



NVIDIA UFM Enterprise User Manual

v6.11.2

Table of Contents

1	About This Document.....	11
1.1	Software Download	11
1.2	Document Revision History	11
2	Release Notes.....	12
2.1	Key Features	12
2.2	Changes and New Features	12
2.2.1	Unsupported Functionalities/Features	14
2.3	Installation Notes	14
2.3.1	Supported Devices	14
2.3.2	System Requirements	15
2.3.3	Software Update from Prior Versions	17
2.4	Bug Fixes in This Release	18
2.5	Known Issues in This Release	19
2.6	Changes and New Features History	19
2.7	Bug Fixes History.....	26
2.8	Known Issues History	30
3	Overview	36
3.1	Scale-Out Your Fabric with Unified Fabric Manager	36
3.2	UFM Benefits.....	36
3.3	Main Functionality Modules	37
3.3.1	Fabric Dashboard	37
3.3.2	Fabric Segmentation (PKey Management).....	37
3.3.3	Fabric Discovery and Physical View	37
3.3.4	Central Device Management	37
3.3.5	Monitoring.....	38
3.3.6	Configuration	38
3.3.7	Fabric Health	38
3.3.8	Logging	38
3.3.9	High Availability.....	38
3.4	Fabric Topology with UFM	38
3.5	UFM Communication Requirements	40
3.5.1	UFM Server Communication with Clients.....	40

3.5.2	UFM Server Communication with InfiniBand Switches	41
3.5.3	UFM Server Communication with InfiniBand Hosts	42
3.5.4	UFM Server High Availability (HA) Active–Standby Communication	43
3.6	UFM Software Architecture.....	44
3.6.1	Graphical User Interface.....	45
3.6.2	Client Tier API	45
3.6.3	Client Tier SDK Tools	45
3.6.4	UFM Server.....	45
3.6.5	Subnet Manager	45
3.6.6	NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) [™] Aggregation Manager	46
3.6.7	Performance Manager	46
3.6.8	Device Manager	46
3.6.9	UFM Switch Agent	46
3.6.10	Communication Protocols.....	46
3.7	Getting Familiar with UFM's Data Model	46
3.7.1	Overview of Data Model	46
4	UFM Regular Installation	48
5	UFM Installation and Initial Configuration	49
5.1	UFM Installation Steps.....	49
5.1.1	Downloading UFM Software and License File	49
5.1.2	Installing UFM Server Software	51
5.2	Running UFM Server Software.....	56
5.2.1	Running UFM Server Software in Management Mode	56
5.2.2	Running UFM Software in High Availability Mode.....	56
5.2.3	Running UFM Software in Monitoring Mode	57
5.2.4	HTTP/HTTPS Port Configuration.....	58
5.2.5	Launching UFM Web UI Session	59
5.2.6	User Authentication	59
5.2.7	Licensing	59
5.2.8	Showing UFM Processes Status	60
5.3	Upgrading UFM Software.....	60
5.4	Uninstalling UFM	60
5.5	Appendix - UFM Migration	61

5.5.1	Overview	61
5.5.2	Backup UFM configuration	61
5.5.3	Restore UFM Configuration	62
5.6	Docker Installation.....	63
5.6.1	General Prerequisites	63
5.6.2	Prerequisites for Upgrading UFM Docker Container	64
5.6.3	Step 1: Loading UFM Docker Image	64
5.6.4	Step 2: Installing UFM Docker.....	64
5.6.5	Installation Modes	65
5.6.6	Upgrading From Existing UFM Container	69
5.6.7	Logging Into UFM Web UI	71
5.7	Historical Telemetry Collection in UFM.....	71
5.7.1	Storage Considerations.....	71
6	UFM Software Installation Prerequisites.....	72
6.1	Prerequisites for UFM Server Software Installation.....	72
6.2	Additional Prerequisites for UFM High Availability (HA) Installation	72
7	UFM System Requirements	73
8	UFM Server Health Monitoring	74
8.1	UFM Health Configuration	74
8.1.1	UFM Core Files Tracking.....	77
8.2	Example of Health Configuration.....	77
8.2.1	Event Burst Management	78
8.3	Recovery from Consecutive Failures	78
9	UFM Web UI	79
9.1	Fabric Dashboard	79
9.1.1	Dashboard Views and Panel Management	79
9.1.2	Dashboard Timeline Snapshots.....	80
9.1.3	Dashboard Panels	81
9.1.4	Top N Servers/Switches by Rx or Tx Bandwidth.....	81
9.1.5	Top N Congested Servers/Switches by Rx/Tx Bandwidth	85
9.1.6	Top N Utilized PKeys.....	88
9.1.7	Top N Alarmed Servers/Switches	89
9.1.8	Inventory Summary	93
9.1.9	Fabric Utilization	94


9.1.10	Recent Activities	95
9.1.11	Traffic Map	97
9.2	Network Map	104
9.2.1	Network Map Components	105
9.2.2	Selecting Map Elements	105
9.2.3	Map Information and Settings	106
9.2.4	Map View Tab	109
9.2.5	Map Zoom In Tab	109
9.2.6	Map Layouts	111
9.2.7	Information View Tab	113
9.2.8	Link Analysis	114
9.2.9	Topology Compare	119
9.2.10	Properties Tab	120
9.2.11	Network Map Elements Actions	121
9.3	Managed Elements	124
9.3.1	Devices Window	125
9.3.2	Ports Window	147
9.3.3	Virtual Ports Window	151
9.3.4	Unhealthy Ports Window	152
9.3.5	Cables Window	154
9.3.6	Groups Window	154
9.3.7	Inventory Window	157
9.3.8	PKeys Window	157
9.3.9	HCA's Window	162
9.4	Logical Elements	162
9.4.1	Environments	163
9.4.2	Networks	166
9.4.3	Logical Servers	175
9.5	Events & Alarms	179
9.6	Telemetry	181
9.7	System Health	181
9.7.1	UFM Health Tab	181
9.7.2	UFM Logs Tab	183
9.7.3	UFM Snapshot Tab	184

9.7.4	Fabric Health Tab	185
9.7.5	Daily Reports Tab.....	188
9.7.6	Topology Compare Tab	202
9.7.7	Fabric Validation Tab	205
9.7.8	IBDiagnet Tab	208
9.8	Jobs	212
9.9	Settings.....	213
9.9.1	Events Policy.....	213
9.9.2	Device Access.....	218
9.9.3	Network Management	219
9.9.4	Subnet Manager Tab	221
9.9.5	Non-Optimal Links	231
9.9.6	User Management Tab	232
9.9.7	Email.....	234
9.9.8	Remote Location	236
9.9.9	Data Streaming.....	237
9.9.10	Topology Compare	238
9.9.11	Token-based Authentication.....	239
9.9.12	Plugin Management.....	240
9.9.13	User Preferences	241
10	UFM Plugins	243
10.1	rest-rdma Plugin	243
10.1.1	Deployment Server	243
10.1.2	How to Run.....	244
10.1.3	Examples	245
10.2	NDT Plugin	246
10.2.1	Overview	246
10.2.2	Deployment	246
10.2.3	Authentication	246
10.2.4	REST API	246
10.2.5	NDT Format	247
10.2.6	Other.....	248
10.3	UFM Telemetry Fluent Streaming (TFS) Plugin	249
10.3.1	Overview	249

10.3.2	Deployment	249
10.3.3	Authentication	249
10.3.4	Rest API	249
10.4	UFM Events Fluent Streaming (EFS) Plugin.....	250
10.4.1	Overview	250
10.4.2	Deployment	250
10.4.3	Authentication	250
10.4.4	Rest API.....	250
10.5	GRPC-Streamer Plugin.....	250
10.5.1	Authentication	250
10.5.2	Create a Session to UFM from GRPC	251
10.5.3	Create New Subscription.....	251
10.5.4	Edit Known Subscription	252
10.5.5	Get List of Known Subscribers	253
10.5.6	Delete a Known Subscriber	253
10.5.7	Run a Known Subscriber Once.....	253
10.5.8	Run Streamed Data of a Known Subscriber	254
10.5.9	Run a New Subscriber Once	254
10.5.10	Run New Subscriber Streamed Data.....	255
10.5.11	Run A Serialization on All the Running Streams	256
10.5.12	Stop a Running Stream	256
10.5.13	Run a subscribe stream	256
10.5.14	Get the variables from a known subscriber.....	257
11	Troubleshooting	258
11.1	Split-Brain Recovery in HA Installation	258
12	Appendixes.....	259
12.1	Appendix - Diagnostic Utilities.....	259
12.1.1	InfiniBand Diagnostics Commands.....	259
12.1.2	Diagnostic Tools	260
12.1.3	Utilities Descriptions	261
12.2	Appendix - Supported Port Counters and Events.....	275
12.2.1	InfiniBand Port Counters.....	275
12.2.2	Supported Traps and Events	277
12.3	Appendix - Used Ports.....	282

12.4	Appendix - Configuration Files Auditing	283
12.5	Appendix - IB Router	284
12.5.1	IB Router Scripts	285
12.5.2	IB Router Configuration	288
12.6	Appendix - NVIDIA SHARP Integration.....	288
12.6.1	NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™	288
12.6.2	NVIDIA SHARP Aggregation Manager	289
12.6.3	Running NVIDIA SHARP AM in UFM	290
12.6.4	Operating NVIDIA SHARP AM with UFM.....	290
12.6.5	Monitoring NVIDIA SHARP AM by UFMHealth	291
12.6.6	Managing NVIDIA SHARP AM by UFM High Availability (HA)	291
12.6.7	NVIDIA SHARP AM Logs.....	292
12.6.8	NVIDIA SHARP AM Version	292
12.7	Appendix - AHX Monitoring	292
12.7.1	Overview	292
12.8	Appendix - UFM SLURM Integration	292
12.8.1	Prerequisites.....	292
12.8.2	Automatic Installation	292
12.8.3	Manual Installation	293
12.8.4	UFM SLURM Config File.....	293
12.8.5	Configuring UFM for NVIDIA SHARP Allocation	294
12.8.6	Prolog and Epilog.....	294
12.8.7	Integration Files.....	294
12.8.8	Running UFM-SLURM Integration.....	294
12.9	Appendix - Device Management Feature Support.....	295
12.10	Appendix - UFM Event Forwarder	298
12.11	Appendix - UFM Multisite Portal Integration.....	300
12.11.1	Configuring Multisite Agent Credentials.....	300
13	Document Revision History	301
14	High Availability	305
14.1	Overview of High Availability	305
14.1.1	HA-Related Events	305
14.1.2	HA-Related Considerations	305

14.2	High Availability Functionality	306
15	Table Enhancements.....	307
15.1	Look and Feel Improvements	307
15.2	Displayed Columns.....	307
15.3	Export All Data as CSV.....	308
16	Time Zone Converter	309
17	Cable Transceiver Temperatures.....	310
17.1	GUI Views	310
17.1.1	Alarms	310
17.1.2	Event Policy.....	310
17.2	Appendix - SM Partitions.conf File Format.....	310

 You can download a PDF version [here](#).

1 About This Document

NVIDIA® UFM® Enterprise is a powerful platform for managing InfiniBand scale-out computing environments. UFM enables data center operators to efficiently monitor and operate the entire fabric, boost application performance and maximize fabric resource utilization.

1.1 Software Download

To download the UFM software, please visit [NVIDIA's Licensing Portal](#).

If you do not have a valid license, please fill out the [NVIDIA Enterprise Account Registration](#) form to get a UFM evaluation license.

1.2 Document Revision History

For the list of changes made to this document, refer to [Document Revision History](#).


2 Release Notes


NVIDIA® UFM® is a powerful platform for managing InfiniBand scale-out computing environments. UFM enables data center operators to efficiently monitor and operate the entire fabric, boost application performance and maximize fabric resource utilization. NVIDIA® UFM®-SDN Appliance is a powerful platform for managing InfiniBand scale-out computing environments. UFM enables data center operators to efficiently monitor and operate the entire fabric, boost application performance and maximize fabric resource utilization.

2.1 Key Features

UFM provides a central management console, including the following main features: UFM-SDN Appliance provides a central management console, including the following main features:


- Fabric dashboard including congestion detection and analysis
- Advanced real-time health and performance monitoring
- Fabric health reports
- Threshold-based alerts
- Fabric segmentation/isolation
- Quality of Service (QoS)
- Routing optimizations
- Central device management
- Task automation
- Logging
- High availability
- Daily report: Statistical information of the fabric during the last 24 hours
- Event management
- Switch auto-provisioning
- UFM-SDN Appliance in-service software upgrade
- Fabric validation tests
- Client certificate authentication
- IPv6 on management ports


 Prior to installation, please verify that all prerequisites are met. Please refer to [System Requirements](#).

 The Logical Server Model Management feature is going to be deprecated in UFM v6.12.0.

2.2 Changes and New Features


This section lists the new and changed features in this software version.

 For an archive of changes and features from previous releases, please refer to [Changes and New Features History](#).

 The items listed in the table below apply to all UFM license types.

Feature	Description
UFM Discovery and Device Management	<ul style="list-style-type: none"> InBand autoscovcovery of switches' IP addresses using ibdiagnet Discovering the device's PSID and FW version using ibdiagnet by default instead of using an SM vendor plugin
CPU Affinity	Enabling the user to control CPU affinity of UFM's major processes
gRPC API	Added support for streaming UFM REST API data over gRPC as part of new UFM plugin. Refer to gRPC-Streamer Plugin
Telemetry	<ul style="list-style-type: none"> Added support for flexible counters infrastructure (ability to change counter sets that are sampled by the UFM) Updated the set of available counters for Telemetry (removed General counters from default view: Row BER, Effective BER and Device Temperature. Now available through the secondary telemetry instance). Refer to Secondary Telemetry
EFS UFM Plugin	Added support for streaming UFM events data to FluentD destination as part of a new UFM plugin. Refer to UFM Telemetry Fluent Streaming (TFS) Plugin
General UI Enhancements	<ul style="list-style-type: none"> Displayed columns of all tables are persistent per user, with the option to restore defaults. Refer to Displayed Columns Improved look and feel in Network Map. Refer to Network Map Added Reveal Uptime to the general tab in the devices information tabs. Refer to Device General Tab
High Availability Deployment	<ul style="list-style-type: none"> Added support for joining a new UFM device into the HA pair without stopping the UFM HA (in case of a secondary UFM node permanent failure). For more information, refer to Installing UFM Server Software for High Availability Changed UFM HA package installation command parameters. For more information, refer to Installing UFM Server Software for High Availability
REST APIs	<p>Added support for PKey filtering for default session data. Refer to Get Default Monitoring Session Data by PKey Filtering.</p> <p>Added support for filtering session data by groups. Refer to Monitoring Sessions REST API.</p> <p>Added support for resting all unhealthy ports at once. Refer to Mark All Unhealthy Ports as Healthy at Once</p> <p>Added support for presenting system uptime in UFM REST API. Refer to Systems REST API.</p>
Deployment Installation	UFM installation is now based on Conda-4.12 (or newer) for python3.9 environment and third party packages deployments.
NVIDIA SHARP Software	Updated NVIDIA SHARP software version to v3.1.1.
UFM Logical Elements	UFM Logical Elements (Environments, Logical Servers, Networks) views are deprecated and will no longer be available starting from UFM v6.12.0 (January 2023 release)

 Integrated with MFT version mft-4.22.1-417.

 For bare metal installation of UFM, it is required to install MLNX_OFED 5.X (or newer) before the UFM installation.

Please make sure to use the UFM installation package that is compatible with your setup as detailed in [Bare Metal Deployment Requirements](#).

2.2.1 Unsupported Functionalities/Features

The following distributions are no longer supported in UFM:

- RH7.0-RH7.7 / CentOS7.0-CentOS7.7
- SLES12 / SLES 15
- EulerOS2.2 / EulerOS2.3
- Mellanox Care (MCare) Integration
- UFM on VM (UFM with remote fabric collector)
- Logical server auditing
- UFM high availability script - `/etc/init.d/ufmha` - is no longer supported.



In order to continue working with `/etc/init.d/ufmha` options, use the same options using the `/etc/init.d/ufmd` script.

For example:

Instead of using `/etc/init.d/ufmha model_restart`, please use `/etc/init.d/ufmd model_restart` (on the primary UFM server)

Instead of using `/etc/init.d/ufmha sharp_restart`, please use `/etc/init.d/ufmd sharp_restart` (on the primary UFM server)

The same goes for any other option that was supported on the `/etc/init.d/ufmha` script

2.3 Installation Notes

2.3.1 Supported Devices

2.3.1.1 Supported NVIDIA Externally Managed Switches

Type	Model	Firmware Version
NDR switches	<ul style="list-style-type: none">• MQM9790	31.2010.2110
HDR switches	<ul style="list-style-type: none">• MQM8790	27.2010.3004
EDR switches	<ul style="list-style-type: none">• SB7790• SB7890	15.2008.2946
FDR switches	<ul style="list-style-type: none">• SX6025• SX6015• SX6005	11.1500.0106

2.3.1.2 Supported NVIDIA Internally Managed Switches


Type	Model	Tested OS Version
NDR switches	<ul style="list-style-type: none"> MQM9700 	MLNX-OS 3.10.3002
HDR switches	<ul style="list-style-type: none"> MQ8700 MCS8500 TQ8100-HS2F TQ8200-HS2F 	MLNX-OS 3.10.3100
EDR switches	<ul style="list-style-type: none"> SB7700 SB7780 SB7800 CS7500 CS7510 CS7520 	MLNX-OS 3.6.5010
FDR switches	<ul style="list-style-type: none"> SX6012 SX6018 SX6036 SX6506 SX6512 SX6518 SX6536 SX1012 SX6710 SX6720 SX1700 SX1710 	MLNX-OS 3.6.8008
Long-haul	<ul style="list-style-type: none"> MTX6100 MTX6240 MTX6280 	MLNX-OS 3.6.8008
InfiniBand-Ethernet Gateways	SX6036G (FDR)	MLNX-OS 3.6.8008


2.3.2 System Requirements


2.3.2.1 Bare Metal Deployment Requirements


Platform	Type and Version
OS and Kernel ^(a)	64-bit OS: <ul style="list-style-type: none"> RedHat 7.8: 3.10.0-1127.el7.x86_64 RedHat 7.9: 3.10.0-1160.el7.x86_64 RedHat 8.1: 4.18.0-147.el8.x86_64 RedHat 8.2: 4.18.0-193.el8.x86_64 RedHat 8.4: 4.18.0-80.el8.x86_64 CentOS 7.8: 3.10.0-1127.el7.x86_64 CentOS 7.9: 3.10.0-1160.el7.x86_64 CentOS 8.2: 4.18.0-193.el8.x86_64 CentOS 8 Stream: 5.4.0 FAIR OS 22.08 Ubuntu 18.04: 4.15 Ubuntu 20.04: 5.4.0

Platform	Type and Version
CPU ^(b)	x86_64
HCA ^s	<ul style="list-style-type: none"> • NVIDIA ConnectX®-4 with Firmware 12.12.xxxx and above • NVIDIA ConnectX®-5 with Firmware 16.19.1200 and above • NVIDIA ConnectX®-6 with Firmware 20.24.1000 and above • NVIDIA ConnectX®-7 with Firmware 28.33.1014 and above • NVIDIA BlueField with Firmware 24.33.900 and above • NVIDIA BlueField-2 with Firmware 24.33.900 and above
OFED ^(c)	<ul style="list-style-type: none"> • MLNX_OFED 5.X

 (a) CentOS 8 Stream and RHEL8.4 can be installed without MLNX_OFED; inbox drivers can be used instead.
 (b) CPU requirements refer to resources consumed by UFM. You can also dedicate a subset of cores on a multicore server. For example, 4 cores for UFM on a 16-core server.
 (c) For supported HCAs in each MLNX_OFED version, please refer to MLNX_OFED Release Notes.

 For running SHARP Aggregation Manager within UFM, it is recommended to use MLNX_OFED-5.4.X version or newer.

 Installation of UFM on minimal OS distribution is not supported.

 UFM does not support systems in which NetworkManager service is enabled.
 Before installing UFM on RedHat OS, make sure to disable the service.

2.3.2.2 Docker Installation Requirements

Component	Type and Version
Supported OS	<ul style="list-style-type: none"> • RHEL7 • RHEL8 • Ubuntu18.04 • Ubuntu20.04 • Ubuntu22.04 - TBD: Lenny

2.3.2.3 UFM Server Resource Requirements Per Cluster Size

Fabric Size	CPU Requirements*	Memory Requirements	Disk Space Requirements	
			Minimum	Recommended
Up to 1000 nodes	4-core server	4 GB	20 GB	50 GB
1000-5000 nodes	8-core server	16 GB	40 GB	120 GB

Fabric Size	CPU Requirements*	Memory Requirements	Disk Space Requirements	
			Minimum	Recommended
5000-10000 nodes	16-core server	32 GB	80 GB	160 GB
Above 10000 nodes	Contact NVIDIA Support			

2.3.2.4 UFM GUI Client Requirements

The platform and GUI requirements are detailed in the following tables:

Platform	Details
Browser	Edge, Internet Explorer, Firefox, Chrome, Opera, Safari
Memory	<ul style="list-style-type: none"> • Minimum: 2 GB • Recommended: 4 GB

2.3.2.5 MFT Package Version

Platform	Details
MFT	Integrated with MFT version mft-4.22.1-417

2.3.2.6 UFM SM Version

Platform	Type and Version
SM	UFM package includes SM version 5.13.0


2.3.2.7 UFM NVIDIA SHARP Software Version

Platform	Type and Version
NVIDIA® Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™	UFM package includes NVIDIA SHARP software version 3.1.1


2.3.3 Software Update from Prior Versions

The installer detects versions previously installed on the machine and prompts you to run a clean install of the new version or to upgrade while keeping user data and configuration unchanged.

The upgrade from previous versions maintains the existing database and configuration, allowing a seamless upgrade process.

 Upgrading UFM Enterprise software version is supported up to two previous GA software versions (GA -1 or -2).
For example, if you wish to upgrade to UFM Enterprise v6.11.0, it is possible to do so only from UFM Enterprise v6.9.0 or v6.10.0.


For detailed installation and upgrade instructions, refer to the *UFM Quick Start Guide* or the *UFM User Manual*.

 Due to a possible conflict, SM and MFT installed by the MLNX_OFED must be uninstalled. The installation procedure will detect and print all MLNX_OFED packages that must be removed.


 It is recommended to upgrade to the latest UFM version from the last 2 GA releases that came before it. Upgrading from older UFM versions may result in failures.

2.4 Bug Fixes in This Release

Ref. #	Description
3187979	Description: Wrong behavior in port failover to second host (instead second bonded port)
	Keywords: Wrong behavior, Port, Failover
	Discovered in Release: 6.10.0
3234082	Description: UFM WebUI unresponsive after failover issue
	Keywords: UFM, WebUI, failover
	Discovered in Release: 6.10.0
3199572	Description: Incorrect Tier reporting in the UFM events
	Keywords: Tier, Incorrect Report
	Discovered in Release: 6.10.0
3187979	Description: Wrong behavior in port failover to second host (instead second bonded port)
	Keywords: Wrong behavior, Port, Failover
	Discovered in Release: 6.10.0

 For an archive of bug fixes from previous releases, please refer to [Bug Fixes History](#).

2.5 Known Issues in This Release

 For a list of known issues from previous releases, please refer to [Known Issues History](#).

Ref #	Issue
3240664	Description: This software release does not support upgrading the UFM Enterprise version from the latest GA version (v6.11.0). UFM upgrade is supported in UFM Enterprise v6.9.0 and v6.10.0.
	Workaround: N/A
	Keywords: UFM Upgrade
3242332	Description: Upgrading MLNX_OFED uninstalls UFM
	Workaround: Upgrade UFM to a newer version (v6.11.0 or newer), then upgrade MLNX_OFED
	Keywords: MLNX_OFED, Uninstall, UFM
3237353	Description: Upgrading from UFM v6.10 removes MLNX_OFED crucial packages
	Workaround: Reinstall MLNX_OFED/UFM
	Keywords: MLNX_OFED, Upgrade, Packages
N/A	Description: Running UFM software with external UFM-SM is no longer supported
	Workaround: N/A
	Keywords: External UFM-SM

2.6 Changes and New Features History

Feature	Description
Rev 6.10.0	
System health enhancements	Add support for the periodic fabric health report, and reflected the ports' results in UFM's dashboard
UFM Plugins Management	Add support for plugin management via UFM web UI
UFM Extended Status	<ul style="list-style-type: none"> Add support for showing UFM's current processes status (via shell script) Added REST API for exposing UFM readiness
Failover to Other Ports	Add support for SM and UFM Telemetry failover to other ports on the local machine
UFM Appliance Upgrade	Added a set of REST APIs for supporting the UFM Appliance upgrade
Configuration Audit	Add support for tracking changes made in major UFM configuration files (UFM, SM, SHARP, Telemetry)
UFM Plugins	Add support for new SDK plugins
Telemetry	Add support for statistics processing based on UFM telemetry csv format

UFM High Availability Installation	UFM high availability installation has changed and it is now based on an independent high availability package which should be deployed in addition to the UFM Enterprise standalone package. for further details about the new UFM high availability installation, please refer to - Installing UFM Server Software for High Availability
Rev 6.9.0	
NDR Support	Full E2E NDR including ConnectX-7 HCAs Family (Discovery and Monitoring)
Cable FW burn	Add support for multiple switches with multiple FW images burning
Events	Add support for monitoring and alerting on cable transceiver temperatures over threshold
	Improve SM traps handling (offloading SM traps handling to a separated process)
	Add option for setting events persistency (keeping max last X events) for showing upon UFM startup
	Add option for consolidating similar events on the UFM Web UI Events Log View
SHARP	Add support for failover to secondary bond port in case of IB interface failure
	Add option to override SHARP <code>smx_sock_interface</code> based on UFM <code>fabric_interface</code> (gv.cfg)
	Add option to set SHARP AM <code>ib_port_guid</code> based on UFM <code>fabric_interface</code> (gv.cfg)
SM	Add support for tracking SM configuration changes (configuration history)
	Add support for pkey assignment validation (for user defined pkey assignment only)
Client Certificate Authentication	Add support for client certificate authentication
	Add option to push bootstrap certificate to the UFM via REST API
Configuration Migration (backup / restore)	Add option to migrate UFM configuration from bare metal UFM to a docker container based UFM
MFT Integration Enhancement	Add support for MFT based operation (FW burning, cable info) while <code>m_key</code> / <code>vs_key</code> are configured on SM
Logging	Adding option to configure UFM log folder location
UFM Health	Add option for users to add customized health tests based on scripts (Python / Bash)
Web UI Enhancements	Add support for user defined modular UFM dashboard views (based on available list of pre-defined panels)
	Add support for UFM dashboard timeline (for viewing historical dashboard views)
	Enhance the dashboard inventory view for showing elements (HCAs, Switches, Cables, Gateways, Routers) by version
	Add support for user defined modular UFM telemetry persistent dashboard (Telemetry View)
	Adding option for viewing Web client data based on local client time or UFM server time
	Add option to select UFM look and feel between dark mode and light mode (default is light mode)
	Add support for hierarchical view when presenting the network map elements.
	Add option for selecting the displayed columns for all data tables.

	Add option for exporting all table data into CSV (not only the current displayed page data)
	Improved view of the ports table (port name, speed and width)
	Add option to show disabled/down ports
	Add support for Web UI usage statistics collection
	Add option for sending test email
Telemetry	Add support for updating Telemetry package within installed UFM Enterprise.
UFM Plugins	Add support for running UFM plugins within UFM docker container
	Add support for AHX monitoring plugin
Supported OSs	Add support for installing UFM on Ubuntu18 (Standalone and High availability modes)
	Add support for installing UFM on CentOS 7.9/Redhat7.9
	Add support for installing UFM on FAIR OS 22.03
Rev 6.8	
UFM Telemetry	Changed the Telemetry infrastructure from UFM Telemetry docker container to UFM Telemetry bare metal
	Performance improvements for supporting telemetry on large scale fabrics (up to 216,000 ports fabric)
	Live sessions enhancements - adding support for multiple telemetry sessions based on one UFM Telemetry instance
	Add support for collecting historical telemetry (all fabric ports counters) by default
Unhealthy Ports	Add option (configurable) for automatically Isolating ports which were detected with high BER
	Add option to present unhealthy port table by the connection type (switch-switch or switch-host).
	Add option to mark selected device as unhealthy
UFM Plugins - REST over RDMA	Add support for REST API over RDMA plugin (allowing execution of UFM REST API requests over the InfiniBand fabric)
	Add ability to run Linux command line command, including ibdiagnet, over rdma
UFM Plugins - NDT	Add support for NDT (CSV formatted topology) comparison with UFM fabric detected topology
Fabric Validation Tests	Add context menu options for selected results of fabric validation tests based of UFM model objects (Devices and Ports).
	Add support for Socket-direct mode reporting (Inventory)
	Add support for SHARP Aggregation Manager health tests
	Add support for Tree Topology Analysis support in UFM
Events Policy	Add new category for Events Policy - Security
	Add new UFM events indicating Pkey assignment of guides and removal of guides from Pkey
	Add new UFM events which are triggered when duplicated node or port GUIDs are detected in the fabric
	Add new event for indicating switch down reported by SM

UFM SDK	Add option to get topology via UFM REST API and stream it out to an external destination
Virtualization	Add option to assign selected virtual ports to a specified PKEY (via UFM Web UI)
Cable Information	Showing Link grade in Cable info
Network Map	Add support for network map topology persistency on server side.
UFM Web UI	Add option to copy and paste tables content (GUIDS and LIDS) via UFM Web UI
UFM Authentication	Add support for token based authentication
UFM Slurm Integration	Add several UFM-SLURM Integration Improvements
UFM Docker container	Several docker Enhancement mainly for improving the deployment procedure
SM Configuration	Setting AR (Adaptive Routing) Up Down as the default routing configuration in UFM / SM (for new UFM installations)
UFM REST API	Add Support for CloudX API in UFM for OpenStack integration and allow auto provisioning of the InfiniBand fabric
NDR support	Add support for discovering and monitoring Nvidia NDR switches.
Installation	Updated UFM installation to run without docker dependencies (docker service is no longer required for the UFM installation)
Supported OSs	Add support for installing UFM on CentOS 8 stream, kernel 5.4
UFM High Availability	Add support for independent high availability package (based on Pacemaker and DRBD) which server as the basis for UFM containers high availability deployment
Rev 6.7	
UFM Telemetry-based monitoring	Changed UFM's monitoring mechanism to be based on UFM Telemetry instead of IBPM (for both default and live telemetry sessions)
IB router & IB gateway monitoring	Added support for monitoring of InfiniBand router and gateway ports
SHARP aggregation manager events	Added support for showing SHARP aggregation manager events in UFM
SHARP over UCX	Added support fAdded support for automatically isolating ports with high BER (with monitoring being performed based on the Symbol BER)or running SHARP aggregation manager over UCX
Periodic topology check	Added support for periodic run of topology comparisons and reporting of topology changes against preset topology
Visual topology difference	Added option to view visual-representation of topology changes in the network topology map (as compared to a "master" or user-defined topology)

System dump for externally managed switches	Added support for collecting system dump for externally managed switches
Syslog settings via web UI	Added support for configuring UFM syslog settings via UFM web UI
Upgrade for group of switches	Added support for software/firmware upgrade for a group of switches
NDR switches readiness	Added support for discovery and management of NDR switches
Transition to file-based storage	Transitioned from Mysql to SQLite DB for persistent model objects
Counters over threshold	Added support for showing telemetry counters over a predefined threshold when using historical statistic collection
HDR cables burning	Added support for burning HDR cable transceivers for selected switches
Dragonfly+ topology analysis	Added fabric validation test to validate an existing Dragonfly+ topology
Form-based authentication	Added support for enhanced authentication mechanism for UFM REST API
Web UI enhancements	<ul style="list-style-type: none"> • Context switch for events & alarms • Zoom-in and filtering options for network map • Updated live session members
Uploading ibdiagnet results	Added option to upload periodic ibdiagnet results to any remote destination over SCP or SFTP
Telemetry API enhancements	Added option to retrieve short counter format or specified counters only for monitoring session data REST API
SLURM integration enhancements	Added support for token-based authentication, instead of basic authentication, to connect UFM
High BER ports list	Added support for displaying all ports with high BER (from the Ports view) as well as the ability to mark them as unhealthy
OpenSM GUID list	Added support for new OpenSM traps (UFM Events) which indicate activity in the fabric of unexpected OpenSM
UFM docker enhancements	Added support for UFM docker installer container to simplify UFM container installation and upgrade procedures (for both standalone & HA deployments)
REST API	Links API has been updated with two additional fields: source_port_name, destination_port_name.
BlueField DPUs support	Added support for management of BlueField DPU devices in the fabric

Topology map enhancements	Added support for selection and running of actions on multiple elements in network map
REST API	The response format returned by the API endpoint at /ufmRest/resources/systems has changed. Please check this link for the updated API response format.
Rev 6.6.0	
Licensing	Added support for UFM subscription license
Periodic ibdiagnet	Added ability to execute ibdiagnet periodically and collect the generated logs
Sysdump	Added ability to perform sysdump on internally managed switches
	Added ability to perform sysdump on hosts
Event streaming	Added ability to stream UFM events via FluentBit plugin
Virtualization	Added support for port virtualization including virtualization events
Telemetry	Added support for new telemetry capabilities and showing historical data reports
Multiple rail optimization	Added support for multiple rail optimization validation test
MCARE	Added support for MCARE integration with UFM over REST API
Supported OS	Added support for Red Hat and CentOS versions 7.7, 7.8, 8.1, and 8.2
MLNX_OFED	MLNX_OFED v5.1 integration for both regular and docker container deployments
Log history	Added support for showing history of UFM, OpenSM, and Events logs
Multi-HCA grouping	Added support for grouping Windows Multi-HCA
Congestion map	Added support for traffic and congestion map for used-defined port group
IB Gateway	Added support for IB Gateway discovery
IB Router	Added support for IB Router discovery
Topology comparison	Enhanced topology diff reports
Look and feel	Updated look and feel to NVIDIA theme
Rev 6.5.2	
New licensing mechanism	Added support for the new UFM subscription license (keeping backward compatibility with old license file)
Periodic ibdiagnet execution	Added the option to execute ibdiagnet command (using any supported flag) via UFM web UI
System dump for switches and hosts	Added support for running and uploading system dump from internally managed switches and hosts via UFM web UI
Pagination	Added support for paginating web UI tables for better responsiveness
PKey versioning	Added support for PKey versioning to indicate PKey related changes

Integration with MCare	Add support for UFM-Mellanox Care integration over UFM REST APIs
New supported OSs	Add support for installing UFM on RHEL 7.7 and RHEL 7.8 and SLES 12 SP5
Rev 6.5.1	
Large scale support improvements	Improved the handling of IB Performance Monitoring (IBPM) statistic data and generation of events in UFM for large scale fabrics
	Offloaded handling of topology changes of large scale fabrics to a new process in UFM
UFM Safe Startup	Set all UFM ports to full membership upon UFM startup so that all UFM IB applications (e.g. OpenSM, IBPM, ibdiagnet) have full access to the IB fabric
IBPM Resiliency	If UFM's fabric interface is configured as a bond, UFM restarts the IBPM on the secondary interface (the new active interface) if the active interface fails
Rev 6.5.0	
Large scale support improvements	Added support for running UFM in large scale setup (up to 40K nodes)
Multi-port SM	Added an option to run UFM-SM on multiple pre-configured ports
Python3 support	Unified UFM code to run using Python3 code for all supported distributions (RH7, RH8, SLES12, SLES15)
Python virtual environment support	Used Python virtual environment to avoid UFM installation conflicts with system packages
Cable lifecycle events	Added support for new cables lifecycle events (e.g. cable added, removed, changed location and duplicated)
Updating port speed via UFM	Added REST API to control the rate limit of physical and virtual ports
Enhanced SM configuration via UFM	Added REST API for updating SM congestion control and adaptive routing parameters
IB Gateway support	Added support for discovery and monitoring of IB Gateway
Ports display	Present all disabled ports as well for each device in the right ports tab
Externally managed switch reset option	Added support for resetting externally managed switches
MetroX-2 system support	Added support for MetroX-2 systems TQ8100-HS2F and TQ8200-HS2F
UFM-SHARP resources allocation integration	Added REST API to allocate and deallocate SHARP resources
UFM Multisite Portal	Single pane of glass to manage multiple UFM's in one console

Mlxlink support	Added option to display enhanced cable information, for selected port, using mlxlink
-----------------	--

2.7 Bug Fixes History

Ref. #	Description
3107006	Description: Using GET All Modules REST API (GET /ufmRest/resources/modules), returns N/A in device_name.
	Keywords: Modules, N/A, device_name
	Discovered in Release: 6.9
3076817	Description: Upgrading to the latest UFM version (UFMAPL_4.8.0.6_UFM_6.9.0.7), the UFM WEB UI shows log and error messages with "invalid date."
	Keywords: WEB UI, "invalid date"
	Discovered in Release: 6.9
3060127	Description: UFM WEB UI - Ports REST API returns tier parameters as N/A in response
	Keywords: WEB UI, tier, N/A
	Discovered in Release: 6.9
3052660	Description: UFM monitoring mode is not working
	Keywords: Monitoring, mode
	Discovered in Release: 6.9
3031121	Description: Network map showing a link between QM8790 and Manta Ray leaf having BW of >20,000 Gb/s
	Keywords: Network Map, BW, 20,000
	Discovered in release: 6.8.0
3003366	Description: UFM Starting and Stopping On Its Own Since Merge
	Keywords: Start, Stop
	Discovered in release: 6.7.0
2968236	Description: Fabric health Old Alerts and events do not clear
	Keywords: Fabric Health, Alerts, clear
	Discovered in release: 6.8.0
2957984	Description: BER Not Being Read or Reported
	Keywords: BER, Not, Reported
	Discovered in release: 6.8.0
3032227	Description: UFM UFMAPL_4.7.0.3_UFM_6.8.0.6 lists one of my skyways as "host" instead of "gateway"
	Keywords: skyway, gateway, host
	Discovered in release: 6.8.0
2966472	Description: UFM Fabric health BER_CHECK warnings

Ref. #	Description
	Keywords: Fabric Health, BER, check
	Discovered in release: 6.8.0
2801258	Description: UFM failed to serve incoming REST API requests
	Keywords: REST API, hang, unresponsive
	Discovered in release: 6.7.0
2782069	Description: UFM APL 4.6 BER not reported (None) in event logs
	Keywords: BER, events, log
	Discovered in release: 6.7.0
2744757	Description: UFM health test: CheckSMConnectivityOnStandby should consider multiple GUIDs on a port
	Keywords: UFM Health, SM connectivity, multiple guides
	Discovered in release: 6.7.0
2830281	Description: UFM (container) is not starting after server reboot
	Keywords: UFM Container, reboot
	Discovered in release: 6.7.0
2804807	Description: UFM WEB GUI becomes Unresponsive and Event/REST API log stops printing
	Keywords: Web UI, unresponsive
	Discovered in release: 6.7.0
2699393	Description: IPMI console login connects to CentOS (UM docker OS) instead of Ubuntu (host OS) after UFM docker installation.
	Keywords: IPMI; CentOS; Login
	Discovered in release: 6.6.1
2638032	Description: Wrong module (line/spine) label appears in effective BER event.
	Keywords: Module; Effective; BER; Event
	Discovered in release: 6.4.1
2618603	Description: UFM failover is not working when bond0 is configured with IPoIB.
	Keywords: Failover, Bond; IPoIB
	Discovered in release: 6.6.1
2615514	Description: UFM software no longer supports license type "UFM APPLIANCE".
	Keywords: License; UFM Appliance
	Discovered in release: 6.5.2
2589617	Description: UFM stopped to discover topology on SuperPOD environment.
	Keywords: Stopped; discover
	Discovered in release: 6.5.2
2335141	Description: Memory leak discovered in ModelMain.py process.
	Keywords: Memory leak

Ref. #	Description
	Discovered in Release: 6.5.1
	Fixed in Release: 6.5.2
2300082	Description: CMP python error
	Keywords: Python, error
	Discovered in Release: 6.5.1
	Fixed in Release: 6.5.2
2373665	Description: UFM license check of UFM permanent license generates invalid license status at the UFM Health Report.
	Keywords: Permanent license; UFM health report
	Discovered in Release: 6.5.1
	Fixed in Release: 6.5.2
2125784	Description: Some commands appear for users with monitor privileges which are not functional. It is recommended not to use this user role.
	Keywords: Monitor, permissions, user
	Discovered in Release: 4.2.0
	Fixed in Release: 6.5.1
-	Description: Performance degradation caused by OpenSM changing the default rate limit of management PKey (0x7fff) to 2.5 GB/s instead of 10GB/s.
	Keywords: OpenSM, Degradation, rate limit
	Discovered in version: 4.2.0
	Fixed in Release: 6.5.1
-	Description: Each HCA is discovered and represented as a separate host. A host with multiple HCAs will be represented as multiple host instances.
	Keywords: Fabric Topology
	Fixed in Release: 6.5.1
1967348	Description: Email sender address cannot contain more than one period (".") in the domain name.
	Keywords: Email, sender, period
	Discovered in Release: 6.3
	Fixed in Release: 6.4
2069425	Description: SMTP server username cannot have more than 20 characters.
	Keywords: Email
	Discovered in Release: 6.3
	Fixed in Release: 6.4
1914379	Description: MellanoxCare service can now communicate with UFM (valid only when http communication is configured between MCare and UFM).
	Keywords: MellanoxCare, http, https
	Discovered in Release: 6.2
	Fixed in Release: 6.3

Ref. #	Description
1783048	Description: Opening UFM web UI in monitoring mode is now supported.
	Keywords: Web UI, monitoring mode
	Discovered in Release: 6.2
	Fixed in Release: 6.3
1691882	Description: UFM Agent now is now part of the UFM web UI.
	Keywords: UFM Agent
	Discovered in Release: 6.1
	Fixed in Release: 6.3
1793244	Description: UFM/module temperature thresholds notifications.
	Keywords: Temperature thresholds
	Discovered in Release: 6.1
	Fixed in Release: 6.3
1678669	Description: Fixed an issue where UFM HA prerequisite script was checking for wrong Virtual IP port argument.
	Keywords: UFM HA, prerequisite, Virtual IP, port
	Discovered in Release: 6.1
	Fixed in Release: 6.2
1706226	Description: Fixed an issue where MLNX_OS credentials were missing at the device "access_credentials" menu (the issue was detected on old Java based GUI). At the new UFM Web UI - MLNX_OS credentials are represented by HTTP credentials.
	Keywords: MLNX_OS, credentials
	Discovered in Release: 6.1
	Fixed in Release: 6.2
1486595	Description: Fixed an issue where CentOS 7.5 was not recognized as RHEL 7 flavor upon installation.
	Keywords: Installation, CentOS, RHEL
	Discovered in: 6.0
	Fixed in: 6.1
1358248	Description: Fixed the issue where ibdiagnet's unresponsiveness when using the get_physical_info flag caused UFM to hang.
	Keywords: ibdiagnet
	Discovered in: 5.10
	Fixed in: 6.0
1294010	Description: Fixed the issue where partition configuration was lost after upgrading to UFM version 5.9.6 and restarting the server.
	Keywords: partitions.conf, PKey, configuration
	Discovered in: 5.9.6
	Fixed in: 5.10

Ref. #	Description
1276539	Description: Updated report execution command in order to avoid the following false warning of wrong link speed during topology comparison.
	Keywords: Topology compare report
	Discovered in: 5.9.6
	Fixed in: 5.10
1131286	Description: Fixed a memory leak of UFM's main process when running multiple reports periodically.
	Keywords: Memory leak, reports
	Discovered in: 5.9
	Fixed in: 5.9.6
1064349	Description: Fixed an issue where UFM reported false alarm about OpenSM irresponsiveness (sminfo command returned with failure).
	Keywords: OpenSM, sminfo
	Discovered in: 5.8
	Fixed in: 5.9.6
987236	Description: Fixed a web UI security issue by changing the SSL certificate RSA keys' size to 2048 bit (instead of 1024).
	Keywords: Web UI, security, certificate, apache
	Discovered in: 5.8
	Fixed in: 5.9
965302	Description: Fixed UFM HA installation with non-standard file mode creation mask (umask 000).
	Keywords: HA, umask
	Discovered in: 5.8
	Fixed in: 5.9

2.8 Known Issues History

Ref #	Issue
3144732	Description: By default, a managed Ubuntu 22 host will not be able to send system dump (sysdump) to a remote host as it does not include the sshpass utility.
	Workaround: In order to allow the UFM to generate system dump from a managed Ubuntu 22 host, install the sshpass utility prior to system dump generation.
	Keywords: Ubuntu 22, sysdump, sshpass
3129490	Description: HA uninstall procedure might get stuck on Ubuntu 20.04 due to multipath daemon running on the host.
	Workaround: Stop the multipath daemon before running the HA uninstall script on Ubuntu 20.04.
	Keywords: HA uninstall, multipath daemon, Ubuntu 20.04

Ref #	Issue
3147196	Description: Running the upgrade procedure on bare metal Ubuntu 18.04 in HA mode might fail.
	Workaround: For instructions on how to apply the upgrade for bare metal Ubuntu 18.04, refer to High Availability Upgrade for Ubuntu 18.04 .
	Keywords: Upgrade, Ubuntu 18.04, Docker Container, failure
3145058	Description: Running upgrade procedure on UFM Docker Container in HA mode might fail.
	Workaround: For instructions on how to apply the upgrade for UFM Docker Container in HA, refer to Upgrade Container Procedure .
	Keywords: Upgrade, Docker Container, failure
3061449	Description: Upon upgrade of UFM all telemetry configurations will be overridden with the new telemetry configuration of the new UFM version.
	Workaround: If the telemetry configuration is set manually, the user should set up the configuration after upgrading the UFM for the changes to take effect. Telemetry manual configuration should be set on the following telemetry configuration file right after UFM upgrade: <code>/opt/ufm/conf/telemetry_defaults/launch_ibdiagnet_config.ini</code> .
	Keywords: Telemetry, configuration, upgrade, override.
3053455	Description: UFM “Set Node Description” action for unmanaged switches is not supported for Ubuntu18 deployments
	Workaround: N/A
	Keywords: Set Node Description, Ubuntu18
3053455	Description: UFM Installations are not supported on RHEL8.X or CentOS8.X
	Workaround: N/A
	Keywords: Install, RHEL8, CentOS8
3052660	Description: UFM monitoring mode is not working
	Workaround: In order to make UFM work in monitoring mode, please edit telemetry configuration file: <code>/opt/ufm/conf/telemetry_defaults/launch_ibdiagnet_config.ini</code> Search for <code>arg_12</code> and set empty value: <code>arg_12=</code> Restarting the UFM will run the UFM in monitoring mode. Before starting the UFM make sure to set: <code>monitoring_mode = yes</code> in <code>gv.cfg</code>
	Keywords: Monitoring, mode
3054340	Description: Setting non-existing log directory will fail UFM to start
	Workaround: Make sure to set a valid (existing) log directory when setting this parameter (<code>gv.cfg</code> à <code>log_dir</code>)
	Keywords: Log, Dir, fail, start
-	Description: Restoring HA standby node and configuring UFM HA with external UFM-Subnet Managers are not supported on Ubuntu bare-metal deployments
	Workaround: N/A
	Keywords: HA standby node, bare-metal
2887364	Description: After upgrading to UFM6.8, in case UFM failed over to the secondary node, trying to get cable information for selected port will fail.

Ref #	Issue
	<p>Workaround: On the secondary UFM node, copy the following files to /usr/bin/ folder:</p> <ul style="list-style-type: none"> • /usr/flint • /usr/flint_ext • /usr/mlxcables • /usr/mlxcables_ext • /usr/mlxlink • /usr/mlxlink_ext <p>trying to get cable information on the secondary UFM node should work now.</p> <p>Keywords: upgrade, failover, cable information</p>
2784560	<p>Description: Intentional stop for master container and start it again or reboot of master server will damage the HA failover option</p> <p>Workaround: manually restart UFM cluster</p> <p>Keywords: UFM Container; Reboot, Failover</p>
2872513	<p>Description: after rebooting master container, Failover will be triggered twice (once to the standby and then back again to the master container)</p> <p>Workaround: N/A</p> <p>Keywords: UFM Container, reboot, failover</p>
2863388	<p>Description: Fail to get cables info for NDR Split Port.</p> <p>Workaround: N/A</p> <p>Keywords: Cable, NDR, Split</p>
N/A	<p>Description: In case of using SM mkey per port, several UFM operations might fail (get cable info, get system dump, switch FW upgrade)</p> <p>Workaround: N/A</p> <p>Keywords: SM, mkey per port</p>
2702950	<p>Description: Internet connection is required to download and install SQLite on the old container during software the upgrade process.</p> <p>Workaround: N/A</p> <p>Keywords: Container; upgrade</p>
2694977	<p>Description: Adding a large number of devices (~1000) to a group or a logical server, on large scale setup takes ~2 minutes.</p> <p>Workaround: N/A</p> <p>Keywords: Add device; group; logical server; large scale</p>
2710613	<p>Description: Periodic topology compare will not report removed nodes if the last topology change included only removed nodes.</p> <p>Workaround: N/A</p> <p>Keywords: Topology comparison</p>
2698055	<p>Description: UFM, configured to work with telemetry for collecting historical data, is limited to work only with the configured HCA port. If this port is part of a bond interface and a failure occurs on the port, collection of telemetry data via this port stops.</p> <p>Workaround: Reconfigure telemetry with the new active port and restart it within UFM.</p> <p>Keywords: Telemetry; history; bond; failure</p>

Ref #	Issue
2705974	<p>Description: If new ports are added after UFM startup, the default session REST API (GET /ufmRest/monitoring/session/0/data) will not include port statistics for the newly added ports.</p> <p>Workaround: Reset the main UFM.</p> <ul style="list-style-type: none"> For UFM standalone - <code>/etc/init.d/ufmd model_restart</code> For UFM HA - <code>/etc/init.d/ufmha model_restart</code> <p>Keywords: Default session; REST API; missing ports</p>
2714738	<p>Description: Intentional stop for master container and start it again or reboot of master server will damage the HA failover option</p> <p>Workaround: manually Restart UFM cluster</p> <p>Keywords: UFM Container; Reboot, Failover</p>
2872513	<p>Description: after rebooting master container, Failover will be triggered twice (once to the standby and then back again to the master container)</p> <p>Workaround: N/A</p> <p>Keywords: UFM Container, reboot, failover</p>
2863388	<p>Description: Fail to get cables info for NDR Splitted Port.</p> <p>Workaround: N/A</p> <p>Keywords: Cable, NDR, Split</p>
N/A	<p>Description: In case of using SM mkey per port, several UFM operations might fail (get cable info, get system dump, switch FW upgrade)</p> <p>Workaround: N/A</p> <p>Keywords: SM, mkey per port,</p>
-	<p>Description: The UFM which is configured to work with telemetry for collecting historical data, is limited to work only with the configured HCA port - if this port is part of the bond interface and failure occurs, all telemetry data via this port will be stopped.</p> <p>Workaround: If a historical telemetry port is apart of the bond and a failure occurs, user should reconfigure the telemetry with a new active port and restart it within UFM.</p> <p>Keywords: telemetry, history, bond, failure</p> <p>Discovered in release: 6.7</p>
2459320	<p>Description: Docker upgrade to UFM6.6.1 from UFM6.6.0 is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Docker; upgrade</p> <p>Discovered in release: 6.6.1</p>
-	<p>Description: SHARP Aggregation Manager over UCX is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: UCX; SHARP AM</p> <p>Discovered in release: 6.6.1</p>
2288038	<p>Description: When the user try to collect system dump for UFM Appliance host, the job will be completed with an error with the following summary: "Running as a none root user Please switch to root user (super user) and run again."</p> <p>Workaround: N/A</p>

Ref #	Issue
	Keywords: System dump, UFM Appliance host
	Discovered in release: 6.5.2
2100564	Description: For modular dual-management switch systems, switch information is not presented correctly if the primary management module fails and the secondary takes over.
	Workaround: To avoid corrupted switch information, it is recommended to manually set the virtual IP address (box IP address) for the switch as the managed switch IP address (manual IP address) within UFM.
	Keywords: Modular switch, dual-management, virtual IP, box IP
	Discovered in release: 6.4.1
2135272	Description: UFM does not support hosts equipped with multiple HCAs of different types (e.g. a host with ConnectX®-3 and ConnectX-4/5/6) if multi-NIC grouping is enabled (i.e. <code>multinic_host_enabled = true</code>).
	Workaround: All managed hosts must contain HCAs of the same type (either using ConnectX-3 HCAs or use ConnectX-4/5/6 HCAs).
	Keywords: Multiple HCAs
	Discovered in release: 6.4.1
2063266	Description: Firmware upgrade for managed hosts with multiple HCAs is not supported. That is, it is not possible to perform FW upgrade for a specific host HCA.
	Workaround: Running software (MLNX_OFED) upgrade on that host will automatically upgrade all the HCAs on this host with the firmware bundled as part of this software package.
	Keywords: FW upgrade, multiple HCAs
	Discovered in release: 6.4.1
-	Description: Management PKey configuration (e.g. MTU, SL) can be performed only using PKey management interface (via GUI or REST API).
	Workaround: N/A
	Keywords: PKey, Management PKey, REST API
	Discovered in release: 6.4
2092885	Description: UFM Agent is not supported for SLES15 and RHEL8/CentOS8.
	Workaround: N/A
	Keywords: UFM Agent
	Discovered in release: 6.4
-	Description: CentOS 8.0 does not support IPv6.
	Workaround: N/A
	Keywords: IPv6
	Discovered in release: 6.4
1895385	Description: QoS parameters (<code>mtu</code> , <code>sl</code> and <code>rate_limit</code>) change does not take effect unless OpenSM is restarted.
	Workaround: N/A
	Keywords: QoS, PKey, OpenSM
	Discovered in release: 6.3

Ref #	Issue
-	<p data-bbox="304 259 1388 293">Description: Logical Server Auditing feature is supported on RedHat 7.x operating systems only.</p> <p data-bbox="304 304 496 338">Workaround: N/A</p> <p data-bbox="304 349 719 383">Keywords: Logical Server, auditing, OS</p> <p data-bbox="304 394 587 427">Discovered in release: 5.9</p>
-	<p data-bbox="304 445 1066 479">Description: Configuration from lossy to lossless requires device reset.</p> <p data-bbox="304 490 1249 524">Workaround: Reboot all relevant devices after changing behavior from lossy to lossless.</p> <p data-bbox="304 535 632 568">Keywords: Lossy configuration</p>

3 Overview

3.1 Scale-Out Your Fabric with Unified Fabric Manager

NVIDIA's Unified Fabric Manager (UFM®) is a powerful platform for managing scale-out computing environments. UFM enables data center operators to efficiently monitor and operate the entire fabric, boost application performance and maximize fabric resource utilization.

While other tools are device-oriented and involve manual processes, UFM's automated and application-centric approach bridges the gap between servers, applications and fabric elements, thus enabling administrators to manage and optimize from the smallest to the largest and most performance-demanding clusters.

3.2 UFM Benefits

Benefit	Description
Central Console for Fabric Management	UFM provides all fabric management functions in one central console. The ability to monitor, troubleshoot, configure and optimize all fabric aspects is available via one interface. UFM's central dashboard provides a one-view fabric-wide status view.
In-Depth Fabric Visibility and Control	UFM includes an advanced granular monitoring engine that provides real-time access to switch and host data, enabling cluster-wide monitoring of fabric health and performance, real-time identification of fabric-related errors and failures, quick problem resolution via granular threshold-based alerts and a fabric utilization dashboard.
Advanced Traffic Analysis	Fabric congestion is difficult to detect when using traditional management tools, resulting in unnoticed congestion and fabric under-utilization. UFM's unique traffic map quickly identifies traffic trends, traffic bottlenecks, and congestion events spreading over the fabric, which enables the administrator to identify and resolve problems promptly and accurately.
Enables Multiple Isolated Application Environments on a Shared Fabric	Consolidating multiple clusters into a single environment with multi-tenant data centers and heterogeneous application landscapes requires specific policies for the different parts of the fabric. UFM enables segmentation of the fabric into isolated partitions, increasing traffic security and application performance.
Service-Oriented Automatic Resource Provisioning	UFM uses a logical fabric model to manage the fabric as a set of business-related entities, such as time critical applications or services. The logical fabric model enables fabric monitoring and performance optimization on the application level rather than just at the individual port or device level. Managing the fabric using the logical fabric model provides improved visibility into fabric performance and potential bottlenecks, improved performance due to application-centric optimizations, quicker troubleshooting and higher fabric utilization.

Benefit	Description
Quick Resolution of Fabric Problems	UFM provides comprehensive information from switches and hosts, showing errors and traffic issues such as congestion. The information is presented in a concise manner over a unified dashboard and configurable monitoring sessions. The monitored data can be correlated per job and customer, and threshold-based alarms can be set.
Seamless Failover Handling	Failovers are handled seamlessly and are transparent to both the user and the applications running on the fabric, significantly lowering downtime. The seamless failover makes UFM in conjunction with other Mellanox products, a robust, production-ready solution for the most demanding data center environments.
Open Architecture	UFM provides an advanced Web Service interface and CLI that integrate with external management tools. The combination enables data center administrators to consolidate management dashboards while flawlessly sharing information among the various management applications, synchronizing overall resource scheduling, and simplifying provisioning and administration.

3.3 Main Functionality Modules

3.3.1 Fabric Dashboard

UFM's central dashboard provides a one-view fabric-wide status view. The dashboard shows fabric utilization status, performance metrics, fabric-wide events, and fabric health alerts.

The dashboard enables you to efficiently monitor the fabric from a single screen and serves as a starting point for event or metric exploration.

3.3.2 Fabric Segmentation (PKey Management)

In the PKey Management view you can define and configure the segmentation of the fabric by associating ports to specific defined PKeys. You can add, remove, or update the association of ports to the related PKeys and update the qos_parameters for pkey (mtu, rate, service_level).

3.3.3 Fabric Discovery and Physical View

UFM discovers the devices on the fabric and populates the views with the discovered entities. In the physical view of the fabric, you can view the physical fabric topology, model the data center floor, and manage all the physical-oriented events.

3.3.4 Central Device Management

UFM provides the ability to centrally access switches and hosts, and perform maintenance tasks such as firmware and software upgrade, shutdown and restart.

3.3.5 Monitoring

UFM includes an advanced granular monitoring engine that provides real time access to switch and server data. Fabric and device health, traffic information and fabric utilization are collected, aggregated and turned into meaningful information.

3.3.6 Configuration

In-depth fabric configuration can be performed from the Settings view, such as routing algorithm selection and access credentials.

The Event Policy Table, one of the major components of the Configuration view, enables you to define threshold-based alerts on a variety of counters and fabric events. The fabric administrator or recipient of the alerts can quickly identify potential errors and failures, and actively act to solve them.

3.3.7 Fabric Health

The fabric health tab contains valuable functions for fabric bring-up and on-going fabric operations. It includes one-click fabric health status reporting, UFM Server reporting, database and logs' snapshots and more.

3.3.8 Logging

The Logging view enables you to view detailed logs and alarms that are filtered and sorted by category, providing visibility into traffic and device events as well as into UFM server activity history.

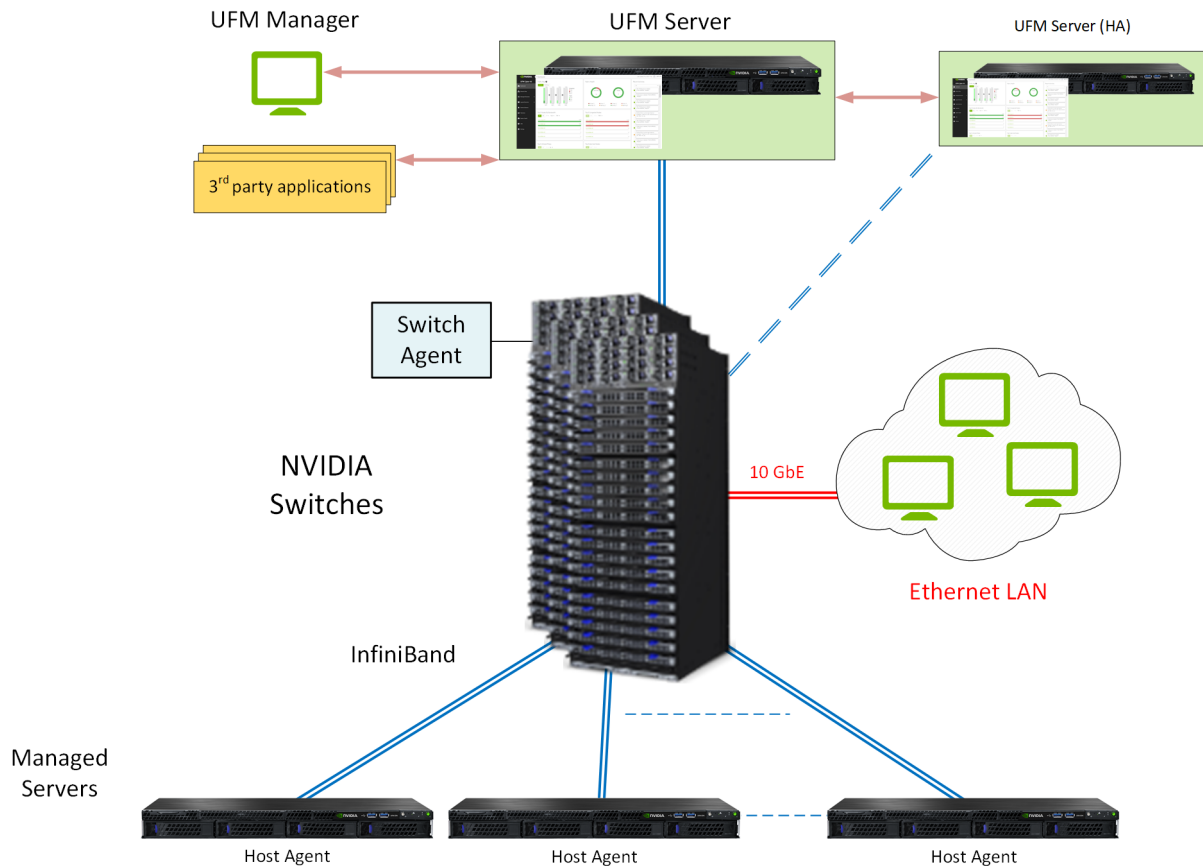
3.3.9 High Availability

In the event of a failover, when the primary (active) UFM server goes down or is disconnected from the fabric, UFM's High Availability (HA) capability allows for a secondary (standby) UFM server to immediately and seamlessly take over fabric management tasks. Failovers are handled seamlessly and are transparent to both the user and the applications running in the fabric. UFM's High Availability capability, when combined with Mellanox's High Availability switching solutions allows for non-disruptive operation of complex and demanding data center environments.

3.4 Fabric Topology with UFM

NVIDIA® UFM® is a host-based solution, providing all management functionality required for managing fabrics.

Fabric Topology with UFM



UFM Server is a server on which UFM is installed and has complete visibility over the fabric to manage routing on all devices.

UFM HA Server is a UFM installed server on a secondary server for High Availability deployment.

Managed Switching Devices are fabric switches, gateways, and routers managed by UFM.

Managed Servers are the compute nodes in the fabric on which the various applications are running, and UFM manages all servers connected to the fabric.

UFM Host Agent is an optional component that can be installed on the Managed Servers. UFM Host Agent provides local host data and host device management functionality.

