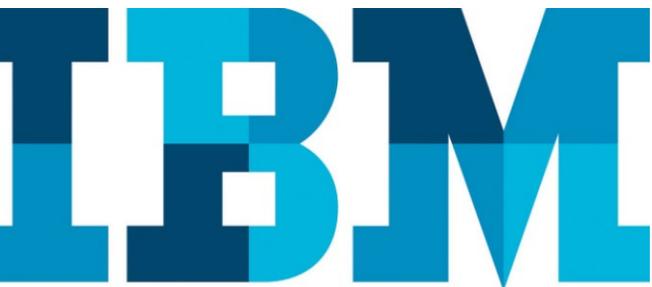


Troubleshooting link-related issues with Ethernet connections on IBM Spectrum Virtualize

Identify and fix network connectivity issues

Table of contents

<i>Introduction.....</i>	<i>2</i>
<i>Possible Ethernet link issues</i>	<i>2</i>
<i>Forward error correction (FEC) related link issues</i>	<i>4</i>
<i>Summary</i>	<i>6</i>
<i>Get more information.....</i>	<i>7</i>
<i>About the authors</i>	<i>7</i>



Overview

Challenge

IBM spectrum Virtualize supports different types of Ethernet traffic with a variety of adapters from different vendors. Currently, troubleshooting Ethernet link-related issues such as inactive link port, unreachable IP, and so on is a challenge and there is no troubleshooting guide available for identifying and resolving such issues.

Solution

This paper provides troubleshooting steps for identifying and fixing Ethernet link related issues on IBM Spectrum Virtualize.

Introduction

IBM® Spectrum® Virtualize supports various Ethernet traffic such as: host attach, clustering, IP replication, and storage virtualization with backend connected storage.

Ethernet link speed and standards are rapidly evolving, and so are the interconnect standards between the Ethernet ports. Given that, there are many ways in which an Ethernet port on the host or the storage system can be connected to an Ethernet switch. And this includes the possibility of connecting the mismatched speeds at the two ends. Apart from mismatched speeds there could also be other issues that cause an incompatibility and hence lead to a link down issue.

The objective of this white paper is to outline the ways to identify and troubleshoot link-related issues on IBM Spectrum Virtualize.

Possible Ethernet link issues

This section lists some of the possible Ethernet link issues and ways to identify and fix them.

Adapter's port link light is off

Possible steps to troubleshoot:

- Verify by changing the switch port to which this adapter port is connected.
- Verify by changing the cable.
- Verify by changing Small Form-factor Pluggable (SFP) in case of optical fiber interconnect.
- Verify mismatched forward error correction (FEC) settings at both ends, that is, switch and adapter port or FEC Auto-Negotiation must be set on both the ends for the connection to work. For more details, refer to section "[Forward error correction \(FEC\) related link issues](#)".
- If nothing works check if the adapter is at fault.

Link is up but target is unreachable from host

Steps to troubleshoot:

- IBM Spectrum Virtualize provides the `svctask ping` command to verify the network connectivity between the initiator and the target IPs. Use the following command for IPv4 addresses:
`svctask ping -srcip4 <source IP v4> <destination IP v4>`

Or, use the following command for IPv6 addresses:

```
svctask ping -srcip6 <source IP v6> <destination IP v6>
```

- Ensure that the maximum transmission unit (MTU) settings for all the ports for the entire path (including switch, storage, and host ports) are matching.

Use the `svcinfo lspportip` command to find the MTU settings for a storage port for IBM Spectrum Virtualize version earlier than 8.5.0.

Use the `svcinfo lsportethernet` command to find MTU settings for a storage port for IBM Spectrum Virtualize version 8.5.0 or higher.

- Ensure that the VLAN ID is identical for all the ports for the entire path (including switch, storage, and host ports). Use the `svcinfo lspportip` command to view the VLAN ID on IBM Spectrum Virtualize version earlier than 8.5.0 and the `svcinfo lspip` command to view the VLAN ID on IBM Spectrum Virtualize version 8.5.0 or later.

Use the `ifconfig` or `ip` command to check the VLAN details on the Linux® host and `ipconfig` for the Microsoft® Windows® host.

