



NVIDIA UFM Enterprise Quick Start Guide v6.15.2

Table of Contents

1	UFM Installation Steps.....	5
1.1	Downloading UFM Software and License File	5
1.1.1	Obtaining License	5
1.1.2	Downloading UFM Software	6
1.2	Installing UFM Server Software	7
1.2.1	Prerequisites for UFM Server Software Installation.....	8
1.2.2	Additional Prerequisites for UFM High Availability Installation	9
1.2.3	Installing UFM Server on Bare Metal Server.....	10
1.2.4	Installing UFM Docker Container Mode	14
1.3	Activating Software License.....	20
1.4	Initial Configuration	21
1.4.1	Configuring Fabric Interface	21
1.5	Running UFM Server Software.....	21
1.5.1	Running UFM Server Software in Management Mode	22
1.5.2	Running UFM Software in High Availability Mode.....	22
1.5.3	Running UFM Software in Monitoring Mode	22
1.5.4	UFM Web UI Supported Browsers	23
1.5.5	Launching UFM Web UI Session	23
1.5.6	Licensing	24
1.5.7	More Information	26
2	Upgrading UFM Software.....	27
2.1	Upgrading UFM on Bare Metal Server	27
2.1.1	Upgrading UFM on Bare Metal - Standalone Server Upgrade	27
2.1.2	Upgrading UFM on Bare Metal - High Availability Upgrade	28
2.2	Upgrading UFM on Docker Container.....	30
2.2.1	Upgrading UFM on Docker Container in Standalone Mode	30
2.2.2	Upgrading UFM Container in High Availability Mode	31
3	Uninstalling UFM	34
3.1	Standalone Mode.....	34
3.2	High Availability	34
3.3	Docker Deployment.....	34
4	Appendix - Configuring UFM Over Static IPv4 Address	35

- 5 Appendix - Replacing the Standby Node in HA Mode 36
- 5.1 Replacing the Standby Node 36
- 6 Appendix - Restoring UFM Data 37
- 7 Document Revision History 38

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

About This Document

This guide is intended for field personnel and provides step-by-step instructions for installing and running the NVIDIA® UFM® Enterprise software, and installing and activating the license.

Related Documentation

- [NVIDIA End User License Agreement](#)
- [NVIDIA UFM Enterprise User Manual](#)
- [NVIDIA UFM REST API](#)
- [UFM High-Availability User Guide](#)

Technical Support

Customers who purchased NVIDIA products directly from NVIDIA are invited to contact us through the following methods:

- E-mail: enterprisesupport@nvidia.com
- Enterprise Support page: <https://www.nvidia.com/en-us/support/enterprise>

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding technical support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Document Revision History

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

1 UFM Installation Steps

To install the UFM software:

1. [Download the UFM software and license file](#)
2. [Install the UFM server software and activate the license file](#)
3. [Perform initial configuration](#)
4. [Run the UFM server software](#)

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.

1.1.1 Obtaining License

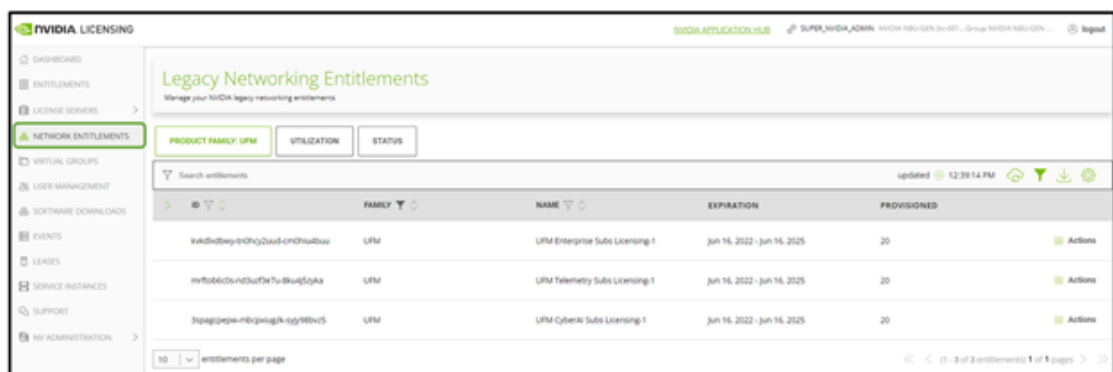
UFM is licensed per managed device according to the UFM license agreement.

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.



The screenshot shows the NVIDIA Licensing portal interface. The main heading is "Legacy Networking Entitlements" with a sub-heading "Manage your NVIDIA legacy networking entitlements". There are tabs for "PRODUCT FAMILY: UFM", "UTILIZATION", and "STATUS". Below the tabs is a search bar and a table of entitlements. The table has columns for FAMILY, NAME, EXPIRATION, and PROVISIONED. There are three rows of data, each with an "Actions" button.

FAMILY	NAME	EXPIRATION	PROVISIONED
UFM	UFM Enterprise Subs Licensing-1	Jun 16, 2022 - Jun 16, 2025	20
UFM	UFM Telemetry Subs Licensing-1	Jun 16, 2022 - Jun 16, 2025	20
UFM	UFM CyberAI Subs Licensing-1	Jun 16, 2022 - Jun 16, 2025	20

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.

Manage License File ✕

Make changes to the license allotment and generate a new file

ID	NAME	PROVISIONED	EXPIRATION
kvkdlxdbwy-tn0hcy2uud-cm0hiu4buu	UFM Enterprise Subs Licensing-1	20	Jun 16, 2022 - Jun 16, 2025

minx-ufm-kvkdlxdbwy-tn0hcy2uud-cm0hiu4buu-20220711143558.lic
license file generated Jul 11, 2022 5:37 PM last downloaded Jul 11, 2022 5:37 PM

MAC Address

Secondary MAC Address (optional)

GENERATE LICENSE FILE **DOWNLOAD LICENSE FILE**

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.

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.

