	Unique MAC address used to identify the system in LACP negotiations.	Priority for this system. Lower value is higher priority. Range is from 1 to 65535.

Command Default System priority is 32768. MAC address is automatically assigned from the backplane pool.

Command Modes XR Config

Command History	Release	Modification
	Release 7.0.12	This command was introduced.

Usage Guidelines The parameters are the system MAC address and the priority of the system. The MAC address must be unique to the system (if it matches a partner system, LACP negotiations fail). The combination of the MAC address and the set system priority determine the priority of the LACP bundles.

Task ID	Task ID	Operation
	bundle read,	write

Example

The following example shows how to configure the MAC address on an LACP system:

```
RP/0/RP0/CPU0:router (config) lACP system mac 000c.15c0.bd15
```

mtu

To adjust the maximum transmission unit (MTU) value for packets on the interface, use the **mtu** command in interface configuration mode. To return the interface to the default MTU for the interface type, use the **no** form of this command.

mtu *bytes*

Syntax Description

bytes Maximum number of bytes in a Layer 2 frame. Range is from 64 through 65535.

Command Default

The default MTU for each interface is as follows:

- Ethernet—1514 bytes
- POS—4474 bytes
- Tunnel—1500 bytes
- Loopback—1514 bytes
- ATM—4470 bytes

Command Modes

Interface configuration

Command History

Release	Modification
Release 7.0.12	This command was introduced.

Usage Guidelines

Use the **mtu** command to set a specific MTU value for an interface, or use the **no mtu** command to return the interface to the default MTU value for that interface type. The MTU value can be increased or decreased using the **mtu** command, subject to minimum and maximum MTU limits for the interface type.

If the MTU value is not configured, then each interface will have a default MTU value that is specific to the interface type. The default MTU value is generally the largest Layer 2 frame size possible for the interface type.

The default/configured MTU value on an atm interface includes the L2 header.

The MTU size consists of L2 header that includes either SNAP(8bytes)/MUX(0)/NLPID(2) header or the AAL5 SDU. The AAL5 SDU includes the L3 datagram and the optional Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) header.

The Ethernet interface is the Layer 3 datagram plus 14 bytes. For ATM main interface, the MTU is L3 datagram + 0 bytes.

For ATM L3 sub interface, mtu is as follows:

- SNAP - L3 datagram + 8 bytes
- NLPID - L3 datagram + 2 bytes
- MUX - L3datagram + 0 bytes
- When no pvc is configured under sub interface - L3datagram + 0 bytes

You can use the **show interfaces** command to determine if the MTU value has been changed. The **show interfaces** command output displays the MTU size for each interface in the MTU (byte) field. Note that the MTU size that is displayed includes the Layer 2 header bytes used for each encapsulation type.



Note You can use the **show interfaces** command to determine if the MTU value has been changed. The **show interfaces** command output displays the MTU size for each interface in the MTU (byte) field. Note that the MTU size that is displayed includes the Layer 2 header bytes used for each encapsulation type.



Note Changing the MTU on an interface triggers a change on the protocols and encapsulations configured on that interface, although some protocol-specific configurations can override the interface MTU. For example, specifically changing the interface MTU configuration does not affect the IP MTU configuration, but may affect the resulting MTU on that node.



Note For the 10x10GigE CPAK (10 ports with only 8 profiles), it is not possible to support 10 different MTUs on 10 different 10GigE ports. One of the profiles needs to be reserved for the default MTU, in case you need to change the configured MTU back to the default MTU. Therefore on the 10x10g CPAK, you can configure different MTU sizes on 7 ports and the other 3 ports have the default MTU size. If you configure the 8th port, the configuration command succeeds but an error appears on the console.

