



Configuring Interface Buffers

This chapter provides information about interfaces buffers, its features, and how to configure the interface buffers.

- [Finding Feature Information, on page 2](#)
- [Feature History for Interface Buffers, on page 3](#)
- [Information About Interface Buffers, on page 4](#)
- [Configuring Interface Buffers, on page 19](#)
- [Configuration Examples for Interface Buffers, on page 23](#)
- [Verifying Interface Buffer Configuration, on page 24](#)
- [Troubleshooting Interface Buffer Credits, on page 26](#)

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 Interface Buffers

Table 1: Feature History for Interface Buffers


Feature Name	Release	Feature Information
Buffer-to-Buffer Credit Recovery	8.4(1)	Support for buffer-to-buffer credit recovery for NP ports.
Buffer-to-Buffer Credit Recovery	8.2(1)	Support for buffer-to-buffer credit recovery for F ports.
Enhanced Receiver Ready	8.1(1)	This feature was introduced. The following commands were introduced: <ul style="list-style-type: none">• show flow-control er_rdy• switchport vl-credit• system fc flow-control er_rdy

Information About Interface Buffers

Fibre Channel interfaces use buffer to buffer credits to ensure all packets are delivered to their destination without frame drops even if there is congestion in the network.

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a Fibre Channel link-level flow-control mechanism that ensures every frame (a Fibre Channel packet) that is sent has enough buffer space to be received. Each Fibre Channel link can be viewed as two unidirectional links each with their own set of BB_credits. During link initialization, each side informs the other side of the number of receive (Rx) BB_credits it has via the Exchange Link Parameters (ELP) and Accept (ELP) on an E port and FLOGI and Accept (FLOGI) on an F or NP port. When the Rx BB_credit number is received, it is stored as the transmit (Tx) BB_credit number. This way each side's Rx BB_credit number is the other side's Tx BB credit number for each direction on the link.



Each buffer location holds exactly one Fibre Channel frame regardless of size. As a frame is to be transmitted, the sender checks the remaining Tx BB_credit number. If it is greater than 0, a frame can be transmitted. The sender then decrements Tx BB_credit remaining number and transmits the frame. After the frame is received

and processed such that the receiver's buffer location is cleared, the receiver transmits a R_RDY primitive (BB_credit). When the BB credit is received, the sender increments its Tx BB_credit remaining number. This mechanism guarantees that a sender never transmits a frame that the receiver does not have a buffer to hold it in.

**Note**

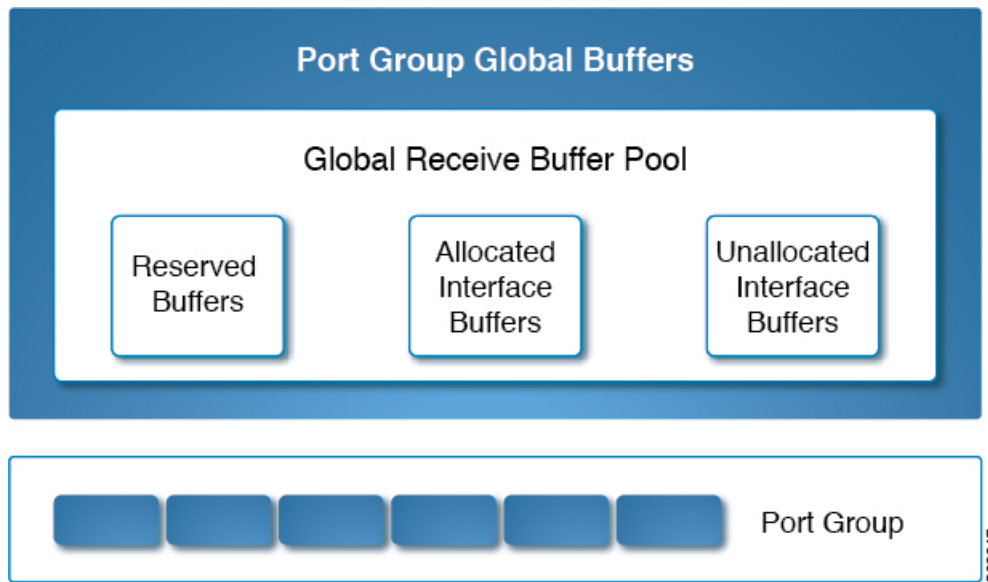
- Cisco MDS switches support mechanisms to avoid loss of R_RDYs which may cause issues on links. For more information, see the [Buffer-to-Buffer Credit Recovery, on page 16](#) section.
- Reconfiguring BB_credits on an active link is a disruptive operation.
- The number of Rx BB_credits does not have to match on each side of the link.
- Only the Rx BB_credits can be configured on an individual interface because those are the only credits that an interface has control over.
- If the transmitter decrements the remaining Tx BB_credit remaining number and hits zero, the *Tx transition to zero counter* will be incremented by one. This typically indicates some level of congestion at the receiving device. Although it could also indicate that there are insufficient buffers for the speed and distance of the link.
- If the receiver does not transmit R_RDYs to the sender of frames, then once the number of frames equals to the Rx BB_credits are received the sender must stop sending since it has hit 0 Tx BB_credit remaining. The receiver will also hit 0 Rx BB_credits remaining and will increment its *Rx transition to zero* counter by one.
- Long-distance links may must have the number of BB_credits increased on both sides to ensure maximum performance.

Global Receive Buffer Pool

A port group is a set of contiguous ports that share common resources such as bandwidth and buffer credits from a global pool of buffers.

The global pool of buffers includes the global receive buffer pool. The global receive buffer pool includes the following buffer groups:

- Reserved internal buffers
- Allocated buffers for each Fibre Channel interface (user configured or assigned by default)
- Unallocated buffers, if any, to be used for additional buffers when required



Extended Buffer-to-Buffer Credits

Extended buffer-to-buffer credits are made possible by allocating extra buffers to specific interfaces. These extra buffers are taken from the unallocated buffer pool.



Note

The ENTERPRISE_PKG license is required to use extended buffer-to-buffer credits on 16 Gbps, 32 Gbps, and 64 Gbps switching modules.