The UFM Host Agent provides the following functionality:

- Discovery of IP address, CPU, and memory parameters on host
- Collection of CPU/Memory/Disk performance statistics on host
- Upgrading HCA Firmware and OFED remotely
- Creating an IP interface on top of the InfiniBand partition

UFM Switch Agent is an embedded component in NVIDIA switches that allows IP address discovery on the switch and allows UFM to communicate with the switch. For more information, please refer to [Device Management Feature Support](#).


3.5 UFM Communication Requirements

This chapter describes how the UFM server communicates with InfiniBand fabric components.

3.5.1 UFM Server Communication with Clients

The UFM Server communicates with clients over IP. The UFM Server can belong to a separate IP network, which can also be behind the firewall.

UFM Server Communication with Clients



3.5.1.1 UFM Server Communication with UFM Web UI Client

Communication between the UFM Server and the UFM web UI client is HTTP(s) based. The only requirement is that TCP port 80 (443) must not be blocked.

3.5.1.2 UFM Server Communication with SNMP Trap Managers

The UFM Server can send SNMP traps to configured SNMP Trap Manager(s). By default, the traps are sent to the standard UDP port 162. However, the user can configure the destination port. If the specified port is blocked, UFM Server traps will not reach their destination.


3.5.1.3 Summary of UFM Server Communication with Clients

Affected Service	Network	Address / Service / Port	Direction
Web UI Client	Out-of-band management*	HTTP / 80 HTTPS / 443	Bi-directional
SNMP Trap Notification	Out-of-band management*	UDP / 162 (configurable)	UFM Server to SNMP Manager

*If the client machine is connected to the IB fabric, IPoIB can also be used.

3.5.2 UFM Server Communication with InfiniBand Switches

UFM Server Communication with InfiniBand Switches



3.5.2.1 UFM Server InfiniBand Communication with Switch

The UFM Server must be connected directly to the InfiniBand fabric (via an InfiniBand switch). The UFM Server sends the standard InfiniBand Management Datagrams (MAD) to the switch and receives InfiniBand traps in response.

3.5.2.2 UFM Server Communication with Switch Management Software (Optional)

The UFM Server auto-negotiates with the switch management software on Mellanox Grid Director switches. The communication is bound to the switch Ethernet management port.

The UFM Server sends a multicast notification to MCast address 224.0.23.172, port 6306 (configurable). The switch management replies to UFM (via port 6306) with a unicast message that contains the switch GUID and IP address. After auto-negotiation, the UFM Server and switch management use XML-based messaging.

The following Device Management tasks are dependent on successful communication as described above:

- Switch IP discovery
- FRU Discovery (PSU, FAN, status, temperature)
- Software and firmware upgrades

The UFM Server manages IB Switch Devices over SNMP (default port 161 - configurable) and / or SSH (default port 22 - configurable).

3.5.2.3 UFM Server Communication with Externally Managed Switches (Optional)

UFM server uses Ibdagnet tool to discover chassis information (PSU, FAN, status, temperature) of the externally managed switches.


By monitoring chassis information data, UFM can trigger selected events when module failure occurs or a specific sensor value is above threshold.

3.5.2.4 Summary of UFM Server Communication with InfiniBand Switches

Affected Service	Network	Address / Service / Port	Direction
InfiniBand Management / Monitoring	InfiniBand	Management Datagrams	Bi-directional
Switch IP Address Discovery (auto-negotiation with switch management software)	Out-of-band management	Multicast 224.0.23.172, TCP / 6306 (configurable)	Multicast: UFM Server to switch TCP: Bi-directional
Switch Chassis Management / Monitoring	Out-of-band management	TCP / UDP / 6306 (configurable) SNMP / 161 (configurable) SSH / 22 (configurable)	Bi-directional

3.5.3 UFM Server Communication with InfiniBand Hosts

UFM Server Communication with InfiniBand Hosts



3.5.3.1 UFM Server InfiniBand Communication with HCAs


The UFM Server must be connected directly to the InfiniBand fabric. The UFM Server sends the standard InfiniBand Management Datagrams (MADs) to the Host Card Adapters (HCAs) and receives InfiniBand traps.

3.5.3.2 UFM Server Communication with Host Management (Optional)

The UFM Server auto-negotiates with the UFM Agent on a Host. The UFM Host Agent can be bound to the management Ethernet port or to an IPoIB interface (configurable). The UFM Server sends a multicast notification to MCast address 224.0.23.172, port 6306 (configurable). The UFM Agent replies to UFM (port 6306) with a unicast message that contains the host GUID and IP address. After auto-negotiation, the UFM Server and UFM Agent use XML-based messaging.

The following Device Management tasks are dependent on successful communication as described above:

- Host IP discovery
- Host resource discovery and monitoring: CPU, memory, disk
- Software and firmware upgrades

 UFM 3.6 supports in-band HCA FW upgrade. This requires enabling FW version and PSID discovery over vendor-specific MADs. For more information, see the UFM User Manual.


The UFM Server connects to the hosts over SSH (default port 22 - configurable) with root credentials, which are located in the UFM Server database.

3.5.3.3 Summary of UFM Server Communication with InfiniBand Hosts

Affected Service	Network	Address / Service / Port	Direction
InfiniBand Management / Monitoring	InfiniBand	Management Datagrams	Bi-directional
Host IP Address Discovery (auto-negotiation with UFM Host Agent)	Out-of-band management or IPoIB	Multicast 224.0.23.172, TCP / 6306 (configurable)	Multicast: UFM Server to UFM Agent TCP: Bi-directional
Host OS Management / Monitoring	Out-of-band management or IPoIB	TCP / UDP / 6306 (configurable) SSH / 22 (configurable)	Bi-directional

3.5.4 UFM Server High Availability (HA) Active—Standby Communication

UFM Server HA Active—Standby Communication



3.5.4.1 UFM Server HA Active–Standby Communication

UFM Active–Standby communication enables two services: heartbeat and DRBD.

- *heartbeat* is used for auto-negotiation and keep-alive messaging between active and standby servers. *heartbeat* uses port 694 (udp).
- DRBD is used for low-level data (disk) synchronization between active and standby servers. DRBD uses port 8888 (tcp).


Affected Service	Network	Address / Service / Port	Direction
UFM HA heartbeat	Out-of-band management*	UDP / 694	Bi-directional
UFM HA DRBD	Out-of-band management*	TCP / 8888	Bi-directional

*An IPoIB network can be used for HA, but this is not recommended, since any InfiniBand failure might cause split brain and lack of synchronization between the active and standby servers.

3.6 UFM Software Architecture

The following figure shows the UFM high-level software architecture with the main software components and protocols. Only the main logical functional blocks are displayed and do not necessarily correspond to system processes and threads.

UFM High-Level Software Architecture



3.6.1 Graphical User Interface

UFM User Interface is a web application based on JavaScript and Angular JS, which is supported by any Web Browser. The Web application uses a standard REST API provided by the UFM server.

3.6.2 Client Tier API

Third-party clients are managed by the REST API.

3.6.3 Client Tier SDK Tools

Support for UFM’s API and a set of tools that enhance UFM functionality and interoperability with third-party applications are provided as part of UFM.

3.6.4 UFM Server

UFM server is a central data repository and management server that manages all physical and logical data. UFM-SDN Appliance receives all data from the Device and Network tiers and invokes Device and Network tier components for management and configuration tasks. UFM-SDN Appliance uses a database for data persistency. The UFM-SDN Appliance is built on the Python twisted framework.

3.6.5 Subnet Manager

Subnet Manager (SM) is the InfiniBand “Routing Engine”, a key component used for fabric bring-up and routing management.

UFM uses the Open Fabric community OpenSM Subnet Manager. UFM uses a plug-in API for runtime management and fabric data export.

3.6.6 NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ Aggregation Manager

NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) is a technology that improves the performance of mathematical and machine learning applications by offloading collective operations from the CPU to the switch network.

Aggregation Manager (AM) is a key component of NVIDIA SHARP software, used for NVIDIA SHARP resources management.

For further information about NVIDIA SHARP AM, refer to [Appendix - NVIDIA SHARP Integration](#).

3.6.7 Performance Manager

The UFM Performance Manager component collects performance data from the managed fabric devices and sends the data to the UFM-SDN Appliance for fabric-wide analysis and display of the data.

3.6.8 Device Manager

The Device Manager implements the set of common device management tasks on various devices with varying management interfaces. The Device Manager uses SSH protocol and operates native device CLI (command-line interface) commands.

3.6.9 UFM Switch Agent

UFM Switch Agent is an integrated part of NVIDIA switch software. The agent supports system parameter discovery and device management functionality on switches.

3.6.10 Communication Protocols

UFM uses the following communication protocols:

- Web UI communicates with the UFM server utilizing Web Services carried on REST API.
- The UFM server communicates with the switch Agent located on managed switches by proprietary TCP/UDP-based discovery and monitoring protocol and SSH.
- Monitoring data is sent by the switch Agent to UFM server Listener by a proprietary TCP-based protocol.

3.7 Getting Familiar with UFM's Data Model

3.7.1 Overview of Data Model

UFM enables the fabric administrator to manage the fabric based on business-oriented requirements, as opposed to managing the fabric based on device-oriented and port-oriented management systems.

The business-centric capability is provided by UFM's data model which treats the physical fabric topology as an abstraction. You can define groups of fabric resources. For example, servers represent a certain application, a job running on the fabric, or a reserved computing resource pool for specific customers of a multi-tenant fabric.

All UFM functionality is associated with the data model. For example, monitored data and fabric events are correlated to the logical groups, fabric and host configuration is performed according to the model, and performance optimization is derived from the data model.





3.7.1.1 UFM Model Basics

The fabric managed by UFM consists of a set of physical and logical objects including their connections. The Object Model has a hierarchical object-oriented tree structure with objects as the tree elements. Each object defines an abstraction for physical or logical fabric elements.

3.7.1.2 Physical Model

The Physical Model represents the physical resources and connectivity topology of the Network. UFM enables discovery, monitoring and configuration of the managed physical objects.

Physical Objects

Icon	Name	Description
N/A	Port Object	Represents the external physical port on switch or on Host Channel Adapter (HCA). A port is identified by its number. UFM provides InfiniBand standard management and monitoring capabilities on the port level.
N/A	Module Object	Represents the Field Removable Unit, Line card, and Network card on switch or HCA on host. For Mellanox Switches, Line and Network Cards are modeled as modules.
	Link Object	Represents the physical connection between two active ports.
	Computer Object	Represents the computer (host) connected to the Fabric. The UFM Agent installed on the host provides extended monitoring and management capabilities. Hosts without agents are limited to InfiniBand standard management and monitoring capabilities.
	Switch Object	Represents the switch chassis in the Fabric. A Switch object is created for every Mellanox Switch. Switches of other vendors are represented as InfiniBand Switches and limited by InfiniBand standard management and monitoring capabilities.
	Rack Object	Represents the arbitrary group of switches or computers. When linked devices are shown as a group, the link is shown between the group and the peer object.

4 UFM Regular Installation

5 UFM Installation and Initial Configuration

UFM® software includes Server and Agent components. UFM Server software should be installed on a central management node. For optimal performance, and to minimize interference with other applications, it is recommended to use a dedicated server for UFM. The UFM Agent is an optional component and should be installed on fabric nodes. The UFM Agent should not be installed on the Management server.

The following sections provide step-by-step instructions for installing and activating the license file, installing the UFM server software, and installing the UFM Agent.

- [UFM Installation Steps](#)
- [Running UFM Server Software](#)
- [Upgrading UFM Software](#)
- [Uninstalling UFM](#)
- [Appendix - UFM Migration](#)
- [Docker Installation](#)
- [Historical Telemetry Collection in UFM](#)

5.1 UFM Installation Steps

- [Downloading UFM Software and License File](#)
- [Installing UFM Server Software](#)

5.1.1 Downloading UFM Software and License File

Before you obtain a license for the UFM® software, prepare a list of servers with the MAC address of each server on which you plan to install the UFM software. These MAC addresses are requested during the licensing procedure.

5.1.1.1 Obtaining License

UFM is licensed per managed device according to the UFM license agreement.

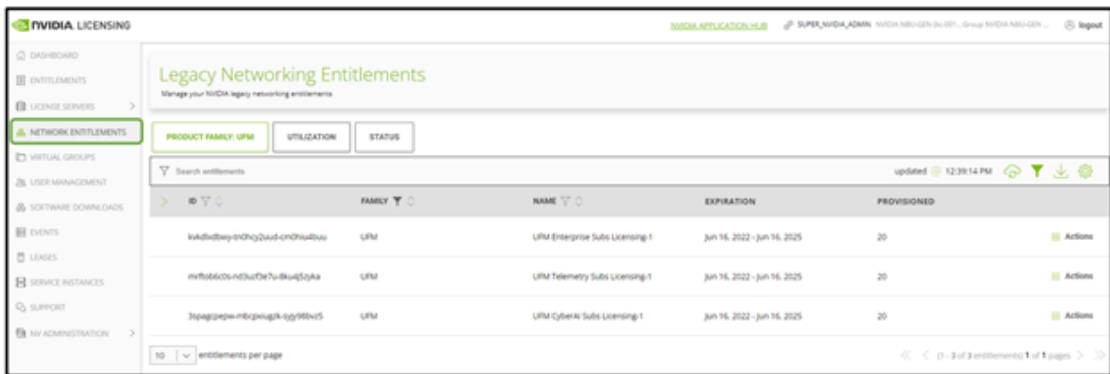
When you purchase UFM, you will receive an email with instructions on obtaining your product license. A valid UFM license is a prerequisite for the installation and operation of UFM.

UFM licenses are per managed node and are aggregative. If you install an additional license, the system adds the previous node number and the new node number and manages the sum of the nodes. For example, if you install a license for 10 managed nodes and an additional license for 15 nodes, UFM will be licensed for up to 25 managed nodes.

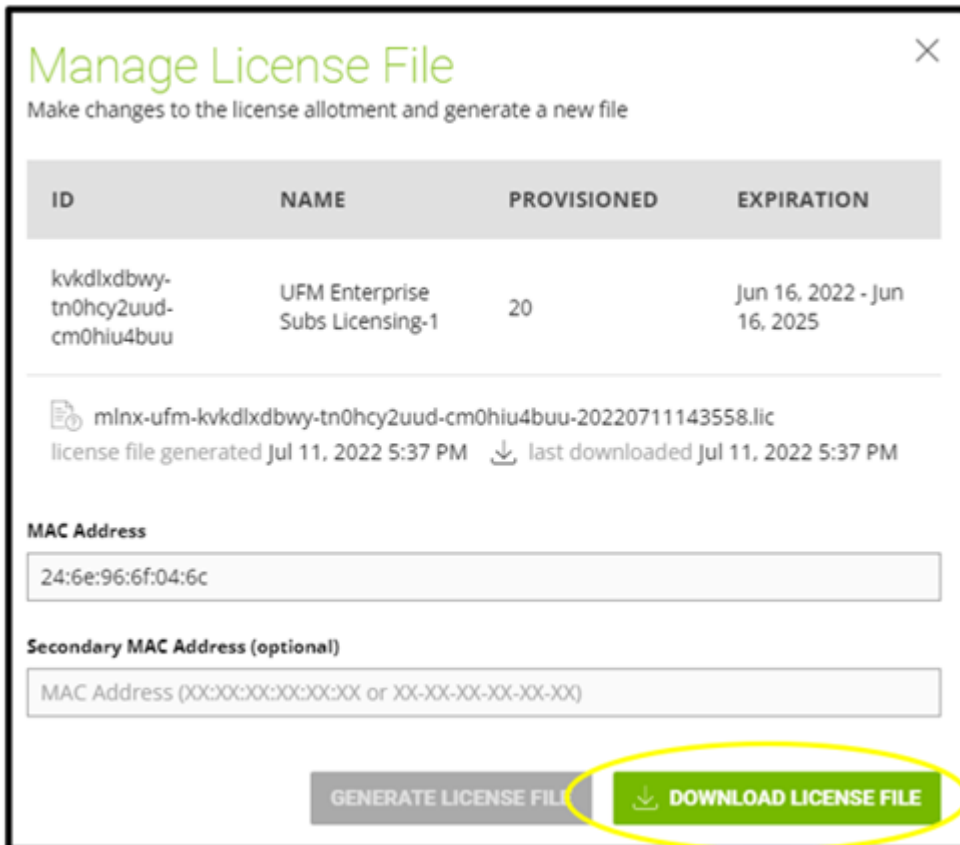
To obtain the license:

1. Go to NVIDIA's [Licensing and Download Portal](#) and log in as specified in the licensing email you received.
 - If you did not receive your NVIDIA Licensing and Download Portal login information, contact your product reseller.
2. If you purchased UFM directly from NVIDIA and you did not receive the login information, contact enterprisesupport@nvidia.com. Click on the Network Entitlements tab. You'll see a list with the serial licenses of all your software products and software product license

information and status.



3. Select the license you want to activate and click on the “Actions” button.
4. In the MAC Address field, enter the MAC address of the delegated license-registered host. If applicable, in the HA MAC Address field, enter your High Availability (HA) server MAC address. If you have more than one NIC installed on a UFM Server, use any of the MAC addresses.



5. Click on Generate License File to create the license key file for the software.
6. Click on Download License File and save it on your local computer.

If you replace your NIC or UFM server, repeat the process of generating the license to set new MAC addresses. You can only regenerate a license two times. To regenerate the license after that, contact NVIDIA Sales Administration at enterprisesupport@nvidia.com.

5.1.1.2 Downloading UFM Software

⚠ Due to internal packaging incompatibility, this release has two different packages for each of the supported distributions:

- One for UFM deployments over MLNX_OFED 5.X (or newer)


Please make sure to use the UFM installation package compatible to your setup.

This software download process applies to software updates and first-time installation.

If you own the UFM Media Kit and this is your first-time installation, skip this section.

To download the UFM software:

1. Click on Software Downloads, filter the product family to UFM, and select the relevant version of the software. Click on Download.



2. Save the file on your local drive.
3. Click Close.

5.1.2 Installing UFM Server Software

The default UFM installation directory is `/opt/ufm`.

UFM Server installation options are:

- Standalone
- High Availability (HA) - Delivered in a separate package as of UFM v6.10.0.
- Docker Container

The following processes might be interrupted during the installation process:

- httpd (apachi2 in Ubuntu)
- dhcpd

After installation:

1. Activate the software license.
2. [Perform initial configuration.](#)



Before you run UFM, ensure that all ports used by the UFM server for internal and external communication are open and available. For the list of ports, see [Used Ports](#).

5.1.2.1 Installed Packages

A of UFM Enterprise v6.11.0, installation is based on Conda-4.12 (or newer) for Python3.9 environment and third-party packages deployments. The below-listed packages can be used for all supported operating systems.

Conda binaries	Conda Python Environment
_libgcc_mutex=0.1=main	appdirs==1.4.4
_openmp_mutex=5.1=1_gnu	apscheduler==3.9.1
c-ares=1.18.1=h7f8727e_0	asgiref==3.5.2
ca-certificates=2022.07.19=h06a4308_0	asn1crypto==1.5.1
curl=7.84.0=h5eee18b_0	attrs==21.4.0
krb5=1.19.2=hac12032_0	automat==20.2.0
ld_impl_linux-64=2.38=h1181459_1	bcrypt==3.2.2
libcurl=7.84.0=h91b91d3_0	cached-property==1.5.2
libedit=3.1.20210910=h7f8727e_0	cachetools==5.1.0
libev=4.33=h7f8727e_1	cairocffi==1.0.0
libffi=3.3=he6710b0_2	cairosvg==2.5.2
libgcc-ng=11.2.0=h1234567_1	carbon==1.1.10
libgomp=11.2.0=h1234567_1	certifi==2022.5.18
libnghttp2=1.46.0=hce63b2e_0	cffistring==1.15.0
libssh2=1.10.0=h8f2d780_0	chardet==4.0.0
libstdcxx-ng=11.2.0=h1234567_1	charset-normalizer==2.0.12
ncurses=6.3=h5eee18b_3	click==8.1.3
openssl=1.1.1q=h7f8727e_0	constantly==15.1.0
pip=22.1.2=py39h06a4308_0	cryptography==37.0.2
python=3.9.12=h12debd9_1	cssselect==1.1.0
readline=8.1.2=h7f8727e_1	cssselect2==0.6.0
sqlite=3.39.2=h5082296_0	daemonize==2.5.0
tk=8.6.12=h1ccaba5_0	defusedxml==0.7.1
wheel=0.37.1=pyhd3eb1b0_0	distro==1.7.0
xz=5.2.5=h7f8727e_1	djano==3.0.14
zlib=1.2.12=h7f8727e_2	djano-piston3==0.3rc2
	djano-tagging==0.4.3

Conda binaries	Conda Python Environemnt
	docker==5.0.3
	ecdsa==0.17.0
	flask==1.1.1
	graphite-web==1.1.10
	hyperlink==21.0.0
	idna==3.3
	importlib-metadata==4.11.3
	incremental==21.3.0
	inotify==0.2.10
	ipaddress==1.0.23
	ipy==1.1
	isodate==0.6.1
	itsdangerous==1.1.0
	jinja2==2.10.3
	jsonschema==4.5.1
	lxml==4.8.0
	markupsafe==1.1.1
	more-itertools==8.13.0
	mysqlclient==2.1.0
	netaddr==0.8.0
	netifaces==0.11.0
	nose==1.3.7
	ntlm-auth==1.5.0
	numpy==1.22.4
	paramiko==2.11.0
	pbr==5.9.0
	pillow==9.1.1
	platformdirs==2.5.2
	ply==3.11
	psutil==5.9.0
	pyasn1==0.4.8
	pyasn1-modules==0.2.8
	pycairo==1.21.0
	pycparser==2.21
	pycrypto==2.6.1
	pycryptodomex==3.14.1

Conda binaries	Conda Python Environemnt
	pydes==2.0.1
	pydo==2.0.5
	pygal==3.0.0
	pyhamcrest==2.0.3
	pyinotify==0.9.6
	pynacl==1.5.0
	pyopenssl==22.0.0
	pyparsing==3.0.9
	pyrsistent==0.18.1
	pyserial==3.5
	pysmi==0.3.4
	pysnmp==4.4.12
	python-dateutil==2.8.2
	python-hostlist==1.21
	python-magic==0.4.27
	python-mimeparse==1.6.0
	pytz==2022.1
	pytz-deprecation-shim==0.1.0.post0
	PyYAML==6.0
	requests==2.27.1
	requests-file==1.5.1
	requests-ntlm==1.1.0
	requests-toolbelt==0.9.1
	service-identity==21.1.0
	setproctitle==1.1.10
	setuptools==62.3.2
	six==1.16.0
	soappy-py3==0.52.27
	south==0.8.4
	sqlparse==0.4.2
	stdeb==0.10.0
	subprocess32==3.5.4
	tinycss==0.4
	tinycss2==1.1.1
	twisted==22.4.0
	txamqp==0.8.2

Conda binaries	Conda Python Environment
	typing-extensions==4.2.0
	tzdata==2022.1
	tzlocal==4.2
	ujson==5.3.0
	urllib3==1.26.9
	webencodings==0.5.1
	websocket-client==1.3.2
	werkzeug==0.16.0
	wheel==0.37.1
	whisper==1.1.8
	wstools==0.4.8
	wstools-py3==0.54.4
	zeep==4.1.0
	zipp==3.8.0
	zope-interface==5.4.0
	aiohttp==3.8.1
	aiosignal==1.2.0
	async_timeout==4.0.2
	asynctest==0.13.0
	frozenset==1.2.0
	idna_ssl==1.1.0
	multidict==5.2.0
	yaml==1.7.2

5.1.2.2 Installing UFM Server Software as Standalone

For instructions on installing the UFM server software as a standalone for InfiniBand, please refer to [Quick Start Guide](#).

5.1.2.3 Installing UFM Server Software for High Availability

If high availability (HA) is required, install UFM for HA on a server that is designated to be the master. For instructions, please refer to [Quick Start Guide](#).

For more information, see [High Availability](#).

5.1.2.4 Installing UFM Server Software for Docker

Please refer to [Docker Installation](#).

5.1.2.5 Activating Software License

For instructions on how to activate the software license, please refer to the [UFM Quick Start Guide](#).


5.2 Running UFM Server Software

➤ *Before you run UFM, do the following:*

- Perform initial configuration.
- Ensure that all ports used by the UFM server for internal and external communication are open and available. For the list of ports, see [Used Ports](#).

You can run the UFM server software in the following modes:

- Management
- Monitoring
- High Availability
- High Availability with failover to an external SM

 In Management or High Availability mode, ensure that all Subnet Managers in the fabric are disabled *before* running UFM. Any remaining active Subnet Managers will prevent UFM from running.

5.2.1 Running UFM Server Software in Management Mode

After installing, run the UFM Server by invoking:

```
systemctl start ufm-enterprise.service
```

 `/etc/init.d/ufmd` - Available for backward compatibility.

Log files are located under `/opt/ufm/files/log` (the links to log files are in `/opt/ufm/log`).

5.2.2 Running UFM Software in High Availability Mode

On the Master server, run the UFM Server by invoking:

```
ufm_ha_cluster start
```

You can specify additional command options for the `ufmha` service.

ufm_ha_cluster Command Options

Command	Description
start	Starts UFM HA cluster.
stop	Stops UFM HA cluster.

Command	Description
failover	Initiates failover (change mastership from local server to remote server).
takeover	Initiates takeover (change mastership from remote server to local server).
status	Shows current HA cluster status.
cleanup	Cleans the HA configurations on this node.
help	Displays help text.

5.2.3 Running UFM Software in Monitoring Mode

Run UFM in Monitoring mode while running concurrent instances of Subnet Manager on NVIDIA switches. Monitoring and event management capabilities are enabled in this mode. UFM non-monitoring features such as provisioning and performance optimization are disabled in this mode.


The following table describes whether features are enabled or disabled in Monitoring mode.


Features Enabled/Disabled in Monitoring Mode

Feature	Enabled/Disabled in Monitoring Mode
Fabric Discovery	Enabled
Topology Map	Enabled
Fabric Dashboard	Enabled
Fabric Monitoring	Enabled
Alerts and Thresholds (inc. SNMP traps)	Enabled
Fabric Logical Model	Enabled
Subnet Manager and plugins	Disabled
Subnet Manager Configuration	Disabled
Automatic Fabric Partitioning	Disabled
Central Device Management	Disabled
Quality of Service	Disabled
Failover (High Availability mode)	Disabled
Traffic Aware Routing Algorithm	Disabled
Device Management	Disabled
Integration with Schedulers	Disabled
Unhealthy Ports	Disabled

In Monitoring mode, UFM periodically discovers the fabric and updates the topology maps and database.

For Monitoring mode, connect UFM to the fabric using port ib0 only. The fabric must have a subnet manager (SM) running on it (on another UFM, HBSM, or switch SM).

 When UFM is running in Monitoring mode, the internal OpenSM is not sensitive to changes in OpenSM configuration (opensm.conf).

 When running in Monitoring mode, the following parameters are automatically overwritten in the /opt/ufm/files/conf/opensm/opensm_mon.conf file on startup:

- event_plugin_name osmufmpi
- event_plugin_options --vendinfo -m 0

Any other configuration is not valid for Monitoring mode.

 To run in Monitoring mode:

1. In the /opt/ufm/conf/gv.cfg configuration file:


- Set monitoring_mode to yes
- If required, change mon_mode_discovery_period (the default is 60 seconds)
- Set reset_mode to no_reset

We recommend this setting when running multiple instances of UFM so that each port counter is not reset by different UFM instances. For more information, see [Resetting Physical Port Counters](#).

2. Restart the UFM Server.

The Running mode is set to Monitoring, and the frequency of fabric discovery is updated according to the setting of mon_mode_discovery_period.

Note that a monitor icon will appear at the top of the navigation bar indicating that monitoring mode is enabled:



5.2.4 HTTP/HTTPS Port Configuration

After installation, you can configure the web server to communicate in the secure protocol HTTP/S. For further information, please refer to the [Launching a UFM Web UI Session](#) section.

Port 8088 is an internal port that is used by the UFM server (a port that is not exposed to the user by the Apache Web Server). Apache web server listens on port 80 and forwards the incoming traffic to the local port 8088. Port 8088 is configurable, port 80 is not.

To configure using HTTP/S protocol instead of the default HTTP, add the following to the configuration file at /opt/ufm/conf/gv.cfg::

```
# WebServices Protocol (http/https) and Port
ws_port = 8088
ws_protocol = https
```

UFM installation configures HTTPS protocol in the webserver as follows:

- Configures listening on port 443
- Configures default virtual host
- Creates/uses local certificates

5.2.5 Launching UFM Web UI Session

For instructions, please refer to the [UFM Quick Start Guide](#).

5.2.6 User Authentication

UFM User Authentication is based on standard Apache User Authentication. Each Web Service client application must authenticate against the UFM server to gain access to the system.

The UFM software comes with one predefined user:

- Username: admin
- Password: 123456

You can add, delete, or update users via [User Management Tab](#).

5.2.7 Licensing

UFM license is subscription-based featuring the following subscription options:

- 1-year subscription
- 3-year subscription
- 5-year subscription
- Evaluation 30-day trial license



UFM will continue to support old license types, but they are no longer available to obtain.

2 months before the expiration of your subscription license, UFM will warn you that your license will expire soon. After the subscription expires, UFM will continue to work with the expired license for two months beyond its expiration.

During this extra two-month period, UFM will generate a critical alarm indicating that the UFM license has expired and that you need to renew your subscription. Failing to do so within that 2-month period activates UFM Limited Mode. Limited mode blocks all REST APIs and access to the UFM web UI.

UFM enables functionality based on the license that was purchased and installed. This license determines the functionality and the maximum allowed number of nodes in the fabric.

To renew your UFM subscription, purchase a new license and install the new license file by downloading the license file to a temp directory on the UFM master server and then copying the license file to `/opt/ufm/files/licenses/` directory.



UFM may not detect new license files if downloaded directly to `/opt/ufm/files/licenses`. If UFM does not detect the new license file, a UFM restart may be required.

If several licenses are installed on the server (more than one license file exists under `/opt/ufm/files/licenses/`), UFM uses only the strongest license and takes into consideration the expiration date, and the managed device limits on it, regardless of any other licenses that may exist on the server.

For instructions on how to view your license, please refer to the [UFM Quick Start Guide](#).

5.2.8 Showing UFM Processes Status

A script under `/opt/ufm/scripts` calls `show_ufm_status.sh`, which allows the user to view the current status of UFM's main processes.

Running the command with the `-e` (`extended_processes`) option shows the main and sub-processes being handled by the UFM.

```
[root@r-ufm77 gvm_github]# /opt/ufm/scripts/show_ufm_status.sh
=====
                        UFM Main Processes
=====
ModelMain      Process is : [ Running ]
Opensm         Process is : [ Running ]
SHARP          Process is : [ Running ]
Unhealthy Ports Process is : [ Running ]
Daily Report   Process is : [ Running ]
UFM Health     Process is : [ Running ]
UFM Telemetry  Process is : [ Running ]
=====

[root@r-ufm77 gvm_github]# /opt/ufm/scripts/show_ufm_status.sh -e
=====
                        UFM Main Processes
=====
ModelMain      Process is : [ Running ]
Opensm         Process is : [ Running ]
SHARP          Process is : [ Running ]
Unhealthy Ports Process is : [ Running ]
Daily Report   Process is : [ Running ]
UFM Health     Process is : [ Running ]
UFM Telemetry  Process is : [ Running ]
=====
                        UFM ModelMain Child Processes
=====
SMClientConsumer Process is : [ Running ]
SMTrapHandler    Process is : [ Running ]
SysinfoJsonAgent Process is : [ Running ]
Telemetry Agent  Process is : [ Running ]
Telemetry History Process is : [ Running ]
=====
```

5.3 Upgrading UFM Software

After UFM® installation, UFM detects existing UFM versions previously installed on the machine and prompts you to run a clean install of the new version or to upgrade. We recommend backing up the UFM configuration before upgrading the UFM as specified in section UFM Database and Configuration File Backup.

For Standalone Server and High Availability upgrades, please refer to the [UFM Quick Start Guide](#).

5.4 Uninstalling UFM

UFM Server can be uninstalled by running an `uninstall` script as described in the [UFM Quick Start Guide](#).

5.5 Appendix - UFM Migration

5.5.1 Overview

UFM migration enables backup and restores UFM configuration files.

5.5.2 Backup UFM configuration

By default, the following folders (placed in `/opt/ufm/files`) are being backed up:

- conf
- dashboardViews
- licenses
- networkViews
- scripts
- sqlite
- templates/user-defined
- ufmhealth/scripts
- userdata
- users_preferences



The user may also backup the UFM historical telemetry data ("-t" argument).

5.5.2.1 UFM (Bare Metal)

```
/opt/ufm/scripts/ufm_backup.sh --help
usage: ufm_backup.py [-h] [-f BACKUP_FILE] [-t]
```

5.5.2.1.1 Optional Arguments

-h	--help	show this help message and exit
-f	--backup-file BACKUP_FILE	full path of zip file to be generated
-t	--telemetry	backup UFM historical telemetry

5.5.2.2 UFM Docker Container

1. Backup UFM configuration. Run:

```
docker exec ufm /opt/ufm/scripts/ufm_backup.sh
```

2. Copy the backup file from UFM docker container to the host. Run:

```
docker cp ufm:/root/<backup file> <path on host>
```

5.5.2.3 UFM Appliance

1. Backup UFM configuration. Run:

```
ufm data backup [with-telemetry]
```

2. Upload the backup file to a remote host. Run:

```
ufm data upload <backup file> <upload URL>
```



More details can be found in the log file `/tmp/ufm_backup.log`.

5.5.3 Restore UFM Configuration



All folders which are a part of the UFM backup are restored (filter is done during the backup stage).

5.5.3.1 UFM Bare Metal

```
/opt/ufm/scripts/ufm_restore.sh --help  
usage: ufm_restore.py [-h] -f BACKUP_FILE [-u] [-v]
```

5.5.3.1.1 Optional Arguments

-h	--help	show this help message and exit
-f BACKUP_FILE	--backup-file BACKUP_FILE	full path of zip file generated by backup script
-u	--upgrade	upgrades the restored UFM files
-v	--verbose	makes the operation more talkative

5.5.3.2 UFM Docker Container

1. Stop UFM. Run:

```
docker exec ufm /etc/init.d/ufmd stop
```

2. Copy the backup file from the host into UFM docker container. Run:

```
docker cp <backup file> ufm:/tmp/<backup file>
```

3. Restore UFM configuration. Run:

```
docker exec ufm /opt/ufm/scripts/ufm_restore.sh -f /tmp/<backup file> [--upgrade]
```

4. Start UFM. Run:

```
docker exec ufm /etc/init.d/ufmd start
```

5.5.3.3 UFM Appliance

1. Stop UFM. Run:

```
no ufm start
```

2. Copy the backup file from a remote host into UFM appliance. Run:

```
ufm data fetch <download URL>
```

3. Restore UFM configuration. Run:

```
ufm data restore <backup file>
```

4. Start UFM. Run:

```
ufm start
```



When restoring the UFM configuration from host to a container, the following parameters in `/opt/ufm/files/conf/gv.cfg` may be reset the following:

- fabric_interface
- ufma_interfaces
- mgmt_interface



UFM configuration upgrade during restore is not supported in UFM Appliance GEN2/GEN2.5

More details can be found in the log files `/tmp/ufm_restore.log` and `/tmp/ufm_restore_upgrade.log`

5.6 Docker Installation

5.6.1 General Prerequisites

- MLNX_OFED must be installed on the server that will run UFM Docker
- For UFM to work, you must have an InfiniBand port configured with an IP address and in "up" state.



For InfiniBand support, please refer to [NVIDIA Inbox Drivers](#), or MLNX_OFED guides.

- Make sure to stop the following services before running UFM Docker container, as it utilizes the same default ports that they do: Pacemaker, httpd, OpenSM, and Carbon.

- If firewall is running on the host, please make sure to add an allow rule for UFM used ports (listed below):

⚠ If the default ports used by UFM are changed in UFM configuration files, make sure to open the modified ports on the host firewall.

- 80 (TCP) and 443 (TCP) are used by WS clients (Apache Web Server)
- 8000 (UDP) is used by the UFM server to listen for REST API requests (redirected by Apache web server)
- 6306 (UDP) is used for multicast request communication with the latest UFM Agents
- 8005 (UDP) is used as a UFM monitoring listening port
- 8888 (TCP) is used by DRBD to communicate between the UFM Primary and Standby servers
- 2022 (TCP) is used for SSH

5.6.2 Prerequisites for Upgrading UFM Docker Container

- Supported versions for upgrade are UFM v.6.7.0 and above.
- UFM files directory from previous container version mounted on the host.

5.6.3 Step 1: Loading UFM Docker Image

To load the UFM docker image, pull the latest image from docker hub:

```
docker pull mellanox/ufm-enterprise:latest
```

⚠ You can see full usage screen for ufm-installation by running the container with `-h` or `-help` flag:

```
docker run --rm mellanox/ufm-enterprise-installer:latest -h
```


5.6.4 Step 2: Installing UFM Docker


5.6.4.1 Installation Command Usage

```
docker run -it --name=ufm_installer --rm \
-v /var/run/docker.sock:/var/run/docker.sock \
-v /etc/systemd/system:/etc/systemd_files/ \
-v [UFM_FILES_DIRECTORY]:/installation/ufm_files/ \
-v [LICENSE_DIRECTORY]:/installation/ufm_licenses/ \
mellanox/ufm-enterprise:latest \
--install [OPTIONS]
```


Modify the variables in the installation command as follows:

- `[UFM_FILES_DIRECTORY]` : A directory on the host to mount UFM configuration files.

 **UFM_FILES_DIRECTORY** must have read/write permissions for other users because UFM needs write data during runtime.

 Example: If you want UFM files on the host to be under `/opt/ufm/files/` you must set this volume to be: `-v /opt/ufm/files/:/installation/ufm_files/`

- `[UFM_LICENSES_DIR]` : UFM license file or files location.

 Example: If your license file or files are located under `/downloads/ufm_license_files/` then you must set this volume to be `-v /downloads/ufm_license_files/:/installation/ufm_licenses/`

- `[OPTIONS]` : UFM installation options. For more details see the table below.

5.6.4.1.1 Command Options

Flag	Description	Default Value
<code>-f --fabric-interface</code>	IB fabric interface name.	ib0
<code>-g --mgmt-interface</code>	Management interface name.	eth0
<code>-h --help</code>	Show help	N/A


5.6.5 Installation Modes

UFM Enterprise installer supports several deployment modes:

5.6.5.1 Stand Alone (SA) Installation

1. Create a directory on the host to mount and sync UFM Enterprise files with read/write permissions. For example: `/opt/ufm_files/`.
2. Copy only your UFM license file(s) to a temporary directory which we're going to use in the installation command. For example: `/tmp/license_file/`
3. Run the UFM installation command according to the following example which will also configure UFM fabric interface to be `ib1`:

```
docker run -it --name=ufm_installer --rm \  
-v /var/run/docker.sock:/var/run/docker.sock \  
-v /etc/systemd/system:/etc/systemd_files/ \  
-v /opt/ufm/files:/installation/ufm_files/ \  
-v /tmp/license_file:/installation/ufm_licenses/ \  
mellanox/ufm-enterprise:latest \  
--install \  
--fabric-interface ib1
```

 The values below can be updated in the command to your needs:

- /opt/ufm/files/
- /tmp/license_file/
- For example, if you want UFM files to be mounted in another location on your server, create that directory and replace the path in the command.

4. Reload system

```
systemctl daemon-reload
```

5. To Start UFM Enterprise service run:

```
systemctl start ufm-enterprise
```

5.6.5.2 High Availability

5.6.5.2.1 Pre-deployments requirements

- Install pacemaker, pcs, and drbd-utils on both servers
- A partition for DRBD on each server (with the same name on both servers) such as /dev/sdd1 . Recommended partition size is 10-20 GB, otherwise DRBD sync will take a long time to complete.
- CLI command hostname -i must return the IP address of the management interface used for pacemaker sync correctly (update /etc/hosts/ file with machine IP)
- Create the directory on each server under /opt/ufm/files/ with read/write permissions on each server. This directory will be used by UFM to mount UFM files, and it will be synced by DRBD.

5.6.5.2.2 Installing UFM Containers

On the main server, install UFM Enterprise container with the command below:

```
docker run -it --name=ufm_installer --rm \
-v /var/run/docker.sock:/var/run/docker.sock \
-v /etc/systemd/system:/etc/systemd_files/ \
-v /opt/ufm/files:/installation/ufm_files/ \
-v /tmp/license_file:/installation/ufm_licenses/ \
mellanox/ufm-enterprise:latest \
--install
```

On each the standby (secondary) server, install UFM Enterprise container like the following example with the command below:

```
docker run -it --name=ufm_installer --rm \
-v /var/run/docker.sock:/var/run/docker.sock \
-v /etc/systemd/system:/etc/systemd_files/ \
-v /opt/ufm/files:/installation/ufm_files/ \
mellanox/ufm-enterprise:latest \
--install
```

5.6.5.2.3 Downloading UFM HA Package

Download the UFM-HA package on both servers using the following command:

```
wget https://www.mellanox.com/downloads/UFM/ufm_ha_4.0.0-8.tgz
```

5.6.5.2.4 Installing UFM HA Package

1. [On Both Servers] Extract the downloaded UFM-HA package under /tmp/
2. [On Both Servers] Go to the extracted directory /tmp/ufm_ha_XXX and run the installation script:

```
./install.sh -l /opt/ufm/files/ -d /dev/sda5 -p enterprise
```

Option	Description
-l	Location For DRBD. Please always use /opt/ufm/files/
-d	Partition (disk) name for DRBD
-p	Product Name. For UFM Enterprise this must always be “enterprise”

5.6.5.2.5 Configuring UFM HA


There are two methods to configure the HA cluster:

- [Configure HA with SSH Trust](#) - Requires passwordless SSH connection between the servers.
- [Configure HA without SSH Trust](#) - Does not require passwordless SSH connection between the servers, but asks you to run configuration commands on both servers.

5.6.5.2.5.1 Configure HA with SSH Trust

1. [On the Master Server] Run the following command:

```
configure_ha_nodes.sh \  
--cluster-password 123456 \  
--main-hostname ufm-host01 \  
--main-ip 192.168.10.1 \  
--main-sync-interface enp2s0f0 \  
--standby-hostname ufm-host02 \  
--standby-ip 192.168.10.2 \  
--standby-sync-interface enp2s0f0 \  
--virtual-ip 192.168.10.5
```

 **configure_ha_nodes.sh** requires SSH connection to the standby server. If SSH trust is not configured, then you are prompted to enter the SSH password of the standby server during configuration runtime

Option	Description
<code>--cluster-password</code>	UFM HA cluster password for authentication by the pacemaker.
<code>--main-hostname</code>	Master (main) server hostname
<code>--main-ip</code>	Master (main) server IP address

Option	Description
<code>--main-sync-interface</code>	Port name (interface) on a master (main) server that will be used in DRBD sync
<code>--standby-hostname</code>	Standby server hostname
<code>--standby-ip</code>	Standby server IP address
<code>--standby-sync-interface</code>	Port name (interface) on standby server that will be used in DRBD sync
<code>--virtual-ip</code>	UFM HA cluster Virtual IP
<code>--no-vip</code>	Configure HA without virtual IP

2. Depending on the size of your partition, wait for the configuration process to complete and DRBD sync to finish.

5.6.5.2.5.2 Configure HA without SSH Trust

If you cannot establish an SSH trust between your HA servers, you can use `ufm_ha_cluster` directly to configure HA. You can see all the options for configuring HA in the Help menu:

```
ufm_ha_cluster config -h
```

Usage:

```
ufm_ha_cluster config [<options>]
```

Option	Description
<code>-r</code> <code>--role <node role></code>	Node role (master or standby).
<code>-n</code> <code>--peer-node <node-hostname></code>	Peer node name.
<code>-s</code> <code>--peer-sync-ip <ip address></code>	Peer node sync IP address
<code>-c</code> <code>--sync-interface</code>	Local interface to be used for DRBD sync
<code>-i</code> <code>--virtual-ip <virtual-ip></code>	Cluster virtual IP (should be used for master only)
<code>-p</code> <code>--hacluster-pwd <pwd></code>	HA cluster user password.
<code>-h</code> <code>--help</code>	Show this message
<code>-N</code> <code>--no-vip</code>	Configure HA without virtual IP

To configure HA, follow the below instructions:



Please change the variables in the commands below based on your setup.

1. [On Both Servers] Run the following command to set the cluster password:

```
ufm_ha_cluster set-password -p <cluster_password>
```

2. [On Standby Server] Run the following command to configure Standby Server:

```
ufm_ha_cluster config -r standby -n <master_hostname> -s <master_ip_address> -c  
<standby_sync_interface_name> -p <cluster_password>
```

3. [On Master Server] Run the following command to configure Master Server:

```
ufm_ha_cluster config -r master -n <standby_hostname> -s <standby_ip_address> -c  
<master_sync_interface_name> -i <virtual_ip_address> -p <cluster_password>
```

Starting HA Cluster

- To start UFM HA cluster:

```
ufm_ha_cluster start
```

- To check UFM HA cluster status:

```
ufm_ha_cluster status
```

- To stop UFM HA cluster:

```
ufm_ha_cluster stop
```

- To uninstall UFM HA, first stop the cluster and then run the uninstallation command as follows:

```
/opt/ufm/ufm_ha/uninstall_ha.sh
```

5.6.6 Upgrading From Existing UFM Container



Upgrade the UFM container based on the existing UFM configuration files that are mounted on the server. It is important to use that same directory as a volume for the UFM installation command.

In the below example /opt/ufm_files is used.

5.6.6.1 Upgrading UFM Container in SA Mode

1. Stop the UFM Enterprise service. Run:

```
systemctl stop ufm-enterprise
```

2. Remove the old docker image. Run:

```
docker rmi mellanox/ufm-enterprise:latest
```

3. Load the new UFM docker image. Run:

```
docker pull mellanox/ufm-enterprise:latest
```

4. Run the docker upgrade command:

```
docker run -it --name=ufm_installer --rm \  
-v /var/run/docker.sock:/var/run/docker.sock \  
-v /etc/systemd/system:/etc/systemd_files/ \  
-v /opt/ufm/files:/opt/ufm/shared_config_files/ \  
mellanox/ufm-enterprise:latest --upgrade
```

5. Reload system manager configuration:

```
systemctl daemon-reload
```

6. Start UFM Enterprise service:

```
systemctl start ufm-enterprise
```

5.6.6.2 Upgrading UFM Container in HA Mode

1. Stop HA Cluster on the master node. Run:

```
ufm_ha_cluster stop
```

2. Remove the old docker image from both servers. Run:

```
docker rmi mellanox/ufm-enterprise:latest
```

3. Load the new docker image on both servers. Run:

```
docker pull mellanox/ufm-enterprise:latest
```


4. Run the docker command to upgrade UFM on the master node. Run:

```
docker run -it --name=ufm_installer --rm \  
-v /var/run/docker.sock:/var/run/docker.sock \  
-v /etc/systemd/system:/etc/systemd_files/ \  
-v /opt/ufm/files:/opt/ufm/shared_config_files/ \  
mellanox/ufm-enterprise:latest --upgrade
```

5. Download and extract the latest UFM HA package. Run

```
wget https://www.mellanox.com/downloads/UFM/ufm_ha_4.0.0-8.tgz
```

6. Install the extracted UFM HA package:

 In the below command, please modify the partition name based on the already configured DRBD partition.

```
./install.sh -l /opt/ufm/files/ -d /dev/sda5 -p enterprise
```

7. Start UFM HA cluster. Run:

```
ufm_ha_cluster start
```

5.6.7 Logging Into UFM Web UI


To open UFM WEB UI, open the following URL in your browser: [http://\[SERVER_IP\]/ufm/](http://[SERVER_IP]/ufm/) and type the default credentials.

5.7 Historical Telemetry Collection in UFM

5.7.1 Storage Considerations

UFM periodically collects fabric port statistics and saves them in its SQLite database. Before starting up UFM Enterprise, please consider the following disk space utilization for various fabric sizes and duration.

The measurements in the table below were taken with sampling interval set to once per 30 seconds.

 Be aware that the default sampling rate is once per 300 seconds. Disk utilization calculation should be adjusted accordingly.

Number of Nodes	Ports per Node	Storage per Hour	Storage per 15 Days	Storage per 30 Days
16	8	1.6 MB	576 MB (0.563 GB)	1152 MB (1.125 GB)
100	8	11 MB	3960 MB (3.867 GB)	7920 MB (7.734 GB)
500	8	50 MB	18000 MB (17.58 GB)	36000 MB (35.16 GB)
1000	8	100 MB	36000 MB (35.16 GB)	72000 MB (70.31 GB)

6 UFM Software Installation Prerequisites

Before installing UFM software, verify the prerequisites for standalone and high availability installation.

6.1 Prerequisites for UFM Server Software Installation

Please refer to the [UFM Quick Start Guide](#).

6.2 Additional Prerequisites for UFM High Availability (HA) Installation

Please refer to the [UFM Quick Start Guide](#).

7 UFM System Requirements

Please refer to [Installation Notes](#) for information on system prerequisites.

8 UFM Server Health Monitoring

The UFM Server Health Monitoring module is a standalone module that monitors UFM resources and processes according to the settings in the `/opt/ufm/files/conf/UFMHealthConfiguration.xml` file.

For example:

- Each monitored resource or process has its own failure condition (number of retries and/or timeout), which you can configure.
- If a test fails, UFM will perform a *corrective operation*, if defined for the process, for example, to restart the process. You can change the configured corrective operation. If the corrective operation is set to "None", after the defined number of failures, the *give-up* operation is performed.
- If a test reaches the configured threshold for the number of retries, the health monitoring initiates the *give-up* operation defined for the process, for example, UFM failover or stop.
- By default, events and alarms are sent when a process fails, and they are also recorded in the internal log file.

Each process runs according to its own defined schedule, which you can change in the configuration file.

Changes to the configuration file take effect only after a UFM Server restart. (It is possible to kill and run in background the process `nohup python /opt/ufm/ufmhealth/UfmHealthRunner.pyo &.`)

You can also use the configuration file to improve disk space management by configuring:

- How often to purge MySQL binary log files.
- When to delete compressed UFM log files (according to free disk space).

The settings in the `/opt/ufm/files/conf/UFMHealthConfiguration.xml` file are also used to generate the UFM Health Report.

The following section describes the configuration file options for UFM server monitoring.

8.1 UFM Health Configuration

The UFM health configuration file contains three sections:

- Supported Operations—This section describes all the operations that can be used in tests, and their parameters.
- Supported Tests—This section describes all the tests. Each test includes:
 - The main test operation.
 - A corrective operation, if the main operation fails.
 - A give-up operation, if the main operation continues to fail after the corrective operation and defined number of retries.

The number of retries and timeout is also configured for each test operation.

- Test Schedule - This section lists the tests in the order in which they are performed and their configured frequency.

The following table describes the default settings in the `/opt/ufm/files/conf/UFMHealthConfiguration.xml` file for each test. The tests are listed in the order in which they are performed in the default configuration file.

You might need to modify the default values depending on the size of your fabric.


For example, in a large fabric, the SM might not be responsive for *sminfo* for a long time; therefore, it is recommended to increase the values for timeout and number of retries for *SMResponseTest*.

Recommended configurations for *SMResponseTest* are:

- For a fabric with 5000 nodes:
 - Number of retries = 12
 - Frequency = 10
- For a fabric with 10000 nodes:
 - Number of retries = 12
 - Frequency = 20

Test Name / Description	Test Operation	Corrective Operation (if Test Operation fails)	No. Retries / Give-up Operation	Test Frequency
CpuUsageTest Checks total CPU utilization.	CPUTest Tests that overall CPU usage does not exceed 80% (this percentage is configurable).	None If UFM Event Burst Management is enabled, it is automatically initiated when the test operation fails	1 Retry None	1 minute
AvailableDiskSpaceTest Checks available disk space.	FreeDiskTest Tests that disk space usage for <i>/opt/ufm</i> does not exceed 90% (this percentage is configurable).	CleanDisk Delete compressed UFM log files under <i>/opt/ufm</i>	3 Retries None	1 hour
CheckIBFabricInterface Checks state of active fabric interface.	IBInterfaceTest Tests that active fabric interface is up.	BringUpIBFabricInterface Bring up the fabric interface	3 Retries SMOrUFMFailoverOrDoNothing	35 seconds
CheckIBFabricInterfaceStandby (HA only) Checks state of fabric interface on standby.	IBInterfaceTestOnStandby Tests that fabric interface on standby is up.	None	1 Retry None	1 minute
MemoryTest Checks total memory usage.	MemoryUsageTest Tests that memory usage does not exceed 90% (this percentage is configurable).	None	1 Retry None	1 minute
SMProcessTest Checks status of the OpenSM service.	SMRunningTest Tests that the SM process is running.	RestartProcess Restart the SM process	1 Retry UFMFailoverOrDoNothing	10 seconds
SMResponseTest Checks responsiveness of SM (when SM process is running).	SMTTest Tests SM responsiveness by sending the <i>sminfo</i> query to SM.	None	9 Retries UFMFailoverOrDoNothing	10 seconds

Test Name / Description	Test Operation	Corrective Operation (if Test Operation fails)	No. Retries / Give-up Operation	Test Frequency
IbpmTest Checks status of the IBPM (Performance Manager) service.	ProcessIsRunningTest Tests that the IBPM service is running.	RestartProcess Restart the IBPM service	3 Retries None	1 minute
ModelMainTest Checks status of the main UFM service	ProcessIsRunningTest Tests that the UFM service is running.	RestartProcess Restart the UFM service	3 Retries UFMFailoverOrDoNot hing	20 seconds
HttpdTest Checks status of the httpd service.	ProcessIsRunningTest Tests that the httpd service is running.	RestartProcess Restart the httpd service	3 Retries None	20 seconds
MySqlTest Checks status of the MySql service.	ConnectToMySql Tests that the MySql service is running.	None	1 Retry UFMFailoverOrDoNot hing	20 seconds
CleanMySql Purges MySql Logs	AlwaysFailTest Fails the test in order to perform the corrective action.	PurgeMySqlLogs Purge all MySql Logs on each test	1 Retry None	24 hours
UFMServerVersionTest Checks UFM software version and build.	UfmVersionTest Returns UFM software version information.	None	1 Retry None	24 hours
UFMServerLicenseTest Checks UFM License information.	UfmLicenseTest Returns UFM License information.	None	1 Retry None	24 hours
UFMServerHAConfiguration Test (HA only) Checks the configuration on master and standby.	UfmHAConfigurationTest Returns information about the master and standby UFM servers.	None	1 Retry None	24 hours
UFMMemoryTest Checks available UFM memory.	UfmMemoryUsageTest Tests that UFM memory usage does not exceed 80% (this percentage is configurable).	None	1 Retry None	1 minute
UFMCpuUsageTest Checks UFM CPU utilization.	CPUtest Tests that UFM CPU usage does not exceed 60% (this percentage is configurable).	None	1 Retry None	1 minute
CheckDrbdTcpConnectionPerformanceTest (HA only) Checks the tcp connection between master and standby	TcpConnectionPerformanceTest Tests that bandwidth is greater than 100 Mb/sec and latency is less than 70 usec (configurable).	None	2 Retry None	10 minute

 The Supported Operations section of the configuration file includes additional optional operations that can be used as corrective operations or give-up operations.

8.1.1 UFM Core Files Tracking

To receive a notification every time OpenSM or ibpm creates a core dump, please refer to the list of all current core dumps of OpenSM and ibpm in the UFM health report.

 To receive core dump notifications, do the following:

1. Set the `core_dumps_directory` field in the `gv.cfg` file to point to the location where all core dumps are created (by default, this location is set to `/tmp`).

```
core_dumps_directory = /tmp
```

2. Set the naming convention for the core dump file. The name must include the directory configured in the step above.

The convention we recommend is:

```
echo "/tmp/%t.core.%e.%p.%h" > /proc/sys/kernel/core_pattern
```

3. Make sure core dumps directory setting is persistent between reboots. Add the `kernel.core_pattern` parameter with the desired file name format to the `/etc/systctl.conf` file. Example:

```
kernel.core_pattern=/tmp/%t.core.%e.%p.%h
```

4. Configure the core file size to be unlimited.

```
ulimit -c unlimited
```

5. (Only on UFM HA master) Update the UFM configuration file `gv.cfg` to enable core dump tracking.

```
track_core_dumps = yes
```

8.2 Example of Health Configuration

The default configuration for the overall memory test in the `opt/ufm/files/conf/UFMHealthConfiguration.xml` file is:

```
<Test Name="MemoryTest" NumOfRetriesBeforeGiveup="3" RetryTimeoutInSeconds="10">
  <TestOperation Name="MemoryUsageTest">
    <Parameters>
      <Parameter Name="ThresholdInPercents" Value="90"/>
    </Parameters>
  </TestOperation>
  <CorrectiveOperation Name="None"/>
  <GiveupOperation Name="None"/>
</Test>
```

This configuration tests the available memory. If memory usage exceeds 90%, the test is repeated up to 3 times at 10 second intervals, or until memory usage drops to below 90%. No corrective action is taken and no action is taken after 3 retries.

To test with a usage threshold of 80%, and to initiate UFM failover or stop UFM after three retries, change the configuration to:

```
<Test Name="MemoryTest" NumOfRetriesBeforeGiveup="3" RetryTimeoutInSeconds="10">
  <TestOperation Name="MemoryUsageTest">
    <Parameters>
      <Parameter Name="ThresholdInPercents" Value="80"/>
    </Parameters>
  </TestOperation>
  <CorrectiveOperation Name="None"/>
  <GiveupOperation Name="UFMFailoverOrStop"/>
</Test>
```

8.2.1 Event Burst Management

UFM event burst management can lower the overall CPU usage following an event burst by suppressing events. Event burst management is configured in the *gv.cfg* configuration file.

When the overall CPU usage exceeds the threshold configured by the *CpuUsageTest* in the */opt/ufm/files/conf/UFMHealthConfiguration.xml* file, a High CPU Utilization event occurs.

This event initiates the UFM event burst management, which:

- Suppresses events. The default level of suppression enables critical events only.
- If, after a specified period of time (30 seconds, by default), no further High CPU Utilization event occurs, the UFM server enables all events.

To modify Event burst management configuration, change the following parameters in the *gv.cfg* file:

```
# The events' level in case events are suppressed (the possible levels are disable_all_events,
enable_critical_events, and enable_all_events)
# The entire feature can be turned off using the level "enable_all_events"
suppress_events_level = enable_critical_events
# The amount of time in seconds which events are suppressed
suppress_events_timeout = 30
```

8.3 Recovery from Consecutive Failures

UFM Server Health Monitor might restart or trigger a failover in order to recover from specific failures. In case a re-start or failover fails, UFM Server Health Monitor tries the operation again. Upon a number of consecutive failure attempts to restart or failover, UFM Server Health Monitor stops trying to restart Model Main and allows OpenSM to run without intervention. The behavior maximum number of consecutive restart attempts is defined in the configuration file */opt/ufm/files/conf/UFMHealthConfiguration.xml*:

```
<Parameter Name="RestartAttempts" Value="8"/>
<Failover MaxAllowedAttempts="6"/>
```

9 UFM Web UI

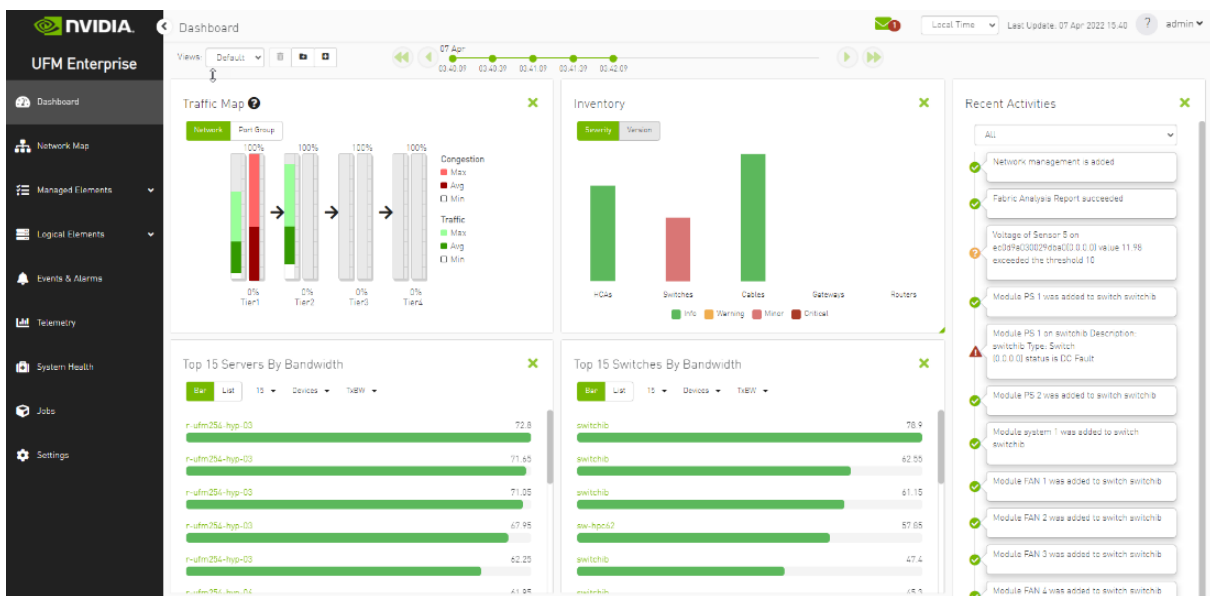
This section is constituted by the following sub-sections:

- [Fabric Dashboard](#)
- [Network Map](#)
- [Managed Elements](#)
- [Logical Elements](#)
- [Events & Alarms](#)
- [Telemetry](#)
- [System Health](#)
- [Jobs](#)
- [Settings](#)

9.1 Fabric Dashboard

The dashboard window summarizes the fabric's status, including events, alarms, errors, traffic and statistics.

Fabric Dashboard View




The Fabric Dashboard view consists of the following six dashboards, which provide real-time information about the fabric.

9.1.1 Dashboard Views and Panel Management

UFM is installed with a default view of the most important panels. These panels are resizable and draggable. Users can customize their default view or create new views altogether

The dashboard views and panels are managed by a set of action buttons appearing at the top of the main dashboard screen:




Clicking on the Add Panel button will show a model to select which panels you wish to add to the current dashboard view.

The screenshot shows the 'New Panels' selection dialog. On the left, a list of categories is shown with their respective counts: All (12), Health (2), Monitoring (7), and Events and Alarms (3). The main area displays a grid of dashboard panels. Two panels are selected, indicated by a checkmark in the top right corner of each panel's header: 'Inventory' and 'Traffic Map'. At the bottom, a status bar shows '2 Panels Selected' and two buttons: 'Add Panels' and 'Cancel'.

9.1.2 Dashboard Timeline Snapshots

Once the user is logged into the UFM Enterprise, the UFM will start recording snapshots of the dashboard panel data every 30 seconds.

The user is able to navigate between these snapshots and load the dashboard data of a specific data snapshot.



9.1.3 Dashboard Panels

The Fabric Dashboard view consists of the following 12 panels, which are categorized into 3 main categories and provide real-time information about the fabric.


