



Congestion Avoidance

This chapter provides information about slow devices that cause congestion (slow-drain devices) in a network and how to avoid slow-drain devices.

- [Finding Feature Information, on page 2](#)
- [Feature History for Congestion Avoidance, on page 3](#)
- [Information About Congestion Avoidance, on page 4](#)
- [Configuring Congestion Avoidance, on page 5](#)
- [Configuration Examples for Congestion Avoidance, on page 12](#)
- [Verifying Configuration Examples for Slow Drain Device Detection and Congestion Avoidance, on page 14](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the New and Changed chapter or the Feature History table below.

Feature History for Congestion Avoidance

There are no new features for this chapter in the Cisco MDS 9000 Series NX-OS Interfaces Configuration Guide for Cisco MDS NX-OS Release 7.3(0)D1(1).

Information About Congestion Avoidance

Congestion avoidance focuses on minimizing or completely avoiding the congestion that results from frames being queued to congested ports.

Cisco MDS switches have multiple features designed to void congestion in SAN:

- Congestion-drop timeout threshold (Fibre Channel and FCoE): The congestion-drop timeout threshold determines the amount of time a queued Fibre Channel or FCoE frame will stay in the switch awaiting transmission. Once the threshold is reached the frame is discarded as a *timeout drop*. The lower the value the quicker these queued frames are dropped and the result buffer freed. This can relieve some back pressure in the switch, especially on ISLs. By default it is 500 ms but can be configured as low as 200 ms in 1 ms increments. It is configured using the **system timeout congestion-drop** (Fibre Channel) and **system timeout fcoe congestion-drop** (FCoE) commands.
- No-credit-drop timeout threshold (Fibre Channel only): No-credit-drop timeout threshold is used to time when a Fibre Channel port is at zero Tx credits. Once a Fibre Channel port hits zero Tx credits the timer is started. If the configured threshold is reached then all frames queued to that port will be dropped regardless of their actual age in the switch. Furthermore, as long as the port remains at zero Tx credits, all newly arriving frames are immediately dropped. This can have a dramatic effect on relieving congestion especially on upstream ISLs. This allows unrelated flows to move continuously. This is off by default. If configured, it should be set to a value that is lower than the configured (or defaulted) Fibre Channel congestion-drop timeout. It is configured via the **system timeout no-credit-drop** command. The no-credit timeout functionality is only used for edge ports because these ports are directly connected to the slow-drain devices.
- Pause-drop timeout threshold (FCoE only): Pause-drop timeout threshold is used to time when a FCoE port is in a continuous state of Rx pause (unable to transmit). After an FCoE port receives a PFC pause with a non-zero quanta, the timer is started. If the port continues to receive PFC pauses with a non-zero quanta such that it remains in the Rx pause state continuously for the pause-drop threshold, then all frames queued to that port will be dropped regardless of their actual age in the switch. Furthermore, as long as the port remains in a Rx pause state, all newly arriving frames are immediately dropped. This can have a dramatic effect on relieving congestion especially on the upstream ISLs. This allows unrelated flows to move continuously. This is on by default with a value of 500 ms. If configured, it should be set to a value that is lower than the configured (or defaulted) FCoE congestion-drop timeout. It is configured via the **system timeout fcoe pause-drop** commands (available from Cisco MDS NX-OS Release 8.2(1) onwards). The FCoE pause-drop timeout functionality is only used for edge ports, because these ports are directly connected to the slow-drain devices.
- Port monitor with portguard actions of flap and error disable: For more information, see the [Port Monitor](#) section.



Note

The *no-credit timeout* functionality is only used for edge ports because these ports are directly connected to the slow-drain devices. The no-credit timeout functionality is not supported on Generation 1 modules.

Configuring Congestion Avoidance

Slow-drain devices are devices that do not accept frames at the configured rate. The presence of these slow-drain devices leads to traffic congestion in the Fibre Channel or Fibre Channel over Ethernet (FCoE) fabric. This traffic congestion can affect the unrelated flows in the fabric that use the same ISLs for its traffic as the slow-drain device. This is true even though the destination devices are not slow-drain devices.

From Cisco MDS NX-OS Release 4.2(1), slow-drain device detection and congestion avoidance is supported on all Fibre Channel switching modules.

