Mellanox InfiniBand Sep 2013

InfiniBand debug

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™

Agenda

- IB bring-up
  - Topology matching
  - Fabric clean-up
  - SM optimization
- Fabric Debug

IB Bring-Up: Procedure Overview

- Make sure all nodes are
  - Responding
  - Have the same OFED version
  - All HCAs are attached to driver
  - HCAs have expected FW version
- Run SM
- Make sure all ports are in active state
- Loop while (fabric cleaning in idle mode)
- Match topology
- Loop while (fabric cleaning under stress)
- Optimize SM

Fabric clean-up

- Fabric clean is used to filter out bad HW including bad cables, bad ports or bad connectivity
- fabric cleanup algorithm
  1. Zero all counters (ibdiagnet –pc)
  2. Wait X time
  3. Check for errors exceeding allowed threshold during this X time (ibdiagnet –lw 4x –ls 10 –P all=1)
  4. Fix problematic links (re-sit or swap cables, replace switch ports or HCAs)
  5. Go to 1

What are we looking for?

- Fabric configuration Issues (SM)
- Environmental
- Communication errors
- Switch/Module (IPR/FCR) status
- Hardware/Cable
- Fabric topology Issues

Recommended Practice for Cleaning the Fabric

- Zero Port counters
  - *Very important to always start with a clean baseline!*
- Run stress test across fabric

  Pallas (Intel) MPI Benchmark, mpi_bandwidth, mpi_latency, perf_main, etc.
- Identify issues through Port counters
  - *Congestion, Bad links, Packet Loss, etc.*
- Locate and fix problems
  - *Cable Faults, HCA reseating, etc.*

Capture Log Files for support

- **Switch related**
- LOGs.tar from the export logs command
- Portcounters.csv
- ibnetdiscover output

- **Host related**
- Dmesg output
- Any onscreen errors
- Uname –ar
- Cat /etc/issue
- /var/log/messages
- lsmod
- lspci
- ibv_devinfo
- ib-setup

Basic Tier 1 debug HCA

#	Operation description:	Operation performed:	Have this operation resolved the issue?
1	Reseat the card.	☐ Yes ☐ No	☐ Yes ☐ No
2	Replace the cable with a known good cable.	☐ Yes ☐ No	☐ Yes ☐ No
3	Connect the cable to another known working port/s.	☐ Yes ☐ No	☐ Yes ☐ No
4	Swap the card with a known good card.	☐ Yes ☐ No	☐ Yes ☐ No Check “No” if the issue migrated with the faulty card

5	**Reboot the server.**	**☐ Yes ☐ No**	**☐ Yes ☐ No**
6	Ports LED indicators state:		
7	Device information which the card is connected to:		
8	Active SM location:		

Basic Tier 1 debug HCA

9	**Environment information**	
	cat /etc/issue	Click here to enter text.
	uname –a	Click here to enter text.
	cat /proc/cpuinfo | grep ‘model name’ | uniq	Click here to enter text.
	ofed_info | head -1	Click here to enter text.
	ifconfig –a	Click here to enter text.
	ethtool <interface>	Click here to enter text.
	ethtool –i <interface_of_Mellanox_port_num>	Click here to enter text.
	ibdev2netdev	Click here to enter text.
10	Card Detection	
	lspci | grep –i Mellanox	Click here to enter text.

11	**Mellanox Firmware Tool (MFT)**	
	Download and install MFT: http://www.mellanox.com/content/pages.php?pg=management_tools&menu_section=34 Refer to the User Manual for the installation instructions. Once installed, run: mst start mst status	
	flint –d <mst_device> q	Click here to enter text.
12	Ports Information	
	ibstat	Click here to enter text.
	Ibv_devinfo	Click here to enter text.
13	Firmware Version Upgrade Download latest firmware version using the PSID/board ID: http://www.mellanox.com/supportdownloader/ flint –d <mst_device> -i <firmware_bin_file> b	☐ Yes ☐ No
14	Collect log file /var/log/messages dmesg > system.log	