Task ID

Task ID Operations

```
interface read,
write
```

Examples

In this example, the MTU value for all interfaces is verified. The MTU value is shown in the next-to-last column:

```
RP/0/RP0/CPU0:router# show interfaces all brief
```

Intf Name	Intf State	LineP State	Encap Type	MTU (byte)	BW (Kbps)
Nu0	up	up	Null	1500	Unknown
TenGigE6/0/0/0	up	up	HDLC	4474	2488320
TenGigE6/0/0/1	up	up	HDLC	4474	2488320
TenGigE6/0/0/2	admin-down	admin-down	HDLC	4474	2488320
TenGigE6/0/0/3	admin-down	admin-down	HDLC	4474	2488320
Mg0//CPU0/0	up	up	ARPA	1514	100000

```
RP/0/RP0/CPU0:router# configure
```

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 6/0/0/0
```

```
RP/0/RP0/CPU0:router(config-if)# mtu 1000
```

After the **mtu** command is used to decrease the MTU Layer 2 frame size for the POS interface on 6/0/0/0 to 1000 bytes, the **show interfaces all brief** command is used again to verify that the MTU Layer 2 frame size has been changed:

```
RP/0/RP0/CPU0:router# show interfaces all brief
```

Intf Name	Intf State	LineP State	Encap Type	MTU (byte)	BW (Kbps)
Nu0	up	up	Null	1500	Unknown
PO6/0/0/0	up	up	HDLC	1000	2488320
PO6/0/0/1	up	up	HDLC	4474	2488320
PO6/0/0/2	admin-down	admin-down	HDLC	4474	2488320
PO6/0/0/3	admin-down	admin-down	HDLC	4474	2488320
Mg0//CPU0/0	up	up	ARPA	1514	100000

rewrite ingress tag

To specify the encapsulation adjustment that is to be performed on the frame ingress to the service instance, use the **rewrite ingress tag** command in the interface configuration mode. To delete the encapsulation adjustment that is to be performed on the frame ingress to the service instance, use the **no** form of this command.

```
rewrite ingress tag pop { 1 | 2 } symmetric | { push dot1ad vlan-id dot1q vlan-id |
dot1q vlan-id symmetric } | { translate { 1-to-1 dot1ad vlan-id | dot1q vlan-id symmetric
} | 1-to-2 dot1ad vlan-id dot1q vlan-id symmetric | 2-to-1 { dot1ad vlan-id | dot1q
vlan-id } symmetric | 2-to-2 dot1ad vlan-id dot1q vlan-id symmetric }
```

Syntax Description	
<i>vlan-id</i>	VLAN ID, can be given as single ID.
push dot1q <i>vlan-id</i>	Pushes one 802.1Q tag with <i>vlan-id</i> .
push dot1ad <i>vlan-id</i>	Pushes one Dot1ad tag with <i>vlan-id</i> .
pop 1	One tag is removed from the packet. This command can be combined with a push (pop N and subsequent push <i>vlan-id</i>).
pop 2	Two tags are removed from the packet. This command can be combined with a push (pop N and subsequent push <i>vlan-id</i>).
translate 1-to-1 dot1q <i>vlan-id</i> or translate 1-to-1 dot1ad <i>vlan-id</i>	Replaces the incoming tag (defined in the encapsulation command) into a different 802.1Q or dot1ad tag at the ingress service instance.
translate 1-to-2 dot1q <i>vlan-id</i> or translate 1-to-2 dot1ad <i>vlan-id</i>	Replaces the incoming tag defined by the encapsulation command by a pair of 802.1Q or dot1ad tags.
translate 2-to-2 dot1q <i>vlan-id</i> or translate 2-to-2 dot1ad <i>vlan-id</i>	Replaces the pair of tags defined by the encapsulation command by a pair of VLANs defined by this rewrite.
translate 2-to-1 dot1q <i>vlan-id</i> or translate 2-to-1 dot1ad <i>vlan-id</i>	Replaces a pair of tags defined in the encapsulation command by <i>vlan-id</i> .
symmetric	A rewrite operation is applied on both ingress and egress. The operation on egress is the inverse operation as ingress. Note Symmetric is the default behavior. Hence, it cannot be disabled.

Command Default The frame is left intact on ingress.

Command Modes Interface configuration

Command History	Release	Modification
	Release 7.2.12	This command was introduced.

Usage Guidelines The **symmetric** keyword is accepted only when a single VLAN is configured in encapsulation. If a list of VLANs is configured in encapsulation, the **symmetric** keyword is accepted only for push rewrite operations; all other rewrite operations are rejected.

The **pop** command assumes the elements being popped are defined by the encapsulation type.

The **rewrite ingress tag translate** command assume the tags being translated from are defined by the encapsulation type. In the 2-to-1 option, the “2” means 2 tags of a type defined by the **encapsulation** command. The translation operation requires at least “from” tag in the original packet. If the original packet contains more tags than the ones defined in the “from”, then the operation should be done beginning on the outer tag.

Examples

The following example shows how to specify the encapsulation adjustment that is to be performed on the frame ingress to the service instance:

