Fibre Channel Traffic Issues with MXP−MR−10DME−C Linecards

Document ID: 117543

Contributed by Sunil Poonia, Cisco TAC Engineer.

Mar 25, 2014

Contents

Introduction

This document describes problems that are observed when Storage Area Network (SAN) switches use an MXP−MR Series Linecard in order to transport Fibre Channel (FC) traffic. This document is intended to consolidate all known issues, defects, and their solutions.

Note: Reference the 11.12 MXP_MR_10DME_C and MXP_MR_10DME_L Cards section of the ***Cisco ONS 15454 DWDM Configuration Guide, Release 9.8*** for additional information about this card.

Note: Log into the Cisco Technical Support Website for more information or log into the Cisco World Contacts web page in order to obtain a directory of toll−free Technical Support numbers for your country.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Multi−Service Transport Platform (MSTP) systems, concepts, and hardware
- Cisco transport controller
- Terms used in FC traffic

Components Used

The information in this document is based on these hardware and software versions:

- MXP−MR (15454−10DME−C) Linecard
- Optical Networking System (ONS) 15454 MSTP

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

The Cisco MXP−MR−10DME−C linecard is used in order to aggregate a mix of client SAN service inputs (GE, FICON, and FC) into one OUT−2/STM−64/OC−192 DWDM tunk−side signal.

Problems have been observed with FC traffic through this linecard. These problems can be fluctuations on the switch interface, errors on MXP−MR−10DME−C client interfaces, errors on the SAN switch interface, or interoperability issues.

For greater clarification of the information described in this document, contact the Cisco Technical Assistance Center (TAC).

MXP−MR−10DME Issues with FC Traffic

Interoperability Issues

This section describes interoperability issues between Cisco 9500 Series Multilayer Data Switches (MDS9500) and ONS 15454−10DME Series Linecards that use 4−G FC interfaces.

For ONS Dense Wave Division Multiplexing (DWDM) connections that use the Distance Extension (DE) feature (also known as buffer−to−buffer credit spoofing), the Fibre Channel Buffer−to−Buffer State Change Notification (FCBBSCN) option must be disabled on the Inter−Switch Link Protocols (ISLs).

Note: Reference the Configuring Interface Buffers section of the ***Cisco MDS 9000 Family NX−OS Interfaces Configuration Guide*** for more information.

For the ONS with the 10DME linecard, at two Gb/s the ONS drops the buffer−to−buffer (B2B) recovery frames and does not pass them, even with DE enabled. However, at four Gb/s the ONS does pass the B2B frames. This causes the interoperability issue with the FCBBSCN feature.

In order to resolve this behavior, complete one of these two steps:

1. Configure the MDS9500 with ***no switchport fcbbscn***.

2. Disable the DE feature on the 10DME linecard ports.

Increment of *fcStatsRxRecvrReady* and *fcStatsTxRecvrReady*

If an increment for only ***fcStatsRecvrReady*** and ***fcStatsTxRecvrReady*** is observed from the Performance tab on MXP−MR−10DME−C linecards and no other parameter increments are seen when traffic hits, use the solution described in this section.

Problem Summary

The loss of a Receiver−Ready (***R_RDY***) transmission word prevents the release of a credit buffer. By default, the switch ports initialize links with the Exchange Link Parameters (ELP) mode ***1***. However, gateways expect initialization with ELP mode ***2***, which is also referred to as ***ISL R_RDY*** mode. Therefore, in order to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode ***2***.

As each host transmits a frame to the switch, the switch reads the SID and Domain ID (DID) in the frame header. If the DID of the destination address is the same as that of the switch (intra−switch communications), the frame buffer is copied to the destination port, and a credit ***R_RDY*** is sent to the host. The switch only needs to read word zero and word one of the FC frame in order to perform what is known as *cut−through routing*. A frame might begin to emerge from the output port before it is entirely received by the input port. The entire frame does not need to be buffered in the switch.