1.2 Installing UFM Server Software

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

UFM Server installation options are:

- [Installing UFM Server on Bare Metal Server](#)
 - [Installing UFM on Bare Metal Server- Standalone Mode](#)
 - [Installing UFM on Bare Metal Server - High Availability Mode](#)
- [Installing UFM Docker Container Mode](#)
 - [Installing UFM on Docker Container - Standalone Mode](#)
 - [Installing UFM on Docker Container - High Availability Mode](#)

The following processes might be interrupted during the installation process:

- httpd (Apache2 in Ubuntu)
- dhcpd

To install UFM over static IPv4 configuration (instead of DHCP) please refer to [Appendix - Configuring UFM Over Static IPv4 Address](#) before installation.

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 in the UFM User Manual.

1.2.1 Prerequisites for UFM Server Software Installation

Verify that a supported version of Linux is installed on your machine. For details, see UFM System Requirements.

The following table lists the packages that must be installed on your machine (according to the system OS) before you install the UFM server software.

RedHat 7	RedHat 8	RedHat 9	Ubuntu 18.04	Ubuntu 20.04	Ubuntu 22.04
acl	acl	acl	acl	acl	acl
apr-util-openssl	apr-util-openssl	apr-util-openssl	apache2	apache2	apache2
bc	bc	bc	bc	bc	bc
cairo	gnutls	gnutls	chrpath	chrpath	chrpath
gnutls	httpd	httpd	cron	cron	cron
httpd	iptables	iptables-nft	gawk	gawk	gawk
iptables	jansson	jansson	lftp	lftp	lftp
lftp	lftp	lftp	libcurl4	libcurl4	libcurl4
libxml2	libnsl	libnsl	logrotate	logrotate	logrotate
libxslt	libxml2	libxml2	python3	python3	python3
mod_session	libxslt	libxslt	qperf	qperf	qperf
mod_ssl	mod_session	mod_session	rsync	rsync	rsync
net-snmp	mod_ssl	mod_ssl	snmpd	snmpd	snmpd
net-snmp-libs	net-snmp	net-snmp	sqlite3	sqlite3	sqlite3
net-snmp-utils	net-snmp-libs	net-snmp-libs	sshpas	sshpas	sshpas
net-tools	net-snmp-utils	net-snmp-utils	ssl-cert	ssl-cert	ssl-cert
php	net-tools	net-tools	sudo	sudo	sudo
psmisc	php	php	telnet	telnet	telnet
python3	psmisc	psmisc	zip	zip	zip
python3-libs	python36	python3			
qperf	qperf	qperf			
rsync	rsync	rsync			
sqlite	sqlite	sqlite			
sshpas	sshpas	sshpas			
sudo	sudo	sudo			
telnet	telnet	telnet			
zip	zip	zip			

In addition, ensure the following before you begin installation:

- The computer hostname is not defined as 127.0.0.1 and localhost is defined as 127.0.0.1.
- The hostname must NOT appear on the loopback address line. An example of the loopback address is: 127.0.0.1 localhost.localdomain localhost.
- Disable the firewall service (/etc/init.d/iptables stop), or ensure that the required ports are open (see the prerequisite script, refer to [Used Ports](#)).
- SELinux is disabled.
- If more than one fabric is managed by different UFM instances, set up different management network spaces for each fabric (not the same LAN).
- Uninstall any previously installed Subnet Manager from the UFM server machine.
- MLNX_OFED 5.x version is installed prior to installing UFM.
- As of UFM v.6.12.0, it is NOT mandatory to configure the IPoIB fabric interface with an IP address.

In cases where the IP is configured, it is mandatory that the IP is permanently configured and that it starts automatically upon server reboot (the IPoIB fabric interface should be active even if the network is down).

The user can set a persistent IP address using Netplan (mainly for Ubuntu systems) or modifying the interface network script (RedHat systems).

- The default MLNX_OFED installation includes opensm. Remove the MLNX_OFED opensm before UFM installation like the following examples:

RedHat:

```
rpm -e opensm-3.3.9.MLNX_20111006_e52d5fc-0.1
```

Ubuntu:

```
apt purge opensm
```

By default, ib0 and eth0 are configured as primary access points for the UFM management. If different management and/or InfiniBand interfaces (including bond interfaces) are used as the primary access points, you should modify the configuration file by running the script `/opt/ufm/scripts/change_fabric_config.sh` as described in the section Configuring General Settings in `gv.cfg`.

Change the UFM Agent interface to the Ethernet and/or IPoIB interfaces used for communication with UFM Agent:

```
ufma_interfaces = ib0,eth0
```

1.2.2 Additional Prerequisites for UFM High Availability Installation

- Reliable and high-capacity out-of-band IP connectivity between the UFM Primary and Secondary servers (1 Gb Ethernet is recommended). This connectivity is used for DRBD synchronization.

- Format two identical servers with dedicated disk partitions for UFM replication. Since the UFM configuration file is replicated to the standby server, both master and standby servers must have the same interfaces.
- Allocate exactly the same size partition on both servers (master and slave) for the replicated data. See UFM Server Requirements for the recommended partition size. Partitions should not be mounted and must be zeroed (the file system should not be installed on the partitions). For disk partitioning, see the Linux user manual (man fdisk).
- We recommend establishing a passwordless SSH (via /root/.ssh/authorized_keys file) between the two servers before the installation.
- In fabrics consisting of multiple tiers of switches, it is recommended that the management ports (ib0) of the primary and secondary UFM server be connected to different fabric switches on the same tier (the outermost edge in CLOS 5 designs).
This is because by default, UFM manages the IB fabric via ib0, port 1 of the HCA. Failure or disconnect of ib0, the IB management port, causes a failure condition in UFM resulting in HA failover.
When the management ports (ib0) of the primary and secondary UFM server are connected to the same switch, a failure of this switch will result in a disconnect of both UFM servers from the fabric, and therefore UFM will not be able to manage the fabric.

Subnet Manager is running over the native InfiniBand layer, therefore bonding the IpoIB interfaces will not provide high availability. For additional information, please refer to section UFM Failover to Another Port.

The UFM installation includes the InfiniBand Performance Management module (IBPM). This module is responsible for reporting performance information back to UFM and upper layer applications. When available, this process is offloaded to the non-management port (default ib1) of the UFM server. Failure or disconnect of the non-management port (ib1) on the primary UFM server will not cause UFM to failover. By default, the UFM Health Monitoring process is configured to try to restart the IBPM. For more information, see UFM Health Configuration in the UFM User Manual.