Basic Tier 1 debug 1U Switch

#	Operation description:	Operation performed:	Have this operation resolved the issue?
1	Power cycle the switch.	☐ Yes ☐ No	☐ Yes ☐ No
2	If the issue relates to port/s functionality - Replace the cable/s with known good cable/s (of the same type).	☐ Yes ☐ No	☐ Yes ☐ No
3	If the issue relates to port/s functionality - Connect the port/s to other known working port/s (destinations).	☐ Yes ☐ No	☐ Yes ☐ No
4	If the issue relates to port/s functionality – Perform loopback test with other working port/s on the same switch.	☐ Yes ☐ No	☐ Yes ☐ No Check “No” if the port/s remained faulty while connected to other port/s on the same switch

Basic Tier 1 debug 1U Switch

5	**Switch Front Status LED indicators state: Click here to enter text.**	
6	Switch Rear Status LED indicators state: Click here to enter text.	
7	PSU module Status LED indicators state: Click here to enter text.	
8	FAN module Status LED indicators state: Click here to enter text.	
9	Active SM location (switch/Open SM/UFM): Click here to enter text.	
10	Relevant port/s Status LED indicators state: Click here to enter text.	

Basic Tier 1 debug 1U Switch

11	**Software Version Upgrade**	☐ Yes ☐ No
	SX60xx series: Verify software version currently installed: *show version* If software upgrade is needed, please use the following link in order to download the latest software version available and documentation: http://support.mellanox.com/SupportWeb/Switches/infiniband_switches/SX60XX Follow the upgrade instruction in the User Manual, chapter 4.3.	Check “Yes” if the switch software version installed equals to latest

Host Troubleshooting

- **Bad cabling**
- **IPoIB Interface problem**
- **Missing Configuration**
- **HCA problem**
- **SM problem**

Let’s start by checking the basics

Troubleshooting (Cont)

1. Be sure cables are plugged in properly.
2. Check that the SM is running

Login to the master Switch CLI

- Run the command sm-info show and make sure that sm mode is enabled and sm state is master
- Run the command sm-info show few times , make sure sm activity counter is progressing
- In case the sm state is not master it means that other switch or node in the fabric is running another SM that may be the master

Troubleshooting
Make sure HCA is working

1. **Run lspci check that you see mellanox HCA is identifiede on the PCI bus**
2. **If not reseat HCA or the raiser card**
3. **Replace HCA with another**
4. **Check that the Host links are active**

Troubleshooting (Cont)
Networking

1. Check that the IPoIB interface is up
2. Run ifconfig –a to view all network interface it might be that the ib0 or ib1 is there but not activated
3. Run ifconfig make sure your IPoIB interface is configured (in case its not use ib-config to configure it <ib-config –h for help>)
4. If not ACTIVE, you should have no LED's on the HCA
5. Check that SM in running

Troubleshooting (Cont) Networking

6. Check that you can ping between nodes on IPoIB
7. Run the command ifconfig and make sure the following line appears exactly at your IB interface:
8. If the RUNNING The IPOIB host is not joined to the IB multicast group. In this case check the SM health.
9. Check for IP problems such as duplicate IP, wrong routing table or wrong destination address
10. If not check to see you have latest firmware on the switches and HCA ASIC
11. Run Cable and Link Tests

17

ibdiagnet

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™

Cluster utilities - ibdiagnet / ibdiagpath

- Integrated diagnostic tools
  - Queries cluster topology and indicates any port errors, link width, or link speed mismatch.
  - Automates calls to many “low level” operations
- Easy to use
  - Similar flags, logs and reports for both tools
  - Report using meaningful names when topology file is provided

Ibdiagnet Tool

- Ibdiagnet is an integrated Infiniband fabric diagnostics command line tool.
- It scans the IB fabric using directed / lid route packets and extracts the available information regarding its connectivity and devices status
- It then checks for errors in the following scopes:
  - Ports (Counters thresholds, port state)
  - Nodes (Firmware versions, LID assignmets)
  - Links (Links speed and width, Cables info)
  - Fabric (Topology matching, Subnet Manager, Routing)
- Errors are reported to screen and saved in a log file

ibdiagnet

ibdiagnet scans the fabric using directed / lid route packets and extracts all the available information regarding its connectivity and devices.
It then checks errors on ports, nodes, links and cluster scopes and reports them. ibdiagnet is included in the ibutils package which is part of Mellanox OFED.