All ports on the 16 Gbps and 32 Gbps switching modules support extended buffer-to-buffer credits. There is a limitation on the maximum number of extended buffer-to-buffer credits you can assign to a port. If necessary, you can configure interfaces to use minimum credits to make more extended buffer-to-buffer credits available to other ports.

For long-distance ISLs, the extended buffer-to-buffer credits feature allows you to configure the receive buffers up to the level you need and within hardware resource limits. When necessary, you can reduce the buffers on one port and assign them to another port in the same port group. However, you must first release the buffers from the other ports before configuring larger extended buffer-to-buffer credits for a port.

Extended BB_credits are typically used on long-distance ISL ports (E ports). If you require additional BB_credits on a port or a group of ports, buffers may need to be made available.

To allow most buffers to be available, perform these steps:

1. Configure ISL ports over different port-groups and modules.
2. Configure ports that are connected to end devices (F ports) from **mode auto** to **mode F** by using the **switchport mode f** command.

In general, you can configure any port in a port group to dedicated rate mode. To do this, you must first release the buffers from the other ports before configuring larger extended buffer-to-buffer credits for a port. This will reduce the number of buffers allocated to those ports 500–32 (on most switch types) and add those saved buffers into the unallocated pool.

Without changing the default mode or speed, you can assign the remaining available BB_credits within the port group, using the **switchport fcrxbbcredit extended credits** command. Use the **show port-resources module slot** command to verify the updated BB_credits allocation among the interfaces.

For example, if there are 300 extended BB_credits available, we can assign these 300 BB_credits to a port that is having 500 BB_credits so that the port now has 800 BB_credits provided all other ports in the port group are configured with speed auto and mode auto.

```
switch# configure terminal
switch(config)# interface fc1/1
switch(config-if)# switchport fcrxbbcredit extended 800
```

When we configure ports in a port group to F Port mode, the reserved buffer credits for such ports are reduced from 500 BB_credits to 32 BB_credits and the remaining BB_credits are assigned to the unallocated interface buffer pool. If you need more BB_credits for that particular port group buffer pool, you can reduce the number of BB_credits being used by the F ports, using the **switchport fcrxbbcredit credits** command.

An alternate option is to configure the remaining ports in a port group to minimum credits to free up all BB_credits from these ports for extended BB_credits use.



Note The ENTERPRISE_PKG license is required to use extended buffer-to-buffer credits on 16 Gbps and 32 Gbps switching modules. All ports on the 16 Gbps and 32 Gbps switching modules support extended buffer-to-buffer credits. There are no limitations for how many extended buffer-to-buffer credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can configure interfaces to use minimum credits to make more extended buffer-to-buffer credits available to other ports.

Default BB Credit Buffers

Table 2: Default BB Credit Buffers

	Mode		
Speed		Auto	Fixed
	Auto	500	E ports: 500
			F ports: 32
	Fixed	E ports: 500	E ports: 500
		F ports: 32	F ports: 32
			Note 1 extra buffer is always consumed by each F and 2 extra buffers are always consumed by each E ports from the port group buffer pool.



Note 16 extra buffers are always consumed by each E and F ports from the port group buffer pool.

Buffer-to-Buffer Credit Allocation

This section describes how buffer credits are allocated on Cisco MDS 9000 Series Multilayer switches.

32 Gbps Switching Modules or Switches

Table 3: 32 Gbps Switching Modules or Switches Buffer-to-Buffer Credit Allocation

Number of port groups	<ul style="list-style-type: none"> • Cisco MDS 9220i: 1 • Cisco MDS 9132T: 2 • Cisco MDS 9148T: 3 • Cisco MDS 9396T: 6 • Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module: 3
-----------------------	--

Default buffer-to-buffer credits	Auto/Mode E ports: 500
	F ports: 32
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port group, when all other ports in the port group are NOT configured with minimum BB credit of 1 or 2, using extend BB credits	300
Maximum configurable global buffers available per port using extended buffers, when all other ports in the port group are configured with minimum BB credit of 1 or 2, using extended BB credits	8170

16 Gbps Switching Modules or Switches

Table 4: 16 Gbps Switching Modules or Switches Buffer-to-Buffer Credit Allocation

Number of port groups	<ul style="list-style-type: none"> • Cisco MDS 9396S: 24 • Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module: 12 • Cisco MDS 9700 24/10-Port SAN Extension Module: 6
Default buffer-to-buffer credits	Auto/Mode E ports: 500
	F ports: 32
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port group, when all other ports in the port group are NOT configured with minimum BB credit of 1 or 2, using extend BB credits	2150

Maximum configurable global buffers available per port using extended buffers, when all other ports in the port group are configured with minimum BB credit of 1 or 2, using extended BB credits	4095
--	------

Cisco MDS 9250i and Cisco MDS 9148S Fabric Switch

Table 5: Cisco MDS 9250i and Cisco MDS 9148S Fabric Switch Buffer-to-Buffer Credit Allocation

Number of port groups	<ul style="list-style-type: none"> • Cisco MDS 9250i: 10 • Cisco MDS 9148S: 12
Default buffer-to-buffer credits	Auto/Mode E ports: 64
	F ports: 64
Minimum configurable buffers per port	Auto/Mode E ports: 2
	F ports: 1
Extended Buffer-to-Buffer Credit Allocation	
Maximum configurable global buffers available per port using extended credits, when all other ports in the port group are configured with minimum BB credit of 1, using extended BB credits	253



Note The number of BB credits allocated for Cisco MDS 9250i and MDS 9148s is 64 BB credits per port but can be extended to 253 BB credits when other ports in the port group are moved to minimum credits without the need of an *Enterprise_PKG* license.

Examples: Buffer-to-Buffer Credit Allocation



Note In the command outputs, if the bandwidth is displayed as 32 Gbps, then the output is from either Cisco MDS 9700 switch with Cisco MDS 9700 48-Port 32-Gbps Fibre Channel Switching Module, Cisco MDS 9220i, MDS 9132T, MDS 9148T, or MDS 9396T switches.

In the command outputs, if the bandwidth is displayed as 16 Gbps, then the output is from either Cisco MDS 9700 switch with Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module, Cisco MDS 9148S, or Cisco MDS 9250i switches.