1.2.3 Installing UFM Server on Bare Metal Server

- [Installing UFM on Bare Metal Server- Standalone Mode](#)
- [Installing UFM on Bare Metal Server - High Availability Mode](#)

1.2.3.1 Installing UFM on Bare Metal Server- Standalone Mode

To install the UFM server software as a standalone for InfiniBand:

1. Create a temporary directory (for example */tmp/ufm*).
2. Open the UFM software zip file that you downloaded. The zip file contains the following installation files:
 - RedHat 7/CentOS 7/OEL 7: *ufm-X.X-XXX.el7.x86_64.tgz*
 - RedHat 8/CentOS 8: *ufm-X.X-XXX.el8.x86_64.tgz*
 - Ubuntu 18.04: *ufm-X.X-XXX.Ubuntu18.x86_64.tgz*
 - Ubuntu 20.04: *ufm-X.X-XXX.Ubuntu20.x86_64.tgz*
 - Ubuntu 22.04: *ufm-X.X-XXX.Ubuntu22.x86_64.tgz*

3. Extract the installation file for your system's OS to the temporary directory that you created.
4. From within the temporary directory, run the following command as root:

```
./install.sh
```

Running with the option "-o ib" is no longer required. For automatic installation, use the -q flag.

For "quiet" installation -q flag can be added (automatically answer yes for each question the installer asks).

Export MULTISUBNET_CONSUMER=1 environment variable before running the installation script to install the UFM server in Multisubnet Consumer mode.

The UFM software is installed. You can now remove the temporary directory.

1.2.3.2 Installing UFM on Bare Metal Server - High Availability Mode

Before installing UFM server software in High Availability mode, ensure that the [Additional Prerequisites for UFM High Availability Installation](#) are met.

The UFM High-Availability configuration requires dual-link connectivity based on two separate interfaces between the two UFM HA nodes. This configuration comprises of a primary link that is exclusively reserved for DRBD operations and a secondary link designated for backup purposes. Crucially, it is imperative that communication between the servers is established in a bidirectional manner across both interfaces and validated through user-initiated testing, such as a 'ping' command or other suitable alternatives, before HA configuration can be implemented. In cases where only one link is available among the two UFM HA nodes/servers, manually configure UFM with a single link. Refer to [Configure HA without SSH Trust \(Single Link Configuration\)](#).

UFM HA package requires a dedicated partition with the same name for DRBD on both servers. This guide uses `/dev/sda5` as an example.

1. On both servers, Install UFM Enterprise in Stand Alone (SA) mode.

Do not start UFM service.

2. Install the latest pcs and drbd-utils drivers on both servers.

For Ubuntu:

```
apt install pcs pacemaker drbd-utils
```

For CentOS/Red Hat:

```
yum install pcs pacemaker drbd84-utils kmod-drbd84
```

OR

```
yum install pcs pacemaker drbd90-utils kmod-drbd90
```

3. Download UFM-HA latest package from using this command:

```
wget https://download.nvidia.com/ufm/ufm_ha_5.3.0-17.tgz
```

For Sha256:

```
wget https://download.nvidia.com/ufm/ufm_ha_5.3.0-17.tgz.sha256sum
```

For more information on the UFM-HA package and all installation and configuration options, please refer to [UFM High-Availability User Guide](#)

4. Extract the downloaded UFM-HA package on both servers under `/tmp/.`
5. Go to the directory you extracted `/tmp/ufm_ha_XXX` and run the installation script. For example, if your DRBD partition is `/dev/sda5` run:

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

6. Configure the HA cluster. There are the three methods:

- [Configure HA with SSH Trust \(Dual Link Configuration\)](#) - Requires passwordless SSH connection between the servers.
- [Configure HA without SSH Trust \(Dual Link Configuration\)](#) - Does not require passwordless SSH connection between the servers, but asks you to run configuration commands on both servers.
- [Configure HA without SSH Trust \(Single Link Configuration\)](#) - Can be used in cases where only one link is available among the two UFM HA nodes/servers.

Configure HA with SSH Trust (Dual Link Configuration)

- a. On the master server only, configure the HA nodes. To do so, from `/tmp`, run the `configure_ha_nodes.sh` command as shown in the below example

```
configure_ha_nodes.sh \  
--cluster-password 12345678 \  
--master-primary-ip 10.10.10.1 \  
--standby-primary-ip 10.10.10.2 \  
--master-secondary-ip 192.168.10.1 \  
--standby-secondary -ip 192.168.10.2 \  
--no-vip
```

The script `configure_ha_nodes.sh` is located under `/usr/local/bin/`, therefore, by default, you do not need to use the full path to run it.

The `--cluster-password` must be at least 8 characters long.

To set up a Virtual IP for UFM and gain access to UFM through this IP, regardless of which server is running UFM, you may employ the `--no-vip` OR `--virtual-ip` command and provide an IP address as an argument. This can be achieved by navigating to `https://<Virtual-IP>/ufm` on your web browser.

When using back-to-back ports with local IP addresses for HA sync interfaces, ensure that you add your IP addresses and hostnames to the `/etc/hosts` file. This is needed to allow the HA configuration to resolve hostnames correctly based on the IP addresses you are using.

`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

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

Configure HA without SSH Trust (Dual Link Configuration)

If you cannot establish an SSH trust between your HA servers, you can use `ufm_ha_cluster` directly to configure HA. To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config -r standby \  
--local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

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

```
ufm_ha_cluster config -r master --local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

You must wait until after configuration for DRBD sync to finish, depending on the size of your partition. To check the DRBD sync status, run:

```
ufm_ha_cluster status
```

Configure HA without SSH Trust (Single Link Configuration)

This is not the recommended configuration and, in case of network failure, it might cause HA cluster split brain.

If you cannot establish an SSH trust between your HA servers, you can use `ufm_ha_cluster` directly to configure HA. To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config \  
-r standby \  
-e 10.212.145.5 \  
-l 10.212.145.6 \  
--enable-single-link
```

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

```
ufm_ha_cluster config -r master \  
-e 10.212.145.6 \  
-l 10.212.145.5 \  
-i 10.212.145.50 \  
--enable-single-link
```

You must wait until after configuration for DRBD sync to finish, depending on the size of your partition. To check the DRBD sync status, run:

```
ufm_ha_cluster status
```

Starting HA Cluster

- To start UFM HA cluster:

```
ufm_ha_cluster start
```

- To check UFM HA cluster status:

```
ufm_ha_cluster status
```

Stopping UFM HA cluster:

```
ufm_ha_cluster stop
```

For complete details on high availability, refer to [NVIDIA UFM High-Availability User Guide](#).