- Health:
 - Inventory
 - Fabric Health
- Monitoring:
 - Traffic Map
 - Levels Traffic Map
 - Top X Servers by bandwidth
 - Top X Switches by bandwidth
 - Top X congested servers
 - Top X congested switches
 - Top X utilized Pkeys
- Events and Alarms:
 - Recent Activities
 - Top X alarmed servers
 - Top X alarmed switches

9.1.4 Top N Servers/Switches by Rx or Tx Bandwidth


The Top N servers/switches by Rx or Tx Bandwidth component shows the top elements that are transmitting or receiving the most bandwidth per second. These elements are classified top-down according to the defined Transmit (Tx) or Receive (Rx) bandwidth (MB/sec Rate).


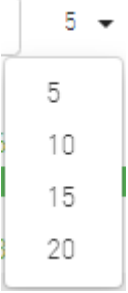

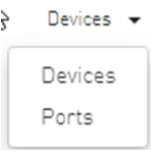

Bandwidth is measured as a rate in bytes/sec.

- Transmitted (Tx) bandwidth is measured by N server/switch ports in MB/sec
- Received (Rx) bandwidth is measured by N server/switch ports in MB/sec

 N can be 5, 10, 15, or 20.

The following table lists the icons of this component:

Options	Description
List view 	Shows the top N elements as a list Each element is shown in a row with the name of the element and the bandwidth rate

Options	Description
<p>Bar view</p> 	<p>Shows the top N nodes as a bar graph</p> <ul style="list-style-type: none"> • X axis shows the rate as a value • Y axis shows the Node (server) name
<p>Drop-down menu</p> 	<p>Selects the number of items to display Default: 10 nodes</p>
<p>Monitoring attributes</p> 	<p>Selects the attribute for monitoring:</p> <ul style="list-style-type: none"> • TxBW - Transmit Bandwidth • RxBW - Receive Bandwidth
<p>View by port/element</p> 	<p>Switches view to top 5 elements by bandwidth or top 5 ports by bandwidth. Nodes view is presented by default.</p> <ul style="list-style-type: none"> • Clicking a specific port in the ports view under the port column redirects to the ports table and highlights that particular port • Clicking a specific device in the devices view under the device column redirects to the Devices table and highlights that particular node
<p>Filter toggle</p> 	<p>Toggles the filter textbox</p>

Top Servers/Switches by Bandwidth—Bar View



Top Servers/Switches by Bandwidth—List View

Top 15 Servers By Bandwidth ✕

Bar List
15 Devices TxBW

5

Device	TxBW BandWidth (Gbps) ↓
r-ufm254-hyp-04	75.35
r-ufm254-09	74.6
r-ufm254-011	65.95
r-ufm254-04	64.7
r-ufm254-012	63.2

1 to 5 of 15 << < Page 1 of 3 > >>

Right-clicking a device displays a list of the actions that can be performed. These actions (shown in the following screenshot) are the same actions available in the devices table (see [Devices Actions](#) table under [Devices Window](#)).

Top 15 Servers By Bandwidth ✕

Bar List
15 Devices TxBW

5

Device	TxBW BandWidth (Gbps)
r-ufm254-hyp-03	38.8
r-ufm254-hy	40.1
ufm-host87	79.05
r-ufm254-01	47.6
r-ufm254-02	72.8

- Mark As Unhealthy ▶
- Firmware Upgrade
- Add To Group ▶
- Remove From Group ▶
- Suppress Notifications
- Add To Monitor Session

o 5 of 15 << < Page 1 of 3 > >>

Right-clicking a port displays a list of the actions that can be performed. These actions (shown in the following screenshot) are the same actions available in the Ports table (see [Ports Window](#) for more information).

The screenshot shows a window titled "Top 15 Servers By Bandwidth" with a close button (X) in the top right. Below the title are controls for view type (Bar/List), count (15), sort criteria (Ports/TxBW), and a refresh button (5). The main area is a table with two columns: "Port" and "TxBW BandWidth (Gbps)". The table lists five server entries. A context menu is open over the first row, showing the following actions: "Go To Peer", "Reset", "Mark As Unhealthy", and "Disable". At the bottom right of the table, there are pagination controls: "1 to 5 of 15", navigation arrows, and "Page 1 of 3".

Port	TxBW BandWidth (Gbps)
r-ufm254-hyp-03 HCA-1 (port #11)	13.85
r-ufm254-hyp-03 HCA-1 (port #12)	77.6
ufm-host87 HCA-1 (port #11)	52.95
r-ufm254-01	34.8
r-ufm254-02	65.95

9.1.5 Top N Congested Servers/Switches by Rx/Tx Bandwidth



The Top N Congested devices by Rx or Tx Bandwidth component shows the top congested devices, classified top-down according to the defined Transmit (Tx) or Receive (Rx) bandwidth.

Bandwidth is measured as congestion bandwidth rate (CBW) by percentage.


- For Tx, congestion is measured by N HCA ports.
- For Rx, congestion is measured by N switch ports connected to HCAs.

N can be 5, 10, 15, or 20.

Top N Congested Servers by Bandwidth—List View



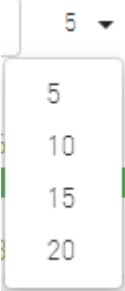


Top N Congested Servers/Switches by Bandwidth—Bar View



The following table describes the options available in this component.

Top N Congested Devices by Rx/Tx Bandwidth

Options	Description
Bar view 	Shows the top N congested devices as a bar graph <ul style="list-style-type: none"> • X axis shows the rate as a percentage • Y axis shows the congested Node (server) name
List view 	Shows the top N congested nodes as a list Each congested node is shown in a row with the name of the node and its picture. It also shows the bandwidth rate
Drop-down menu 	Enables selecting the number of top N congested nodes Default: 10 nodes

Options	Description
<p>View by port/element</p> <p>› Devices ▾</p> <div style="border: 1px solid #ccc; padding: 5px; width: fit-content;"> <p>Devices</p> <p>Ports</p> </div>	<p>Switches view to Top 5 elements By Bandwidth or Top 5 Ports By Bandwidth. Devices view is presented by default.</p> <ul style="list-style-type: none"> Clicking a specific port in the Ports view under the Port column redirects to the Ports table and highlights that particular port Clicking a specific device in the Nodes view under the Device column redirects to the Devices table and highlights that particular node
<p>Monitoring attributes</p> <p>· TxBW ▾</p> <div style="border: 1px solid #ccc; padding: 5px; width: fit-content;"> <p>TxBW</p> <p>RxBW</p> </div>	<ul style="list-style-type: none"> RCBW - Receive Congested Bandwidth (percentage) TCBW - Transmit Congested Bandwidth (percentage)

9.1.6 Top N Utilized PKeys

Top N Utilized PKeys displays the top utilized PKeys based on the number of the PKey members.

⚠ N can be 5, 10, 15, or 20.

Top N Utilized PKeys—List View

Top 5 Utilized PKeys


Bar
List
5 ▾

5 ▾
🔍

Pkey	# of GUIDs
0x22	3
0x23	2
0x77	1
0x7fff	0




1 to 4 of 4
|<
<
Page 1 of 1
>
>|

Top N Utilized PKeys—Bar View



The following table describes the options available in this component.

Top N Utilized PKeys

Options	Description
Bar view 	Shows the top N Utilized PKeys as a bar graph <ul style="list-style-type: none"> • X axis shows the number of members • Y axis shows the names of the PKeys
List view 	Shows the top N Utilized PKeys as a list Each PKey is shown in a row with the name of the PKey and the number of its members
Drop-down menu 	Enables selecting the number of top N Utilized PKeys Default: 10 Utilized PKeys

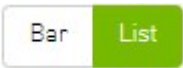

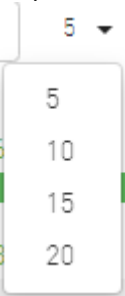

9.1.7 Top N Alarmed Servers/Switches

The Top N Alarmed Servers/Switches component shows the top nodes with alarms classified in a descending order. Alarmed nodes are measured according to the following:

- Severity - only the top nodes, in order of severity:
 - Critical
 - Minor
 - Warning
 - Normal
- Alarm - numbers (N can be 5, 10, 15, or 20)

The following table lists the components.

Top N Alarmed Servers/Switches

Options	Description
List view 	Shows the top N alarmed servers/switches as a list. Each alarmed device is shown in a row with the name of the node and the number of alarms.
Bar view 	Shows the top N alarmed devices as a bar graph. <ul style="list-style-type: none"> • X axis shows the number of alarms • Y axis shows the names of the alarmed nodes (servers)
Drop down menu 	Enables selecting the number of top N alarmed nodes. Selects the number of items to display. Default: 10 alarmed nodes
Filter toggle 	Toggles the Filter textbox

Top Alarmed Servers/Switches—List View

Top 5 Alarmed Servers ✕

Bar **List** 5 ▾

5 ▾

Device	Alarms
r-ufm254-hyp-03	9
r-ufm254-hyp-04	9
ufm-host87	7

1 to 3 of 3 ⏪ < Page 1 of 1 > ⏩

Top 5 Alarmed Switches ✕



Bar **List** 5 ▾

5 ▾

Device	Alarms
sw-hpc62	9
switchib	8


1 to 2 of 2 ⏪ < Page 1 of 1 > ⏩


Top N Alarmed Servers/Switches—Bar View



9.1.8 Inventory Summary

The Fabric Inventory Summary component shows a summary of your fabric inventory (HCAs, Switches, Gateways, Routers and Cables) categorized by the element's severity or firmware version.





Clicking on one bar element with specific severity/firmware version will redirect you to the clicked element's table.

9.1.9 Fabric Utilization


The Fabric Utilization component shows the number of alarmed objects, categorized by the alarm's severity. They are as follows:

1. Warning
2. Minor
3. Normal
4. Critical

If Server X has 2 minor alarms, 1 warning alarm and 2 critical alarms, and Server Y has 0 minor alarms, 2 warning alarms and 1 critical alarm, the Fabric Resource Utilization pie chart will show 2 servers in the critical slice, 2 servers in the warning slice and 1 server in the minor slice.

You can filter for both switches and nodes of a specific severity level by clicking the specific pie slice indicating the severity.

In the example below, the Devices table lists all the switches of severity level "Minor" after clicking the red (Minor) slice from the Switches pie chart.



Devices Local Time Last Update: 07 Apr 2022 17:01 admin

Showing 2 out of 5. Click to reset all filters

Severity	Name	GUID	Type	Model	IP	Firmware Version
Minor	sw-p0c02	6-7d4e400300a5a2e0	switch	M587600	N/A	15.1200.102
Minor	switch-9	6res0e7e03002768e0	switch	EDR	N/A	19.2009.1029

Viewing 1-2 of 2

9.1.10 Recent Activities

The Recent Activities component lists the recent events detected by the UFM system.


Recent Activities X

All ▾

- ✓ Network management is added
- ✓ Fabric Analysis Report succeeded
- ? Voltage of Sensor 5 on ec0d9a030029dba0(0.0.0.0) value 11.98 exceeded the threshold 10
- ✓ Module PS 1 was added to switch switchib
- ! Module PS 1 on switchib Description: switchib Type: Switch (0.0.0.0) status is DC Fault
- ✓ Module PS 2 was added to switch switchib
- ✓ Module system 1 was added to switch switchib
- ✓ Module FAN 1 was added to switch switchib
- ✓ Module FAN 2 was added to switch switchib
- ✓ Module FAN 3 was added to switch switchib
- ✓ Module FAN 4 was added to switch switchib

You can filter for the events you would like to see in one list using the drop-down menu that provides the following options:

- All - shows all recent activities
- All issues - shows all non-Info activities
- Info - shows all activities with Info severity or higher
- Minor - shows you all activities with Minor severity or higher
- Warning - shows you all activities with Warning severity or higher
- Critical - shows you all activities with Critical severity




9.1.11 Traffic Map

The Traffic Map dashboard shows the normal traffic versus congested traffic distributed on switch tiers and on port groups. This view, together with the Top N Congestion dashboard, gives a full status of the traffic congestion of the fabric.

9.1.11.1 Network Traffic Map

Four double bars represent the transmitted bandwidth (normalized transmit data) and normalized congested bandwidth (CBW), both measured in bytes/sec with minimum, average, and maximum bandwidth values.



An explanatory window on traffic map opens once clicked on the ? icon.

Traffic Map Guide

Mellanox's unique Traffic Map provides a valuable real-time aggregate view of the fabric performance by showing the overall bandwidth utilization per switching tier coupled with congestion information.

Reading the Traffic Map Chart

The Traffic Map contains four tiers; each tier is represented by a green and a red bar, as shown in the following Traffic Map Chart :

Color coding for each tier is as follows:

- The green the percentage of overall bandwidth generated by the specific tier. This bar is divided in light and dark green colors.
 - The light green indicates the peak port utilization.
 - The dark green indicates the average utilization.
- The red bar indicates the percentage of congestion (also referred to as lost bandwidth) in the specific tier. This bar is divided in red and dark red colors.
 - Red indicates the peak port congestion.
 - Dark red indicates the average congestion.

Close

The percentage of total theoretical bandwidth (TBW) is calculated based on the underlying InfiniBand technology (SDR, DDR, QDR, FDR or EDR). The speed can be viewed when checking the ports.

- The vertical axis shows the following:
 - Bandwidth (BW) is represented by a green bar and is measured in percentages
 - Congested Bandwidth (CBW) is represented by a red bar and is measured in percentages
 - Minimum, average, and maximum bandwidth are represented in each bar by a subset color
- The horizontal axis represents the tiers.

The bottom of the dashboard represents the tier-related transmitted traffic, which is divided into four segments by measurement ports:

 - Tier 1 - represents the traffic injected by all adapters
 - Tier 2 - represents the traffic sent from the edge switches to the core of the fabric (in case of a single Director switch, this tier indicates traffic utilization inside the Director between the line and fabric boards)

- Tier 3 - represents the traffic sent from the core to the edge switches
- Tier 4 - represents the traffic sent from the edge switch to the adapters



The illustrations at the bottom of the tiers show a four-tier topology:
Server [tier 1] Switch [tier 2] Director Switch [tier 3] Switch [tier 4] Server.

9.1.11.2 Levels Network Traffic Map

Different representation of the fabric traffic map that based on the devices/ports levels.



The level of the device/port is the distance between the device and the nearest server/gateway.

Levels Calculations:

- The levels calculations are configurable from the `gv.cfg` file under TopologyLevels section enable item and it is disabled by default.
- The levels names are configurable from the `gv.cfg` file under TopologyLevels section levels item and by default we are defining up to 4 levels levels equals server, leaf, spine, core

- Server: hosts and gateways.
- Leaf: switches and routers that are directly connected to the server
- Spine: switches and routers that are directly connected to the leaf
- Core: switches and routers that are directly connected to the spine

If the fabric has more than 4 levels, the level value will be L + distance e.g., L4, L5, L(N), and if levels was empty, the levels will start from L0, L1, L2, etc.

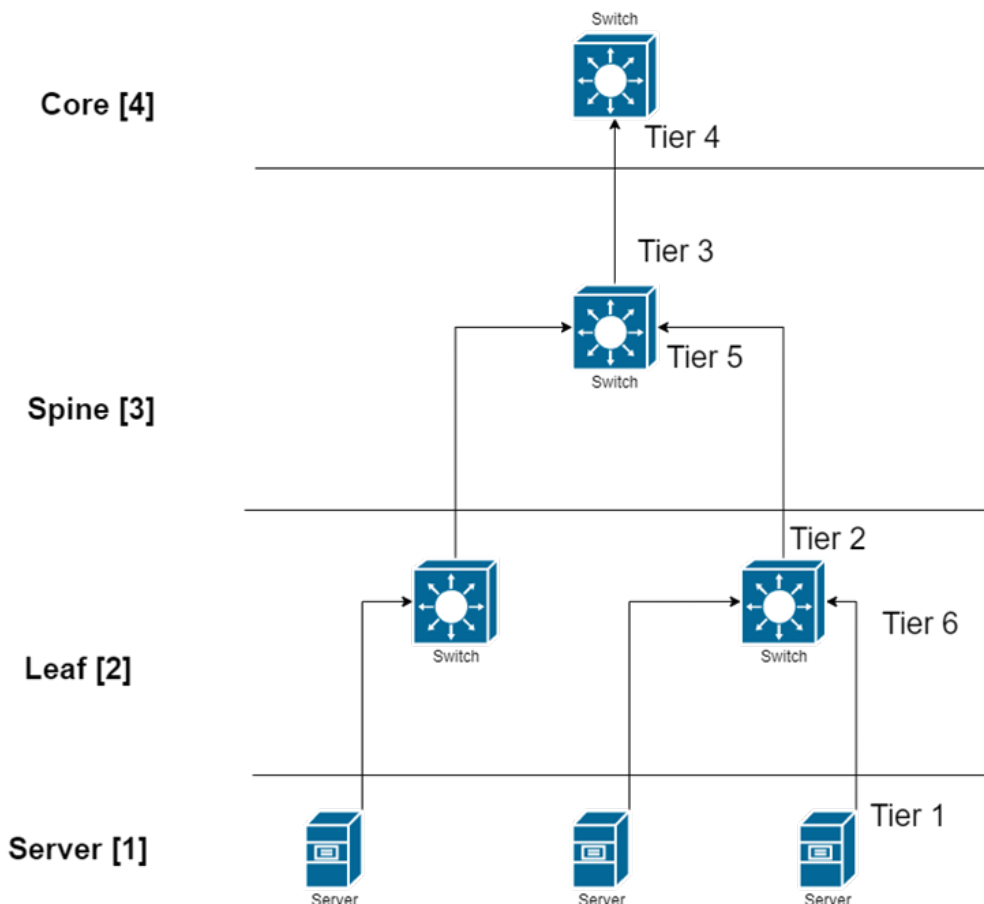
The levels calculations are done at either the discovery stage or once the topology changes.

Ports Tiers calculations based on the levels:

If the levels calculations is enabled, the port's tier will be calculated as the following steps:


1. Get the level for both port's parent device and port's peer parent device
2. Decide whether the port's data flow is the up or down direction, by checking the order of the parent and peer parent level:
 - a. If the parent's level order is less than or equals the parent peer level, then the port's flow is up and tier is the parent level order
 - b. If the port's flow is down and the tier is the distance between the host to the root device and the distance between the root to the parent device

Example:



If the level calculations are disabled, the tier calculations will be done as mentioned in this section.

9.1.11.3 Port Group Traffic Map




9.1.11.4 Traffic Map Bar Chart

- Bandwidth Bars

The bandwidth graph shows how traffic is traversing the fabric and how traffic is being transmitted between the servers. For example, the following considerations could be evaluated:

- The size of the difference between max bandwidth and min bandwidth.
- The traffic that is flowing in the middle tiers and whether it would be more efficient to move the traffic to the edges to save the uplinks.

Bandwidth levels are measured in percentages, as shown below:

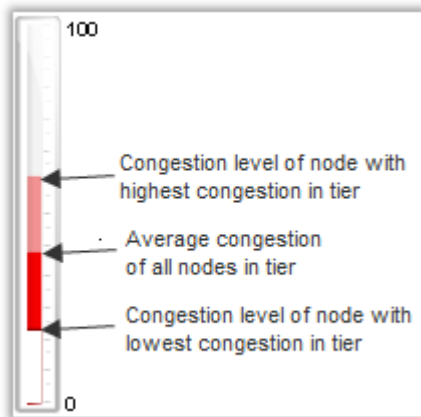


- Congestion Bars

The Congestion graph shows where congestion starts. For example, the following considerations could be evaluated:

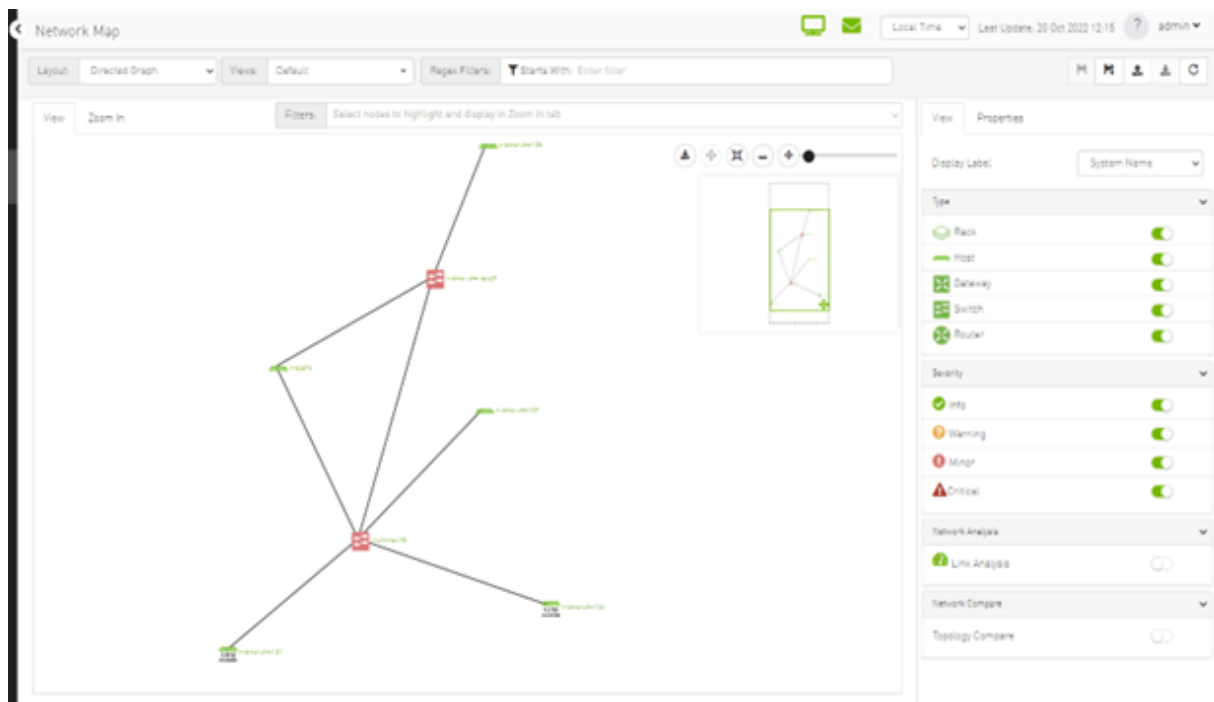
- If congestion is in the first or second tier, there is probably a routing problem
- If there is no red bar, it means that there is no congestion or no routing problems

Congestion levels are measured in percentages, as shown:










9.2 Network Map

The Network Map window shows the fabric, its topology, elements and properties. UFM performs automatic fabric discovery and displays the fabric elements and their connectivity. In the Network Map window, you can see how the fabric and its elements are organized (e.g., switches and hosts).



9.2.1 Network Map Components

Component	Icon	Description
Switches		Represents third party switches discovered/managed by UFM
Hosts		Represents the computer (host) connected to the discovered/managed switches
Routers		Represents third party routers discovered/managed by UFM
Gateways		Represents third party gateways discovered/managed by UFM
Links		Represents the connections between devices on the fabric
Racks		Represents all nodes (hosts) physically connected to a switch

 The level of severity of devices affects the color they are displayed in. For further information, refer to table "[Device Severity Levels](#)".

- To zoom in/out of the map, scroll the mouse wheel up and down or using the slider on the right top corner
- To move around in the map, press and hold down the left key while you move sideways and up/down
- To see the hosts inside a rack, right-click the Rack icon and click "Expand Hosts"



9.2.2 Selecting Map Elements

Users are able to select elements from the Network Map. Right-clicking an element opens a context menu which allows users to perform actions on it.

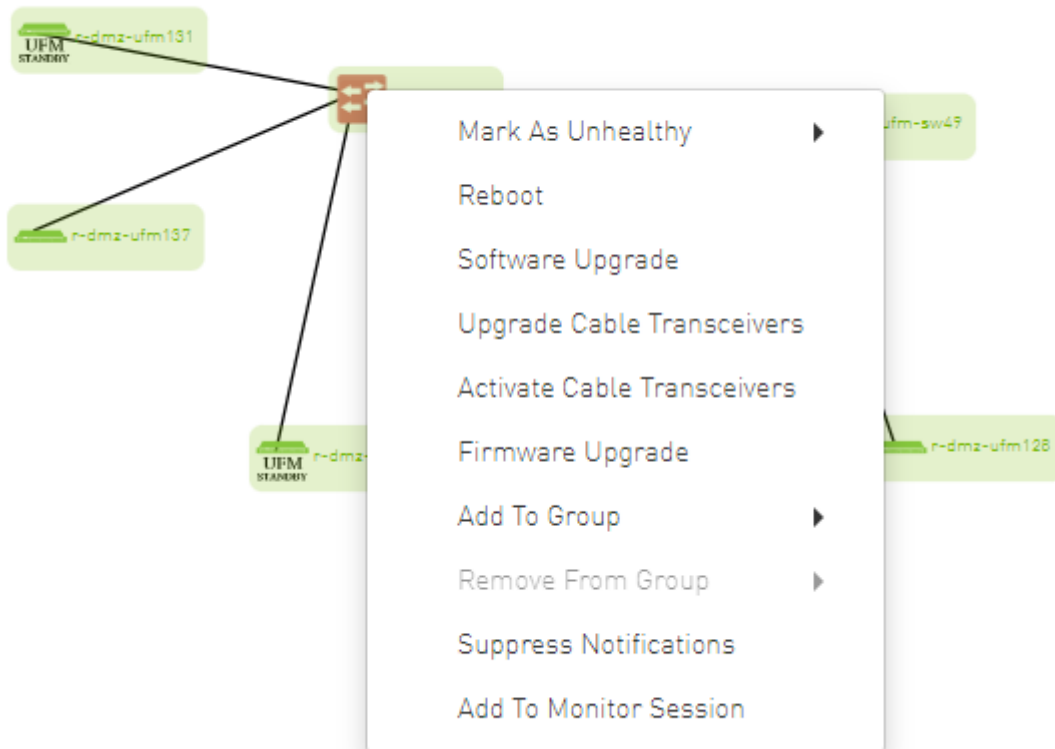
It is possible to select multiple elements at once using any of the following methods:

- By holding down Ctrl or Shift and dragging their mouse across the map.

 Please note that Ctrl starts new selection, while Shift adds to the current selection.

- By holding down Shift and clicking a new element on the map.

Multi-select makes it possible for users to perform actions on multiple devices with one right-click rather than repeating the same process per device.








9.2.3 Map Information and Settings

The right pane of the Network Map view enables you to control the view settings, as well as obtain further information on selected elements from the map.





View Properties

Display Label System Name


Type

 Rack	<input checked="" type="checkbox"/>
 Host	<input checked="" type="checkbox"/>
 Gateway	<input checked="" type="checkbox"/>
 Switch	<input checked="" type="checkbox"/>
 Router	<input checked="" type="checkbox"/>

Severity

 Info	<input checked="" type="checkbox"/>
 Warning	<input checked="" type="checkbox"/>
 Minor	<input checked="" type="checkbox"/>
 Critical	<input checked="" type="checkbox"/>






Network Analysis



 Link Analysis	<input type="checkbox"/>
---	--------------------------


The customized views created using the type and severity filters, selected fabric nodes, zoom level, and Expand/Collapse All Racks options can be saved for later access. These customized views can be saved and accessed using the bar available on top of the Network Map:

Views: Default



Regex Filters: Starts With: Enter filter


    

- "Save As" icon () saves newly created customized views
- "Save" icon () saves edits performed on existing views


- “Import” icon () import map from local device. The file format should be txt




- “Export” icon () export network as text file
- To reload/refresh the network map, use the refresh icon ().
- Drop down menu gives access to all previously saved views



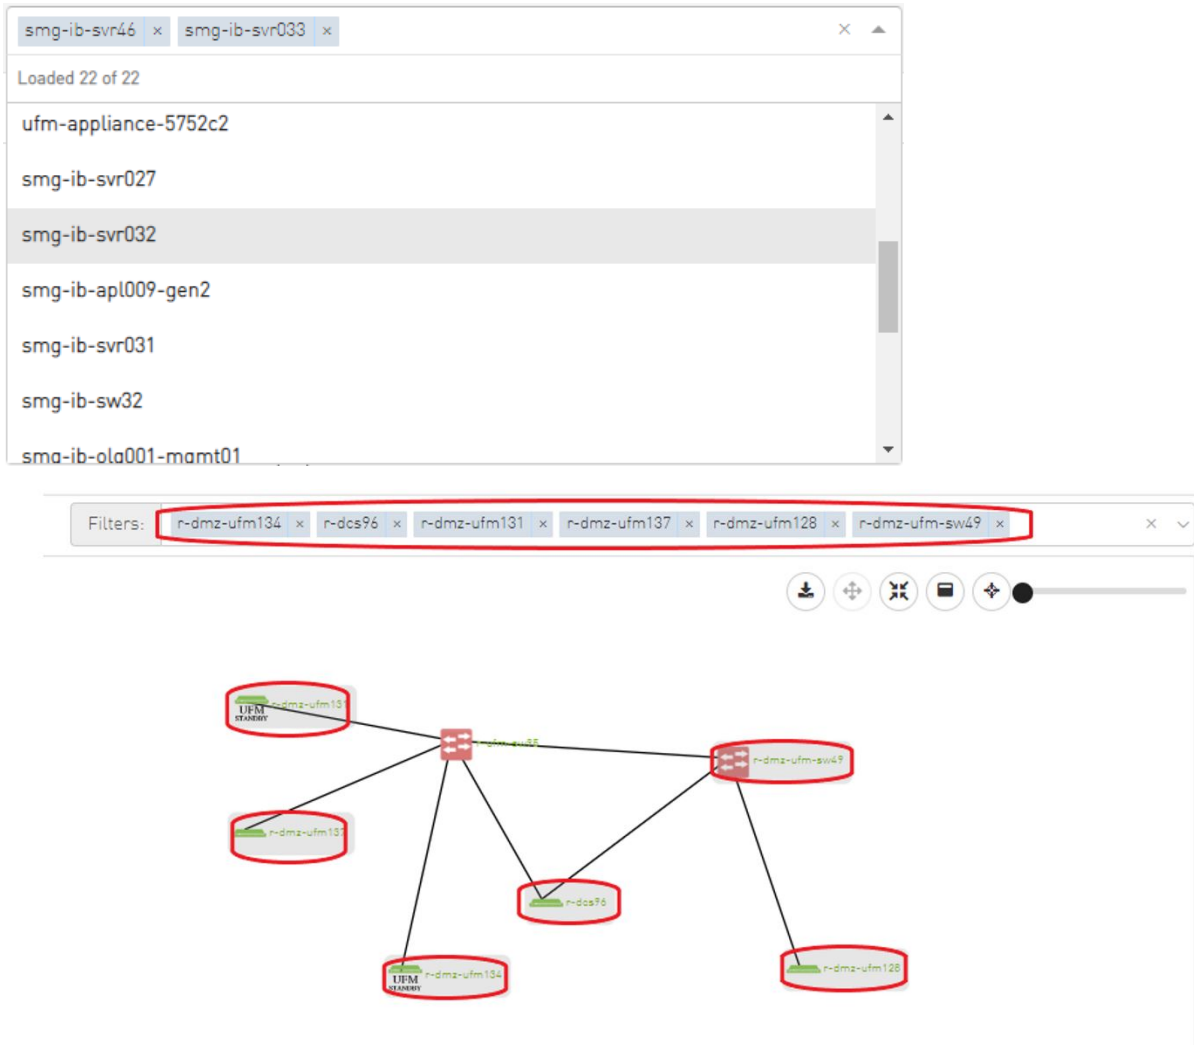
- “Default” view is a predefined view where nodes are positioned randomly, all filters are enabled, and all racks are collapsed. Changes made to this view cannot be saved unless under a new view name using the “Save As” icon.



- Saved views can be deleted using the “x” button.



You can select a node from the dropdown menu located above the Network Map view in order to highlight/display them in the “Zoom In” tab.



9.2.4 Map View Tab

The Network Map "View" tab displays the fabric containing all nodes (e.g. switches, racks including the hosts, etc).

If your fabric consists of more than 500 nodes, please note that:

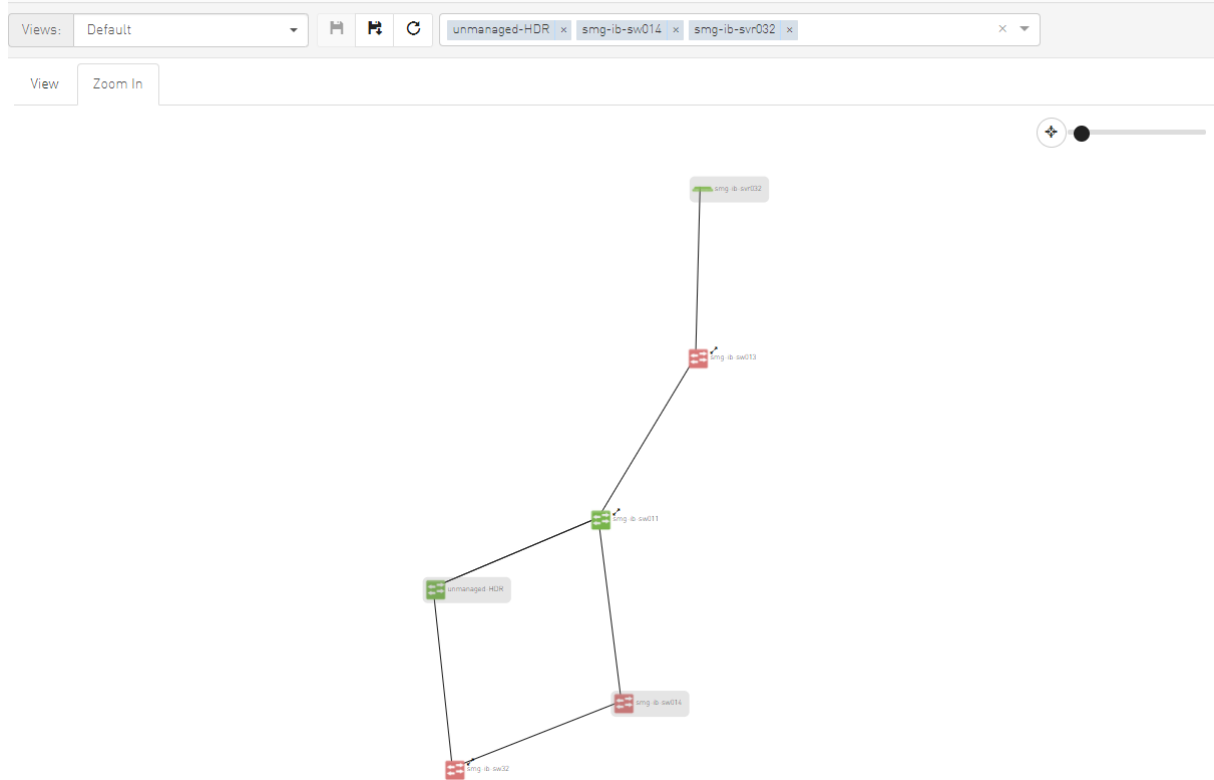
- The "View" tab will show only the switches in your fabric. Therefore, "Expand all racks" and "Rack filter" functions will be disabled.
- Link analysis will be disabled.

To have a better experience in this instance, you can switch to the "Zoom In" tab.


9.2.5 Map Zoom In Tab

The Network Map "Zoom In" tab displays only the selected nodes from the dropdown menu above the map view and the nodes directly connected to the selected nodes.

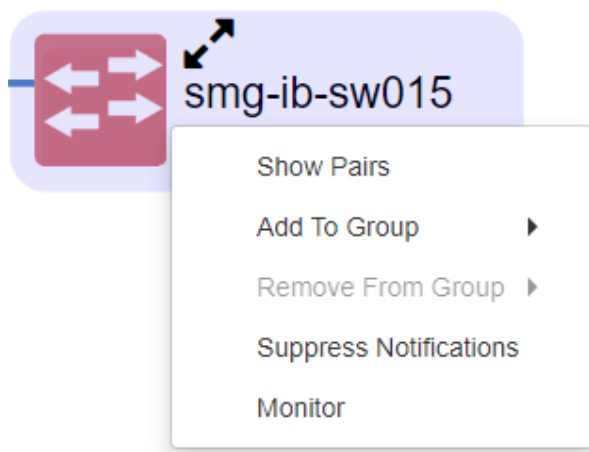
Network Map



If some switches still have hidden connected nodes, you will see the following icon:



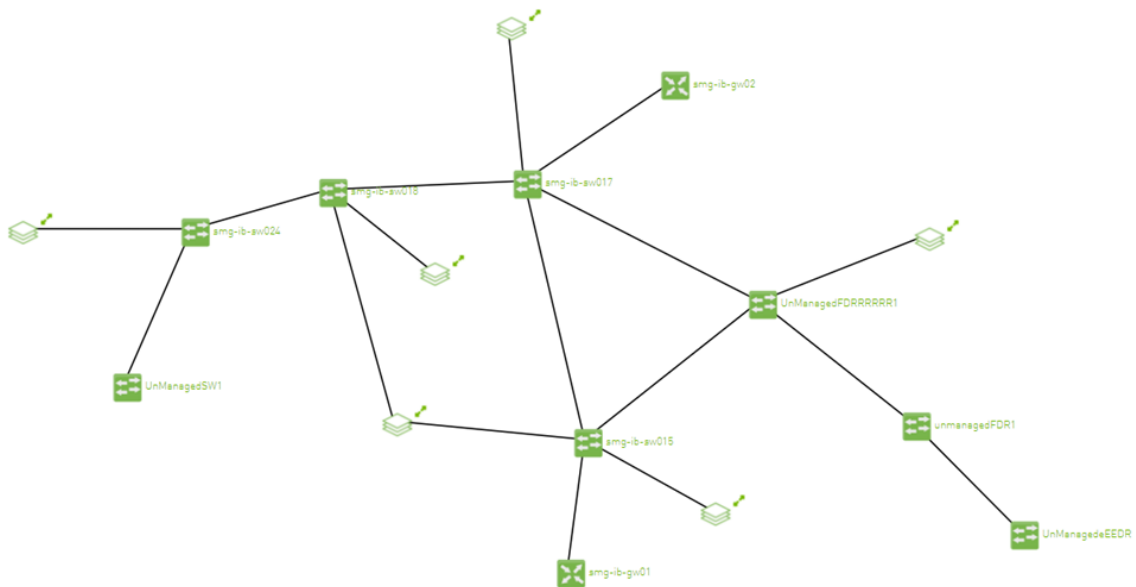
To reveal the hidden nodes connected to this switch, you can right-click it and select "Show Pairs" which adds this switch to the selected nodes list and shows the direct connected nodes to this switch.




9.2.6 Map Layouts

Layout controls nodes positions in the map. UFM network map supports two types of layouts:


- Directed layout: the nodes are distributed depending on the connections between them so that the connected nodes will be near each other without conflict.



- Hierarchical layout: the nodes are distributed as layers; each layer will contain nodes that have the same level value.




You can switch between layouts from the dropdown menu located above the Network Map view.



The default layout for small fabric (less than 30 nodes) is hierarchical and for large fabric is directed.

9.2.7 Information View Tab





- Enables searching for one or more elements in the map, by typing either their name or their GUID in the Search field. Note that the search mechanism is not case-sensitive.
- Enables displaying the elements either by their name, GUID, or IP.
- Enables viewing all hosts of all racks in the fabric using the "Expand All Racks" button.



- Enables customizing the view of the map by filtering for certain elements to appear in the map using the Type (see table "[Network Map Components](#)") and Severity (see table "[Device Severity Levels](#)") filters. Example:

The screenshot shows a network map interface with a 'Zoom In tab' at the top left. The main area displays a hierarchical network diagram. The root node is a red square labeled 'n-ufm-gw35'. It has five child nodes: 'n-dmz-ufm131' (green square), 'n-dmz-ufm137' (green square), 'n-dmz-ufm134' (green square), 'n-dcs70' (green square), and 'n-dmz-ufm128' (green square). The 'n-dmz-ufm131' node has a sub-label 'TIEM station'. The 'n-dmz-ufm134' node has a sub-label 'TIEM station'. The 'n-dcs70' node has a sub-label 'n-dcs70'. The 'n-dmz-ufm128' node has a sub-label 'n-dmz-ufm128'. Above the diagram are several icons for map interaction: a download icon, a pan icon, a zoom in icon, a zoom out icon, a refresh icon, and a search icon. To the right of the diagram is a 'Properties' panel with two tabs: 'View' and 'Properties'. The 'View' tab is active. It contains a 'Display Label' dropdown set to 'System Name'. Below this are two sections: 'Type' and 'Severity'. The 'Type' section has five rows, each with an icon, a label, and a toggle switch: 'Rack' (green square icon, toggle on), 'Host' (green square icon, toggle on), 'Gateway' (green square icon with 'X', toggle on), 'Switch' (green square icon with 'X', toggle on), and 'Router' (green square icon with 'X', toggle on). The 'Severity' section has four rows, each with an icon, a label, and a toggle switch: 'Info' (green checkmark icon, toggle on), 'Warning' (orange question mark icon, toggle on), 'Minor' (red exclamation mark icon, toggle on), and 'Critical' (red exclamation mark icon, toggle on). Below these are two more sections: 'Network Analysis' and 'Network Compare'. 'Network Analysis' has one row: 'Link Analysis' (green square icon with 'X', toggle off). 'Network Compare' has one row: 'Topology Compare' (green square icon with 'X', toggle off).

Device Severity Levels

Component	Description
	Info
	Critical
	Minor
	Warning

9.2.8 Link Analysis

Link analysis allows the user to display the link analytics according to a selected static counter, and define the conditions on which the analysis is based. The links are colored according to the specified conditions. It is possible to define up to five conditions per counter.

The counter's conditions are applied on four values:

- The source values of the selected counter
- The destination value of the selected counter
- The source value of the opposite of the selected counter
- The destination value of the opposite of the selected counter






The worst matched value between these four is taken into consideration.

The "Network Analysis" section on the right side under the View tab contains a radio button to enable/disable the link analysis.





View Properties

Display Label System Name


Type

 Rack	<input checked="" type="checkbox"/>
 Host	<input checked="" type="checkbox"/>
 Gateway	<input checked="" type="checkbox"/>
 Switch	<input checked="" type="checkbox"/>
 Router	<input checked="" type="checkbox"/>

Severity

 Info	<input checked="" type="checkbox"/>
 Warning	<input checked="" type="checkbox"/>
 Minor	<input checked="" type="checkbox"/>
 Critical	<input checked="" type="checkbox"/>

Network Analysis


 Link Analysis

Counter:

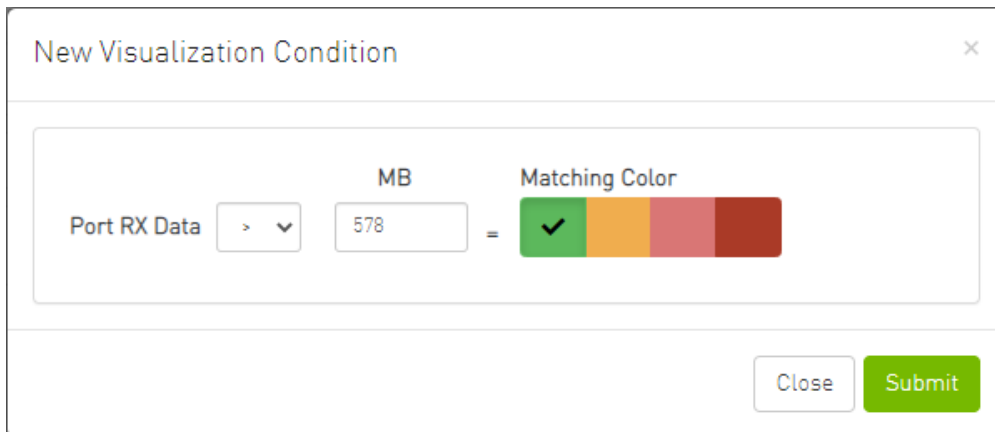
Port RX Data +

To define a condition:

1. Select the desired counter, and click the + button.



2. Select the appropriate operator, and define the desired threshold and color on the form that pops up. This color is applied on the link if the link monitoring value matches the respective condition.



⚠ The colors are sorted from the lowest to the highest priority (i.e from left to right, green to red).

- ⚠** The counter's conditions are sorted based on the threshold values:
- Ascending if the operator is greater than (>)
 - Descending if the operator is smaller than (<)

Last matched condition's color are taken into consideration in the link coloring.

3. Once the condition is set, the network map lights up the links that meet your condition.

The screenshot shows a network map with a central switch (n-ufm-gw99) and several other devices. The links are color-coded: green for 'Port RX Data > 0 Gb' and orange for 'Port RX Data > 140 Gb'. The properties panel on the right shows settings for 'Type', 'Severity', and 'Network Analysis', with 'Link Analysis' enabled and the two conditions listed below.

Type	Severity	Network Analysis
Rack	Info	Link Analysis
Host	Warning	Port RX Data > 0 Gb
Gateway	Minor	Port RX Data > 140 Gb
Switch	Critical	
Router		

⚠ Note how the added conditions are listed in the Network Analysis section, if Link Analysis is enabled, and they are colored accordingly.

View

Properties

Link 1

Link/Port Properties




Property	Source	Destination
System GUID	0x0002c903007b78b0	0xb8599f0300fc6de4
Port	1	3
MTU	4096	4096
Width	4X	4X
Speed	FDR	FDR
Port RX Data	20379.85 Gb	5.9 Gb
Port TX Data	18.05 Gb	6134.55 Gb
Port RX Data Rate	0 Gb/s	0 Gb/s
Port TX Data Rate	0 Gb/s	0 Gb/s
Port RX Packets	1285841763 Packets	7796207 Packets
Port TX Packets	22720574 Packets	386937725 Packets
Port RX Packets Rate	2.9 Packets/s	2.9 Packets/s
Port TX Packets Rate	2.9 Packets/s	2.9 Packets/s

Cable Info



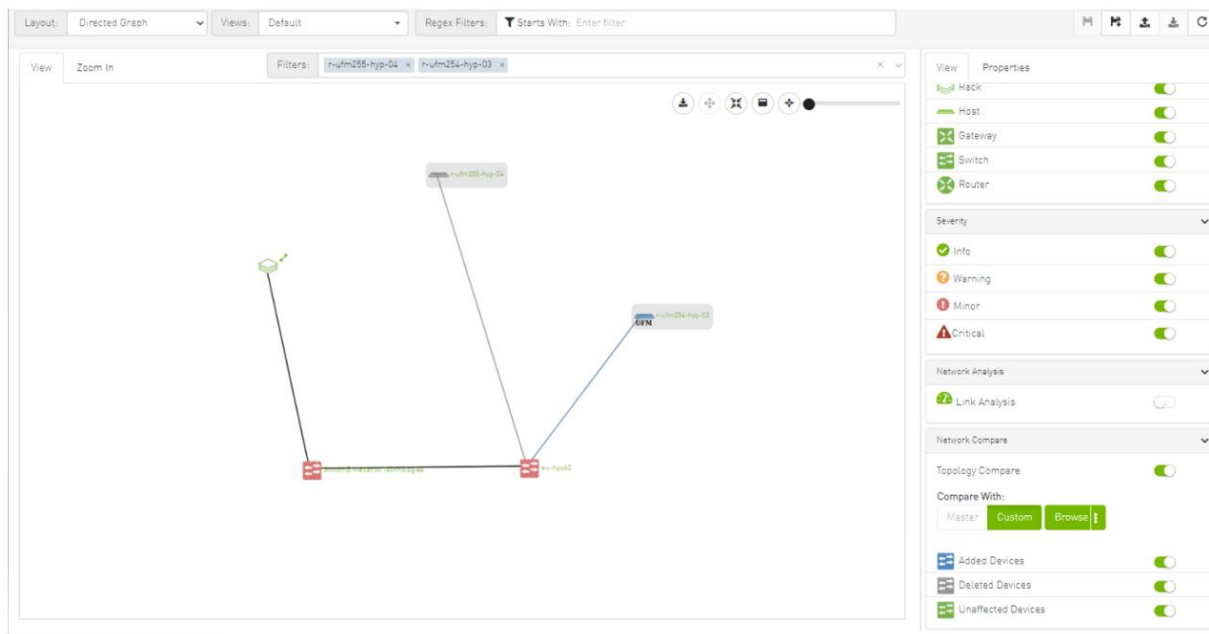
Property	Value
Part Number	MCP1600-E001
Length	1 m
Serial Number	MT1625VS05686
Identifier	QSFP+
Technology	Copper cable- unequalized
Revision	A2

 Notice how the monitored counter is presented in boldface, and the background color is presented with the worst matched condition.

Please note that if the current layout and view are saved, the defined conditions are saved inside the view being saved.

9.2.9 Topology Compare

It is possible to enable the [Topology Compare](#) feature from the View tab in the right-hand pane. When the radio button is enabled, it is possible to compare the current topology with the master topology or with a custom topology whose `.topo` file you may upload.

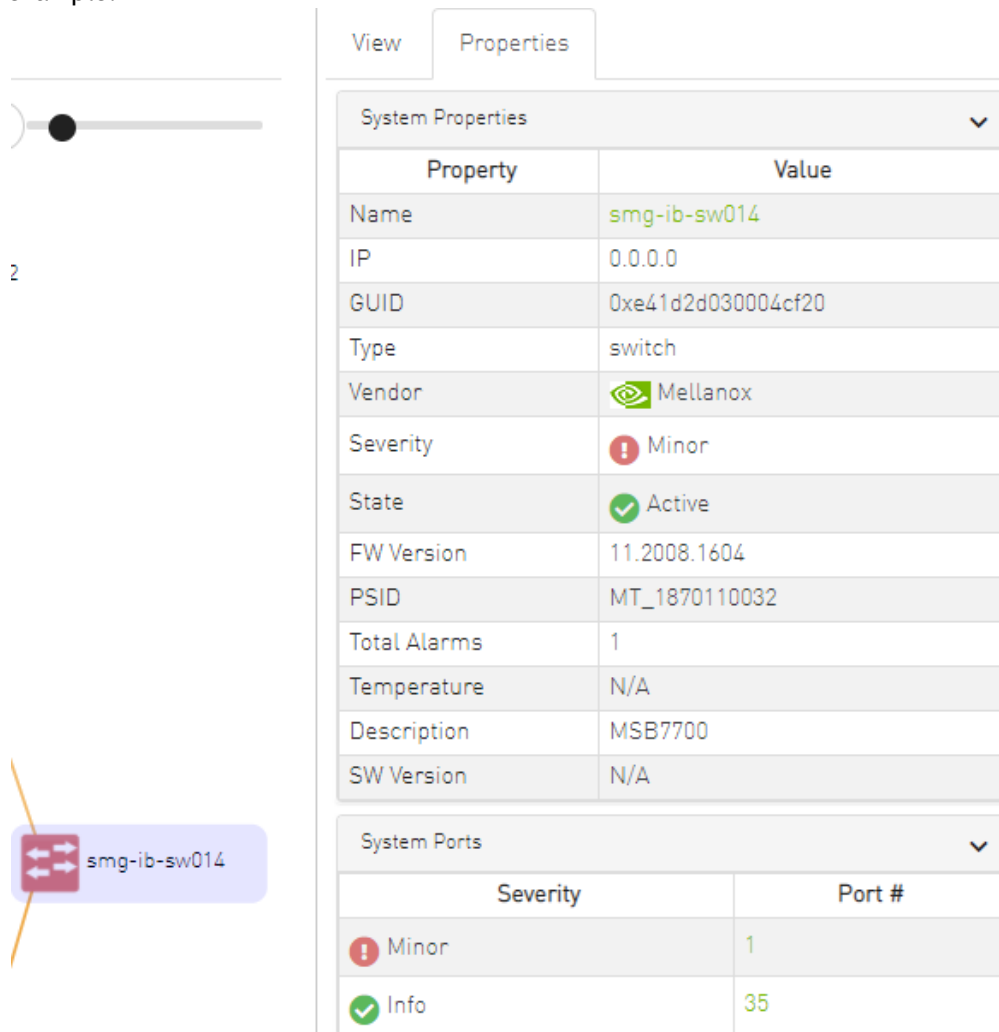


Topology compare key:

- A blue node signifies an added node
- A gray host signifies a deleted node
- A gray and black line signifies that some links were deleted and others were unchanged
- A gray and blue line signifies that some links were deleted, and others were added
- A gray, blue, and black line signifies that some links were deleted, some were added, and some were unchanged
- A blue and black line signifies that some links were added, and some were unchanged




9.2.10 Properties Tab

- Provides details on a specific system selected from the map, as shown in the following example:





View Properties

System Properties

Property	Value
Name	smg-ib-sw014
IP	0.0.0.0
GUID	0xe41d2d030004cf20
Type	switch
Vendor	 Mellanox
Severity	 Minor
State	 Active
FW Version	11.2008.1604
PSID	MT_1870110032
Total Alarms	1
Temperature	N/A
Description	MSB7700
SW Version	N/A

System Ports

Severity	Port #
 Minor	1
 Info	35

- Provides link/port properties and cable info on a specific link selected from the map, including destination and source ports, as shown in the following example:

View Properties

Link 1

Collect System Dump

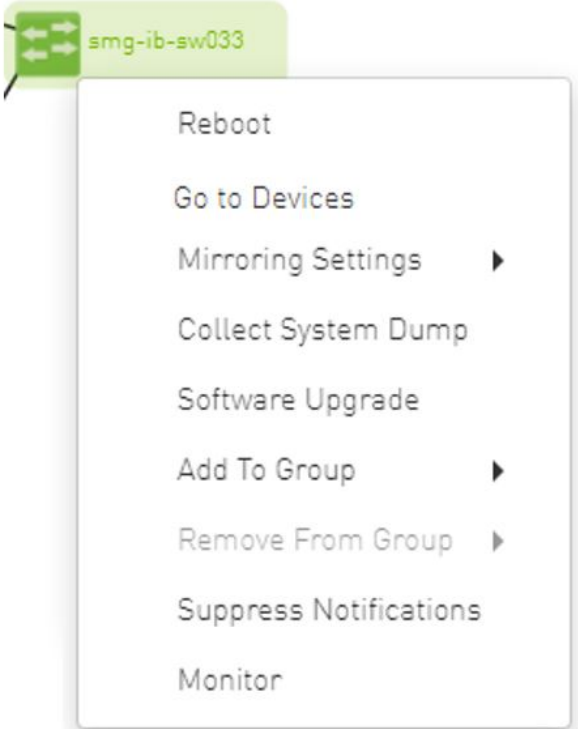
Link/Port Properties		
Property	Source	Destination
System GUID	0x0008f105002020fb	0x248a070300f88fe0
Port	18	1
MTU	4096	4096
Width	4X	4X
Speed	EDR	EDR
Port RX Data	614 MB	164 MB
Port TX Data	164 MB	614 MB
Port RX Data Rate	0 MB/s	0 MB/s
Port TX Data Rate	0 MB/s	0 MB/s
Port RX Packets	1662888 Packets	597647 Packets
Port TX Packets	597646 Packets	1662723 Packets
Port RX Packets Rate	0.45 Packets/s	0.25 Packets/s
Port TX Packets Rate	0.25 Packets/s	0.45 Packets/s

Cable Info	
Property	Value
Part Number	MCP1600-E00A
Length	1 m
Serial Number	MT1714VS00778
Identifier	QSFP+
Technology	Copper cable- unequalized
Revision	A2

9.2.11 Network Map Elements Actions

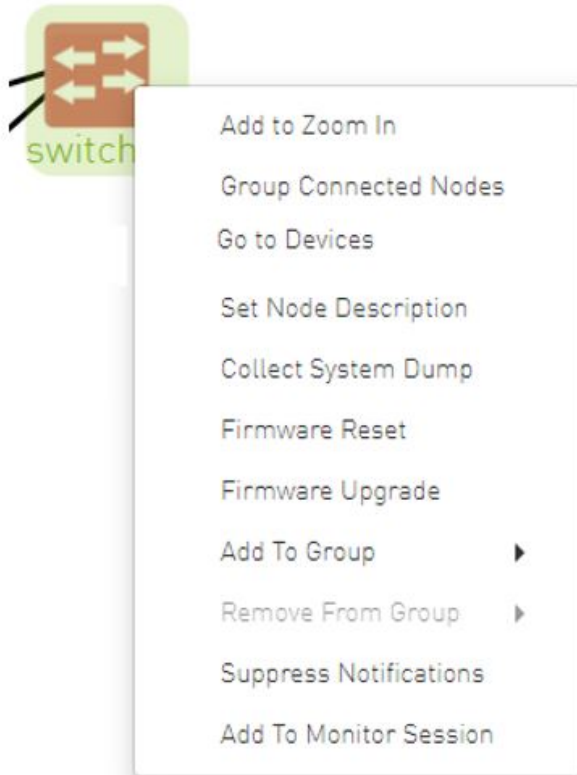
In the Network Map, a right-click on any of the elements enables performing a set of actions depending on the element type and its capabilities. See the list of available actions for each element type in the tables below.

9.2.11.1 Supported Actions for Internally Managed Switches



Element Type	Supported Actions	Description
Managed Switch	Reboot	Reboot the switch software
	Mirroring Settings	Set the mirroring configuration for the switch
	Collect System Dump	Collect system dump from the device
	Software Upgrade	Perform switch software upgrade
	Add to Group	Add switch to logical group
	Remove from Group	Remove switch from logical group
	Suppress Notification	Suppress all event notifications for the switch
	Monitor	Configure and activate switch monitoring
	Go to Devices	Go to devices page and select the device

9.2.11.2 Supported Actions for Externally Managed Switches



Element Type	Supported Actions	Description
Externally Managed Switch	Set Node Description	Sets description for specific node
	Firmware Reset	Perform switch firmware reset
	Firmware Upgrade	Perform switch firmware upgrade
	Add to Group	Add switch to logical group
	Remove from Group	Remove switch from logical group
	Suppress Notification	Suppress all event notifications for the switch
	Monitor	Configure and activate switch monitoring
	Go To Devices	Go to devices page and select the device


9.2.11.3 Supported Actions for Hosts



Element Type	Supported Actions	Description
Hosts	Firmware Upgrade	Perform switch firmware upgrade
	Add to Group	Add host to logical group
	Remove from Group	Remove host from logical group
	Suppress Notification	Suppress all event notifications for the host
	Monitor	Configure and activate host monitoring

9.3 Managed Elements

The UFM Managed Elements window allows you to obtain information on the fabric physical elements, such as devices, ports and cables.

 All information provided in a tabular format in UFM web UI can be exported into a CSV file.

- [Devices Window](#)
- [Ports Window](#)
- [Virtual Ports Window](#)
- [Unhealthy Ports Window](#)
- [Cables Window](#)
- [Groups Window](#)
- [Inventory Window](#)
- [PKeys Window](#)
- [HCAs Window](#)

9.3.1 Devices Window

The Devices window shows data pertaining to the physical devices in a tabular format.

Severity	Name	GUID	Type	Model	IP	Firmware Version
Minor	r-dmz-ufm-sw49	0x0002c903007b78b0	switch	SX6036	fcfc.fcfc:209.36:202.e...	9.4.5110
Minor	r-ufm-sw95	0xb8599f0300fc6de4	switch	MQM8700	fcfc.fcfc:209.36:ba59...	27.2022.612
Info	r-dmz-ufm134	0x1070fd03000b22f8	host		192.168.1.153	22.34.282
Info	r-dcs96	0x1070fd030071ea4e	host		0.0.0.0	20.31.1014
Info	r-dmz-ufm131	0x1070fd03000b22c4	host		0.0.0.0	22.34.282
Info	r-dmz-ufm137	0x1070fd03000b22cc	host		0.0.0.0	22.32.1062
Info	r-dmz-ufm128	0xe41d2d03005cf34c	host		0.0.0.0	12.22.252

Devices Window Data


Data Type	Description
Health	Health of the device reflecting the highest alarm severity. Please refer to the Health States table.
Name	Name of the device <div style="border: 1px solid yellow; padding: 5px; margin-top: 10px;"> If UFM Agent is running on a device, the following icon will appear next to the device name: </div>
GUID	System GUID of the device
Type	Type of the device: switch, node, IB router, and getaway
IP	IP address of the device
Vendor	The vendor of the device
Firmware Version	The firmware version installed on the device

Health States

Icon	Name	Description
	Normal	Information/notification displayed during normal operating state or a normal system event.
	Critical	Critical means that the operation of the system or a system component fails.
	Minor	Minor reflects a problem in the fabric with no failure.

Icon	Name	Description
	Warning	Warning reflects a low priority problem in the fabric with no failure. A warning is asserted when an event exceeds a predefined threshold.

A right-click on the device name displays a list of actions that can be performed on it.



Devices Actions

Action	Description
Firmware Upgrade	Perform a firmware upgrade on the selected device
Firmware Reset	Reboot the device. This action is only applicable to unmanaged hosts (servers).
Set Node Description	Configure a description to this node
Collect System Dump	Collect the system dump log for a specific device
Add to Group	Add the selected device to a devices group
Remove from Group	Remove the selected device from a devices group
Suppress Notifications	Suppress all event notifications for the device
Add to Monitor Session	Configure and activate host monitoring
Show in Network Map	Move to Zoom In tab in network map and add the selected device to filter list

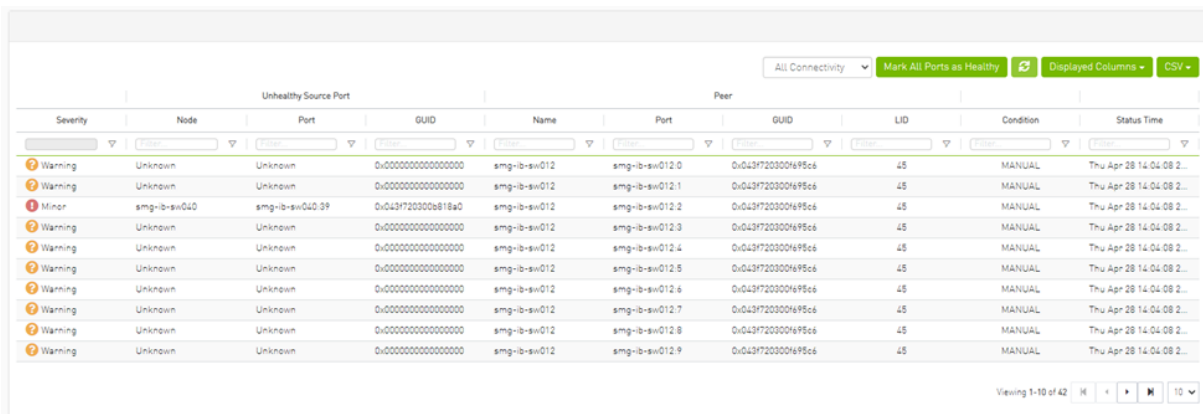

! Collecting system dump for hosts, managed by UFM, is available only for hosts which are set with a valid IPv4 address and installed with MLNX_OFED.

9.3.1.1 Mark Device as Unhealthy

From the Devices table, it is possible to mark devices as healthy or unhealthy using the context menu (right-click).