```
Router# configure
Router(config)# interface hundredGigE 0/0/0/24.1 l2transport
Router(config-if)# encapsulation dot1q 10
Router(config-if)# rewrite ingress tag push dot1q 200 symmetric
```

The following example shows how to remove one outer tag from the packet:

```
Router# configure
Router(config)# interface hundredGigE 0/0/0/24.1 l2transport
Router(config-if)# encapsulation dot1q 10
Router(config-subif)# rewrite ingress tag pop 1 symmetric
```

The following example shows how to replace the incoming tag (defined in the encapsulation command) into a different dot1ad tag at the ingress service instance:

```
Router# configure
Router(config)# interface hundredGigE 0/0/0/24.1 l2transport
Router(config-if)# encapsulation dot1q 10
Router(config-subif)# rewrite ingress tag translate 1-to-1 dot1ad 2 symmetric
```

Related Commands	Command	Description
	encapsulation dot1q, on page 7	Defines the matching criteria to map 802.1Q frames ingress on an interface to the appropriate service instance.
	encapsulation dot1ad dot1q, on page 6	Defines the matching criteria to be used in order to map single-tagged 802.1ad frames ingress on an interface to the appropriate service instance.

show im dampening

To display the state of all interfaces on which dampening has been configured, use the **show im dampening** command in XR EXEC mode .

```
show im dampening [{interface type | ifhandle handle}]
```

Syntax Description	
interface type	(Optional) Interface type. For more information, use the question mark (?) online help function.
ifhandle handle	(Optional) Identifies the caps node whose Interface Manager (IM) dampening information you want to display.

Command Default If you do not specify an interface, then the system displays brief details about all dampened interfaces.

Command Modes EXEC

Command History	Release	Modification
	Release 7.0.12	This command was introduced.

Usage Guidelines If you do not specify an interface, then the system displays brief details about all dampened interfaces.

The physical hardware (layer 1) is not the only part of an interface that can change state. L2 keepalive failure event is one of the many instances that can have a similar impact on routing protocols despite the underlying interface state staying UP. To take account of such events, when dampening is configured on an interface, it is applied independently to every layer. They all use the same parameters as the interface but they have their own penalty value which is incremented when that layer changes state.

Capsulations that may be dampened in this way include these:

- L2 basecaps, such as HDLC and PPP, which may flap if keepalives are not received due to events such as intermittent packet loss.
- L3 capsulations (for example ipv4, ipv6). These may be brought down if another link has a conflicting IP address configured.
- Other locations where negotiation takes place with a peer router, as in the case of PPP control protocols such as IPCP. If the negotiation fails, then the caps is brought down.

Task ID	Task ID	Operations
	interface	read

Examples

This example shows the output from the **show im dampening** command issued with default values:

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 0/4/0/0
RP/0/RP0/CPU0:router(config-if)# no shutdown
```



```

RP/0/RP0/CPU0:router(config-if)# dampening
RP/0/RP0/CPU0:router# show im dampening

Interface                Proto          Caps          Penalty Suppressed
-----
TenGigE0/4/0/0          0              0              0      NO

RP/0/RP0/CPU0:router# show im dampening interface TenGigE 0/4/0/0

TenGigE0/4/0/0 (0x05000d00)
Dampening enabled: penalty 0, not suppressed
  underlying state: Up
  half_life: 1      reuse: 750
  suppress: 3000    max-suppress-time: 4
  restart-penalty: 0

RP/0/RP0/CPU0:router# show interfaces TenGigE 0/4/0/0

TenGigE0/4/0/0 is up, line protocol is down
  Dampening enabled: penalty 0, not suppressed
  half_life: 1      reuse: 750
  suppress: 3000    max-suppress-time: 4
  restart-penalty: 0
  Hardware is Ten Gigabit Ethernet
  Description: ensoft-gsr5 TenGigE 4\2
  Internet address is Unknown
  MTU 4474 bytes, BW 155520 Kbit
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation HDLC, crc 16, controller loopback not set, keepalive set (10 sec)
  Last clearing of "show interface" counters never
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 total input drops
  0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  48 packets output, 1504 bytes, 0 total output drops
  Output 0 broadcast packets, 0 multicast packets
  0 output errors, 0 underruns, 0 applique, 0 resets
  0 output buffer failures, 0 output buffers swapped out