From Cisco MDS NX-OS Release 5.2(1), slow-drain device detection and congestion avoidance is supported on all FCoE switch modules.

From Cisco MDS NX-OS Release 5.2(1), slow-drain detection and congestion avoidance functionality for edge ports was enhanced.

From Cisco MDS NX-OS Release 6.2(13), many new features were added to slow-drain device detection and congestion avoidance.

From Cisco MDS NX-OS Release 8.1(1), slow-drain device detection and congestion isolation feature was added.

Multiple features are available on the Cisco MDS 9000 Series Multilayer Switches to detect slow drain and avoid the resulting effects.

Configuring the Congestion Drop Timeout Value for FCoE

When an FCoE frame takes longer than the congestion drop timeout period to be transmitted by the egress port, the frame is dropped. This dropping of frames is useful in controlling the effect of slow egress ports that are paused almost continuously (long enough to cause congestion), but not long enough to trigger the pause timeout drop. Frames dropped due to the congestion drop threshold are counted as egress discards against the egress port. Egress discards release buffers in the upstream ingress ports of a switch, allowing the unrelated flows to move continuously through the ports.

The congestion drop timeout value is 500 ms by default for all port types. We recommend that you retain the default timeout for core ports, and consider configuring a lower value for edge ports. The congestion drop timeout value should be equal to or greater than the pause drop timeout value for that port type.

To configure the congestion drop timeout value for FCoE, perform these steps:

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Depending on the Cisco MDS NX-OS release version you are using, use one of the following commands to configure the system-wide FCoE congestion drop timeout, in milliseconds, for core or edge ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases

```
switch(config)# system default interface congestion timeout milliseconds mode {core | edge}
```

The FCoE congestion drop timeout range is from 100 to 1000 ms.

Note To prevent premature packet drops, the minimum value recommended for FCoE congestion drop timeout is 200 ms.

(Optional) Depending on the Cisco MDS NX-OS release version you are using, use one of the following commands to revert to the default FCoE congestion drop timeout value of 500 milliseconds:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases
`switch(config)# no system default interface congestion timeout milliseconds mode {core | edge}`

Configuring Pause Drop Timeout for FCoE

When an FCoE port is in a state of continuous pause during the FCoE pause drop timeout period, all the frames that are queued to that port are dropped immediately. As long as the port continues to remain in the pause state, the newly arriving frames destined for the port are dropped immediately. These drops are counted as egress discards on the egress port, and free up buffers in the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

To reduce the effect of a slow-drain device on unrelated traffic flows, configure a lower-pause drop timeout value than the congestion frame timeout value, for edge ports. This causes the frames that are destined for a slow port to be dropped immediately after the FCoE pause drop timeout period has occurred, rather than waiting for the congestion timeout period to drop them.

By default, the FCoE pause drop timeout is enabled on all ports and the value is set to 500 ms. We recommend that you retain the default timeout core ports and consider configuring a lower value for edge ports.

To configure the FCoE pause drop timeout value, perform these steps:

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to configure the system-wide FCoE pause drop timeout value, in milliseconds, for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases
`switch(config)# system default interface pause timeout milliseconds mode {core | edge}`

The range is from 100 to 500 milliseconds.

(Optional) Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to enable the FCoE pause drop timeout to the default value of 500 milliseconds for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases
`switch(config)# system default interface pause mode {core | edge}`

(Optional) Depending on the Cisco MDS NX-OS release version that you are using, use one of the following commands to disable the FCoE pause drop timeout for edge or core ports:

- Cisco MDS NX-OS Release 8.1(1) and earlier releases
`switch(config)# no system default interface pause mode {core | edge}`

Configuring the Congestion Drop Timeout Value for Fibre Channel

When a Fibre Channel frame takes longer than the congestion timeout period to be transmitted by the egress port, the frame is dropped. This option of the frames being dropped is useful for controlling the effect of slow egress ports that lack transmit credits almost continuously; long enough to cause congestion, but not long enough to trigger the no-credit timeout drop. These drops are counted as egress discards on the egress port, and release buffers into the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

By default, the congestion timeout value is 500 ms for all port types. We recommend that you retain the default timeout for core ports and configure a lower value (not less than 200 ms) for edge ports. The congestion timeout value should be equal to or greater than the no-credit frame timeout value for that port type.

