

# **Additional Configuration (Optional)**

# Table of contents

General Settings in gv.cfg
Enabling SHARP Aggregation Manager
Enabling Predefined Groups
Enabling Multi-NIC Host Grouping
Defining Node Description Black-List
Running UFM Over IPv6 Network Protocol
Adding SM Plugin (e.g. lossymgr) to event_plugin_name Option
Multi-port SM
Configuring UDP Buffer
Virtualization
Static SM LID
Configuring Log Rotation
Configuring UFM Logging
Configuring UFM Over Static IPv4 Address
Configuring Syslog
Configuring Unhealthy Ports
Configuration Examples in gv.cfg
Managing Dynamic Telemetry
SM Trap Handler Configuration
Setting CPU Affinity on UFM
Quality of Service (QoS) Support
UFM Failover to Another Port

Configuring Managed Switches Info Persistency
Configuring Partial Switch ASIC Failure Events
Enabling Network Fast Recovery
Disabling Rest Roles Access Control
Enabling/Disabling Authentication
Kerberos Authentication
UFM Authentication Server
Azure AD Authentication
Managing UFM Configuration Files Based on Fabric Size
Configuration File and Parameters
Setting up Telemetry in UFM
Enabling UFM Telemetry
Changing UFM Telemetry Default Configuration
Supporting Generic Counters Parsing and Display
Supporting Multiple Telemetry Instances Fetch
Low-Frequency (Secondary) Telemetry
Low-Frequency (Secondary) Telemetry - Exposing IPv6 Counters
Stopping Telemetry Endpoint Using CLI Command
Exposing Switch Aggregation Nodes Telemetry
Exposing Performance Histogram Counters for Egress Queue Depth Indications (Secondary) Telemetry

# General Settings in gv.cfg

Configure general settings in the conf/gv.cfg file.



When running UFM in HA mode, the gv.cfg file is replicated to the standby server.

#### **Enabling SHARP Aggregation Manager**

SHARP Aggregation Manager is disabled by default. To enable it, set:

[Sharp] sharp\_enabled = true



Upon startup of UFM or SHARP Aggregation Manager, UFM will resend all existing tenant allocations to SHARP AM.

#### **Enabling Predefined Groups**

enable\_predefined\_groups = true

#### (i) Note

By default, pre-defined groups are enabled. In very large-scale fabrics, pre-defined groups can be disabled in order to allow faster startup of UFM.

#### **Enabling Multi-NIC Host Grouping**

multinic\_host\_enabled = true



Upon first installation of UFM 6.4.1 and above, multi-NIC host grouping is enabled by default. However, if a user is upgrading from an older version, then this feature will be disabled for them.

#### j) Note

It is recommended to set the value of this parameter before running UFM for the first time.

#### **Defining Node Description Black-List**



Node descriptions from the black-list should not be used for Multi-NIC grouping.

During the process of host reboot or initialization/bringup, the majority of HCAs receive a default label rather than an actual, real description. To prevent the formation of incorrect multi-NIC groups based on these default labels, this feature offers the option to establish a blacklist containing possible node descriptions that should be avoided when grouping Multi-NIC HCAs during host startup. Once a legitimate node description is assigned to the host, the HCAs are organized into multi-NIC hosts based on their respective descriptions. It is recommended to configure this parameter before initiating the UFM for the first time.

For instance, nodes initially identified with descriptions listed in the exclude\_multinic\_desc will not be initially included in Multi-NIC host groups until they obtain an updated, genuine node description.

Modify the exclude\_multinic\_desc parameter in the cv.fg file:

exclude\_multinic\_desc = localhost,generic\_name\_1,generic\_name\_2

#### **Running UFM Over IPv6 Network Protocol**

The default multicast address is configured to an IPv4 address. To run over IPv6, this must be changed to the following in section UFMAgent of gv.cfg.

```
[UFMAgent]
...
# if ufmagent works in ipv6 please set this multicast address to FF05:0:0:0:0:0:0:0:15F
mcast_addr = FF05:0:0:0:0:0:0:15F
```

# Adding SM Plugin (e.g. lossymgr) to event\_plugin\_name Option

The following options allow users to set the SM plugin options via the UFM configuration. Once SM is started by UFM, it will start the SM plugin with the specified options.

# Event plugin name(s)
event\_plugin\_name osmufmpi lossymgr

Add the plug-in options file to the event\_plugin\_options option:

# Options string that would be passed to the plugin(s) event\_plugin\_options --lossy\_mgr -f <lossy-mgr-options-file-name>

These plug-in parameters are copied to the opensm.conf file in Management mode only.

### **Multi-port SM**

SM can use up to eight-port interfaces for fabric configuration. These interfaces can be provided via /opt/ufm/conf/gv.cfg. The users can specify multiple IPoIB interfaces or bond interfaces in /opt/ufm/conf/gv.cfg, subsequently, the UFM translates them to GUIDs and adds them to the SM configuration file (/opt/ufm/conf/opensm/opensm.conf). If users specify more than eight interfaces, the extra interfaces are ignored.

[Server]

# disabled (default) | enabled (configure opensm with multiple GUIDs) | ha\_enabled (configure multiport SM with high availability)

multi\_port\_sm = disabled

# When enabling multi\_port\_sm, specify here the additional fabric interfaces for OpenSM conf

- # Example: ib1,ib2,ib5 (OpenSM will support the first 8 GUIDs where first GUID will
- # be extracted the fabric\_interface, and remaining GUIDs from additional\_fabric\_interfaces

additional\_fabric\_interfaces =

#### i Note

UFM treats bonds as a group of IPoIB interfaces. So, for example, if bond0 consists of the interfaces ib4 and ib8, then expect to see GUIDs for ib4 and ib8 in opensm.conf.



Duplicate interface names are ignored (e.g. ib1,ib1,ib1,ib2,ib1 = ib1,ib2).

### **Configuring UDP Buffer**

This section is relevant only in cases where telemetry\_provider=ibpm. (By default, telemetry\_provider=telemetry).