B2B flow control occurs through a sending port with its available credit supply and waits to have the credits replenished by the port on the opposite end of the link. These B2B credits are used by Class 2 and Class 3 services and rely on the FC ***R_RDY*** control word that is sent from the receiving−link port to the sender.

The rate of frame transmission is regulated by the receiving−link port based on the ability of the buffers to hold received frames.

Disable R_RDY Mode EMC SAN Switch Ports

Here is an example of the old switch interface configuration:

- Traffic Isolation is disabled
- Trunking is disabled
- ***ISL_R_RDY*** mode is enabled

This is the new configuration that resolves the interoperability issue previously described:

- Traffic Isolation is disabled
- Trunking is disabled
- ***ISL_R_RDY*** mode is disabled

Conclusion

The problem is based on a known issue with EMC switches and MXP−MR−10DME linecards.

In order to resolve this issue, these changes are made on the EMC switch:

- Traffic Isolation is disabled
- Trunking is disabled
- ***ISL_R_RDY*** mode is disabled
- Both switches are set for ELP mode ***2***

Note: These errors can also occur if the switch does not support DE or if it does not support the required distance of DWDM. Check with the switch vendor for more information, because the switch software/hardware might need an upgrade.

Intermittent Output Discards Observed on Switch Interface

Summary

In this scenario, output discards of packets are observed on the switch interface and no alarms/conditions are observed on ONS 15454 systems. On the MXP−MR−10DME−C linecard, increments of ***fcStatsRxRecvrReady*** and ***8b10bInvalidOrderedSetsDispErrorsSum*** are observed.

Resolution Summary

Verify the software version of the affected ONS 15454 system. If the software is Version 8.50, 8.51, or 8.52, then there is a software defect that causes the issue.

In order to resolve this issue, ONS must be upgraded to software Version 9.1.0.

Cisco Bug ID CSCsr75681

These are the defect symptoms:

- Packets are lost and then traffic resumes.
- Output discards are observed at the interface of the MDS9513 switch.
- No errors are reported on the CTC.

The condition of the defect is that the 10DME linecard is attached to the MDS9513 switch and configured with:

- 4−G FC
- DE on E−E

After the ONS software upgrade, check the Field Programmable Gate Array (FPGA) version on the MXP−MR−10DME−C linecard:

1. Log into Cisco Transport Controller (CTC).
2. Navigate to ***Card View*** > ***Maintenance*** > ***Info***.

The latest versions of the FPGAs are ***RAILTO_SOUTH 1.41*** and ***RIALTO_NORTH 2.35,*** tested in Version 9.222 and available in all new releases. Reference the FPGA section for more information.

If the newest FPGA versions are not available after the software upgrade, then perform a Force FPGA update. Reference the Upgrading the Cisco ONS 15454 to Release 9.1 Cisco article for more information about upgrades to this system.

Note: Log into the Cisco Technical Support Website for more information or log into the Cisco World Contacts web page in order to obtain a directory of toll−free Technical Support numbers for your country.

Increment of *mediaIndStatsTxFramesBadCRC* Errors on an MXP−MR−10DME Client Interface

Summary

Both switch interfaces report intermittent CRC errors. On the MXP−MR−10DME linecard, an increment of ***TxBadCRC*** errors is observed from the client port.

Resolution Summary

Verify the software version of the ONS 15454 node. If the software is Version 7.0 through 8.52, then the system is affected by a software defect.

In order to fix this issue, the ONS node must be upgraded to software Version 9.1.0.

Cisco Bug ID CSCsm50360

These are the defect symptoms:

- Single bit error events (error rate around 1E−12) are generated by the 10DME linecard in the egress (TX−out) direction.
- Errors are captured by the ***mediaIndStatTXFramesBadCRC*** counter.

The condition of the defect is that 4−G FC mode is used on ***Port 1*** or ***Port 5*** (***Port 1*** is mostly affected).

After the ONS software upgrade, check the Field Programmable Gate Array (FPGA) version on the MXP−MR−10DME−C linecard:

1. Log into Cisco Transport Controller (CTC).
2. Navigate to ***Card View > Maintenance > Info***.

The latest versions of the FPGAs are ***RAILTO_SOUTH 1.41*** and ***RIALTO_NORTH 2.35,*** tested in Version 9.222 and available in all new releases. Reference the FPGA section for more information.