```

This sample output shows a POS interface with PPP basecaps and IPCP. The subsequent output for **show im dampening interface <ifname>** contains a table of any capsulations which have their own penalty as shown below:

```

RP/0/RP0/CPU0:router# show im dampening

Interface                Protocol        Capsulation    Pen  Sup
-----
HundredGigabitEthernet0/0/0/0          629  NO
HundredGigabitEthernet0/0/0/1          2389 YES
POS0/2/0/0                              0   NO
POS0/2/0/0          <base>          ppp            0   NO
POS0/2/0/0          ipv4            ipcp           0   NO

RP/0/RP0/CPU0:router# show im dampening interface TenGigaE 0/1/0/0

```

show im dampening

```
TenGigE 0/1/0/0 (0x01180020)
Dampening enabled: Penalty 1625, SUPPRESSED (42 secs remaining)
Underlying state: Down
half-life: 1 reuse: 1000
suppress: 1500 max-suppress-time: 4
restart-penalty: 0
```

Protocol	Capsulation	Pen	Suppression	U-L State
ipv6	ipv6	1625	YES 42s remaining	Down



Note When dampening is configured on an interface it is also applied independently to all encapsulations on that interface. For example, the ppp or hdlc basecaps state can flap even while the interface stays up and if keepalives fail. The **show im dampening interface** command contains one line for each such encapsulation as well as the interface itself.

Table 1: show im dampening Field Descriptions

Field	Description
Dampening	Indicates the dampening state and penalty value: not suppressed, suppressed.
underlying state	Underlying state of the interface: up, down, administratively down (if an interface has been configured to be “shutdown”).
half_life	This is the time (in minutes) at which the penalty on the interface would be half that of the original penalty (of 1000) when the interface transitions from UP to DOWN. It ranges from 1 to 45 minutes and the default is 1 minute.
reuse	Penalty value below which a stable interface is unsuppressed. It ranges from 1 to 20000 and the default value is 750.
suppress	Limit at which an unstable interface is suppressed when the penalty value exceeds the suppress value. It ranges from 1 to 20000 and the default value is 2000.
max-suppress-time	Maximum time (in minutes) that an interface can be suppressed. The default is 4 minutes.
restart-penalty	Penalty assigned to the interface when it flaps.

Related Commands

Command	Description
dampening, on page 4	Turns on event dampening.

show interfaces

To display statistics for all interfaces configured on the router or for a specific node, use the **show interfaces** command in XR EXEC mode.

```
show interfaces [{type interface-path-id | all | local | location node-id}] [{accounting | brief | description | detail | summary}]
```

Syntax	Description
<i>type</i>	(Optional) Specifies the type of interface for which you want to display statistics. For more information, use the question mark (?) online help function.
<i>interface-path-id</i>	Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.
all	(Optional) Displays interface information for all interfaces. This is the default.
local	(Optional) Displays interface information for all interfaces in the local card.
location <i>node-id</i>	(Optional) Displays information about all interfaces on the specified node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation.
accounting	(Optional) Displays the number of packets of each protocol type that have been sent through the interface.
brief	(Optional) Displays brief information of each interface (one line per interface).

description	(Optional) Displays the status, protocol, and description of each interface (one line per interface).
detail	(Optional) Displays detailed information about each interface. This is the default.
summary	(Optional) Displays a summary of interface information by interface type.