1.2.4 Installing UFM Docker Container Mode

1.2.4.1 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

1.2.4.2 Prerequisites for Upgrading UFM Docker Container

- Supported versions for upgrade are UFM v.6.10.0 and above.
- UFM files directory from previous container version mounted on the host.

1.2.4.3 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
```

If an Internet connection is not available, perform the following:

- Copy the UFM image to your machine.
- Load the image from the file using this command:

```
docker image load -i <image-path>
```

1.2.4.4 Step 2: Installing UFM Docker

1.2.4.4.1 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 /opt/ufm/files:/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_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.

1.2.4.4.1.1 Command Options

Flag	Description	Default Value
-f --fabric-interface	IB fabric interface name.	ib0
-g --mgmt-interface	Management interface name.	eth0
-h --help	Show help	N/A
-m --multisubnet-consumer	UFM Multisubnet Consumer mode	N/A

1.2.4.5 Installation Modes

UFM Enterprise installer supports several deployment modes:

- [Installing UFM on Docker Container - Standalone Mode](#)
- [Installing UFM on Docker Container - High Availability Mode](#)

1.2.4.6 Installing UFM on Docker Container - Standalone Mode

1. 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/
2. 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
```

3. Reload systemd:


```
systemctl daemon-reload
```

4. To Start UFM Enterprise service run:

```
systemctl start ufm-enterprise
```

1.2.4.7 Installing UFM on Docker Container - High Availability Mode

1.2.4.7.1 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.

1.2.4.7.2 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 the standby (secondary) server, install the 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
```

1.2.4.7.3 Downloading UFM HA Package

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

```
wget https://download.nvidia.com/ufm/ufm_ha_5.3.0-17.tgz
```

For SHA256:

```
wget https://download.nvidia.com/ufm/ufm_ha_5.3.0-17.tgz.sha256sum
```

1.2.4.7.4 Installing UFM HA Package

For more information on the UFM-HA package and all installation and configuration options, please refer to [UFM High Availability User Guide](#).

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. For example, if your DRBD partition is /dev/sda5 run the following command:

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

1.2.4.7.5 Configuring UFM HA

There are the three methods to configure the HA cluster:

- [Configure HA with SSH Trust \(Dual Link Configuration\)](#) - Requires passwordless SSH connection between the servers.
- [Configure HA without SSH Trust \(Dual Link Configuration\)](#) - Does not require passwordless SSH connection between the servers, but asks you to run configuration commands on both servers.
- [Configure HA without SSH Trust \(Single Link Configuration\)](#) - Can be used in cases where only one link is available among the two UFM HA nodes/servers.

1.2.4.7.5.1 Configure HA with SSH Trust (Dual Link Configuration)

1. On the master server only, configure the HA nodes. To do so, from /tmp, run the `configure_ha_nodes.sh` command as shown in the below example

```
configure_ha_nodes.sh \  
--cluster-password 12345678 \  
--master-primary-ip 10.10.50.1 \  
--standby-primary-ip 10.10.50.2 \  
--master-secondary-ip 192.168.10.1 \  
--standby-secondary-ip 192.168.10.2 \  
--no-vip
```

The script `configure_ha_nodes.sh` is located under `/usr/local/bin/`, therefore, by default, you do not need to use the full path to run it.

The `--cluster-password` must be at least 8 characters long.

When using back-to-back ports with local IP addresses for HA sync interfaces, ensure that you add your IP addresses and hostnames to the `/etc/hosts` file. This is needed to allow the HA configuration to resolve hostnames correctly based on the IP addresses you are using.

`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

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

```
ufm_ha_cluster status
```

1.2.4.7.5.2 Configure HA without SSH Trust (Dual Link Configuration)

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
```

To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config -r standby -e <peer ip address> -l <local ip address> -p <cluster_password>
```

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

```
ufm_ha_cluster config -r master -e <peer ip address> -l <local ip address> -p <cluster_password> -i <virtual ip address>
```

Configure HA without SSH Trust (Single Link Configuration)

This is not the recommended configuration and, in case of network failure, it might cause HA cluster split brain.

If you cannot establish an SSH trust between your HA servers, you can use `ufm_ha_cluster` directly to configure HA. To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config \  
-r standby \  
-e 10.212.145.5 \  
-l 10.212.145.6 \  
--enable-single-link
```

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

```
ufm_ha_cluster config -r master \  
-e 10.212.145.6 \  
-l 10.212.145.5 \  
-i 10.212.145.50 \  
--enable-single-link
```

You must wait until after configuration for DRBD sync to finish, depending on the size of your partition. To check the DRBD sync status, run:

```
ufm_ha_cluster status
```

IPv6 Example:

```
ufm_ha_cluster config -r standby -l fcfc:fcfc:209:224:20c:29ff:fee7:d5f2 -e fcfc:fcfc:209:224:20c:29ff:feeb:4962 --enable-single-link -p some_secret
```

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
```

1.3 Activating Software License

1. Before starting the UFM software, copy your license file(s) downloaded from [NVIDIA Licensing and Download Portal](#) (*volt-ufm-<serial-number>.lic*) to the master server under the */opt/ufm/files/licenses* directory. We recommend that you back up the license file(s). In High Availability mode, the license files are replicated to the standby machine automatically. Your software is now activated.
2. Run the UFM software as described in the following sections.

When a UFM license is not provided for activation upon the first UFM installation, the UFM runs on an auto-generated evaluation license which expires after 30 days from the first start-up of the UFM.

1.4 Initial Configuration

After installing the UFM® server software and before running UFM, perform the following:

- Mandatory Configuration:
 - Configure General Settings in `gv.cfg`
- [Additional Configurations](#) Options:
 - General Configuration options
 - Quality of Service
 - Activate and Enable Lossy Configuration Manager (Advanced License Only)
 - Activate and Enable Congestion Control Manager (Advanced License Only)

1.4.1 Configuring Fabric Interface

In most common cases, UFM is run in management mode; the UFM SM manages the InfiniBand fabric. In such cases, the only mandatory configuration is setting the `fabric_interface` parameter.

The fabric interface should be set to one of the InfiniBand IPoB interfaces, which connect the UFM/SM to the fabric:

```
fabric_interface = ib0
```

- By default, `fabric_interface` is set to `ib0`
- `fabric_interface` must be up and running before UFM startup. Otherwise, UFM will not be able to run.

For additional configuration options, please refer to the [UFM Enterprise Additional Configurations](#).

1.5 Running UFM Server Software

Before running UFM®:

- Perform initial configuration. Refer to [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](#) in the UFM User Manual.

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

- Management
- High Availability
- Monitoring
- Monitoring mode 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.

1.5.1 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`).

1.5.2 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.
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.

1.5.3 Running UFM Software in Monitoring Mode

Currently, running UFM in Monitoring mode is not fully supported as it was in the old Java-based UI. The unsupported functionalities in Monitoring mode are not reflected in the UFM web-based UI.

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).

Note that when running in Monitoring mode, the following parameters are automatically overwritten in the `/opt/ufm/files/conf/opensm/opensm.conf` file upon 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* in the UFM User Manual.
2. Ensure that the LMC value set in the `opensm` configuration file `/opt/ufm/files/conf/opensm/opensm.conf` is the same as the LMC value configured in the active SM.
3. 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`.
4. Go to the *About* window to verify that the Running mode is Monitoring.

For information about running Monitoring mode with failover to an external SM, see *Running UFM in Monitoring Mode with Failover to an External SM*. For information about Monitoring mode parameters and changing running mode from Monitoring mode to Management mode, see the *UFM User Manual*.

1.5.4 UFM Web UI Supported Browsers

UFM Web UI is supported on all the following web browsers: Internet Explorer, Firefox, Chrome and Opera.

For optimal UFM Web UI performance, make sure you are using the latest version available of Google Chrome.

For more information, see *UFM User Manual*.

1.5.5 Launching UFM Web UI Session

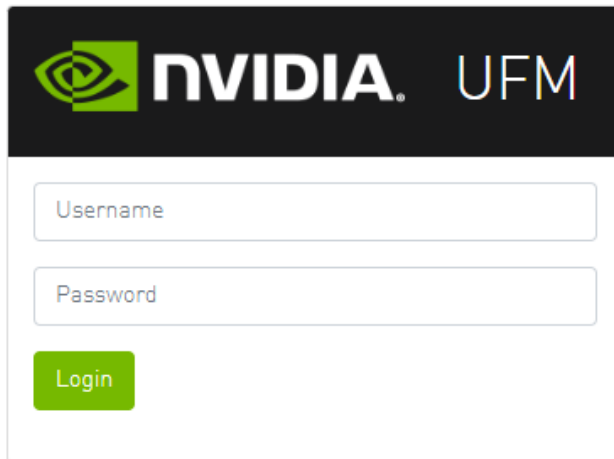
Before accessing the UFM Web UI:

- If required, you can change the configuration of the connection (port and protocol) between the UFM server and the APACHE server in the file `gv.cfg`:
 - `ws_protocol` = `http` or `https`
 - Setting the parameter `ws_protocol` to `http` allows unsecured access
 - Setting the parameter `ws_protocol` to `https` denies unsecured access.
 - `ws_port` = port number

➤ To launch a UFM Web UI session, do the following:

1. Launch the Web UI by entering the following URL in your browser:
`http://<UFM_server_IP>/ufm`

https://<UFM_server_IP>/ufm



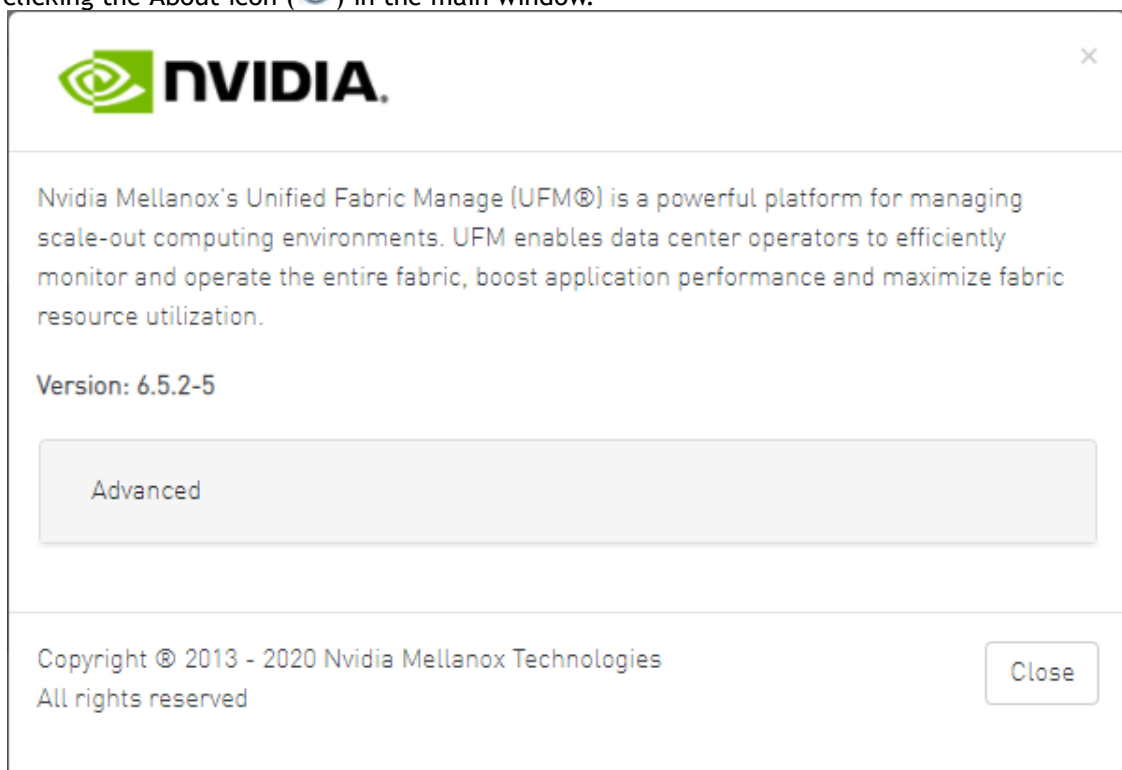
The image shows the login page for NVIDIA UFM. It features a black header with the NVIDIA logo and 'UFM' text. Below the header are two input fields: 'Username' and 'Password'. A green 'Login' button is positioned below the password field.

2. In the Login page, enter your User Name and your predefined user Password and click Login.


Once you have entered your user name and password, the main window shows the UFM Dashboard. For more information, see the Fabric Dashboard section in the *UFM User Manual*.

1.5.6 Licensing

1. After installing and activating your software, you can view your licenses in the Web UI by clicking the About icon (🌐) in the main window.



2. To view the advanced license information, click the Advanced button. The advanced license details will be displayed below.



- Product Functionality is updated only after startup. If you replace the UFM license, UFM continues to work in the previous mode until the UFM server is restarted.

➤ To view license information from the CLI:

Run CLI Command "ufmlicense" to display information about all installed licenses on the UFM server under /opt/ufm/files/licenses. This includes invalid and expired license information.

There are two UFM HA licenses where each license includes 2 different MACs: one for the primary machine and one for the standby machine.

In a given time, for each license, only one MACs is detected to be "Valid" (exists on the local machine) where the other MAC is detected as "Invalid" (exist on the standby machine).


See below *output example when running the CLI command* `ufmlicense` in SA and HA Modes.

HA Mode Output Example:

```
[root@ip-10-224-16-49-dg11 ~]# docker exec -ti ufm bash
root@ ip-10-224-16-49-dg11~# ufmlicense
|-----|
|-----|
|NVIDIA Corp|xxxxxxx-xxxxxxx|UFM Enterprise|Subscription |e4:43:4b:18:3c:e0|2025-08-29 |128
|3 Years |Invalid |
|-----|
|-----|
|NVIDIA Corp|xxxxxxx-xxxxxxx|UFM Enterprise|Subscription |e4:43:4b:18:49:a0|2025-08-29 |128
|3 Years |Valid |
|-----|
|-----|
|NVIDIA Corp|xxxxxxx-xxxxxxx|UFM Enterprise|Subscription |e4:43:4b:18:3c:e0|2025-07-13 |128
|3 |Invalid |
|-----|
|-----|
|BBK Electronics Corp Ltd|xxxxxxx-xxxxxxx|UFM Enterprise|Subscription |e4:43:4b:18:49:a0|
2025-07-13 |128 |3 |Valid |
```

SA Mode Output Example:

```
root@ufm-production:~# ufmlicense
|-----|
| Customer ID | SN | swName | Type | MAC Address | Exp. Date | Limit|
|-----|
|-----|
|495760397 |123456778 |UFM Enterprise |Evaluation |NA |2090-11-21 |1024 |
Advanced |Valid |
|-----|
|-----|
```

 To remove a license:

Delete the license file from /opt/ufm/files/licenses.

1.5.7 More Information

See the [UFM User Manual](#) for information on the following procedures:

- Configuring UFM failover to another port
- HTTP/HTTPS port configuration
- UFM SNMP agent configuration
- User Management

2 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 the section UFM Database and Configuration File Backup.

Upgrading the UFM Enterprise software version is supported up to two previous GA software versions (GA -1 or GA -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.

- [Upgrading UFM on Bare Metal Server](#)
- [Upgrading UFM on Docker Container](#)

2.1 Upgrading UFM on Bare Metal Server

2.1.1 Upgrading UFM on Bare Metal - Standalone Server Upgrade

You can upgrade the UFM standalone server software for InfiniBand from the previous UFM version.

To upgrade the UFM server software:

1. Create a temporary directory (for example `/tmp/ufm`).
2. Open the UFM software zip file that you downloaded. The zip file contains the following installation files for:
 - RedHat 7/CentOS 7/OEL 7: `ufm-X.X -XXX.el7.x86_64.tgz`
 - RedHat 8/CentOS 8/OEL 8: `ufm-X.X -XXX.el8.x86_64.tgz`
 - Ubuntu 18.04: `ufm-X.X -XXX.ubuntu18.x86_64.tgz`
 - Ubuntu 20.04: `ufm-X.X -XXX.ubuntu20.x86_64.tgz`
 - Ubuntu 22.04: `ufm-X.X -XXX.ubuntu22.x86_64.tgz`
3. Extract the installation file for your system's OS to the temporary directory that you created.
4. Stop the UFM server. Run:

```
systemctl stop ufm-enterprise
```

5. From within the temporary directory, run the following command as root:

```
./upgrade.sh
```

A configuration backup ZIP file will be created in the running directory (e.g. `/tmp/ufm`). The backup file name is `ufm_X.X.X_bkp.zip` (X.X.X is the previous version).

- a. Upgrade from the previous version: the existing UFM data and configuration are preserved.

- b. In case upgrade.sh script stops before completion (e.g. missing prerequisite), the upgrade procedure can be resumed by fixing the issue (e.g. installing missing prerequisite) and rerunning ./upgrade.sh again.

6. Restart the UFM server. Run:

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

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

7. After the upgrade, remove the temporary directory

2.1.2 Upgrading UFM on Bare Metal - High Availability Upgrade

As of UFM version 6.14.0, UFM upgrade on HA supports in-service upgrade, meaning UFM can continue running during the steps of the upgrade, and there is no need to stop UFM before the upgrade (although this is also supported).

You can upgrade the UFM server HA software for InfiniBand from the previous release. The upgrade is performed on both servers.

To upgrade the UFM server software:

1. On the standby server, extract the new UFM Enterprise package to the /tmp folder:

```
tar -xzf ufm-X.X.X-XXXXX.tgz -C /tmp
```

2. On the standby server, enter to the installation folder and upgrade script:

```
standby# cd /tmp/ ufm-X.X.X-X.<OS_NAME>.x86_64.mofed5/
```

3. Run the UFM upgrade script on the standby server:

```
./upgrade.sh
```

4. After the completion of the upgrade script, the UFM code will undergo an upgrade, while the UFM data will remain unchanged. The automatic upgrade of UFM data will take place during the next startup of UFM. To initiate this process, execute a failover from the Master node (or perform a takeover from the Standby node).

