Global Interface Commands

This module describes the global command line interface (CLI) commands for configuring interfaces on the Cisco NCS 5000 Series Router.

For detailed information about Global interfaces concepts, configuration tasks, and examples, refer to the *Interface and Hardware Component Configuration Guide for Cisco NCS 5000 Series Routers*

bandwidth (global)

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

bandwidth *rate*

Syntax Description

rate	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 6.0	This command was introduced.

Usage Guidelines

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 TenGigE Ethernet interface:

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

bundle wait-while

To specify the duration of the wait-while timer for a bundle, use the **bundle wait-while**bundle wait-while command in the bundle interface configuration mode. To disable waiting, use the **no**form of the command.

bundle wait-while *time*
nobundle wait-while *time*

Syntax Description

time	Wait-while time in milliseconds, range is 0-2000.

Command Default

2000 milliseconds

Command Modes

Bundle Interface Configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

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	Operation
bundle	read, write
interface	read, write

In the following example, the while-while time is configured for 20 seconds:

```
RP/0/(config-if)bundle wait-while 20
```

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   XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines   No specific guidelines impact the use of this command.

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		
half-life	(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.
reuse	(Optional) Penalty value below which a stable interface is unsuppressed. Range is from 1 through 20000. Default value is 750.
suppress	(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.
max-suppress-time	(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 6.0	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
```

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	*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 interfaces are configured

Command Modes

XR Config mode

Command History

Release	Modification
Release 6.0	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).

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
```

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	Unique MAC address used to identify the system in LACP negotiations.	
priority	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 mode

Command History

Release	Modification
Release 6.0	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 an interface, use the **mtu** command in interface configuration mode.

To return 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
- Tunnel—1500 bytes
- Loopback—1514 bytes

Command Modes

Interface configuration

Command History

Release	Modification
Release 6.0	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 Ethernet interface is the Layer 3 datagram plus 14 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.

mtu

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.

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        Intf        LineP         Encap  MTU      BW
            Name        State       State         Type  (byte)   (Kbps)
--------------------------------------------------------------------------------
            Nu0         up          up            Null   1500    Unknown
    TenGigE0/6/0/0      up          up            HDLC   4474    2488320
    TenGigE0/6/0/1      up          up            HDLC   4474    2488320
    TenGigE0/6/0/2      admin-down  admin-down    HDLC   4474    2488320
    TenGigE0/6/0/3      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 0/6/0/0
RP/0/RP0/CPU0:router(config-if)# mtu 1000
```

replace

To substitute any configuration in the router with new settings, use the **replace** command in XR Config mode.

replace **interface** *interface_name_before* **with** *interface_name_after*

replace **pattern** *string_before* **with** *string_after* ［ **dry-run** ］

Syntax Description		
	interface	Specifies the details of interface configuration replacement follows.
	interface_name_before	Specifies the name of an interface in the router that you want to replace.
		For more information about the syntax for the router, use the question mark (**?**) online help function.
	interface_name_after	Specifies the new interface name that replaces the current interface name specified in the *interface_name_before* variable.
		For more information about the syntax for the router, use the question mark (**?**) online help function.
	pattern	Specifies that the details of string replacement follow.
	string_before	Specify the configuration string in the router that you need to replace. The *string_before* can be any regular expression that specifies a match pattern in text.
		Note You must specify the *string_before* in a single quote.
	string_after	Specify the new string that replaces the configuration matching the *string_before* variable.
		Note You must specify the *string_after* in a single quote.
	dry-run	Displays the configuration after the pattern replacement without preparing the config changes for a commit. This option facilitates verifying the pattern replacement changes and provides an extra layer of protection to avoid accidentally committing unwanted configuration changes.

Command Default No default behavior or values.

Command Modes XR Config

Command History

Release	Modification
Release 7.0.1	This command was introduced.

Usage Guidelines No specific guidelines impact the use of this command.

Task ID

Task ID	Operations
interface	read, write

Examples

The following example shows how to use the **replace** command:

```
Router# config
Router(config)# replace interface gigabitEthernet 0/0/0/0 with loopback 450
Loading.
4 bytes parsed in 1 sec (3)bytes/sec
```

```
Router# config
Router(config)# replace pattern '10\.20\.30\.40' with '100.200.250.225'
Loading.
232 bytes parsed in 1 sec (230)bytes/sec
```

Examples

The following example details configuration changes on using the **replace** command:

Original Configuration:

```
Router(config-ospf-ar-if)#show configuration
Building configuration...
!! IOS XR Configuration 0.0.0
interface GigabitEthernet0/0/0/0
 description first
 ipv4 address 10.20.30.40 255.255.0.0
 shutdown
!
router ospf 10
 cost 100
 area 200
  cost 200
  interface GigabitEthernet0/0/0/0
   transmit-delay 5
  !
```

Using **replace** command:

```
Router(config-ospf-ar-if)# replace interface gigabitEthernet 0/0/0/0 with loopback 450

Building configuration...
Loading.
232 bytes parsed in 1 sec (230)bytes/sec
```

Configuration changes on using **replace** command:

```
Router(config-ospf-ar-if)#show configuration
Building configuration...
!! IOS XR Configuration 0.0.0
interface Loopback450
 description first
 ipv4 address 10.20.30.40 255.255.0.0
 shutdown
```

```
 !
no interface GigabitEthernet0/0/0/0
router ospf 10
 area 200
  interface Loopback450
   transmit-delay 5
  !
  no interface GigabitEthernet0/0/0/0
```

Examples

The following example shows how to use the **dry-run** option in the **replace** command:

```
Router# config
Router(config)# replace pattern 'vrf thr' with 'vrf three' dry-run
no vrf thr
vrf three
 address-family ipv4 unicast
  import route-target
   65321:3
  !
  export route-target
   65321:3
  !
 !
exit
router static
 no vrf thr
 vrf three
  address-family ipv4 unicast
   192.168.3.0/24 vrf one 192.168.1.1
   192.168.3.0/24 vrf two 192.168.2.2
  !
exit
end
Router(config)# commit
No configuration changes to commit.
```

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

XR EXEC mode

Command History

Release	Modification
Release 6.0	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:

- L3 capsulations (for example ipv4, ipv6). These may be brought down if another link has a conflicting IP address configured.

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 HundredGigE 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
```

Note When dampening is configured on an interface it is also applied independently to all capsulations 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 capsulation 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.

Field	Description
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.

show interfaces

To display statistics, state and other information such as mac address etc. for all interfaces configured on the router or for a specific node, use the **show interfaces** command in XR EXEC mode.

show interfaces [**summary** | [*type interface-path-id* | **all** | **local**] [**non-dynamic**] [**brief** | **description** | **detail** | **sparse** | **accounting**]] [**location** *node-id*]

Syntax Description

type	(Optional) Specifies the type of interface for which you want to display statistics. For more information, use the question mark (**?**) online help function.
interface-path-id	Physical interface or virtual interface.
	Note Use the **show interfaces brief** command to see a list of all interfaces 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 *node-id*	(Optional) Displays information about all interfaces on the specified node. The *node-id* argument is entered in the *rack/slot/module* 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.
sparse	(Optional) Displays sparse per interface information excluding interface statistics.
non-dynamic	(Optional) Displays interface information excluding dynamic interfaces.

Command Default

By default, executing **show interface** command without any option works similar as **show interface all**, and it displays the information for all interfaces in the system.

Command Modes

XR EXEC mode

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines

The **show interfaces** command displays statistics, state and other information such as mac address etc. for the network interfaces.

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 Executing **show interfaces** command without filters obtains and displays interface statistics for all interfaces. Hence, it is recommended to execute the command with filters to select specific interfaces or interface types of interest for a faster response. Executing **show interfaces** with options **brief** or **sparse** excludes interface statistics thus providing a faster response.