To configure the congestion frame timeout value for the Fibre Channel, perform these steps:

-
- Step 1** Enter configuration mode:
- ```
switch# configure terminal
```
- Step 2** Configure the Fibre Channel congestion drop timeout value, in milliseconds, for the specified port type:
- ```
switch(config)# system timeout congestion-drop milliseconds mode {E | F}
```
- The range is 100-500 ms in multiples of 10.
- Step 3** (Optional) Revert to the default value for the congestion timeout for the specified port type:
- ```
switch(config)# system timeout congestion-drop default mode {E | F}
```
- 

## Configuring the No-Credit Frame Timeout Value for Fibre Channel

When a Fibre Channel egress port has no transmit credits continuously for the no-credit timeout period, all the frames that are already queued to that port are dropped immediately. As long as the port remains in this condition, newly arriving frames destined for that port are dropped immediately. These drops are counted as egress discards on the egress port, and release buffers in the upstream ingress ports of the switch, allowing unrelated flows to continue moving through them.

No-credit dropping can be enabled or disabled. By default, frame dropping is disabled and the frame timeout value is 500 ms for all port types. We recommend that you retain the default frame timeout for core ports and configure a lower value (300 ms) for edge ports. If the slow-drain events continue to affect unrelated traffic flows, the frame timeout value for the edge ports can be lowered to drop the previous slow-drain frames. This frees the ingress buffers for frames of unrelated flows, thus reducing the latency of the frames through the switch.



### Note

- The no-credit frame timeout value should always be less than the congestion frame timeout for the same port type, and the edge port frame timeout values should always be lower than the core port frame timeout values.
  - The slow-port monitor delay value should always be less than the no-credit frame timeout value for the same port type.
-

For pre-16-Gbps-capable modules and systems, the no-credit timeout value can be 100 to 500 ms in multiples of 100 ms. On these systems, the no-credit condition is checked only at 100-ms intervals. At this point, if the no-credit condition exists, dropping starts. Depending on the timing of the actual onset of the no-credit condition, the task of checking port dropping can be delayed by up to 100 ms later than the configured value.

On 16-Gbps and later modules and systems, the no-credit timeout value can be 1 to 500 ms in multiples of 1 ms. Dropping starts immediately after the no-credit condition comes into existence for the configured timeout value.

To configure the no-credit timeout value, perform these steps:

---

**Step 1** Enter configuration mode:

```
switch# configure terminal
```

**Step 2** Specify the no-credit timeout value:

```
switch(config)# system timeout no-credit-drop milliseconds mode F
```

(Optional) Revert to the default no-credit timeout value (500 ms):

```
switch(config)# system timeout no-credit-drop default mode F
```

(Optional) Disable the no-credit drop timeout value:

```
switch(config)# no system timeout no-credit-drop mode F
```

---

## Configuring the Slow-Port Monitor Timeout Value for Fibre Channel

The slow-port monitor functionality is similar to the no-credit frame timeout and drop functionality, except that it does not drop frames; it only logs qualifying events. When a Fibre Channel egress port has no transmit credits continuously for the slow-port monitor timeout period, the event is logged. No frames are dropped unless the no-credit frame timeout period is reached and no-credit frame timeout drop is enabled. If the no-credit frame timeout drop is not enabled, no frames are dropped until the congestion frame timeout period is reached.

Slow-port monitoring is implemented in the hardware, with the slow-port monitor functionality being slightly different in each generation of hardware. The 8-Gbps modules report a single slow-port monitor event for each 100-ms window in which the slow-port monitor threshold has crossed one or more times. They do not have the ability to report the exact number of slow-port events. The advanced 8-Gbps and 16-Gbps modules and switches are not restricted and can detect each instance of the slow-port monitor threshold being crossed. The slow-port monitor log is updated at 100-ms intervals. A log entry for a slow port on an 8-Gbps module can increment by a maximum of one. A log for a slow-port event on an advanced 8-Gbps or 16-Gbps module or system increments the exact number of times the threshold is reached.

Slow port monitor can also generate an alert and syslog message via port monitor.