Command Default No default behavior or values

Command Modes XR EXEC mode

Command History	Release	Modification
	Release 7.0.12	This command was introduced.

Usage Guidelines The **show interfaces** command displays statistics for the network interfaces. The resulting display shows the interface processors in slot order.

For example, if you type the **show interfaces** command without an interface type, you receive information for all the interfaces installed in the networking device. Only by specifying the interface *type*, *slot*, and *port* arguments can you display information for a particular interface.

If you enter a **show interfaces** command for an interface type that has been removed from the networking device, an error message is displayed: “Interface not found.”

The output displayed depends on the network for which an interface has been configured.



Note The 5-minute input and output rates should be used only as an approximation of traffic per second during a given 5-minute period. These rates are exponentially weighted averages with a time constant of 5 minutes. A period of four time constants must pass before the average is within 2 percent of the instantaneous rate of a uniform stream of traffic over that period.

Task ID	Task ID Operations
	interface read

Examples This example shows the output from the **show interfaces** command. The output displayed depends on the type and number of interface cards in the networking device.

```
RP/0/RP0/CPU0:router# show interfaces HundredGigE 0/3/0/35
```

```
HundredGigE0/3/0/35 is up, line protocol is up
```

```

Interface state transitions: 1
Hardware is HundredGigE, address is e666.9aa0.223c (bia e666.9aa0.223c)
Description: **To RouterX Hu0/7/0/2**
Internet address is 192.168.1.29/30
MTU 1514 bytes, BW 100000000 Kbit (Max: 100000000 Kbit)
  reliability 255/255, txload 239/255, rxload 238/255
Encapsulation ARPA,
Full-duplex, 100000Mb/s, unknown, link type is force-up
output flow control is off, input flow control is off
Carrier delay (up) is 10 msec
loopback not set,
Last link flapped 3w3d
ARP type ARPA, ARP timeout 04:00:00
Last input 00:00:00, output 00:00:00
Last clearing of "show interface" counters never
30 second input rate 93725392000 bits/sec, 32528692 packets/sec
30 second output rate 93726416000 bits/sec, 32527860 packets/sec
 68118736643563 packets input, 24783244282360579 bytes, 0 total input drops
 0 drops for unrecognized upper-level protocol
Received 0 broadcast packets, 0 multicast packets
   0 runts, 0 giants, 0 throttles, 0 parity
174 input errors, 174 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
68115867305777 packets output, 24782409845763776 bytes, 0 total output drops
Output 0 broadcast packets, 0 multicast packets
 0 output errors, 0 underruns, 0 applique, 0 resets
 0 output buffer failures, 0 output buffers swapped out
 0 carrier transitions

```

This example shows bundle member links whose link interface status is “err-disable” and line protocol state is “admin-down” after the bundle interface has been administratively shut down using the **shutdown** command:

```
RP/0/RP0/CPU0:router# show interfaces brief
```

```
Thu May 6 06:30:55.797 DST
```

Intf Name	Intf State	LineP State	Encap Type	MTU (byte)	BW (Kbps)
BE10	down	down	ARPA	1514	0
BE100	up	up	ARPA	1514	100000000
BE101	up	up	ARPA	1514	100000000
Lo0	up	up	Loopback	1500	0
Nu0	up	up	Null	1500	0
Fo0/3/0/26	admin-down	admin-down	ARPA	1514	40000000
Hu0/3/0/0	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/1	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/2	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/3	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/4	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/5	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/6	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/7	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/8	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/9	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/10	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/11	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/12	admin-down	admin-down	ARPA	1514	100000000
Hu0/3/0/13	down	down	ARPA	1514	100000000
Hu0/3/0/14	up	up	ARPA	1514	100000000
Hu0/3/0/15	up	up	ARPA	1514	100000000
Hu0/3/0/16	admin-down	admin-down	ARPA	1514	100000000

```

Hu0/3/0/17 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/18 up up ARPA 1514 100000000
Hu0/3/0/19 up up ARPA 1514 100000000
Hu0/3/0/20 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/21 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/22 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/23 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/24 up up ARPA 1514 100000000
Hu0/3/0/25 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/27 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/28 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/29 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/30 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/31 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/32 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/33 admin-down admin-down ARPA 1514 100000000
Hu0/3/0/34 down down ARPA 1514 100000000
Hu0/3/0/35 up up ARPA 1514 100000000
Mg0/RP0/CPU0/0 up up ARPA 1514 1000000
Mg0/RP1/CPU0/0 up up ARPA 1514 1000000