- Ensure that the prefix or the subnet and gateway settings are correct at both the ends. Subnets are used to convert large networks into smaller isolated networks.

For example, if a subnet is 255.255.255.0 or prefix is 24 then the right range for the IPs will be 192.168.10.0 to 192.168.10.254

For subnet 255.255.255.192, there will be four different network ranges: 192.168.123.1-62, 192.168.123.65-126, 192.168.123.129-190, 192.168.123.193-254

- Ensure that there are no duplicate IPs in the network.
- Verify that there are no firewalls set, restricting the communication between storage and host from the host TCP port 3260 and IP replication traffic between clusters that uses the TCP port 3265 for IP replication.

Forward error correction (FEC) related link issues

FEC is a commonly used error detection and correction technique for data transmission over high-speed networks.

To begin with, the initiator (node or port) sends the message with redundancy, which is achieved by adding some checksums. These checksums can be derived using different algorithms such as Base-R / FC FEC, Reed-Solomon (RS) and so on. While receiving the message, the target (node or port) can take the necessary action based on the same algorithm used by the initiator and generates the frame by removing the redundant bits.

For 25 Gb networks, FEC is a must but for 10 Gb networks, FEC is not a required.

FECs play a significant role in link initialization as it must be negotiated between the sender and the receiver during initialization. Link will be up only when a match is achieved and both sides agree to a specific supported FEC type. Else, the link will remain down.

FEC negotiation depends on following factors:

- Adapter port FEC setting
- Adapter's capability to auto-negotiate with switch port
- Switch port FEC setting or its capability to auto-negotiate

For IBM Spectrum Virtualize, in case of iWARP (example: for feature AHB7 for IBM FlashSystem® 9200) there are considerations explained in the following sections.

IBM Spectrum Virtualize with version 8.4.0.0 or earlier

For successful link initialization, FEC negotiation must succeed. However, if negotiation fails, the port becomes inactive. It can be due to any of the following reasons for the Chelsio T6-225 adapter port becoming inactive in IBM Spectrum Virtualize with version earlier than 8.4.0.0:

- By default, auto-negotiation of FEC is disabled on Chelsio T6-225 adapter
- When a Chelsio port is connected to a switch using a copper cable, where auto-negotiation is enabled, the port will show the link as active only if the switch port has AN enabled.
- When a Chelsio port is connected to a switch using an optic cable, where auto-negotiation is disabled, Chelsio T6 chooses FEC as Base-R or RS based on the cable type. And only if the switch port has a matching FEC type, the port will be active else it will remain inactive.

Solution:

- In case of copper cable, have auto-negotiation or FEC-RS on the switch port.
- In case of optic cable, the FEC mismatch problem is fixed with Spectrum virtualize version 8.4.0.0
- A user can upgrade to an interim fix version 8.3.1.4, if the installed code is 8.3.0.0 or 8.3.1.0 With IBM Spectrum Virtualize version 8.4.0.0 or later, IBM introduced a new iWARP driver and firmware that supports auto-negotiation on optic cables as well and hence works well with all FEC types.

Configuring FEC types on switch

FEC is a configurable field on several Ethernet switches. Few switches are capable of configuring it on the *per port basis* and few have a common setting mechanism for all ports at the same time. Switch configuration depends on the switch model and its vendor. This paper includes the details of Cisco Nexus-3232C 100G switch for explaining FEC settings.

Perform the following steps to configure FEC and verify the setting on a Cisco 3232C switch:

1. Identify what FEC is set on a specific switch port using the show interface `<interface id>` command.
For example: show interface Eth1/11/1

```
C100G_2(config-if)# show interface Eth1/11/1
Ethernet1/11/1 is down (Link not connected)
admin state is up, Dedicated Interface
Hardware: 25000 Ethernet, address: e00e.dadc.31b4 (bia e00e.dadc.31b4)
MTU 1500 bytes, BW 25000000 Kbit, DLY 10 usec
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, medium is broadcast
Port mode is trunk
auto-duplex, auto-speed, media type is 100G
Speed is turned on
Auto-Negotiation is turned off, FEC mode is Off
Auto-mdix is turned off
Rate mode is dedicated
Switchport monitor is off
EtherType is 0x8100
EEE (efficient-ethernet) : n/a
Last link flapped 00:22:51
Last clearing of "show interface" counters never
33 interface resets
30 seconds input rate 0 bits/sec, 0 packets/sec
30 seconds output rate 0 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 0 bps, 0 pps; output rate 0 bps, 0 pps
RX
532768962737 unicast packets  22050679 multicast packets  40680011 broadcast packets
532831693427 input packets  564710805354899 bytes
17198482134 jumbo packets  0 storm suppression packets
0 runs  0 giants  1 CRC  0 no buffer
2 input error  0 short frame  0 overrun  0 underrun  0 ignored
0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
0 input with dribble  0 input discard
0 Rx pause
TX
199394364122 unicast packets  178999662 multicast packets  271768198 broadcast packets
199845131982 output packets  190510203561637 bytes
15311135208 jumbo packets
0 output error  0 collision  0 deferred  0 late collision
0 lost carrier  0 no carrier  0 babble  215010 output discard
0 Tx pause
C100G_2(config-if)#
```

Note: This example shows how to set different values for FEC on the initiator. The same values should be set on the target as well to make the link active.

2. FEC negotiation will be disabled if its mode is **off** on both the ends. Change the FEC mode to off on the switch port as shown in the following screen capture.

```
C100G_2(config)# interface Eth1/11/1
C100G_2(config-if)# fec
cl74 cl91 off
C100G_2(config-if)# fec off
C100G_2(config-if)#
```

3. Set FEC to CL74 if Base-R algorithm is to be used for communication on the switch port.

```
C100G_2(config-if)# interface Eth1/11/1
C100G_2(config-if)# fec cl74
C100G_2(config-if)# show interface Eth1/11/1 | grep FEC
  Auto-Negotiation is turned off, FEC mode is CL74
C100G_2(config-if)#
```

4. Set FEC to CL91 if Reed-Solomon algorithm is to be used for communication on the switch port.

```
C100G_2(config-if)# interface Eth1/11/1
C100G_2(config-if)# fec cl91
C100G_2(config-if)# show interface Eth1/11/1 | grep FEC
  Auto-Negotiation is turned off, FEC mode is CL91
C100G_2(config-if)#
```

Summary

This paper includes the crisp details of how to troubleshoot Ethernet network connectivity issues with IBM Spectrum Virtualize. It also describes the problems and the solution associated with Forward Error Correction negotiations.

Get more information

To learn more about IBM Spectrum Virtualize, refer to the following websites:

- Forward Error Correction
https://en.wikipedia.org/wiki/Error_correction_code#Forward_error_correction
- Chelsio release notes
<https://service.chelsio.com/beta/drivers/ChelsioUwire-FBSD-3.3.0.1/Release%20Notes.txt>
- Chelsio user guide
<https://service.chelsio.com/store2/T5/Unified%20Wire/Linux/ChelsioUwire-3.8.0.2/Chelsio-UnifiedWire-Linux-UserGuide.pdf>
- IBM Storwize V7000 documentation
<https://www.ibm.com/docs/en/flashsystem-7x00>
- IBM FlashSystem 9200
<https://www.ibm.com/downloads/cas/US-ENUS120-011-CA/name/US-ENUS120-011-CA.PDF>

About the authors

Aakanksha Mathur leads the test effort for IBM Spectrum Virtualize product chain at IBM System Labs, India. She has been working in the storage domain for more than 13 years with IBM System Labs. She can be reached at aamathur@in.ibm.com

Shrirang Bhagwat is a lead developer in the Spectrum Virtualize team at IBM Systems Labs, India. His area of expertise is Ethernet based Block Storage Devices and has lead development for the IP Replication feature. You can reach Shrirang at shbhagwa@in.ibm.com

Subhojit Roy is a senior technical staff member working at IBM India Labs, Pune. He works as a development architect for IBM Spectrum Virtualize products and is leading IBM's Ethernet mission. He can be reached at subhojit.roy@in.ibm.com



© Copyright IBM Corporation 2021
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.



Please recycle