To configure the slow-port monitor timeout value, perform these steps:

---

**Step 1** Enter configuration mode:

```
switch# configure terminal
```



**Step 2** Specify the slow-port monitor timeout value:

```
switch(config)# system timeout slowport-monitor milliseconds mode {E | F}
```

Valid values for the slow-port monitor timeout are:

- 32-Gbps and 16-Gbps modules or switches—1 to 500 ms in 1-ms increments.
- 8-Gbps and advanced 8-Gbps modules—1 to 100 ms in 1-ms increments.

(Optional) Revert to the default slow-port monitor timeout value (50 ms) for the specified port type:

```
switch(config)# system timeout slowport-monitor default mode {E | F}
```

(Optional) Disable the slow-port monitor:

```
switch(config)# no system timeout slowport-monitor default mode {E | F}
```

---

## Displaying Credit Loss Recovery Actions

When a port is at zero transmit credits for 1 full second (F ports) and 1.5 seconds (E ports), it is called a credit loss condition. Cisco MDS initiates credit loss recovery by transmitting a Link Credit Reset (LCR). If the end device responds with a Link Credit Reset Response (LCRR), the port is back at its fully agreed number of B2B credits in both directions. If an LRR is not received, the port is shut down.

When the port detects the credit loss condition and recovers, some of the following actions might occur:

1. An SNMP trap with interface details can be sent, indicating the credit loss event.
2. The port can be error disabled.
3. The port can be flapped.

When the configured threshold is exceeded, one or more of these actions can be combined together. These actions can be turned on or off depending on the situation. The Port Monitor feature provides the CLI to configure the thresholds and action.

The 1 second (F ports) and 1.5 seconds (E ports) timers that are set for the switch to initiate CLR are fixed and cannot be changed.

To verify a port monitor policy to generate SNMP alerts and take other actions in the quantity and timing of these events, perform these steps:

---

**Step 1** Display the last 10 credit loss events per interface per module:

```
switch# show process creditmon credit-loss-events [module x]
```

**Step 2** Display a chronological log of credit loss events per module:

```
switch# show process creditmon credit-loss-event-history [module x]
```

**Note** When a port sees the credit loss condition and fails to recover, the port flaps. This function is already a part of the portguard, and you can configure the supported actions using the Portguard feature.

---

## Configuring the Transmit Average Credit-Not-Available Duration Threshold and Action in Port Monitor

Cisco MDS monitors its ports that are at zero transmit credits for 100 ms or more. This is called transmit average credit-not-available duration. The Port Monitor feature can monitor this using the TX Credit Not Available counter. When the transmit average credit-not-available duration exceeds the threshold set in the port monitor policy, an SNMP trap with interface details is sent, indicating the transmit average credit not available duration event along with a syslog message. Additionally, the following events may be configured:

- A warning message is displayed.
- The port can be error disabled.
- The port can be flapped.

The Port Monitor feature provides the CLI to configure the thresholds and actions. The threshold configuration is configured as a percentage of the interval. The thresholds can be 0 to 100 percent in multiples of 10, and the interval can be 1 second to 1 hour. The default is 10 percent of a 1-second interval and generates a SNMP trap and syslog message when the transmit-average-credit-not-available duration hits 100 ms.

The following edge port monitor policy is active by default. No port monitor policy is enabled for core ports by default.

```
switch# show port-monitor slowdrain
```

```
Policy Name : slowdrain
Admin status : Not Active
Oper status : Not Active
Port type : All Access Ports
```

| Counter                 | Threshold | Interval | Rising<br>Threshold | event | Falling<br>Threshold | event | PMON        | Portguard |
|-------------------------|-----------|----------|---------------------|-------|----------------------|-------|-------------|-----------|
| Credit Loss Reco        | Delta     | 1        | 1                   | 4     | 0                    | 4     | Not enabled |           |
| TX Credit Not Available | Delta     | 1        | 10%                 | 4     | 0%                   | 4     | Not enabled |           |

The following example shows how to configure a new policy similar to the slowdrain policy with the tx-credit not available threshold set to 200 ms:



**Note** The default *slowdrain* port monitor policy cannot be modified and a new policy needs to be configured.

```
switch# configure
switch(config)# port-monitor name slowdrain_tx200ms
switch(config-port-monitor)# logical-type edge
switch(config-port-monitor)# no monitor counter all
switch(config-port-monitor)# monitor counter credit-loss-reco
switch(config-port-monitor)# monitor counter tx-credit-not-available
switch(config-port-monitor)# counter tx-credit-not-available poll-interval 1 delta
switch(config-port-monitor)# rising-threshold 20 event 4 falling-threshold 0 event 4
switch(config-port-monitor)# no port-monitor activate slowdrain
switch(config)# port-monitor activate slowdrain_tx200ms
switch(config)# end
```

```
switch# show port-monitor active
```