Common usage (example):

ibdiagnet -pc –r -ls 14 -lw 4x --get_cable_info --pm_pause_time [Time for test in sec (e.g 1200)] -o /var/ibdiagnet2_`date +%F_%H_%M_%S

- **-pc-** Perform a clear counters fabric wise
- **-r** -Check for routing issues
- **-lw** <1x|4x|12x> **-ls** <2.5|5|10|14>
  - Link speed and width checked on every port in the network
- **--get_cable_info** - Read the cable info – type, length, manufacturer, etc…
- **--pm_pause_time** <T>
  - Time to sleep before resume collecting counters
- **-o** <out-dir> Output directory

ibdiagnet -pc –r -ls 14 -lw 4x --get_cable_info --pm_pause_time [Time for test in sec (e.g 1200)] -o /var/ibdiagnet2_

- ibdiagnet -pc –r -ls 10 -lw 4x --get_cable_info --pm_pause_time 200

Output files

- `ibdiagnet.log` - A dump of all the application reports generate according to the provided flags
- `ibdiagnet.lst` - List of all the nodes, ports and links in the fabric
- `ibdiagnet.fdbs` - A dump of the unicast forwarding tables of the fabric switches
- `ibdiagnet.mcfdbs` - A dump of the multicast forwarding tables of the fabric switches
- `ibdiagnet.sm` - List of all the SM (state and priority) in the fabric
- `ibdiagnet.pm` - A dump of the pm Counters values, of the fabric links
- `ibdiagnet.db_csv` - A dump of the internal subnet database.

Ibdiagnet usage (Fabric Cleaning)

- Ibdiagnet is particularly useful in finding misconfigured links (speed/width, topology mismatches, and marginal link/cable issues.
- Typical usage:
  - Clear all port counters using ‘ibdiagnet –pc’
  - Stress the cluster
  - Check cluster using ‘ibdiagnet -P all=1 -ls 10 -lw 4x -pc --get_cable_info --pm_pause_time
    - Checks for link speed, link width, and port error counters greater than 1