There are two options for marking a device as unhealthy:

- Isolate
- No Discover



Server: `conf/opensm/opensm-health-policy.conf` content:

```

0xe41d2d030003e3b0 34 UNHEALTHY isolate
0xe41d2d030003e3b0 19 UNHEALTHY isolate
0xe41d2d030003e3b0 3 UNHEALTHY isolate
0xe41d2d030003e3b0 26 UNHEALTHY isolate
0xe41d2d030003e3b0 0 UNHEALTHY isolate
0xe41d2d030003e3b0 27 UNHEALTHY isolate
0xe41d2d030003e3b0 7 UNHEALTHY isolate
0xe41d2d030003e3b0 10 UNHEALTHY isolate
0xe41d2d030003e3b0 11 UNHEALTHY isolate
0xe41d2d030003e3b0 22 UNHEALTHY isolate
0xe41d2d030003e3b0 18 UNHEALTHY isolate
0xe41d2d030003e3b0 29 UNHEALTHY isolate
0xe41d2d030003e3b0 8 UNHEALTHY isolate
0xe41d2d030003e3b0 5 UNHEALTHY isolate
0xe41d2d030003e3b0 17 UNHEALTHY isolate
0xe41d2d030003e3b0 23 UNHEALTHY isolate
0xe41d2d030003e3b0 15 UNHEALTHY isolate
0xe41d2d030003e3b0 24 UNHEALTHY isolate
0xe41d2d030003e3b0 2 UNHEALTHY isolate
0xe41d2d030003e3b0 16 UNHEALTHY isolate
0xe41d2d030003e3b0 13 UNHEALTHY isolate
0xe41d2d030003e3b0 14 UNHEALTHY isolate
0xe41d2d030003e3b0 32 UNHEALTHY isolate
0xe41d2d030003e3b0 33 UNHEALTHY isolate
0xe41d2d030003e3b0 35 UNHEALTHY isolate
0xe41d2d030003e3b0 20 UNHEALTHY isolate
0xe41d2d030003e3b0 21 UNHEALTHY isolate
0xe41d2d030003e3b0 28 UNHEALTHY isolate
0xe41d2d030003e3b0 1 UNHEALTHY isolate
0xe41d2d030003e3b0 9 UNHEALTHY isolate
0xe41d2d030003e3b0 4 UNHEALTHY isolate
0xe41d2d030003e3b0 31 UNHEALTHY isolate
0xe41d2d030003e3b0 30 UNHEALTHY isolate
0xe41d2d030003e3b0 36 UNHEALTHY isolate
0xe41d2d030003e3b0 12 UNHEALTHY isolate
0xe41d2d030003e3b0 25 UNHEALTHY isolate
0xe41d2d030003e3b0 6 UNHEALTHY isolate

```

`/opt/ufm/files/log/opensm-unhealthy-ports.dump` content:



9.3.1.2 Mark Device as Healthy

The screenshot shows a table with columns: S..., Name, GUID, Type, Model, IP, and Firmware Ve... The table contains several rows of device information. The row for 'r-dcs96' is highlighted in black, and a context menu is open over it. The menu options are: Copy Cell, Mark As Healthy, Firmware Upgrade, Show In Network Map, Add To Group, Remove From Group, Suppress Notifications, and Add To Monitor Session. The 'Mark As Healthy' option is highlighted in the menu.

Server `/opt/ufm/files/conf/opensm/opensm-health-policy.conf` content:

```

0xe41d2d030003e3b0 15 HEALTHY
0xe41d2d030003e3b0 25 HEALTHY

```



```

0xe41d2d030003e3b0 35 HEALTHY
0xe41d2d030003e3b0 0 HEALTHY
0xe41d2d030003e3b0 11 HEALTHY
0xe41d2d030003e3b0 21 HEALTHY
0xe41d2d030003e3b0 28 HEALTHY
0xe41d2d030003e3b0 7 HEALTHY
0xe41d2d030003e3b0 17 HEALTHY
0xe41d2d030003e3b0 14 HEALTHY
0xe41d2d030003e3b0 24 HEALTHY
0xe41d2d030003e3b0 34 HEALTHY
0xe41d2d030003e3b0 3 HEALTHY
0xe41d2d030003e3b0 10 HEALTHY
0xe41d2d030003e3b0 20 HEALTHY
0xe41d2d030003e3b0 31 HEALTHY
0xe41d2d030003e3b0 6 HEALTHY
0xe41d2d030003e3b0 16 HEALTHY
0xe41d2d030003e3b0 27 HEALTHY
0xe41d2d030003e3b0 2 HEALTHY
0xe41d2d030003e3b0 13 HEALTHY
0xe41d2d030003e3b0 23 HEALTHY
0xe41d2d030003e3b0 33 HEALTHY
0xe41d2d030003e3b0 30 HEALTHY
0xe41d2d030003e3b0 9 HEALTHY
0xe41d2d030003e3b0 19 HEALTHY
0xe41d2d030003e3b0 26 HEALTHY
0xe41d2d030003e3b0 36 HEALTHY
0xe41d2d030003e3b0 5 HEALTHY
0xe41d2d030003e3b0 12 HEALTHY
0xe41d2d030003e3b0 22 HEALTHY
0xe41d2d030003e3b0 32 HEALTHY
0xe41d2d030003e3b0 1 HEALTHY
0xe41d2d030003e3b0 8 HEALTHY
0xe41d2d030003e3b0 18 HEALTHY
0xe41d2d030003e3b0 29 HEALTHY
0xe41d2d030003e3b0 4 HEALTHY

```

`/opt/ufm/files/log/opensm-unhealthy-ports.dump` content:

```
# NodeGUID, PortNum, NodeDesc, PeerNodeGUID, PeerPortNum, PeerNodeDesc, {BadCond1, BadCond2, ...}, timestamp
```

9.3.1.3 Upgrading Software and Firmware for Hosts and Externally Managed Switches

9.3.1.3.1 Software/Firmware Upgrade via FTP

Software and firmware upgrade over FTP is enabled by the UFM Agent. UFM invokes the Software/Firmware Upgrade procedure locally on switches or on hosts. The procedure copies the new software/firmware file from the defined storage location and performs the operation on the device. UFM sends the set of attributes required for performing the software/firmware upgrade to the agent.

The attributes are:

- File Transfer Protocol - default FTP
 - The Software/Firmware upgrade on InfiniScale III ASIC-based switches supports FTP protocol for transmitting files to the local machine.
 - The Software/Firmware upgrade on InfiniScale IV-based switches and hosts supports TFTP and protocols for transmitting files to the local machine.
- IP address of file-storage server
- Path to the software/firmware image location

The software/firmware image files should be placed according to the required structure under the defined image storage location. Please refer to section [Devices Window](#).
- File-storage server access credentials (User/Password)

9.3.1.3.2 In-Band Firmware Upgrade

You can perform in-band firmware upgrades for externally managed switches and HCAs. This upgrade procedure does not require the UFM Agent or IP connectivity, but it does require current PSID recognition. Please refer to section [PSID and Firmware Version In-Band Discovery](#). This feature requires that the Mellanox Firmware Toolkit (MFT), which is included in the UFM package, is installed on the UFM server. UFM uses flint from the MFT for in-band firmware burning.

Before upgrading, you must create the firmware repository on the UFM server under the directory `/opt/ufm/files/userdata/fw/`. The subdirectory should be created for each PSID and one firmware image should be placed under it. For example:

```
/opt/ufm/files/userdata/fw/  
  MT_0D80110009  
    fw-ConnectX2-rel-2_9_1000-MHQH29B-XTR_A1.bin  
  MT_0F90110002  
    fw-IS4-rel-7_4_2040-MIS5023Q_A1-A5.bin
```

9.3.1.3.3 Directory Structure for Software or Firmware Upgrade Over FTP

Before performing a software or firmware upgrade, you must create the following directory structure for the upgrade image. The path to the `<ftp user home>/<path>/` directory should be specified in the upgrade dialog box.

```
<ftp user home>/<path>/  
  InfiniScale3 - For anafa based switches Software/Firmware upgrade images  
    voltaire_fw_images.tar - firmware image file  
    ibswmpr-<S/w version>.tar - software image file  
  InfiniScale4 - For InfiniScale IV based switches Software/Firmware upgrade images  
    firmware_2036_4036.tar - Firmware image file  
    upgrade_2036_4036.tgz - Software image file  
  OFED /* For host SW upgrade*/  
    OFED-<OS label>.tar.bz2  
  <PSID>* - For host FW upgrade  
    fw_update.img
```

The `<PSID>` value is extracted from the `mstflint` command:

```
mstflint -d <device> q
```

The device is extracted from the `lspci` command. For example:

```
# lspci  
06:00.0 InfiniBand: Mellanox Technologies MT25208 InfiniHost III Ex  
# mstflint -d 06:00.0 q | grep PSID  
PSID: VLT0040010001
```

9.3.1.3.4 PSID and Firmware Version In-Band Discovery

The device PSID and device firmware version are required for in-band firmware upgrade and for the correct functioning of Subnet Manager plugins, such as Congestion Control Manager and Lossy Configuration Management. For most devices, UFM discovers this information and displays it in the Device Properties pane. The PSID and the firmware version are discovered by the Vendor-specific MAD.

By default, the `gv.cfg` file value for `event_plugin_option` is set to `(null)`. This means that the plugin is disabled and `opemsm` does not send MADs to discover devices' PSID and FW version. Therefore, values for devices' PSID and FW version are taken from `ibdiagnet` output (section `NODES_INFO`).

The below is an example of the default value:

```
event_plugin_options = (null)
```

To enable the vendor-specific discovery by `opemsm`, in the `gv.cfg` configuration file, change the value of `event_plugin_option` to `(--vendinfo -m 1)`, as shown below:

```
event_plugin_options = --vendinfo -m 1
```

If the value is set to `--vendinfo -m 1`, the data should be supplied by `opemsm`, and in this case the `ibdiagnet` output is ignored.



In some firmware versions, the information above is currently not available.

9.3.1.3.5 Switch Management IP Address Discovery

From NVIDIA switch FM version 27.2010.3942 and up, NVIDIA switches support switch management IP address discovery using MADs. This information can be retrieved as part of `ibdiagnet run` (`ibdiagnet` output), and assigned to discover switches in UFM.

There is an option to choose the IP address of which IP protocol version that is assigned to the switch: IPv4 or IPv6.

The `discovered_switch_ip_protocol` key, located in the `gv.cfg` file in section `[FabricAnalysis]`, is set to 4 by default. This means that the IP address of type IPv4 is assigned to the switch as its management IP address. In case this value is set to 6, the IP address of type IPv6 is assigned to the switch as its management IP address.

After changing the `discover_switch_ip_protocol` value in `gv.cfg`, the UFM Main Model needs to be restarted for the update to take effect. The discovered IP addresses for switches are not persistent in UFM - every UFM Main Model restarts the values of management IP address which is assigned from the `ibdiagnet` output.

9.3.1.3.6 Upgrading Server Software

The ability to update the server software is applicable only for hosts (servers) with the UFM Agent.

To upgrade the software:

1. Select a device.
2. From the right-click menu, select Software Update.
3. Enter the parameters listed in the following table.

Parameter	Description
Protocol	Update is performed via FTP protocol
IP	Enter the host IP
Path	Enter the parent directory of the FTP directory structure for the Upgrade image. The path should not be an absolute path and should not contain the first slash (/) or trailer slash.
User	Name of the host username
Password	Enter the host password

4. Click Submit to save your changes.

9.3.1.3.7 Upgrading Firmware

You can upgrade firmware over FTP for hosts and switches that are running the UFM Agent, or you can perform an in-band upgrade for externally managed switches and HCAs.

Before you begin the upgrade ensure that the new firmware version is in the correct location. For more information, please refer to section [In-Band Firmware Upgrade](#).

To upgrade the firmware:

1. Select a host or server.
2. From the right-click menu, select Firmware Upgrade.
3. Select protocol In Band.
4. For upgrade over FTP, enter the parameters listed in the following table.

Parameter	Description
IP	Enter device IP
Path	Enter the parent directory of the FTP directory structure for the Upgrade image. The path should not be an absolute path and should not contain the first slash (/) or trailer slash.
Username	Name of the host username
Password	Enter the host password

5. Click submit to save your changes.



The firmware upgrade takes effect only after the host or externally managed switch is restarted.

9.3.1.3.8 Upgrade Cables Transceivers Firmware Version

The main purpose of this feature is to add support for burning of multiple cables transceiver types on multiple devices using linkx tool which is part of flint. This needs to be done from both ends of the cable (switch and HCA/switch).

To upgrade cables transceivers FW version:

1. Navigate to managed elements page
2. select the target switches and click on Upgrade Cable Transceivers option

S...	Name	GUID	Type	Model	IP	Firmware Ve...
✓	smg-ib-sim001	0xb8599f0300c...	host		0.0.0.0	18.32.524
✓	smg-ib-svr031	0x98039b0300...	host		0.0.0.0	20.31.2006
✓	smg-ib-apl022...	0x98039b0300...	host		0.0.0.0	20.32.1010
?	smg-ib-svr032	0x1070fa03007...	host		0.0.0.0	28.33.810
!	smg-ib-sw...	0x98039b0300...	switch	MQM8700	10.209.24.136	27.2000.2046
!	smg-ib-olg...			CS7520	10.209.27.99	mismatched
?	smg-ib-sw...			MQM9700	10.209.24.121	31.2010.2036
!	smg-ib-sw...			MQM8700	10.209.24.10	27.2010.2010
!	smg-ib-sw...			MQM8700	10.209.24.57	27.2010.1202
!	smg-ib-sw...			MSB7700	10.209.27.36	11.2008.3328


3. A model will be shown containing list of the active firmware versions for the cables of the selected switches, besides the version number, a badge will show the number of matched switches:

Upgrade Cable Transceivers

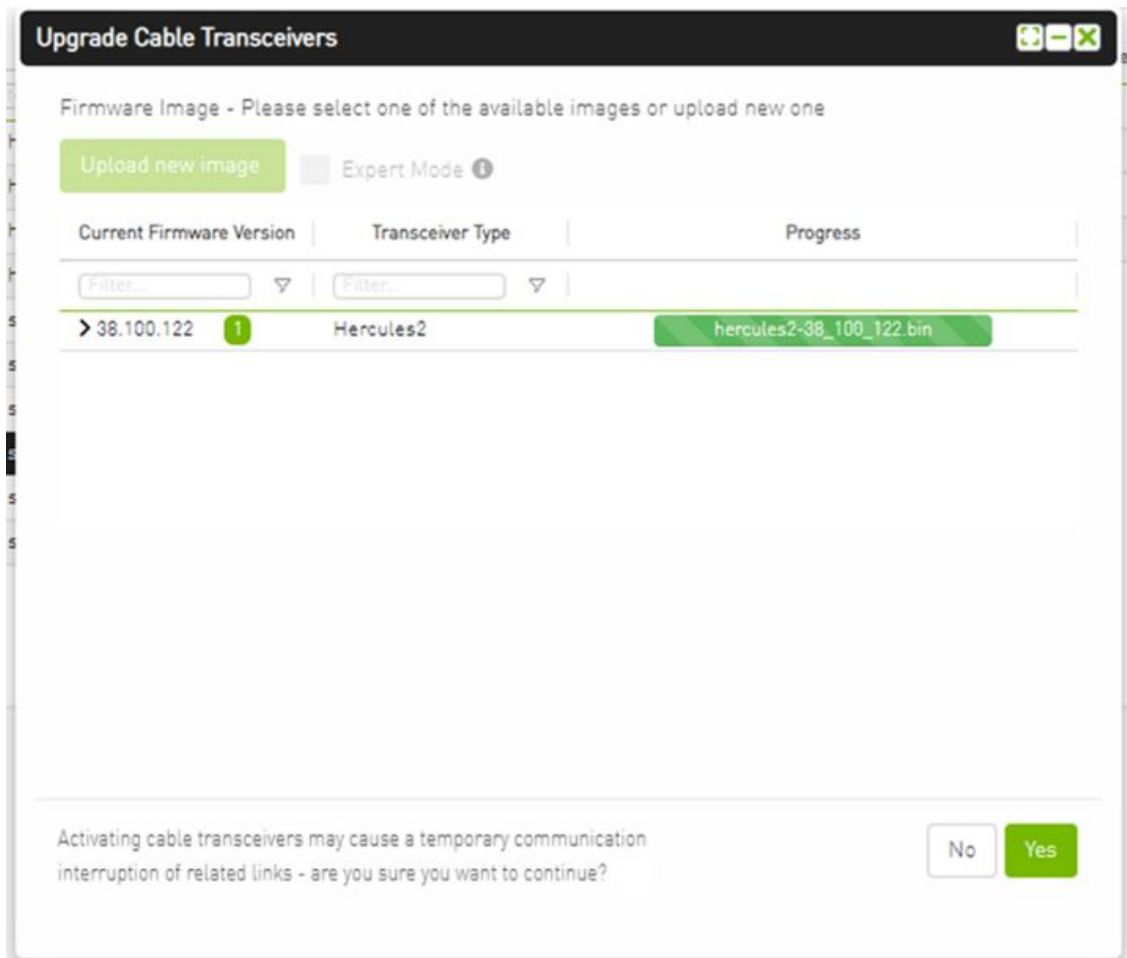
Firmware Image - Please select one of the available images or upload new one

Expert Mode ⓘ

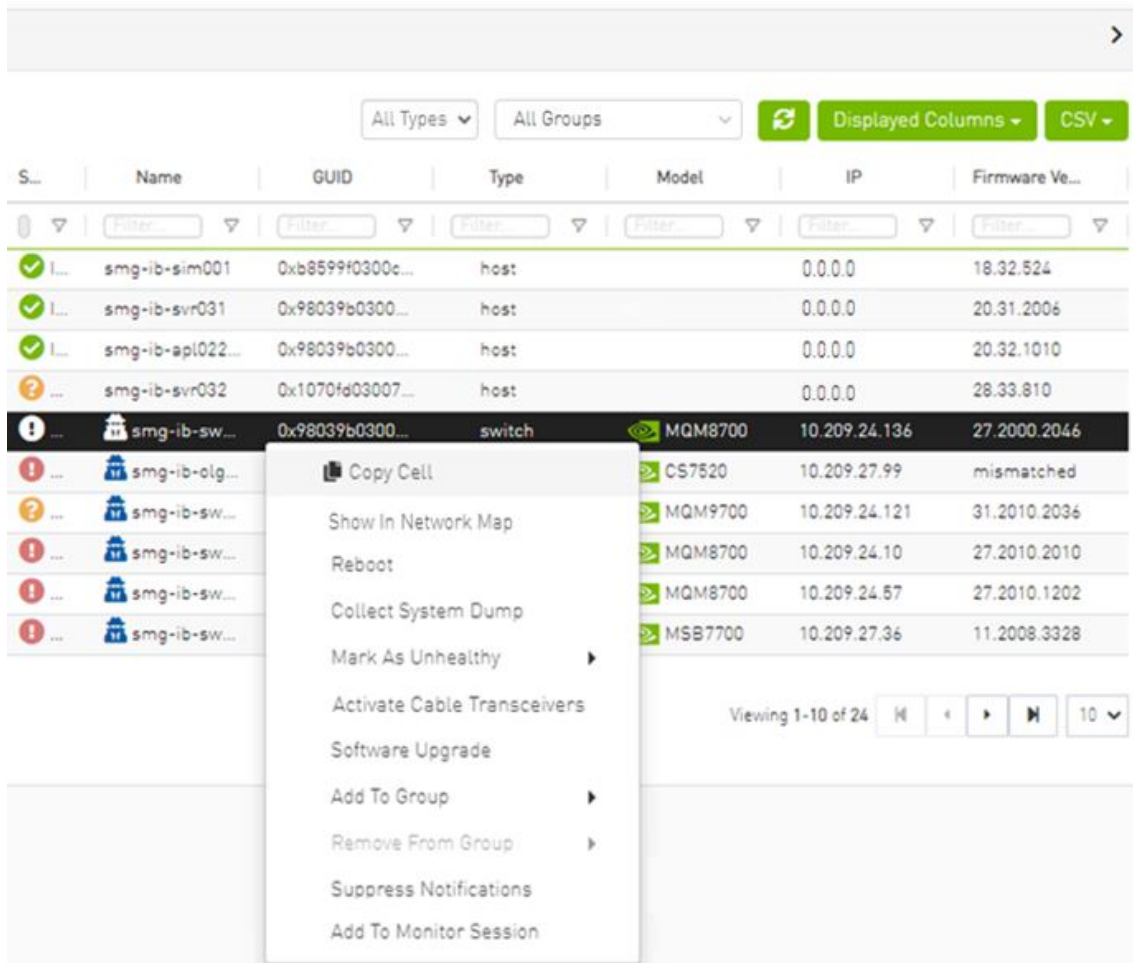
Current Firmware Version	Transceiver Type	Image
<input type="text" value="Filter..."/> ▾	<input type="text" value="Filter..."/> ▾	<input type="text" value="Filter..."/> ▾
> 38.100.122 1	Hercules2	No Selected Image ▾



4. After the user clicks Submit, the GUI will start sending the selected binaries with the relevant switches sequentially, and a model with a progress bar will be shown (this model can be minimized):




5. After the whole action is completed successfully, you will be able to see the following message at the model bottom The upgrade cable transceivers completed successfully, do you want to activate it? by clicking the yes button it will run a new action on all the burned devices to activate the new uploaded binary image.
6. Another option to activate burned cables transceivers you can go to the Groups page and right click on the predefined Group named Devices Pending FW Transceivers Reset or you can right click on the upgraded device from managed element page and select Activate cable Transceivers action.



9.3.1.4 Device Information Tabs

Selecting a device from the Devices table reveals the Device Information table on the right side of the screen. This table provides information on the device's ports, cables, groups, events, alarms, inventory, and device access.



9.3.1.4.1 General Tab

Provides general information on the selected device.

0xb8599f0300fc6de4 - Device Information

General Ports Cables Groups Alarms Events Inventory Device Access

Property	Value
Name	r-ufm-sw95
Type	switch
IP	fcfc:fcfc:209:36:ba59:9fff:fef6:7db4
Model	MQM8700
Up Time	92d 0h 30m 50.368s

9.3.1.4.2 Ports Tab

This tab provides a list of the ports connected to this device in a tabular format.

0x98039b0300a8b71e - Device Information

General Ports Cables Groups Alarms Events Inventory Device Access


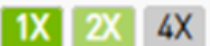
Active Displayed Columns CSV

Severity	State	System	Port	LID	Peer Node
Info	✓	smg-ib-sw032	3	5	smg-ib-sw036
Minor	✓	smg-ib-sw032	5	5	smg-ib-sw036
Info	✓	smg-ib-sw032	16	5	smg-ib-sw056

Viewing 1-3 of 3

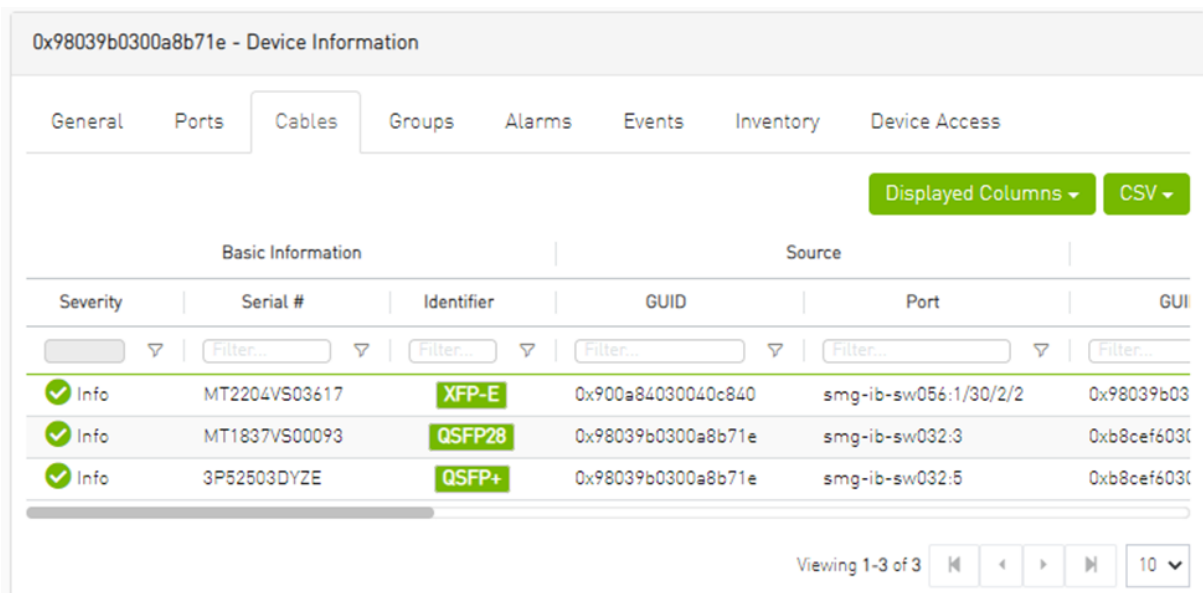
Ports Data

Data Type	Description
Port Number	The number of ports on device.

Data Type	Description
Node	The node name/GUID/IP that the port belongs to. Note that you can choose the node label (name/GUID/IP) using the drop-down menu available above the Ports data table.
Health	Health of the port reflecting the highest alarm severity. Please refer to the Health States table.
State	Indicates whether the port is connected (active or inactive).
LID	The local identifier (LID) of the port.
MTU	Maximum Transmission Unit of the port.
Speed 	Lists the highest value of active, enabled and supported speeds in icons indicating their status: <ul style="list-style-type: none"> Dark green - active speed Light green - enabled speed Grey - supported yet disabled speed
Width 	Lists the highest value of active, enabled and supported widths in icons indicating their status: <ul style="list-style-type: none"> Dark green - active width Light green - enabled width Grey - supported yet disabled width
Peer	The GUID of the device the port is connected to.
Peer Port	The name of the port that is connected to this port.

9.3.1.4.3 Cables Tab

This tab provides a list of the cables connected to this device in a tabular format.



Cables Data

Data Type	Description
Basic Information	

Health	Health of the cable reflecting the highest alarm severity. Please refer to the Health States table.
Serial Number	Serial number of the cable.
Identifier	Identifier of the cable.
Source Port Information	
Source GUID	GUID of the source port the cable is connected to.
Source Port	The number of the source port the cable is connected to.
Destination Port Information	
Destination GUID	GUID of the destination port the cable is connected to.
Destination Port	The number of the destination port the cable is connected to.
Advanced Information	
Revision	Revision of the cable.
Link Width	The maximum link width of the cable.
Part Number	Part number of the cable.
Technology	The transmitting medium of the cable: copper/optical/etc.
Length	The cable length in meters.

9.3.1.4.4 Groups Tab

This tab provides a list of the groups to which the selected device belongs.

The screenshot displays the 'Groups' tab for a device with ID 0x98039b0300a8b71e. The interface includes several tabs: General, Ports, Cables, Groups (selected), Alarms, Events, Inventory, and Device Access. Below the tabs, there are controls for 'All' (dropdown), 'Displayed Columns' (button), and 'CSV' (button). The main content area is a table with the following columns: Severity, Name (with an upward arrow), Description, and Type. Each cell in the table has a 'Filter...' dropdown. The table lists three groups:

Severity	Name ↑	Description	Type
⚠ Critical	1U Switches	Includes all 1U Switches that exi...	General
⚠ Critical	Alarmed Devices	Devices with alarms	General
⚠ Critical	Switches	Includes all Switches that exist i...	General

At the bottom right, there is a pagination control showing 'Viewing 1-3 of 3' and navigation buttons (back, forward, first, last) along with a page size dropdown set to '10'.

Groups Data

Data Type	Description
Severity	Aggregated severity level of the group (the highest severity level of all group members).
Name	Name of the group.
Description	Description of the group.
Type	Type of the group: General/Rack.

9.3.1.4.5 Alarms Tab

This tab provides a list of all UFM alarms related to the selected device.

0x043f720300b818a0 - Device Information

General Ports Cables Groups **Alarms** Events Inventory Device Access

Clear All Alarms Refresh Displayed Columns CSV

Severity	Date/Time ↓	Source	Reason	C
Minor	2022-04-28 14:28:46	default(12) / Switch: smg-ib-s	Found a [50.0] link that oper...	26
Warning	2022-04-28 14:09:55	default(12) / Switch: smg-ib-s	Peer Port Mellanox Technol...	1
Critical	2022-04-28 14:08:24	default(12) / Switch: smg-ib-s	smg-ib-sw040: (system guid...	5
Warning	2022-04-28 14:04:48	default(12) / Switch: smg-ib-s	Peer Port smg-ib-sw012:2 is...	1

Viewing 1-4 of 4

Alarms Data

Data Type	Description
Alarms ID	Alarm identifier.
Source	Source object (device/port) on which the alarm was triggered.
Severity	The severity of the alarm.
Description	Description of the alarm.
Date/Time	The time when the alarm was triggered.
Reason	Reason for the alarm.
Count	Number of instances that the alarm occurred on the related source object.

9.3.1.4.6 Events Tab

This tab provides a list of the UFM events that are related to the selected device.

0x043f720300b818a0 - Device Information

General Ports Cables Groups Alarms **Events** Inventory Device Access

Clear All Events Refresh Displayed Columns CSV

Severity	Date/Time ↓	Source	Source Type	Descri
Info	2022-04-28 14:16:42	default(12) / Switch: smg-ib-s	Switch	Action reboot on
Info	2022-04-28 14:10:13	default(12) / Switch: smg-ib-s	Switch	System Image G
Info	2022-04-28 14:10:13	default(12) / Switch: smg-ib-s	Switch	Capability Mask
Info	2022-04-28 14:09:24	default(12) / Switch: smg-ib-s	Switch	smg-ib-sw040: [
Warning	2022-04-28 14:08:24	Source 043f720300b818a0_39	Link	Link went down:
Warning	2022-04-28 14:08:24	Source 043f720300b818a0_41	Link	Link went down:
Info	2022-04-28 14:07:41	default(12) / Switch: smg-ib-s	Switch	Action reboot st:
Info	2022-04-28 14:04:14	default(12) / Switch: smg-ib-s	Switch	Switch Upgrade
Info	2022-04-28 14:02:42	default(12) / Switch: smg-ib-s	Switch	Switch SW upgrn
Info	2022-04-28 14:02:42	default(12) / Switch: smg-ib-s	Switch	Action sw_upgrs


Viewing 1-10 of 11

Events Data

Data Type	Description
Severity	Event severity - Info, Warning, Error, Critical or Minor.
Event Name	The name of the event.
Source	The source object (device/port) on which the event was triggered.
Date/Time	The time when the event was triggered.
Category	The category of the event indicated by icons. Hovering over the icon will display the category name.
Description	Description of the event. Full description can be displayed by hovering over the text.

9.3.1.4.7 Inventory Tab

This tab provides a list of the device's modules with information in a tabular format.

 This tab is available for switches only.

0xec0d9a0300b41cd0 - Device Information

General Ports Cables Groups Alarms Events **Inventory** Device Access

Displayed Columns CSV

Severity	Status	Serial Number	System Name	Description	Type	Soft
Info	DC Fault	MT1746X21023	unmanagedEDR	PS - 1	PS	N/A
Info	OK	MT1746X21024	unmanagedEDR	PS - 2	PS	N/A
Info	OK	MT1747X01215	unmanagedEDR	SYSTEM	SYSTEM	N/A
Info	OK	MT1747X00087	unmanagedEDR	FAN - 1	FAN	N/A
Info	OK	MT1747X00087	unmanagedEDR	FAN - 2	FAN	N/A
Info	OK	MT1747X00088	unmanagedEDR	FAN - 3	FAN	N/A
Info	OK	MT1747X00088	unmanagedEDR	FAN - 4	FAN	N/A
Info	OK	MT1747X00101	unmanagedEDR	FAN - 5	FAN	N/A
Info	OK	MT1747X00101	unmanagedEDR	FAN - 6	FAN	N/A
Info	OK	MT1747X00100	unmanagedEDR	FAN - 7	FAN	N/A

Viewing 1-10 of 12

Inventory Data

Data Type	Description
Health	Health of the module reflecting the highest alarm severity. Please refer to the Health States table.
Status	The module status.
Serial Number	Serial number of the module.
Name	Name of the device.
Description	Description of the module.
Type	Type of the module: spine/line/etc.
Firmware Version	Firmware version installed on the module.
Hardware Version	Hardware version of the module.
Temperature	Temperature of the module.

9.3.1.4.8 HCAs Tab

This tab provides a list of the device's HCAs with information in a tabular format.

 This tab is available for hosts only.

0xec0d9a0300bf551c - Device Information

General Ports Cables Groups Alarms Events **HCA's** Device Access

Displayed Columns

Severity	System	Name <input type="text" value="Name"/>	GUID	Type	Port 1	Name <input type="text" value="Name"/>	Port 2
<input type="checkbox"/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>	<input type="text" value="Filter..."/>
<input checked="" type="checkbox"/> Info	smg-ib-svr45		0xec0d9a0300bf551c	ConnectX-5	smg-ib-svr45	HCA-3	smg-ib-
<input checked="" type="checkbox"/> Info	smg-ib-svr45		0x98039b03009ffb22	ConnectX-6	smg-ib-svr45	HCA-1	smg-ib-

Viewing 1-2 of 2

Data Type	Description
Health	Health of the HCA reflecting the highest alarm severity. Please refer to the Health States table.
Name	HCA Index
GUID	HCA GUID
Type	HCA Type
Port GUID	HCA ports GUIDs
PSID	HCA PSID
FW Version	HCA firmware version

9.3.1.4.9 Device Access Tab

This tab allows for managing the access credentials of the selected device for remote accessibility. To be able to set access credentials for the device, a device IP must be set either by installing UFM Agent on the device, or by manually setting the IP under IP Address Settings (IP is now supported with v4 and v6).

0xe41d2d030021d450 - Device Information


General Ports Cables Groups Alarms Events Inventory **Device Access**

IP Address Settings

Mode Auto Manual

Static IP v4 v6

Device Access is not available right now, try enabling ufm agent or set manual IP from IP Address Settings Above

 After manually setting the IP address of NVIDIA® Mellanox® InfiniScale IV® and SwitchX® based switches, UFM will first validate the new IP before setting it.

To edit your device access credentials

1. Select the preferred protocol tab:
 - SSH - allows you to define the SSH parameters to open an SSH session on your device (available for nodes and switches)
 - IPMI - allows you to set the IPMI parameters to open an IPMI session on your device for remote power control (available for nodes only)
 - HTTP - allows you to define the HTTP parameters to open an HTTP session on your device (available for switches only)

2. Click Update to save your changes.

0x98039b0300a8b71e - Device Information

General Ports Cables Groups Alarms Events Inventory **Device Access**

IP Address Settings >

SSH ▾

Credentials

Override Global Settings

User:

Password:

Confirmation:

Connection

Port

Timeout

Manual IP v4


HTTP >

Device Access Credentials Parameters

Field	Description
User	Fill in or edit the computer user name.
Password	Enter the device password.
Confirmation	Enter the device password a second time to confirm.
Manual IP	Enter the device IP address (could be IPv4/IPv6).
Port	Enter the port number.
Timeout	Enter the connection timeout (in seconds) for the device specific protocol (SSH/ HTTP/IPMI).

9.3.1.4.10 Virtual Networking Tab

This tab displays a map containing the HCAs for the selected device, and the ports and virtual ports it is connected to.



9.3.2 Ports Window

Provides a list of all ports in UFM.

All Ports High BER Ports

Active Displayed Columns CSV

Severity	State	System	Name	P. Name	LID	Peer Node	Name	Peer Name	Peer LID	MTU	Speed	Width
Warning	✓	r-hyp-sw-01	1	9	r-ufm254-hyp-01	HCA-1/1	1	4096		SDR	4X	
Info	✓	r-hyp-sw-01	23	9	ufm-host86	HCA-1/1	3	4096		EDR	4X	
Minor	✓	r-hyp-sw-01	36	9	SwitchIB Mellanox Technologies	36	2	4096		FDR	EDR	4X
Info	✓	r-ufm254-hyp-01	HCA-1/1	1	r-hyp-sw-01	1	9	4096		SDR	EDR	4X
Info	✓	r-ufm254-hyp-02	HCA-1/1	10	SwitchIB Mellanox Technologies	1	2	4096		FDR	EDR	4X
Minor	✓	SwitchIB Mellanox Technologies	1	2	r-ufm254-hyp-02	HCA-1/1	10	4096		FDR	EDR	4X
Info	✓	SwitchIB Mellanox Technologies	36	2	r-hyp-sw-01	36	9	4096		FDR	EDR	4X
Info	✓	ufm-host86	HCA-1/1	3	r-hyp-sw-01	23	9	4096		EDR	4X	

Viewing 1-8 of 8

The table can be filtered by port state. The filter contains two options:

- Active - only active ports

- All - all ports

Severity	State	System	Name	P. Name	LID	Peer Node	Name	Peer ...	Peer LID	MTU	Speed	Width	
Warning	✓	r-hyp-sw-01		1	9	r-ufm254-hyp-01	HCA-1/1		1	4096	SDR	4X	
Info	✓	r-hyp-sw-01		23	9	ufm-host86	HCA-1/1		3	4096	EDR	4X	
Minor	✓	r-hyp-sw-01		36	9	SwitchIB Mellanox Technologies		36	2	4096	FDR	EDR	4X
Info	✓	r-ufm254-hyp-01		HCA-1/1	1	r-hyp-sw-01		1	9	4096	SDR	EDR	4X
Info	✓	r-ufm254-hyp-02		HCA-1/1	10	SwitchIB Mellanox Technologies		1	2	4096	FDR	EDR	4X
Minor	✓	SwitchIB Mellanox Technologies		1	2	r-ufm254-hyp-02	HCA-1/1		10	4096	FDR	EDR	4X
Info	✓	SwitchIB Mellanox Technologies		36	2	r-hyp-sw-01		36	9	4096	FDR	EDR	4X
Info	✓	ufm-host86		HCA-1/1	3	r-hyp-sw-01		23	9	4096	EDR	4X	

When right-clicking one of the available ports, the following actions appear:

Severity	State	System	Name	P. Name	LID	Peer Node	Name	Peer ...	Peer LID	MTU	Speed	Width	
Warning	✓	r-hyp-sw-01		1	9	r-ufm254-hyp-01	HCA-1/1		1	4096	SDR	4X	
Info	✓	r-hyp-sw-01		23	9	ufm-host86	HCA-1/1		3	4096	EDR	4X	
Minor	✓	r-hyp-sw-01		36	9	SwitchIB Mellanox Technologies		36	2	4096	FDR	EDR	4X
Info	✓	r-ufm254-hyp-01		HCA-1/1	1	r-hyp-sw-01		1	9	4096	SDR	EDR	4X
Info	✓	r-ufm254-hyp-02		HCA-1/1	10	SwitchIB Mellanox Technologies		1	2	4096	FDR	EDR	4X
Minor	✓	SwitchIB Mellanox Technologies		1	2	r-ufm254-hyp-02	HCA-1/1		10	4096	FDR	EDR	4X
Info	✓	SwitchIB Mellanox Technologies		36	2	r-hyp-sw-01		36	9	4096	FDR	EDR	4X
Info	✓	ufm-host86		HCA-1/1	3	r-hyp-sw-01		23	9	4096	EDR	4X	

Clicking "Cable Information" opens up a window which provides data on operational, module, and troubleshooting information as shown in the following:

Cable Information - 7cfe900300f73be0_1

Operational Info Module Info Troubleshooting Info

Property	Value
Group Opcode	N/A
Recommendation	No issue was observed.
Status Opcode	0

Cable Information - 7cfe900300f73be0_1

Operational Info | **Module Info** | Troubleshooting Info

Property	Value
Vendor Serial Number	MT1515VS07837
Vendor Part Number	MCP1600-E001
Vendor Name	Mellanox
Attenuation (5g,7g,12g) [dB]	4,5,9
Bias Current [mA]	N/A
Cable Technology	Copper cable unequalized
Cable Type	Passive copper cable
CDR RX	N/A
CDR TX	N/A
Compliance	N/A
Digital Diagnostic Monitoring	No
FW Version	N/A
Identifier	QSFP+
LOS Alarm	N/A
OUI	Mellanox
Power Class	1.5 W max
Rev	A2
Rx Power Current [dBm]	N/A
Temperature [C]	N/A
Transfer Distance [m]	1
Tx Power Current [dBm]	N/A
Voltage [mV]	N/A
Wavelength [nm]	N/A

Cable Information - 7cfe900300f73be0_1

Operational Info | **Module Info** | Troubleshooting Info


Property	Value
Auto Negotiation	ON
FEC	Standard LL RS-FEC - RS(271,257)
Loopback Mode	No Loopback
Physical state	LinkUp
Speed	IB-EDR
State	Active
Width	0x
Enabled Link Speed	0x0000003f (EDR,FDR,FDR10,QDR,DDR,SDR)
Supported Cable Speed	0x0000003f (EDR,FDR,FDR10,QDR,DDR,SDR)

9.3.2.1 Physical Grade and Eye Opening Information

Eye opening information contains the following data:


- Physical Grade: [Grade0, Grade1, Grade2, Grade3]
- Height Eye Opening [mV]: [Height0, Height1, Height2, Height3]
- Phase Eye Opening [psec]: [Phase0, Phase1, Phase2, Phase3]

A new tab called Eye Information was added under cable information modal in ports table.



9.3.2.2 Auto-isolation of High-BER Ports


The High BER Ports tab lists all high-BER ports in the fabric.



The flags `high_ber_ports_auto_isolation` must be configured in the `gv.cfg` file to enable this feature.

For each port discovered as a high-BER port, a new event is triggered in the Events table.


Marking the high-BER port as unhealthy suppresses all events and notifications related to the auto-isolated port.



9.3.3 Virtual Ports Window

 This page is only available if [Virtualization is enabled in gy.cfg](#).

Provides a list of all virtual ports in UFM.



Right-clicking a virtual port allows navigation to the physical port mapped it is mapped to.

Virtual Port State	System Name	Port Name	Virtual Port GUID	Virtual Port LID
✓	H_2303	H_2303 HCA-1	0x0000001230009209	100000
✓	H_2303	H_2303 HCA-1	0x0000002340009209	100001
✓	H_2303	H_2303 HCA-1	0x0000003450009209	100002
✓	H_2303	H_2303 HCA-1	0x0000004560009209	100003

Clicking "Go to port" navigates to the [Virtual Networking tab](#) of the Device Information screen.


9.3.4 Unhealthy Ports Window


The Unhealthy Ports tab shows all the unhealthy nodes in the fabric.

After the Subnet Manager examines the behavior of subnet nodes (switches and hosts) and discovers that a node is “unhealthy” according to the conditions specified below, the node is displayed in the Unhealthy Ports window. Once a node is declared as “unhealthy”, Subnet Manager can either ignore, report, isolate or disable the node. The user is provided with the ability to control the actions performed and the phenomena that declares a node “unhealthy.” Moreover, the user can “clear” nodes that were previously marked as “unhealthy.”

The information is displayed in a tabular form and includes the unhealthy port’s state, source node, source port, source port GUID, peer node, peer port, peer GUID, peer LID, condition, and status time.

Severity	Node	Port	GUID	Name	Port	GUID	LID	Condition	Status Time
Info	smg-ib-sv012	smg-ib-sv012.2	0x043f720300695c6	smg-ib-sv040	smg-ib-sv040.39	0x043f7203006818a0	33	FLAPPING	Thu Apr 28 14:04:08 2...
Minor	smg-ib-sv012	smg-ib-sv012.40	0x043f720300695c6	smg-ib-sv022	smg-ib-sv022.36	0x7cfe9003009a05b0	39	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.16	0x043f720300695c6	smg-ib-sv056	smg-ib-sv056.1/20/1/1	0x90a8400040c840	12	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.31	0x043f720300695c6	smg-ib-apl022-gen3	smg-ib-apl022-gen3...	0x98039a03009fcdce	53	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.32	0x043f720300695c6	smg-ib-apl022-gen3	smg-ib-apl022-gen3...	0x98039a03009fcdce	54	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.26	0x043f720300695c6	smg-ib-vrt003	smg-ib-vrt003 HCA-1	0x98039a03009fcdce	14	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.33	0x043f720300695c6	smg-ib-apl021-gen3	smg-ib-apl021-gen3...	0xb8599d0005681a0	1	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.34	0x043f720300695c6	smg-ib-apl021-gen3	smg-ib-apl021-gen3...	0xb8599d0005681a1	35	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-ib-sv012	smg-ib-sv012.29	0x043f720300695c6	smg-ib-sv036	smg-ib-sv036.33/1	0xb8ce16000604afe	56	FLAPPING	Thu Apr 28 14:10:11 2...


 The feature requires OpenSM parameter `hm_unhealthy_ports_checks` to be set to TRUE (default).

 This feature is not available in the "Monitoring Only Mode."


The following are the conditions that would declare a node as "unhealthy":

- Reboot - If a node was rebooted more than 10 times during last 900 seconds
- Flapping - If several links of the node found in Initializing state in 5 out of 10 previous sweeps
- Unresponsive - A port that does not respond to one of the SMPs and the MAD status is TIMEOUT in 5 out of 7 previous SM sweeps
- Noisy Node - If a node sends traps 129, 130 or 131 more than 250 traps with interval of less than 60 seconds between each two traps
- Seterr - If a node respond with bad status upon SET SMPs (PortInfo, SwitchInfo, VLArb, SL2VL or Pkeys)
- Illegal - If illegal MAD fields are discovered after a check for MADs/fields during `receive_process`
- Manual - Upon user request mark the node as unhealthy/healthy
- Link Level Retransmission (LLR) - Activated when retransmission-per-second counter exceeds its threshold

All conditions except LLR generate Unhealthy port event, LLR generates a High Data retransmission event.

 To clear a node from the Unhealthy Ports Tab, do the following:

1. Go to the Unhealthy Ports window under Managed Elements.
2. From the Unhealthy Ports table, right click the desired port it and mark it as healthy.



Severity	Node	Port	GUID	Name	Port	GUID	LID	Condition	Status Time
Info	smg-lb-sw012	smg-lb-sw012.2	0x043f7203000695cc	smg-lb-sw040	smg-lb-sw040.39	0x043f7203000695cc	33	FLAPPING	Thu Apr 28 14:04:08 2...
Minor	smg-lb-sw012	smg-lb-sw012.40	0x043f7203000695cc	smg-lb-sw012	smg-lb-sw012.36	0x7c1e70030009a0b0	39	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-lb-sw012	smg-lb-sw012.16	0x043f7203000695cc	smg-lb-sw012	1/307/1/1	0x900a840300040c840	12	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-lb-sw012	smg-lb-sw012.31	0x043f7203000695cc	smg-lb-sw012	gen3 ...	0x98039a030009f0cee	53	FLAPPING	Thu Apr 28 14:10:11 2...
Warning	smg-lb-sw012	smg-lb-sw012.32	0x043f7203000695cc	smg-lb-sw012	gen3 ...	0x98039a030009f0cee	54	FLAPPING	Thu Apr 28 14:10:11 2...

 To mark a node as permanently healthy, do the following:

1. Open the `/opt/ufm/files/conf/health-policy.conf.user_ext` file.
2. Enter the node and the port information and set it as "Healthy."
3. Run the `/opt/ufm/scripts/sync_hm_port_health_policy_conf.sh` script.

9.3.4.1 Unhealthy Port Connectivity Filter

It is possible to filter the Unhealthy Ports table by connectivity (all, host-to-switch, or switch-to-host).

Filtering the Unhealthy Ports table is possible from the dropdown options at the top of the table which includes

- All Connectivity
- Switch to Switch
- Host to Switch

Severity	Node	Port	GUID	Name	Port	Peer	LID	Condition	Status Time
Info	smg-ib-sw012	smg-ib-sw012.2	0x043720300695c6	smg-ib-sw040	smg-ib-sw040.39	0x043720300618a0	39	FLAPPING	Thu Apr 28 14:04:08.2...
Minor	smg-ib-sw012	smg-ib-sw012.40	0x043720300695c6	smg-ib-sw022	smg-ib-sw022.36	0x7cfe9003009a05b0	39	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.18	0x043720300695c6	smg-ib-sw056	smg-ib-sw056.1/30/1/1	0x900a84003040c840	12	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.31	0x043720300695c6	smg-ib-apl022-gen3	smg-ib-apl022-gen3...	0x98039b00009fcdce	53	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.32	0x043720300695c6	smg-ib-apl022-gen3	smg-ib-apl022-gen3...	0x98039b00009fcdce	54	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.26	0x043720300695c6	smg-ib-virt003	smg-ib-virt003.HCA-1	0x98039b00009fcdce	14	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.33	0x043720300695c6	smg-ib-apl021-gen3	smg-ib-apl021-gen3...	0xb859903005681a0	1	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.34	0x043720300695c6	smg-ib-apl021-gen3	smg-ib-apl021-gen3...	0xb859903005681a1	35	FLAPPING	Thu Apr 28 14:10:11.2...
Warning	smg-ib-sw012	smg-ib-sw012.29	0x043720300695c6	smg-ib-sw036	smg-ib-sw036.33/1	0xb859903005681a1	56	FLAPPING	Thu Apr 28 14:10:11.2...

9.3.5 Cables Window

Provides a list of all cables in UFM. For more information, see [Device's Cables Tab](#).

Severity	Serial #	Identifier	GUID	Port	GUID	Port	Revision	Link Width	Part #	Technology	Firmware...	Length
Info	MT2153V50...	XFP-E	0x900a8403...	smg-ib-sw056.1/1/1/1	0x900a8403...	smg-ib-sw056.1/2/1/1	A3	4X	MCP4Y10-N...	Copper cabl...	N/A	0.5 m
Info	MT2153V50...	XFP-E	0x900a8403...	smg-ib-sw056.1/1/2/1	0x900a8403...	smg-ib-sw056.1/2/2/1	A3	4X	MCP4Y10-N...	Copper cabl...	N/A	0.5 m
Info	MT2204V50...	XFP-E	0x900a8403...	smg-ib-sw056.1/30/2/2	0x98039b003...	smg-ib-sw032.16	A1	4X	MCP7Y70-H...	Copper cabl...	N/A	2 m
Info	MT2204V50...	XFP-E	0x900a8403...	smg-ib-sw056.1/30/2/1	0xb80c1603...	smg-ib-sw035.16	A1	4X	MCP7Y70-H...	Copper cabl...	N/A	2 m
Info	MT1439V52...	QSFP+	0x7cfe9003...	smg-ib-sw022.28	0x248a0703...	smg-ib-cig001-mgmt01.L1/U2/3	A3	4X	MC2207130...	Copper cabl...	N/A	2 m
Info	MT1515V50...	QSFP+	0x7cfe9003...	smg-ib-sw022.11	0x7cfe9003...	smg-ib-sw022.29	A2	4X	MCP1600-E...	Copper cabl...	N/A	1 m
Info	MT2204V50...	XFP-E	0x0437203...	smg-ib-sw012.16	0x900a8403...	smg-ib-sw056.1/30/1/1	A1	4X	MCP7Y70-H...	Copper cabl...	N/A	2 m
Info	MT1611V50...	QSFP28	0x0437203...	smg-ib-sw012.40	0x7cfe9003...	smg-ib-sw022.36	A2	4X	MCP1600-C...	Copper cabl...	N/A	2 m
Info	MT1515V50...	QSFP+	0x248a0703...	smg-ib-cig001-mgmt01.L2/U2/11	0xec099a03...	unmanagedEDR.21	A2	4X	MCP1600-E...	Copper cabl...	N/A	2 m
Info	MT1605V50...	QSFP+	0x248a0703...	smg-ib-cig001-mgmt01.L2/U2/3	0xec099a03...	unmanagedEDR.26	A2	4X	MCP1600-E...	Copper cabl...	N/A	3 m

Right-clicking a cable from the list allows users to Collect System Dump for the endpoints of the link.

9.3.6 Groups Window

The Groups window allows users to create new groups of devices and provides information about existing groups.

All predefined groups have Read permissions only, except Suppressed_Devices to/from which the user is also able to add/remove members or devices.

The following predefined groups auto-populate upon UFM startup: Switches, 1U_Switches, Modular_Switches, Gateway_Devices, and Hosts.

To create a group of devices, do the following:

1. Click “New” under “Groups.”

Severity	Name	Description	Type
	TU Switches	Includes all TU Switches that exist in the fabric	General
	Alarmed Devices	Devices with alarms	General
	Devices Pending FW Transceivers Reset	Includes all Devices that pending FW transceivers reset to active burned ...	General
	Gateway Devices	Includes all Gateway Devices that exist in the fabric	General
	Modular Switches	Includes all Modular Switches that exist in the fabric	General
	Routers	Includes all Router Devices that exist in the fabric	General
	Servers	Includes all Hosts that exist in the fabric	General
	Servers With DPU	Includes all Devices that has DPU that exist in the fabric	General
	Suppressed Devices	No event notifications issued	General
	Switches	Includes all Switches that exist in the fabric	General

Viewing 1-10 of 10

2. In the New Group wizard, fill in the required information under the General tab: Name (must be between 4-20 characters), Type (General/Rack/Port), and Description (optional), and click Next.

New Group

1 General 2 Members


Name

Type

Description


Next

3. Under Members tab, move the members of the new group from the Available list to the Selected list.



4. Click “Finish” and the new group will appear under the Groups window.

Group members details - port’s hostname, port’s GUID, and device’s IP address - can be viewed when selecting the group from the list of all groups available.



Group Actions

Right-clicking a group enables performing the following actions:

- Edit - groups can be modified either by editing the group description under General tab, or substituting group members under Members tab
 - Delete - existing groups can be deleted from the list
 - **Remove All Members** - all members of an existing group can be removed at once
 - Collect System Dump - sysdump may be generated for all members of an existing group
- The user can filter group by type (General, Rack, Super Switch and Port)

Severity	Name ↑	Description	Type
Minor	1U Switches	Includes all 1U Switches that	General
Minor	Alarmed Devices	Devices with alarms	General
Minor	Devices Pending FW Transceivers Reset	Includes all Devices that pending FW transe...	General
Info	Gateway Devices	Includes all Gateway Devices that exist in the...	General
Info	Modular Switches	Includes all Modular Switches that exist in th...	General
Info	Routers	Includes all Router Devices that exist in the f...	General
Minor	Servers	Includes all Hosts that exist in the fabric	General
Info	Servers With DPU	Includes all Devices that has DPU that exist i...	General
Info	Suppressed Devices	No event notifications issued	General
Minor	Switches	Includes all Switches that exist in the fabric	General

Viewing 1-10 of 10

9.3.7 Inventory Window

Provides a list of all modules in UFM. For more information, see [Device's Inventory Tab](#).

Severity	Status	Serial Number	System	Name	Description	Type	Software Version	Part Number	Temperature
Info	OK	X1LM0930003	smg-b-sw040		SYSTEM	SYSTEM	3.10.1202-xB6_64	S597A41873	37
Info	OK	X1LM0930003	smg-b-sw040		MGMT - 1	MGMT	N/A	S597A41873	N/A
Info	OK	N/A	smg-b-sw040		FAN - 1	FAN	N/A	N/A	N/A
Info	OK	N/A	smg-b-sw040		FAN - 3	FAN	N/A	N/A	N/A
Info	OK	N/A	smg-b-sw040		FAN - 2	FAN	N/A	N/A	N/A
Info	OK	N/A	smg-b-sw040		FAN - 5	FAN	N/A	N/A	N/A
Info	OK	N/A	smg-b-sw040		FAN - 4	FAN	N/A	N/A	N/A
Info	OK	N/A	smg-b-sw040		FAN - 6	FAN	N/A	N/A	N/A
Warning	fatal	X1LM08P0029	smg-b-sw040		PS - 2	PS	N/A	SP87A44110	N/A
Info	OK	X1LM08P0028	smg-b-sw040		PS - 1	PS	N/A	SP87A44110	N/A

Viewing 1-10 of 47

9.3.8 PKeys Window

The PKeys window allows users to create new groups of ports and provides information about existing PKeys.

This window offers one predefined PKey (highlighted in the list of PKeys): Management key 0x7fff with Read permissions only.

For further information about InfiniBand partitioning (Pkeys management), please refer to the [Partitioning Appendix](#).

9.3.8.1 Creating New PKey

To create a PKey:

1. Click the “New” button under “PKeys”.
Please note that the yellow highlighted PKeys are predefined ones.

PKey	Partition	IP Over IB
0x7fff	management	✓
0x7ff	api_pkey_0x7ff	✓

2. In the New PKey wizard, fill in the required information under the General tab:
 - Name—must be between 0x1 and 0x7fff, inclusive
 - Index-0 attribute—True/False
 - IP Over IB attribute—True/False

New PKey

General Members


Pkey: 0x PKey Name

Index-0

IP Over IB

Next

3. Click “Next.”
4. Under Members tab, select the device of which ports you would like to group in one PKey, and move the members (ports) from the Available list to the Selected list. For each member (port) you may specify a membership type (Full/limited).




5. Click “Finish”. The new PKey will become available under the PKey window.

When selecting a PKey from the PKeys table, PKey Information table will appear on the right side of the screen. This table provides information on the PKey's members and QoS settings.


9.3.8.2 PKey Members Tab

Provides details on the PKey members: port's hostname (node), device's IP address, port GUID, port number, membership and index-0 attributes values.



9.3.8.3 PKey QoS Tab

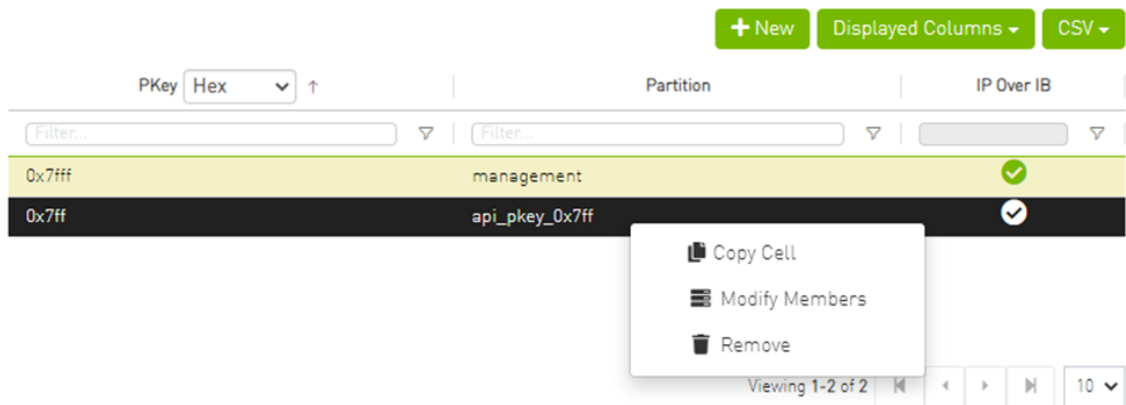
Displays the current partitioning parameter settings of the selected PKey: MTU Limit, Service Level and Rate limit. These settings can be modified by the user.





9.3.8.4 PKey Actions

Right-clicking one PKey from the list enables performing the following actions:

- **Modify Members** - PKeys can be modified either by editing the attributes under General tab, or updating the members under Members tab. Including updating ports memberships.
- **Remove** - existing PKeys can be deleted from the list.



 For information on partitioning, refer to [Appendix - Partitioning](#).

 Note that restarting OpenSM is required for the QoS parameters change to take effect.

9.3.8.5 Support Pkey with Virtual Ports

Creating a pkey with virtual ports is supported, so pkey can contain the following types of port:

- Physical
- Virtual
- Both physical and virtual

The create new pkey wizard dropdown includes port types.

1 General 2 Members

Click on a device to select the members ports from the Available list:

Devices

System Name ↑

Filter...

- r-ufm254-hyp-03
- r-ufm254-hyp-04
- ufm-host87

Viewing 1-3 of 3

Available Ports Show: Physical

GUID ↑

Filter...

- 0x0c42a103007aca90

Viewing 1-1 of 1

Selected Ports

GUID ↑ Memb... Full

Filter... Filter...

No items were found

Viewing 0-0 of 0

Previous

Finish

1 General 2 Members

Click on a device to select the members ports from the Available list:

Devices

System Name ↑

Filter...

- r-ufm254-hyp-03
- r-ufm254-hyp-04
- ufm-host87

Viewing 1-3 of 3

Available Ports Show: Virtual

GUID ↑

Filter...

- 0x1122334477667700
- 0x1122334477667701
- 0x1122334477667710
- 0x1122334477667711

Viewing 1-4 of 4

Selected Ports

GUID ↑ Memb... Full

Filter... Filter...

No items were found

Viewing 0-0 of 0

Previous

Finish

1 General 2 Members

Click on a device to select the members ports from the Available list:

Devices

System Name ↑

Filter...

- r-ufm254-hyp-03
- r-ufm254-hyp-04
- ufm-host87

Viewing 1-3 of 3

Available Ports Show: Both ↓

GUID ↑

Filter...

- 0x0c42a103007aca90
- 0x1122334477667700
- 0x1122334477667701
- 0x1122334477667710
- 0x1122334477667711

Viewing 1-5 of 5

Selected Ports

GUID ↑ Memb... Full ↓

Filter... Filter...

No items were found

Viewing 0-0 of 0

Previous
Finish

9.3.9 HCAs Window

Provides a list of all the HCAs of the hosts in UFM. For more information, see section "[HCAs Tab](#)".

Severity	System	Name	GUID	Type	Port 1	Name	Port 2	Name	PSID	FW Version
Info	smg-lib-svr45		0xe0d9a0300ef551c	ConnectX-5	smg-lib-svr45	HCA-3	smg-lib-svr45	HCA-4	MT_0000000008	16.32.566
Info	smg-lib-gw01-lib-gw		0x0c42a1030098b138	ConnectX-6	smg-lib-gw01-lib-gw	HCA-7	N/A		MT_0000000491	20.30.1004
Info	smg-lib-vrt003		0x9809b03009f14e	ConnectX-6	smg-lib-vrt003	HCA-1	N/A		MT_0000000228	20.29.550
Info	smg-lib-svr036		0x7cfe903000e8a54	ConnectX-4	smg-lib-svr036	HCA-1	smg-lib-svr036	HCA-2	MT_2190110032	12.28.2006
Info	smg-lib-sim001		0x1070903000a0980	BlueField2	smg-lib-sim001	HCA-1	smg-lib-sim001	HCA-2	MT_0000000872	24.33.900
Info	smg-lib-svr027		0x248a0703008fa280	ConnectX-4	smg-lib-svr027	HCA-1	smg-lib-svr027	HCA-2	MT_2190110032	12.28.2006
Info	smg-lib-apl021-gen3		0xb899f03005681a0	ConnectX-6	smg-lib-apl021-gen3	miu5_0	smg-lib-apl021-gen3	miu5_1	MT_0000000224	20.32.1010
Info	smg-lib-svr46		0xe0d9a03000a41ab2	ConnectX-5	smg-lib-svr46	HCA-3	N/A		MT_0000000010	16.32.566
Info	smg-lib-apl009-gen2		0x248a07030009f18ba	ConnectX-4	N/A		smg-lib-apl009-gen2	HCA-2	MT_2190110032	12.28.2006
Info	smg-lib-svr001		0x9809b03000e71ec0	ConnectX-6	smg-lib-svr001	HCA-1	N/A		IBM0000000027	20.31.2006

Viewing 1-10 of 23

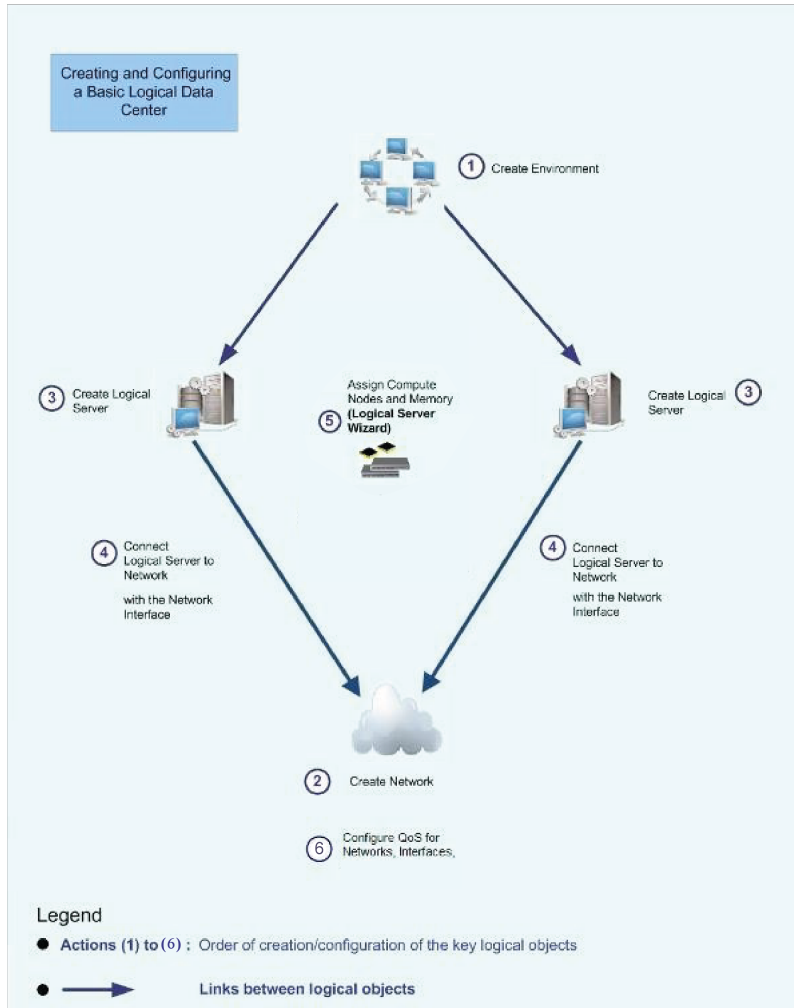
9.4 Logical Elements

All information provided in a tabular format in UFM web UI can be exported into a CSV file.