Policy Name : slowdrain\_tx200ms  
 Admin status : Active  
 Oper status : Active  
 Port type : All Edge Ports

| Counter             | Threshold | Interval | Threshold | Rising event | Threshold | Falling event | Portguard   | Warning     | PMON        |
|---------------------|-----------|----------|-----------|--------------|-----------|---------------|-------------|-------------|-------------|
| Credit Loss         |           |          |           |              |           |               |             |             |             |
| Reco Delta          | 1         | 1        | 4         | 0            | 4         | Not enabled   | Not enabled |             |             |
| TX Credit           |           |          |           |              |           |               |             |             |             |
| Not Available Delta |           |          | 1         | 20%          | 4         | 0%            | 4           | Not enabled | Not enabled |

# Configuration Examples for Congestion Avoidance



## Note

- From Cisco MDS NX-OS Release 8.1(1), mode E is treated as logical-type core and mode F is treated as logical-type edge.
- The port *Logical type* is displayed as the *Port type*.

This example shows how to check the currently active port-monitor policy:

```
switch# show port-monitor active
Policy Name : sample
Admin status : Active
Oper status : Active
Port type : All Ports

```

| Counter       | Threshold | Interval | Rising<br>Threshold | event | Falling<br>Threshold | event | Warning<br>Threshold | PMON<br>Portguard |
|---------------|-----------|----------|---------------------|-------|----------------------|-------|----------------------|-------------------|
| Link          |           |          |                     |       |                      |       |                      |                   |
| Loss          | Delta     | 10       | 6                   | 4     | 5                    | 4     | Not enabled          | Flap              |
| Sync          |           |          |                     |       |                      |       |                      |                   |
| Loss          | Delta     | 60       | 5                   | 4     | 1                    | 4     | Not enabled          | Not enabled       |
| Signal        |           |          |                     |       |                      |       |                      |                   |
| Loss          | Delta     | 60       | 5                   | 4     | 1                    | 4     | Not enabled          | Not enabled       |
| Invalid       |           |          |                     |       |                      |       |                      |                   |
| Words         | Delta     | 60       | 1                   | 4     | 0                    | 4     | Not enabled          | Not enabled       |
| Invalid       |           |          |                     |       |                      |       |                      |                   |
| CRC's         | Delta     | 30       | 20                  | 2     | 10                   | 2     | Not enabled          | Not enabled       |
| State         |           |          |                     |       |                      |       |                      |                   |
| Change        | Delta     | 60       | 5                   | 4     | 0                    | 4     | Not enabled          | Not enabled       |
| TX            |           |          |                     |       |                      |       |                      |                   |
| Discards      | Delta     | 60       | 200                 | 4     | 10                   | 4     | Not enabled          | Not enabled       |
| LR RX         | Delta     | 60       | 5                   | 4     | 1                    | 4     | Not enabled          | Not enabled       |
| LR TX         | Delta     | 60       | 5                   | 4     | 1                    | 4     | Not enabled          | Not enabled       |
| Timeout       |           |          |                     |       |                      |       |                      |                   |
| Discards      | Delta     | 60       | 200                 | 4     | 10                   | 4     | Not enabled          | Not enabled       |
| Credit        |           |          |                     |       |                      |       |                      |                   |
| Loss Reco     | Delta     | 1        | 1                   | 4     | 0                    | 4     | Not enabled          | Not enabled       |
| TX Credit     |           |          |                     |       |                      |       |                      |                   |
| Not Available | Delta     | 3        | 40%                 | 4     | 2%                   | 4     | Not enabled          | Not enabled       |
| RX Datarate   | Delta     | 60       | 80%                 | 4     | 20%                  | 4     | Not enabled          | Not enabled       |
| TX Datarate   | Delta     | 60       | 80%                 | 4     | 20%                  | 4     | Not enabled          | Not enabled       |
| ASIC Error    |           |          |                     |       |                      |       |                      |                   |
| Pkt to xbar   | Delta     | 300      | 5                   | 4     | 0                    | 4     | Not enabled          | Not enabled       |

```