If the newest FPGA versions are not available after the software upgrade, then perform a Force FPGA update. Reference the Upgrading the Cisco ONS 15454 to Release 9.1 Cisco article for more information about upgrades to this system.

Note: Log into the Cisco Technical Support Website for more information or log into the Cisco World Contacts web page in order to obtain a directory of toll−free Technical Support numbers for your country.

Known Software Defects

ONS software Versions 7.0 through 8.50 have defects related to the MXP−MR−10DME linecard and FC traffic. This section describes the known defects.

Cisco Bug ID CSCsc36494

Description

Cisco bug ID title: ***Manual Y cable switches with squelching turned off can cause a Fibre channel link with brocade switches to go down***.

Known Affected Software Releases

This defect is observed in software Versions 7.0 and later.

Conditions

Complete these steps in order to reproduce this issue:

1. Set up the MXP_MR_10DME linecards so that they are Y−cable protected. Provision squelching to be off. DE is enabled.
2. Ensure that the path between the pair of Y−cable protected linecards has no distance introduced, but the protect path has a delay of 800 km introduced.
3. Begin FC traffic with brocade switches.
4. Perform a user−initiated, manual Y−cable switch from CTC.

After a few switchovers, the FC link goes down. SIGLOSS and GFP−CSF alarms are observed.

Workaround

Cisco recommends that you provision squelching to be on when interworking with brocade switches. If for some reason, squelching must be off with brocade switches, Cisco recommends that you use a ***FORCE*** command in order to perform a Y−cable switch.

Expected Resolution

There is no known resolution for this issue. Cisco recommends that you apply the workaround.

Cisco Bug ID CSCsh71385

Description

Cisco bug ID title: ***10DME−C:FC traffic unexpected degradation when interoperating w/ Brocade***.

With 1− or 2−G FC traffic, throughput occurs as expected in one direction, while the opposite direction shows issues. The initial throughput is normal but slowly decreases to zero due to credit lost. The switch normally triggers LR in order to restore the link. The MXP−10DME linecard might delete the ***VC−RDY*** when the interpacket gap is ARB instead of Idle.

Known Affected Software Releases

Software Versions 7.0 through 8.50 are affected by this defect.

Conditions

The traffic is FC (1− or 2−G FC) and is Class 3 Virtual Class (***VC−RDY*** is used instead of ***R−RDY***). The problem occurs when the MXP−10DME linecard performs negative rate compensation (interpacket gap removal).

Workaround

Use the ISC (***R_rdy***) mode on the switch.

Expected Resolution

There is a fix for this issue included in software Version 8.52.

Cisco Bug ID CSCsj42162

Description

Cisco bug ID title: ***Traffic not restored in a chain of mxp−mr−10dme***.

Some packets are corrupted with CRC errors (***mediaIndStatsRxFramesBadCRC*** increments on Payload/Statistics) in the MXP−MR−10DME linecards when the source Ethernet signal is dropped and then reintroduced in a daisy−chain setup of MXP−MR−10DME linecards (connected back and forth with each other).

Known Affected Software Releases

This defect is observed in software Version 7.3.

Conditions

There must be a daisy−chain setup with at least four MXP−MR−10DME linecards, and the source Ethernet signal must be dropped and reintroduced.

Workaround

The Admin state of every SFP in the path of the lost carrier must be moved to ***OOS−DSBLD*** and then to the ***IS*** state.

Expected Resolution

There is a fix for this defect included in software Version 8.51.

Cisco Bug ID CSCsm50360

Description

Cisco bug ID title: ***10dme Egress CRC errors in 4FC mode***.

There are single bit error events (error rate around 1E−12) that are generated by the 10DME linecard in the egress (TX out) direction. Errors are captured by the ***mediaIndStatTXFramesBadCRC*** counter.