To work with large-scale fabrics, users should set the set\_udp\_buffer flag under the [IBPM] section to "yes" for the UFM to set the buffer size (default is "no").

```
# By deafult, UFM does not set the UDP buffer size. For large scale fabrics
# it is recommended to increase the buffer size to 4MB (4194304 bits).
set_udp_buffer = yes
# UDP buffer size
udp_buffer_size = 4194304
```

#### Virtualization

This allows for supporting virtual ports in UFM.

[Virtualization] # By enabling this flag, UFM will discover all the virtual ports assigned for all hypervisors in the fabric enable = false # Interval for checking whether any virtual ports were changed in the fabric interval = 60

#### **Static SM LID**

Users may configure a specific value for the SM LID so that the UFM SM uses it upon UFM startup.

[SubnetManager]

# 1- Zero value (Default): Disable static SM LID functionality and allow the SM to run with any LID.

# Example: sm\_lid=0

# 2- Non-zero value: Enable static SM LID functionality so SM will use this LID upon UFM startup. sm\_lid=0

#### (j) Note

To configure an external SM (UFM server running in sm\_only mode), users must manually configure the opensm.conf file (/opt/ufm/conf/opensm/opensm.conf) and align the value of master\_sm\_lid to the value used for sm\_lid in gv.cfg on the main UFM server.

### **Configuring Log Rotation**

This section enables setting up the log files rotate policy. By default, log rotation runs once a day by cron scheduler.

[logrotate] #max_files specifies the number of times to rotate a file before it is deleted (this definition will be applied to
#SM and SHARP Aggregation Manager logs, running in the scope of UFM).
#A count of 0 (zero) means no copies are retained. A count of 15 means fifteen copies are retained
(default is 15)
max_files = 15
#With max_size, the log file is rotated when the specified size is reached (this definition will be applied
to
#SM and SHARP Aggregation Manager logs, running in the scope of UFM). Size may be specified in
bytes (default),
#kilobytes (for example: 100k), or megabytes (for example: 10M). if not specified logs will be rotated
once a day.
max_size = 3

# **Configuring UFM Logging**

The [Logging] section in the gv.cfg enables setting the UFM logging configurations.

Field	Default Value	Value Options	Description	
level	WARNI NG	CRITICAL, ERROR, WARNING, INFO, DEBUG	The definition of the maub logging level for UFM components.	
smclient_leve l	WARNI NG	CRITICAL, ERROR, WARNING, INFO, DEBUG	The logging level for SM client log messages	
event_log_lev el	INFO	CRITICAL, ERROR, WARNING, INFO, DEBUG	The logging level for UFM events log messages	
rest_log_level	INFO CRITICAL, ERROR WARNING, INFO,		Logging level for REST API related log messages	

Field	Default Value	Value Options	Description	
		DEBUG		
authenticatio n_service_log _level	INFO CRITICAL, ERROR, WARNING, INFO, DEBUG Messages		logging level for UFM authentication log messages	
log_dir	/opt/uf m/files/ log	N/A	It is possible to change the default path to the UFM log directory. The configured log_dir must have read, write and execute permission for ufmapp user (ufmapp group). In case of HA, UFM should be located in the directory which is replicated between the UI master and standby servers. A change of the default UFM log directory m affect UFM dump creation and inclusion of UFM logs in dump.	
max_history_ lines	100000	N/A	The maximum number of lines in log files to be shown in UI output for UFM logging.	
[Logging] # Optional logging levels: CRITICAL, ERROR, WARNING, INFO, DEBUG. level = WARNING smclient_level = WARNING event_log_level = INFO rest_log_level = INFO authentication_service_log_level = INFO				

# The configured log\_dir must have read, write and execute permission for ufmapp user (ufmapp group).

log\_dir = /opt/ufm/files/log

max\_history\_lines = 100000

# **Configuring UFM Over Static IPv4 Address**

Follow this procedure to to run UFM on a static IP configuration instead of DHCP:

1. Modify the defined management Ethernet interface network script to be static. Run:

# vi /etc/sysconfig/network-scripts/ifcfg-enp1s0

Update the required interface with the static IP configuration (IP address, netmask, broadcast, and gateway):

NAME="enp1s0"DEVICE="enp1s0" ONBOOT="yes" BOOTPROTO="static" IPADDR="10.209.37.153" NETMASK="255.255.252.0" BROADCAST="10.209.39.255" GATEWAY="10.209.36.1" TYPE=Ethernet DEFROUTE="yes"

2. Add host entries to the /etc/hosts file. Run:

# vi /etc/hosts

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4

::1 localhost localhost.localdomain localhost6 localhost6.localdomain6

10.209.37.153 <hostname>

#### 3. Check hostname. Run:

# vi /etc/hostname <hostname>

#### 4. Set up DNS resolution at /etc/resolv.conf. Run:

# vi /etc/resolv.conf search mtr.labs.mlnx nameserver 8.8.8.8

5. Restart network service. Run:

service network restart

6. Check Configuration. Run:

# hostname <hostname> # hostname -i 10.209.37.153

# **Configuring Syslog**

This configuration enables the UFM to send log messages to syslog, including remote syslog. The configuration described below is located in the [Logging] section of the gv.cfg file.