```

This example shows how to configure the Fibre Channel congestion drop timeout value of 210 milliseconds for mode E:

```
switch# configure terminal
switch(config)# system timeout congestion-drop 210 mode E
```

This example shows how to configure the Fibre Channel congestion drop timeout to the default value of 200 milliseconds for mode E:

```
switch# configure terminal
switch(config)# system timeout congestion-drop default mode E
```

This example shows how to configure the Fibre Channel no-credit drop timeout value of 100 milliseconds for mode F:

```
switch# configure terminal
switch(config)# system timeout no-credit-drop 100 mode F
```

This example shows how to configure the Fibre Channel no-credit drop timeout to the default value of 500 milliseconds for mode F:

```
switch# configure terminal
switch(config)# system timeout no-credit-drop default mode F
```



---

**Note** The no-credit drop timeout value is disabled by default.

---

This example shows how to disable the Fibre Channel no-credit drop timeout for mode F:

```
switch# configure terminal
switch(config)# no system timeout no-credit-drop default mode F
```

This example shows how to configure the Fibre Channel hardware slowport monitoring value of 10 milliseconds for mode F:

```
switch# configure terminal
switch(config)# system timeout slowport-monitor 10 mode F
```

This example shows how to configure the Fibre Channel hardware slowport monitoring to the default value of 50 milliseconds for mode F:

```
switch# configure terminal
switch(config)# system timeout slowport-monitor default mode F
```



---

**Note** The slowport monitoring value is disabled by default.

---

This example shows how to disable the Fibre Channel hardware slowport monitoring for mode F:

```
switch# configure terminal
switch(config)# no system timeout slowport-monitor default mode F
```

# Verifying Configuration Examples for Slow Drain Device Detection and Congestion Avoidance

## Displaying the Congestion Frame Timeout Value for FCoE

The following commands display the congestion frame timeout value for FCoE (pause counter log and pause event log, respectively, with timeout value):

(Optional) Display the pause counter log with time-stamp information:

```
switch# show logging onboard flow-control pause-count
```

(Optional) Display the pause event log with time-stamp information:

```
switch# show logging onboard flow-control pause-events
```

## Displaying the Pause Frame Timeout Value for FCoE

The following commands display the pause frame timeout value for FCoE:

(Optional) Display the pause counter log with time-stamp information:

```
switch# show logging onboard flow-control pause-count
```

(Optional) Display the pause counters per module per interface with time-stamp information:

```
switch# show logging onboard flow-control pause-events
```

(Optional) Display the timeout drops per module per interface with time-stamp information:

```
switch# show logging onboard flow-control timeout-drops [module x] [last mm minutes] [last hh hours] [last dd days]
```

## Displaying the Congestion Drop Timeout Value for Fibre Channel

The following command displays the timeout drops per module per interface with time-stamp information:

```
switch# show logging onboard flow-control timeout-drops [module x] [last mm minutes] [last hh hours] [last dd days]
```

## Displaying the No-Credit Frame Timeout Value for Fibre Channel

The following command displays various error statistics per module per interface with time-stamp information:

```
switch# show logging onboard [module x] [starttime mm/dd/yy-hh:mm:ss] error-stats
```

The following counters indicate that the no-credit drop threshold has been reached:

- FCP\_CNTR\_FORCE\_TIMEOUT\_ON
- AK\_FCP\_CNTR\_FORCE\_TIMEOUT\_ON
- FCP\_SW\_CNTR\_FORCE\_TIMEOUT\_ON

The following counters indicate that a credit has been received on the interface, and the port no longer drops packets because of the no-credit drop condition:

- FCP\_CNTR\_FORCE\_TIMEOUT\_OFF
- AK\_FCP\_CNTR\_FORCE\_TIMEOUT\_OFF
- FCP\_SW\_CNTR\_FORCE\_TIMEOUT\_OFF

## Verifying Congestion Detection and Avoidance

The following commands display slow-port monitor events:



**Note** These commands are applicable for both supervisor and module prompts.

Display slow-port monitor events per module:

```
switch# show process creditmon slowport-monitor-events [module x [port y]]
```