The following example displays the default buffers when the switchport mode and speed are set to auto:

```
switch(config)# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 300 Available dedicated
  buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
  buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
-----
Interfaces in the  B2B Credit Bandwidth      Rate Mode
Port-Group        Buffers      (Gbps)
-----
fc1/1              500          32.0      dedicated
fc1/2              500          32.0      dedicated
fc1/3              500          32.0      dedicated
fc1/4              500          32.0      dedicated
fc1/5              500          32.0      dedicated
fc1/6              500          32.0      dedicated
fc1/7              500          32.0      dedicated
fc1/8              500          32.0      dedicated
fc1/9              500          32.0      dedicated
fc1/10             500          32.0      dedicated
fc1/11             500          32.0      dedicated
fc1/12             500          32.0      dedicated
fc1/13             500          32.0      dedicated
fc1/14             500          32.0      dedicated
fc1/15             500          32.0      dedicated
fc1/16             500          32.0      dedicated
```

The following example displays the buffer allocation when one port is set to E port mode, remaining ports are set to F Port mode, and all ports are set to speed auto:

```
switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 7320 Available dedicated
  buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
  buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
-----
Interfaces in the  B2B Credit Bandwidth      Rate Mode
Port-Group        Buffers      (Gbps)
-----
fc1/1              500          32.0      dedicated
fc1/2              32           32.0      dedicated
fc1/3              32           32.0      dedicated
fc1/4              32           32.0      dedicated
fc1/5              32           32.0      dedicated
fc1/6              32           32.0      dedicated
fc1/7              32           32.0      dedicated
fc1/8              32           32.0      dedicated
fc1/9              32           32.0      dedicated
fc1/10             32           32.0      dedicated
fc1/11             32           32.0      dedicated
fc1/12             32           32.0      dedicated
fc1/13             32           32.0      dedicated
```

```

fc1/14          32          32.0          dedicated
fc1/15          32          32.0          dedicated
fc1/16          32          32.0          dedicated

```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set to F Port mode with speed auto, 16000, or 32000:

```

switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
-----
Interfaces in the   B2B Credit   Bandwidth   Rate Mode
Port-Group         Buffers      (Gbps)
-----
fc1/1              7820        32.0        dedicated
fc1/2              32          32.0        dedicated
fc1/3              32          32.0        dedicated
fc1/4              32          32.0        dedicated
fc1/5              32          32.0        dedicated
fc1/6              32          32.0        dedicated
fc1/7              32          32.0        dedicated
fc1/8              32          32.0        dedicated
fc1/9              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated
fc1/1              32          32.0        dedicated

```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set to F Port mode with speed 8000:

```

switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 152.0 Gbps
-----
Interfaces in the   B2B Credit   Bandwidth   Rate Mode
Port-Group         Buffers      (Gbps)
-----
fc1/1              7580        32.0        dedicated
fc1/2              32          8.0         dedicated
fc1/3              32          8.0         dedicated
fc1/4              32          8.0         dedicated
fc1/5              32          8.0         dedicated
fc1/6              32          8.0         dedicated

```

```

fc1/7          32          8.0          dedicated
fc1/8          32          8.0          dedicated
fc1/9          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated
fc1/1          32          8.0          dedicated

```

The following example displays the buffer allocation when two ports are set to E port mode with extended buffers, remaining ports are set to F, and all ports are set to speed auto:

```

switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 0 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 512.0 Gbps
-----
Interfaces in the   B2B Credit   Bandwidth   Rate Mode
Port-Group         Buffers      (Gbps)
-----
fc1/1               3926         32.0        dedicated
fc1/2               3926         32.0        dedicated
fc1/3               32           32.0        dedicated
fc1/4               32           32.0        dedicated
fc1/5               32           32.0        dedicated
fc1/6               32           32.0        dedicated
fc1/7               32           32.0        dedicated
fc1/8               32           32.0        dedicated
fc1/9               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated
fc1/1               32           32.0        dedicated

```

The following example displays the buffer allocation when one port is set to E port mode with extended buffers, speed auto, and remaining ports are set out of service:

```

switch# show port-resources module 1
Module 1
Available dedicated buffers for global buffer #0 [port-group 1] are 94 Available dedicated
buffers for global buffer #1 [port-group 2] are 300 Available dedicated buffers for global
buffer #2 [port-group 3] are 300

Port-Group 1
Total bandwidth is 512.0 Gbps
Allocated dedicated bandwidth is 32.0 Gbps
-----
Interfaces in the Port-Group B2B Credit Bandwidth Rate Mode
                        Buffers      (Gbps)
-----

```

```

fc1/1      8191 32.0      dedicated
fc1/2 (out-of-service)
fc1/3 (out-of-service)
fc1/4 (out-of-service)
fc1/5 (out-of-service)
fc1/6 (out-of-service)
fc1/7 (out-of-service)
fc1/8 (out-of-service)
fc1/9 (out-of-service)
fc1/10 (out-of-service)
fc1/11 (out-of-service)
fc1/12 (out-of-service)
fc1/13 (out-of-service)
fc1/14 (out-of-service)
fc1/15 (out-of-service)
fc1/16 (out-of-service)

```

The following example displays how to allocate maximum BB_credits on Cisco 9148S and 9250i switches:

The following example shows that the port-group 2 on the switch includes ports fc1/5-8 and each port has 64 credits:

```

switch# show port-resources module 1
.
.
.

Port-Group 2
Available dedicated buffers are 0
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers      (Gbps)
-----
fc1/5                            64          16.0      dedicated
fc1/6                            64          16.0      dedicated
fc1/7                            64          16.0      dedicated
fc1/8                            64          16.0      dedicated

```

To allocate maximum BB_credits to port fc1/5, perform these steps:

1. Configure ports fc1/6-8 in the port-group to a minimum BB_Credit of 1:

```

switch# configure
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# interface fc1/6-8
switch(config-if)# switchport fcrxbbcredit 1

```

2. Configure port fc1/5 with the maximum BB_credits of 253:

```

switch# configure
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# interface fc1/5
switch(config-if)# switchport fcrxbbcredit 253

```

3. Verify the BB_credits allocation on port fc1/5:

```
switch# show port-resources module 1
.
.
.

Port-Group 2
Available dedicated buffers are 0
-----
Interfaces in the Port-Group      B2B Credit  Bandwidth  Rate Mode
                                Buffers      (Gbps)
-----
fc1/5                            253         16.0       dedicated
fc1/6                             1           16.0       dedicated
fc1/7                             1           16.0       dedicated
fc1/8                             1           16.0       dedicated
```

Long-Distance ISLs

When long-distance ISLs are required, you must have sufficient BB_credits configured to ensure that the ISL can run at maximum capacity. The simplest formula or the rule of thumb for computing BB_credits for long-distance ISL assumes a full sized Fibre Channel frame of approximately 2 KB and factors in the interface operating speed and one way distance of the ISL.