✎

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.

```
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:

```
Router# show interfaces brief

Thu May  6 06:30:55.797 DST

            Intf      Intf      LineP              Encap  MTU      BW
            Name      State     State              Type   (byte)  (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
```

This example shows the output from the **show interfaces** *interface-path-id* **sparse** command.

```
Router#show interfaces gigabitEthernet 0/1/0/0 sparse
Wed Apr 5 18:05:34.000 IST

Interface name : GigabitEthernet0/1/0/0
Operational state : administratively down
Admin state : administratively down
MAC address : 02:7d:42:e9:bd:36
Burned In Address : 027d.42e9.bd36
IPv4 Address : 2.2.2.2/24
Max. Bandwidth (Kbit) : 1000000
Effective Bandwidth (Kbit) : 1000000
MTU (in bytes) : 1514
Duplexity : Full-duplex
Link type : force-up
```

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	Displays the number of times the interface has changed the state. **Note** • 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.
line protocol state	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. **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).
Hardware	Displays the current hardware type.
address is *n.n.n.n/n*	Displays the Layer 2 address (MAC address for Ethernet interfaces). **Note** Enter the **mac-address** command to configure the hardware address.

Field	Description
bia	Displays the burned-in address (BIA) for the interface. The BIA is the default L2 (MAC) address for the interface. **Note** The BIA is not configurable.
description	Displays the user-defined string that is associated with the interface. **Note** Enter the **description** command to configure the description associated with the interface.
Internet address	Displays the Layer 3 (IP) address for the interface. **Note** Enter the **ipv4 address** command to configure the internet address for the interface.
MTU	Displays the maximum transmission unit (MTU) for the interface. The MTU is the maximum packet size that can be transmitted over the interface. **Note** The MTU field indicates the interface MTU. Enter the **mtu** command to configure a lower MTU value at the Layer 3 level.
BW	Displays the bandwidth of the interface in kbps.
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.

Field	Description
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 *keepalive* 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.
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.

Field	Description
ARP timeout	ARP timeout in *hours*:*mins*:*secs*. 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 *hours*:*mins*:*secs*.
Input rate	Average number of bits and packets received per second during the load-interval. If the interface is not in promiscuous mode, it senses network traffic that it sends and receives (rather than all network traffic). **Note** Load duration is based on load-interval configured under the interface. The default load duration is 5 minutes, if load-interval is not configured under the interface. **Note** The input rate should be used only as an approximation of traffic per second during a given load duration. This rate is exponentially weighted average with a time constant of load duration. 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.
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. **Note** This does not include FCS bytes.
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.

Field	Description
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.
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. **Note** This does not include FCS bytes.
total output drops	Number of packets that were dropped before being transmitted

Field	Description
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-specifc.
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.

shutdown (global)

To disable an interface (to force an interface to be administratively down), use the **shutdown** command in interface configuration mode. To enable an interface that has been shut down, use the **no** form of this command.

shutdown

Syntax Description   This command has no keywords or arguments.

Command Default   The interface is enabled by default and is disabled only when shutdown is configured.

Note   When you add an interface to the system, or when all the configuration for an interface is lost or deleted, the interface is put in the shutdown state by the system adding the interface.

Command Modes   Interface configuration

Command History

Release	Modification
Release 6.0	This command was introduced.

Usage Guidelines   Use the **shutdown** command to move the state of an interface to administratively down, which stops traffic flowing through the interface. This state does not stop other action from happening on the interface such as changes in configuration, protocols, capsulations, and so forth.

The **shutdown** command also marks the interface as unavailable. To check whether the state of an interface is down, use the **show interfaces** command in XR EXEC mode, which displays the current state of the interface. An interface that has been shut down is shown as administratively down in the display from the **show interfaces** command.

Task ID

Task ID	Operations
interface	read, write

Examples   In this example, TenGigE interface 0/4/0/2 is turned off:

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 0/4/0/2
RP/0/RP0/CPU0:router(config-if)# shutdown
```