When designing your model, it is recommended to go about it in the following order:

1. Create an environment.
2. Create a network.
3. Create logical servers.
4. Connect each logical server to a network with a logical server interface.
5. Assign compute nodes using the Logical Server wizard.
6. Configure QoS for networks, interfaces.

The following figure represents the design concept:



Logical Elements allows you to:

- Manage the fabric according to specific needs (e.g. business needs)
- Enable Fabric partitioning and setting QoS policy
- Automate configuration and change management

9.4.1 Environments

The Environments view allows the user to list/manage all existing environment details (e.g. severity, name, description, state).

Severity	Name	Description	State
Info	env1	N/A	created
Info	env2	N/A	created

Viewing 1-2 of 2

When users select an environment, they are able to show/list the logical server details (e.g. severity, name, state virtual NICs, requested computes, used computes) which exists inside this environment:

Severity	Name	State	Virtual NIC(s)	Requested ...	Used Comp...
Info	logical1	allocated	2	3	3

Clicking the logical server name, redirects the user into the logical server view, and the selected logical server is chosen. For more details, please refer to [Logical Servers](#).

Property	Value
Name	logical1
Description	N/A
Environment	env1
OS Type	Linux
Error State	none

9.4.1.1 Creating New Environment

To create a new environment, click the New button located above the environments table.

New Environment
✕

Name

Description

Environment's fields:

1. In the Name field, enter a name for your new Environment.
2. Optional: In the Description field, enter a description for your new Environment
3. Click Submit.

Your new Environment is created. You can see it under the Environments table.

9.4.1.2 Environment Actions

+ New
Displayed Columns ▾
CSV ▾

Severity	Name	Description	State
Info	env1	N/A	created
Info	env2		created

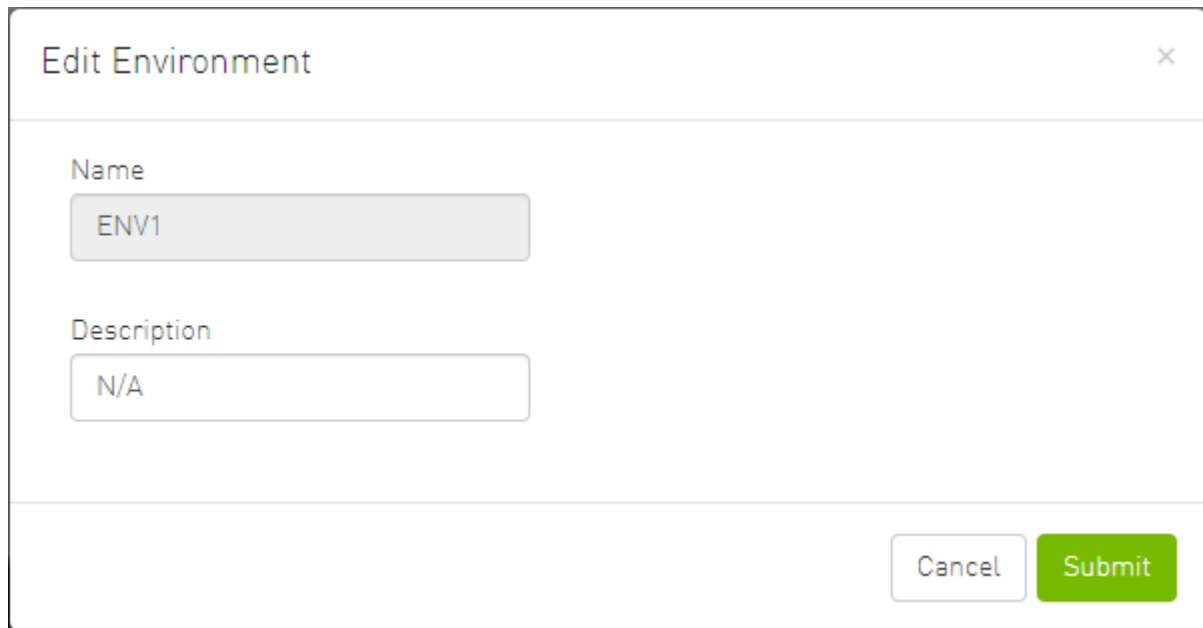
- Copy Cell
- Edit
- Delete

Viewing 1-2 of 2

⏪
⏩
⏴
⏵
10 ▾

9.4.1.2.1 Editing Environments

Click Edit to edit the selected environment.



The screenshot shows a dialog box titled "Edit Environment" with a close button (X) in the top right corner. It contains two input fields: "Name" with the value "ENV1" and "Description" with the value "N/A". At the bottom right, there are two buttons: "Cancel" and "Submit".

9.4.1.2.2 Deleting Environments

Click Delete to delete the selected environment



The screenshot shows a confirmation dialog box with the text "Are you sure you want to delete the environment: ENV1 ?". At the bottom right, there are two buttons: "No" and "Yes".

9.4.2 Networks

The Networks view lists the existing global and local networks and allows managing them.

Severity	Name	PKey	Description	State	Scope
Info	net1	0x1	N/A	created	Global

Viewing 1-1 of 1

9.4.2.1 Adding New Network

This section describes how to create a new global or local network.

When creating a network or network interface, the following SM files are edited as a result of network and network interface configuration:

- partitions.conf - the partitions.conf file is changed when a network is created, deleted, or modified.
- qos-policy.conf - the qos-policy.conf file is changed when a network, logical server, or network interface is created, deleted or updated.

After creating a network, you may attach it to a logical server, which creates a network interface that enables you to use partitions and configure QoS on logical server members.

To add a new network, click the New button.

New Network

General

Scope

Global Local

Name

Network name

Description

Description

PKey

0x 0

PKey Membership

Full Partial

IP Configuration

QoS

IP Services

Cancel Submit

The "New Network" window contains four sections presented in the following subsections.

9.4.2.1.1 General

General ▼

Scope

Global Local

Environment

+

Name

Description

PKey

PKey Membership

Full Partial


This section includes 5 fields:

- Scope: Identify the scope of the new network. If "Local" is selected, a new field appears called "Environment" to select which in which environment the new network should be.
- Name: Name of the network map (required)
- Description: Description of the network map (optional)
- Pkey: PKey of the network map (required)
- Pkey Membership

9.4.2.1.2 IP Configuration

This part is optional.

- IP Subnet: Enter the network IP subnet.

 Using class A network addresses (255.0.0.0) for logical networks may cause high memory consumption.

- Network Mask: Enter the network mask.
- Default Gateway: Enter the network default gateway.

IP Configuration ▼

IP Subnet


Network Mask

Default Gateway

9.4.2.1.3 QoS

This part is optional.

UFM allows fabric traffic prioritization by providing four predefined Service Levels (SL). Each SL defines different queuing priority of the traffic in the fabric. The SL is configured centrally and is applied to all fabric ports. Prioritization occurs when traffic with different SL levels is competing for bandwidth on the same port at the same time.

 To benefit from QoS capabilities in the fabric, please enable QoS from the Settings page.

QoS ▼

MTU Limit

Service Level


Rate Limit

QoS is provisioned to the SM via the `partitions.conf` and `qos-policy.conf` configuration files. You cannot remove or manually modify these files.

QoS parameters are associated with network interfaces.

The UFM software defines the attributes in the `qos-policy.conf` file. When each port group is associated with logical server members, a QoS-level is associated with the QoS parameters set, and the matching rule represents the network interface object. The MTU limit is defined for the network object.

A partition is specified by the network. The QoS is generally defined by the network but can be overwritten by the network interface for the specific logical server, providing a more granular definition of QoS for the specified logical server.

 Before setting QoS, make sure to select the preferred algorithm. If you do not pre-select the algorithm, UFM automatically applies QoS settings to the default algorithm (MINHOP). For more information about configuring the algorithm, please refer to [UFM Routing Protocols](#).

QoS Field	Description
MTU limit	The Maximum Transmission Unit (number of bytes) is defined for network object
Service Level	Select a predefined service level Possible values: 0-15
Rate Limit	Rate limit in Mbps. This value is converted to a standard InfiniBand enumerator (rate_limit, which has fixed values), and provisioned to the SM via the partitions.conf and qos-policy.conf files.


9.4.2.1.4 IP Services

This part is optional.

UFM allows you to specify one of the following IP distribution (configuration) methods:

- Static - UFM Agent creates a new interface with static IP addresses
- External - UFM does not create an interface on hosts. Host configuration is user-defined.

This section contains the IP configuration method (i.e. static or external).

 The default is external, and it will be disabled in case the IP configuration field is empty

IP Services ▼

IP Configuration Method

External ▼

IP Services ▼

IP Configuration Method

Static ▼

Domain Name

Primary DNS

Secondary DNS

The static IP option reveals 3 fields:

- Domain name: the name of the network domain
- The Primary DNS and Secondary DNS fields must have a valid IPv4 format

9.4.2.2 Network Actions

Severity	Name	PKey	Description	State	Scope
Info	net1	0x1	N/A	created	Global

Viewing 1-1 of 1

9.4.2.2.1 Editing Network

Click Edit to edit the selected network map.

Edit Network ✕

General ▾

Scope
 Global Local

Name

Description

PKey

PKey Membership
 Full Partial

IP Configuration ▸

QoS ▸

IP Services ▸

9.4.2.2.2 Deleting Network

Click Delete to delete the selected network.

Are you sure you want to delete the Network: **network1** ?

9.4.3 Logical Servers

The Logical Server object allows you to define a logical server or cluster, allocate resources, and add network interfaces to connect logical servers to the network (partitioning). The resources automatically allocated by UFM inherit the properties of the network in which they reside. Specific resources may be allocated manually.

Logical server activity can be monitored by activating Logical Server Auditing.

When creating a logical server group:

- The UFM server machine cannot be defined as a logical server resource
- UFM does not allow the UFM server to be part of central device management actions, such as reboot, shutdown, and software upgrade

The Logical Servers view lists all existing logical server details (e.g severity, name, description, and state) and allows managing them.

Severity	Name	State	Virtual NIC(s)	Requested Computes	Used Computes
Info	logical1	allocated	2	3	3
Info	logical2	allocated	2	1	1

Clicking on any logical server opens up an Element Information view with the following tabs:

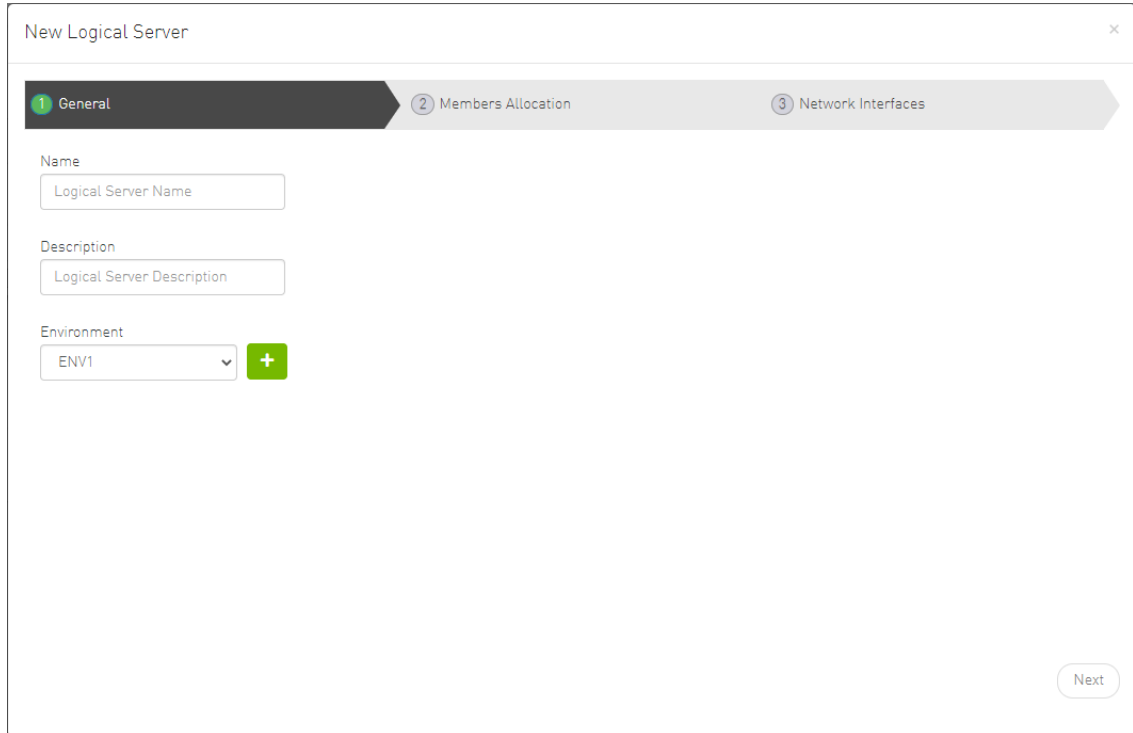
- General
- Members
- Network Interfaces
- Events - the flag `ls_monitoring` must be enabled to view this tab
- Monitoring - the flag `ls_monitoring` must be enabled to view this tab

Property	Value
Name	logical2
Description	N/A
Environment	env1
OS Type	Linux
Error State	none

9.4.3.1 Creating New Logical Server

To create a new logical server, click the New button located above the logical server table. A wizard pops open with 3 steps:

- **General:**



The screenshot shows a wizard titled "New Logical Server" with a close button (X) in the top right corner. The wizard is divided into three steps: 1. General (highlighted in dark grey), 2. Members Allocation (light grey), and 3. Network Interfaces (light grey). The General step contains three input fields: "Name" with the placeholder text "Logical Server Name", "Description" with the placeholder text "Logical Server Description", and "Environment" with a dropdown menu showing "ENV1" and a green plus button to its right. A "Next" button is located at the bottom right of the wizard.

Contains three fields:

- **Name (mandatory):** Name of the new logical server
- **Description (optional):** Description of the new logical server
- **Environment:** Select to which environment the new logical server is be added. Clicking the + button by the drop-down menu provides the ability to create a new environment.

- **Member Allocation:**

Contains two methods to allocate members to the new logical server:

- **Manually:** Select the members manually from a table view

New Logical Server

1 General 2 Members Allocation 3 Network Interfaces

Manually Automatically

<input type="checkbox"/>	Name	GUID	IP
<input type="checkbox"/>	smg-ib-apl002-gen1	0x0002c903001c5f50	0.0.0.0
<input type="checkbox"/>	smg-ib-apl009-gen2	0x248a0703003f18ba	0.0.0.0
<input type="checkbox"/>	smg-ib-svr030	0x98039b03008855a6	0.0.0.0
<input type="checkbox"/>	smg-ib-svr033	0x248a0703008fa200	0.0.0.0
<input type="checkbox"/>	smg-ib-sim001	0xf452140300188540	0.0.0.0
<input type="checkbox"/>	smg-ib-apl004-gen2	0x248a0703008fa15c	0.0.0.0
<input type="checkbox"/>	ufm-appliance-5752c2	0x0002c903000eae670	0.0.0.0
<input type="checkbox"/>	smg-ib-svr032	0xe41d2d0300ef5fa8	0.0.0.0

Viewing 1-8 of 8

Previous Next

- **Automatically:** Specify how many members to allocate and member allocation is done automatically

New Logical Server

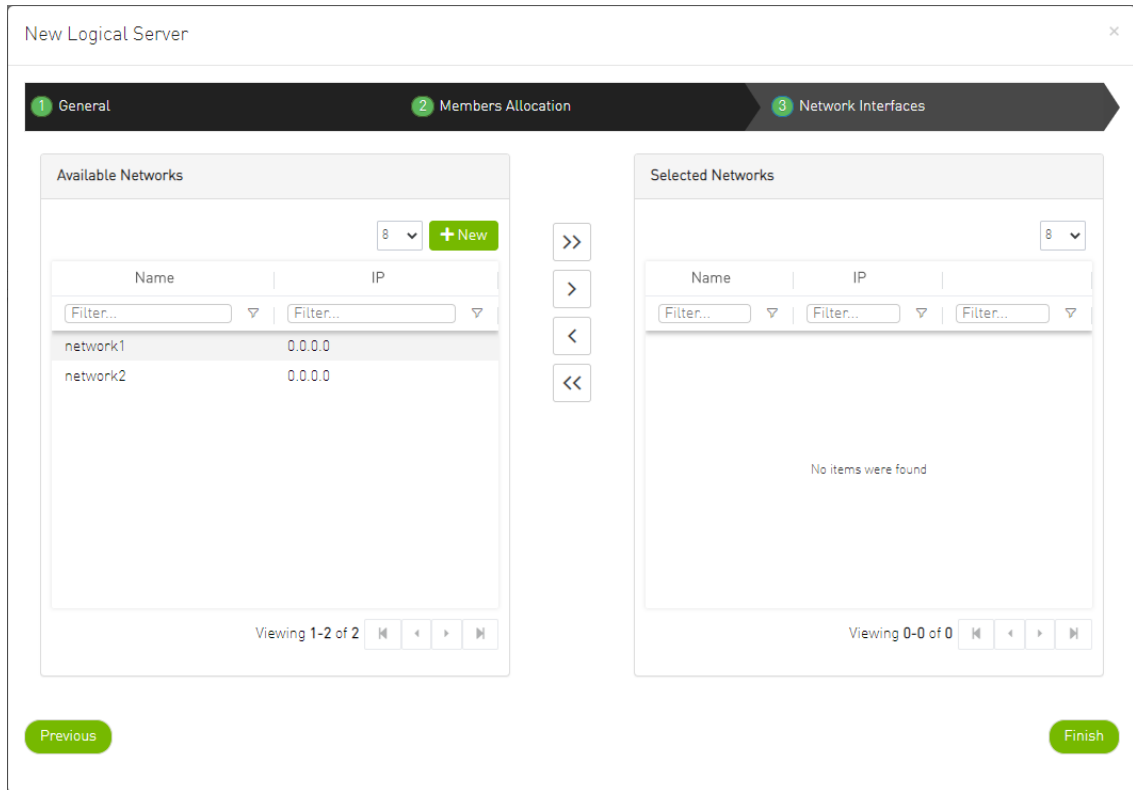
1 General 2 Members Allocation 3 Network Interfaces

Manually Automatically

Units
 (Available Systems: 8)


Previous Next

- Network Interfaces:



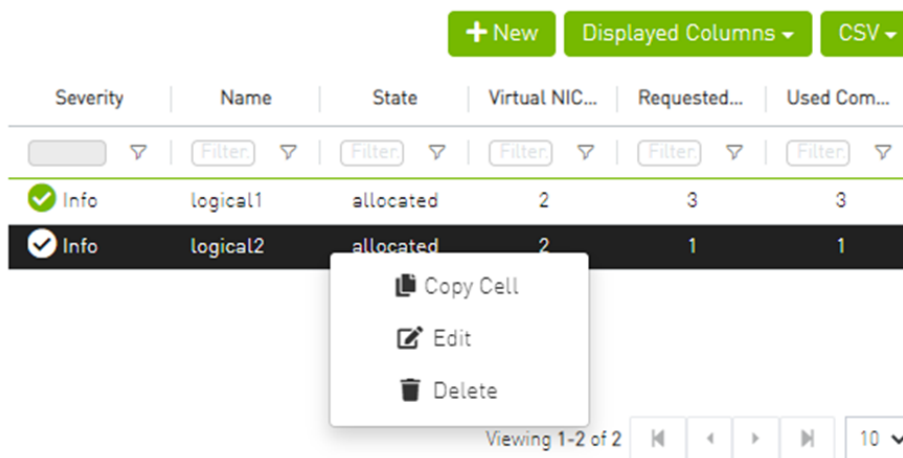
This allows to bind/link between the new logical server and existing networks.

- Clicking the New button above the available networks table allows the user to create a new network.
- Clicking the Edit hyperlink, allows the user to edit the selected network.

 For more details, please refer to [Networks](#).

After completing this wizard, click Finish to create the new logical server.

9.4.3.2 Logical Server Actions



9.4.3.2.1 Editing Logical Servers

Click Edit to edit the selected logical server.

Edit Logical Server

1 General 2 Members Allocation 3 Network Interfaces

Name
LS2

Description
N/A

Environment
ENV2 +

Next

9.4.3.2.2 Deleting Logical Servers

Click Delete to delete the logical server.

Are you sure you want to delete the logical server: LS2 ?

No Yes

9.5 Events & Alarms



All information provided in a tabular format in UFM web UI can be exported into a CSV file.

UFM allows you to identify any problem including ports and device connectivity using events and alarms. Problems can be detected both prior to running applications and during standard operation.

Events trigger alarms (except for "normal" events. i.e., Info events) when they exceed a predefined threshold. Events and alarms can be configured under Events Policy tab under Settings window. For more information, refer to [Events Policy Tab](#).

Events & Alarms Local Time Last Update: 28 Apr 2022 16:46 admin

Alarms

Clear All Alarms Refresh Displayed Columns CSV

Severity	Date/Time ↓	Alarm Name	Source	Source Type	Reason	Count
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw032 / 5	IBPort	Found a [50.0] link that operates in [25.0] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-olg001-mgmt	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw035 / 1	IBPort	Found a [50.0] link that operates in [25.0] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw035 / 23	IBPort	Found a 4x link that operates in 2x width mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw035 / 24	IBPort	Found a 4x link that operates in 2x width mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw035 / 26	IBPort	Found a [50.0] link that operates in [25.0] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	default[12] / Switch: smg-ib-s	IBPort	Found a [50.0] link that operates in [25.0] speed mode.	53
Minor	2022-04-28 16:43:46	Non-opti...	Switch: smg-ib-sw022 / 28	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	default[12] / Switch: smg-ib-s	IBPort	Found a [25.0] link that operates in [2.5] speed mode.	56
Minor	2022-04-28 16:43:46	Non-opti...	default[12] / Switch: smg-ib-s	IBPort	Found a [50.0] link that operates in [25.0] speed mode.	53

Viewing 1-10 of 77

Events

Clear All Events Refresh Displayed Columns CSV

Severity	Date/Time ↓	Event Name	Source	Source Type	Description	Category
Info	2022-04-28 16:41:29	Network Interface...	logical2(0/0)	LogicalServer	Network Interface env1_logical2_manage...	
Info	2022-04-28 16:41:29	Logical Server Ad...	env1(1)	Environment	Logical Server logical2 is added	
Info	2022-04-28 16:41:29	Compute Resourc...	logical2(1/1)	LogicalServer	Compute Resource logical2/1 [smg-ib-svr...	
Info	2022-04-28 16:41:29	Logical Server Re...	logical2(1/1)	LogicalServer	Logical Server allocated 1 Resources	
Info	2022-04-28 16:41:29	Network Interface...	logical2(1/1)	LogicalServer	Network Interface env1_logical2_net1 is a ...	
Critical	2022-04-28 16:38:38	Module status FA...	default[12] / Switch: smg-ib-sw	Switch	Module PS 2 on smg-ib-sw040[10.209.24...	
Info	2022-04-28 16:32:22	Environment Added	Grid	Grid	Environment env2 is added	
Info	2022-04-28 16:31:35	Network Interface...	logical1(0/0)	LogicalServer	Network Interface env1_logical1_manage...	
Info	2022-04-28 16:31:35	Logical Server Ad...	env1(0)	Environment	Logical Server logical1 is added	
Info	2022-04-28 16:31:35	Compute Resourc...	logical1(1/1)	LogicalServer	Compute Resource logical1/1 [smg-ib-svr...	

Viewing 1-10 of 100

Users can enable the events persistency mechanism from the `gv.cfg`. This allows the user to see the events in the case of restarting the UFM or in HA mode.

- ⚠ Alternatively you can run the following commands:
- `ufm events persistency enable`
 - `ufm events max-restored`

The persistency is deactivated by default and can be enabled by the following controlled parameters in the config file:

- `max_restored_events = 50 #` - will determine the number of events to restore
- `events_persistency_enabled = true #` - will set to true for the feature to work

9.6 Telemetry

Error: null

9.7 System Health

The System Health window enables running and viewing reports and logs for monitoring and analyzing UFM server and fabric health through the following tabs: UFM Health, UFM Logs, UFM Snapshot, Fabric Health, Daily Reports and Topology Compare.

- [UFM Health Tab](#)
- [UFM Logs Tab](#)
- [UFM Snapshot Tab](#)
- [Fabric Health Tab](#)
- [Daily Reports Tab](#)
- [Topology Compare Tab](#)
- [Fabric Validation Tab](#)
- [IBDiagnet Tab](#)

9.7.1 UFM Health Tab

Through UFM Health tab, you can create reports that run a series of checks on the UFM server.

Each check that is run for a report triggers a corresponding event. Events are also triggered when a report starts and ends. For more information, see [Events & Alarms](#).

To run a new report, click “Run New Report”. Results will be displayed inline automatically.

System Health

UFM Health UFM Logs UFM Snapshot Fabric Health Daily Reports Topology Compare Fabric Validation IBDiagnet

UFM Health Report

Date 2020-10-11 17:21:00
Created By admin

Show Problems Only

✓ UFM Configuration	Completed Successfully. See details below >
✓ UFM Processes	Completed Successfully. See details below >
✓ Memory Monitoring	Completed Successfully. See details below >
✓ CPU Monitoring	Completed Successfully. See details below >
✓ Disk Monitoring	Completed Successfully. See details below >
✓ Fabric Interface	Completed Successfully. See details below >
✓ Core Dumps List	Completed Successfully. See details below >

You can expand the results of each check or expand the results of all checks at once by clicking the “Expand All” button.

To view only the errors of the report results, click the “Show Problems Only” checkbox.

The following tables describe the checks included in the report.

UFM Health Report Checks

UFM Configuration	
Check	Description
Release Number	UFM software version and build.
License Type	Type of license, permanent or evaluation.
License Customer Number	The customer number provided by NVIDIA.
License UID	The UFM serial number provided by NVIDIA.
License Expiration Date	License expiration date for limited licenses.
License Functionality	Level of functionality enabled for the end-user, standard or advanced.
License Devices Limit	The maximum number of devices that UFM is licensed to manage. Note that it displays the current active and valid UFM licenses (not the sum of all valid licenses devices)
Running Mode	UFM running mode, Standalone or High Availability (HA). When UFM is in HA mode, additional information is displayed for the master and standby servers.

UFM Processing	
Check	Description
OpenSM	Status of the OpenSM service.
ibpm	Status of the ibpm (Performance Manager) service.
ModelMain	Status of the main UFM service.
httpd	Status of the httpd service.
MySQL	Status of the MySQL service.

Memory Monitoring	
Check	Description
Total memory usage	Percentage of total memory usage.
UFM memory usage	Percentage of UFM memory usage

CPU Monitoring	
Check	Description
Total CPU Capacity	Percentage of CPU capacity available
CPUs Number	Number of CPUs
Total CPU utilization	Percentage of total CPU utilization.
UFM CPU utilization	Percentage of UFM CPU utilization.

Disk Monitoring	
Check	Description
Disk <diskname>	Percentage of disk usage.

Fabric Interface	
Check	Description
Fabric Interface	Name and state of fabric interface.

9.7.2 UFM Logs Tab

UFM logging records events and actions that can serve to identify fabric and UFM server issues and assist in troubleshooting.

The logs are categorized into three files according to the activities they record: Event logs, SM logs, and UFM logs.

To view the log files, select the desired log file from the drop-down menu. Log data will be displayed:

System Health

UFM Health UFM Logs UFM Snapshot Fabric Health Daily Reports Topology Compare Fabric Validation IBDiagnet

Event Logs Time Last 24 hours 10000 Search...

Log View

```

2020-11-09 13:15:27.382 [84852] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 14:15:48.621 [84853] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 14:15:51.566 [84854] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 14:20:48.702 [84855] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 15:04:31.752 [84856] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:04:34.737 [84857] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 15:06:34.147 [84858] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:06:37.035 [84859] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 15:09:28.227 [84860] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:09:31.035 [84861] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 15:14:07.896 [84862] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:15:06.817 [84863] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:15:43.199 [84864] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:16:33.799 [84865] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:17:17.746 [84866] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:17:41.635 [84867] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:18:04.345 [84868] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:20:23.340 [84869] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 15:20:26.226 [84870] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 15:25:23.597 [84871] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 16:23:33.584 [84872] [352] INFO [Logical_Model] Grid [Grid]: Network management is added
2020-11-09 16:23:36.510 [84873] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 16:28:33.650 [84874] [605] CRITICAL [Maintenance] Grid [Grid]: Fabric Analysis Report failed, Return code: 1
2020-11-09 16:37:17.833 [84875] [352] INFO [Logical_Model] Grid [Grid]: Network management is added

```

In the Logs window, you can do the following:

- Refresh the data using the Refresh button on the right-hand side of the screen
- Search for a specific value using the Search bar
- Limit the display to a specific time period using the Time drop-down menu
- Limit the display to a specific number of lines using the drop-down menu (use "All" option to display all lines)

9.7.2.1 Event Logs

Event Logs show the history of fabric events detected and initiated by the UFM server. The timestamp and severity of an event is indicated as well as the cause of the event and additional

relevant information. *The Event log is kept on the UFM server in the /opt/ufm/log/events.log file.* Events can be configured whether to appear in the log files under the Events Policy tab in the Settings window. For more information, see [Events Policy](#).

See "[Appendix - Supported Port Counters and Events](#)" for a comprehensive list of Events.

9.7.2.2 Subnet Manager (SM) Logs

SM Logs show messages of the Subnet Manager and communication plug-in.

The log verbosity is defined by selecting the Log Levels in the Subnet Manager tab under Settings window. For more information, see [Subnet Manager Tab](#).

9.7.2.3 UFM Logs

UFM Logs is a general log of UFM Server. The log saves a history of user actions, events, polling results and other server activities and errors. Log verbosity is defined on start-up in the configuration file /opt/ufm/conf/gv.cfg:

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

The default verbosity level is WARNING.

9.7.3 UFM Snapshot Tab

You can export and save UFM database information and configuration files in a predefined location. In this way you can create a full snapshot before upgrading.

By default, the snapshot includes UFM database and UFM configuration files. You can also save troubleshooting information, so that you can send all information required for debugging to Mellanox Support. The additional troubleshooting information includes system snapshot files and UFM log files.

To create a snapshot, click the "Create Snapshot" button.



The screenshot shows the 'System Health' interface with several tabs: UFM Health, UFM Logs, UFM Snapshot (selected), Fabric Health, Daily Reports, Topology Compare, Fabric Validation, and IBDiagnet. Below the tabs, there is a 'Create Snapshot' section with the text 'Create Snapshot of UFM Database and Configuration Files.' and a checkbox labeled 'Include Troubleshooting Information' which is currently unchecked. A green 'Create Snapshot' button is located to the right of the checkbox.

To save the troubleshooting information for debugging purposes, check the Include Troubleshooting Information checkbox.

UFM will create the snapshot and save the data to the predefined location. By default, the snapshot files are stored under /opt/ufm/backup directory. You can change the location of the snapshot files in the gv.cfg configuration file in the backup folder location section.

For example:

```
#backup folder location
backup_folder=/opt/ufm/backup
```


9.7.4 Fabric Health Tab

Through Fabric Health tab, you can create reports that run a series of checks on the fabric.

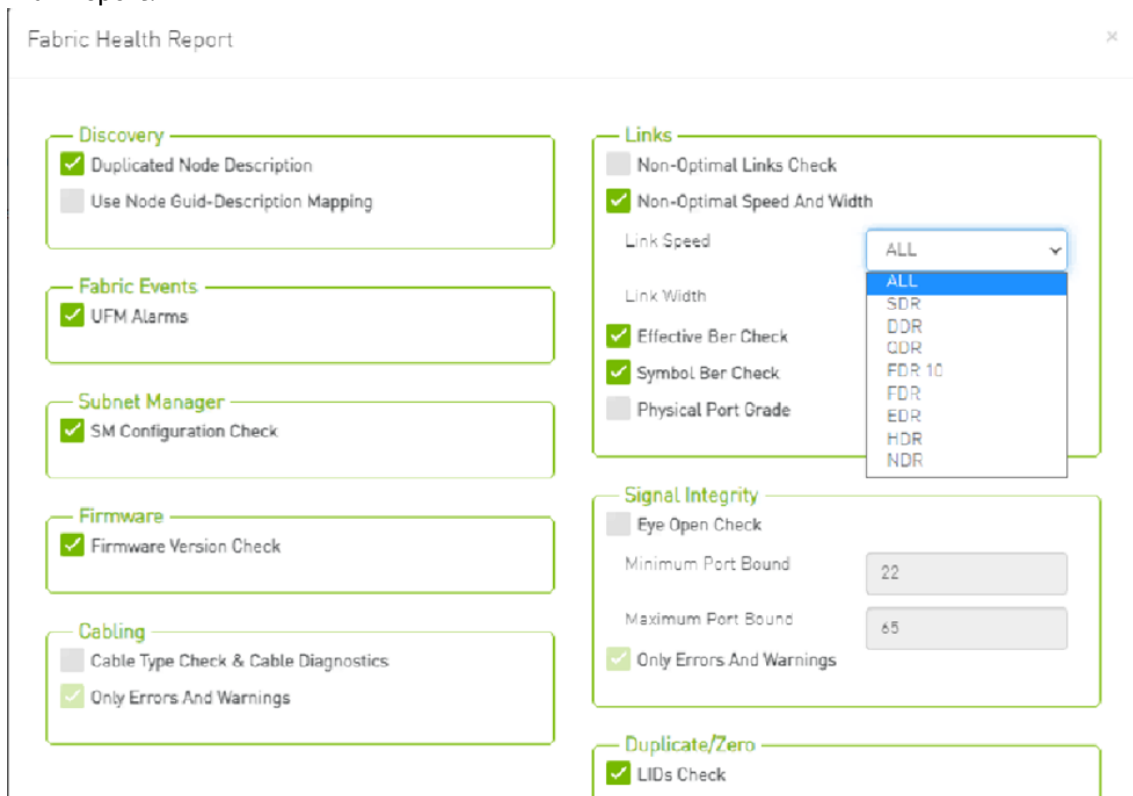
Each check that is run for a report triggers a corresponding event. Events are also triggered when a report starts and ends. For more information, see [Events & Alarms](#).

 To run a new report, do the following:


1. Click “Run New Report.”



2. Select the desired fabric health checks to run in the Fabric Health Report window and click “Run Report.”



Results will be displayed automatically:

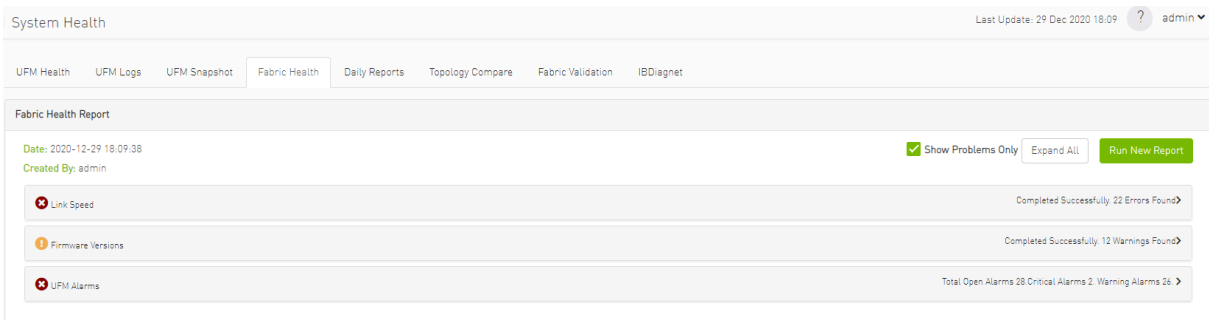


The report displays, the following:

- A report summary table of the errors and warnings generated by the report.
- A fabric summary of the devices and ports in the fabric.
- Details of the results of each check run by the report.

You can expand the view of each check or expand the view of all checks at once by clicking “Expand All.”

To view only the errors of the report results, click the “Show Problems Only” checkbox.



The following table describes the checks included in the report.

Fabric Health Report Checks

Check	Description	To run, select:
Duplicate/Zero LID Check	Lists all ports with same LID or zero LID value.	LIDs Check Default: Selected
Duplicated Node Description	Lists all nodes with same node description. Does not include switches with the same description.	Duplicated Node Description Default: Selected

Check	Description	To run, select:
Use Node GUID-Description Mapping	Enables the usage of a mapping file (between node GUID and node description) when running duplicate node description analysis of the fabric. This file is located on the UFM server side at: <i>/opt/ufm/conf/sm_guid_desc_mapping.cfg</i> , and uses the following format (node_guid → description): <i>0x248a070300702710 "Desc1"</i> <i>0x248a0703007026f0 "Desc2"</i> <i>0x0002c90300494100 "Desc3"</i>	Use Node GUID-Description Mapping Default: Unchecked Note: In order for this checkbox to be available, the Duplicated Node Description checkbox should also be selected. Otherwise, this checkbox will be greyed-out.
SM Check	Checks that: <ul style="list-style-type: none"> • There is one and only one active (master) Subnet Manager in the fabric. • The master is selected according to highest priority and lowest port GUID. The report lists all SMs in the fabric with their attributes.	SM Configuration Check Default: Selected
Bad Links Check	Performs a full-fabric discovery and reports “non-responsive” ports with their path.	Non-Optimal Links Check Default: Selected
Link Width	Checks if link width is optimally used. <ul style="list-style-type: none"> • When a width is selected, the report lists the active links that do not meet the optimum for the selection. • When no width is selected (All), the test checks whether the enabled width on both sides of the link equals the configured maximum (confirms that auto-negotiation was successful). 	None-Optimal Speed and Width Default: Selected Link Width: The default is ALL.
Link Speed	Checks if link speed is optimally used. <ul style="list-style-type: none"> • When a speed is selected, the report lists the active links that do not meet the optimum for the selection. • When no speed is selected (All), the test checks whether the enabled speed on both sides of the link equals the configured maximum (confirms that auto-negotiation was successful). 	None-Optimal Speed and Width Default: Selected Link Speed: The default is ALL.
Effective Ber Check	Provides a BER test for each port, calculates BER for each port and check no BER value has exceeded the BER thresholds. In the results, this section will display all ports that has exceeded the BER thresholds. Note that there are two levels of threshold: Warning threshold (default=1e-13) and Error threshold (default=1e-8).	Effective Ber Check Default: Selected
Effective Port Grade	Provides a grade per port lane in the fabric, which indicates the current port lane quality.	Physical Port Grade Default: Not Selected

Check	Description	To run, select:
Firmware Check	Checks for firmware inconsistencies. For each device model in the fabric, the test finds the latest installed version of the firmware and reports devices with older versions.	Firmware Version Check Default: Selected
Eye Open Check	(For QDR only) Lists Eye-Opener information for each link. When minimum and maximum port bounds are specified, the report lists the links with eye size outside of the specified bounds.	Eye Open Check Default: Selected Minimum and Maximum port bound: By default no bounds are defined.
Cable Information	Reports cable information as stored in EPROM on each port: cable vendor, type, length and serial number.	Cable Type Check & Cable Diagnostics Default: NOT selected because this test might take a long time to complete (40 msec per port)
UFM Alarms	Lists all open alarms in UFM.	UFM Alarms Default: Selected

9.7.5 Daily Reports Tab

The Daily Report feature collects, analyzes, and reports the most significant issues of the fabric in the last 24 hours (from 00:00 to 24:00). The reports present statistical information such as Summary of Traffic, Congestions and UFM events that occurred during the last 24 hours. These statistics are sent to a pre-defined recipients list on a daily basis. It is also possible to specify a non-24-hour range, by updating the UFM configuration file—see section [Other Daily Report Configurations](#) for details.

The following are the formats of the Daily Report:

- Interactive—opened via the browser. The charts are displayed in SVG format. This format can be accessed from the UFM Web UI and is also sent by email as an attachment (see [Daily Report View in the Web UI](#) section below).
- Static—opened via mail client (Outlook, Gmail, Hotmail, etc). The charts are displayed in PNG format.

9.7.5.1 Activating and Deactivating the Daily Report

Daily Report can be activated/deactivated via the `/opt/ufm/conf/gv.cfg` file.

 Daily Reports mechanism is activated by default.

 To deactivate the Daily Report, do the following:

1. Open the `/opt/ufm/conf/gv.cfg` file.
2. Find the DailyReport section.
3. Set the `daily_report_enabled` option to false.

```
daily_report_enabled = false
```

➤ *To re-activate the Daily Report:, do the following:*

1. Open the `/opt/ufm/conf/gv.cfg` file.
2. Find the DailyReport section.
3. Set the `daily_report_enabled` option to true.

```
daily_report_enabled = true
```

9.7.5.2 Saving Daily Reports

UFM saves the interactive Daily Reports under the `/opt/ufm/files/reports/Daily` directory. Each report will be saved under a directory with its respective date. For example, report for Sept. 28th, 2014 will be located under: `/opt/ufm/files/reports/Daily/2014-09-28/` By default, the maximum number of reports that will be saved is 365 (one per day).

➤ *To configure the maximum number of reports to save, do the following:*

1. Open the `/opt/ufm/conf/gv.cfg` file.
2. Find the DailyReport section.
3. Set the `max_reports` option to the desired value. A count of 0 (zero) means no copies are retained. (default and max is 365).
4. Restart UFM.

9.7.5.3 Other Daily Report Configurations

All the Daily Report configuration parameters can be found in the "DailyReport" section in `gv.cfg` configuration file.

The following are additional Daily Report configurations options:

- `top_x` option specifies the number of results in the "Top X" charts. Max number can be 20. (Default value is 10). `top_x` value will be applied to all charts existing in the Daily Report.
- `mail_send_interval` option specifies the epoch in minutes after midnight that the report can be emailed. By default, if UFM was down during midnight, and was restarted after 1:00, the report of the previous day will be generated and saved, but will not be emailed. This can be changed by editing the `mail_send_interval`. (default value is 60 minutes, meaning that the report will be send only between 00:00 to 1:00).

- `log_level` option specifies the Daily Report log verbosity. Default value is INFO (optional values: INFO, WARNING and ERROR).
- `attach_fabric_health_report` option indicates whether or not to add the fabric health report as attachment to the mail. Default value is true (optional values: true or false).
- `fabric_health_report_timeout` specifies the max time in seconds, to wait for fabric health report generation. Default value is 900 seconds (15 minutes).
In case of large fabrics, fabric health report might take longer than the default 15 minutes. User can enlarge the timeout for fabric health report to complete.
- `max_attached_file_size` specifies the maximum file size in Bytes for each email attachment that can be sent. Default value is 2 Megabytes.
If the size of a certain file has exceeded this value, the file will not be sent as an attachment in the Daily Report mail.

```
[DailyReport]
# top_x specifies the number of results per each top x chart.
# max number can be 20.(default is 10)
top_x=10
# max_reports specifies the number of reports to save.
# A count of 0 (zero) means no copies are retained.(default and max is 365)
max_reports = 365
#time interval in minutes after midnight
#when passed mail will not be sent
mail_send_interval=60
log_level = INFO
daily_report_enabled = true
attach_fabric_health_report = true
fabric_health_report_timeout = 900
# max attached file size in bytes, default is 2M (2097152 Bytes)
max_attached_file_size = 2097152
```

- `max_attached_file_size` specifies the maximum file size in Bytes for each email attachment that can be sent. Default value is 2 Megabytes.
- The `start_hour` and `end_hour` options enable selecting a sub-range of the day, during which, the relevant report data will be collected. Since by default this option is configured to collect data from the last 24 hours, the default `start_hour` is set to 0 (or 00), and the default `end_hour` is set to 24.

If these options are configured to different values, the generated report will include data from the specified interval only. The `start_hour` values range is 00 to 23, and the `end_hour` values range is 00 to 24. The specified `end_hour` must be greater than the specified `start_hour`. If, for example, the `start_hour` is configured to 08, and the `end_hour` is configured to 10, the generated report will include data collected between 08:00-10:00 (excluding 10:00).

9.7.5.4 Report Content


9.7.5.4.1 Sidebar







The Sidebar includes general information regarding the fabric, such as: the site name, number of switches and hosts in the fabric, and the dates on which the report was generated.

Navigation between the charts can be done via the menu charts on the sidebar.

Fabric
Events (by severity)
Normalized Traffic and Congestion
Hosts Utilization
Most active events
Hosts
Top Senders (Hosts only)
Hosts with most events
Hosts with most critical events
Most congested hosts
Hosts with most link down events
Switches
Switches with most events
Switches with most critical events
Most congested switches
Switches with most link down events

9.7.5.4.2 Daily Report Highlights


The top of the report shows highlight activities of the network, such as: the host with the most events, the most congested host and switch, and top sender host. To see the related chart of each highlight, click the corresponding  icon in the "Link to chart column."

Highlights		
	Highlight	Link to chart
Switch with most events	'switch-630744'	
Host with most events	'r-ufm135 HCA-1'	
Total events during the last 24 hours	total: 110973, critical events: 14877, warning events: 14784, minor events: 81312.	
Most congested host	'r-ufm87 HCA-1' (20.0% congestion)	
Top sender host	'r-ufm86 HCA-1' (46.0% BW and 0% congestion)	
Highest traffic patterns	Highest traffic hour: 09:00-10:00 (46.0% BW), Most congested hour: 23:00-24:00 (10.0% congestion)	
Number of unhealthy ports	0	N/A

9.7.5.4.3 Available Charts

9.7.5.4.3.1 Events by Severity

Events by Severity displays in a graphical view the distribution of all the UFM events that occurred during each hour. Events are separated into the following severity levels: Critical, Minor, and Warning.




Hovering over the bars in the interactive report displays the amount of events per hour.

9.7.5.4.3.2 Normalized Traffic and Congestion

Normalized Traffic and Congestion displays in a graphical view the normalized traffic and congestions of the fabric. This graph displays the accumulated data for the Senders in the fabric (not including switches).

Congestion normalization is based on the number of delayed packets (packets that wait in the queue) and bandwidth loss.

The graph displays the percentage of the traffic utilization in green and the percentage of the congestion in red.



⚠️ Hovering over the bars in the interactive report displays the percentage of the traffic/ congestion per hour.


9.7.5.4.3.3 Hosts Utilization Distribution

Hosts Utilization Distribution displays in a graphical view the groups of hosts, where each host belongs to a specific group according to its utilization status.

To see the hosts in each group, click on the pie chart (at the interactive report).

The utilization groups are:


- Very low—up to 20% utilized
- Low—20-40% utilized
- Moderate—40-60% utilized
- High—60-80% utilized
- Very high—80-100% utilized



⚠️ Hovering over the slices in the interactive report displays the percentage of hosts in this group.

9.7.5.4.3.4 Most Active Events

Most Active Events displays in a graphical view the most active events, ordered by the number of occurrences during the last 24 hours.




⚠️ Hovering over the bars in the interactive report displays the number of occurrences for each active event, and hovering on each event's name displays a tooltip with the event's description.

9.7.5.4.3.5 Top Senders

Top Senders displays in a graphical view the normalized traffic and congestions of the top sender hosts. Congestion normalization is based on the number of the delayed packets (packets that wait in queue) and bandwidth loss.


The graph displays the percentage of the traffic utilization in green and the percentage of the congestion in red.



⚠️ Hovering over the bars in the interactive report displays the percentage of the traffic/ congestion for a selected host.

9.7.5.4.3.6 Hosts with Most Events


Hosts with Most Events displays in a graphical view the hosts with the most events. Events are separated into the following severity levels: Critical, Minor, and Warning.



⚠️ Hovering over the bars in the interactive report displays the amount of events per severity for a selected host.

9.7.5.4.3.7 Hosts with Most Critical Events

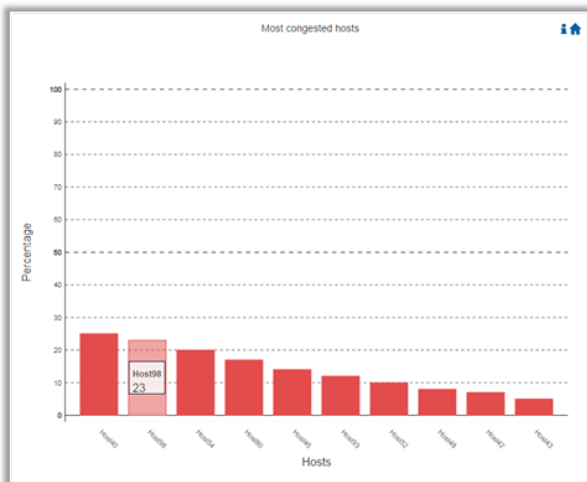
Hosts with Most Critical Events displays in a graphical view the hosts with the most critical events.



⚠️ Hovering over the bars in the interactive report displays the amount of critical events for a selected host.

9.7.5.4.3.8 Most Congested Hosts


Most Congested Hosts displays in a graphical view the normalized congestions of the most congested hosts. Congestion normalization is based on the number of the delayed packets (packets that wait in queue) and bandwidth loss.



⚠️ Hovering over the bars in the interactive report displays the percentage of the congestion for a selected host.

9.7.5.4.3.9 Hosts with Most Link Down Events


Hosts with Most Link Down Events displays in a graphical view the list of the hosts with the most link down events during the last 24 hours.



⚠️ Hovering over the bars in the interactive report displays the amount of link-down events for a selected host.

9.7.5.4.3.10 Switches with Most Events


Switches with Most Events displays in a graphical view the switches with the most events. Events are separated into the following severity levels: Critical, Minor, and Warning.



⚠️ Hovering over the bars in the interactive report displays the amount of events per severity for a selected switch.

9.7.5.4.3.11 Switches with Most Critical Events


Switches with Most Critical Events displays in a graphical view the switches with the most critical events.



⚠️ Hovering over the bars in the interactive report displays the amount of critical events for a selected switch.

9.7.5.4.3.12 Most Congested Switches


Most Congested Switches displays in a graphical view the normalized congestions of the most congested switches. Congestion normalization is based on the number of delayed packets (packets that wait in queue) and bandwidth loss.




⚠️ Hovering over the bars in the interactive report displays the percentage of the congestion for a selected switch.


9.7.5.4.3.13 Switches with Most Link Down Events

Switches with Most Link Down Events displays in a graphical view the list of the switches with the most link down events during the last 24 hours.



⚠️ Hovering over the bars in the interactive report displays the amount of link-down events for a selected switch.

⚠️ Clicking on the “help” icon  in the upper right corner of each chart, in the interactive report, will display a short description of the chart.

Clicking on the “home” icon  in the upper right corner of each chart, in the interactive report, will move the display to the beginning of the report.


⚠️ On charts: “Events by Severity”, “Hosts with Most Events”, and “Switches with Most Events”, if the maximum value in the Y-axis is less than 5, an “m” unit will appear and stand for “milli”.

⚠️ For all charts, if the value is higher than 1000 in the Y-axis, a “k” unit will appear and stand for “killo”.

9.7.5.4.4 Daily Report View in the Web UI

In this tab, you can select the UFM daily reports that you wish to view and you can specify the recipients to which these daily reports will be sent.


➤ To view a specific daily report, click the relevant report date from the list of available daily reports.




The specified report content will be displayed when clicking the report (see [Activating and Deactivating the Daily Report](#)).

➤ To configure the Recipients list for the daily reports, do the following:


1. Click Recipients List under System Health → Daily Reports tab.



2. Click New.



3. In the Recipients List window, enter valid recipient email addresses, comma-separated, and click Submit.



The new recipient/recipients will be added to the Daily Reports Recipients list.

Email
user@user.com

Viewing 1-1 of 1

These recipients will automatically start receiving the UFM daily reports.

9.7.6 Topology Compare Tab

9.7.6.1 Overview

The Topology Compare tab allows two methods of topology comparison:

- Periodic Comparison
- Custom Comparison

9.7.6.1.1 Periodic Comparison

Periodic comparison allows users to compare the current fabric topology with a preset master topology. The master topology may be set either by selecting the current topology or uploading a predefined custom topology.

The screenshot shows the UFM interface for topology comparison. At the top, there are tabs for 'Periodic Comparisons' and 'Custom Comparisons'. Below this, the 'Master Topology Snapshot' is shown as '/opt/ufm/files/periodicTopo/master.topo' with a 'Last Update' of '2022-04-27 20:23:01'. On the right, there are three buttons: 'Update Master Topology', 'Download Topology', and 'Settings'. The main area is split into two panels. The left panel, 'Topology Compare Reports', contains a table with columns 'ID' and 'Date/Time'. The table has 8 rows, with the first row (ID: 6, Date/Time: 2022-04-28 6:00:07) selected. The right panel, 'Topology Compare Report Details', shows the details for the selected report. It includes a 'Date' of '2022-04-28 6:00:07' and 'Created By: UFM'. A summary bar indicates 'Total: 1 Additional cables detected'. Below this, there is a table with columns 'Severity' and 'Detected Differences'. A warning is shown: 'Warning: Unplanned cable connection between S7cfe900300a5a2a0/N7cfe900300a5a2a8/P1 and sw-hpcc2/U1/P37'. At the bottom of the details panel, it says 'Viewing 1-1 of 1'.

When a report is selected from the "Topology Compare Reports" table, its result are displayed on the right side under "Topology Compare Report Details".

- To update the master topology with the latest (current) topology or a custom topology saved in external file, click the "Updated Master Topology" dropdown button.

This screenshot shows the 'Update Master Topology' dropdown menu. The menu is open, showing two options: 'With Latest Topology' and 'With Custom Topology'. The background shows the 'Topology Compare Report Details' panel with a summary bar indicating 'Total: 1 nodes have non-parsible NodeDescription', 'Total: 5 Additional cables detected', and 'Total: 6 Additional nodes detected'. The 'Update Master Topology', 'Download Topology', and 'Settings' buttons are visible at the top right.

- To download the current topology as a `.topo` file, click the "Download Topology" button.
- The Settings button navigates to the [Topology Compare tab](#) of the Settings view which allows users to configure periodic comparison settings.

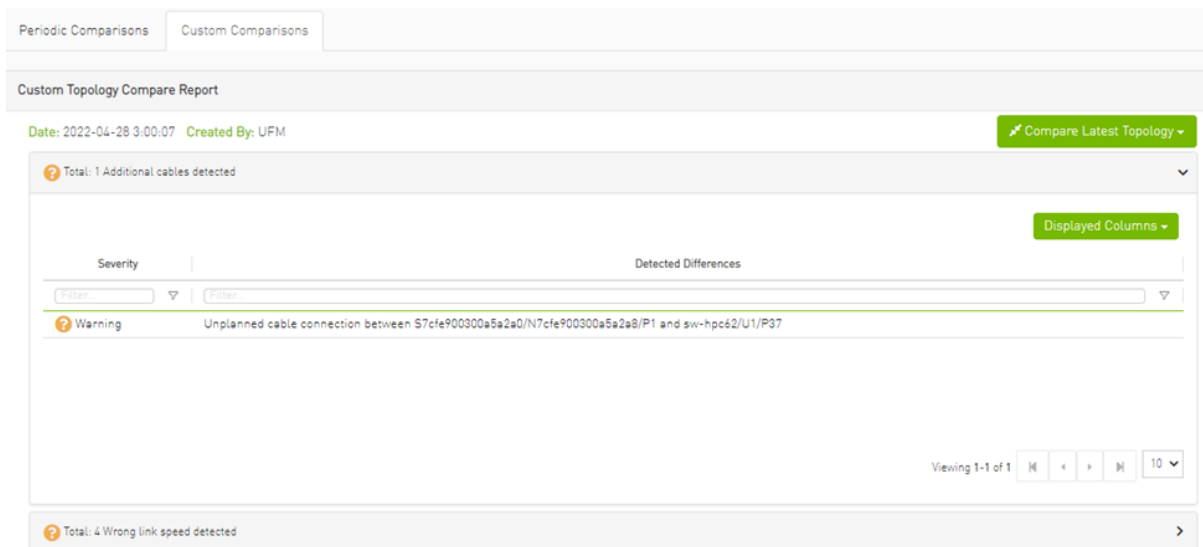
9.7.6.1.2 Custom Comparison

Custom comparison compares user-defined topology with the current fabric topology. UFM compares the current fabric topology to a topology snapshot (of the same setup) and reports any differences between them.

To be able to use the UFM topology comparison mechanism, first you need to create a TOPO file that defines the current topology of the fabric.

i Ideally, the topology snapshot (`.topo` file) should be taken after the setup bring-up phase has been completed so that no more topology changes are expected to take place.

Once the TOPO file is created, you can use the topology comparison mechanism to compare the current fabric topology to the one in the TOPO file and view their differences (if found).



To compare the current topology with the master topology or a custom topology (external file), make a selection from the "Compare Latest Topology" dropdown button and upload the `.topo` file to compare against.

9.7.6.2 Topology Comparison Flow

➤ *To create the topology file for later comparison with the current topology, do the following:*

1. Verify that the following path for ibdiagnet ibnl directory exists: `/opt/ufm/tmp/ibdiagnet.out/tmp/ibdiag_ibnl`. If the path does not exist, make sure to create it manually.
2. Run the following command on the UFM server machine to create the topology file (`mytopo.topo`). Note that the file extension must be `.topo` for UFM to recognize it.


```
/opt/ufm/opensm/bin/ibdiagnet -w /tmp/mytopo.topo
--out_ibnl_dir /opt/ufm/tmp/ibdiagnet.out/tmp/ibdiag_ibnl
```

Once command execution is completed, the new topology file (`/tmp/mytopo.topo`) will be created and can be used for later comparison with the current fabric topology. Also, several `.ibnl` files that were (optionally) created will be found in the defined output directory (`/opt/ufm/tmp/ibdiagnet.out/tmp/ibdiag_ibnl`). These `.ibnl` files will be used when comparing any topology file to the current fabric topology.

At any time during your UFM session, you can view the last generated report through the UFM web UI or in HTML format in a browser window.

➤ *To perform topology comparison, do the following:*

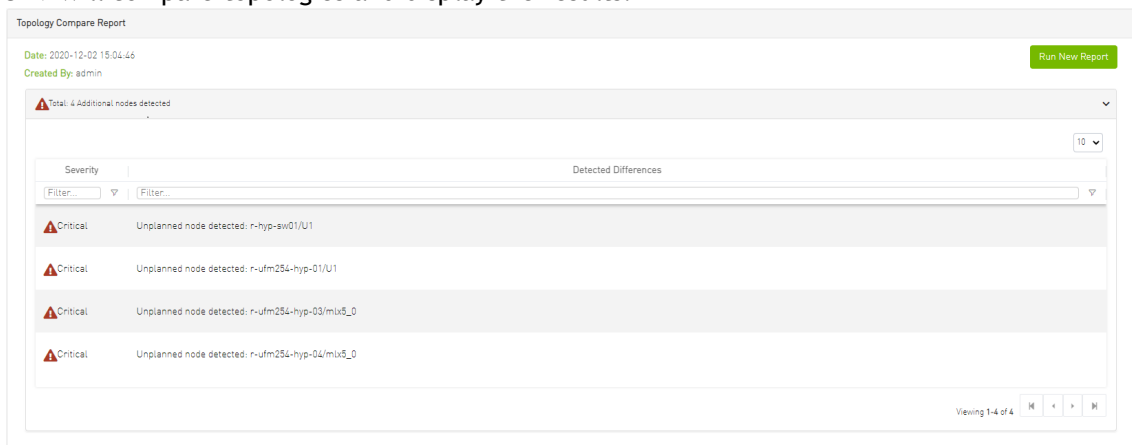
1. Click Run Now Report under System Health à Topology Compare.



2. Browse for the required topology setup file in the *Load Topology File* dialog box.




3. Click Load.
UFM will compare topologies and display the results.



9.7.7 Fabric Validation Tab

The Fabric Validation tab displays the fabric validation tests and gives the ability to run the test and receive/view the summary as a job output. Summary of the job contains all errors and warnings that were found during the test execution.



Test	Description
Check Lids	Checks for bad lids. Possible lid errors are: <ul style="list-style-type: none"> • zero lid • lid duplication
Check Links	Checks for connectivity issues where all ports connected are not in the same state (active)
Check Subnet Manager	Checks for errors related to subnet manager. Possible SM errors are: <ul style="list-style-type: none"> • Failed to get SMInfo Mad • SM Not Found • SM Not Correct (master SM with wrong priority) • Many master SMs exists
Check Duplicate Nodes	Checks for duplications in nodes description
Check Duplicate Guides	Checks for GUIDs duplications
Check Routing	Checks for failures in getting routing MADs
Check Link Speed	Checks for errors related to link speed. Possible link speed errors are: <ul style="list-style-type: none"> • Different speed between ports • Wrong configuration - 'enable' not part of the 'supported' • Unexpected speed
Check Link Width	Checks for errors related to link width. Possible link width errors are: <ul style="list-style-type: none"> • Different width between ports • Wrong configuration - 'enable' not part of the 'supported' • Unexpected width
Check Partition Key	Checks for errors related to PKey. Possible PKey errors are: <ul style="list-style-type: none"> • Failed to get Pkey Tables • Mismatching pkeys between ports
Check Temperature	Checks for failure in getting temperature sensing.

Test	Description
Check Cables	Checks for errors related to cables. Possible cable errors are: <ul style="list-style-type: none"> This device does not support cable info capability Failed to get cable information (provides a reason)
Check Effective BER	Checks that the Effective BER does not exceed the threshold
Dragonfly Topology Validation	Validate if the topology is Dragonfly
SHARP Fabric Validation	Checks for SHARP Configurations in the fabric
Tree Topology Validation	Checks if the fabric is a tree topology
Socket Direct Mode Reporting	Presents the inventory of fabric HCAs that are using socket direct

To run a specific test, click the play button. The job will be displayed once completed.