Note If the average frame size is less than 2 KB, the number of BB_credits must be increased.

Interface Speed	Minimum Number of BB_Credits Required Per Km (One Way)
1 Gbps	0.5 BB_Credit
2 Gbps	1 BB_Credit
4 Gbps	2 BB_credits
8 Gbps	4 BB_credits
16 Gbps	8 BB_credits
32 Gbps	16 BB_credits

As per the table, to operate a 16-Gbps Fibre Channel ISL over 50 km, you would multiply the one way distance (50) times the minimum number of BB_credits per km (8). That is, a 50 km 16-Gbps ISL requires 400 BB_credits when the average frame size is approximately 2 KB. This is the minimum number of BB_credits that are required for the link to function at its best when utilized to its maximum. To accommodate a smaller average frame size than the maximum (full sized) value, more BB_credits would be required proportionally. Since each buffer is for a Fibre Channel frame irrespective of its size, when Fibre Channel frames are not full sized, more BB_credits are required to achieve full link utilization. In this case, an approximate yet simple formula for calculating BB_credits is the following:

$$BB_credits = (\text{Minimum number of BB_credits required per km for interface speed} \times \text{One-way distance (km)}) / ((\text{Average receive frame size (bytes)} / 2150 \text{ bytes}))$$

The following example displays the BB_credits calculation for a 16 Gbps link that is 50 km long with an average input frame size of approximately 1 KB (1075 bytes):

$$(8 \text{ BB_credits per km at 16 Gbps} \times 50 \text{ km}) / (1075 / 2150) = 800 \text{ BB_credits}$$

To take into consideration the actual average input frame size first determine the average frame size by dividing the total input bytes by the total frames input. The average frame size must be determined for the input direction (Rx side) on an interface since the receive BB_credits are being set. The total bytes and frames can be viewed in the **show interface counters** command output.

```
switch# show interface fc 2/7 counters
fc2/7
 5 minutes input rate 1048060640 bits/sec, 131007580 bytes/sec, 94786 frames/sec
 5 minutes output rate 253368512 bits/sec, 31671064 bytes/sec, 47717 frames/sec
 14079632456 frames input, 18624775031572 bytes
 0 discards, 0 errors, 0 CRC/FCS
 0 unknown class, 0 too long, 0 too short
8089598629 frames output, 6040401816628 bytes
 0 discards, 0 errors
 0 timeout discards, 0 credit loss
 0 input OLS, 0 LRR, 0 NOS, 0 loop inits
 0 output OLS, 0 LRR, 0 NOS, 0 loop inits
 0 link failures, 0 sync losses, 0 signal losses
15031 Transmit B2B credit transitions to zero
 0 Receive B2B credit transitions to zero
11192 2.5us TxWait due to lack of transmit credits
Percentage TxWait not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
500 receive B2B credit remaining
481 transmit B2B credit remaining
481 low priority transmit B2B credit remaining
Last clearing of "show interface" counters: 2d09h
```

In the above example, the calculation for the average frame size is:

$$18624775031572 \text{ (bytes)} / 14079632456 \text{ (frames)} = 1323 \text{ bytes/frame which is approximately 1.3 KB/frame}$$

To complete the calculation:

$$(8 \text{ BB_credits per km at 16 Gbps} \times 50 \text{ km}) / (1323 \text{ average bytes per frame} / 2150 \text{ bytes}) = 650 \text{ BB_credits}$$

Thus, for an 8 Gbps link traversing 50 km carrying an average frame size of 1323 bytes would take a minimum of 650 BB_credits.



Note The required BB_credits for the other end of the same link may be different due to a different average frame size in the opposite direction. The average frame size must be calculated in a similar way from the adjacent interface.

For more information on how to change BB_credits, see the [Extended Buffer-to-Buffer Credits, on page 6](#) section and see the **switchport fcrxbbcredit std_bufs** and **switchport fcrxbbcredit extended ext_bufs** command outputs.

Buffer-to-Buffer Credit Recovery

Although Fibre Channel standards require low bit and frame error rates, there is a likelihood of errors occurring. When these errors affect certain Fibre Channel primitives, credit loss might occur. When credits are lost,

performance degradation might occur. When all credits are lost, transmission of frames in that direction stops. The Fibre Channel standards introduces a feature for two attached ports to detect and correct such scenarios nondisruptively. This feature is called *buffer-to-buffer credit recovery*.

A credit can be lost in either of these scenarios:

- An error corrupts the start-of-frame (SoF) delimiter of a frame. The receiving port fails to recognize the frame and subsequently does not send a corresponding receiver ready (R_RDY) primitive to the sender. The sending port does not replenish the credit to the receiving port.
- An error corrupts an R_RDY primitive. The receiving port fails to recognize the R_RDY and does not replenish the corresponding credit to the sending port.

The Buffer-to-Buffer Credit Recovery feature can help recover from the two specified scenarios. It is a per-hop feature and is negotiated between two directly attached peer ports when the link comes up, by exchanging parameters. Buffer-to-buffer credit recovery is enabled when a receiver acknowledges a nonzero buffer-to-buffer state change number (BB_SC_N).

Buffer-to-buffer credit recovery functions as follows:

1. The local port and peer port agree to send checkpoint primitives to each other for frames and R_RDYs, starting from the time the link comes up.
2. If a port detects frame loss, it sends the corresponding number of R_RDYs to replenish the lost credits at the peer port.
3. If a port detects R_RDY loss, the port internally replenishes the lost credits to the interface buffer pool.