```

Table 2: show interfaces Field Descriptions

Field	Description
Interface name	Displays the name of the current interface. In the example, the interface name is TenGigE0/1/0/1.
Interface state	Displays the state of the interface. In the example, the interface is in the administratively down state.
Interface state transitions	<p>Displays the number of times the interface has changed the state.</p> <p>Note</p> <ul style="list-style-type: none"> • Interface state transitions command counts only if the interface stays up. If the line protocol flaps, then it is not counted. • Interface state transitions counts the state when the line protocol state changes the state from up to down/admin-down or admin-down/down to up. If an interface changes the state from down to admin-down or admin-down to down, the counter is not incremented. • Use the clear state-transitions command to clear the counter for the current or all interfaces.

Field	Description
line protocol state	<p>Displays the state of the Layer 2 line protocol. This field may be different from the interface state if, for example, a keepalive failure has brought down the Layer 2.</p> <p>Note The line protocol state is not the same as the protocol state displayed in the show ip interfaces command, because it is the state of Layer 2 (media) rather than Layer 3 (IP protocol).</p>
Hardware	Displays the current hardware type.
address is <i>n.n.n.n/n</i>	<p>Displays the Layer 2 address (MAC address for Ethernet interfaces).</p> <p>Note Enter the mac-address command to configure the hardware address.</p>
bia	<p>Displays the burned-in address (BIA) for the interface. The BIA is the default L2 (MAC) address for the interface.</p> <p>Note The BIA is not configurable.</p>
description	<p>Displays the user-defined string that is associated with the interface.</p> <p>Note Enter the description command to configure the description associated with the interface.</p>
Internet address	<p>Displays the Layer 3 (IP) address for the interface.</p> <p>Note Enter the ipv4 address command to configure the internet address for the interface.</p>
MTU	<p>Displays the maximum transmission unit (MTU) for the interface. The MTU is the maximum packet size that can be transmitted over the interface.</p> <p>Note The MTU field indicates the interface MTU. Enter the mtu command to configure a lower MTU value at the Layer 3 level.</p>
BW	Displays the bandwidth of the interface in kbps.

Field	Description
reliability	Displays the proportion of packets that are not dropped and do not have errors. Note The reliability is shown as a fraction of 255.
txload	Indicates the traffic flowing out of the interface as a proportion of the bandwidth. Note The txload is shown as a fraction of 255.
rxload	Indicates the traffic flowing into the interface as a proportion of the bandwidth. Note The rxload is shown as a fraction of 255.
Encapsulation	Layer 2 encapsulation installed on the interface.
CRC	Indicates the length of the cyclic redundancy check (CRC), in bytes. Note The CRC is not present for all interface types. Note Enter the pos crc command to configure the CRC.
loopback or controller loopback	Indicates whether the hardware has been configured to be looped back. Note Enter the loopback command to configure the loopback or controller loopback.
keepalive	Displays the configured keepalive value, in seconds. Note Enter the keepalive command to configure the value of the keepalive field. Note The <i>keepalive</i> field may not be present if it is not applicable to the interface type.
Duplexity	Displays the duplexity of the link. Note This field is present only for shared media. Note For some interface types, you can configure the duplexity by entering the full-duplex and half-duplex commands.