Field	Default Value	Value Options	Description
syslog	false	True or False	Enables/disables UFM syslog option
syslog_ addr	/dev/log # for remote rsyslog_hostn ame:514	N/A	UFM syslog configuration (syslog_addr) For working with local syslog, set value to: /dev/log For working with external machine, set value to: host:port Important note: the default remote syslog server port is 514 As the UFM log messages could be sent to remote server, change the rsyslog configuration on the remote server

Field	Default Value	Value Options	Description
			The /etc/rsyslog.conf file should be edited and two sections should be uncommented as shown below: # Provides UDP syslog reception \$ModLoad imudp \$UDPServerRun 514 # Provides TCP syslog reception \$ModLoad imtcp \$InputTCPServerRun 514 Restart the remote syslog service, run: service rsyslog restart
ufm_sy slog	false	True or False	Sets syslog option for main UFM process logging messages - False - Not to send. True: Send
smclie nt_sysl og	false	True or False	Sets syslog option for OpenSM logging messages - False - Not to send. True: Send
event_s yslog	false	True or False	Sets syslog option for events logging messages - False - Not to send. True: Send
rest_sy slog	false	True or False	Sets syslog option for UFM REST API logging messages - False - Not to send. True: Send
authen tication _syslog	false	True or False	Set syslog option for UFM authentication logging messages. False - Not to send. True: Send
syslog_ level	WARNING	CRITICAL, ERROR, WARNING, INFO, DEBUG	Sets global syslog messages logging level. The syslog level is common for all the UFM components. The syslog level that is sent to syslog is the highest among the syslog level and component log level defined in the above section.
syslog_ facility	LOG_USER	LOG_KERN, LOG_USER, LOG_MAIL,	Includes the remote syslog package header value for log message facility.

Field	Default Value	Value Options	Description
		LOG_DAEMON,	
		LOG_AUTH,	
		LOG_SYSLOG, LOG_LPR,	
		LOG_NEWS, LOG_UUCP	
		,LOG_CRON,	
		LOG_AUTHPRIV,	
		LOG_FTP,	
		LOG_NTP,LOG_SECURITY	
		, LOG_CONSOLE,	
		LOG_SOLCRON	

syslog = false #syslog configuration (syslog\_addr) # For working with local syslog, set value to: /dev/log # For working with external machine, set value to: host:port syslog\_addr = /dev/log # The configured log\_dir must have read, write and execute permission for ufmapp user (ufmapp group). log\_dir = /opt/ufm/files/log # Main ufm log. ufm\_syslog = false smclient\_syslog = false event\_syslog = false rest\_syslog = false authentication\_syslog = false syslog\_level = WARNING # Syslog facility. By default - LOG\_USER # possible facility codes: LOG\_KERN,LOG\_USER,LOG\_MAIL,LOG\_DAEMON,LOG\_AUTH,LOG\_SYSLOG, # LOG\_LPR,LOG\_NEWS,LOG\_UUCP,LOG\_CRON,LOG\_AUTHPRIV,LOG\_FTP, LOG\_NTP,LOG\_SECURITY,LOG\_CONSOLE,LOG\_SOLCRON # for reference https://en.wikipedia.org/wiki/Syslog syslog\_facility = LOG\_USER

### **Configuring Unhealthy Ports**

In gv.cfg file there is a section named **UnhealthyPorts** and parameters in this section are used for unhealthy ports managing in UFM.

Unhealthy port state could be defined by used using UI or REST API request or reported by OpenSM or ibutilities.

UFM has an ability to check periodically fabric ports healthiness and to report unhealthy ports out or to perform automatically predefined isolation action for unhealthy ports.

In addition, using exclude\_unhealthy\_ports key in **UnhealthyPorts** section unhealthy ports could be excluded from ibdiagnet report.

By default, the value for this parameter is set as *false*. It means that unhealthy ports will appear in ibdiagnet reports, but if need to exclude unhealthy port from ibdiagnet reports

this parameter should be set to true and UFM server should be restarted so this action will take effect.

UFM starting flow will configure indiagnet configuration file with appropriate parameters and unhealthy ports will not appear in UFM health and Fabric health reports.

[UnhealthyPorts] enable\_ibdiagnet = true log\_level = INFO syslog = false# scheduling\_mode possible values: fixed\_time/interval. # If fixed\_time - ibdiagnet will run every 24 hours on the specified time - <fixed\_time>. # If interval - ibdiagnet will run first time after <start\_delay> minutes from UFM startup and every <interval> hours (default scheduling mode). scheduling\_mode = interval # First ibdiagnet start delay interval (minutes) start\_delay = 5 # ibdiagnet run interval (hours) interval = 3# ibdiagnet run at a fixed time (example: 23:17:35) fixed time = 23:00:00# By enabling this flag all the discovered high ber ports will be marked as unhealthy automatically by UFM high\_ber\_ports\_auto\_isolation = false # Auto isolation mode - which type of ports should be isolated. # Options: switch-switch, switch-host, all (default: switch-switch). auto\_isolation\_mode = switch-switch # Trigger Partial Switch ASIC Failure whenever number of unhealthy ports exceed the defined percent of the total number of the switch ports. switch\_asic\_fault\_threshold = 20

# exclude unhealthy ports from ibdiagnet reports
exclude\_unhealthy\_ports=false

### **Configuration Examples in gv.cfg**

The following show examples of configuration settings in the gv.cfg file:

• Polling interval for Fabric Dashboard information

ui\_polling\_interval = 30

• [**Optional**] UFM Server local IP address resolution (by default, the UFM resolves the address by gethostip). UFM Web UI should have access to this address.

ws\_address = <specific IP address>

• HTTP/HTTPS Port Configuration

# WebServices Protocol (http/https) and Port ws\_port = 8088 ws\_protocol = http

• Connection (port and protocol) between the UFM server and the APACHE server

```
ws_protocol = <http or https>
ws_port = <port number>
```

For more information, see Launching a UFM Web UI Session.

• SNMP get-community string for switches (fabric wide or per switch)

# default snmp access point for all devices [SNMP] port = 161 gcommunity = public

• Enhanced Event Management (Alarmed Devices Group)

[Server] auto\_remove\_from\_alerted = yes

• Log verbosity

[Logging] # optional logging levels #CRITICAL, ERROR, WARNING, INFO, DEBUG level = INFO

For more information, see "<u>UFM Logs</u>".

- Settings for saving port counters to a CSV file
  - [CSV] write\_interval = 60 ext\_ports\_only = no

For more information, see "Saving the Port Counters to a CSV File".

• Max number of CSV files (UFM Advanced)

[CSV] max\_files = 1 For more information, see "<u>Saving Periodic Snapshots of the Fabric (Advanced License Only)</u>".