Display the slow-port monitor events on the Onboard Failure Logging (OBFL):

```
switch# show logging onboard slowport-monitor-events
```



**Note** The slow-port monitor events are logged periodically into the OBFL.

The following example displays the credit monitor or output of the **creditmon slow-port monitor-events** command for the 16-Gbps and 32-Gbps modules and switches:

```
switch# show process creditmon slowport-monitor-events
```

```

Module: 06 Slowport Detected: YES
=====
Interface = fc6/3

admin	slowport	oper	Timestamp
delay	detection	delay	
(ms)	count	(ms)	

1	46195	1	1. 10/14/12 21:46:51.615
1	46193	50	2. 10/14/12 21:46:51.515
1	46191	50	3. 10/14/12 21:46:51.415
1	46189	50	4. 10/14/12 21:46:51.315
1	46187	50	5. 10/14/12 21:46:51.215
1	46185	50	6. 10/14/12 21:46:51.115
1	46183	50	7. 10/14/12 21:46:51.015
1	46181	50	8. 10/14/12 21:46:50.915
1	46179	50	9. 10/14/12 21:46:50.815
1	46178	50	10. 10/14/12 21:46:50.715

```



**Note** For 16-Gbps modules, 32-Gbps modules, and Cisco MDS 9700, 9148S, 9250i, and 9396S switches, if **no-credit-drop** timeout is configured, the maximum value of **tx-slowport-oper-delay** as shown in slow-port monitor events is limited by the **no-credit-drop timeout**. So, the maximum value for **tx-slowport-oper-delay** can reach the level of the **no-credit-drop** timeout even if the actual slow-port delay from the device is higher because the frames are forcefully dropped by the hardware when **tx-slowport-oper-delay** reaches the level of the **no-credit-drop** timeout.

The following example displays the output of the **creditmon slowport-monitor-events** command for the Cisco MDS 9500 switches (8-Gbps modules):

```
switch# show process creditmon slowport-monitor-events
 Module: 04 Slowport Detected: YES
=====
Interface = fc4/13

admin	slowport	Timestamp
delay	detection	
(ms)	count	

1	194	1. 04/29/15 17:19:13.345
1	193	2. 04/29/15 17:19:13.245
1	192	3. 04/29/15 17:19:13.145
1	191	4. 04/29/15 17:19:13.045
1	190	5. 04/29/15 17:19:12.945
1	189	6. 04/29/15 17:19:12.845
1	188	7. 04/29/15 17:19:12.745
1	187	8. 04/29/15 17:19:12.645
1	186	9. 04/29/15 17:19:12.545
1	185	10. 04/29/15 17:19:12.445

=====
```



**Note** The Cisco MDS 9500 Series 8-Gbps modules can only detect whether the slow-port monitor has reached the threshold (admin delay) or not for every 100-ms polling interval. The modules cannot determine the actual length of time, whether it is higher than the threshold, or when the port is in the zero-transmit-credits-remaining condition. Additionally, the modules cannot determine if the slow-port monitor has reached the threshold (admin delay) multiple times. The modules can record only one event per 100-ms polling interval.

The following example displays output of the **creditmon slow-port-monitor-events** command for the Cisco MDS 9500 switches (advanced 8-Gbps modules):



**Note** The Cisco MDS 9500 Series advanced 8-Gbps modules utilize the transmit wait functionality to implement slow-port monitoring. Hence, **tx-slowport-oper-delay** is the total amount of time the port was in the zero-transmit-credits-remaining condition during the 100-ms polling interval. No specific duration of time is indicated.



```

switch# show process creditmon slowport-monitor-events module 1
 Module: 01 Slowport Detected: YES
=====
Interface = fc1/5

admin	slowport	txwait	Timestamp
delay	detection	oper	
(ms)	count	delay	
		(ms)	

10	888	93	1. 04/30/15 21:33:42.561
10	887	81	2. 04/30/15 21:33:42.461
10	886	76	3. 04/30/15 21:33:42.361
10	885	99	4. 04/30/15 21:33:42.261
10	884	99	5. 04/30/15 21:33:42.161

=====

```



**Note** For advanced 8-Gbps modules, the transmit-wait value does not increment after the **no-credit-drop** threshold has been reached because the frames are forcefully dropped by the hardware, and no more frames are queued for transmit. Consequently, when slow-port monitor is used in conjunction with **no-credit-drop**, the **tx-slowport-oper-delay** value, as shown in the output of the slow-port monitor events command may be lower than expected.