Field	Description
Speed	Speed and bandwidth of the link in Mbps. This field is present only when other parts of the media info line are also displayed (see duplexity and media type).
Media Type	Media type of the interface.
output flow control	Whether output flow control is enabled on the interface.
input flow control	See output flow control.
ARP type	Address Resolution Protocol (ARP) type used on the interface. This value is not displayed on interface types that do not use ARP.
ARP timeout	ARP timeout in <i>hours:mins:secs</i> . This value is configurable using the arp timeout command.
Last clearing of counters	Time since the following counters were last cleared using the clear counters exec command in <i>hours:mins:secs</i> .
5 minute input rate	<p>Average number of bits and packets received per second in the last 5 minutes. If the interface is not in promiscuous mode, it senses network traffic that it sends and receives (rather than all network traffic).</p> <p>Note The 5-minute period referenced in the command output is a load interval that is configurable under the interface. The default value is 5 minutes.</p> <p>Note The 5-minute input should be used only as an approximation of traffic per second during a given 5-minute period. This rate is exponentially weighted average with a time constant of 5 minutes. A period of four time constants must pass before the average will be within two percent of the instantaneous rate of a uniform stream of traffic over that period.</p>
packets input	Number of packets received on the interface that were successfully delivered to higher layers.
bytes input	Total number of bytes successfully received on the interface.

Field	Description
total input drops	Total number of packets that were dropped after they were received. This includes packets that were dropped due to configured quality of service (QoS) or access control list (ACL) policies. This does not include drops due to unknown Layer 3 protocol.
drops for unrecognized upper-level protocol	Total number of packets that could not be delivered because the necessary protocol was not configured on the interface.
Received broadcast packets	Total number of Layer 2 broadcast packets received on the interface. This is a subset of the total input packet count.
Received multicast packets	Total number of Layer 2 multicast packets received on the interface. This is a subset of the total input packet count.
runts	Number of received packets that were too small to be handled. This is a subset of the input errors count.
giants	Number of received packets that were too large to be handled. This is a subset of the input errors count.
throttles	Number of packets dropped due to throttling (because the input queue was full).
parity	Number of packets dropped because the parity check failed.
input errors	Total number of received packets that contain errors and hence cannot be delivered. Compare this to total input drops, which counts packets that were not delivered despite containing no errors.
CRC	Number of packets that failed the CRC check.
frame	Number of packets with bad framing bytes.
overrun	Number of overrun errors experienced by the interface. Overruns represent the number of times that the receiver hardware is unable to send received data to a hardware buffer because the input rate exceeds the receiver's ability to handle the data.
ignored	Total number of ignored packet errors. Ignored packets are those that are discarded because the interface hardware does not have enough internal buffers. Broadcast storms and bursts of noise can result in an increased number of ignored packets.
abort	Total number of abort errors on the interface.

Field	Description
packets output	Number of packets received on the interface that were successfully delivered to higher layers.
bytes output	Total number of bytes successfully received on the interface.
total output drops	Number of packets that were dropped before being transmitted
Received broadcast packets	Number of Layer 2 broadcast packets transmitted on the interface. This is a subset of the total input packet count.
Received multicast packets	Total number of Layer 2 multicast packets transmitted on the interface. This is a subset of the total input packet count.
output errors	Number of times that the receiver hardware was unable to handle received data to a hardware buffer because the input rate exceeded the receiver's ability to handle the data.
underruns	Number of underrun errors experienced by the interface. Underruns represent the number of times that the hardware is unable to transmit data to a hardware buffer because the output rate exceeds the transmitter's ability to handle the data.
applique	Number of applique errors.
resets	Number of times that the hardware has been reset. The triggers and effects of this event are hardware-specific.
output buffer failures	Number of times that a packet was not output from the output hold queue because of a shortage of MEMD shared memory.
output buffers swapped out	Number of packets stored in main memory when the output queue is full; swapping buffers to main memory prevents packets from being dropped when output is congested. The number is high when traffic is bursty.
carrier transitions	Number of times the carrier detect (CD) signal of a serial interface has changed state.

show interfaces