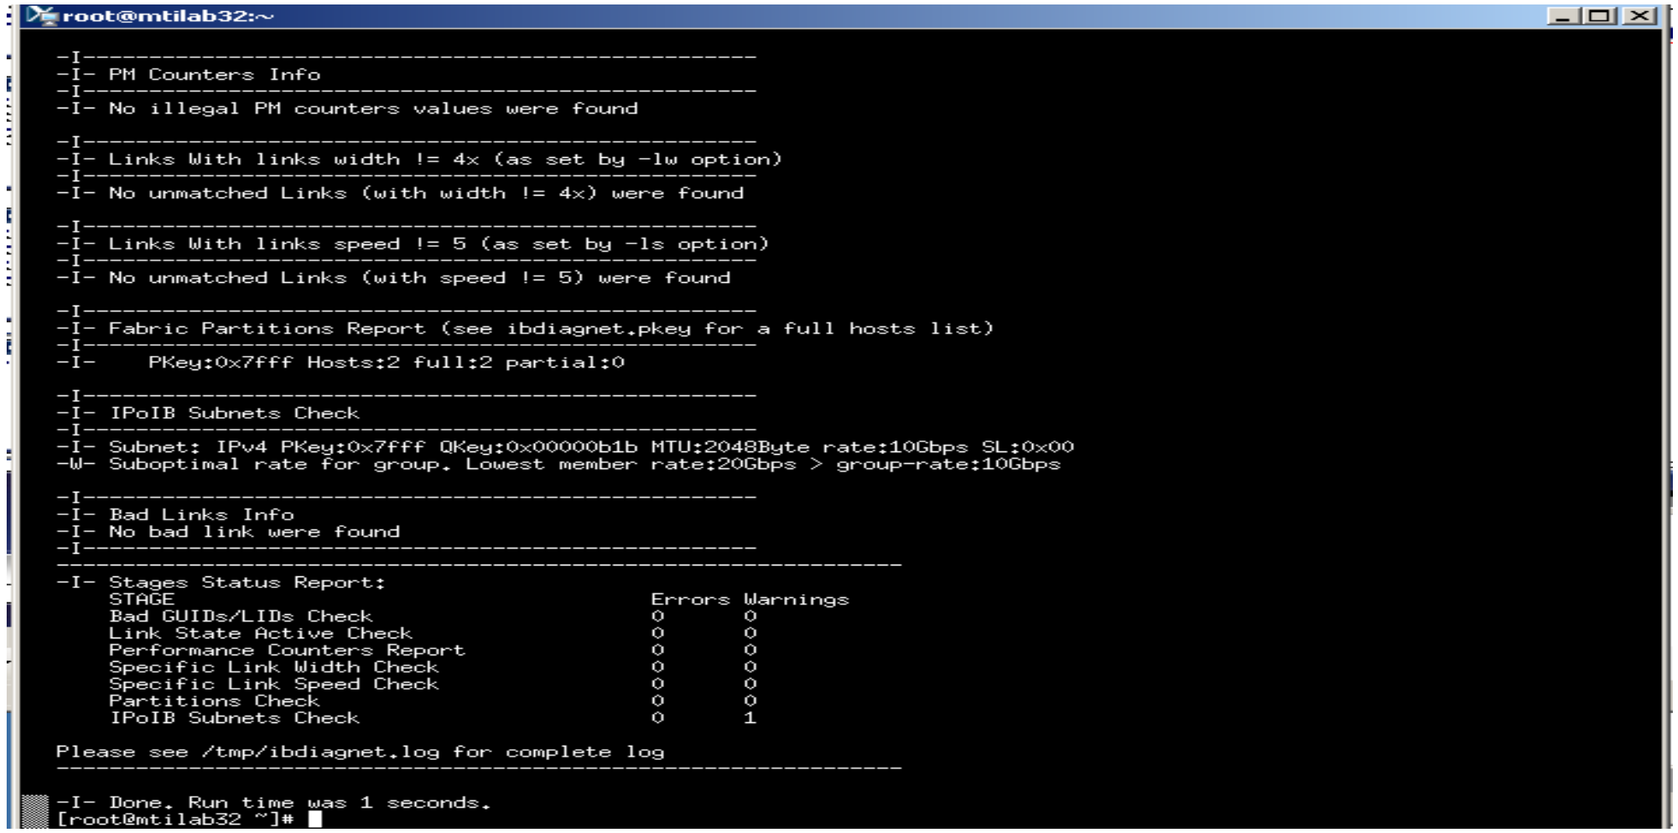

Cluster utilities - ibnetdiscover

- Reports a complete topology of cluster
- Shows all interconnect connections reporting:
  - Port LIDs
  - Port GUIDs
  - Host names
  - Link Speed
- GUID to name file can be used for more readable topology in regards to switch devices

Error counter review

- SymbolErrors
  - **Total number of minor link errors. Usually an 8b/10b error due to a bit error**
- Link Recovers
  - **Total number of times the Port Training state machine has successfully completed the link error recovery process.**
- LinkDowned
  - **Total number of times the Port Training state machine has failed the link error recovery process and downed the link.**
- RcvErrors
  - **Total number of packets containing an error that were receive on the port. Usually due to a CRC error caused by a bit error within the packet.**
- RcvSwRelayErrors
  - **Total number of packets received on the port that were discarded because they could not be forwarded by the switch relay. This counter should typically be ignored since Anafa-II has a bug that counts these when it gets a multicast packet on a port where that port also belongs to the multicast group of the packet.**
- XmtDiscards
  - **Total number of outbound packets discarded by the port because the port is down or congested. Usually due to the output port HOQ lifetime being exceeded.**
- VL15Dropped
  - **Number of incoming VL15 packets dropped due to resource limitations (e.g., lack of buffers) in the port**
- XmtData,RcvData
  - **Total number of 32-bit data words transmitted and received on the port.**
- XmtPkts,RcvPkts
  - **Total number of data packets transmitted and received on the port.**

Error counter review

- SymbolErrors
  - **Total number of minor link errors. Usually an 8b/10b error due to a bit error**
- Link Recovers
  - **Total number of times the Port Training state machine has successfully completed the link error recovery process.**
- LinkDowned
  - **Total number of times the Port Training state machine has failed the link error recovery process and downed the link.**
- RcvErrors
  - **Total number of packets containing an error that were receive on the port. Usually due to a CRC error caused by a bit error within the packet.**
- RcvSwRelayErrors
  - **Total number of packets received on the port that were discarded because they could not be forwarded by the switch relay. This counter should typically be ignored since Anafa-II has a bug that counts these when it gets a multicast packet on a port where that port also belongs to the multicast group of the packet.**
- XmtDiscards
  - **Total number of outbound packets discarded by the port because the port is down or congested. Usually due to the output port HOQ lifetime being exceeded.**
- VL15Dropped
  - **Number of incoming VL15 packets dropped due to resource limitations (e.g., lack of buffers) in the port**
- XmtData,RcvData
  - **Total number of 32-bit data words transmitted and received on the port.**
- XmtPkts,RcvPkts
  - **Total number of data packets transmitted and received on the port.**