The screenshot shows a user interface for running tests. On the left, a 'Tests' panel lists various tests, with 'Check Lids' selected. On the right, the 'Check Lids' job details are shown, including a 'Fabric Summary' table with the following data:

Category	Count
Total Nodes	56
IB Switches	15
IB Channel Adapters	30
IB Aggregation Nodes	11
IB Routers	0

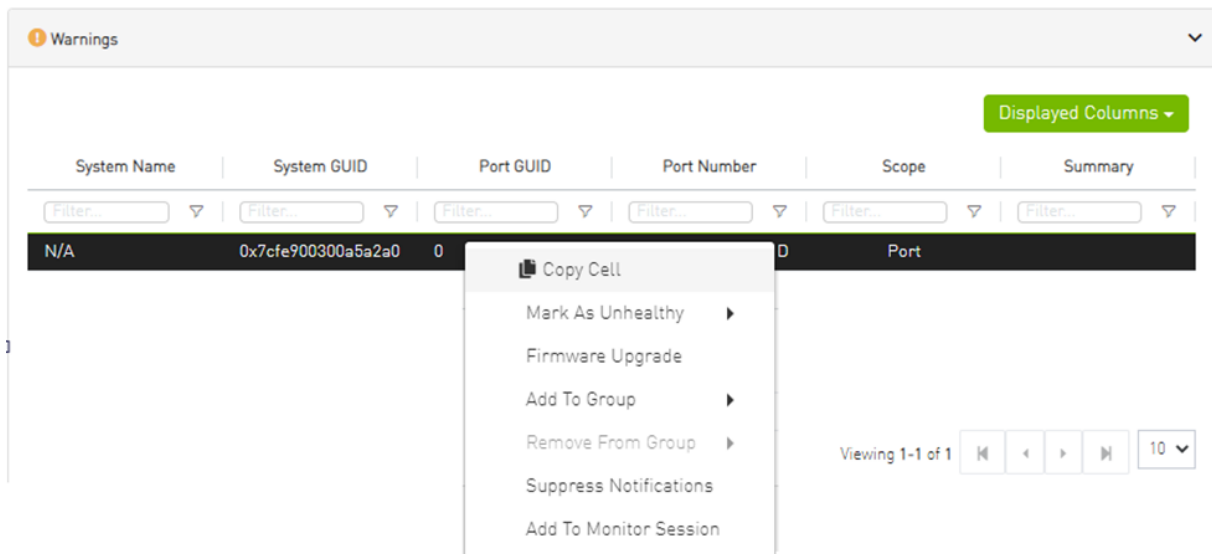

The job will also be displayed in the Jobs window.

9.7.7.1 Add Model Objects to Validation Test

Some validation tests contain data related to devices or ports like device GUID and port GUID.

Depending on that information a context menu for each related device/port can be shown.

If the data is related to a port the context menu will contain both port and device options.



9.7.8 IBDiagnet Tab

The periodic IBDiagnet tab allows users to create scheduled ibdiagnet tasks on their devices using any of the defined parameters.

Users can also configure a remote location (local/remote) to save the ibdiagnet output to. To create a new ibdiagnet command:

1. Click the New button on the top right of the IBDiagnet tab to open the “New IBDiagnet Command” wizard.

New IBDiagnet Command

1 Parameters 2 Run

Name
IBDiagnet_CMD_1609284355963

Category	Status	Flag Name	Value
Filter...		Filter...	
General			
Link Validation			
	<input checked="" type="checkbox"/>	--ls	2.5
	<input type="checkbox"/>	--lw	1x
Port Counters			
	<input type="checkbox"/>	--pc	
	<input checked="" type="checkbox"/>	--pm_pause_time	1
	<input type="checkbox"/>	--per_slvl_cntrs	
	<input type="checkbox"/>	--sc	
	<input type="checkbox"/>	--scr	
	<input type="checkbox"/>	--extended_speeds	SW

Additional Parameters

Type additional flags for ibdiagnet run

Next

2. Select the desired ibdiagnet flags for your command by selecting the listed flags (categories are expandable), or by manually adding the desired flags into the Additional Parameters box below, and then click Next.

New IBDiagnet Command

1 Parameters 2 Run


Name
IBDiagnet_CMD_1601490607733

Category	Status	Flag Name	Value
Filter...			
General			
Link Validation			
	<input checked="" type="checkbox"/>	--ls	2.5
	<input checked="" type="checkbox"/>	--lw	1x
Port Counters			
	<input type="checkbox"/>	--pc	
	<input checked="" type="checkbox"/>	--pm_pause_time	1
	<input type="checkbox"/>	--per_slvl_cntrs	
	<input type="checkbox"/>	--sc	
	<input type="checkbox"/>	--scr	
	<input type="checkbox"/>	--extended_speeds	sw

Additional Parameters

Type additional flags for ibdiagnet run

Next

 It is possible to use the filters at the top of the Category and Flag Name columns in order to search for flags.

3. In the Run screen:
 - a. Select the location of the ibdiagnet results. UFM can export ibdiagnet command run results to a local location on the UFM server, or to a [configurable remote location](#).

- b. Select whether you would like to save this run for later (Save), run it immediately (Save and Run Now), or schedule it for a later time (Schedule) and then click Finish.


New IBDiagnet Command x

1 Parameters
2 Run

Location
 Local Remote

Output Path: /opt/ufm/files/periodicibdignet

Running Mode
 Save
 Save and Run Now
 Schedule


Save

Summary

Previous
Finish

⚠ Note that you can see the summary of your chosen flags for this run in the Summary panel.

You will then be able to see run results on the tab which will display where the output is saved on the server.

Output Path: /opt/ufm/files/periodicibdignet

IBDiagnet ↻

+ New
Displayed Columns ▾
CSV ▾


Name	Task State	Last Run ↓	Last Run Output
IBDiagnet_CMD_1651155713770	Disabled	✓ 28/04/2022 17:22:15	/opt/ufm/files/periodicibdignet/IBDiag...

Viewing 1-1 of 1 ⏪ ⏩ 10 ▾


It is also optional to edit/activate/deactivate/delete a running task using right-click.

Under gv.cfg, it is possible to configure other parameters.

```
[PeriodicIbdiagnet]
# Directory location where outputs are written
periodic_ibdiagnet_dir_location=/opt/ufm/files/periodicIbdiagnet
# Minimum time between two tasks (in minutes)
minimum_task_interval=60
# Maximum number of tasks running simultaneously
max_optional_tasks=5
# Maximum number of outputs to save per task (oldest gets deleted)
max_saved_outputs=5
# Percentage threshold for disk usage from which UFM deletes old task results
disk_usage_threshold=80
```

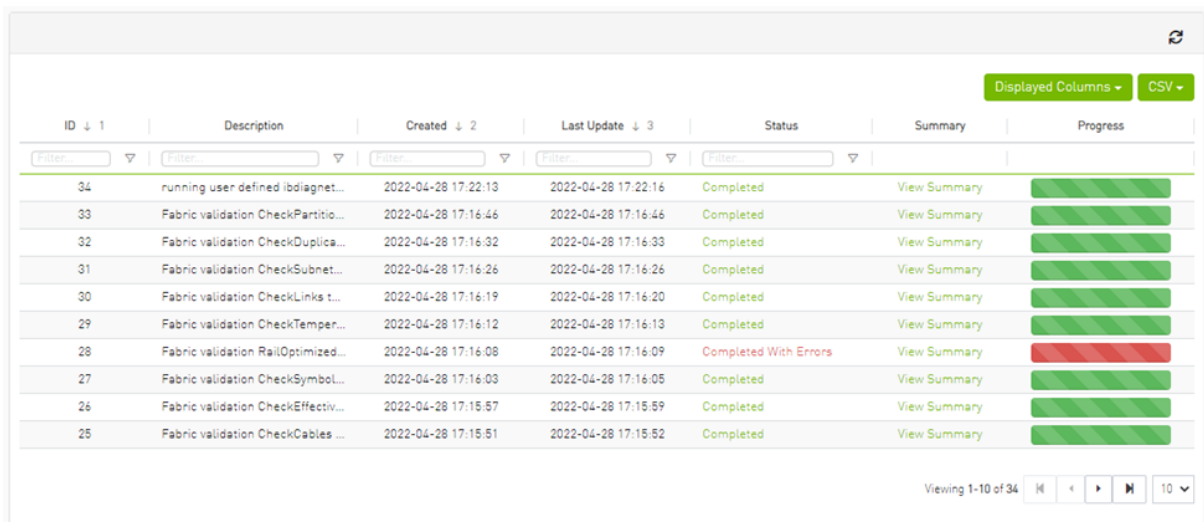
 UFM restart is required for these changes to take effect.











9.8 Jobs

 All information provided in a tabular format in UFM web UI can be exported into a CSV file.

The Jobs window displays all of UFM running Jobs. A Job is a running task defined by the user and applied on one or more of the devices (provisioning, software upgrade, firmware upgrade, reboot, etc.).

UFM users can monitor the progress of a running job, as well as the time it was created, its last update description and its status. The status value can be “Running” (during operation) “Completed with Errors”, in case an error has occurred, and “Completed.”



ID ↓ 1	Description	Created ↓ 2	Last Update ↓ 3	Status	Summary	Progress
34	running user defined ibdiagnet...	2022-04-28 17:22:13	2022-04-28 17:22:16	Completed	View Summary	
33	Fabric validation CheckPartitio...	2022-04-28 17:16:46	2022-04-28 17:16:46	Completed	View Summary	
32	Fabric validation CheckDuplica...	2022-04-28 17:16:32	2022-04-28 17:16:33	Completed	View Summary	
31	Fabric validation CheckSubnet...	2022-04-28 17:16:26	2022-04-28 17:16:26	Completed	View Summary	
30	Fabric validation CheckLinks t...	2022-04-28 17:16:19	2022-04-28 17:16:20	Completed	View Summary	
29	Fabric validation CheckTemper...	2022-04-28 17:16:12	2022-04-28 17:16:13	Completed	View Summary	
28	Fabric validation RailOptimized...	2022-04-28 17:16:08	2022-04-28 17:16:09	Completed With Errors	View Summary	
27	Fabric validation CheckSymbol...	2022-04-28 17:16:03	2022-04-28 17:16:05	Completed	View Summary	
26	Fabric validation CheckEffectiv...	2022-04-28 17:15:57	2022-04-28 17:15:59	Completed	View Summary	
25	Fabric validation CheckCables ...	2022-04-28 17:15:51	2022-04-28 17:15:52	Completed	View Summary	

When selecting a job from the main Jobs table, its related sub jobs will be displayed in the Sub Jobs table below.

ID ↓ 1	Description	Created ↓ 2	Last Update ↓ 3	Status	Summary	Progress
34	running user defined ibdiagnet...	2022-04-28 17:22:13	2022-04-28 17:22:16	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
33	Fabric validation CheckPartitio...	2022-04-28 17:16:46	2022-04-28 17:16:46	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
32	Fabric validation CheckDuplica...	2022-04-28 17:16:32	2022-04-28 17:16:33	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
31	Fabric validation CheckSubnet...	2022-04-28 17:16:26	2022-04-28 17:16:26	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
30	Fabric validation CheckLinks t...	2022-04-28 17:16:19	2022-04-28 17:16:20	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
29	Fabric validation CheckTemper...	2022-04-28 17:16:12	2022-04-28 17:16:13	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
28	Fabric validation RailOptimized...	2022-04-28 17:16:08	2022-04-28 17:16:09	Completed With Errors	View Summary	<div style="width: 100%; height: 10px; background-color: red;"></div>
27	Fabric validation CheckSymbol...	2022-04-28 17:16:03	2022-04-28 17:16:05	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
26	Fabric validation CheckEffectiv...	2022-04-28 17:15:57	2022-04-28 17:15:59	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>
25	Fabric validation CheckCables ...	2022-04-28 17:15:51	2022-04-28 17:15:52	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>

Viewing 1-10 of 34

ID ↓ 1	Related Object	Description	Created ↓ 2	Last Update ↓ 3	Status	Summary	Progress
34.1	Site	running user defi...	2022-04-28 17:22:13	2022-04-28 17:22:16	Completed	View Summary	<div style="width: 100%; height: 10px; background-color: green;"></div>

9.9 Settings



All information provided in a tabular format in UFM web UI can be exported into a CSV file.

This window enables configuring the following UFM server and fabric-related settings:

- [Events Policy](#)
- [Device Access](#)
- [Network Management](#)
- [Subnet Manager Tab](#)
- [Non-Optimal Links](#)
- [User Management Tab](#)
- [Email](#)
- [Remote Location](#)
- [Data Streaming](#)
- [Topology Compare](#)
- [Token-based Authentication](#)
- [Plugin Management](#)
- [User Preferences](#)

9.9.1 Events Policy

The Events Policy tab allows you to define how and when events are triggered for effective troubleshooting and fabric maintenance.

Event	Category	Mail	GUI	Alarm	Syslog	Log File	SNMP	Threshold	TTLSec	Severity
GID Address In Service		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1	300	Info
GID Address Out of Se...		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1	300	Warning
New MCast Group Cre...		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1	300	Info
MCast Group Deleted		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1	300	Info
Symbol Error		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	200	300	Warning
Link Error Recovery		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1	300	Minor
Link Down		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	0	300	Warning
Port Receive Errors		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	5	300	Warning
Port Receive Remote ...		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	5	300	Minor
Port Receive Switch R...		<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	9999	300	Minor

Events are reported by setting the following parameters:

Option	Description/Instructions
Event	Event description.
Category	Event category, such as Communication Error and Hardware represented by icons.
Mail	When selected, the corresponding events will be sent a list of recipients according to Configuring Email-on-Events .
Web UI	When selected, the corresponding events are displayed in the Events & Alarms window in the Web UI.
Alarm	Select the Alarm option to trigger an alarm for a specific event. When selected, the alarms will appear in the Events & Alarms window in the Web UI.
Syslog	When checked along with the Log file option, the selected events will be written to Syslog.
Log File	Select the Log File option if you would like the selected event to be reported in a log file.
SNMP	The UFM Server will send events to third-party clients by means of SNMP traps. Select the event SNMP check box option to enable the system to send an SNMP trap for the specific event. The SNMP trap will be sent to the port defined in Configuration file located under: /opt/ufm/conf/gv.cfg. For further information, refer to SNMP Settings (Advanced License Only) .
Threshold	An event will be triggered when the traffic/error rate exceeds the defined threshold. For example: when PortXmit Discards is set to 5 and the counter value grows by 5 units or more between two sequential reads, an event is generated.
TTL (Sec)	TTL (Alarm Time to Live) sets the time during which the alarm on the event is visible on UFM Web UI. TTL is defined in seconds. CAUTION: Setting the TTL to 0 makes the alarm permanent, meaning that the alarm does not disappear from the Web UI until cleared manually.
Action	The action that will be executed in case the event which has triggered the action can be none or isolated (make the port unhealthy or isolated). This attribute can be set only for ports event policy.
Severity	Select the severity level of the event and its alarm from the drop-down list: Info, Warning, Minor, and Critical.

Additional Events Policy Table Options (for Advanced License)

Option	Description/Instructions
SNMP	

- ⚠ • Category column in the Events Policy table indicates to which category the event belongs. These categories are defined in the event configuration file and cannot be modified. Categories are: Hardware, Fabric Configuration, Communication Error, Fabric Notification, Maintenance, Logical Model, Fabric Topology, Gateway, Module Status, and UFM Server.
- Event logs can still be checked even if the events.log file checkbox was not checked during Syslog configuration.
- For a certain event to be sent to Syslog, both the Syslog and the Log File checkboxes must be checked. Otherwise, the selected events will not be sent to Syslog.

See [Appendix - Supported Port Counters and Events](#) for detailed information on port counters and events.

9.9.1.1 SNMP Settings (Advanced License Only)

When UFM is running, the Web UI Policy Table shows the SNMP traps. You can then modify and save an SNMP Trap flag for each event. SNMP settings are enabled only after the installation of the UFM Advanced license.

UFM sends SNMP Trap using version SNMPV2 to the default port 162.

➤ *To set the SNMP properties:*

1. Open the `/opt/ufm/conf/gv.cfg` configuration file.
2. Under the [Notifications] line (see the following example):
 - a. Set the (snmp_listeners) IP addresses and ports
 - b. Port is optional - the default port number is 162
 - c. Use a comma to separate multiple listeners

Format:

```
snmp_listeners = <IP Address 1>[:<port 1>][,<IP Address 2>[:<port 2>]...]
```

Example:

```
[Notifications]
snmp_listeners = host1, host2:166
```

9.9.1.2 Configuring Email-on-Events

UFM enables you to configure each event to be sent by email to a list of pre-defined recipients. Every 5 minutes (configurable) UFM will collect all “Mail” selected events and send them to the list of pre-defined recipients. By default, the maximum number of events which can be sent in a single email is 100 (configurable, should be in the range of 1-1000)

The order of events in the email body can be set as desired. The available options are: order by severity or order by time (by default: order by severity)


➤ To change email-on-events setting, do the following:

1. Edit the `/opt/ufm/conf/gv.cfg` file.
2. Go to section “[Events]” and set the relevant parameters:
 - `sending_interval` (default=5)—Time interval for keeping events (minimum 10 seconds, maximum 24 hours)
 - `sending_interval_unit` (default = minute)—Optional units: minute, second, hour
 - `cyclic_buffer` (default=false)—If the cyclic buffer is set to true, older events will be dropped, otherwise newer events will be dropped (if reaches max count)
 - `max_events` (default=100)—Maximum number of events to be sent in one mail (buffer size), should be in the range of 1-1000
 - `group_by_severity` (default=true)—Group events in mail by severity or by time

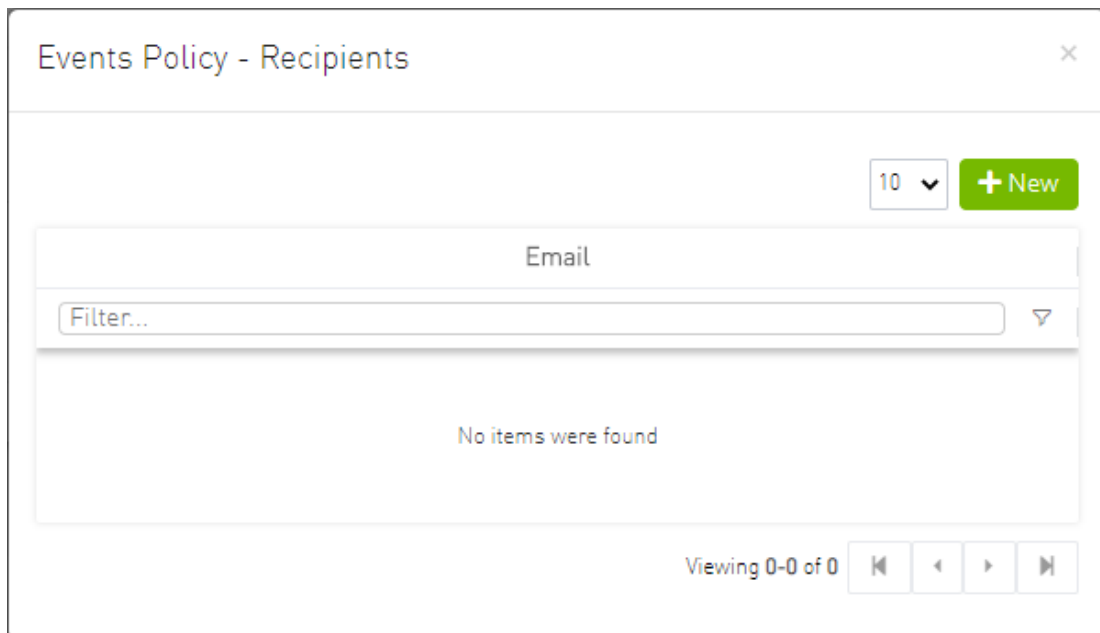
➤ To receive the email-on-events, do the following:

⚠ Configure SMTP settings under Settings window → Email tab - see [Email Tab](#).

1. Configure the Recipients List under Settings → Events Policy.



2. Click New.



3. In the Recipients List window, enter valid recipient email addresses, comma-separated, and click Submit.

New Recipients

Recipients

comma separated email addresses list

Close Submit

The new recipients are then added to the Events Policy Recipients list. These recipients automatically start receiving emails on the events for which the Mail checkbox is checked in the table under Events Policy.

9.9.1.3 Security

9.9.2 Device Access

The screenshot shows the 'Settings' page with the 'Device Access' tab selected. The 'Switch SSH' section is expanded, showing the following configuration options:

- Credentials:**
 - User:
 - Password:
 - Confirmation:
- Connection:**
 - Port:
 - Timeout:

An 'Update' button is located at the bottom right of the 'Switch SSH' section. Below this section are three collapsed sections: 'Server SSH', 'HTTP', and 'IPMI', each with a right-pointing arrow.

You can configure default access parameters for remote administration via the following protocols:

- Switch/Server SSH - allows you to define the SSH parameters to open an SSH session on your device
 - IPMI - allows you to set the IPMI parameters to open an IPMI session on your device for remote power control
 - HTTP - allows you to define the HTTP parameters to open an HTTP session on your device
- Default credentials are applicable to all switches and servers in the fabric.



The default SSH (CLI) switch credentials match the Grid Director series switch. To change the credentials for IS5030/IS5035 edit the [SSH_Switch] section in the gv.cfg file.

Define access parameters for the remote user as described in the following table.

Site Access Credential Parameters

Parameter	Description
User	The name of the user allowed remote access.
Password	Enter the user password.
Confirmation	Re-enter the password.
Port	Each communication protocol has a default port for connection. You can modify the port number, if required.
Timeout	Each communication protocol has a default timeout, i.e. the maximum time, in seconds, to wait for a response from the peer. You can modify the timeout, if required.

9.9.3 Network Management

UFM achieves maximum performance with latency-critical tasks by implementing traffic isolation, which minimizes cross-application interference by prioritizing traffic to ensure critical applications get the optimal service levels.

9.9.3.1 UFM Routing Protocols

UFM web UI supports the following routing engines:

- MINHOP - based on the minimum hops to each node where the path length is optimized (i.e., shortest path available).
- UPDN - also based on the minimum hops to each node but it is constrained to ranking rules. Select this algorithm if the subnet is not a pure Fat Tree topology and deadlock may occur due to a credit loops in the subnet.
- DNUP - similar to UPDN, but allows routing in fabrics that have some channel adapter (CA) nodes attached closer to the roots than some switch nodes.
- File-Based (FILE) - The FILE routing engine loads the LFTs from the specified file, with no reaction to real topology changes.
- Fat Tree - an algorithm that optimizes routing for congestion-free "shift" communication pattern.

Select Fat Tree algorithm if a subnet is a symmetrical or almost symmetrical fat-tree. The Fat Tree also optimizes K-ary-N-Trees by handling non-constant K in cases where leafs (CAs) are not fully staffed, and the algorithm also handles any Constant Bisectional Bandwidth (CBB) ratio. As with the UPDN routing algorithm, Fat Tree routing is constrained to ranking rules.

- Quasi Fat Tree - PQFT routing engine is a closed formula algorithm for two flavors of fat trees
- Quasi Fat Tree (QFT)
- Parallel Ports Generalized Fat Tree (PGFT)
PGFT topology may use parallel links between switches at adjacent levels, while QFT uses parallel links between adjacent switches in different sub-trees. The main motivation for that is the need for a topology that is not just optimized for a single large job but also for smaller concurrent jobs.
- Dimension Order Routing (DOR) - based on the Min Hop algorithm, but avoids port equalization, except for redundant links between the same two switches. The DOR algorithm

provides deadlock-free routes for hypercubes, when the fabric is cabled as a hypercube and for meshes when cabled as a mesh.

- Torus-2QoS - designed for large-scale 2D/3D torus fabrics. In addition, you can configure Torus-2QoS routing to be *traffic aware*, and thus optimized for neighbor-based traffic.
- Routing Engine Chain (Chain) - an algorithm that allows configuring different routing engines on different parts of the IB fabric.
- Adaptive Routing (AR) - enables the switch to select the output port based on the port's load. This option is not available via UFM Web UI.
 - AR_UPDN
 - AR_FTREE
 - AR_TORUS
 - AR_DOR
- Dragonfly+ (DFP, DPF2)

9.9.3.2 Configuring Routing Protocol

Network Management tab enables setting the preferred routing protocol supported by the UFM software, as well as routing priority.

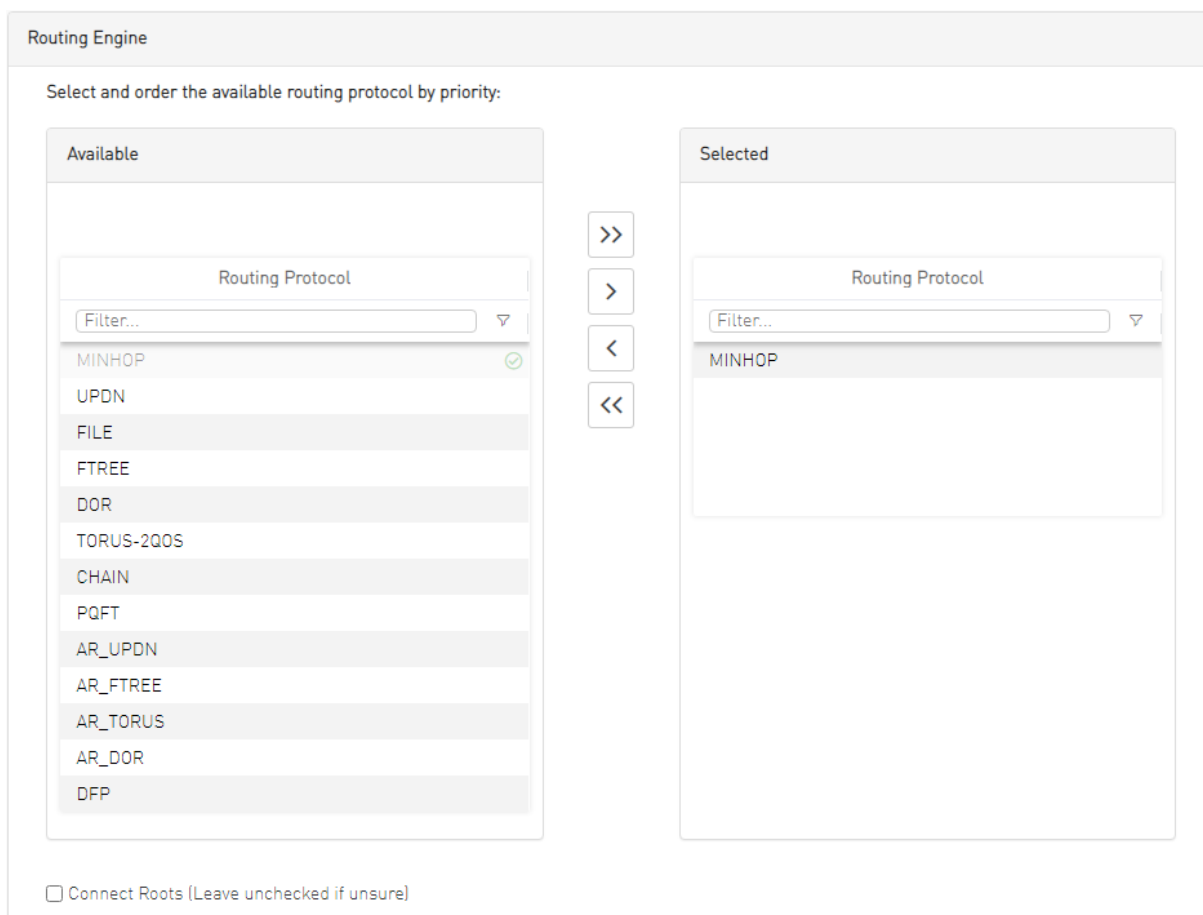
To set the desired routing protocol, move one routing protocol or more from the Available list to the Selected list, and click "Save" in the upper right corner.

Routing Information	
Lid Matrix Dump File	/opt/ufm/files/conf/opensm/lid_matrix.conf
LFTS File	/opt/ufm/files/conf/opensm/lfts.conf
Root Guid File	/opt/ufm/files/conf/opensm/root_guid.conf
Compute Nodes File	N/A
Node IDs File	N/A
Guid Routing Order File	N/A
Active Routing Engine	minhop

The protocol at the top of the list has the highest priority and will be chosen as the Active Routing Engine. If the settings for this protocol are not successful, UFM takes the next available protocol.

Routing Information is listed on the top of the screen:

Field/Box	Description
LID Matrix Dump File	File holding the LID matrix dump configuration
LFTS File	File holding the LFT routing configuration
Root GUID File	File holding the root node GUIDS (for fat-tree or Up/Down)
Compute Nodes File	File holding GUIDs of compute nodes for fat-tree routing algorithm
GUID Routing Order File	File holding the routing order GUIDs (for MinHop and Up/Down)
Node IDs File	File holding the node IDs
Active Routing Engine	The current active routing algorithm used by the managing OpenSM



9.9.4 Subnet Manager Tab

UFM is a management platform using a user-space application for InfiniBand fabric management. This application is developed within the context of an open-source environment. This application serves as an InfiniBand Subnet Manager and a Subnet Administration tool.

The UFM Subnet Manager (SM) is a centralized entity running on the server that discovers and configures all the InfiniBand fabric devices to enable traffic flow throughout the fabric.

To view and configure SM parameters in the *Subnet Manager* tab, select the relevant tab according to the required configuration.

For more information, please refer to [Appendix - Enhanced Quality of Service](#).

9.9.4.1 SM Keys Configuration

The SM Keys tab enables you to view the Subnet Manager Keys. You cannot change the configuration in this tab.

Keys	MKey	0x 0
Limits	SA Key	0x 1
Lossy	Subnet Prefix	0x fe80000000000000
SL2VL	SM Key	0x 1
Sweep	MKey Lease Period	60 (sec)
Handover	LMC	0
Threading	No Partition Enforcement	false
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

Field	Description	Default
MKey	A field that allows you to view or edit the M_Key value sent to all ports to qualify all the set (PortInfo). Authentication is performed by the management entity at the destination port and is achieved by comparing the key contained in the SMP with the key (the M_Key Management key) residing at the destination port.	0x0000000000000000
SA Key	Shows the SM_Key value to qualify the receive SA queries as 'trusted'.	0x0000000000000001
Subnet Prefix	An identifier of the subnet. The subnet prefix is used as the most significant 64 bit of the GID of each InfiniBand node in the subnet.	0xfe80000000000000
SM Key	Read-only field that displays the Key of the Subnet Manager (SM).	0x0000000000000001
MKey Lease Period	A field that allows you to view or edit the lease period used for the M_Key on this subnet in [sec].	0
LMC	Defines the LID Mask Control value for the SM. Possible values are 0 to 7. LID Mask Control (LMC) allows you to assign more than one LID per port. NOTE: Changes to the LMC parameter require a UFM restart.	0
No Partition Enforcement	Disables partition enforcement by switches.	Disabled

9.9.4.2 SM Limits Configuration

The SM Limits tab enables you to view and set the Subnet Manager Limits.


Keys	Packet Life Time	0x 12
Limits	Subnet Timeout	18
Lossy	Maximal Operational VL	VL0-VL3
SL2VL	Head Of Queue Life Time	0x 12
Sweep	Leaf Head Of Queue Life Time	0x 10
Handover	VL Stall Count	0x 7
Threading	Leaf VL Stall Count	0x 7
Logging	Force Link Speed	Max Supported
Misc	Local Physical Error Threshold	0x 8
QoS	Overrun Errors Threshold	0x 8
Congestion Control		
Adaptive Routing		

To configure SM Limits, set the fields as described in the table below, and click “Save.”

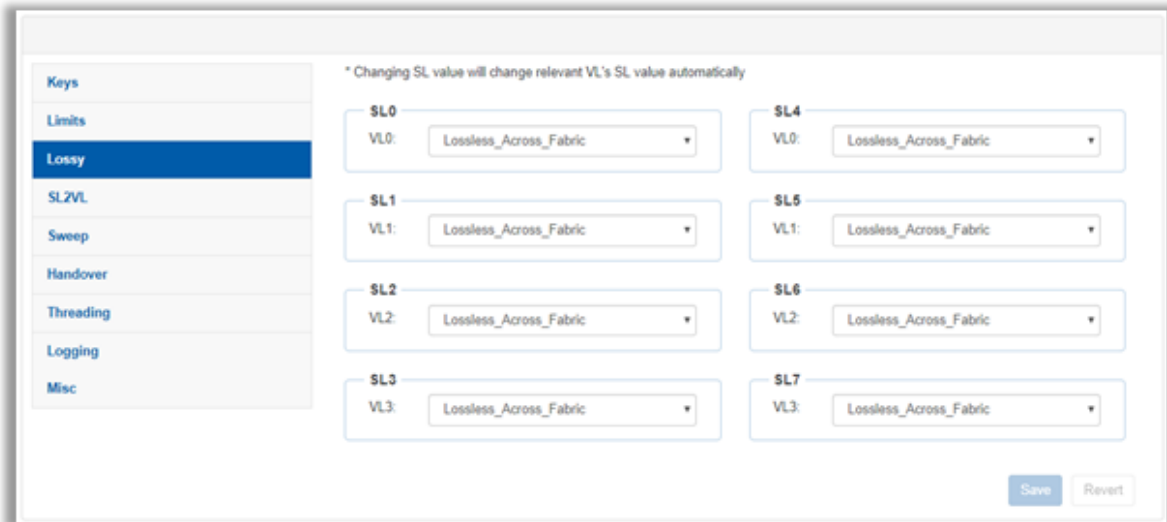
Field	Description	Default
Packet Life Time	A field that allows you to view and/or edit the code of maximum lifetime a packet in a switch. The actual time is $4.096 \text{ usec} * 2^{\langle \text{packet_life_time} \rangle}$. The value 0x14 disables this mechanism	0x12
Subnet Timeout	A field that allows you to view and/or edit the subnet_timeout code that will be set for all the ports. The actual timeout is $4.096 \text{ usec} * 2^{\langle \text{subnet_timeout} \rangle}$	18
Maximal Operational VL	A field that allows you to view and/or edit the limit of the maximal operational VLS: <ul style="list-style-type: none"> • 0: NO_CHANGE • 1: VL0 1 • 2: VL0_VL1 • 3: VL0_VL3 • 4: VL0_VL7 • 5: VL0_VL14 	3
Head of Queue Life Time	A field that allows you to view and/or edit the code of maximal time a packet can wait at the head of transmission queue. The actual time is $4.096 \text{ usec} * 2^{\langle \text{head of queue lifetime} \rangle}$. The value 0x14 disables this mechanism.	0x12
Leaf Head of Queue Life Time	A field that allows you to view and/or edit the maximum time a packet can wait at the head of queue on a switch port connected to a CA or gateway port.	0x10
VL Stall Count	A field that allows you to view the number of sequential packets dropped that cause the port to enter the VLStalled state. The result of setting this value to zero is undefined.	0x07

Field	Description	Default
Leaf VL Stall Count	This field allows you to view the number of sequential packets dropped that cause the port to enter the VLStalled state. This value is for switch ports driving a CA or gateway port. The result of setting the parameter to zero is undefined.	0x07
Force Link Speed	A parameter that allows you to modify the PortInfo:LinkSpeedEnabled field on switch ports. If 0, do not modify. <ul style="list-style-type: none"> Values are: 1: 2.5 Gbps 3: 2.5 or 5.0 Gbps 5: 2.5 or 10.0 Gbps 7: 2.5 or 5.0 or 10.0 Gbps 2,4,6,8-14 Reserved 15: set to PortInfo:LinkSpeedSupported 	15 By default, UFM sets the enabled link speed equal to the supported link speed.
Local Physical Error Threshold	A field that allows you to view and/or edit the threshold of local phy errors for sending Trap 129.	0x08
Overrun Errors Threshold	A field that allows you to view and/or edit the threshold of credit overrun errors for sending Trap 130.	0x08

9.9.4.3 SM Lossy Manager Configuration

 This tab is available to users with an advanced license only.

The SM Lossy tab enables you to view and set the Lossy Configuration Manager options after Lossy Configuration has been enabled.



9.9.4.4 SM SL2VL Mapping Configuration

The SM SL2VL tab enables you to view the SL (service level) to VL (virtual lane) mappings and the configured Lossy Management. You cannot change the configuration in this tab.

However, you can change it in the previous [SM Lossy Manager Configuration \(Advanced License only\)](#) tab.

Keys	
Limits	
Lossy	
SL2VL	
Sweep	
Handover	
Threading	
Logging	
Misc	
QoS	
Congestion Control	
Adaptive Routing	

Qos Option Type	SL0	SL1	SL2	SL3	SL4	SL5	SL6	SL7
Default	0	1	2	3	0	1	2	3
Hca	0	1	2	3	0	1	2	3
Switch Port 0	0	1	2	3	0	1	2	3
Switch External Ports	0	1	2	3	0	1	2	3
Router	0	1	2	3	0	1	2	3

9.9.4.5 SM Sweep Configuration

The Sweep tab enables you to view and/or set the Subnet Manager Sweep Configuration parameters.

Keys	
Limits	
Lossy	
SL2VL	
Sweep	
Handover	
Threading	
Logging	
Misc	
QoS	
Congestion Control	
Adaptive Routing	

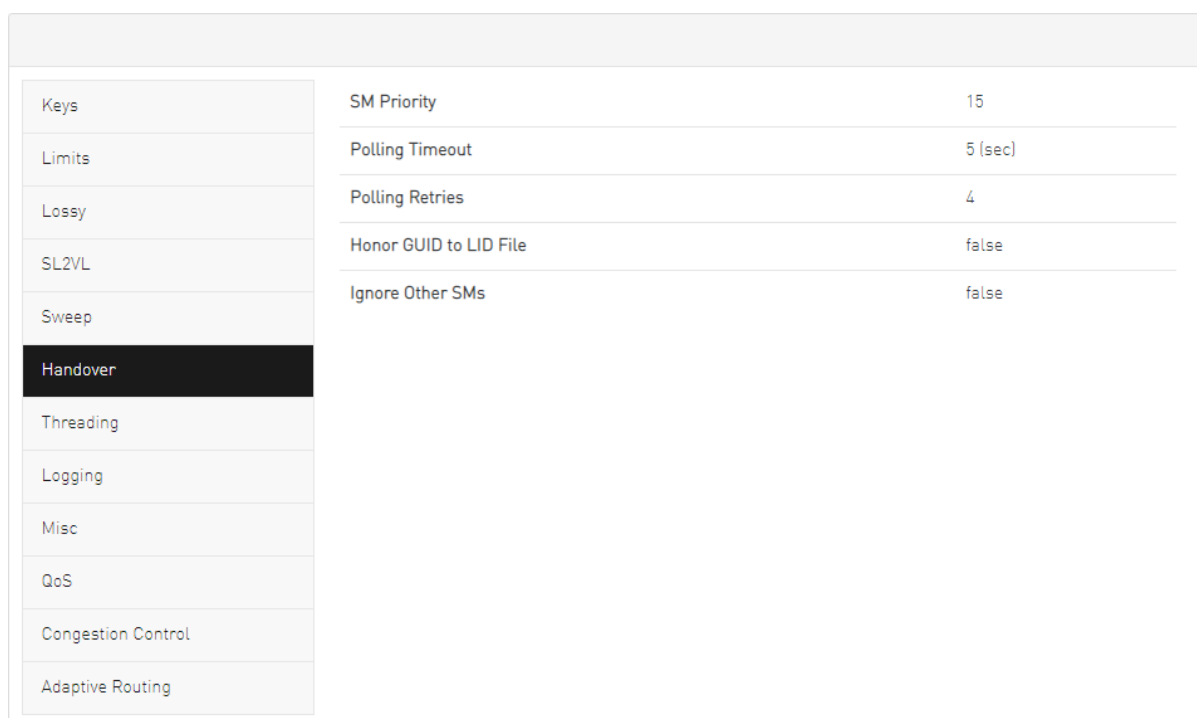
Sweep Interval	<input type="text" value="10"/>	seconds
Reassign Lids	<input type="checkbox"/>	
Sweep On Trap	<input checked="" type="checkbox"/>	
Force Heavy Sweep	<input type="checkbox"/>	false

To configure SM Sweep, set the fields as described in the table below and click "Save."

Field/Box	Description	Default
Sweep Interval	A field that allows you to view and/or edit the number of seconds between light sweeps (0 disables it).	10
Reassign LIDs	If enabled, causes all LIDs to be reassigned.	Disabled
Sweep on Trap	If enabled, traps 128 and 144 will cause a heavy sweep.	Enabled
Force Heavy Sweep	If enabled, forces every sweep to be a heavy sweep.	Disabled

9.9.4.6 SM Handover Configuration

The SM Handover tab enables you to view the Subnet Manager Handover Configuration parameters. You cannot change the configuration in this tab.



Keys	SM Priority	15
Limits	Polling Timeout	5 (sec)
Lossy	Polling Retries	4
SL2VL	Honor GUID to LID File	false
Sweep	Ignore Other SMs	false
Handover		
Threading		
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

Field/Box	Description	Default
SM Priority	A field that shows the SM priority used for determining the master. Range is 0 (lowest priority) to 15 (highest). Note: Currently, these settings may not be changed.	15
Polling Timeout	A field that shows the timeout in [sec] between two polls of active master SM.	Range=10000
Polling Retries	Number of failing polls of remote SM that declares it "not operational."	4
Honor GUID to LID File	If enabled, honor the guid2lid file when coming out of standby state, if the file exists and is valid.	Disabled
Ignore other SMs	If enabled, other SMs on the subnet are ignored.	Disabled

9.9.4.7 SM Threading Configuration

The SM Threading tab enables you to view the Subnet Manager Timing and Threading Configuration parameters. You cannot change the configuration in this tab.

Keys	Max Wire SMPs	8
Limits	Transaction Timeout	200 (ms)
Lossy	Max Message FIFO Timeout	10000
SL2VL	Single Thread	false
Sweep		
Handover		
Threading		
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

Field/Box	Description	Default
Max Wire SMPs	A field that shows the maximum number of SMPs sent in parallel.	4
Transaction Timeout	A field that shows the maximum time in [msec] allowed for a transaction to complete.	200
Max Message FIFO Timeout	A field that shows the maximum time in [msec] a message can stay in the incoming message queue.	10000
Single Thread	When enabled, a single thread is used for handling SA queries.	Disabled

9.9.4.8 SM Logging Configuration

The SM Logging tab enables you to view and/or set the Subnet Manager Logging Configuration parameters.

Keys	Log File	/opt/ufm/files/log/opensm.log
Limits	Log Max Value	<input type="text" value="4096"/> (MB)
Lossy	Dump Files Directory	/opt/ufm/files/log/
SL2VL	Force Log Flush	<input type="checkbox"/>
Sweep	Accumulate Log File	<input checked="" type="checkbox"/>
Handover	Log Levels	<input checked="" type="checkbox"/> Error <input checked="" type="checkbox"/> Info <input type="checkbox"/> Verbose <input type="checkbox"/> Debug <input type="checkbox"/> Funcs <input type="checkbox"/> Frames <input type="checkbox"/> Routing <input type="checkbox"/> Sys
Threading		
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

To configure SM Logging, set the fields as described in the table below and click “Save.”

Field/Box	Description	Default
Log File	Path of the Log file to be used.	cond/opt/ufm/files/log/opensm.log
Log Max Size	A field that allows you to view and/or edit the size limit of the log file in MB. If overrun, the log is restarted.	4096
Dump Files Directory	The directory that holds the SM dump file.	/opt/ufm/files/log
Force Log Flush	Force flush to the log file for each log message.	Disabled
Accumulate Log File	If enabled, the log accumulates over multiple SM sessions.	Enabled
Log Levels	Available log levels: Error, Info, Verbose, Debug, Funcs, Frames, Routing, and Sys.	Error and Info

9.9.4.9 SM Miscellaneous Settings


The Misc tab enables you to view additional Subnet Manager Configuration parameters. You cannot change the configuration in this tab.

Keys	Node Names Map File	N/A
Limits	SA Database File	N/A
Lossy	No Clients Reregistration	false
SL2VL	Disable MultiCast	false
Sweep	Exit On Fatal Event	true
Handover		
Threading		
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

Field/Box	Description	Default
Node Names Map File	A field that allows you to view and/or set the node name map for mapping nodes to more descriptive node descriptions.	None
SA Database File	SA database file name	None
No Clients Reregistration	If enabled, disables client re-registration.	Disabled
Disable Multicast	If enabled, the SM disables multicast support and no multicast routing is performed.	Disabled
Exit on Fatal Event	If enabled, the SM exits on fatal initialization issues.	Enabled

9.9.4.10 SM QoS Configuration

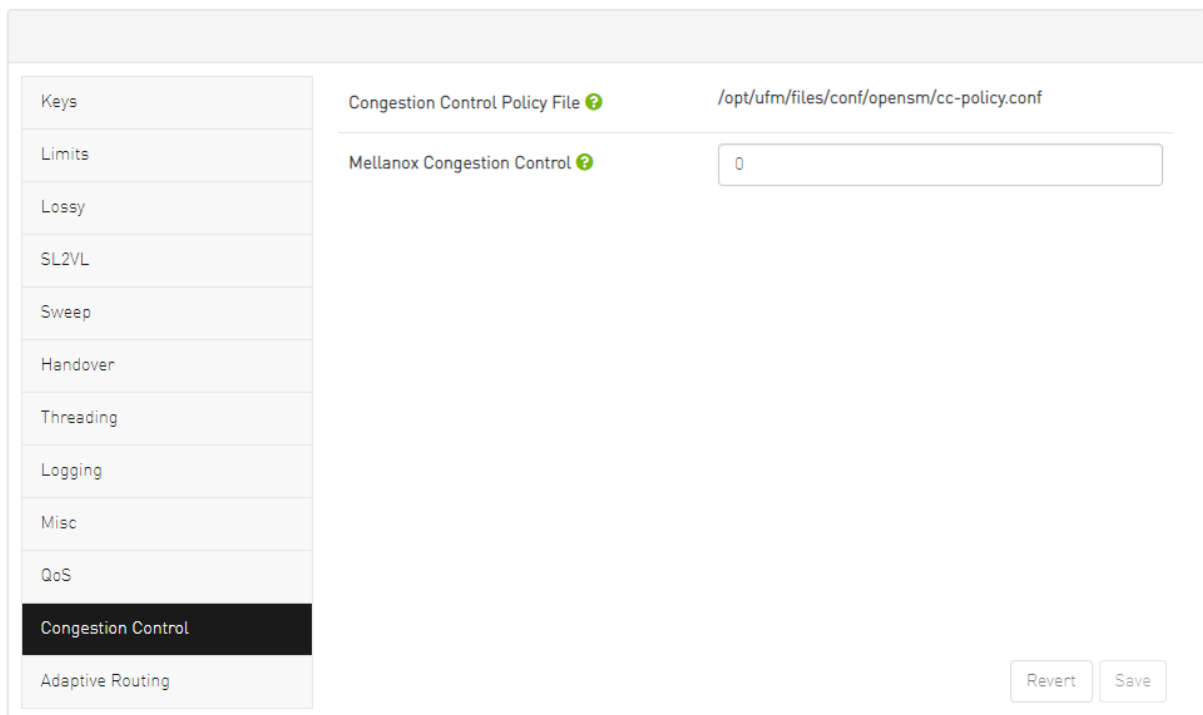
The QoS tab allows you to enable or disable QoS functionality. QoS is disabled by default.



9.9.4.11 SM Congestion Control Configuration

The Congestion Control tab allows you to enable, disable, or ignore congestion control.

- 0 - Ignore (default)
- 1 - Enable
- 2 - Disable



9.9.4.12 SM Adaptive Routing Configuration

The Adaptive Routing tab allows you to configure adaptive routing parameters.

Keys	DFP Down Up Turns Mode	<input type="text" value="0"/>
Limits		
Lossy	DFP Max Cas On Spine	<input type="text" value="2"/>
SL2VL		
Sweep		
Handover	Adaptive Routing SL Mask	<input type="text" value="0x FFFF"/>
Threading		
Logging		
Misc		
QoS		
Congestion Control		
Adaptive Routing		

Revert Save

9.9.5 Non-Optimal Links

A non-optimal link is a link between two ports that is configured to operate at a certain speed and width and is operating at a lower rate. The Non-optimal links feature helps you identify potential link failures and reduce fabric inefficiencies.

Non-optimal links can be any of the following:

- NDR links that operate in HDR, EDR, FDR, QDR, DDR or SDR mode
- HDR links that operate in EDR, FDR, QDR, DDR or SDR mode
- EDR links that operate in FDR, QDR, DDR or SDR mode
- FDR links that operate in QDR, DDR or SDR mode
- QDR links that operate in DDR or SDR mode
- 4X links that operate in 1X mode

The Non-Optimal Links window allows you to set the preferred action for non-optimal links.

Settings

Events Policy Device Access Network Management Subnet Manager **Non-Optimal Links**

Non-optimal Links Configuration

Non-optimal link is a link that is configured to operate in certain speed and width and is operating in a lower rate. This helps to identify potential link failures and helps reduce fabric inefficiencies.

Non-optimal Links Behavior:

Reset all Non-optimal Links

Disable all Non-optimal Links

To set the non-optimal links policy:

From the drop-down menu, select the action for Non-optimal Links behavior.

The drop-down menu defines the default behavior. Options are: Ignore (default), Disable, and Reset.

Option	Description
Ignore	Ignore the non-optimal links
Reset	Reset all non-optimal links ports
Disable	Disable all non-optimal links ports

Reset all Non-Optimal Links allows users to reset all current non-optimal links ports on-demand.

Disable all Non-Optimal Links allows users to disable all current non-optimal links ports on-demand.

9.9.6 User Management Tab

UFM User Authentication is based on standard Apache User Authentication. Each Web Service client application must authenticate against the UFM Server to gain access to the system. UFM implements any kind of third-party authentication supported by the Apache Web Server.

The default user (admin) has System Administration rights. A user with system Administration rights can manage other users' accounts, including creation, deletion, and modification of accounts. The system's default user is the admin user.

 To add a new user account, do the following:

1. Click the “New” button.

The screenshot shows a web interface with a navigation bar at the top containing tabs for "Events Policy", "Device Access", "Network Management", "Subnet Manager", "Non-Optimal Links", and "User Management". Below the navigation bar, there are two sub-tabs: "Topology Compare" and "Access Tokens". The main content area displays a table with the following structure:

ID ↓	Name	Group
1	admin	System Admin

At the top left of the table area is a green "+ New" button. At the top right is a "Displayed Columns" dropdown menu. Below the header row, there are three filter input fields, each labeled "Filter...". At the bottom of the table area, there is a pagination control showing "Viewing 1-1 of 1" and navigation buttons for first, previous, next, and last, along with a page size dropdown set to "10".

2. Fill in the required fields in the dialog box.

The screenshot shows a dialog box titled "Create A User" with a close button (X) in the top right corner. The dialog contains four input fields:


- User Name**: A text input field.
- Group**: A dropdown menu with "System Admin" selected.
- Password**: A text input field.
- Confirm Password**: A text input field.

A green "Create" button is located at the bottom right of the dialog box.

Each user can be assigned to one of the following Group (role) options:

- System Admin - users can perform all operations including managing other users accounts.
- Fabric Admin - users can perform fabric administrator actions such as update SM configuration, update global credentials, manage reports, managing unhealthy ports, and manage PKeys, etc.
- Fabric Operator - users can perform fabric operator actions such as device management actions (enable/disable port, add/remove devices to/from groups, reboot device, upgrade software, etc.)
- Monitoring Only - users can perform monitoring actions such as view the fabric configuration, open monitoring sessions, define monitoring templates, and export monitoring data to CSV files, etc.

To edit existing users accounts, right-click the account from the list of user accounts and perform the desired action (Change Password/Remove).



9.9.7 Email

SMTP configuration is required to set both the [Daily Reports Tab](#) and the Email-on-Events features.

1. In the SMTP Configuration dialogue window, enter the following information:

Settings

Events Policy Device Access Network Management Subnet Manager Non-Optimal Links User Management **Email** System Dump

SMTP Configurations

SMTP Server

SMTP Port

Sender Name

Sender Address

Use Authentication

Use SSL

Username

Password

Save

Attribute	Description
SMTP Server	The IP or host name of the SMTP server. Examples: <ul style="list-style-type: none"> If mail service is installed, localhost is a valid value for this field, but usually it cannot send mails outside the local domain. smtp.gmail.com
SMTP Port	Default value - 25
Sender Name	The name that will be displayed in the email header
Sender Address	A valid email address that will be displayed in the email header
Use Authentication	By default, this field is unchecked. If checked, you must supply a username and password in the respective fields
Use SSL	Default value is false - not using SSL
Username	SMTP account username
Password	SMTP account password

2. Click "Save." All configuration of the SMTP server will be saved in the UFM Database.

Click "Send Test Email" to test the configuration and the following model will appear:

Send Test Email
✕

Recipients

comma separated email addresses list

Subject

UFM Test Email

Message

Receiving this email means that your UFM SMTP configurations is correct.

Close
Send

Attribute	Description
Recipients	User can choose email from event policy and daily report recipients or enter any email
Subject	Email subject
Message	Email message

The System Health window enables running and viewing reports and logs for monitoring and analyzing UFM server and fabric health through the following tabs: UFM Health, UFM Logs, UFM Snapshot, Fabric Health, Daily Reports and Topology Compare.

9.9.8 Remote Location

Remote location tab is used to set a predefined remote location for the results of System Dump action on switches and hosts and for IBDiagnet executions.

Events Policy Device Access Network Management Subnet Manager Non-Optimal Links User Management Email Remote Location Data Stream

Remote Location

Protocol

Remote location is used to save result of System Dump and IBDiagnet.
By default this location will be used.
Path: N/A

Server

Path

Username

Password

Field	Description
Protocol	The protocol to use to move the dump file to the external storage (scp/sftp)
Server	Hostname or IP address of the server
Path	The path where dump files are saved
Username	Username for the server
Password	Respective password

After configuring these parameters, it would be possible for users to collect sysdumps for specific devices, groups, or links (through Network Map/Cables Window) by right-clicking the item and selecting System Dump.

9.9.9 Data Streaming

This section allows users to configure System Logs settings via web UI.

Data Streaming Configurations

System Logs

Status: Disabled Enabled

Mode: Local Remote


Destination: :

System logs level:

Streaming Data

UFM logs

Event logs (allows selecting which events to stream from [Events policy](#))

Field	Description
Status	Enable/disable exporting UFM logs to system logs
Mode	Export logs to local or remote system logs
Destination	Remote server IP/hostname and port
System Logs Level	Log level to export
Streaming Data	Logs to export to system logs. <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;">  Events logs are selected one by one from Events Policy settings when the system logs feature is enabled. </div>

9.9.10 Topology Compare

This tab controls the settings for the [Periodic Topology Comparison](#) feature.

Events Policy | Device Access | Network Management | Subnet Manager | Non-Optimal Links | User Management | Email | Remote Location | Data Streaming | **Topology Compare**

Topology Compare Settings

Comparison Interval (For comparing the current topology with master topology)

Days

Stable Topology Period (For offering user to update the master topology for comparison)


Hours

- Comparison Interval - determines how often the current topology is compared against the master topology

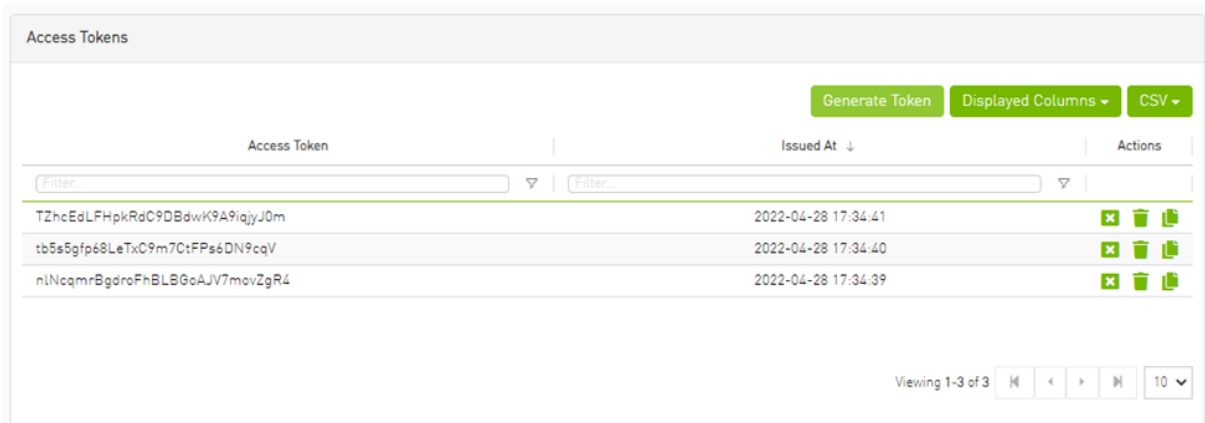
- Stable Topology Period - determines how long a topology must be stable before it is designated the new master topology




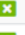





9.9.11 Token-based Authentication

Token-based authentication is a protocol which allows users to verify their identity, and in return receive a unique access token. During the life of the token, users then access the UFM APIs that the token has been issued for, rather than having to re-enter credentials each time they need to use any UFM API.




 Under the Settings section there is a tab titled called “Access Tokens”.


The functionality of the added tab is to give the user the ability to create new tokens & manage the existing ones (list, copy, revoke, delete):



Access Token	Issued At	Actions
TZhcEdLFHpkRdC9DBdwK9A9iqjyJ0m	2022-04-28 17:34:41	  
tb5s5gfp68LeTx09m70tFPs6DN9cqV	2022-04-28 17:34:40	  
nINcqmRBgdreFhBLBGoAJV7movZgR4	2022-04-28 17:34:39	  

Actions:

Name	Icon	Description
Revoke		Revoke a specific token.  The revoked token will no longer be valid.
Delete		Delete a specific token.
Copy		Copy specific token into the clipboard.

 Each user is able to list and manage only the tokens that have been created by themselves. Only the users with system_admin role will be able to create tokens.

9.9.12 Plugin Management


Plugin management allows users to manage UFM plugins without using CLI commands. Under "Settings", there is a tab titled "Plugin Management".

The functionality of the "Plugin Management" tab is to give the user the ability to add, remove, disable and enable plugins.


Name	Enabled	Tag	Port	Shared Volumes	Status
ahxmonitor	✓	latest	8910	/opt/ufm/files/log/log/opt/ufm/files/conf/opt/ufm/files/conf	stop
ndt	✗	NA	NA	NA	stop

Actions:

- Add - Used to add a selected plugin, opens a model to select the needed tag.




- Remove - Used to remove a selected plugin.



- Disable - Used to disable a selected plugin, so the plugin is disabled once the UFM is disabled.



- Enable - Used to enable a selected plugin, so the plugin is enabled once the UFM is enabled.

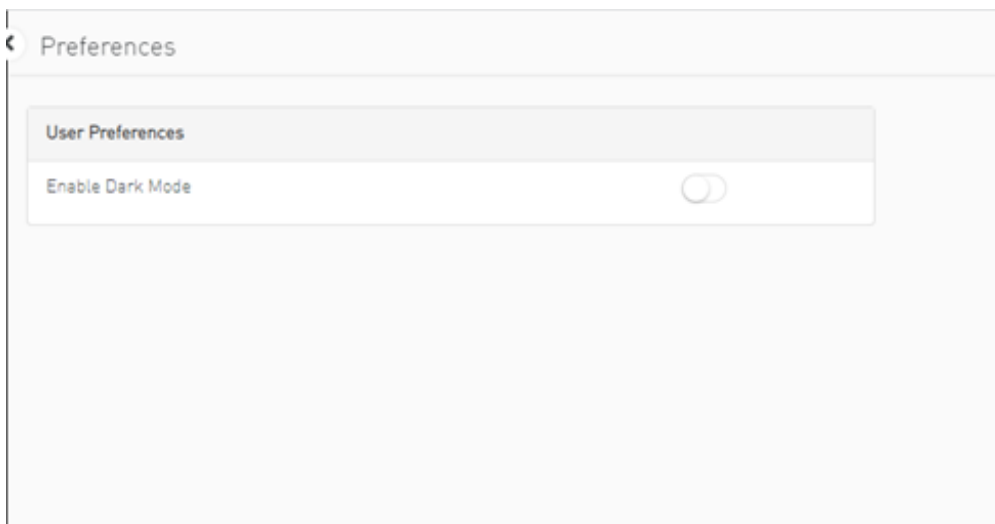
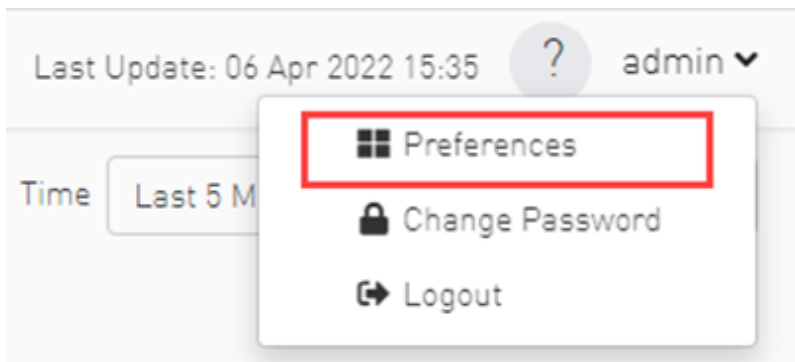


- Add ahxmonitor - Used to add a selected plugin; the action opens a modal to select the requested tag.




9.9.13 User Preferences

This page allows user to change UI preferences in general.



When user enables dark mode, the UFM is presented in dark theme.



10 UFM Plugins

- [rest-rdma Plugin](#)
- [NDT Plugin](#)
- [UFM Telemetry Fluent Streaming \(TFS\) Plugin](#)
- [UFM Events Fluent Streaming \(EFS\) Plugin](#)
- [GRPC-Streamer Plugin](#)

10.1 rest-rdma Plugin

rest-rdma is a utility to send REST requests over IB to the UFM server. rest-rdma is distributed as a docker container that could serve as server and as client.

10.1.1 Deployment Server

10.1.1.1 Deploy Plugin on UFM Appliance

1. Log into your UFM as admin.
2. Enter config mode. Run:

```
enable
config terminal
```



Make sure that UFM is running with `show ufm status`. If UFM is down then run with `ufm start`.

3. Ensure that rest-rdma plugin is disabled with `show ufm plugin command`
4. Pull the plugin container with `docker pull mellanox/ufm-plugin-rest-rdma:[version]`
5. Run `ufm plugin rest-rdma add tag [version]` to enable the plugin
6. Check that plugin is up and running with `docker pull mellanox/ufm-plugin-rest-rdma:[version]`

10.1.1.2 Deploy Plugin on Bare Metal Server

1. Verify that ufm is installed and running
2. Pull image from docker hub:
`docker pull mellanox/ufm-plugin-rest-rdma:[version]`
3. To load image run:
`/opt/ufm/scripts/manage_ufm_plugins.py add -p rest-rdma`

10.1.1.3 Deployment Client

To pull image from docker hub:

```
docker pull mellanox/ufm-plugin-rest-rdma:[version]
```

To start container as client (on any host in the same fabric as UFM server) run:

```
docker run -d --network=host --privileged --name=ufm-plugin-rest-rdma --rm -v /tmp/ibdiagnet:/tmp/ibdiagnet
mellanox/ufm-plugin-rest-rdma:[version] client
```

To check that plugin is up and running run:

```
docker ps
```

10.1.2 How to Run


10.1.2.1 Server

In server mode `ufm_rdma.py` will started automatically and will be restarted if exit. If `ufm_rdma.py` server is not running - enter to the docker and run the following commands to start the server:

```
cd /opt/ufm/src/ufm-plugin-ufm-rest
./ufm_rdma.py -r server
```


10.1.2.2 Client

There are three options to run client:


 From inside the docker, using custom script from the hosting server or using `docker exec` command from hosting server.

1. From inside the docker:
 - a. Enter to the docker using `docker exec -it ufm-plugin-rest-rdma bash`
 - b. then `cd /opt/ufm/src/ufm-plugin-rest-rdma`
 - c. Use `-h` help option to see available parameters
`./ufm_rdma.py -h`
2. From hosting server run script located at `/opt/ufm/ufm-plugin-ufm-rest/ufm-rest-rdma_client.sh` inside docker

 that could be copied using command
`cp <containerId>:/opt/ufm/ufm-plugin-ufm-rest/ufm-rest-rdma_client.sh /host/path/target`

 Example:
`./ufm-rest-rdma_client.sh -u admin -p password -t simple -a GET -w ufmRest/app/ufm_version`


- a. To see available options run:
`./ufm-rest-rdma_client.sh -h`
3. From hosting server using `docker exec` command.