(j) Note
The access credentials that are defined in the following sections of the conf/gv.cfg file are used only for initialization:
<ul> <li>SSH_Server</li> </ul>
<ul> <li>SSH_Switch</li> </ul>
○ TELNET
o IPMI
• SNMP
<ul> <li>MLNX_OS</li> </ul>
To modify these access credentials, use the UFM Web UI. For more information, see " <u>Device Access</u> ".

- Configuring the UFM communication protocol with MLNX-OS switches. The available protocols are:
  - http
  - https (default protocol for secure communication)

For configuring the UFM communication protocol after fresh installation and prior to the first run, set the MLNX-OS protocol as shown below.

#### Example:

[MLNX\_OS] protocol = https Once UFM is started, all UFM communication with MLNX-OS switches will take place via the configured protocol.

# For changing the UFM communication protocol while UFM is running, perform the following:

- 1. Set the desired protocol of MLNX-OS in the conf/gv.cfg file (as shown in the example above).
- 2. Restart UFM.
- 3. Update the MLNX-OS global access credentials configuration with the relevant protocol port. Refer to "<u>Device Access</u>" for help.

For the http protocol - default port is 80.

For the https protocol - default port is 443.

4. Update the MLNX-OS access credentials with the relevant port in all managed switches that have a valid IP address.

### **Managing Dynamic Telemetry**

The management of dynamic telemetry instances involves the facilitation of user requests for the creation of multiple telemetry instances. As part of this process, the UFM enables users to establish new UFM Telemetry instances according to their preferred counters and configurations. These instances are not initiated by the UFM but rather are monitored for their operational status through the use of the UFM Telemetry bring-up tool

For more information on the supported REST APIs, please refer to <u>UFM Dynamic</u> <u>Telemetry Instances REST API</u>.

The configuration parameters can be found in the gv.cfg configuration file under the DynamicTelemetry section.

Name	Description	Defa ult value
max_instances	Maximum number of simultaneous running UFM Telemetries.	5
new_instance_ delay	Delay time between the start of two UFM Telemetry instances, in minutes.	5
update_discov ery_delay	The time to wait before updating the discovery file of each telemetry instance if the fabric has changed, in minutes.	10
endpoint_time out	Telemetry endpoint timeout, in seconds.	5
bringup_timeo ut	Telemetry bringup tool timeout, in seconds.	60
initial_exposed _port	Initial port for the available range of ports (range(initial_exposed_port, initial_exposed_port + max_instances)).	9003
instances_sess ions_compatib ility_interval	Every instances_sessions_compatibility_interval minutes the UFM verifies compliance between instances and sessions to avoid zombie sessions. if 0 is configured this process won't be activate	10

# **SM Trap Handler Configuration**

The SMTrap handler is the SOAP server that handles traps coming from OpenSM.

There are two configuration values related to this service:

- osm\_traps\_debounce\_interval defines the period the service holds incoming traps
- osm\_traps\_throttle\_val Once osm\_traps\_debounce\_interval elapses, the service transfers osm\_traps\_throttle\_val to the Model Main



# **Setting CPU Affinity on UFM**

This feature allows setting the CPU affinity for the major processes of the UFM (such as ModelMain, SM, SHARP, Telemetry).

In order to increase the UFM's efficiency, the number of context-switches is reduced. When each major CPU is isolated, users can decrease the number of context-switches, and the performance is optimized.

The CPU affinity of these major processes is configured in the following two levels:

- Level 1- The major processes initiation.
- Level 2- Preceding initiation of the model's main subprocesses which automatically uses the configuration used in level 1 and designates a CPU for each of the sub-processes.

According to user configuration, each process is assigned with affinity.

By default, this feature is disabled. In order to activate the feature, configure Is\_cpu\_affinity\_enabled with true, check how many CPUs you have on the machine, and set the desired affinity for each process.

For example:

[CPUAffinity] Is\_cpu\_affinity\_enabled=true Model\_main\_cpu\_affinity=1-4 Sm\_cpu\_affinity=5-13 SHARP\_cpu\_affinity=14-22 Telemetry\_cpu\_affinity=22-23 The format should be a comma-separated list of CPUs. For example: 0,3,7-11.

The ModelMain should have four cores, and up to five cores. The SM should have as many cores as you can assign. You should isolate between the ModelMain cores and the SM cores.

SHARP can be assigned with the same affinity as the SM. The telemetry should be assigned with three to four CPUs.

### **Quality of Service (QoS) Support**

Infiniband Quality of Service (QoS) is disabled by default in the UFM SM configuration file.

To enable it and benefit from its capabilities, set the qos flag to TRUE in the /opt/ufm/files/conf/opensm/opensm.conf file.

Example:

# Enable QoS setup qos FALSE



The QoS parameters settings should be carefully reviewed before enablement of the qos flag. Especially, sl2vl and VL arbitration mappings should be correctly defined.

For information on Enhanced QoS, see Appendix – SM Activity Report.

#### **UFM Failover to Another Port**

When the UFM Server is connected by two or more InfiniBand ports to the fabric, you can configure UFM Subnet Manager failover to one of the other ports. When failure is detected on an InfiniBand port or link, failover occurs without stopping the UFM Server or other related UFM services, such as mysql, http, DRDB, and so on. This failover process prevents failure in a standalone setup, and preempts failover in a High Availability setup, thereby saving downtime and recovery.



Network Configuration for Failover to IB Port

To enable UFM failover to another port:

- Configure bonding between the InfiniBand interfaces to be used for SM failover. In an HA setup, the UFM active server and the UFM standby server can be connected differently; but the bond name must be the same on both servers.
- Set the value of fabric\_interface to the bond name. using the /opt/ufm/scripts/change\_fabric\_config.sh command as described in <u>Configuring</u> <u>General Settings in gv.cfg</u>. If ufma\_interface is configured for IPoIB, set it to the bond name as well. These changes will take effect only after a UFM restart. For example, if bond0 is configured on the ib0 and ib1 interfaces, in gv.cfg, set the parameter fabric\_interface to bond0.
- If IPoIB is used for UFM Agent, add bond to the ufma\_interfaces list as well.