Buffer-to-buffer credit recovery implementation is as follows:

1. Buffer-to-buffer state change SOF (BB_SCs) primitives are transmitted every $2^{BB_SC_N}$ number of frames sent. This enables an attached port to determine if any frames are lost. If frames loss is detected, the receiver of the BB_SCs transmits the appropriate number of R_RDYs to compensate for the lost frames.
2. Buffer-to-buffer state change R_RDY (BB_SCr) primitives are transmitted every $2^{BB_SC_N}$ number of R_RDY primitives sent. This enables an attached port to determine if any R_RDY primitives are lost. If R_RDY primitive loss is detected, the receiver of the BB_SCr increments the number of transmit credits by the appropriate number to compensate for the lost R_RDYs.

The Buffer-to-Buffer Credit Recovery feature can be used on any nonarbitrated loop link. This feature is most useful on unreliable links, such as Metropolitan Area Networks (MANs) or WANs, but can also help on shorter, high-loss links, such as a link with a faulty fiber connection.



Note The Buffer-to-Buffer Credit Recovery feature is not compatible with the distance extension (DE) feature, also known as buffer-to-buffer credit spoofing. If you use intermediate optical equipment, such as dense wavelength-division multiplexing (DWDM) or Fibre Channel bridges that use DE on Inter-Switch Links (ISLs) between switches, then buffer-to-buffer credit recovery on both sides of an ISL must be disabled using the **no switchport fcbbn** command.

The following are the guidelines and restrictions for the Buffer-to-Buffer Credit Recovery feature:

- E ports
 - This feature is enabled by default on ISLs (E ports).

- This feature works on an ISL between a Cisco switch and a peer switch from any vendor, provided this feature is supported on the peer switch.
- This feature is supported only on links that are in R_RDY flow control mode. It is not supported on links that are in ER_RDY flow control mode.
- F ports
 - This feature is enabled by default on F ports starting from Cisco MDS NX-OS Release 8.2(1).
 - This feature works on an F port between a Cisco switch and a peer device from any vendor, provided this feature is supported on the peer device.



Note Some host bus adapters (HBAs) do not support the Buffer-to-Buffer Credit Recovery feature. Others support this feature at only certain speeds. Check with your HBA vendor about the exact configurations supported.

- NP ports
 - The adjacent N-PortID Virtualization (NPIV) F port must also support this feature. Prior to Cisco MDS NX-OS Release 8.4(1), N-PortID Virtualization (NPIV) ports do not support buffer-to-buffer credit recovery for Cisco N-Port Virtualizer (Cisco NPV) switch logins.
 - This feature is enabled by default on NP ports starting from Cisco MDS NX-OS Release 8.4(1).

The count of times the buffer-to-buffer credits have been recovered for both types of recovery can be displayed using the **show interface counters detailed** command:

Prior to Cisco MDS NX-OS Release 8.4(1a) and earlier release:

```
switch# show interface fc1/1 counters detailed
fc1/1
...
0 BB_SCs credit resend actions, 0 BB_SCr Tx credit increment actions
```

From Cisco MDS NX-OS Release 8.4(2) and later release:

```
switch# show interface fc1/1 counters detailed
fc1/1
...

Congestion Stats:
Tx Timeout discards: 0
Tx Credit loss: 0
BB_SCs credit resend actions: 0
BB_SCr Tx credit increment actions: 0
```

Receive Data Field Size

By default, the maximum data field size is configured for Fibre Channel interfaces and cannot be reconfigured.

Configuring Interface Buffers

Configuring Buffer-to-Buffer Credits



Note When you configure port mode to auto or E for all the ports in the global buffer pool, you must reconfigure buffer credits on one or more ports (other than the default mode).

To configure a single pool of buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be in R_RDY flow-control mode.

Before you begin

Enable the Receiver Ready (R_RDY) mode on ISLs before configuring the shared buffer-to-buffer credit pool. For more information, see [Disabling Extended Receiver Ready](#).

-
- Step 1** Enter configuration mode:
switch# **configure terminal**
- Step 2** Select a Fibre Channel interface and enter interface configuration submode:
switch(config)# **interface fc slot/port**
- Step 3** Set the buffer-to-buffer credits as a single pool on an interface:
switch(config-if)# **switchport fcxbbcredit credits mode {E | Fx}**
(Optional) Reset the buffer-to-buffer credits on the interface to the default value:
switch(config-if)# **switchport fcxbbcredit default**
-

Configuring Buffer-to-Buffer Credits for Virtual Links



Note When you configure port mode to auto or E, and rate mode to dedicated for all the ports in the global buffer pool, you must reconfigure buffer credits on one or more ports (other than the default mode).

To configure per-virtual-link buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be an ISL in ER_RDY flow-control mode.

Before you begin

Enable the Extended Receiver Ready (ER_RDY) mode on ISLs before configuring the virtual-link credits. For more information, see [Enabling Extended Receiver Ready](#).

-
- Step 1** Enter configuration mode:
switch# **configure terminal**
- Step 2** Select a Fibre Channel interface and enter interface configuration submode:
switch(config)# **interface fc slot/port**
- Step 3** Set the buffer-to-buffer credits per virtual-link on an ISL:
switch(config-if)# **switchport vl-credit v10 credits v11 credits v12 credits v13 credits**
- Step 4** (Optional) Reset the buffer-to-buffer credits on the ISL to the default value:
switch(config-if)# **switchport vl-credit default**
-

Configuring Extended Buffer-to-Buffer Credits



Note You cannot configure regular buffer-to-buffer credits after configuring the extended buffer-to-buffer credits.

To configure a single pool of extended buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be in R_RDY flow-control mode.

Before you begin

Enable the Receiver Ready (R_RDY) mode on ISLs before configuring the shared buffer-to-buffer credit pool. For more information, see [Disabling Extended Receiver Ready](#).