The following example displays the transmit-wait statistics for a particular interface for the Cisco MDS 9500 switches (advanced 8-Gbps modules and 16-Gbps modules):

```

switch# show interface fc1/1 counters
or
switch# show interface fc1/1 counters details

switch(config)# show int fc1/81 counters
fc1/81
 5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 5 minutes output rate 280288 bits/sec, 35036 bytes/sec, 15 frames/sec
5206879 frames input, 11142684612 bytes
 0 class-2 frames, 0 bytes
 5206879 class-3 frames, 11142684612 bytes
 0 class-f frames, 0 bytes
 171 discards, 175 errors, 0 CRC/FCS
 0 unknown class, 0 too long, 4 too short
2498081 frames output, 5345868788 bytes
 0 class-2 frames, 0 bytes
 2498081 class-3 frames, 5345868788 bytes
 0 class-f frames, 0 bytes
 7260927715 discards, 0 errors
 7260927715 timeout discards, 0 credit loss
 2 input OLS, 272 LRR, 0 NOS, 0 loop inits
 3 output OLS, 3 LRR, 2 NOS, 0 loop inits
 2 link failures, 0 sync losses, 0 signal losses
2498321 Transmit B2B credit transitions to zero
 275 Receive B2B credit transitions to zero
 54867361792 2.5us TxWait due to lack of transmit credits
Percentage Tx credits not available for last 1s/1m/1h/72h: 50%/50%/92%/52%
 32 receive B2B credit remaining
 0 transmit B2B credit remaining

```

```

0 low priority transmit B2B credit remaining
Last clearing of "show interface" counters :never

```

```

=====
switch(config)# show int fc1/81 counters details
fc1/81
5206879 frames, 11142684612 bytes received
0 class-2 frames, 0 bytes received
0 class-2 discards
0 F_BSY frames, 0 F_RJT frames
 generated against class-2 frames
0 port reject frames
5206879 class-3 frames, 11142684612 bytes received
0 class-f frames, 0 bytes received
171 discards, 175 errors received
7273423181 discards, 0 errors transmitted
2499069 frames, 5347983108 bytes transmitted
0 class-2 frames, 0 bytes transmitted
2499069 class-3 frames, 5347983108 bytes transmitted
171 class-3 frames discarded
0 class-f frames, 0 bytes transmitted
0 class-f frames discarded
0 multicast packets received, 0 transmitted
0 broadcast packets received, 0 transmitted
5206879 unicast packets received, 2499069 transmitted
7273423181 timeout discards, 0 credit loss
2 link failures, 0 sync losses, 0 signal losses
0 primitive sequence protocol errors
31822 invalid transmission words
0 invalid CRCs, 0 Delimiter Errors
0 address identifier errors
0 link reset received while link is active
272 link reset transmitted while link is active
2 Offline Sequence errors received
3 Offline Sequence errors transmitted
0 frames received that are shorter than
 the minimum allowable frame length
 regardless of the CRC/FCS error
0 frames received that are longer than
 the maximum frame length and also have a
 CRC/FCS error
54879203328 2.5us TxWait due to lack of transmit credits
0 frames received with length greater
 than what was agreed to in FLOGI/PLOGI
4 frames received with length less than
 the minimum indicated by the frame header
272 link reset responses received
3 link reset responses transmitted
0 non-operational sequences received
2 non-operational sequences transmitted
0 fragmented frames received
171 frames received with EOF aborts
0 unknown class frames received
0 8b10b disparity errors
0 frames discarded
0 Exchange Link Parameters switch fabric
 internal link service request failures
2499309 Transmit B2B credit transitions to zero
275 Receive B2B credit transitions to zero
0 Enhanced Inter Switch Link (EISL) frames
 discarded
0 framing errors
0 F8 type LIP sequence errors received
0 F8 type LIP sequence errors issued

```

```
0 Non F8 type LIP sequence errors received
0 Non F8 type LIP sequence errors issued
0 fec corrected blocks
0 fec uncorrected blocks
Percentage Tx credits not available for last 1s/1m/1h/72h: 50%/50%/92%/52%
```