Known Affected Software Releases

This defect is observed in software Versions 7.0 through 8.51.

Conditions

The 4−G FC Mode is used on port ***1*** or ***5***. Port ***1*** is mostly affected.

Workaround

There is no known workaround for this defect.

Expected Resolution

A fix for this defect is included in software Version 8.52.

Cisco Bug ID CSCso92457

Description

Cisco bug ID title: ***4GFC switching times of the order of minutes***.

After an MXP−MR−10DME linecard trunk switch occurs, traffic is permanently down or comes back after four to five minutes.

Known Affected Software Releases

This defect is observed in software Version 9.0.

Conditions

Complete these steps in order to reproduce the issue:

1. Provision the MXP−MR−10DME linecards with 4−G FC DE enabled and ensure that they are Y−cable protected.
2. Extract the trunk RX fiber.
3. Ensure that the MXP−MR−10DME linecard raises LOS−P on the trunk and switches accordingly.
4. Verify that traffic is back after 30 seconds (link renegotiation with the switch).

Workaround

There is no known workaround for this defect.

Expected Resolution

A fix for this defect is included in software Version 9.1.

Cisco Bug ID CSCsq46283

Description

Cisco bug ID title: ***MXP−MR−10DME : FC4G IS −> OOSMT −> IS causes slow continuous packet loss.***

Some packets are continuously lost for ten to fifteen minutes.

Known Affected Software Releases

This defect is observed in software Version 8.52.

Conditions

Complete these steps in order to reproduce the issue:

1. Provision the MXP−MR−10DME with 4−G FC or 4−G Ficon and DE enabled.
2. Ensure the port is configured similar to this: IS > OOS, MT > IS.
3. Verify that some packets are lost.

Workaround

Configure the port similar to this: OOS, DSBLD > IS.

Expected Resolution

A fix for this defect is included in software Version 9.00.

Cisco Bug ID CSCsr41096

Description

Cisco bug ID title: ***Interoperability 4G FC with Brocade Silkworm.***

There is no alarm on the CTC when the SAN switch port comes offline/online.

Known Affected Software Releases

This defect is observed in software Version 8.52.

Conditions

The linecard is setup for 4−G FC traffic and uses mode ***E,*** and DE is either enabled or disabled. The problem is related to a particular sequence with small packets of any size (36 bytes, or packets with zero byte payloads).

Workaround

From the MDS switch, set the TE port to ***Trunking Mode: ON***. There is no known workaround on the brocade.

Expected Resolution

A fix for this defect is included in software Version 9.0.

Cisco Bug ID CSCsr75681

Description

Cisco bug ID title: ***MXP−MR−10DME−C:FC4G from MDS9513:DE ON:packet drop***.

Packets are lost, and traffic resumes. Output discards are observed at the interface of the MDS 9513 Series Switch. No errors are reported on the CTC.

Known Affected Software Releases

This defect is observed in software Version 8.52.

Conditions

The 10DME linecard is attached to the MDS 9513 Series Switch and is configured similar to this:

- 4−G FC is used
- DE is enabled
- E−E configuration is used

Note: This problem is also observed when DE is disabled on the MXP linecard.

Workaround

There is no known workaround for this defect.

Expected Resolution

A fix for this defect is included in software Version 9.0.

Cisco Bug ID CSCsr93501

Description

Cisco bug ID title: ***Unexpected throughput drop with DE ON and few credits on Brocade/Qlogic.***

The link cannot achieve full−rate throughput. No frames are lost. The MXP−MR−10DME linecard Performance Monitoring reports ***fcStatsZeroTxCredits*** equal to ***fcStatsRxRecvrReady*** and almost equal to ***fcStatsTxRecvrReady***. This means that the MXP−MR−10DME linecard works with ***0 Tx credits***, as if the link is congested.

Known Affected Software Releases

This defect is observed in software Version 8.52.