-
- Step 1** Enter configuration mode:
switch# **configure terminal**
- Step 2** Enable extended Rx B2B credit configuration:
switch(config)# **feature fcrxbbcredit extended**
- Step 3** Select a Fibre Channel interface and enter interface configuration submode:
switch(config)# **interface fc slot/port**
- Step 4** Set the extended buffer-to-buffer credits as a single pool on an interface:
switch(config-if)# **switchport fcrxbbcredit extended extend_bufs**
- Note** If the ER_RDY flow-control mode is enabled using the **system fc flow-control er_rdy** command, then the configured credits are allocated to individual virtual lanes. For example, if the **switchport fcrxbbcredit extended 1000** command is configured for an interface, the extended buffers for the virtual lanes are configured as **switchport vl-credit extended v10 16 v11 16 v12 47 v13 921**.
- Step 5** (Optional) Reset the extended buffer-to-buffer credits on the interface to the default value:

```
switch(config-if)# switchport fcrxbcredit extended default
```

Configuring Extended Buffer-to-Buffer Credits for Virtual Links



Note You cannot configure regular buffer-to-buffer credits after configuring the extended buffer-to-buffer credits.

To configure per-virtual-link extended buffer-to-buffer credits for a Fibre Channel interface, perform these steps. The interface must be an ISL in ER_RDY flow-control mode.

Before you begin

Enable the Extended Receiver Ready (ER_RDY) mode on ISLs before configuring the virtual link credits. For more information, see [Enabling Extended Receiver Ready](#).

- Step 1** Enter configuration mode:
switch# **configure terminal**
- Step 2** Enable extended Rx B2B credit configuration:
switch(config)# **feature fcrxbcredit extended**
- Step 3** Select a Fibre Channel interface and enter interface configuration submode:
switch(config)# **interface fc slot/port**
- Step 4** Set the extended buffer-to-buffer credits per virtual link on an ISL:
switch(config-if)# **switchport vl-credit extended vl0 credits vl1 credits vl2 credits vl3 credits**
- Step 5** (Optional) Reset the extended buffer-to-buffer credits on the ISL to the default value:
switch(config-if)# **switchport vl-credit extended default**
-

Configuring Buffer-to-Buffer Credit Recovery

Buffer-to-buffer credit recovery is enabled by default on all Fibre Channel ports.

To disable or enable the buffer-to-buffer credit recovery on a port, perform these steps:

- Step 1** Enter configuration mode:
switch# **configure terminal**
- Step 2** Select the interface and enter interface configuration submode:
switch(config)# **interface fc slot/port**

Step 3 Disable buffer-to-buffer credit recovery on the interface:

```
switch(config-if)# no switchport fcbbscn
```

Step 4 (Optional) To enable buffer-to-buffer credit recovery on an interface if it was disabled:

- Cisco MDS NX-OS Release 8.4(1) and earlier releases

```
switch(config-if)# switchport fcbbscn
```

Note The BB_SC_N value is set to the default value of 14.

- Cisco MDS NX-OS Release 8.4(2) and later releases

```
switch(config-if)# switchport fcbbscn value value
```

Caution This command causes traffic disruption on the specified interface.

Configuring Receive Data Field Size



Note From Cisco MDS NX-OS 8.2(1), the **switchport fcxbuFSIZE** command is obsolete on the Cisco MDS 9700 48-port 16-Gbps Fibre Channel Switching Module and the Cisco MDS 9700 48-port 32-Gbps Fibre Channel Switching Module. The receive data field size is permanently set to 2112 bytes. Any receive data field size configuration from earlier Cisco MDS NX-OS versions is ignored.

To configure the receive data field size, perform these steps:

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Select a Fibre Channel interface and enter interface configuration submode:

```
switch(config)# interface fc slot/port
```

Step 3 Set the data field size for the selected interface:

```
switch(config-if)# switchport fcxbuFSIZE bytes
```

Step 4 (Optional) Reset the receive data field size on the interface to the default value:

```
switch(config-if)# no switchport fcxbuFSIZE
```

Configuration Examples for Interface Buffers

This example shows how to enable buffer-to-buffer credit recovery on an interface if it is disabled:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcbbscn
```

This example shows how to configure default credits on an interface:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit default
```

This example shows how to configure 50 receive buffer credits on an interface:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit 50
```

This example shows how to configure 4095 extended buffer credits to an interface:

```
switch# configure terminal
switch(config)# fcrxbbcredit extended enable
switch(config)# interface fc 1/1
switch(config-if)# switchport fcrxbbcredit extended 4095
```

This example shows how to assign buffer-to-buffer credits per virtual link on an ISL:

```
switch# configure terminal
switch(config)# interface fc 1/1
switch(config-if)# switchport vl-credit v10 12 v11 10 v12 29 v13 349
```

This example shows how to assign extended buffer-to-buffer credits per virtual link on an ISL:

```
switch# configure terminal
switch(config)# fcrxbbcredit extended enable
switch(config)# interface fc 1/1
switch(config-if)# switchport vl-credit extended v10 20 v11 25 v12 40 v13 349
```

Verifying Interface Buffer Configuration

This example shows which of the interfaces on a specified module are in R_RDY flow-control mode:

```
switch# show flow-control r_rdy module 3
fc3/17
fc3/18
```

This example shows how to verify the buffer-to-buffer credit information for all interfaces:

```
sswitch# show interface bbcredit
fc2/1 is down (SFP not present)
.
.
.
fc2/17 is trunking
Transmit B2B Credit is 255
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 375
12 receive B2B credit remaining
255 transmit B2B credit remaining
fc2/21 is down (Link failure or not-connected)
.
.
.
fc2/31 is up
Transmit B2B Credit is 0
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 48
12 receive B2B credit remaining
0 transmit B2B credit remaining
```

This example shows how to verify buffer-to-buffer credit information for a specific Fibre Channel interface:

```
switch# show interface fc2/31 bbcredit
fc2/31 is up
Transmit B2B Credit is 0
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 48
12 receive B2B credit remaining
0 transmit B2B credit remaining
```

This example shows how to verify the type of buffers and data field size a port supports:

```
switch# show interface fc1/1 capabilities
fc1/1
Min Speed is 2 Gbps
Max Speed is 16 Gbps
FC-PH Version (high, low) (0,6)
Receive data field size (max/min) (2112/256) bytes
Transmit data field size (max/min) (2112/128) bytes
Classes of Service supported are Class 2, Class 3, Class F
Class 2 sequential delivery supported
Class 3 sequential delivery supported
Hold time (max/min) (100000/1) micro sec
BB state change notification supported
```