When failure is detected on an InfiniBand port or link, UFM initiates the give-up operation that is defined in the Health configuration file for OpenSM failure. By default:

- UFM discovers the other ports in the specified bond and fails over to the first interface that is up (SM failover)
- If no interface is up:
  - In an HA setup, UFM initiates UFM failover
  - In a standalone setup, UFM does nothing

If the failed link becomes active again, UFM will select this link for the SM only after SM restart.

#### **Configuring Managed Switches Info Persistency**

UFM uses a periodic system information-pulling mechanism to query managed switches inventory data. The inventory information is saved in local JSON files for persistency and tracking of managed switches' status.

Upon UFM start up, UFM loads the saved JSON files to present them to the end user via REST API or UFM WebUI.

After UFM startup is completed, UFM pulls all managed switches data and updates the JSON file and the UFM model periodically (the interval is configurable). In addition, the JSON files are part of UFM system dump.

The following parameters allow configuration of the feature via gv.cfg fie:

[SrvMgmt] # how often UFM should send json requests for sysinfo to switches (in seconds) systems\_poll = 180 # To create UFM model in large setups might take a lot of time. # This is an initial delay (in minutes) before starting to pull sysinfo from switches. systems\_poll\_init\_timeout = 5 # to avoid sysinfo dump overloading and multiple writing to host # switches sysinfo will be dumped to disc in json format every set in this variable # sysinfo request. If set to 0 - will not be dumped, if set to 1 - will be dumped every sysinfo request # this case (as example defined below) dump will be created every fifth sysinfo request, so if system\_poll is 180 sec (3 minutes) sysinfo dump to the file will e performed every 15 minutes. sysinfo\_dump\_interval = 5

### **Configuring Partial Switch ASIC Failure Events**

UFM can identify switch ASIC failure by detecting pre-defined portion of the switch ports, reported as unhealthy. By default, this portion threshold is set to 20% of the total switch ports. Thus, the UFM will trigger the partial switch ASIC event in case the number of unhealthy switch ports exceeds 20% of the total switch ports.

You can configure UFM t o control Partial Switch ASIC Failure events. To configure, you may use the gv.cfg file by updating the value of switch\_asic\_fault\_threshold parameter under the UnhealthyPorts section. For an example, in case the switch has 32 ports, once 7 ports are detected as unhealthy ports, the UFM will trigger the partial switch ASIC event. Example:

😮 Warning	2023-01-25 10:41:22	Unhealthy IB Port	default(2) / Switch: sw-ufm-qr	IBPort	Peer Port is considered by SM as unhealthy due to MANUAL.
🕜 Warning	2023-01-25 10:41:02	Unhealthy IB Port	default(2) / Switch: sw-ufm-qr	IBPort	Peer Port "r-ufm51 HCA-1" is considered by SM as unhealthy due to MANUAL.
Critical	2023-01-25 10:41:02	Partial Switch ASIC Failure	default / Switch: sw-ufm-qm0	Switch	Number of switch unhealthy ports has been exceeded the defined threshold which is [4] perce
🥑 Info	2023-01-25 10:40:43	MCast Group Deleted	default(2)	Site	Mcast group is deleted: ff12601bffff0000, 1ff18fe80

# **Enabling Network Fast Recovery**

#### j Note

To enable the Network Fast Recovery feature, ensure that all switches in the fabric use the following MLNX-OS/firmware versions:

- MLNX-OS version 3.10.6004 and up
- Quantum firmware versions:
  - Quantum FW v27.2010.6102 and up
  - Quantum2 FW v31.2010.6102 and up

Fast recovery is a switch-firmware based facility for isolation and mitigation of link-related issues. This system operates in a distributed manner, where each switch is programmed with a simple set of rule-based triggers and corresponding action protocols. These rules permit the switch to promptly react to substrandard links within its locality, responding at a very short reaction time - as little as approximately 100 milliseconds. The policy is provided and managed via the UFM & SM channel. Moreover, every autonomous action taken by a switch in the network is reported to the UFM.

The immediate reactions taken by the switch enable SHIELD and pFRN. These mechanisms collaborate to rectify routing within the proximity of the problematic link before it can disrupt transactions at the transport layer. Importantly, this process occurs rapidly, effectively limiting the spreading of congestion to a smaller segment of the network.

To use the Network Fast Recovery feature, you need to enable the designated trigger in the gv.cfg file. By doing this, you can specify which triggers the UFM will support.

As stated in the gv.cfg file, the feature is disabled by default and the below are the supported fields and options:

[NetworkFastRecovery]

is\_fast\_recovery\_enabled = false

# This will be supported by the Network Fast Recovery.

network\_fast\_recovery\_conditions = SWITCH\_DECISION\_CREDIT\_WATCHDOG,SWITCH\_DECISION\_RAW\_BER,SWITCH\_DECISION\_

Parameter	Description
SWITCH_DECISION_CREDIT_WATCHDOG	TBD
SWITCH_DECISION_RAW_BER	
SWITCH_DECISION_EFFECTIVE_BER	
SWITCH_DECISION_SYMBOL_BER	

The "Unhealthy Ports" page provides visibility of these ports. If desired, the user can mark a port as healthy, triggering a restart of that specific port on the switch.

The trigger that initiated the isolation of ports can be viewed under the "Condition" column, as seen below.