Cluster utilities

- ibswitches
  - Lists all switches in cluster

```
[root@ws203 sbin]# ibswitches
Switch  : 0x0008f1040041082a ports 24 "ISR9024 Voltaire" enhanced port 0 lid 1 lmc 0
Switch  : 0x0008f104003f1090 ports 24 "ISR9288/ISR9096 Voltaire sLB-24D" base port 0 lid 2 lmc 0
Switch  : 0x0008f10400403349 ports 24 "ISR9096 Voltaire sFB-4D" enhanced port 0 lid 4 lmc 0
Switch  : 0x0008f104003f1091 ports 24 "ISR9288/ISR9096 Voltaire sLB-24D" base port 0 lid 3 lmc 0
```

- ibhosts
  - Lists all HCAs in cluster

```
Ca      : 0x0008f1040396a490 ports 2 "ws201 HCA-1"
Ca      : 0x0008f1040396b740 ports 2 "ws200 HCA-1"
Ca      : 0x0008f1040396e6cc ports 2 "ws203 HCA-1"
```

- Ibtracert

Shows path between 2 lids

```
[root@ws203 sbin]# ibtracert 9 7
From ca {0x0008f1040396e6cc} portnum 2 lid 9-9 "ws203 HCA-1"
[2] -> switch port {0x0008f1040041082a}[8] lid 1-1 "ISR9024 Voltaire"
[14] -> switch port {0x0008f104003f1090}[15] lid 2-2 "ISR9288/ISR9096 Voltaire sLB-24D"
[10] -> switch port {0x0008f10400403349}[13] lid 4-4 "ISR9096 Voltaire sFB-4D"
[17] -> switch port {0x0008f104003f1091}[11] lid 3-3 "ISR9288/ISR9096 Voltaire sLB-24D"
[16] -> ca port {0x0008f1040396b741}[1] lid 7-7 "ws200 HCA-1"
To ca {0x0008f1040396b740} portnum 1 lid 7-7 "ws200 HCA-1"
```

Cluster utilities

- Iblinkinfo – Reports link info for each port in an IB fabric, node by node.

Iblinkinfo can be used in a hybrid fabric to identify sub optimal links.

Use: iblinkinfo | grep Could

```
570   22[  ] ==( 4X  10.0 Gbps (FDR10) Active/  LinkUp)==>     617    9[  ] "MF0;i1-ib1:SXX536/L36/U1" ( Could be 14.0625 Gbps)
570   28[  ] ==( 4X  10.0 Gbps (FDR10) Active/  LinkUp)==>     617    3[  ] "MF0;i1-ib1:SXX536/L36/U1" ( Could be 14.0625 Gbps)
58    8[  ] ==( 4X         10.0 Gbps Active/  LinkUp)==>     487    1[  ] "hydraio12 HCA-1" ( Could be 14.0625 Gbps)
58   12[  ] ==( 4X  10.0 Gbps (FDR10) Active/  LinkUp)==>     150    1[  ] "hydraio24 HCA-1" ( Could be 14.0625 Gbps)
```

- 3- allows a basic subset of standard SMP queries including the following: node info, node description, switch info, port info.
  - Common ops:
  - NodeInfo (NI) <addr>
  - NodeDesc (ND) <addr>
  - PortInfo (PI) <addr> [<portnum>]
  - SwitchInfo (SI) <addr>
  - PKeyTable (PKeys) <addr> [<portnum>]

Reading OpenSM log

- Find the SM: sminfo

sm lid 573 sm guid 0x2c90300fe2ed1, activity count 26181972 priority 15 state 3 SMINFO_MASTER

Query node description: smpquery nd 573

Node Description:.....sm2 HCA-1

Thank You