```

Maximum BB state change notifications 14
Rate Mode change not supported

Rate Mode Capabilities Dedicated
Receive BB Credit modification supported yes
FX mode Receive BB Credit (min/max/default) (1/500/32)
ISL mode Receive BB Credit (min/max/default) (2/500/500)
Performance buffer modification supported yes
FX mode Performance buffers (min/max/default) (1/0/0)
ISL mode Performance buffers (min/max/default) (1/0/0)

Out of Service capable yes
Beacon mode configurable yes
Extended B2B credit capable yes
On demand port activation license supported no

```

This example shows how to verify the operational receive data field size for a port:

```

switch# show interface fc 4/1
fc4/1 is down (SFP not present)
Hardware is Fibre Channel
Port WWN is 20:c1:8c:60:4f:c9:53:00
Admin port mode is auto, trunk mode is on
snmp link state traps are enabled
Port vsan is 1
Receive data field Size is 2112
Beacon is turned off
Logical type is Unknown(0)
5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
4 frames input,304 bytes
0 discards,0 errors
0 invalid CRC/FCS,0 unknown class
0 too long,0 too short
4 frames output,304 bytes
0 discards,0 errors
0 input OLS,0 LRR,0 NOS,0 loop inits
0 output OLS,0 LRR, 0 NOS, 0 loop inits
Last clearing of "show interface" counters : never

```

This example shows how to verify credit mode and credit allocation for an ISL:

```

switch# show interface fc9/1
.
.
.
Port flow-control is ER_RDY

Transmit B2B Credit for v10 is 15
Transmit B2B Credit for v11 is 15
Transmit B2B Credit for v12 is 40
Transmit B2B Credit for v13 is 430
Receive B2B Credit for v10 is 15
Receive B2B Credit for v11 is 15
Receive B2B Credit for v12 is 40
Receive B2B Credit for v13 is 430
.
.
.

```

Troubleshooting Interface Buffer Credits

Use the **show interface counters detailed** and the **show logging onboard interrupt-stats** commands to view the number of times a port sent extra R_RDYs or incremented transmit buffer to buffer credits to restore credit counts:

```
switch# show logging onboard interrupt-stats
.
.
.
-----
INTERRUPT COUNTS INFORMATION FOR DEVICE: FCMAC
-----
Interface|                               |          |          |          |
Range    |          Interrupt Counter Name | Count   | MM/DD/YY | HH:MM:SS |
         |                               |         |          |          |
-----|-----|-----|-----|-----|
fc1/1    | IP_FCMAC_INTR_ERR_BB_SCR_INCREMENT | 1       | 01/01/17 | 20:00:00 |
fc1/1    | IP_FCMAC_INTR_ERR_BB_SCS_RESEND   | 1       | 01/01/17 | 10:00:00 |
.
.
.
```

The BB_SCR credit recoveries use the underlying IP_FCMAC_INTR_ERR_BB_SCR_INCREMENT counter and the counter indicates the number of R_RDYs that were lost. The IP_FCMAC_INTR_ERR_BB_SCS_RESEND counter specifies the number of frames that were lost.

Use the **show interface port/slot counters** command to determine the interval the switch was unable to transmit frames since the counters were last cleared:

```
switch# show interface fc1/13 counters
.
.
.
    6252650 2.5us Txwaits due to lack of transmit credits
.
.
.
```

Txwait value can be converted to seconds using the following formula:

TxWait value in seconds = ((TxWait value in 2.5 μs ticks) x 2.5)/(1,000,000)

Using this formula, we can see that the switch was unable to transmit frames for more than 15 seconds.

Use the **show interface port/slot counters** command to determine the duration for which the Tx BB credits were zero for the last 1 second, 1 minute, 1 hour, and 72 hours:

```
switch# show interface fc1/13 counters
.
.
.
    Percentage Tx credits not available for last 1s/1m/1h/72h: 1%/5%/3%/2%
.
.
.
```

Use the **show logging onboard txwait module number** command to check duration for which the remaining Txwait BB credits were zero over the span of 20 seconds:

```
switch# show logging onboard txwait module 2
```

```
-----
Module: 2 txwait count
-----
```

```
-----
Show Clock
-----
2019-04-08 13:56:52
Notes:
  - Sampling period is 20 seconds
  - Only txwait delta >= 100 ms are logged
```

```
-----
```

Interface	Delta TxWait Time	Congestion	Timestamp
	2.5us ticks seconds		
Eth2/2 (VL3)	882562 2	11%	Tue Sep 11 08:52:34 2018
Eth2/1 (VL3)	4647274 11	58%	Tue Sep 11 08:52:14 2018
Eth2/2 (VL3)	7529479 18	94%	Tue Sep 11 08:52:14 2018
Eth2/1 (VL3)	7829159 19	97%	Tue Sep 11 08:51:54 2018
Eth2/2 (VL3)	7923544 19	99%	Tue Sep 11 08:51:54 2018
Eth2/1 (VL3)	5299754 13	66%	Tue Sep 11 08:50:34 2018
Eth2/2 (VL3)	362484 0	4%	Tue Sep 11 08:50:34 2018
Eth2/1 (VL3)	7924925 19	99%	Tue Sep 11 08:50:14 2018
Eth2/2 (VL3)	2566450 6	32%	Tue Sep 11 08:50:14 2018
Eth2/1 (VL3)	7935558 19	99%	Tue Sep 11 08:49:54 2018
Eth2/2 (VL3)	6762560 16	84%	Tue Sep 11 08:49:54 2018
Eth2/1 (VL3)	7908259 19	98%	Tue Sep 11 08:49:34 2018
Eth2/2 (VL3)	5264976 13	65%	Tue Sep 11 08:49:34 2018
Eth2/1 (VL3)	7925639 19	99%	Tue Sep 11 08:49:14 2018

Use the **show logging onboard error-stats** command to list the ports with zero remaining Tx BB credits for 100 ms:

```
switch# show logging onboard error-stats
```

```
-----
Module: 1
-----
```

```
-----
Show Clock
-----
2018-08-28 12:28:15
```

```
-----
Module: 1 error-stats
-----
```

```
-----
ERROR STATISTICS INFORMATION FOR DEVICE: FCMAC
-----
```

Interface Range	Error Stat Counter Name	Count	Time Stamp MM/DD/YY HH:MM:SS
fc7/2	IP_FCMAC_CNT_STATS_ERRORS_RX_BAD_WORDS_FROM_DECODER	35806503	03/17/19 11:32:44
fc7/2	FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO	2	03/17/19 11:32:44
fc7/1	FCP_SW_CNTR_TX_WT_AVG_B2B_ZERO	1	03/17/19 11:32:44
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	1	03/15/19 22:10:25
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	16	03/15/19 18:32:44
fc7/15	F16_TMM_TOLB_TIMEOUT_DROP_CNT	443	03/15/19 15:39:42
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	12	03/15/19 13:37:59
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	8	03/15/19 13:29:59
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	4	03/15/19 13:26:19
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	3	01/01/17 13:12:14
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	25	03/14/19 21:13:34
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	21	03/14/19 21:06:34
fc7/15	FCP_SW_CNTR_RX_WT_AVG_B2B_ZERO	17	03/14/19 20:58:34

Use the **show interface port/slot bbcredit** command to verify BB credits information:

```
switch# show interface fc1/1 bbcredit
fc1/1 is up
  Transmit B2B Credit is 16
  Receive B2B Credit is 16
    17 receive B2B credit remaining
    16 transmit B2B credit remaining
```

Use the **show interface slot/port bbcredit** command to check for discrepancy in the buffer-to-buffer credit values for a specific Fibre Channel interface:

```
switch# show interface fc2/1 bbcredit
fc2/1 is trunking
  Transmit B2B Credit is 500
  Receive B2B Credit is 500
  Receive B2B Credit performance buffers is 0
    500 receive B2B credit remaining
    500 transmit B2B credit remaining
    500 low priority transmit B2B credit remaining
    500 low priority transmit B2B credit remaining
```

Use the **show interface port/slot counters** command to display the Tx and Rx BB credit transitions to zero information:

```
switch# show interface fc1/13 counters
.
.
.
  33 Transmit B2B credit transitions to zero
  394351077 Receive B2B credit transitions to zero
.
.
.
```

Use the **show interface port/slot counters detailed** command to check for credit loss recovery:

**Note**

- In the **show interface port/slot counters [detailed]** command output, the *Transmit B2B credit transitions to zero* counter increments every time the transmit buffer-to-buffer credits goes to zero. When the ISLs are configured in the TX credit double-queue mode using the **system default tx-credit double-queue** command, some TX B2B credits are reserved for high-priority traffic and remaining credits are used for low-priority traffic from the total TX B2B credits configuration. Hence, when ISLs are in TX credit double-queue mode, this counter does not increment though the low-priority credits go to zero because the high-priority credits are still available.
- This command output is applicable for Cisco MDS NX-OS Release 8.4(2) and later releases. The command output varies if you are using Cisco MDS NX-OS Release 8.4(1a) or earlier releases.

```

switch# show interface fc1/4 counters detailed
fc1/4
  Rx 5 min rate bit/sec:                0
  Tx 5 min rate bit/sec:                0
  Rx 5 min rate bytes/sec:              0
  Tx 5 min rate bytes/sec:              0
  Rx 5 min rate frames/sec:             0
  Tx 5 min rate frames/sec:             0

Total Stats:
  Rx total frames:                      9
  Tx total frames:                      21
  Rx total bytes:                       716
  Tx total bytes:                       1436
  Rx total multicast:                   0
  Tx total multicast:                   0
  Rx total broadcast:                   0
  Tx total broadcast:                   0
  Rx total unicast:                     9
  Tx total unicast:                     21
  Rx total discards:                    0
  Tx total discards:                    0
  Rx total errors:                      0
  Tx total errors:                      0
  Rx class-2 frames:                    0
  Tx class-2 frames:                    0
  Rx class-2 bytes:                     0
  Tx class-2 bytes:                     0
  Rx class-2 frames discards:            0
  Rx class-2 port reject frames:         0
  Rx class-3 frames:                    9
  Tx class-3 frames:                    21
  Rx class-3 bytes:                     716
  Tx class-3 bytes:                     1436
  Rx class-3 frames discards:            0
  Rx class-f frames:                    0
  Tx class-f frames:                    0
  Rx class-f bytes:                     0
  Tx class-f bytes:                     0
  Rx class-f frames discards:            0

Link Stats:
  Rx Link failures:                     0
  Rx Sync losses:                       0
  Rx Signal losses:                     0
  Rx Primitive sequence protocol errors: 0

```

```

Rx Invalid transmission words:                0
Rx Invalid CRCs:                              0
Rx Delimiter errors:                          0
Rx fragmented frames:                          0
Rx frames with EOF aborts:                     0
Rx unknown class frames:                       0
Rx Runt frames:                                0
Rx Jabber frames:                              0
Rx too long:                                   0
Rx too short:                                  0
Rx FEC corrected blocks:                       0
Rx FEC uncorrected blocks:                     0
Rx Link Reset(LR) while link is active:        0
Tx Link Reset(LR) while link is active:        0
Rx Link Reset Responses(LRR):                  0
Tx Link Reset Responses(LRR):                  1
Rx Offline Sequences(OLS):                     0
Tx Offline Sequences(OLS):                     1
Rx Non-Operational Sequences(NOS):             0
Tx Non-Operational Sequences(NOS):            0

Congestion Stats:
Tx Timeout discards:                           0
Tx Credit loss:                                0
BB_SCs credit resend actions:                  0
BB_SCr Tx credit increment actions:            0
TxWait 2.5us due to lack of transmit credits: 0
Percentage TxWait not available for last 1s/1m/1h/72h: 0%/0%/0%/0%
Rx B2B credit remaining:                       32
Tx B2B credit remaining:                       16
Tx Low Priority B2B credit remaining:           16
Rx B2B credit transitions to zero:              1
Tx B2B credit transitions to zero:              2

Other Stats:
Zone drops:                                    0
FIB drops for ports 1-16:                      0
XBAR errors for ports 1-16:                    0
Other drop count:                              0

Last clearing of "show interface" counters :   never

```