 To run from inside docker, run:

```
docker exec ufm-plugin-rest-rdma prior to the command.
```

For example: `docker exec ufm-plugin-rest-rdma /opt/ufm/ufm-plugin-ufm-rest/src/ufm_rdma.py -r client -u admin -p password -t simple -a GET -w ufmRest/app/ufm_version`

10.1.3 Examples

 All the examples in this section are relevant for running the ufm-rest-rdma client from inside the docker.

If you must run ufm-rest-rdma using the client script, all quotation marks (") must be wrapped by a backslash (\).

For example, `"running_mode": "once"` must become `\\"running_mode\\":\\"once\\"`.

There are three types of user authentication flows supported by UFM and also by ufm-rest-rdma utility


10.1.3.1 Username/Password Authentication

```
to get UFM version
./ufm_rdma.py -r client -u admin -p password -t simple -a GET -w ufmRest/app/ufm_version

to get ibdiagnet run result
./ufm_rdma.py -r client -u admin -p password -t ibdiagnet -a POST -w ufmRest/reports/ibdiagnetPeriodic -l
'{"general": {"name": "IBDiagnet_CMD_1234567890_199_88", "location": "local", "running_mode": "once"},
"command_flags": {"-pc": ""}}'
```


10.1.3.2 Client Certificate Authentication

```
need to pass path to client certificate file and name of UFM server machine:
./ufm_rdma.py -r client -t simple -a GET -w ufmRest/resources/modules -d /path/to/certificate/file/ufm-client.pfx
-s ufm.azurehpc.core.azure-test.net
```

 Client certificate file should be located INSIDE docker container.

10.1.3.3 Token Authentication

```
need to pass it for authentication
./ufm_rdma.py -r client -k OGUy7TwLvTmFkXyTkcsEWD9KKNvq6f -t simple -a GET -w ufmRestV3/app/ufm_version
```

 Token could be generated using UFM UI.



If a token is used for client authentication, `ufmRestV3` must be used.

10.2 NDT Plugin

10.2.1 Overview

NDT plugin is a self-contained Docker container with REST API support managed by UFM. NDT plugin provides NDT topo diff capability. This feature allows the user to compare IB fabric managed by UFM and NDT files which are used by Microsoft for description of IB clusters network topology.

Main usage cases:

- Get confidence on the IB fabric connectivity during cluster bring-up.
- Get confidence on the specific parts of IB fabric after component replacements.
- Automatically detect any changes in topology.

10.2.2 Deployment

The following are the possible ways NDT plugin can be deployed:

1. On UFM Appliance
2. On UFM Software

Detailed instructions on how to deploy NDT plugin could be found on page [mellanox/ufm-plugin-ndt](https://mellanox.com/ufm-plugin-ndt).

10.2.3 Authentication

Following authentication types are supported:

- basic (/ufmRest)
- client (/ufmRestV2)
- token (/ufmRestV3)

10.2.4 REST API

The following REST APIs are supported:

- GET /help
- GET /version
- POST /upload_metadata
- GET /list
- POST /compare
- POST /cancel
- GET /reports
- GET /reports/<report_id>
- POST /delete

For detailed information on how to interact with NDT plugin, refer to the [NVIDIA UFM Enterprise > Rest API > NDT Plugin REST API](#).

10.2.5 NDT Format

NDT is a CSV file containing data relevant to the IB fabric connectivity.

NDT plugin extracts the IB connectivity data based on the following five fields:

1. Start device
2. Start port
3. End device
4. End port
5. Link type

10.2.5.1 Switch to Switch NDT

By default, IB links are filtered by:

- Link Type is Data
- Start Device and End Device end with IBn, where n is a numeric value.

For TOR switches, Start port/End port field should be in the format Port N, where N is a numeric value.

For Director switches, Start port/End port should be in the format Blade N_Port i/j, where N is a leaf number, i is an internal ASIC number and j is a port number.

Examples:

Start Device	Start Port	End Device	End Port	Link Type
DSM07-0101-0702-01IB0	Port 21	DSM07-0101-0702-01IB1	Blade 2_Port 1/1	Data
DSM07-0101-0702-01IB0	Port 22	DSM07-0101-0702-01IB1	Blade 2_Port 1/1	Data
DSM07-0101-0702-01IB0	Port 23	DSM07-0101-0702-02IB1	Blade 3_Port 1/1	Data
DSM09-0101-0617-001IB2	Port 33	DSM09-0101-0721-001IB4	Port 1	Data
DSM09-0101-0617-001IB2	Port 34	DSM09-0101-0721-001IB4	Port 2	Data
DSM09-0101-0617-001IB2	Port 35	DSM09-0101-0721-001IB4	Port 3	Data

10.2.5.2 Switch to Host NDT

NDT is a CSV file containing data not only relevant to the IB connectivity.

Extracting the IB connectivity data is based on the following five fields:

1. Start device
2. Start port
3. End device
4. End port

5. Link type

IB links should be filtered by the following:

- Link type is Data
- Start device or End device end with IBN, where N is a numeric value.
 - The other Port should be based on persistent naming convention: ibpXsYfZ, where X, Y and Z are numeric values.

For TOR switches, Start port/End port field will be in the format Port n, where n is a numeric value.

For Director switches, Start port/End port will be in the format Blade N_Port i/j, where N is a leaf number, i is an internal ASIC number and j is a port number.

Examples:

Start Device	Start Port	End Device	End Port	Link Type
DSM071081704019	DSM071081704019 ibp11s0f0	DSM07-0101-0514-01IB0	Port 1	Data
DSM071081704019	DSM071081704019 ibp21s0f0	DSM07-0101-0514-01IB0	Port 2	Data
DSM071081704019	DSM071081704019 ibp75s0f0	DSM07-0101-0514-01IB0	Port 3	Data

10.2.6 Other

Comparison results are forwarded to syslog as events. Example of `/var/log/messages` content:

1. Dec 9 12:32:31 <server_ip> ad158f423225[4585]: NDT: missing in UFM "SAT111090310019/SAT111090310019 ibp203s0f0 - SAT11-0101-0903-19IB0/15"
2. Dec 9 12:32:31 <server_ip> ad158f423225[4585]: NDT: missing in UFM "SAT11-0101-0903-09IB0/27 - SAT11-0101-0905-01IB1-A/Blade 12_Port 1/9"
3. Dec 9 12:32:31 <server_ip> ad158f423225[4585]: NDT: missing in UFM "SAT11-0101-0901-13IB0/23 - SAT11-0101-0903-01IB1-A/Blade 08_Port 2/13"

For detailed information about how to check syslog, please refer to the [NVIDIA UFM-SDN Appliance Command Reference Guide](#) > UFM Commands > UFM Logs.

Minimal interval value for periodic comparison in five minutes.

In case of an error the clarification will be provided.

For example, the request “`POST /compare`” without NDTs uploaded will return the following:

- URL: https://<server_ip>/ufmRest/plugin/ndt/compare
- response code: 400
- Response:

```
{
  "error": [
    "No NDTs were uploaded for comparison"
  ]
}
```


Configurations could be found in “ `ufm/conf/ndt.conf` ”

- Log level (default: INFO)
- Log size (default: 10240000)
- Log file backup count (default: 5)
- Reports number to save (default: 10)
- NDT format check (default: enabled)
- Switch to switch and host to switch patterns (default: see NDT format section)

For detailed information on how to export or import the configuration, refer to the [NVIDIA UFM-SDN Appliance Command Reference Guide](#) > UFM Commands > UFM Configuration Management.

Logs could be found in “ `ufm/logs/ndt.log` ”.

For detailed information on how to generate a debug dump, refer to the [NVIDIA UFM-SDN Appliance Command Reference Guide](#) > System Management > Configuration Management > File System.

10.3 UFM Telemetry Fluent Streaming (TFS) Plugin

10.3.1 Overview

TFS plugin is a self-contained Docker container with REST API support managed by UFM. TFS plugin provides Telemetry counters streaming to FluentD capability. As a fabric manager, the UFM Telemetry holds a real-time network telemetry information of the network topology. This information changes over time and is reflected to the telemetry console. In order to do so, we present a stream of the UFM Telemetry data to the FluentD plugin.

10.3.2 Deployment

The following are the possible ways TFS plugin can be deployed:

1. On UFM Appliance
2. On UFM Software

For complete instructions on how to deploy the TFS plugin, refer to [UFM Telemetry endpoint stream To Fluentd endpoint \(TFS\)](#).

10.3.3 Authentication

The following authentication types are supported:

- basic (/ufmRest)
- client (/ufmRestV2)
- token (/ufmRestV3)

10.3.4 Rest API

The following REST APIs are supported:

- POST /plugin/tfs/conf
- GET /plugin/tfs/conf

For detailed information on how to interact with NDT plugin, refer to the [NVIDIA UFM Enterprise > Rest API > TFS Plugin REST API](#).

10.4 UFM Events Fluent Streaming (EFS) Plugin

10.4.1 Overview

EFS plugin is a self-contained Docker container with REST API support managed by UFM. EFS plugin extracts the UFM events from UFM Syslog and streams them to a remote FluentD destination. It also has the option to duplicate current UFM Syslog messages and forward them to a remote Syslog destination. As a fabric manager, it will be useful to collect the UFM Enterprise events/logs, stream them to the destination endpoint and monitor them.

10.4.2 Deployment

The following are the ways EFS plugin can be deployed:

1. On UFM Appliance
2. On UFM Software

For detailed instructions on how to deploy EFS plugin, refer to [UFM Event Stream to FluentBit endpoint \(EFS\)](#).

10.4.3 Authentication

The following authentication types are supported:

- basic (/ufmRest)
- client (/ufmRestV2)
- token (/ufmRestV3)

10.4.4 Rest API

The following REST APIs are supported:

- PUT /plugin/efs/conf
- GET /plugin/efs/conf

For detailed information on how to interact with EFS plugin, refer to the [NVIDIA UFM Enterprise > Rest API > EFS Plugin REST API](#).

10.5 GRPC-Streamer Plugin

10.5.1 Authentication

The following authentication types are supported:

- Basic (/ufmRest)
- Token (/ufmRestV3)

10.5.2 Create a Session to UFM from GRPC

Description: Creates a session to receive REST API results from the UFM's GRPC server. After a stream or one call, the session is deleted so the server would not save the authorizations.

- Call: CreateSession in the grpc
- Request Content Type - message SessionAuth
- Request Data:

```
message SessionAuth{
  string job_id=1;
  string username = 2;
  string password = 3;
  optional string token = 4;
}
```

- Job_id - The unique identifier for the client you want to have
- Username - The authentication username
- Password - The authentication password
- Token - The authentication token
- Response:

```
message SessionRespond{
  string respond=1;
}
```

- Respond types:
 - Success - Ok.
 - ConnectionError - UFM connection error (bad parameters or UFM is down).
 - Other exceptions - details sent in the respond.
- Console command:

```
client session --server_ip=server_ip --id=client_id --auth=username,password --token=token
```

10.5.3 Create New Subscription

- Description: Only after the server has established a session for this grpc client, add all the requested REST APIs with intervals and delta requests.
- Call: AddSubscriber
- Request Content Type - Message SubscriberParams
- Request Data:

```
message SubscriberParams{
  message APIParams {
    string ufm_api_name = 1;
    int32 interval = 2;
    optional bool only_delta = 3;
  }
  string job_id = 1;
  repeated APIParams apiParams = 2;
}
```

- Job_id - A unique subscriber identifier
- apiParams - The list of apiParams from the above message above:
- ufm_api_name - The name from the known to server request api list - TBD

- interval - The interval between messages conducted in a stream run. Presented in seconds.
- only_delta - Receives the difference between the previous messages in a stream run.
- Response content type:

```
message SessionRespond{
  string respond=1;
}
```

- Respond Types:
 - Created a user with session and added new IP- Ok.
 - Cannot add subscriber that do no have an established session - need to create a session before creating subscriber.
 - The server already have the ID - need to create new session and new subscriber with a new unique ID.
- Console command:

```
client create --server_ip=localhost --id=client_id --apis=events;40;True,links,alarms;10
```

The API's list is separated by commas, and each modifier for the REST API is separated by a semi comma.

If the server is not given a modifier, default ones are used (where only_delta is False and interval is based on the API).

10.5.4 Edit Known Subscription

- Description: Changes a known IP. Whether the server has the IP or not.
- Call: AddSubscriber
- Request Content Type - Message SubscriberParams
- Request Data:

```
message SubscriberParams{
  message APIParams {
    string ufm_api_name = 1;
    int32 interval = 2;
    optional bool only_delta = 3;
  }
  string job_id = 1; //unique identifier for this job
  repeated APIParams apiParams = 2;
}
```

- Job_id - The subscriber unique identifier
- apiParams - A list of apiParams from the above message.
- TBD ufm_api_name - name from the known to server request api list
- interval - The interval between messages conducted in a stream run. Presented in seconds.
- only_delta - Receives the difference between the previous messages in a stream run.
- Response content type:

```
message SessionRespond{
  string respond=1;
}
```

- Respond Types:
 - Created user with new IP- Ok.
 - Cannot add subscriber without an established session - need to create a session before creating subscriber.

- Cannot add subscriber illegal apis - cannot create subscriber with empty API list, call again with correct API list.

10.5.5 Get List of Known Subscribers

- Description: Gets the list of subscribers, including the requested list of APIs.
- Call: ListSubscribers
- Request Content Type: google.protobuf.Empty
- Response:

```
message ListSubscriberParams{
  repeated SubscriberParams subscribers = 1;
}
```

- Console command: server subscribes --server_ip=server_ip

10.5.6 Delete a Known Subscriber

- Description: Deletes an existing subscriber and removes the session.
- Call: DeleteSubscriber
- Request Content Type: Message gRPCStreamerID
- Request Data:

```
message gRPCStreamerID{
  string job_id = 1;
}
```

- Response:protobuf.Empty

10.5.7 Run a Known Subscriber Once

- Description: Runs the Rest API list for a known subscriber once and returns the result in message runOnceRespond, and then delete the subscriber's session.
- Call: RunOnceJob
- Request Content Type: Message gRPCStreamerID
- Request Data:

```
message gRPCStreamerID{
  string job_id = 1;
}
```

- Response content type:

```
message runOnceRespond{
  string job_id=1;
  repeated gRPCStreamerParams results = 2;
}
```

- Job_id- The first message unique identifier.
- Results - list of gRPCStreamerParams contains results from each REST API
- Responses:

- Job id - Cannot run a client without an established session. Empty results - an existing session for this client is not found, and the client is not known to the server.
- Job id - Cannot run the client without creating a subscriber. Empty results - a session was created for the client but the subscription is not created.
- Job_id - Cannot connect to the UFM. empty result - the GRPC server cannot connect to the UFM machine and receive empty results, because it cannot create a subscriber with an empty API list. This means that the UFM machine is experiencing a problem.
- Job_id - The first unique message identifier of the messages. Not empty results - Ok
- Console command:

```
client once_id --server_ip=server_ip --id=client_id
```

10.5.8 Run Streamed Data of a Known Subscriber

- Description: Run a stream of results from the Rest API list for a known Subscriber and return the result as iterator, where each item is message gRPCStreamerParams. at the end, delete the session.
- Call: RunStreamJob
- Request Content Type: Message gRPCStreamerID
- Request Data:

```
message gRPCStreamerID{
  string job_id = 1;
}
```

- Response content type: iterator of messages gRPCStreamerParams

```
message gRPCStreamerParams{
  string message_id = 1; // unique identifier for messages
  string ufm_api_name = 2; // what rest api receive the data from
  google.protobuf.Timestamp timestamp = 3; //what time we created the message, can be converted to Datetime
  string data = 4; // data of rest api call
}
```

- Response:
 - One message only containing "Cannot run a client without a session" - A session has not been established
 - No message - A session and/or a subscriber with this ID does not exist.
 - Messages with interval between with the modifiers - Ok
- Console command:

```
client stream_id --server_ip=server_ip --id=client_id
```

10.5.9 Run a New Subscriber Once

- Description: After ensuring that a session for this specific job ID is established, the server runs the whole REST API list for the new subscriber once and returns the following result in message `runOnceRespond`. This action does not save the subscribe ID or the established session in the server.

- Call: RunOnce
- Request Content Type: Message SubscriberParams
- Request Data:

```
message SubscriberParams{
  message APIParams {
    string ufm_api_name = 1;
    int32 interval = 2;
    optional bool only_delta = 3;
  }
  string job_id = 1; //unique identifier for this job
  repeated APIParams apiParams = 2;
}
```

- Response content type:

```
message runOnceRespond{
  string job_id=1;
  repeated gRPCStreamerParams results = 2;
}
```

- Responses:
 - Job id = Cannot run a client without an established session. Empty results - no session for this client.
 - Job_id = 0 - The GRPC server cannot connect to the UFM machine and receive empty results, or it cannot create a subscriber with an empty API list.
 - Job_id = The messages' first unique identifier, and not an empty result - Ok.
- Console command:

```
client once --server_ip=server_ip --id=client_id --auth=username,password --token=token --apis=events;40;True,links;20;False,alarms;10
```

- The console command creates a session for this specific client.
- A token or the basic authorization is needed, not both.

10.5.10 Run New Subscriber Streamed Data

- Description: After the server checks it has a session for this job ID, Run a stream of results from the Rest API list for a new Subscriber and return the result as iterator, where each item is message gRPCStreamerParams. at the end, delete the session.
- Call: RunPeriodically
- Request Content Type: Message SubscriberParams
- Request Data:

```
message SubscriberParams{
  message APIParams {
    string ufm_api_name = 1;
    int32 interval = 2;
    optional bool only_delta = 3;
  }
  string job_id = 1; //unique identifier for this job
  repeated APIParams apiParams = 2;
}
```

- Response content type: iterator of messages gRPCStreamerParams
- Response:
 - Only one message with data equals to Cant run client without session - no session
 - Messages with intervals between with the modifiers - Ok
- Console command:

```
client stream --server_ip=server_ip --id=client_id --auth=username,password --token=token --apis=events;40;True,links;20;False,alarms;10
```

- console command also create session for that client.
- no need for both token and basic authorization, just one of them.

10.5.11 Run A Serialization on All the Running Streams

- Description: Run a serialization for each running stream. The serialization will return to each of the machines the results from the rest api list.
- Call: Serialization
- Request Content Type: google.protobuf.Empty
- Response: google.protobuf.Empty

10.5.12 Stop a Running Stream

- Description: Cancels running stream using the client id of the stream and stop it from outside, If found stop the stream.
- Call: StopStream
- Request Content Type: Message gRPCStreamerID
- Request Data:

```
message gRPCStreamerID{
  string job_id = 1;
}
```

- Response: google.protobuf.Empty

10.5.13 Run a subscribe stream

- Description: Create a subscription to a client identifier, all new messages that go to that client, will be copied and also sent to this stream.
- Call: Serialization
- Request Content Type: message gRPCStreamerID
- Response: iterator of messages gRPCStreamerParams

```
message gRPCStreamerParams{
  string message_id = 1; // unique identifier for messages
  string ufm_api_name = 2; // what rest api receive the data from
  google.protobuf.Timestamp timestamp = 3; //what time we created the message, can be converted to Datetime
  string data = 4; // data of rest api call
}
```

- the identifier may or may not be in the grpc server.
- Cannot be stop streamed using StopStream.
- Console command:

```
client subscribe --server_ip=server_ip --id=client_id
```


10.5.14 Get the variables from a known subscriber

- Description: Get the variables of known subscriber if found, else return empty variables.
- Call: GetJobParams
- Request Content Type: message gRPCStreamerID
- Response:

```
message SubscriberParams{
  message APIParams {
    string ufm_api_name = 1; //currently the list of api from ufm that are supported are [Jobs, Events,
Links, Alarms]
    int32 interval = 2;
    optional bool only_delta = 3;
  }
  string job_id = 1; //unique identifier for this job
  repeated APIParams apiParams = 2;
}
```

10.5.14.1 Get Help / Version

- Description: Get help and the version of the plugin, how to interact with the server. What stages need to be done to extract the rest apis (Session>run once/stream or Session>AddSubscriber>once_id/stream_id)
- Call: Help or Version
- Request Content Type: google.protobuf.Empty
- Response:

```
message SessionRespond{
  string respond=1;
}
```

11 Troubleshooting

11.1 Split-Brain Recovery in HA Installation

The split-brain problem is a DRBD synchronization issue (HA status shows `DUnknown` in the DRBD disk state), which occurs when both HA nodes are rebooted. For example, in cases of electricity shut-down. To recover, please follow the below steps:

- **Step 1: Manually choose a node where data modifications will be discarded.**
It is called the split-brain victim. Choose wisely; all modifications will be lost! When in doubt, run a backup of the victim's data before you continue.
When running a Pacemaker cluster, you can enable maintenance mode. If the split-brain victim is in the Primary role, bring down all applications using this resource. Now switch the victim to the Secondary role:

```
victim# drbdadm secondary ha_data
```

- **Step 2: Disconnect the resource if it's in connection state `WConnection`:**

```
victim# drbdadm disconnect ha_data
```

- **Step 3: Force discard of all modifications on the split-brain victim:**

```
victim# drbdadm -- --discard-my-data connect resource
```

For DRBD 8.4.x:

```
victim# drbdadm connect --discard-my-data resource
```

- **Step 4: Resync starts automatically if the survivor is in a `WConnection` network state. If the split-brain survivor is still in a `Standalone` connection state, reconnect it:**

```
survivor# drbdadm connect resource
```

Now the resynchronization from the survivor (`SyncSource`) to the victim (`SyncTarget`) starts immediately. There is no full sync initiated, but all modifications on the victim will be overwritten by the survivor's data, and modifications on the survivor will be applied to the victim.

12 Appendixes

- [Appendix - Diagnostic Utilities](#)
- [Appendix - Supported Port Counters and Events](#)
- [Appendix - Used Ports](#)
- [Appendix - Configuration Files Auditing](#)
- [Appendix - IB Router](#)
- [Appendix - NVIDIA SHARP Integration](#)
- [Appendix - AHX Monitoring](#)
- [Appendix - UFM SLURM Integration](#)
- [Appendix - Device Management Feature Support](#)
- [Appendix - UFM Event Forwarder](#)
- [Appendix - UFM Multisite Portal Integration](#)

12.1 Appendix - Diagnostic Utilities



For UFM-SDN Appliance, all the below diagnostics commands have ib prefix.

For example, for UFM-SDN Appliance, the command `ibstat` is `ib ibstat`.

12.1.1 InfiniBand Diagnostics Commands

Command	Description
ibstat	Shows the host adapters status.
ibstatus	Similar to ibstat but implemented as a script.
ibnetdiscover	Scans the topology.
ibaddr	Shows the LID range and default GID of the target (default is the local port).
ibroute	Displays unicast and multicast forwarding tables of the switches.
ibtracert	Displays unicast or multicast route from source to destination.
ibping	Uses vendor MADs to validate connectivity between InfiniBand nodes. On exit, (IP) ping-like output is shown.
ibsysstat	Obtains basic information for the specific node which may be remote. This information includes: hostname, CPUs, memory utilization.
sminfo	Queries the SMIInfo attribute on a node.
smpdump	A general purpose SMP utility which gets SM attributes from a specified SMA. The result is dumped in hex by default.
smpquery	Enables a basic subset of standard SMP queries including the following: node info, node description, switch info, port info. Fields are displayed in human readable format.
perfquery	Dumps (and optionally clears) the performance counters of the destination port (including error counters).
ibswitches	Scans the net or uses existing net topology file and lists all switches.
ibhosts	Scans the net or uses existing net topology file and lists all hosts.

Command	Description
ibnodes	Scans the net or uses existing net topology file and lists all nodes.
ibportstate	Gets the logical and physical port states of an InfiniBand port or disables or enables the port (only on a switch). Note: This tool can change port settings. Should be used with caution.
saquery	Issues SA queries.
ibdiagnet	ibdiagnet scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices.
ibnetsplit	Automatically groups hosts and creates scripts that can be run to split the network into sub-networks each containing one group of hosts.
lbqueryerrors	Queries IB spec-defined errors from all fabric ports. Note: This tool can change reset port counters Should be used with caution.
smparquery	Queries adaptive-routing related settings from a particular switch. Note: This tool can change reset port counters Should be used with caution.

12.1.2 Diagnostic Tools

Model of operation: All utilities use direct MAD access to operate. Operations that require QP 0 mads only, may use direct routed mads, and therefore may work even in subnets that are not configured. Almost all utilities can operate without accessing the SM, unless GUID to lid translation is required.

12.1.2.1 Dependencies

Multiple port/Multiple CA support:

When no InfiniBand device or port is specified (as shown in the following example for "Local umad parameters"), the tools select the interface port to use by the following criteria:

1. The first InfiniBand ACTIVE port.
2. If not found, the first InfiniBand port that is UP (physical link up).

If a port and/or CA name is specified, the tool attempts to fulfill the user's request and will fail if it is not possible.

For example:

```
ibaddr      # use the 'best port'
ibaddr -C mthca1 # pick the best port from mthca1 only.
ibaddr -P 2 # use the second (active/up) port from the first available IB device.
ibaddr -C mthca0 -P 2 # use the specified port only.
```

Common Options & Flags

Most diagnostics take the following flags. The exact list of supported flags per utility can be found in the usage message and can be shown using `util_name -h` syntax.

```
# Debugging flags
-d raise the IB debugging level. May be used several times (-ddd or -d -d -d).
-e show umad send receive errors (timeouts and others)
```

```
-h show the usage message
-v increase the application verbosity level.
  May be used several times (-vv or -v -v -v)
-V show the internal version info.
```

```
# Addressing flags
-D use directed path address arguments.
  The path is a comma separated list of out ports.
  Examples:
  "0" # self port
  "0,1,2,1,4" # out via port 1, then 2, ...
-G use GUID address arguments.
  In most cases, it is the Port GUID.
  Examples:
  "0x08f1040023"
-t <smlid> use 'smlid' as the target lid for SA queries.
```

```
# Local umad parameters:
-C <ca_name> use the specified ca_name.
-P <ca_port> use the specified ca_port.
-t <timeout_ms> override the default timeout for the
  solicited mads.
```

CLI notation: all utilities use the POSIX style notation, meaning that all options (flags) must precede all arguments (parameters).

12.1.3 Utilities Descriptions

ibstatus

A script that displays basic information obtained from the local InfiniBand driver. Output includes LID, SMLID, port state, link width active, and port physical state.

Syntax

```
ibstatus [-h] [devname[:port]]
```

Examples:

```
ibstatus # display status of all IB ports
ibstatus mthca1 # status of mthca1 ports
ibstatus mthca1:1 mthca0:2 # show status of specified ports
```

See also: `ibstat`

ibstat

Similar to the `ibstatus` utility but implemented as a binary and not as a script. Includes options to list CAs and/or ports.

Syntax

```
ibstat [-d(ebug) -l(ist_of_cas) -p(ort_list) -s(hort)] <ca_name> [portnum]
```

Examples:

```
ibstat # display status of all IB ports
ibstat mthca1 # status of mthca1 ports
ibstat mthca1 2 # show status of specified ports
ibstat -p mthca0 # list the port guids of mthca0
ibstat -l # list all CA names
```

See also: `ibstatus`

ibroute

Uses SMPs to display the forwarding tables (unicast (LinearForwardingTable or LFT) or multicast (MulticastForwardingTable or MFT)) for the specified switch LID and the optional lid (mlid) range. The default range is all valid entries in the range 1...FDBTop.

Syntax

```
ibroute [options] <switch_addr> [<startlid> [<endlid>]]
```

Nonstandard flags:

```
-a          show all lids in range, even invalid entries.
-n          do not try to resolve destinations.
-M          show multicast forwarding tables. In this case the range
           parameters are specifying mlid range.
node-name-map  node name map file
```

Examples:

```
ibroute 2          # dump all valid entries of switch lid 2
ibroute 2 15       # dump entries in the range 15...FDBTop.
ibroute -a 2 10 20 # dump all entries in the range 10..20
ibroute -n 2       # simple format
ibroute -M 2       # show multicast tables
```

See also: ibtracert

ibtracert

Uses SMPs to trace the path from a source GID/LID to a destination GID/LID. Each hop along the path is displayed until the destination is reached or a hop does not respond. By using the -m option, multicast path tracing can be performed between source and destination nodes.

Syntax

```
ibtracert [options] <src-addr> <dest-addr>
```

Nonstandard flags:

```
-n          simple format; don't show additional information.
-m <mlid>  show the multicast trace of the specified mlid.
-f <force> force
node-name-map  node name map file
```

Examples:

```
ibtracert 2 23          # show trace between lid 2 and 23
ibtracert -m 0xc000 3 5 # show multicast trace between lid 3
and 5 for mcast lid 0xc000.
```

smpquery

Enables a basic subset of standard SMP queries including the following node info, node description, switch info, port info. Fields are displayed in human readable format.

Syntax

```
smpquery [options] <op> <dest_addr> [op_params]
```

Currently supported operations and their parameters:

```
nodeinfo <addr>
nodedesc <addr>
portinfo <addr> [<portnum>] # default port is zero
switchinfo <addr>
pkeys <addr> [<portnum>]
sl2vl <addr> [<portnum>]
vlarb <addr> [<portnum>]
GUIDInfo (GI) <addr>
MlnxExtPortInfo (MEPI) <addr> [<portnum>]
Combined (-c) : use Combined route address argument
node-name-map : node name map file
extended (-x) : use extended speeds
```

Examples:

```
smpquery nodeinfo 2 # show nodeinfo for lid 2
smpquery portinfo 2 5 # show portinfo for lid 2 port 5
```

smpdump

A general purpose SMP utility that gets SM attributes from a specified SMA. The result is dumped in hex by default.

Syntax

```
smpdump [options] <dest_addr> <attr> [mod]
```

Nonstandard flags:

```
-s show output as string
```

Examples:

```
smpdump -D 0,1,2 0x15 2 # port info, port 2
smpdump 3 0x15 2 # port info, lid 3 port 2
```

ibaddr

Can be used to show the LID and GUID addresses of the specified port or the local port by default. This utility can be used as simple address resolver.

Syntax

```
ibaddr [options] [<dest_addr>]
```

Nonstandard flags:

```
gid_show (-g) : show gid address only
lid_show (-l) : show lid range only
Lid_show (-L) : show lid range (in decimal) only
```

Examples:

```
ibaddr # show local address
ibaddr 2 # show address of the specified port lid
ibaddr -G 0x8f1040023 # show address of the specified port guid
```

sminfo

Issues and dumps the output of an sminfo query in human readable format. The target SM is the one listed in the local port info or the SM specified by the optional SM LID or by the SM direct routed path.



CAUTION: Using sminfo for any purpose other than a simple query might result in a malfunction of the target SM.

Syntax

```
sminfo [options] <sm_lid|sm_dr_path> [sminfo_modifier]
```

Nonstandard flags:

```
-s <state>          # use the specified state in sminfo mad
-p <priority>       # use the specified priority in sminfo mad
-a <activity>       # use the specified activity in sminfo mad
```

Examples:

```
sminfo          # show sminfo of SM listed in local portinfo
sminfo 2        # query SM on port lid 2
```

perfquery

Uses PerfMgt GMPs to obtain the PortCounters (basic performance and error counters) from the Performance Management Agent (PMA) at the node specified. Optionally show aggregated counters for all ports of node. Also, optionally, reset after read, or only reset counters.

```
perfquery [options] [<lid|guid> [[port] [reset_mask]]]
```

Nonstandard flags:

```
-a                Shows aggregated counters for all ports of the destination lid.
-r                Resets counters after read.
-R                Resets only counters.
Extended (-x)    Shows extended port counters
Xmtsl (-X)       Shows Xmt SL port counters
Rcvsl ,(-S)      Shows Rcv SL port counters
Xmtdisc (-D)     Shows Xmt Discard Details
rcvrr, (-E)      Shows Rcv Error Details
extended_speeds (-T) Shows port extended speeds counters
oprcvcounters    Shows Rcv Counters per Op code
flowctlcounters Shows flow control counters
vloppackets      Shows packets received per Op code per VL
vlopdata         Shows data received per Op code per VL
vlxmitflowctlerrors Shows flow control update errors per VL
vlxmitcounters   Shows ticks waiting to transmit counters per VL
swportvlcong     Shows sw port VL congestion
rcvcc            Shows Rcv congestion control counters
slrcvfeen        Shows SL Rcv FECN counters
slrcvbeecn       Shows SL Rcv BECN counters
xmitcc           Shows Xmit congestion control counters
vlxmittlecc      Shows VL Xmit Time congestion control counters
smplctl (-c)     Shows samples control
loop_ports (-l)  Iterates through each port
```

Examples:

```
perfquery          # read local port's performance counters
perfquery 32 1     # read performance counters from lid 32, port 1
perfquery -a 32    # read from lid 32 aggregated performance counters
```



```
perfquery -r 32 1 # read performance counters from lid 32 port 1 and reset
perfquery -R 32 1 # reset performance counters of lid 32 port 1 only
perfquery -R -a 32 # reset performance counters of all lid 32 ports
perfquery -R 32 2 0xf000 # reset only non-error counters of lid 32 port 2
```

ibping

Uses vendor mads to validate connectivity between InfiniBand nodes. On exit, (IP) ping like output is show. `ibping` is run as client/server. The default is to run as client. Note also that a default ping server is implemented within the kernel.

Syntax

```
ibping [options] <dest lid|guid>
```

Nonstandard flags:

```
-c <count> stop after count packets
-f flood destination: send packets back to back w/o delay
-o <oui> use specified OUI number to multiplex vendor MADs
-S start in server mode (do not return)
```

ibnetdiscover

Performs InfiniBand subnet discovery and outputs a human readable topology file. GUIDs, node types, and port numbers are displayed as well as port LIDs and node descriptions. All nodes (and links) are displayed (full topology). This utility can also be used to list the current connected nodes. The output is printed to the standard output unless a topology file is specified.

Syntax

```
ibnetdiscover [options] [<topology-filename>]
```

Nonstandard flags:

```
l Lists connected nodes
H Lists connected HCAs
S Lists connected switches
g Groups
full (-f) Shows full information (ports' speed and width, vlcaps)
show (-s) Shows more information
Router_list (-R) Lists connected routers
node-name-map Nodes name map file
cache filename to cache ibnetdiscover data to
load-cache filename of ibnetdiscover cache to load
diff filename of ibnetdiscover cache to diff
diffcheck Specifies checks to execute for --diff
ports : (-p) Obtains a ports report
max_hops (-m) Reports max hops discovered by the library
outstanding_smpps (-o) Specifies the number of outstanding SMP's which should be issued during the scan
```

ibhosts

Traces the InfiniBand subnet topology or uses an already saved topology file to extract the CA nodes.

Syntax

```
ibhosts [-h] [<topology-file>]
```

Dependencies: `ibnetdiscover`, `ibnetdiscover format`

ibswitches

Traces the InfiniBand subnet topology or uses an already saved topology file to extract the InfiniBand switches.

Syntax

```
ibswitches [-h] [<topology-file>]
```

Dependencies: ibnetdiscover, ibnetdiscover format

ibportstate

Enables the port state and port physical state of an InfiniBand port to be queried or a switch port to be disabled or enabled.

Syntax

```
ibportstate [-d( debug) -e( rr_show) -v( erbose) -D( irect) -G( uid) -s smlid -V( ersion) -C ca_name -P ca_port -t  
timeout_ms] <dest dr_path|lid|guid> <portnum> [<op>]
```

Supported ops: enable, disable, query, on, off, reset, speed, espeed, fdr10, width, down, arm, active, vls, mtu, lid, smlid, lmc, mkey, mkeylease, mkeyprot

Examples:

```
ibportstate 3 1 disable # by lid  
ibportstate -G 0x2C9000100D051 1 enable # by guid  
ibportstate -D 0 1 # by direct route
```

ibnodes

Uses the current InfiniBand subnet topology or an already saved topology file and extracts the InfiniBand nodes (CAs and switches).

Syntax

```
ibnodes [<topology-file>]
```

Dependencies: ibnetdiscover, ibnetdiscover format

ibqueryerrors

Queries or clears the PMA error counters in PortCounters by walking the InfiniBand subnet topology.

```
ibqueryerrors [options]
```

Syntax

```
Options:  
--suppress, -s <err1,err2,...> suppress errors listed  
--suppress-common, -c suppress some of the common counters  
--node-name-map <file> node name map file  
--port-guid, -G <port_guid> report the node containing the port  
specified by <port_guid>  
--, -S <port_guid> Same as "-G" for backward compatibility  
--Direct, -D <dr_path> report the node containing the port specified  
by <dr_path>  
--skip-sl don't obtain SL to all destinations  
--report-port, -r report port link information
```

```

--threshold-file <val> specify an alternate threshold file, default: /etc/infiniband-diags/error_thresholds
--GNDN, -R           (This option is obsolete and does nothing)
--data              include data counters for ports with errors
--switch            print data for switches only
--ca                print data for CA's only
--router            print data for routers only
--details           include transmit discard details
--counters          print data counters only
--clear-errors, -k  Clear error counters after read
--clear-counts, -K Clear data counters after read
--load-cache <file> filename of ibnetdiscover cache to load
--outstanding_smps, -o <val> specify the number of outstanding SMP's
                        which should be issued during the scan
--config, -z <config> use config file, default: /etc/infiniband-diags/ibdiag.conf
--Ca, -C <ca>       Ca name to use
--Port, -P <port>   Ca port number to use
--timeout, -t <ms>  timeout in ms
--m_key, -y <key>   M_Key to use in request
--errors, -e        show send and receive errors
--verbose, -v       increase verbosity level
--debug, -d         raise debug level
--help, -h          help message
--version, -V       show version

```

smparquery

Issues Adaptive routing-related queries to the fabric switch.

Syntax

```

Supported ops (and aliases, case insensitive):
  ARInfo (ARI) <addr>
  ARGroupTable (ARGT) <addr> [<plft>] [<group_table>] [<blocknum>]
  ARLPTTable (ARLT) <addr> [<plft>] [<blocknum>]
  PLFTInfo (PLFTI) <addr>
  PLFTDef (PLFTD) <addr> [<blocknum>]
  PLFTMap (PLFTM) <addr> [<plft>] [<control_map>]
  PortSLToPLFTMap (PLFTP) <addr> [<blocknum>]
  RNSubGroupDirectionTable (DIRT) <addr> [<blocknum>]
  RNgGenStringTable (GSTR) <addr> [<plft>] [<blocknum>]
  RNgGenBySubGroupPriority (GSGP) <addr>
  RNRCvString (RSTR) <addr> [<blocknum>]
  RNXmitPortMask (RNXM) <addr> [<blocknum>]
  PortRNCounters (RNPC) <addr>

```

Options:

```

Main
-C|--Ca <ca>           : Ca name to use
-P|--Port <port>       : Ca port number to use
-D|--Direct            : use Direct address argument
-L|--Lid               : use LID address argument
-h|--help              : help message
-V|--version           : show version
-d|--debug             : Print debug logs

```

saquery

Issues SA queries.

Syntax

```
saquery [-h -d -P -N -L -G -s -g] [<name>]
```

Queries node records by default.

d	Enables debugging
P	Gets PathRecord info
N	Gets NodeRecord info
L (-L)	Returns just the Lid of the name specified
G (-G)	Returns just the Guid of the name specified
S (-S)	Returns the PortInfoRecords with isSM capability mask bit on
G (-g)	Gets multicast group info
L (-l)	Returns the unique Lid of the name specified
O (-O)	Returns name for the Lid specified
m(-m)	Gets multicast member info (if multicast group specified, list
x (-x)	member GIDs only for group specified for example 'saquery -m
c (-c)	0xC000')
S (-S)	Gets LinkRecord info"
I (-I)	Gets the SA's class port info
list (-D)	Gets ServiceRecord info
src-to-dst (<src:dst>)	Gets InformInfoRecord (subscription) info
sgid-to-dgid (<sgid-dgid>)	the node desc of the CA's
node-name-map	Gets a PathRecord for <src:dst> where src and dst are either
smkey <val>	node names or LIDs
slid <lid>	Gets a PathRecord for <sgid-dgid> where sgid and dgid are
dlid <lid>	addresses in IPv6 format
mild <lid>	Specifies a node name map file
sgid <gid>	SA SM_Key value for the query. If non-numeric value (like 'x')
dgid <gid>	is specified then saquery will prompt for a value. Default
gid <gid>	(when not specified here or in ibdiag.conf) is to use SM_Key
mgid <gid>	== 0 (or \"untrusted\")
Reversible", 'r', 1, NULL"	Source LID (PathRecord)
numb_path ", 'n', 1, NULL"	Destination LID (PathRecord)
pkey: P_Key (PathRecord, MCMemberRecord).	Multicast LID (MCMemberRecord)
qos_class (-Q)	Source GID (IPv6 format) (PathRecord)
sl	Destination GID (IPv6 format) (PathRecord)
mtu : (-M)	Port GID (MCMemberRecord)
rate (-R)	Multicast GID (MCMemberRecord)
pkt_lifetime	Reversible path (PathRecord)
qkey (-q) (PathRecord, MCMemberRecord).	Number of paths (PathRecord)
tclass (-T)	QoS Class (PathRecord)
flow_label : (-F)	Service level (PathRecord, MCMemberRecord)
hop_limit : (-H)	MTU and selector (PathRecord, MCMemberRecord)
scope	MTU and selector (PathRecord, MCMemberRecord)
join_state (-J)	Rate and selector (PathRecord, MCMemberRecord)
proxy_join (-X)	Packet lifetime and selector (PathRecord, MCMemberRecord)
service_id	If non-numeric value (like 'x') is specified then saquery will
	prompt for a value.
	Traffic Class (PathRecord, MCMemberRecord)
	Flow Label (PathRecord, MCMemberRecord)
	Hop limit (PathRecord, MCMemberRecord)
	Scope (MCMemberRecord)
	Join state (MCMemberRecord)
	Proxy join (MCMemberRecord)
	ServiceID (PathRecord)

Dependencies: OpenSM libvendor, OpenSM libopensm, libibumad

ibsysstat

```
ibsysstat [options] <dest lid|guid> [<op>]
```

Nonstandard flags:

```
Current supported operations:
ping - verify connectivity to server (default)
host - obtain host information from server
cpu - obtain cpu information from server
-o <oui> use specified OUI number to multiplex vendor mads
-S start in server mode (do not return)
```

ibnetsplit

Automatically groups hosts and creates scripts that can be run in order to split the network into sub-networks containing one group of hosts.

Syntax

- Group:

```
ibnetsplit [-v][-h][-g grp-file] -s <.lst|.net|.topo> <-r head-ports|-d max-dist>
```

- Split:

```
ibnetsplit [-v][-h][-g grp-file] -s <.lst|.net|.topo>  
-o out-dir
```

- Combined:

```
ibnetsplit [-v][-h][-g grp-file] -s <.lst|.net|.topo> <-r head-ports|-d max-dist> -o out-dir
```

Usage

- Grouping:

The grouping is performed if the -r or -d options are provided.

- If the -r is provided with a file containing group head ports, the algorithm examines the hosts distance from the set of node ports provided in the head-ports file (these are expected to be the ports running standby SM's).
- If the -d is provided with a maximum distance of the hosts in each group, the algorithm partition the hosts by that distance.



This method of analyzation may not be suitable for some topologies.

The results of the identified groups are printed into the file defined by the -g option (default ibnetsplit.groups) and can be manually edited. For groups where the head port is a switch, the group file uses the FIRST host port as the port to run the isolation script from.

- Splitting:

- If the -o flag is included, this algorithm analyzes the MinHop table of the topology and identifies the set of links and switches that may potentially be used for routing each group ports. The cross-switch links between switches of the group to other switches are declared as split-links and the commands to turn them off using Directed Routes from the original Group Head ports are written into the out-dir provided by the -o flag.

Both stages require a subnet definition file to be provided by the -s flag. The supported formats for subnet definition are:

- *.net - for ibnetdiscover
- *.lst - for opensm-subnet.lst or ibiagnet.lst
- *.topo - for a topology file

HEAD PORTS FILE

This file is provided by the user and defines the ports by which grouping of the other host ports is defined.

Format:

Each line should contain either the name or the GUID of a single port. For switches the port number shall be 0.

```
<node-name>/P<port-num>|<PGUID>
```

GROUPS FILE

This file is generated by the program if the head-ports file is provided to it. Alternatively it can be provided (or edited) by the user if different grouping is desired. The generated script for isolating or connecting the group should be run from the first node in each group.

Format:

Each line may be either:

```
GROUP: <group name>  
<node-name>/P<port-num>|<PGUID>
```

ibdiagnet

ibdiagnet scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices.

It then produces the following files in the output directory (see below):

- "ibdiagnet2.log" - A log file with detailed information.
- "ibdiagnet2.db_csv" - A dump of the internal tool database.
- "ibdiagnet2.lst" - A list of all the nodes, ports and links in the fabric.
- "[ibdiagnet2.pm](#)" - A dump of all the nodes PM counters.
- "ibdiagnet2.mlnx_cntrs" - A dump of all the nodes Mellanox diagnostic counters.
- "ibdiagnet2.net_dump" - A dump of all the links and their features.
- "ibdiagnet2.pkey" - A list of all pkeys found in the fabric.
- "ibdiagnet2.aguid" - A list of all alias GUIDs found in the fabric.
- "[ibdiagnet2.sm](#)" - A dump of all the SM (state and priority) in the fabric.
- "ibdiagnet2.fdb" - A dump of unicast forwarding tables of the fabric switches.
- "ibdiagnet2.mcfdb" - A dump of multicast forwarding tables of the fabric switches.
- "ibdiagnet2.svl" - A dump of SLVL tables of the fabric switches.
- "ibdiagnet2.nodes_info" - A dump of all the nodes vendor specific general information for nodes who supports it.
- "ibdiagnet2.plft" - A dump of Private LFT Mapping of the fabric switches.
- "[ibdiagnet2.ar](#)" - A dump of Adaptive Routing configuration of the fabric switches.
- "ibdiagnet2.vl2vl" - A dump of VL to VL configuration of the fabric switches.

Load plugins from:

```
/tmp/ibutils2/share/ibdiagnet2.1.1/plugins/
```

You can specify additional paths to be looked in with "IBDIAGNET_PLUGINS_PATH" env variable.

Plugin Name	Result	Comment
libibdiagnet_cable_diag_plugin-2.1.1	Succeeded	Plugin loaded
libibdiagnet_phy_diag_plugin-2.1.1	Succeeded	Plugin loaded

Syntax

```
[-i|--device <dev-name>] [-p|--port <port-num>]  
[-g|--guid <GUID in hex>] [--skip <stage>]  
[--skip_plugin <library name>] [--sc]  
[--scr] [--pc] [-P|--counter <<PM>=<value>>]
```

```

[--pm_pause_time <seconds>] [--ber_test]
[--ber_thresh <value>] [--llr_active_cell <64|128>]
[--extended_speeds <dev-type>] [--pm_per_lane]
[--ls <2.5|5|10|14|25|FDR10|EDR20>]
[--lw <1x|4x|8x|12x>] [--screen_num_errs <num>]
[--smp_window <num>] [--gmp_window <num>]
[--max_hops <max-hops>] [--read_capability <file name>]
[--write_capability <file name>]
[--back_compat_db <version.sub_version>]
[-V|--version] [-h|--help] [-H|--deep_help]
[--virtual] [--mads_timeout <mads-timeout>]
[--mads_retries <mads-retries>] [-m|--map <map-file>]
[--vlr <file>] [-r|--routing] [--r_opt <[vs,][mcast,]>]
[--sa_dump <file>] [-u|--fat_tree]
[--scope <file.guid>] [--exclude_scope <file.guid>]
[-w|--write_topo_file <file name>]
[-t|--topo_file <file>] [--out_ibnl_dir <directory>]
[-o|--output_path <directory>]
Cable Diagnostic (Plugin)
[--get_cable_info] [--cable_info_disconnected]
Phy Diagnostic (Plugin)
[--get_phy_info] [--reset_phy_info]

```

Options

```

-i|--device <dev-name> : Specifies the name of the device of the port
                        : used to connect to the IB fabric (in case
                        : of multiple devices on the local system).
-p|--port <port-num> : Specifies the local device's port number
                       : used to connect to the IB fabric.
-g|--guid <GUID in hex> : Specifies the local port GUID value of the
                        : port used to connect to the IB fabric. If
                        : GUID given is 0 then ibdiagnet displays
                        : a list of possible port GUIDs and waits
                        : for user input.
--skip <stage> : Skip the executions of the given stage.
               : Applicable skip stages (vs_cap_smp
               : vs_cap_gmp | links | pm |
               : speed_width_check | all).
--skip_plugin <library name> : Skip the load of the given library name.
               : Applicable skip plugins:
               : (libibdiagnet_cable_diag_plugin-2.1.1 |
               : libibdiagnet_phy_diag_plugin-2.1.1).
--sc : Provides a report of Mellanox counters
--scr : Reset all the Mellanox counters (if -sc
       : option selected).
--pc : Reset all the fabric PM counters.
-P|--counter <<PM>=<value>> : If any of the provided PM is greater than
                           : its provided value then print it.
--pm_pause_time <seconds> : Specifies the seconds to wait between
                           : first counters sample and second counters
                           : sample. If seconds given is 0 then all
                           : second counters sample will be done.
                           : (default=1).
--ber_test : Provides a BER test for each port.
           : Calculate BER for each port and check no
           : BER value has exceeds the BER threshold.
           : (default threshold="10^-12").
--ber_thresh <value> : Specifies the threshold value for the
                     : BER test. The reciprocal number of the
                     : BER should be provided. Example: for
                     : 10^-12 than value need to be
                     : 1000000000000 or 0xe8d4a51000
                     : (10^12). If threshold given is 0 then all
                     : BER values for all ports will be
                     : reported.
--llr_active_cell <64|128> : Specifies the LLR active cell size
                           : for BER test, when LLR is active in the
                           : fabric.
--extended_speeds <dev-type> : Collect and test port extended speeds
                              : counters. dev-type: (sw | all).
--pm_per_lane : List all counters per lane (when
              : available).
--ls <0|2.5|5|10|14|25|50|100|FDR10> : Specifies the expected link speed.
--lw <1x|4x|8x|12x> : Specifies the expected link width.
--screen_num_errs <num> : Specifies the threshold for printing
                        : errors to screen. (default=5).
--smp_window <num> : Max smp MADs on wire. (default=8).
--gmp_window <num> : Max gmp MADs on wire. (default=128).
--max_hops <max-hops> : Specifies the maximum hops for the
                       : discovery process. (default=64).
--read_capability <file name> : Specifies capability masks
                              : configuration file, giving capability
                              : mask configuration for the fabric.
                              : ibdiagnet will use this mapping for
                              : Vendor Specific MADs sending.
--write_capability <file name> : Write out an example file for
                              : capability masks configuration,
                              : and also the default capability
                              : masks for some devices.
--back_compat_db <version.sub_version> : Show ports section in
                                       : "ibdiagnet2.db_csv" according to
                                       : given version. Default version 2.0.
-V|--version : Prints the version of the tool.
-h|--help : Prints help information (without
           : plugins help if exists).
-H|--deep_help : Prints deep help information
               : (including plugins help).
--virtual : Discover VPorts during discovery
          : stage.
--mads_timeout <mads-timeout> : Specifies the timeout (in

```

```

        milliseconds) for sent and received
        mads. (default=500).
--mads_retries <mads-retries>      : Specifies the number of retries for
                                     every timeout mad. (default=2).
-m|--map <map-file>                : Specifies mapping file, that maps
                                     node guid to name
                                     (format: 0x[0-9a-fA-F]+ "name").
                                     Mapping file can also be specified by
                                     Environment variable
                                     "IBUTILS_NODE_NAME_MAP_FILE_PATH".
--src_lid <src-lid>                 : source lid
--dest_lid <dest-lid>               : destination lid
--dr_path <dr-path>                 : direct route path
-o|--output_path <directory>        : Specifies the directory where the
                                     Output files will be placed.
                                     (default="/var/tmp/ibdiagpath/").

Cable Diagnostic (Plugin)
--get_cable_info                    : Indicates to query all QSFP cables
                                     for cable information. Cable
                                     information will be stored
                                     in "ibdiagnet2.cables".

--cable_info_disconnected           : Get cable info on disconnected
                                     ports.

Phy Diagnostic (Plugin)
--get_phy_info                      : Indicates to query all ports for phy
                                     information.
--reset_phy_info                    : Indicates to clear all ports phy
                                     information.

```

ibdiagpath

ibdiagpath scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices. It then produces the following files in the output directory (see below):

- "ibdiagnet2.log" - A log file with detailed information.
- "ibdiagnet2.db_csv" - A dump of the internal tool database.
- "ibdiagnet2.lst" - A list of all the nodes, ports and links in the fabric.
- "ibdiagnet2.pm" - A dump of all the nodes PM counters.
- "ibdiagnet2.mlnx_cntrs" - A dump of all the nodes Mellanox diagnostic counters.
- "ibdiagnet2.net_dump" - A dump of all the links and their features.

Cable Diagnostic (Plugin):

This plugin performs cable diagnostic. It can collect cable info (vendor, PN, OUI etc..) on each valid QSFP cable, if specified.

It produces the following files in the output directory (see below):

- "ibdiagnet2.cables" - In case specified to collect cable info, this file will contain all collected cable info.

Phy Diagnostic (Plugin)

This plugin performs phy diagnostic.

Load Plugins from:

```
/tmp/ibutils2/share/ibdiagnet2.1.1/plugins/
```

You can specify additional paths to be looked in with "IBDIAGNET_PLUGINS_PATH" env variableLoad plugins from:

Plugin Name	Result	Comment
libibdiagnet_cable_diag_plugin-2.1.1	Succeeded	Plugin loaded
libibdiagnet_phy_diag_plugin-2.1.1	Succeeded	Plugin loaded

Syntax


```

[-i|--device <dev-name>] [-p|--port <port-num>]
[-g|--guid <GUID in hex>] [--skip <stage>]
[--skip_plugin <library name>] [--sc]
[--scr] [--pc] [-P|--counter <<PM>=<value>>]
[--pm_pause_time <seconds>] [--ber_test]
[--ber_thresh <value>] [--llr_active_cell <64|128>]
[--extended_speeds <dev-type>] [--pm_per_lane]
[--ls <2.5|5|10|14|25|FDR10|EDR20>]
[--lw <1x|4x|8x|12x>] [--screen_num_errs <num>]
[--smp_window <num>] [--gmp_window <num>]
[--max_hops <max-hops>] [--read_capability <file name>]
[--write_capability <file name>]
[--back_compat_db <version.sub_version>]
[-V|--version] [-h|--help] [-H|--deep_help]
[--virtual] [--mads_timeout <mads-timeout>]
[--mads_retries <mads-retries>] [-m|--map <map-file>]
[--src_lid <src-lid>] [--dest_lid <dest-lid>]
[--dr_path <dr-path>] [-o|--output_path <directory>]

Cable Diagnostic (Plugin)
[--get_cable_info] [--cable_info_disconnected]
Phy Diagnostic (Plugin)
[--get_phy_info] [--reset_phy_info]

```

Options

<pre> -i --device <dev-name> -p --port <port-num> -g --guid <GUID in hex> --skip <stage> --skip_plugin <library name> --sc --scr --pc -P --counter <<PM>=<value>> --pm_pause_time <seconds> --ber_test --ber_thresh <value> --llr_active_cell <64 128> --extended_speeds <dev-type> --pm_per_lane :List all counters per lane (when available). --ls <2.5 5 10 14 25 FDR10 EDR20> --lw <1x 4x 8x 12x> --screen_num_errs <num> --smp_window <num> --gmp_window <num> --max_hops <max-hops> --read_capability <file name> --write_capability <file name> --back_compat_db <version.sub_version> -V --version -h --help -H --deep_help --virtual --mads_timeout <mads-timeout> --mads_retries <mads-retries> -m --map <map-file> --src_lid <src-lid> --dest_lid <dest-lid> --dr_path <dr-path> -o --output_path <directory> Cable Diagnostic (Plugin) --get_cable_info --cable_info_disconnected Phy Diagnostic (Plugin) --get_phy_info --reset_phy_info </pre>	<pre> :Specifies the name of the device of the port used to connect to the IB fabric (in case of multiple devices on the local system). :Specifies the local device's port number used to connect to the IB fabric. :Specifies the local port GUID value of the port used to connect to the IB fabric. If GUID given is 0 than ibdiagnet displays a list of possible port GUIDs and waits for user input. :Skip the executions of the given stage. Applicable skip stages: (vs_cap_smp vs_cap_gmp links pm speed_width_check all). :Skip the load of the given library name. Applicable skip plugins: (libibdiagnet_cable_diag_plugin-2.1.1 libibdiagnet_phy_diag_plugin-2.1.1). :Provides a report of Mellanox counters :Reset all the Mellanox counters (if -sc option selected). :Reset all the fabric PM counters. :If any of the provided PM is greater than its provided value than print it. :Specifies the seconds to wait between first counters sample and second counters sample. If seconds given is 0 than no second counters sample will be done. (default=1). :Provides a BER test for each port. Calculate BER for each port and check no BER value has exceeds the BER threshold. (default threshold="10^-12"). :Specifies the threshold value for the BER test. The reciprocal number of the BER should be provided. Example: for 10^-12 than value need to be 1000000000000 or 0xe8d4a51000(10^12).If threshold given is 0 than all BER values for all ports will be reported. :Specifies the LLR active cell size for BER test, when LLR is active in the fabric. :Collect and test port extended speeds counters. dev-type: (sw all). :Specifies the expected link speed. :Specifies the expected link width. :Specifies the threshold for printing errors to screen. (default=5). :Max smp MADs on wire. (default=8). :Max gmp MADs on wire. (default=128). :Specifies the maximum hops for the discovery process. (default=64). :Specifies capability masks configuration file, giving capability mask configuration for the fabric. ibdiagnet will use this mapping for Vendor Specific MADs sending. :Write out an example file for capability masks configuration, and also the default capability masks for some devices. :Show ports section in "ibdiagnet2.db_csv" according to given version. Default version 2.0. :Prints the version of the tool. :Prints help information (without plugins help if exists). :Prints deep help information (including plugins help). :Discover VPorts during discovery stage. :Specifies the timeout (in milliseconds) for sent and received mads.(default=500). :Specifies the number of retries for every timeout mad. (default=2). :Specifies mapping file, that maps node guid to name (format: 0x[0-9a-fA-F]+ "name"). Mapping file can also be specified by environment variable "IBUTILS_NODE_NAME_MAP_FILE_PATH". :source lid destination lid :direct route path :Specifies the directory where the output files will be placed. (default="/var/tmp/ibdiagpath/"). :Indicates to query all QSFP cables for cable information. Cable information will be stored in "ibdiagnet2.cables". :Get cable info on disconnected ports. :Indicates to query all ports for phy information. :Indicates to clear all ports phy information. </pre>
--	--

12.2 Appendix - Supported Port Counters and Events

Port counters and events are available in the following views:

- Events and Port Counters area, at the bottom of the UFM window
- Error window (Error tab) in the Manage Devices tab
- In the New Monitoring Session window, in the Monitor tab, when clicking Create New Session
- Event Log in the Log tab (click Show Event Log)

12.2.1 InfiniBand Port Counters

The following tables list and describe the port counters and events currently supported:

- InfiniBand Port Counters
- Calculated Port Counters

<i>InfiniBand Port Counters</i>	
Counter	Description
Xmit Data (in bytes)	Total number of data octets, divided by 4, transmitted on all VLs from the port, including all octets between (and not including) the start of packet delimiter and the VCRC, and may include packets containing errors. All link packets are excluded. Results are reported as a multiple of four octets.
Rcv Data (in bytes)	Total number of data octets, divided by 4, received on all VLs at the port. All octets between (and not including) the start of packet delimiter and the VCRC are excluded and may include packets containing errors. All link packets are excluded. When the received packet length exceeds the maximum allowed packet length specified in C7-45: the counter may include all data octets exceeding this limit. Results are reported as a multiple of four octets.
Xmit Packets	Total number of packets transmitted on all VLs from the port, including packets with errors and excluding link packets.
Rcv Packets	Total number of packets, including packets containing errors and excluding link packets, received from all VLs on the port.
Rcv Errors	Total number of packets containing errors that were received on the port including: <ul style="list-style-type: none"> • Local physical errors (ICRC, VCRC, LPCRC, and all physical errors that cause entry into the BAD PACKET or BAD PACKET DISCARD states of the packet receiver state machine) • Malformed data packet errors (LVer, length, VL) • Malformed link packet errors (operand, length, VL) • Packets discarded due to buffer overrun (overflow)
Xmit Discards	Total number of outbound packets discarded by the port when the port is down or congested for the following reasons: <ul style="list-style-type: none"> • Output port is not in the active state • Packet length has exceeded NeighborMTU • Switch Lifetime Limit exceeded • Switch HOQ Lifetime Limit exceeded, including packets discarded while in VLStalled State.

<i>InfiniBand Port Counters</i>	
Counter	Description
Symbol Errors	Total number of minor link errors detected on one or more physical lanes.
Link Error Recovery	Total number of times the Port Training state machine has successfully completed the link error recovery process.
Link Error Downed	Total number of times the Port Training state machine has failed the link error recovery process and downed the link.
Local Integrity Error	The number of times that the count of local physical errors exceeded the threshold specified by LocalPhyErrors
Rcv Remote Physical Error	Total number of packets marked with the EBP delimiter received on the port.
Xmit Constraint Error	Total number of packets not transmitted from the switch physical port for the following reasons: <ul style="list-style-type: none"> • FilterRawOutbound is true and packet is raw • PartitionEnforcementOutbound is true and packet fails partition key check or IP version check
Rcv Constraint Error	Total number of packets received on the switch physical port that are discarded for the following reasons: <ul style="list-style-type: none"> • FilterRawInbound is true and packet is raw • PartitionEnforcementInbound is true and packet fails partition key check or IP version check
Excess Buffer Overrun Error	The number of times that OverrunErrors consecutive flow control update periods occurred, each having at least one overrun error
Rcv Switch Relay Error	Total number of packets received on the port that were discarded when they could not be forwarded by the switch relay for the following reasons: <ul style="list-style-type: none"> • DLID mapping • VL mapping • Looping (output port = input port)
VL15 Dropped	Number of incoming VL15 packets dropped because of resource limitations (e.g., lack of buffers) in the port
XmitWait	The number of ticks during which the port selected by PortSelect had data to transmit but no data was sent during the entire tick because of insufficient credits or of lack of arbitration.

<i>InfiniBand Calculated Port Counters</i>	
Counter	Description
Normalized XmitData	Effective port bandwidth utilization in % XmitData incremental/ Link Capacity
Normalized Congested Bandwidth	Amount of bandwidth that was suppressed due to congestion (XmitWait incremental/ Time) * Link Capacity Separate counters are used for Tier 4 ports and for the rest of the ports.

12.2.2 Supported Traps and Events

Device events are listed as VDM or CDM in the Source column of the Events table in the UFM GUI. For information about defining event policy, see [Configuring Event Management](#).