Conditions

The system is set up similar to this: TestSet > Brocade > MXP−MR−10DME > MXP−MR−10DME > Brocade > TestSet. The brocade is configured with the ISL port (E port) in ***R_RDY*** mode with the ***portCfgISLMode 1*** command. There are eight credits on port ***F*** and eight credits on port ***E*** reported by the brocade GUI interface. Also, the MXP−MR−10DME linecard DE is enabled.

Workaround

There is no known workaround for this defect.

Expected Resolution

A fix for this defect is included in software Version 9.0.

Note: This problem is almost entirely reproduced with brocade and with Q−logic FC switches (16 credits only). MDS switches are not observed to be affected by this problem.

FPGA

There are two FPGAs on each MXP−MR−10DME−C linecard:

- ***RIALTO_NORTH*** for ports ***1*** to ***4***.
- ***RIALTO_SOUTH*** for ports ***5*** to ***8***.

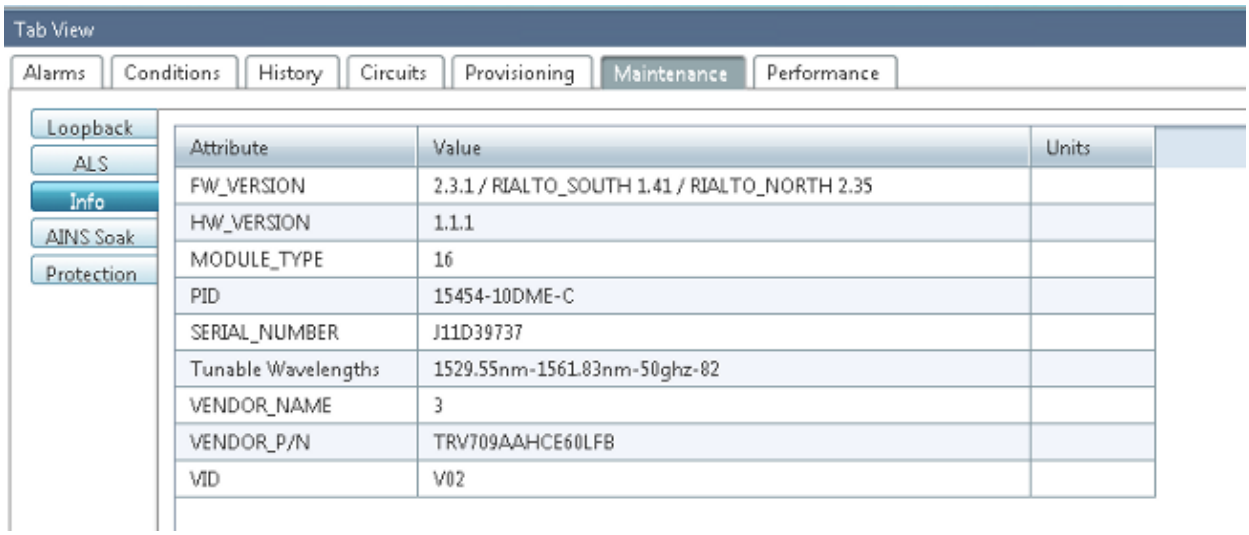

FPGA information can be obtained from CTC. In order to obtain this information while in Card view, click ***Maintenance*** and then ***Info***. Under the ***Info*** tab, the ***FW_VERSION*** (Firmware Version) contains information about both FPGA versions. Ports ***1*** to ***4*** and ports ***5*** to ***8*** are configurable for 1−G FC or 4−G FC traffic.

Note: The FPGA versions shown in the image for both 4−G FC and 1−G FC (and firmware versions) are the latest versions, available in Versions 9.2.2 and later.

Recommendation

Cisco recommends that you upgrade the software versions of systems that run FC traffic to at least Version 9.1.0 in order to avoid the known software defects.

Download Software Version

Reference the Cisco software download page in order to download the latest Cisco ONS 15454 M12 Series MSTPsoftware.

Note: Log into the Cisco Technical Support Website for more information or log into the Cisco World Contacts web page in order to obtain a directory of toll−free Technical Support numbers for your country.

Updated: Mar 25, 2014 Document ID: 117543