```
master# ufm_ha_cluster failover
```

UFM will log the data upgrade to the syslog of the server, in case of issue a backup of the UFM data is saved prior to the upgrade in /opt/ufm/BACKUP directory and can be restored. For more information, refer to [Appendix - Restoring UFM Data](#).

5. Once UFM is operational on the upgraded node (formerly the standby node), proceed to replicate steps 1 to 3 on the non-upgraded node (previously the master node).
6. On both servers, download latest UFM-HA package

```
wget https://www.mellanox.com/downloads/UFM/ufm_ha_5.1.1-6.tgz
```

7. On both servers, extract HA package under `/tmp/` and enter the new directory.
8. On both servers, run the installation command (use the same dedicated DRBD partition for UFM HA. In the following example `/dev/sda5` is used:

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

9. Configure HA. There are two methods:
 - [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.

[Configure HA with SSH Trust](#)

- i. On the master server only, configure the HA nodes. To do so, from `/tmp`, run the `configure_ha_nodes.sh` command as shown in the below example

```
configure_ha_nodes.sh
--cluster-password 12345678 \
--local-primary-ip 10.10.50.1 \
--peer-primary-ip 10.10.50.2 \
--local-secondary-ip 192.168.10.1 \
--peer-secondary-ip 192.168.10.2 \
--no-vip
```

The script `configure_ha_nodes.sh` is located under `/usr/local/bin/`, therefore, by default, you do not need to use the full path to run it.

The `--cluster-password` must be at least 8 characters long.

To set up a Virtual IP for UFM and gain access to UFM through this IP, regardless of which server is running UFM, you may employ the `--no-vip` OR `--virtual-ip` command and provide an IP address as an argument. This can be achieved by navigating to `https://<Virtual-IP>/ufm` on your web browser.

When using back-to-back ports with local IP addresses for HA sync interfaces, ensure that you add your IP addresses and hostnames to the `/etc/hosts` file. This is needed to allow the HA configuration to resolve hostnames correctly based on the IP addresses you are using.

`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

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

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. To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config -r standby \  
--local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

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

```
ufm_ha_cluster config -r master \  
--local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

You must wait until after configuration for DRBD sync to finish depending on the size of your partition.

2.2 Upgrading UFM on Docker 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.

2.2.1 Upgrading UFM on Docker Container in Standalone Mode

1. Stop the UFM Enterprise service. Run:

```
systemctl stop ufm-enterprise
```

2. Remove the existing docker image. Run:

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

3. Load the new UFM Enterprise 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
```

2.2.2 Upgrading UFM Container in High Availability Mode

As of UFM version 6.14.0, UFM upgrade on HA supports in-service upgrade, meaning UFM can continue running during the steps of the upgrade, and there is no need to stop UFM before the upgrade (although this is also supported).

1. Remove the old docker image from the standby server. Run:

```
Stand-by# docker rmi mellanox/ufm-enterprise:latest
```

2. Pull the new UFM Enterprise docker image on the standby server. Run:

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

At this stage, the UFM container has been updated with the latest code. The UFM data, however, will be updated during the next UFM run.

3. Perform a failover to start UFM on the upgraded node. On the master node, run:

```
ufm_ha_cluster failover
```

When UFM starts, it will automatically update the UFM configuration.

4. Repeat steps 1-2 on the un-upgraded node (previous Master node).
5. Download and extract the latest UFM HA package. Run:

```
wget https://download.nvidia.com/ufm/ufm_ha_5.3.0-17.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 --upgrade
```

7. Configure HA. There are two methods:

- [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.

Configure HA with SSH Trust

- i. On the master server only, configure the HA nodes. To do so, from /tmp, run the `configure_ha_nodes.sh` command as shown in the below example

```
configure_ha_nodes.sh
--cluster-password 12345678 \
--local-primary-ip 10.10.50.1 \
--peer-primary-ip 10.10.50.2 \
--local-secondary-ip 192.168.10.1 \
--peer-secondary-ip 192.168.10.2 \
--no-vip
```

The script `configure_ha_nodes.sh` is located under `/usr/local/bin/`, therefore, by default, you do not need to use the full path to run it.

The `--cluster-password` must be at least 8 characters long.

To set up a Virtual IP for UFM and gain access to UFM through this IP, regardless of which server is running UFM, you may employ the `--no-vip` OR `--virtual-ip` command and provide an IP address as an argument. This can be achieved by navigating to `https://<Virtual-IP>/ufm` on your web browser.

When using back-to-back ports with local IP addresses for HA sync interfaces, ensure that you add your IP addresses and hostnames to the `/etc/hosts` file. This is needed to allow the HA configuration to resolve hostnames correctly based on the IP addresses you are using.

`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

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

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. To configure HA, follow the below instructions:

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

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

```
ufm_ha_cluster config -r standby \  
--local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

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

```
ufm_ha_cluster config -r master --local-primary-ip 10.10.50.1 \  
--peer-primary-ip 10.10.50.2 \  
--local-secondary-ip 192.168.10.1 \  
--peer-secondary-ip 192.168.10.2 \  
--hacluster-pwd 123456789 \  
--no-vip
```

IPv6 Example:

```
configure_ha_nodes.sh  
--cluster-password 12345678 \  
--local-primary-ip fcfc:fcfc:209:224:20c:29ff:fee7:d5f2 \  
--peer-primary-ip fcfc:fcfc:209:224:20c:29ff:feeb:4962 \  
--local-secondary-ip fe80::1270:fd03:17:6365 \  
--peer-secondary-ip fe80::1270:fd03:17:5375 \  
--no-vip
```

You must wait until after configuration for DRBD sync to finish depending on the size of your partition.

8. Start UFM HA cluster. Run:

```
ufm_ha_cluster start
```

3 Uninstalling UFM

UFM® Server can be uninstalled by running an uninstall script in the different server modes:

- Standalone Mode
- High Availability
- Docker Deployment

3.1 Standalone Mode

To uninstall the UFM Server:

1. Go to `/opt/ufm`.
2. Run `./uninstall.sh`.

Child interfaces are not deleted.