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
64	GID Address In Service	1	0	Info	1	300	Port	Fabric Notification	
65	GID Address Out of Service	1	0	Warning	1	300	Port	Fabric Notification	
66	New MCast Group Created	1	0	Info	1	300	Port	Fabric Notification	
67	MCast Group Deleted	1	0	Info	1	300	Port	Fabric Notification	
110	Symbol Error	1	1	Warning	200	300	Port	Hardware	
111	Link Error Recovery	1	1	Minor	1	300	Port	Hardware	
112	Link Downed	1	1	Critical	1	300	Port	Hardware	
113	Port Receive Errors	1	1	Minor	5	300	Port	Hardware	
114	Port Receive Remote Physical Errors	0	0	Minor	5	300	Port	Hardware	
115	Port Receive Switch Relay Errors	1	1	Minor	999	300	Port	Fabric Configuration	
116	Port Xmit Discards	1	1	Minor	200	300	Port	Communication Error	
117	Port Xmit Constraint Errors	1	1	Minor	200	300	Port	Communication Error	
118	Port Receive Constraint Errors	1	1	Minor	200	300	Port	Communication Error	
119	Local Link Integrity Errors	1	1	Minor	5	300	Port	Hardware	
120	Excessive Buffer Overrun Errors	1	1	Minor	100	300	Port	Communication Error	
121	VL15 Dropped	1	1	Minor	50	300	Port	Communication Error	

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
122	Congested Bandwidth (%) Threshold Reached	1	1	Minor	10	300	Port	Hardware	
131	Non-optimal link width (1X instead of 4X)	1	1	Minor	1	0	Port	Hardware	
132	Non-optimal link width (1X or 4X instead of 12X)	1	1	Minor	1	0	Port	Hardware	
140	Excessive Buffer Overrun Threshold Reached	1	0	Minor	11	300	Port	Hardware	
141	Flow Control Update Watchdog Timer Expired	1	0	Warning	1	300	Port	Hardware	
144	Capability Mask Modified	1	0	Info	1	300	Port	Fabric Notification	
145	System Image GUID changed	1	0	Info	1	300	Port	Communication Error	
256	Bad M_Key	1	0	Minor	1	300	Port	Security	
257	Bad P_Key	1	0	Minor	1	300	Port	Security	
258	Bad Q_Key	1	0	Minor	1	300	Port	Security	
259	Bad P_Key Switch External Port	1	0	Critical	1	300	Port	Security	
301	Logical Server State Changed	1	0	Info	1	0	Logical Server	Logical Model	
302	Logical Server State Change Failed	1	0	Minor	1	0	Logical Server	Logical Model	
306	Logical Server Added	1	0	Info	1	0	Logical Server	Logical Model	
307	Logical Server Removed	1	0	Info	1	0	Logical Server	Logical Model	
308	Logical Server Resources Allocated	1	0	Info	1	0	Logical Server	Logical Model	
312	Compute Resource Released	1	0	Info	1	0	Logical Server	Logical Model	

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
313	Compute Resource Allocated	1	0	Info	1	0	Logical Server	Logical Model	
314	Logical Server Additional Resources Allocated	1	0	Info	1	0	Logical Server	Logical Model	
315	Logical Server Resources Released	1	0	Info	1	0	Logical Server	Logical Model	
316	Logical Server Compute Resource is Down	1	1	Critical	1	0	Logical Server	Logical Model	
317	Logical Server Compute Resource is Up	1	1	Warning	1	0	Logical Server	Logical Model	
328	Link is Up	1	0	Info	1	0	Link	Fabric Topology	
328	Link is Down	1	0	Warning	1	0	Link	Fabric Topology	
331	Node is Down	1	0	Warning	1	0	Site	Fabric Topology	
332	Node is Up	1	0	Info	1	300	Site	Fabric Topology	
336	Port Action Succeeded	1	0	Info	1	0	Port	Maintenance	
337	Port Action Failed	1	0	Minor	1	0	Port	Maintenance	
338	Device Action Succeeded	1	0	Info	1	0	Port	Maintenance	
339	Device Action Failed	1	0	Minor	1	0	Port	Maintenance	
340	Network Interface Added	1	0	Info	1	0	Logical Server	Logical Model	
341	Network Interface Removed	1	0	Info	1	0	Logical Server	Logical Model	
350	Environment Added	1	0	Info	1	0	Env	Logical Model	
351	Environment Removed	1	0	Info	1	0	Env	Logical Model	
352	Network Added	1	0	Info	1	0	Network	Logical Model	

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
353	Network Removed	1	0	Info	1	0	Network	Logical Model	
370	Gateway Ethernet Link State Changed	1	0	Warning	1	0	Gateway	Gateway	
371	Gateway Reregister Event Received	1	0	Warning	1	0	Gateway	Gateway	
372	Number of Gateways Changed	1	0	Warning	1	0	Gateway	Gateway	
373	Gateway will be Rebooted	1	0	Warning	1	0	Gateway	Gateway	
374	Gateway Reloading Finished	1	0	Info	1	0	Gateway	Gateway	
381	Switch Upgrade Failed	1	0	Info	1	0	Switch	Maintenance	
383	Host Upgrade Failed	1	0	Info	1	0	Computer	Maintenance	
385	Switch FW Upgrade Started	1	0	Info	1	0	Switch	Maintenance	
386	Switch SW Upgrade Started	1	0	Info	1	0	Switch	Maintenance	
388	Host FW Upgrade Started	1	0	Info	1	0	Computer	Maintenance	
389	Host SW Upgrade Started	1	0	Info	1	0	Computer	Maintenance	
391	Switch Module Removed	1	0	Info	1	0	Switch	Fabric Notification	
392	Module Temperature Threshold Reached	1	0	Info	40	0	Module	Hardware	
394	Module Status FAULT	1	1	Critical	1	420	Switch	Module Status	
502	Device Upgrade Finished	1	0	Info	1	300	Device	Maintenance	

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
545	SM is not responding	1	1	Critical	1	300	Grid	Maintenance	
560	User Connected							Security	
561	User Disconnected							Security	
602	UFM Server Failover	1	1	Critical	1	0	Site	Fabric Notification	
701	Non-optimal Link Speed	1	1	Minor	1	0	Port	Hardware	
907	Switch is Down	1	1	Critical	1	0	Site	Fabric Topology	
908	Switch is Up	1	1	Info	1	300	Site	Fabric Topology	
909	Director Switch is Down	1	1	Critical	1	300	Site	Fabric Topology	
910	Director Switch is Up	1	1	Info	1	0	Site	Fabric Topology	
911	Module Temperature Low Threshold Reached	1	1	Warning	60	300	Module	Hardware	
912	Module Temperature High Threshold Reached	1	1	Critical	60	300	Module	Hardware	
913	Module High Voltage	1	1	Warning	10	420	Switch	Module Status	
914	Module High Current	1	1	Warning	10	420	Switch	Module Status	
915	BER_ERROR	1	1	Critical	1e-8	420	Port	Hardware	
916	BER_WARNING	1	1	Warning	1e-13	420	Port	Hardware	
917	SYMBOL_BER_ERROR	1	1	Critical		420	Port	Hardware	
1300	SM_SAKY_VIOLATION	1	1	Warning		5300	Port	Security	
1301	SM_SGID_SPOOFED	1	1	Warning		5300	Port	Security	
1302	SM_RATE_LIMIT_EXCEEDED	1	1	Warning		5300	Port	Security	

Alarm ID	Alarm Name	To Log	Alarm	Default Severity	Default Threshold	Default TTL	Related Object	Category	Description/Message
1303	SM_MULTICAST_GROUPS_LIMIT_EXCEEDED	1	1	Warning		5300	Port	Security	
1304	SM_SERVICES_LIMIT_EXCEEDED	1	1	Warning		5300	Port	Security	
1305	SM_EVENT_SUBSCRIPTION_LIMIT_EXCEEDED	1	1	Warning		5300	Port	Security	
1500	New cable detected	1	0	Info	1	0	Link	Security	
1502	Cable detected in a new location	1	0	Warning	1	0	Link	Security	
1503	Duplicate Cable Detected	1	0	Critical	1	0	Link	Security	
1600	VS/CC Classes Key Violation							Security	

12.3 Appendix - Used Ports

The following is the list of ports used by the UFM Server for internal and external communication:

Port	Purpose
80(tcp), 443(tcp)	Used by WS clients (Apache Web Server)
694(udp)	Used by Heartbeat - communication between UFM Primary and Standby server
3307(tcp)	Used for internal UFM Server communication with MySQL process
2222(tcp)	User for SSH debug console (optional. By default, this port is not used by the UFM server)
8000(udp)	Used for UFM server listening for REST API requests (redirected by Apache web server)
6306(udp)	Used for Multicast requests - communication with latest UFM Agents
8005(udp)	Used as UFM monitoring listening port
8089(tcp)	Used for internal communication between UFM server and MonitoringHistoryEngine
3308 (tcp)	Used for communication between MonitoringHistoryEngine and MonitoringHistory mysql server
8888(tcp)	Used by DRBD - communication between UFM Primary and Standby server

Port	Purpose
15800(tcp)	Used for communication with legacy UFM Agents on Mellanox Grid Director DDR switches
8081(tcp), 8082(tcp)	Used for internal communication with Subnet Manager

12.4 Appendix - Configuration Files Auditing

The main purpose of this feature is to allow users to track changes made to selected configuration files. When activating the feature, all the changes are reflected in specific log files which contain information about the changes and when they took place.

To activate this feature:

In *TrackConfig* section in *gv.cfg*, file value of *track_config* key should be set to true and value of *track_conf_files* key should contain a comma-separated list of defined conf files to be tracked. By default - ALL conf-files are tracked. To activate the feature, after *track_config* key is set to true, the UFM server should be restarted.

Example:

```
[TrackConfig]
# track config files changes
track_config = true
# Could be selected options (comaseparated) UFM, SM, SHARP, Telemetry. Or ALL for all the files.
track_conf_files = ALL
```

The below lists the configuration files that can be tracked:

Conf File Alias	Configuration Files
UFM	/opt/ufm/files/conf/gv.cfg
SM	/opt/ufm/files/conf/opensm/opensm.conf
SHARP	/opt/ufm/files/conf/sharp2/sharp_am.cfg
Telemetry	/opt/ufm/files/conf/telemetry/launch_ibdiagnet_config.ini
ALL	All the above configuration files.

Once the feature is activated and the UFM server is restarted, the UFM generates file which list the changes made in each of the tracked conf files. These files are located in */opt/ufm/files/auditing/* directory and the file naming convention is as follows: original conf file name with *audit.log* suffix.

Example: For *gv.cfg*, the name of the changes-tracking file is *gv.cfg.audit.log*. Changes are stored in auditing files in “linux diff”-like format.

Example:

```
cat /opt/ufm/files/auditing/gv.cfg.audit.log
=== Change occurred at 2022-07-24 07:31:48.679247 ===
---
+++
@@ -45,7 +45,7 @@
mon_mode_discovery_period = 60
check_interface_retry = 5
# The number of times to try if the InfiniBand fabric interface is down. The duration of each retry is 1 second.
```

```

-ibport_check_retries = 90
+ibport_check_retries = 92
ws_address = UNDEFINED
ws_port = 8088
ws_protocol = https

```

12.5 Appendix - IB Router


IB router provides the ability to send traffic between two or more IB subnets thereby potentially expanding the size of the network to over 40k end-ports, enabling separation and fault resilience between islands and IB subnets, and enabling connection to different topologies used by different subnets.

The forwarding between the IB subnets is performed using GRH lookup. The IB router's basic functionality includes:

- Removal of current L2 LRH (local routing header)
- Routing table lookup - using GID from GRH
- Building new LRH according to the destination according to the routing table

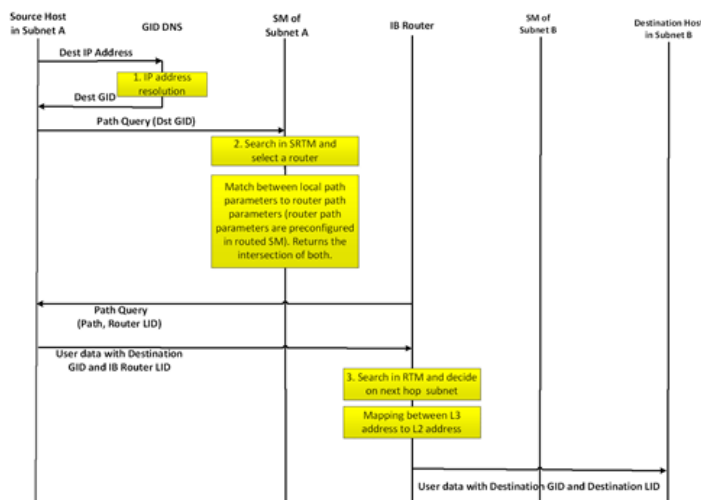
The DLID in the new LRH is built using simplified GID-to-LID mapping (where LID = 16 LSB bits of GID) thereby not requiring to send for ARP query/lookup.

Site-Local Unicast GID Format



For this to work, the SM allocates an alias GID for each host in the fabric where the alias GID = {subnet prefix[127:64], reserved[63:16], LID[15:0]}. Hosts should use alias GIDs in order to transmit traffic to peers on remote subnets.

Host-to-Host IB Router Unicast Flow



12.5.1 IB Router Scripts

The following scripts are supplied as part of UFM installation package.

12.5.1.1 set_num_of_subnets.sh

- Arguments

```
/opt/ufm/scripts/ib_router/set_num_of_subnets.sh --hostname <hostname> --username <username> --password <password> --num-of-subnets <num-of-subnets>
```

- Description - Configures system profile to InfiniBand allowing multiple switch IDs

- Syntax Description

hostname	IB router hostname or IP address
username	IB router username
password	IB router user password
num-of-subnets	Specified number of subnets (AKA SWIDs) to be initialized by the system. Value range: 2-6

- Example

```
/opt/ufm/scripts/ib_router/set_num_of_subnets.sh --hostname 10.6.204.12 --username admin --password admin --num-of-subnets 6
```



As a result of running this script, reboot is performed and all configuration is removed

12.5.1.2 add_interfaces_to_subnet.sh

- Arguments

```
/opt/ufm/scripts/ib_router/add_interfaces_to_subnet.sh --hostname <hostname> --username <username> --password <password> --interface <interface | interface-range> --subnet <subnet>
```

- Description

Maps an interface to a subnet and enables it

- SyntaxDescription

hostname	IB router hostname or IP address
username	IB router username
password	IB router user password
interface interface-range	Single IB interface or range of IB interfaces. Single IB interface: 1/<interface> Range of IB interfaces: 1/<interface>-1/<interface>

subnet	Name of IB subnet (AKA SWID): infiniband-default, infiniband-1...infiniband-5
--------	---

- Example

```
/opt/ufm/scripts/ib_router/add_interfaces_to_subnet.sh --hostname 10.6.204.12 --username admin --password admin --interface 1/1-1/6 --subnet infiniband-1
```

12.5.1.3 remove_interfaces_from_subnet.sh

- Arguments

```
/opt/ufm/scripts/ib_router/remove_interfaces_from_subnet.sh --hostname <hostname> --username <username> --password <password> --interface <interface | interface-range>
```

- Description

Un-maps an interface from a subnet after it has been disabled

- Syntax Description

hostname	IB router hostname or IP address
username	IB router username
password	IB router user password
interface interface-range	Single IB interface or range of IB interfaces. Single IB interface: 1/<interface> Range of IB interfaces: 1/<interface>-1/<interface>

- Example

```
/opt/ufm/scripts/ib_router/remove_interfaces_from_subnet.sh --hostname 10.6.204.12 --username admin --password admin --interface 1/6Example
```

12.5.1.4 add_subnet_to_router.sh

- Arguments

```
/opt/ufm/scripts/ib_router/add_subnet_to_router.sh --hostname <hostname> --username <username> --password <password> --subnet <subnet>
```

- Description

Creates routing on IB subnet interface and enables routing on that interface

- Syntax Description

hostname	IB router hostname or IP address
username	IB router username
password	IB router user password
subnet	Name of IB subnet (AKA SWID): infiniband-default, infiniband-1... infiniband-5

- Example

```
/opt/ufm/scripts/ib_router/add_subnet_to_router.sh --hostname 10.6.204.12 --username admin --password admin --subnet infiniband-3Example
```



As a result of running this script, the set of commands that allow control of IB router functionality is being enabled

12.5.1.5 remove_subnet_from_router.sh

- Arguments

```
/opt/ufm/scripts/ib_router/remove_subnet_from_router.sh --hostname <hostname> --username <username> --password <password> --subnet <subnet>
```

- Description

Destroys routing on IB subnet interface after routing on that interface has been disabled

- Syntax Description

hostname	IB router hostname or IP address
username	IB router username
password	IB router user password
subnet	Name of IB subnet (AKA SWID): infiniband-default, infiniband-1... infiniband-5

- Example

```
/opt/ufm/scripts/ib_router/remove_subnet_from_router.sh --hostname 10.6.204.12 --username admin --password admin --subnet infiniband-defaultExample
```

12.5.1.6 set_ufm_sm_router_support.sh

- Arguments

```
/opt/ufm/scripts/ib_router/set_ufm_sm_router_support.sh [-c <subnet prefix>] [-r] [-h]
```

- Description

[-c <subnet prefix>]: Used for updating OpenSM configuration file with new subnet prefix and forces OpenSM to re-read configuration.

[-r]: Used for resetting OpenSM configuration to default value and canceling IB routing.

- Syntax Description

-c	Configure new IB subnet prefix. Should be followed by new IB router subnet prefix value
-r	Reset to default
-h	Show help

- Example

```
/opt/ufm/scripts/ib_router/set_ufm_sm_router_support.sh -c 0xfec000000001234Examples
```

```
/opt/ufm/scripts/ib_router/set_ufm_sm_router_support.sh -r
```

12.5.2 IB Router Configuration

Step 1: Configure multi-switch. Run:

```
/opt/ufm/scripts/set_num_of_subnets.sh --hostname 10.6.204.12 --username admin --password admin --num-of-subnets 6
```

Step 2: Map interface to a subnet. Run:

```
/opt/ufm/scripts/add_ports_to_subnet.sh --hostname 10.6.204.12 --username admin --password admin --interface 1/1 --subnet infiniband-default
```

Step 3: Create routing on IB subnet interface. Run:

```
/opt/ufm/scripts/add_subnet_to_router.sh --hostname 10.6.204.12 --username admin --password admin --subnet infiniband-default
```

12.6 Appendix - NVIDIA SHARP Integration

12.6.1 NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™

NVIDIA SHARP is a technology that improves the performance of MPI operation by offloading collective operations from the CPU and dispatching to the switch network, and eliminating the need to send data multiple times between endpoints. This approach decreases the amount of data traversing the network as aggregation nodes are reached, and dramatically reduces the MPI operation time.

NVIDIA SHARP software is based on:

- Hardware capabilities in Switch-IB™ 2
- Hierarchical communication algorithms (HCOL) library into which NVIDIA SHARP capabilities are integrated
- NVIDIA SHARP daemons, running on the compute nodes
- NVIDIA SHARP Aggregation Manager, running on UFM

1. These components should be installed from HPCX or MLNX_OFED packages on compute nodes. Installation details can be found in SHARP Deployment Guide.

12.6.2 NVIDIA SHARP Aggregation Manager

Aggregation Manager (AM) is a system management component used for system level configuration and management of the switch-based reduction capabilities. It is used to set up the NVIDIA SHARP trees, and to manage the use of these entities.

AM is responsible for:

- NVIDIA SHARP resource discovery
- Creating topology aware NVIDIA SHARP trees
- Configuring NVIDIA SHARP switch capabilities
- Managing NVIDIA SHARP resources
- Assigning NVIDIA SHARP resource upon request
- Freeing NVIDIA SHARP resources upon job termination

AM is configured by a topology file created by Subnet Manager (SM): `subnet.lst`. The file includes information about switches and HCAs.

12.6.2.1 NVIDIA SHARP AM Prerequisites

In order for UFM to run NVIDIA SHARP AM, the following conditions should be met:

- Managed InfiniBand fabric must include at least one of the following Switch-IB 2 switches with minimal firmware version of 15.1300.0126:
 - CS7500
 - CS7510
 - CS7520
 - MSB7790
 - MSB7800
- NVIDIA SHARP software capability should be enabled for all Switch-IB 2 switches in the fabric (a dedicated logical port #37, for NVIDIA SHARP packets transmission, should be enabled and should be visible via UFM).
- UFM OpenSM should be running to discover the fabric topology.

NVIDIA SHARP AM is tightly dependent on OpenSM as it uses the topology discovered by OpenSM.

- NVIDIA SHARP AM should be enabled in UFM configuration by running:

```
[Sharp]
sharp_enabled = true
```

12.6.2.2 NVIDIA SHARP AM Configuration

By default, when running NVIDIA SHARP AM by UFM, there is no need to run further configuration. To modify the configuration of NVIDIA SHARP AM, you can edit the following NVIDIA SHARP AM configuration file: `/opt/ufm/files/conf/sharp/sharp_am.cfg`.

12.6.3 Running NVIDIA SHARP AM in UFM

➤ To run NVIDIA SHARP AM within UFM, do the following:

1. Make sure that the root GUID configuration file (root_guid.conf) exists in conf/opensm. This file is required for activating NVIDIA SHARP AM.
2. Enable NVIDIA SHARP in conf/opensm/opensm.conf OpenSM configuration file by running "ib sm sharp enable" or by setting the sharp_enabled parameter to 2:


```
# SHArP support
# 0: Ignore SHArP - No SHArP support
# 1: Disable SHArP - Disable SHArP on all supporting switches
# 2: Enable SHArP - Enable SHArP on all supporting switches
sharp_enabled 2
```

3. Make sure that port #6126 (on which NVIDIA SHARP AM is communicating with NVIDIA SHARP daemons) is not being used by any other application. If the port is being used, you can change it by modifying smx_sock_port parameter in the NVIDIA SHARP AM configuration file: conf/sharp2/sharp_am.cfg or via the command "ib sharp port".
4. Enable NVIDIA SHARP AM in conf/gv.cfg UFM configuration file by running the command "ib sharp enable" or by setting the sharp_enabled parameter to true (it is false by default):

```
[Sharp]
sharp_enabled = true
```

5. (Optional) Enable NVIDIA SHARP allocation in conf/gv.cfg UFM configuration file by setting the sharp_allocation_enabled parameter to true (it is false by default):

```
[Sharp]
sharp_allocation_enabled = true
```

 If the field sharp_enabled, and sharp_allocation_enabled are both set as true in gv.cfg, UFM sends an allocation (reservation) request to NVIDIA SHARP Aggregation Manager (AM) to allocate a list of GUIDs to the specified PKey when a new "Set GUIDs for PKey" REST API is called. If an empty list of GUIDs is sent, a PKEY deallocation request is sent to the SHARP AM.

NVIDIA SHARP allocations (reservations) allow SHARP users to run jobs on top of these resource (port GUID) allocations for the specified PKey. For more information, please refer to the *UFM REST API Guide* under Actions REST API → PKey GUIDs → Set/Update PKey GUIDs.

12.6.4 Operating NVIDIA SHARP AM with UFM

If NVIDIA SHARP AM is enabled, running UFM will run NVIDIA SHARP AM, and stopping UFM will stop NVIDIA SHARP AM.

➤ To start UFM with NVIDIA SHARP AM (enabled):

```
/etc/init.d/ufmd start
```

The same command applies to HA, using `/etc/init.d/ufmha`.

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

- To stop UFM with NVIDIA SHARP AM (enabled):

```
/etc/init.d/ufmd stop
```

- To stop only NVIDIA SHARP AM while leaving UFM running:

```
/etc/init.d/ufmd sharp_stop
```

- To start only NVIDIA SHARP AM while UFM is already running:

```
/etc/init.d/ufmd sharp_start
```

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

- To restart only NVIDIA SHARP AM while UFM is running:

```
/etc/init.d/ufmd sharp_restart
```

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

- To display NVIDIA SHARP AM status while UFM is running:

```
/etc/init.d/ufmd sharp_status
```

12.6.5 Monitoring NVIDIA SHARP AM by UFMHealth

UFMHealth monitors SHARP AM and verifies that NVIDIA SHARP AM is always running. When UFMHealth detects that NVIDIA SHARP AM is down, it will try to re-start it, and will trigger an event to the UFM to notify it that NVIDIA SHARP AM is down.

12.6.6 Managing NVIDIA SHARP AM by UFM High Availability (HA)

In case of a UFM HA failover or takeover, NVIDIA SHARP AM will be started on the new master node using the same configuration that was used prior to the failover/takeover.

12.6.7 NVIDIA SHARP AM Logs

NVIDIA SHARP AM log file (sharp_am.log) at /opt/ufm/files/log.

NVIDIA SHARP AM log files are rotated by UFM logrotate mechanism.

12.6.8 NVIDIA SHARP AM Version

NVIDIA SHARP AM version can be found at /opt/ufm/sharp/share/doc/SHARP_VERSION.

12.7 Appendix - AHX Monitoring

AHX Monitoring is a tool that is used to monitor AHX devices.

12.7.1 Overview

AHX monitoring enables monitoring HDR director switch cooling devices (i.e. AHX) and sends events to UFM.

The events are triggered on the switch associated with the cooling device if the monitoring utility encounters an issue.

The monitoring utility runs periodically and communicates with the AHX devices over the Modbus protocol (TCP port 502).

For deployment and configuration, please refer to the AHX Monitoring plugin in [Mellanox Docker HUB](#).

12.8 Appendix - UFM SLURM Integration

Simple Linux Utility for Resource Management (SLURM) is a job scheduler for Linux and Unix-like kernels.

By integrating SLURM with UFM, you can:

- Assign partition keys (pkeys) to SLURM nodes that are assigned for specific SLURM jobs.
- Create SHARP reservations based on SLURM nodes assigned for specific SLURM jobs.

12.8.1 Prerequisites

- UFM 6.9.0 (or newer) installed on a RedHat 7.x
- Python 2.7 on SLURM controller
- UFM-SLURM integration files (provided independently)

12.8.2 Automatic Installation

A script is provided to install the UFM-SLURM integration automatically.

1. Using the SLURM controller, extract the UFM-SLURM integration tar file:

```
tar -xf ufm_slurm_integration.tar.gz
```

2. Run the installation script using root privileges.

```
sudo ./install.sh
```

12.8.3 Manual Installation

To install the UFM-SLURM integration manually:

1. Extract the UFM-SLURM integration tar file:

```
tar -xf ufm_slurm_integration.tar.gz
```

2. Copy the UFM-SLURM integration files to the SLURM controller folder.
3. Change the permissions of the UFM-SLURM integration files to 755.
4. Modify the SLURM configuration file on the SLURM controller, `/etc/slurm/slurm.conf`, and add/modify the following two parameters:

```
PrologSlurmctld=/etc/slurm/ufm-prolog.sh  
EpilogSlurmctld=/etc/slurm/ufm-epilog.sh
```

12.8.4 UFM SLURM Config File

The integration process uses a configuration file located at `/etc/slurm/ufm_slurm.conf`. This file is used to configure settings and attributes for UFM-SLURM integration.

Here are the contents:

Attribute Name	Description
auth_type	Should be <code>token_auth</code> , or <code>basic_auth</code> . If you select <code>basic_auth</code> you need to set <code>ufm_server_user</code> and <code>ufm_server_pass</code> . If you select <code>token_auth</code> you need to set <code>token_auth</code> .
ufm_server_user	Username of UFM server used to connect to UFM if you set <code>auth_type=basic_auth</code>
ufm_server_pass	UFM server user password
token_auth=generated_token	Set <code>generated_token</code> , for more info how to generate token please see section Prolog and Epilog .
ufm_server	IP of UFM server to connect to
log_file_name	Name of integration logging file
partially_alloc	Determines whether or not to allow allocation of nodes



All of these attributes are mandatory.

12.8.5 Configuring UFM for NVIDIA SHARP Allocation

To configure UFM for NVIDIA SHARP allocation/deallocation you must set `sharp_enabled` and `enable_sharp_allocation` to true in `gv.cfg` file.

12.8.5.1 Generate token_auth

If you set `auth_type=token_auth` in UFM SLURM's config file, you must generate a new token by logging into the UFM server and running the following `curl` command:

```
curl -H "X-Remote-User:admin" -XPOST http://127.0.0.1:8000/app/tokens
```

Then you must copy the generated token and paste it into the config file beside the `token_auth` parameter.

12.8.6 Prolog and Epilog

After submitting jobs on SLURM, there are two scripts that are automatically executed:

- `ufm-prolog.sh` - the prolog script is executed when a job is submitted and before running the job itself. It creates the partition key (pkey) assignment and/or NVIDIA SHARP reservation and assigns the SLURM job hosts for them.
- `ufm-epilog.sh` - the epilog script is executed when a job is complete. It removes the partition key (pkey) assignment and/or NVIDIA SHARP reservation and free the associated SLURM job hosts.

12.8.7 Integration Files

The integration use scripts and configuration files to work, which should be copied to SLURM controller `/etc/slurm`. Here is a list of these files:

File Name	Description
<code>ufm-prolog.sh</code>	Bash file which executes jobs related to UFM after the SLURM job is completed
<code>ufm-epilog.sh</code>	Bash file which executes jobs related to UFM before the SLURM job is executed
<code>ufm_slurm.conf</code>	UFM-SLURM integration configuration file
<code>ufm_slurm_prolog.py</code>	Python script file which creates the partition key (pkey) assignment and/or SHARP reservation when the prolog bash script is running
<code>ufm_slurm_epilog.py</code>	Python script file which removes partition key (pkey) assignment and/or SHARP reservation based on the SLURM job hosts.
<code>ufm_slurm_utils.py</code>	Utility Python file containing functions and utilities used by the integration process

12.8.8 Running UFM-SLURM Integration

Using the SLURM controller, execute the following commands to run your batch job:

```
$ sbatch -N4 slurm_demo.sh
Submitted batch job 1
```



N4 is the number of compute nodes used to run the jobs. `slurm_demo.sh` is the job batch file to be run.

The output and result are stored on the working directory `slurm-{id}.out` where `{id}` is the ID of the submitted job.

In the above example, after executing `sbatch` command, you can see that the submitted job ID is 1. Therefore, the output file would be stored in `slurm-1.out`.

Execute the following command to see the output:

```
$ cat slurm-1.out
```

On the UFM side, a partition key (PKey) is assigned with all SLURM job IDs allocated to hosts. In case it was configured in `ufm_slurm.conf` file otherwise will use the default management PKey.

In addition, the UFM-SLURM will automatically create SHARM AM reservation in case UFM SHARP and UFM SHARP Allocation are enabled in UFM.

After the SLURM job is completed, the UFM removes the job-related partition key (pkey) assignment and SHARP reservation.

From the moment a job is submitted by the SLURM server until its completion, a log file named `/tmp/ufm_slurm.log` logs all of the actions and errors that occurred during the execution.

This log file can be changed by modifying the `log_file_name` parameter in `/etc/slurm/ufm_slurm.conf`.

12.9 Appendix - Device Management Feature Support

The following table describes the management features available on supported devices.

Feature	10 Gb Ethernet Gateway Module	Grid Director 4700/4200/4036/4036E v3.5	Managed IS5000 Switches	Managed SX6000 Switches	Externally Managed IS5000 / SX6000 Switches	Gateway BX5020	HP C-Class	Linux Hosts	Windows Hosts
Discovery									
IB L2 Discovery	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Advanced Discovery (IP, hostname, Hosts: CPU, memory, FW version)	Yes	Yes	No	Yes	No	No	No	Yes with UFM Host Agent	No
Ethernet access Management interface	Yes	Yes	Yes	Yes	No	No	No	Yes	Yes
Provisioning/ Configuration									
IB Partitioning (pkey)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
QoS: SL (SM configuration)	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
QoS: Rate Limit (SM configuration)	N/A	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interface/VIF Configuration (IP, hostname, mtu, Bonding)	N/A	N/A	N/A	N/A	N/A	No	N/A	Yes with UFM Host Agent	No
Device Monitoring									
Device Resources: CPU, Memory, Disk	No	Yes	No	No	No	No	No	Yes with UFM Host Agent	No
Get device alerts (Temperature, PS, Fan) Note: This feature is not supported on Switch-X switches.	Yes	Yes	No	Yes	Yes	No	No	No	No
L1 (Physical Port) - Monitoring	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
L2-3 (Interface/VIF) - Monitoring	No	No	No	No	No	No	No	Yes with UFM Host Agent	No

Congestion Monitoring per port (enables congestion map)	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Congestion Monitoring per flow (Advanced Package)	No	Yes	No	No	No	No	No	No	No
Device Management									
Add/remove to/from Rack	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Add/remove to/from Logical Server	N/A	N/A	N/A	N/A	N/A	N/A	N/A	Yes	Yes
View/clear Alarms	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SSH terminal to device	Yes	Yes	Yes	Yes	No	No	No	Yes	Yes
Power On	No	No	No	No	No	No	No	Yes with IPMI	No
Reboot	No	No	No	Yes (SX3606 only)	No	No	No	Yes with IPMI	No
Shutdown	No	No	No	No	No	No	No	Yes with IPMI	No
Port Enable/Disable	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firmware Upgrade (HCA & switch)	No	Yes	No	Yes (Upon SW upgrade - SX6036 only)	No	No	No	Yes	No
Inband Firmware Upgrade (over InfiniBand connection)	No	No	No	No	Yes	No	No	Yes	Yes
Software Upgrade (OFED & switch)	No	Yes	No	Yes (SX3606 only)	No	No	No	Yes with UFM Host Agent	No
Protocols									

Communication UFM Server - Device	IB/SNMP	IB/UDP/SSH	IB	IB/HTTP/SSH	IB	IB	IB	IB, SSH, IPMI, UDP	IB
-----------------------------------	---------	------------	----	-------------	----	----	----	--------------------	----

1. For a full list of supported IS5000 switches, see [Supported IS5000 Switches](#).
2. QoS Rate Limit (SM configuration): On ConnectX HCAs-only, for hosts.
3. XmitWait counter monitoring requires ConnectX HCAs with firmware version 2.6 and above.
4. This feature requires that the IP address is configured.

12.10 Appendix - UFM Event Forwarder

UFM Event Forwarder enables streaming of UFM events via FluentBit forwarder plugin to any external destination.

To deploy the UFM Event Forwarder on a Linux machine:

1. Connect to the Linux host via SSH.
2. Ensure the docker is installed on the host. Run:

```
# docker -version
```

3. Make sure that the docker service is up and running. If it is not, start the docker service. Run:

```
# sudo service docker start
```

4. Pull the UFM Event Forwarder image. Run:

```
# sudo docker pull mellanox/ufm-events-forwarder
```

Alternatively, if you do not have internet connection, contact NVIDIA Support to receive the UFM Event Forwarder docker image and load it to the host. Run:

```
# sudo cp <ufm-events-forwarder image path> /tmp/  
# sudo docker load -i /tmp/<image name>
```

5. If you are running in HA mode, repeat step 1-4 on the standby node.



Steps 6-9 should only be configured on the master node.

6. Enable the event-forwarder in main UFM config file. Run:

```
# vim /opt/ufm/files/conf/gv.cfg  
[Plugins]  
events_forwarder_enabled=true
```

7. Configure UFM to send events via syslog to the FluentBit event forwarder in gv.cfg.

```
[Logging]  
syslog_addr=127.0.0.1:5140  
syslog = true  
ufm_syslog = true  
event_syslog = true  
syslog_level = <severity>
```



<severity> may be set to any of the following values: CRITICAL, ERROR, WARNING, INFO, or DEBUG.

8. Configure the destination IP and port for the FluentBit event forwarder (requires Python 3):

```
# python /opt/ufm/scripts/events-forwarder/configure-fluent-bit.py -i <IP> -p <port>
```

Alternatively, if you have Python 2:

```
# /opt/ufm/venv_ufm/bin/python /opt/ufm/scripts/events-forwarder/configure-fluent-bit.py -i <IP> -p <port>
```

9. Start UFM. Run:

```
# /etc/init.d/ufmd start
```

Alternatively, if you are running in HA:

```
# /etc/init.d/ufmha start
```

10. Verify that UFM Event Forwarder is running successfully. Run:

```
# /etc/init.d/ufmd start
ufmd start
Starting opensm: [ OK ]
Starting MySQL: [ OK ]
Restarting httpd: [ OK ]
Starting snmpd: [ OK ]
Starting UFM main module: [ OK ]
Starting Events-Forwarder: [ OK ]
Starting Daily Report: [ OK ]
Starting UnhealthyPorts: [ OK ]
Starting ibpm: [ OK ]
```



Make sure the status of Events-Forwarder is OK.

Stopping UFM will also stop the Event Forwarder.

```
# /etc/init.d/ufmd stop
ufmd stop
Stopping ibpm: [ OK ]
Stopping Daily Report: [ OK ]
Stopping UnhealthyPorts: [ OK ]
Stopping Events-Forwarder: [ OK ]
Stopping UFM main module: [ OK ]
Stopping MySQL: [ OK ]
Stopping OpenSM: [ OK ]
```

After configuration, the Event Forwarder should always be running on the active node only. After a failover, for example, it will be stopped on the old master and will be started on the new active node.

If the destination IP and port are reconfigured (step 8), the Event Forwarder container should be restarted automatically with the newly applied configuration.

12.11 Appendix - UFM Multisite Portal Integration

NVIDIA® Mellanox® UFM® Enterprise Multisite Portal consolidates fabric information from several UFM servers into one central console. This provides the fabric administrator with a central view of devices, alerts, congestion, and other fabric health and performance information across all sites.

In order to configure UFM to work with the multisite portal, the following parameters must be set in the main UFM configuration file: `gv.cfg`.

```
[multisite]
enabled = true
#site_name is mandatory when multisite is enabled
site_name =
server = 10.213.1.122
port = 443
protocol=https
interval = 60
file=/opt/ufm/data/multisite/summary
max_files=60
```

Parameter	Description
enabled	Enables multisite agent in UFM
site_name	User-defined name which will be presented in the multisite portal
server	IPv4 address of the multisite portal server
port	The port to connect to on the multisite portal server
protocol	The communication protocol to use to connect to the multisite portal. The following options are available: <ul style="list-style-type: none">• https (default)• http• file (to save multisite agent summary information locally)
interval	Determines frequency in which data is sent by the multisite agent (in seconds)
file	Location where local summary data of the multisite agent is maintained
max_files	Maximum number of files to maintain

12.11.1 Configuring Multisite Agent Credentials

In order to configure the username and password of the multisite portal server, users must enter the `scripts` folder and run the following script:

```
cd /opt/ufm/scripts
./update_multisite_agent_creds.sh -u <USER> -p <PASSWORD>
```

For more options of configuring agent credentials, please run:

```
./update_multisite_agent_creds.sh -h
```

13 Document Revision History

Release	Date	Description
6.11.2	Jun 30, 2024	Updated Installation Notes
6.11.1	Dec 1, 2022	Updated the following sections: <ul style="list-style-type: none"> • Changes and New Features to include the upgrade of NVIDIA SHARP SW version • Installation Notes • Known Issues in This Release • Troubleshooting
	Dec 19, 2022	Updated Changes and New Features
	Jan 26, 2023	Updated Bug Fixes in This Release
6.11.0	Nov 21, 2022	<p>Updated the following sections:</p> <ul style="list-style-type: none"> • Added a link to UFM SDK 3.0 under Related Documentation • Changes and New Features • Installation Notes • Bug Fixes in This Release • Known Issues in This Release • Installing UFM HA Package • Network Map with new screenshots and new instructions for Map Information and Settings • Devices Window with new screenshots • PSID and Firmware Version In-Band Discovery • Groups Window with new screenshots • Table Enhancements with new screenshots • UFM Telemetry Fluent Streaming (TFS) Plugin • Enabling UFM Telemetry <p>Added the following sections:</p> <ul style="list-style-type: none"> • CPU Affinity on UFM • Switch Management IP Address Discovery • UFM Events Fluent Streaming (EFS) Plugin • In Telemetry <ul style="list-style-type: none"> • Changing UFM Telemetry Default Configuration • Supporting Generic Counters Parsing and Display • Supporting Multiple Telemetry Instances Fetch • Secondary Telemetry

Release	Date	Description
6.10.0	July 31, 2022	<p>Updated the following sections:</p> <ul style="list-style-type: none"> • Release Notes • UFM Installation and Initial Configuration • Installation Notes • UFM Software Architecture • Network Management • Subnet Manager Tab • Non-Optimal Links • Cable Transceiver Temperatures • Telemetry • Network Management • Docker Installation • Supported Actions for Internally Managed Switches • Appendix - NVIDIA SHARP Integration • Appendix - SM Default Files • Appendix - UFM Subnet Manager Default Properties • Appendix - SM Activity Report • Appendix - Configuration Files Auditing • Appendix - Enhanced Quality of Service • Appendix - Partitioning • Appendix - Diagnostic Utilities • Appendix - Adaptive Routing • Appendix - UFM SLURM Integration <p>Added the following sections:</p> <ul style="list-style-type: none"> • Showing UFM Processes Status • Plugin Management • Appendix - Configuration Files Auditing
	September 2022	<p>Updated:</p> <ul style="list-style-type: none"> • Appendix - UFM Event Forwarder • NDR switches firmware version in Supported NVIDIA Externally Managed Switches. • Licensing • License Devices limit in UFM Health Tab • Operating NVIDIA SHARP AM with UFM • Changes and New Features • Unsupported Functionalities/Features
	October 2022	<p>Updated the examples in Docker Installation</p>

Release	Date	Description
6.9	April 2022	<p>Added:</p> <ul style="list-style-type: none"> • Change UFM Telemetry Default Configuration • Configuring Log Rotation • SMTrap Handler Configuration • Auto-isolation of High-BER Ports • Auto-isolation of High-BER Ports • Time Zone Converter • Table Enhancements • Cable Transceiver Temperatures • Appendix - AHX Monitoring • Appendix - UFM SLURM Integration • User Preferences • UFM Telemetry Fluent Streaming (TFS) Plugin • Appendix - Configuration Files Auditing • Appendix - UFM Migration <p>Updated:</p> <ul style="list-style-type: none"> • Docker Installation • High Availability • Events & Alarms • Initial Configuration • UFM Web UI Main Navigation Buttons • Fabric Dashboard • Network Map • Devices Window • Ports Window • Telemetry • User Management Tab • Supported Traps and Events
6.8	November 30, 2021	<p>Added:</p> <ul style="list-style-type: none"> • Token Based Authentication • NDT Plugin • rest-rdma Plugin • Mark Device as Unhealthy • Mark Device as Healthy • Unhealthy Port Connectivity Filter • Security • Physical Grade and Eye Opening Information • Add Model Objects to Validation Test • Support Pkey with Virtual Ports <p>Updated:</p> <ul style="list-style-type: none"> • Historical Telemetry Collection in UFM • Telemetry • Fabric Validation Tab • Docker Installation

Release	Date	Description
6.7	July 05, 2021	<p>Added:</p> <ul style="list-style-type: none"> • Maximum Live Telemetry Sessions • Topology Compare • Data Streaming • Topology Compare Tab • Docker Installation <p>Updated:</p> <ul style="list-style-type: none"> • PKeys Window • Telemetry • Report Content • Fabric Validation Tab • IBDiagnet Tab • Appendix - UFM Event Forwarder • Appendix - Supported Traps and Events

14 High Availability

14.1 Overview of High Availability

UFM provides High Availability (HA) mechanisms to allow smooth fabric operation even if the UFM server fails or the connection between the UFM server and the rest of the fabric is not operating optimally.

UFM High Availability requires two distinct servers to run UFM software: one server is initially configured as the UFM active server and the other is configured as the UFM standby server. As a result, when the UFM active server fails or communication to the UFM active server ceases functioning, the UFM standby server takes over and becomes the new UFM active server. After such a failover, it is possible to repair the “old active UFM server” and bring it online as a new “UFM standby server.”

 Throughout this document, the following terms are used interchangeably:

Master—Active

Standby—Slave

UFM recovery relies on three mechanisms:

- UFM Database replication (from active to standby server)
- UFM Keep Alive (heartbeat) mechanism
- UFM server failover

For information about installing and running the UFM software for High Availability, see [Installing UFM Server Software for High Availability](#).


14.1.1 HA-Related Events

When the UFM server fails over to the UFM standby server, a UFM Failover event is generated.

14.1.2 HA-Related Considerations

We recommend that you locate the active and standby UFM servers in different sections of the fabric, so a single failure of an edge switch or a line card will not disconnect both UFM servers from the fabric.

We recommend that you bring up the failed UFM server as quickly as possible, to enable the fabric to sustain a possible secondary failure of the new active UFM server.

 **CAUTION:** A secondary failover (from the “new” active server to the “newly” brought up standby server) will succeed only after the UFM database’s initial replication as the “new standby server” has been completed. UFM can sustain a second failover only a few minutes after the new UFM standby server is up and running. This time depends on the size of the replicated partition and link speed (between the active and standby servers).

14.2 High Availability Functionality

The high availability capability is based on standard Linux packages - heartbeat and drbd.

Both heartbeat and drbd are installed on the master and slave nodes:

- drbd synchronizes a replicated partition between the two servers (but the partition itself / dev/drbd0 is visible only on the master node).
/opt/ufm/files is mounted on the drbd device and all data under this directory (partition) is replicated.
- Heartbeat is responsible for starting UFM on master node and stopping it on the slave. Heartbeat sends "keep alive" messages between the two servers, and when the master fails, the slave assumes mastership.



For high availability, use a reliable and high-capacity out-of-band management network (1 Gb Ethernet is recommended). Using inband IPoB will cause the HA split-brain condition if there is an InfiniBand network failure.

A *virtual IP (VIP) address* is an IP address that is not connected to a specific computer or Network Interface card (NIC) on a computer. Incoming packets are sent to the VIP address, but all packets travel through real Network Interfaces.

The VIP address belongs to the master node; failover of the system will result in failover of the virtual IP to the second node as well. When using UFM with HA, it is essential to always use the virtual IP instead of the server's IP to assure UFM operation on the master server.

Always use the virtual IP instead of the server's IP to assure connection to the UFM master server.

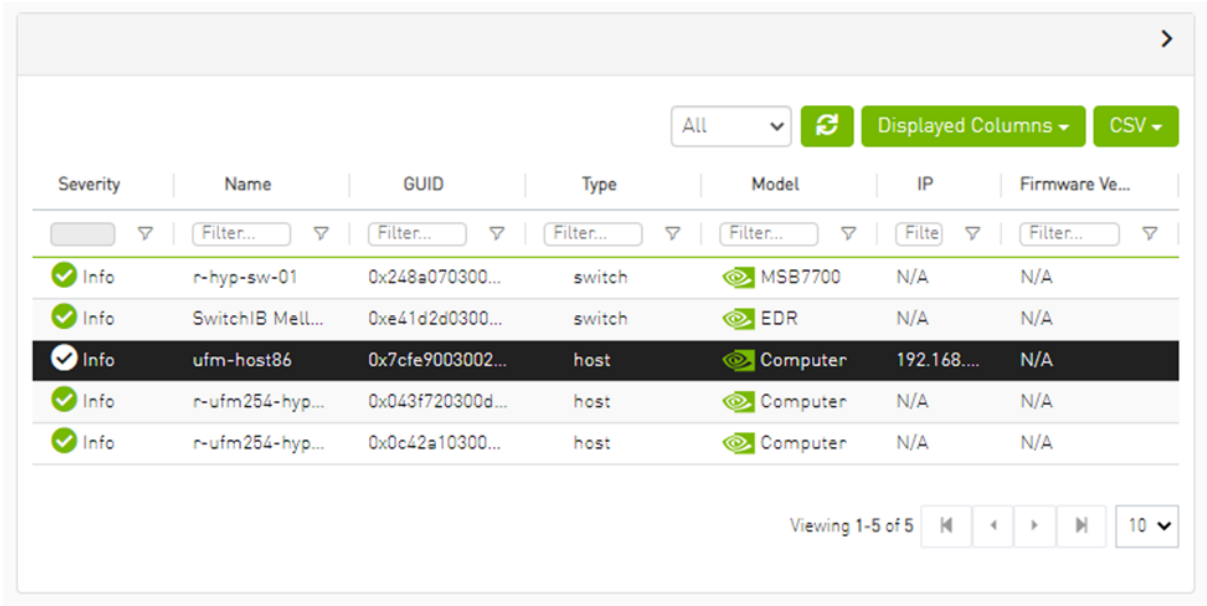
The failover will not happen if the standby server is not ready to take the mastership. UFM Health periodically checks the readiness of the standby server for the following:

- management network connectivity
- DRBD state
- disk space availability
- if the server is connected to the same InfiniBand cluster
- if the management InfiniBand port is Active and the IPoB interface is UP and RUNNING

If any of the condition above is not met, UFM Health will send a critical event. It is strongly recommended to repair the standby server as soon as possible to prevent risk of cluster malfunction.

15 Table Enhancements

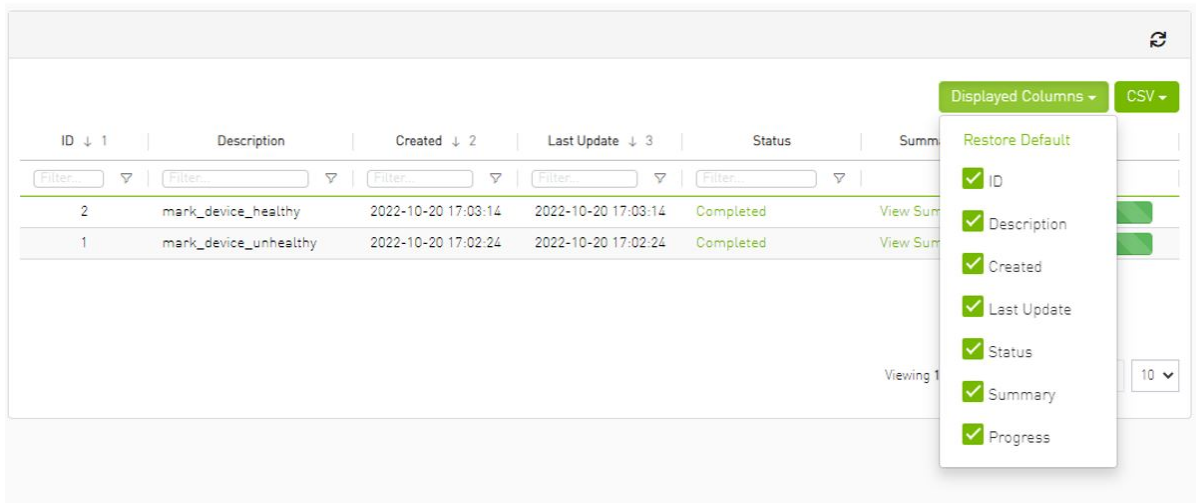
15.1 Look and Feel Improvements



The screenshot shows a table with columns: Severity, Name, GUID, Type, Model, IP, and Firmware Ve... The table contains five rows of data. The third row is highlighted. Above the table are controls for 'All', 'Displayed Columns', and 'CSV'. Below the table are pagination controls showing 'Viewing 1-5 of 5' and a page size of '10'.

Severity	Name	GUID	Type	Model	IP	Firmware Ve...
Info	r-hyp-sw-01	0x248a070300...	switch	MSB7700	N/A	N/A
Info	SwitchIB Mell...	0xe41d2d0300...	switch	EDR	N/A	N/A
Info	ufm-host86	0x7cfe9003002...	host	Computer	192.168....	N/A
Info	r-ufm254-hyp...	0x043f720300d...	host	Computer	N/A	N/A
Info	r-ufm254-hyp...	0x0c42a10300...	host	Computer	N/A	N/A


15.2 Displayed Columns



The screenshot shows a table with columns: ID, Description, Created, Last Update, Status, and Summary. A 'Displayed Columns' dropdown menu is open, showing a list of columns with checkboxes. The table contains two rows of data. The 'Displayed Columns' menu is open, showing a list of columns with checkboxes: ID, Description, Created, Last Update, Status, Summary, and Progress.

ID	Description	Created	Last Update	Status	Summary
2	mark_device_healthy	2022-10-20 17:03:14	2022-10-20 17:03:14	Completed	View Sum...
1	mark_device_unhealthy	2022-10-20 17:02:24	2022-10-20 17:02:24	Completed	View Sum...

- Restore Default
- ID
- Description
- Created
- Last Update
- Status
- Summary
- Progress

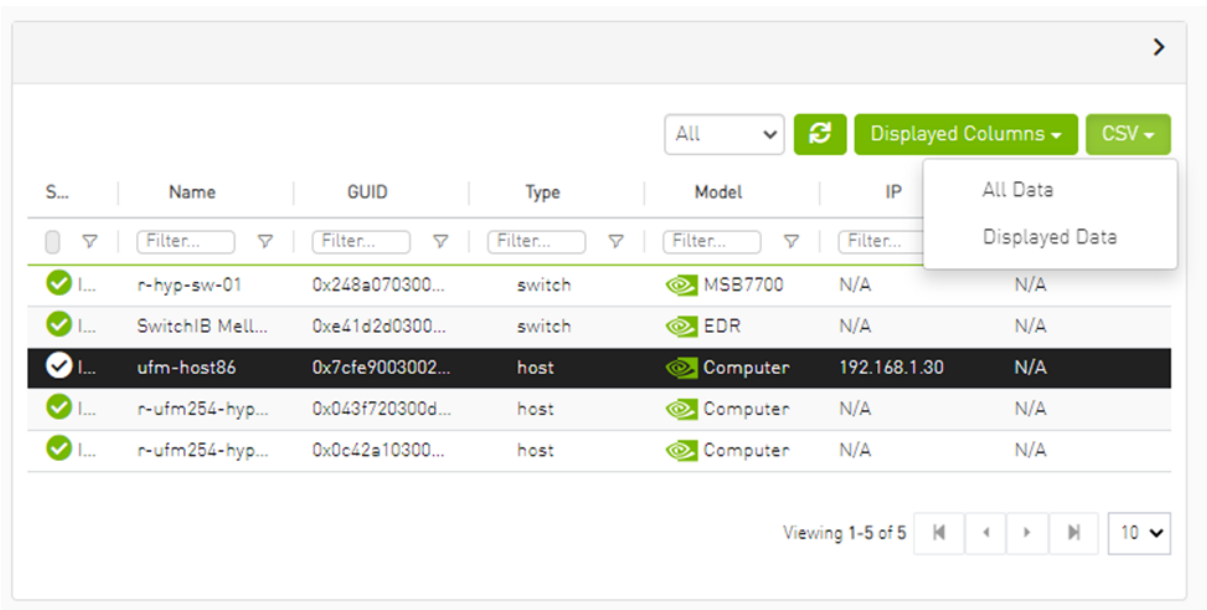


⚠ Displayed columns of all tables are persistent per user, with the option to restore defaults.

15.3 Export All Data as CSV

There are two options for exporting as CSV

- All Data: all data returned from server.
- Displayed Data: only displayed rows.




16 Time Zone Converter

Time zone converter provides the ability to unify all times in UFM like events and alarms, ibdiagnet, telemetry and logs.

The user can switch between local and machine time.

There is a drop-down menu in the status bar to switch between local and server/machine time.



Events & Alarms Local Time Last Update: 09 Apr 2022 17:25 admin

Alarms

Clear All Alarms Displayed Columns CSV

Severity	Date/Time ↓	Alarm Name	Source	Source Type	Reason	Count
Minor	2022-04-09 17:25:09	Non-optimal ...	default(3) / Switch: r-hyp-sw-01 /	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Minor	2022-04-09 17:25:09	Non-optimal ...	default(3) / Switch: SwitchIB Mell	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Minor	2022-04-09 17:25:09	Non-optimal ...	default(3) / Switch: SwitchIB Mell	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Warning	2022-04-05 15:26:47	Unhealthy IB...	default(3) / Switch: r-hyp-sw-01 /	IBPort	Peer Port is considered by SM as unhealthy due to MANUAL.	1
Warning	2022-04-05 15:26:27	Unhealthy IB...	default(3) / Switch: SwitchIB Mell	IBPort	Peer Port is considered by SM as unhealthy due to MANUAL.	1

Events & Alarms Server Time Last Update: 09 Apr 2022 11:31 admin

Alarms

Clear All Alarms Displayed Columns CSV

Severity	Date/Time ↓	Alarm Name	Source	Source Type	Reason	Count
Minor	2022-04-09 11:25:09	Non-optimal ...	default(3) / Switch: r-hyp-sw-01 /	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Minor	2022-04-09 11:25:09	Non-optimal ...	default(3) / Switch: SwitchIB Mell	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Minor	2022-04-09 11:25:09	Non-optimal ...	default(3) / Switch: SwitchIB Mell	IBPort	Found a [25.0] link that operates in [14.0] speed mode.	1180
Warning	2022-04-05 9:26:47	Unhealthy IB...	default(3) / Switch: r-hyp-sw-01 /	IBPort	Peer Port is considered by SM as unhealthy due to MANUAL.	1
Warning	2022-04-05 9:26:27	Unhealthy IB...	default(3) / Switch: SwitchIB Mell	IBPort	Peer Port is considered by SM as unhealthy due to MANUAL.	1

! In the screenshots, the difference between Server Time and Local Time is 6 hours.

17 Cable Transceiver Temperatures

The UFM has alarms that notify the user in cases where an active cable overheats/overcools.

The UFM uses ibdiagnet to get cable temperature analysis and report exceptions via the Alarms view.

17.1 GUI Views

17.1.1 Alarms

Severity	Date/Time ↓	Alarm Name	Source	Sourc...	Reason ▾	Count
Critical	2022-03-12 23:25:09	Cable Temperature High	default[3] / Switch: r-hyp-sw-4	IBPort	Cable High Temperature Alarm reported- current temperature: 116C- threshold: 70C	1
Critical	2022-03-12 23:25:09	Cable Temperature Low	default[3] / Computer: r-ufm2	IBPort	Cable Low Temperature Alarm reported- current temperature: 50C- threshold: 90C	1

17.1.2 Event Policy

Event ▾	Category	Mail	GUI	Alarm	Syslog	Log File	SNMP	Threshold	TTLSec	Severity
Cable Temperature High		<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		0	Critical
Cable Temperature Low		<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		0	Critical

17.2 Appendix - SM Partitions.conf File Format

This appendix presents the content and format of the *SM partitions.conf* file.

```

OpenSM Partition configuration
=====

The default partition will be created by OpenSM unconditionally even
when partition configuration file does not exist or cannot be accessed.

The default partition has P_Key value 0x7fff. OpenSM's port will always
have full membership in default partition. All other end ports will have
full membership if the partition configuration file is not found or cannot
be accessed, or limited membership if the file exists and can be accessed
but there is no rule for the Default partition.

Effectively, this amounts to the same as if one of the following rules
below appear in the partition configuration file:
In the case of no rule for the Default partition:
Default=0x7fff : ALL=limited, SELF=full ;
In the case of no partition configuration file or file cannot be accessed:
Default=0x7fff : ALL=full ;

File Format
=====

Comments:

Line content followed after '\#\ ' character is comment and ignored by
parser.

General file format:

<Partition Definition>:[<newline>]<Partition Properties>;

Partition Definition:
[PartitionName][=PKey][,ipoib_bc_flags][,defmember=full|limited]

PartitionName - string, will be used with logging. When omitted
empty string will be used.
PKey - P_Key value for this partition. Only low 15 bits will
be used. When omitted will be autogenerated.
ipoib_bc_flags - used to indicate/specify IPOIB capability of this partition.

defmember=full|limited - specifies default membership for port guid
list. Default is limited.

```

```

ipoib_bc_flags:
  ipoib_flag|[mgroup_flag]*

  ipoib_flag - indicates that this partition may be used for IPoIB, as
  a result the IPoIB broadcast group will be created with
  the flags given, if any.

Partition Properties:
  [<Port list>|<Mcast Group>]* | <Port list>

Port list:
  <Port Specifier>[,<Port Specifier>]

Port Specifier:
  <PortGUID>[=[full|limited]]

  PortGUID      - GUID of partition member EndPort. Hexadecimal
                 numbers should start from 0x, decimal numbers
                 are accepted too.
  full or limited - indicates full or limited membership for this
                 port. When omitted (or unrecognized) limited
                 membership is assumed.

Mcast Group:
  mgid=gid[,mgroup_flag]*<newline>

  - gid specified is verified to be a Multicast address
  IP groups are verified to match the rate and mtu of the
  broadcast group. The P_Key bits of the mgid for IP
  groups are verified to either match the P_Key specified
  in by "Partition Definition" or if they are 0x0000 the
  P_Key will be copied into those bits.

mgroup_flag:
  rate=<val> - specifies rate for this MC group
             (default is 3 (10Gbps))
  mtu=<val> - specifies MTU for this MC group
             (default is 4 (2048))
  sl=<val> - specifies SL for this MC group
           (default is 0)
  scope=<val> - specifies scope for this MC group
              (default is 2 (link local)). Multiple scope settings
              are permitted for a partition.
  NOTE: This overwrites the scope nibble of the specified
  mgid. Furthermore specifying multiple scope
  settings will result in multiple MC groups
  being created.
  qkey=<val> - specifies the Q_Key for this MC group
             (default: 0x0b1b for IP groups, 0 for other groups)
             WARNING: changing this for the broadcast group may
             break IPoIB on client nodes!!!
  tclass=<val> - specifies tclass for this MC group
              (default is 0)
  FlowLabel=<val> - specifies FlowLabel for this MC group
                 (default is 0)

newline: '\n'

```

Note that values for rate, mtu, and scope, for both partitions and multicast groups, should be specified as defined in the IBTA specification (for example, mtu=4 for 2048).

There are several useful keywords for PortGUID definition:

- 'ALL' means all end ports in this subnet.
- 'ALL_CAS' means all Channel Adapter end ports in this subnet.
- 'ALL_SWITCHES' means all Switch end ports in this subnet.
- 'ALL_ROUTERS' means all Router end ports in this subnet.
- 'SELF' means subnet manager's port.

Empty list means no ports in this partition.

Notes:

White space is permitted between delimiters ('=', ',', ':', ';').

PartitionName does not need to be unique, PKey does need to be unique. If PKey is repeated then those partition configurations will be merged and first PartitionName will be used (see also next note).

It is possible to split partition configuration in more than one definition, but then PKey should be explicitly specified (otherwise different PKey values will be generated for those definitions).

Examples:

```

Default=0x7fff : ALL, SELF=full ;
Default=0x7fff : ALL, ALL_SWITCHES=full, SELF=full ;

NewPartition , ipoib : 0x123456=full, 0x3456789034=limited, 0x2134af2306 ;

YetAnotherOne = 0x300 : SELF=full ;
YetAnotherOne = 0x300 : ALL=limited ;

ShareIO = 0x80 , defmember=full : 0x123451, 0x123452;
# 0x123453, 0x123454 will be limited
ShareIO = 0x80 : 0x123453, 0x123454, 0x123455=full;
# 0x123456, 0x123457 will be limited
ShareIO = 0x80 : defmember=limited : 0x123456, 0x123457, 0x123458=full;

```

```
ShareIO = 0x80 , defmember=full : 0x123459, 0x12345a;  
ShareIO = 0x80 , defmember=full : 0x12345b, 0x12345c=limited, 0x12345d;  
  
# multicast groups added to default  
Default=0x7fff, ipoib:  
  mgid=ff12:401b::0707, sl=1 # random IPv4 group  
  mgid=ff12:601b::16      # MLDv2-capable routers  
  mgid=ff12:401b::16      # IGMP  
  mgid=ff12:601b::2       # All routers  
  mgid=ff12::1, sl=1, Q_Key=0xDEADBEEF, rate=3, mtu=2 # random group  
  ALL=full;
```

Note:

The following rule is equivalent to how OpenSM used to run prior to the
partition manager:

```
Default=0x7fff, ipoib:ALL=full;
```


Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or



Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2024 NVIDIA Corporation & affiliates. All Rights Reserved.