Unhealthy Ports	Health Policy								
						All Connectivity 🔷	Mark All Ports as I	lealthy 🛛 🕄 Displaye	d Columns 👻 🛛 CSV 🗸
		Unhealthy Source Port		Peer					
Severity	Node	Port	GUID	Name	Port	GUID	LID	Condition	Status Time
	V Filter	∇ Filter ∇	Filter 🗸 🗸		<b>▽</b>   [Filter <b>▽</b>	Filter V	Filter	▼ Filter ▼	Filter 7
🕜 Warning	smg-ib-sw012	smg-ib-sw012:16	0x043f720300f695c6	smg-ib-sw056	smg-ib-sw056:117	0x900a84030040c8	17	SWITCH_DECISIO	Mon Mar 06 12:33:
								Viewing 1-1 of 1	<ul> <li>✓</li> <li>✓</li></ul>

#### **Disabling Rest Roles Access Control**

By default, the Rest Roles Access Control feature is enabled. It can be disabled by setting the roles\_access\_control\_enabled flag to false:

[RolesAccessControl] roles\_access\_control\_enabled = true

### **Enabling/Disabling Authentication**

#### **Kerberos Authentication**

By default, <u>Kerberos Authentication</u> is disabled. To enable it, set the kerberos\_auth\_enabled flag to true. Additionally, provide the required configurations such as kerberos\_cred\_key\_path, kerberos\_use\_local\_name and kerberos\_auto\_sign\_up.

[KerberosAuth] # This section responsible to manage kerberos authentication # Set to true to enable the kerberos auth feature, and set to false to disable it. Default is false. kerberos\_auth\_enabled = false # The path of the keytab file containing credentials for GSSAPI authentication. kerberos\_cred\_key\_path = /etc/kadm5.keytab # Set to true to configure the Apache server to map authenticated principal names (which represent different clients) to local usernames, # and set to false to use the principle names as usernames. Default is true (this value will be reflected in the 'GssapiLocalName' directive in Apache). kerberos\_use\_local\_name = true # Set to true to enable auto sign up of users who do not exist in UFM DB. Default is true. kerberos\_auto\_sign\_up = true # The default role assigned to create users if they do not exist when 'kerberos\_auto\_sign\_up' is set to true. kerberos\_default\_role = System\_Admin

**kerberos\_auth\_enabled**: By default, Kerberos authentication remains disabled. To activate it, the user must set this flag to 'true' and then restart UFM.

**kerberos\_cred\_key\_path**: This specifies the path to the keytab file containing credentials for GSSAPI authentication.

**kerberos\_use\_local\_name**: Set to true to configure the Apache server to map authenticated principal names (which represent different clients) to local usernames, and set to false to use the principal names as usernames. Default is true (this value will be reflected in the 'GssapiLocalName' directive in Apache).

**kerberos\_auto\_signup**: For successful authentication via Kerberos, the user must already exist within the UFM database, otherwise, the authentication will be refused by UFM. If this property is set to 'true,' UFM will create the non-existing users in the UFM DB.

**kerberos\_default\_role**: The default role is assigned to create users if they do not exist when 'kerberos\_auto\_sign\_up' is set to true.

Finally, restart the UFM to use Kerberos authentication.

#### **UFM Authentication Server**

By default, <u>UFM Authentication Server</u> is enabled. To disable it, you need to set the "auth\_service\_enabled" parameter to 'false' and then restart the UFM service to initiate the authentication server. Additionally, you can use enable/disable flags for Basic, Session, and Token authentication:

[AuthService] auth\_service\_enabled = true auth\_service\_interface = 127.0.0.1 auth\_service\_port = 8087 # the serving port for the authentication server basic\_auth\_enabled = true session\_auth\_enabled = true token\_auth\_enabled = true

#### **Azure AD Authentication**

By default, <u>Azure AD Authentication</u> is disabled. To enable it, set the azure\_auth\_enabled flag to 'true'. Additionally, provide the required configurations from the Azure AD Application such as TENANT\_ID, CLIENT\_ID and CLIENT\_SECRET which can be found under the "**Overview**" section of the registered application in the Azure portal. Finally, the <u>UFM</u> <u>Authentication Server</u> should be enabled to use the Azure AD Authentication.

[AzureAuth] azure\_auth\_enabled = false # TENANT ID of app registration TENANT\_ID = # Application (client) ID of app registration CLIENT\_ID = # Application's generated client secret CLIENT\_SECRET =

# **Managing UFM Configuration Files Based on Fabric Size**

This function allows users to automate the process of updating the UFM configuration files by parsing a primary configuration file called large\_scale\_subnet.cfg file and applying the values to multiple target files or resetting to default values using the small\_scale\_subnet.cfg.

The below are instructions on how to use a Python script to parse a configuration file (large\_scale\_subnet.cfg) and update the values of specific parameters in multiple target UFM

configuration files (gv.cfg, reports.cfg, opensm.cfg, and sharp\_am.cfg). The script can operate in two modes:

- Large Scale Subnet Mode: This mode directly updates the UFM configuration files based on the parsed configuration from the <a href="https://www.scale.subnet.cfg">large\_scale\_subnet.cfg</a> file.
- **Small Scale Subnet (Default) Mode**: Sets the UFM configuration files to their default values by parsing the small\_scale\_subnet.cfg file.

#### **Configuration File and Parameters**

The primary configuration file contains all the parameters, and their values must be updated over the multiple UFM configuration files.

#### **Primary Configuration Files**

- /opt/ufm/files/conf/ large\_scale\_subnet.cfg
- /opt/ufm/files/conf/ small\_scale\_subnet.cfg

#### **Target UFM Configuration Files**

- /opt/ufm/files/conf/gv.cfg
- /opt/ufm/files/conf/reports.cfg
- /opt/ufm/files/conf/opensm/opensm.conf
- /opt/ufm/files/conf/sharp/sharp\_am.cfg

Example structure of large\_scale\_subnet.cfg and small\_scale\_subnet.cfg:

[GV] [GV.Server] # disabled (default) | enabled (configure opensm with multiple GUIDs) | ha\_enabled (configure multiport SM with high availability). multi\_port\_sm = ha\_enabled # report\_events that will determine which trap to send to ufm all/security/none report\_events = security [GV.FabricAnalysis] # initial\_delay (in minutes) - the initial delay for running fabric analysis for the first time after UFM was started initial\_delay = 10