3. To delete primary interfaces, restart `/etc/init.d/openibd`.

3.2 High Availability

To uninstall the UFM Server in high availability mode:

1. Run the following on the master and slave to clean up the UFM HA configuration:

```
ufm_ha_cluser cleanup
```

2. To uninstall the UFM HA configuration, run:

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

3. To uninstall UFM Enterprise software, run the following on the master and slave:

```
/opt/ufm/uninstall.sh
```

3.3 Docker Deployment

To uninstall the UFM Server in high availability mode:

1. Run the following on the master and slave:

```
ufm_ha_cluser cleanup
```

2. Run:

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

3. Run the following on the master and slave:

```
/opt/ufm/files/uninstall.sh
```

4 Appendix - Configuring UFM Over Static IPv4 Address

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

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

```
# vi /etc/sysconfig/network-scripts/ifcfg-enpls0
```

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

```
NAME="enpls0"DEVICE="enpls0"
ONBOOT="yes"
BOOTPROTO="static"
IPADDR="10.209.37.153"
NETMASK="255.255.252.0"
BROADCAST="10.209.39.255"
GATEWAY="10.209.36.1"
TYPE=Ethernet
DEFROUTE="yes"
```

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

```
# vi /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
10.209.37.153 <hostname>
```

3. Check hostname. Run:

```
# vi /etc/hostname
<hostname>
```

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

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

5. Restart network service. Run:

```
service network restart
```

6. Check Configuration. Run:

```
# hostname
<hostname>
# hostname -i
10.209.37.153
```

5 Appendix - Replacing the Standby Node in HA Mode

5.1 Replacing the Standby Node

- Install the HA package for the new node (standby).
- Disconnect the standby node (the old standby) and run the following command on the master node:

```
ufm_ha_cluster detach
```

- Configure the new standby node; please refer to the relevant section depending on the installation
- Connect the new standby to the cluster by running the command on the master node:

```
ufm_ha_cluster attach -l <local primary ip address> -e <peer primary ip address> -E <peer secondary ip address> -p <clust
```

6 Appendix - Restoring UFM Data

These instructions apply in high availability scenario only.

In the event of an in-service upgrade failure, the previous version of UFM's data will be safeguarded as a backup in the "/opt/ufm/BACKUP" directory, formatted as

"ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip."

To restore the data on the unupgraded node, follow these steps:

1. Copy the backup file from the upgraded node to the unupgraded node using the following command:

```
scp /opt/ufm/BACKUP/ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip root@<unupgraded_node_ip>:/opt/ufm/BACKUP/
```

2. Perform a failover of UFM to the master node, which is mandatory for data mount migration (including '/opt/ufm/files') to the master node: On the Master node, execute:

```
ufm_ha_cluster takeover
```

3. Stop UFM on the unupgraded node:

```
ufm_ha_cluster stop
```

4. Restore UFM configuration files from the backup:

```
/opt/ufm/scripts/ufm_restore.sh -f /opt/ufm/BACKUP/ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip
```

5. Start UFM on the unupgraded node (Note: Only the upgraded node can function until the upgrade issue is resolved, and failovers will not work).

Now, the issue that caused the upgrade failure can be addressed. If the problem is resolved, you can attempt the in-service upgrade again by failing UFM over to the upgraded node.

Alternatively, if needed, you can revert the changes made by reinstalling the old UFM version on the upgraded node.

7 Document Revision History

Revision	Date	Description of Changes
6.15.2	Jan 5, 2024	No changes
6.15.1	Dec 11, 2023	No changes
6.15.0	Nov 5, 2023	No changes
	Nov 30, 2023	Updated: <ul style="list-style-type: none"> • Loading UFM Docker Image - Added instructions on how to load UFM docker image when an Internet connection is not established • Configure HA without SSH Trust (Single Link Configuration) - Added an IPv6 example • Upgrading UFM on Docker Container
6.14.1	Aug 31, 2023	No changes
6.14.0	Aug 10, 2023	Updated: <ul style="list-style-type: none"> • Installation Command Options - Added the <code>-m</code> flag • Installing UFM Server on Bare Metal Server - Added a note about export <code>MULTISUBNET_CONSUMER=1</code> environment variable. • Upgrading UFM on Bare Metal - High Availability Upgrade • Upgrading UFM Container in High Availability Mode Added: <ul style="list-style-type: none"> • Appendix - Restoring UFM Data
6.13.1	May 18, 2023	No changes
	Jul 16, 2023	Added a third method to configure HA cluster in single link configuration. See Installing UFM on Bare Metal Server - High Availability Mode and Installing UFM on Docker Container - High Availability Mode .
6.13.0	May 5, 2023	Updated: <ul style="list-style-type: none"> • Installing UFM Server Software • Upgrading UFM Software
6.12.1	Feb 19, 2023	No changes
	Mar 16, 2024	Updated Initial Configuration
6.12.0	Feb 6, 2023	Updated: <ul style="list-style-type: none"> • Installing UFM Server Software for High Availability • Prerequisites for UFM Server Software Installation
6.11.1	Dec, 2022	Updated the following chapter: <ul style="list-style-type: none"> • High Availability Upgrade for RedHat7/8 and CentOS7/8

Revision	Date	Description of Changes
6.11.0	Nov, 2022	Updated the following chapters: <ul style="list-style-type: none"> • Prerequisites for UFM Server Software Installation • Installing UFM Server Software for High Availability • Initial Configuration • High Availability Upgrade for RedHat7/8 and CentOS7/8
6.10.0	Jul, 2022	Updated all chapters to reflect relevant changes.
	Oct, 2022	Updated Uninstalling UFM
	Nov, 2022	Updated High Availability Upgrade for RedHat7/8 and CentOS7/8
6.9	Apr, 2022	Updated the following chapters: <ul style="list-style-type: none"> • Installing UFM Server Software • Downloading UFM Software and License File • Upgrading UFM Software
6.8	Dec 09, 2021	No changes from UFM 6.7
6.7	Jul 05, 2021	Updated link reference in Deploying UFM as Docker Container
6.6	Nov 02, 2020	Added appendix Configuring UFM Over Static IPv4 Address
6.5.2	Nov 02, 2020	No changes from UFM 6.5.1
6.5.2	Nov 02, 2020	No changes from UFM 6.5.1
6.5.1	Aug 31, 2020	No changes from UFM 6.5

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.