[GV.logrotate]
#max\_files specifies the number of times to rotate a file before it is deleted.
#A count of 0 (zero) means no copies are retained. A count of 10 means fifteen copies are retained
(default is 10)
max\_files = 10

[REPORTS] [REPORTS.FabricHealth] # Fabric health report timeout timeout = 1800

[REPORTS.TopologyCompare] # Topology compare report timeout timeout = 1800

[REPORTS.FabricAnalysis] # Fabric analysis report timeout timeout = 1800

```
[OPENSM]
#Amount of physical port to handle in one shot
virt_max_ports_in_process = 512
max_op_vls = 2
qos = TRUE
# Single MAD Sl2vl for all ports
use_optimized_slvl = TRUE
# Timeout for long MAD config time. might need to change 1000
long_transaction_timeout = 500
routing_engine = ar_updn
use_ucast_cache = TRUE
root_guid_file = /opt/ufm/files/conf/opensm/root_guid.conf
pgrp_policy_file = /opt/ufm/files/conf/opensm/pgrp_policy.conf
```

```
[SHARP]
ib_qpc_sl = 1
fabric_update_interval = 10
lst_file_timeout = 10
```

lst\_file\_retries = 30
max\_tree\_radix = 80
generate\_dump\_files = TRUE
dynamic\_tree\_allocation = TRUE
dynamic\_tree\_algorithm = 1
smx\_keepalive\_interval = 10

#### Script Usage Example in the CLI:

- Large Scale Subnet Mode: /opt/ufm/scripts/set\_ufm\_scale\_profile.sh --mode large\_scale\_subnet --force\_update
- Small Scale Subnet (Default) Mode: /opt/ufm/scripts/ set\_ufm\_scale\_profile.sh --mode small\_scale\_subnet

The force\_update script parameter adds any parameters found in large\_scale\_subnet.cfg and small\_scale\_subnet.cfg that are not present in the UFM configuration files. For example, if a user adds a new parameter called test\_param = 500 to the large\_scale\_subnet.cfg file under the [Server] section and this parameter does not exist in the gv.cfg file, running the script with the --force\_update option will add test\_param = 500 to the [Server] section of the gv.cfg file.

#### **Expected Output**

#### 1. Large Scale Subnet Mode:

- The script reads large\_scale\_subnet.cfg.
- It updates the parameters in the target UFM configuration files based on the parsed data.
- It logs messages for any skipped parameters.

#### 1. Small Scale Subnet - Default Mode:

- The script reads small\_scale\_subnet.cfg.
- It updates the parameters in the target UFM configuration files based on the parsed data.
- It logs messages for any skipped parameters or adds the parameter to the configuration file if the force\_update was True.

#### (i) Note

**Note:** In case of the script running failure, the script will reset the UFM configuration files to their default values.

#### Setting up Telemetry in UFM

Setting up telemetry deploys UFM Telemetry as bare metal on the same machine. Historical data is sent to SQLite database on the server and live data becomes available via UFM UI or REST API.

### **Enabling UFM Telemetry**

The UFM Telemetry feature is enabled by default and the provider is the UFM Telemetry. The user may change the provider via flag in conf/gv.cfg

The user may also disable the History Telemetry feature in the same section.

[Telemetry] history\_enabled=True

# **Changing UFM Telemetry Default Configuration**

There is an option to configure parameters on a telemetry configuration file which takes effect after restarting the UFM or failover in HA mode.The <code>launch\_ibdiagnet\_config.ini</code> default file is located under /opt/ufm/conf/telemetry\_defaults and is copied to the telemetry configuration location ( (/opt/ufm/conf/telemetry) upon startup UFM.

All values taken from the default file take effect at the deployed configuration file except for the following:

Note that normally the user does not have to do anything and they get two preconfigured instances – one for low frequency and one for higher-frequency sampling of the network.

Value	Description					
hca	-					
scope_file	-					
plugin_env_PROM ETHEUS_ENDPOIN T	The port on which HTTP endpoint is configured					
plugin_env_PROM ETHEUS_INDEXES	Configures how data is indexed and stored in memory					
config_watch_enab led=true	Configures network watcher to inform ibdiagnet that network topology has changed (as ibdiagnet lacks the ability to re-discover network changes)					
plugin_env_PROM ETHEUS_CSET_DIR	Specifies where the counterset files, which define the data to be retrieved and the corresponding counter names.					
num_iterations	The number of iterations to run before 'restarting', i.e. rediscovering fabric.					
plugin_env_CLX_RE START_FILE	A file that is 'touched' to indicate that an ibdiagnet restart is necessary					

The following attributes are configurable via the gv.cfg:

- sample\_rate (gv.cfg dashboard\_interval) only if manual\_config is set to false
- prometheus\_port

### **Supporting Generic Counters Parsing and Display**

As of UFM v6.11.0, UFM can support any numeric counters from the HTTP endpoint. The list of supported counters are fetched upon starting the UFM from all the endpoints that are configured.

Some of the implemented changes are as follows:

1. Counter naming – all counters naming convention is extracted from the HTTP endpoint. The default cset file is configured as follows:

"Infiniband\_LinkIntegrityErrors=^LocalLinkIntegrityErrorsExtended\$" to get this name to the UFM.

Counters received as floats should contain an "\_f" suffix such as: Infiniband\_CBW\_f=^infiniband\_CBW\$

- 2. Attribute units To see units of a specific counter on the UI graphs, configure the cset file to have the counter returned as "counter\_name\_u\_unit".
- 3. Telemetry History:

The SQLite history table (/opt/ufm/files/sqlite/ufm\_telemetry.db – telemetry\_calculated), contains the new naming convention of the telemetry counters.

In the case of an upgrade, all previous columns that were configured are renamed following the new naming convention, and then, the data is saved.if a new counter that is not in the table needs to be supported, the table is altered upon UFM start.

- 4. New counter/cset to fetch if there is a new cset/counter that needs to be supported AFTER the UFM already started, preform system restart.
- 5. Created New API/UfmRestV2/telemetry/counters for the UI visualization. This API returns a dictionary containing the counters that the UFM supports, based on the fetched URLs and their units (if known).

### Supporting Multiple Telemetry Instances Fetch

This functionality allows users to establish distinct Telemetry endpoints that are defined to their preferences.

Users have the flexibility to set the following aspects:

- Specify a list of counters they wish to pull. This can be achieved by selecting from an existing, predefined counters set (cset file) or by defining a new one.
- Set the interval at which the data should be pulled.

Upon initiating the Telemetry endpoint, users can access the designated URL to fetch the desired counter data.

To enable this feature, under the [Telemetry] section in gv.cfg,the flag named "additional\_cset\_url" holds the list of additional URLs to be fetched.

the URLs should be separated by " " (with a space) and should follow the following format: http://:/csv/. For example <u>http://10.10.10.10:9001/csv/minimal</u> <u>http://10.10.10:9002/csv/test</u>.

(i) Note

Only csv extensions are supported.

Each UFM Telemetry instance run by UFM can support multiple cset (counters set) in parallel.If the user would like to have a second cset file fetched by UFM and exposed by the same UFM Telemetry instance, the new cset file should be placed under /opt/ufm/files/conf/telemetry/prometheus\_configs/cset/ and configured in gv.cfg to fetch its data as described above.

# Low-Frequency (Secondary) Telemetry

As a default configuration, a second UFM Telemetry instance runs, granting access to an extended set of counters that are not available in the default telemetry session. The default telemetry session is used for the UFM Web UI dashboard and user-defined telemetry views. These additional counters can be accessed via the following API endpoint: http://<UFM\_IP>:9002/csv/xcset/low\_freq\_debug.It is important to note that these exposed counters are not accessible through UFM's REST APIs.All the configurations for the second telemetry can be found under /opt/ufm/files/conf/secondary\_telemetry/, where the defaults are located under /opt/ufm/files/conf/secondary\_telemetry\_defaults/. The second telemetry instance also allows telemetry data to be exposed on disabled ports, although this feature can be disabled if desired.

The relevant flags in the gv.cfg file are as follows:

- secondary\_telemetry = true (To enable or disable the entire feature)
- secondary\_endpoint\_port = 9002 (The endpoint's exposed port)
- secondary\_disabled\_ports = true (If set to true, secondary telemetry will expose data on disabled ports)
- secondary\_slvl\_support = false (if set to true, low-frequency (secondary) Telemetry will collect counters per slvl, the corresponding supported xcset can be found under /opt/ufm/files/conf/secondary\_telemetry/prometheus\_configs/cset/low\_freq\_debug\_

The counters that are supported by default, collected, and exposed can be located in the directory /opt/ufm/files/conf/secondary\_telemetry/prometheus\_configs/cset/low\_freq\_debug\_per\_slvl.xcset.

For the list of low-frequency (secondary) telemetry fields and available counters, please refer to <u>Low-Frequency (Secondary) Telemetry Fields</u>.

#### Low-Frequency (Secondary) Telemetry - Exposing IPv6 Counters

To allow the low-frequency (secondary) telemetry instance to expose counters on its IPv6 interfaces, perform the following:

1. Change the following flag in the gv.cfg:

secondary\_ip\_bind\_addr =0:0:0:0:0:0:0:0

2. Restart UFM telemetry or restart UFM.

#### **Stopping Telemetry Endpoint Using CLI Command**

To stop low-frequency (secondary) telemetry endpoint only using the CLI you may run the following command:

/etc/init.d/ufmd ufm\_telemetry\_secondary\_stop

#### **Exposing Switch Aggregation Nodes Telemetry**

To expose switches SHARP aggregation nodes telemetry, follow the below steps:

• Configure the low-frequency (secondary) telemetry instance. Run:

vi /opt/ufm/files/conf/secondary\_telemetry\_defaults/launch\_ibdiagnet\_config.ini

#### • Set the following:

- arg\_16=--sharp --sharp\_opt dsc
- plugin\_env\_CLX\_EXPORT\_API\_SKIP\_SHARP\_PM\_COUNTERS=0
- Add the wanted attributes to the default xcset or to a new one:
  - New xcset
    - vi /opt/ufm/files/conf/secondary\_telemetry/prometheus\_configs/cset/<name for your choise>.xcset
    - After restarting, query curl http://<UFM\_IP>:9002/csv/xcset/<chosen\_name>
  - Existing xcset
    - vi /opt/ufm/files/conf/secondary\_telemetry/prometheus\_configs/cset/low\_freq\_de

0

• Add the following attributes:

- packet\_sent
- ack\_packet\_sent
- retry\_packet\_sent
- rnr\_event
- timeout\_event
- oos\_nack\_rcv
- rnr\_nack\_rcv
- packet\_discard\_transport
- packet\_discard\_sharp
- aeth\_syndrome\_ack\_packet
- hba\_sharp\_lookup
- hba\_received\_pkts
- hba\_received\_bytes
- hba\_sent\_ack\_packets
- rcds\_sent\_packets
- hba\_sent\_ack\_bytes
- rcds\_send\_bytes
- hba\_multi\_packet\_message\_dropped\_pkts
- hba\_multi\_packet\_message\_dropped\_bytes
- Restart telemetry:

/etc/init.d/ufmd ufm\_telemetry\_stop

#### Exposing Performance Histogram Counters for Egress Queue Depth Indications (Secondary) Telemetry

To enable the secondary telemetry instance to expose performance histogram counters for all VLs, perform the following:

1. Change the following flag in the gv.cfg file:

queue\_depth\_indications\_all\_vls = true

If this flag remains set to false, the secondary telemetry instance will only collect counters for VLs 0 and 1.

2. Restart UFM telemetry or restart UFM.

After the secondary telemetry instance restarts, you can find the collected counters at:

/opt/ufm/conf/secondary\_telemetry/prometheus\_configs/cset/low\_freq\_debug\_per\_slvl.xcset

© Copyright 2024, NVIDIA. PDF Generated on 08/14/2024