



NVIDIA UFM High-Availability User Guide

v6.1.1

Table of Contents

1	Release Notes.....	5
1.1	Changes and New Features.....	5
1.2	Bug Fixes in This Release	5
1.3	Known Issues in This Release	6
1.4	Bug Fixes History.....	6
2	Overview	7
2.1	UFM State	7
2.2	Connectivity Options	7
2.3	HA Cluster Resources	9
2.4	Cluster Network Access	10
2.5	Supported platforms.....	10
3	Prerequisites.....	11
3.1	Pacemaker Packages	11
3.2	DRBD	11
4	Installation and Configuration	12
4.1	Installation.....	12
4.2	Configuration	12
4.2.1	Configure HA with SSH Trust	13
4.2.2	Configure HA without SSH Trust	14
4.3	DRBD Configuration.....	15
4.4	NFS File Sharing	15
4.5	Multi-Nodes Support.....	15
4.5.1	Using File Configuration	15
4.5.2	Configuration File.....	16
4.6	UFM HA Cluster Operations.....	16
4.6.1	Show UFM HA version.....	16
4.6.2	Starting UFM HA Cluster	17
4.6.3	Checking UFM Cluster Status	17
4.6.4	Stopping UFM HA Cluster.....	17
4.6.5	Takeover Services	17
4.6.6	Master Failover	17
4.6.7	Automatic cleanup of failed actions	17

4.6.8	Replacing the Standby Node	18
4.6.9	Uninstalling UFM HA	18
5	Monitoring and Troubleshooting	19
6	Standby Node Replacement	21
6.1	Objectives.....	21
6.2	Prerequisites.....	21
6.3	Replacement Process	21
6.4	Run Command – Full Replacement Workflow	22
6.5	Resuming After a Failure.....	23
6.6	Aborting the Previous Run	23
6.7	Cluster Validation	24
7	UFM High-Level Architecture.....	25
7.1	FR#1.....	25
7.2	FR#2.....	25
7.3	Solution Options.....	25
7.3.1	FR#1.....	25
7.3.2	FR2#.....	26
8	Document Revision History	27

About This Document

This document describes NVIDIA® UFM High-Availability (HA) Architecture, connectivity, configuration options and monitoring procedures.

Software Download

To download the latest UFM High-Availability software package, please visit [NVIDIA's Licensing Portal](#).

Related Documents

Pacemaker	<ul style="list-style-type: none">• https://clusterlabs.org/pacemaker/doc/deprecated/en-US/Pacemaker/2.0/pdf/Clusters_from_Scratch/Pacemaker-2.0-Clusters_from_Scratch-en-US.pdf
DRBD	<ul style="list-style-type: none">• https://linbit.com/drbd/
Split-Brain	<ul style="list-style-type: none">• https://xahteivi.eu/resources/hints-and-kinks/solve-drbd-split-brain-4-steps/

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: [Enterprise Support Services](#)

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 Release Notes

- [Changes and New Features](#)
- [Bug Fixes in This Release](#)
- [Known Issues in This Release](#)
- [Bug Fixes History](#)

1.1 Changes and New Features

Feature	Description
Single Link Support for Replace Standby in UFM HA	Added the <code>--enable-single-link</code> option to the <code>ufm_replace_standby</code> command, enabling support for high availability (HA) configurations that use a single network interface.
Single Link Mode in HA Node Configuration	Updated the <code>configure_ha_nodes.sh</code> script to include a <code>--enable-single-link</code> option, allowing setup of single-interface HA environments.
UFM HA Cluster Restore Improvements	Added the <code>--config-only</code> option to the <code>ufm_ha_cluster_restore</code> command, enabling restoration of only configuration files without performing a full system restore.
Enhanced Visibility into UFM Service Startup	Integrated the UFM HA cluster service with the startup logger utility to provide better visibility into the UFM service startup process. This integration introduces detailed logging for each phase of HA cluster initialization, including system preparation and PCS resource activation, offering clearer insight into service startup progression.

1.2 Bug Fixes in This Release

Ref #	Description
#4617151	Description: Fixed an issue where the UFM HA cluster attach command did not work with HA package version
	Keywords: HA Cluster, HA Package
	Discovered in Release: v5.9.0
#4617150	Description: Fixed an issue with UFM upgrade that was caused by a timeout in starting the cluster.
	Keywords: Upgrade, Timeout, HA Cluster
	Discovered in Release: v5.9.0
#4601427	Description: Fixed an issue of UFM XDR Appliance node name not being adjustable
	Keywords: XDR Appliance, Node, Name
	Discovered in Release: v6.0.0

1.3 Known Issues in This Release

N/A

1.4 Bug Fixes History

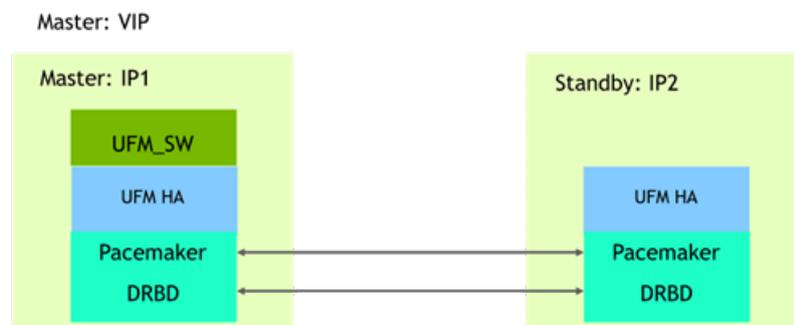
Ref #	Description
4161386	Description: Fixed issue where automatic failover did not occur due to previous failed actions.
	Keywords: Automatic Failover, HA
	Discovered in Release: v5.7.0
4128470	Description: Fixed the issue with HA Dual-Link that required running <code>systemctl restart drbd</code> after rebooting the standby node.
	Keywords: <code>systemctl</code> , <code>drbd</code> , Dual-link, Standy Reboot
	Discovered in Release: v5.7.0
4214538	Description: Fixed issue with UFM is not performing failover when management interface is down.
	Keywords: Failover, Management Interface
	Discovered in Release: 5.8.0
4143960	Description: Fixed the issue where "marking ringid 1 interface FAULTY" appeared in the syslog due to insufficient network bandwidth.
	Keywords: ringid, syslog, HA
	Discovered in Release: 1.8.1

2 Overview

UFM HA provides High-Availability on the host level for UFM products (UFM Enterprise/UFM Appliance Gen 3.0 and UFM Cyber-AI). The solution is based on Pacemaker to monitor host resources, services, and applications; and DRBD to sync file-system states. The HA package can be used with both bare-metal and Dockerized UFM deployments.

UFM HA should be installed on the master and standby nodes. The below figure describes the UFM Enterprise HA Architecture.

UFM ENTERPRISE SW HA



2.1 UFM State

The below files are replicated between the master and standby nodes:

```
/opt/ufm/files/*
```

Examples: log files, events, SQLite DB files (configuration, Telemetry history, persistent states topology groups).

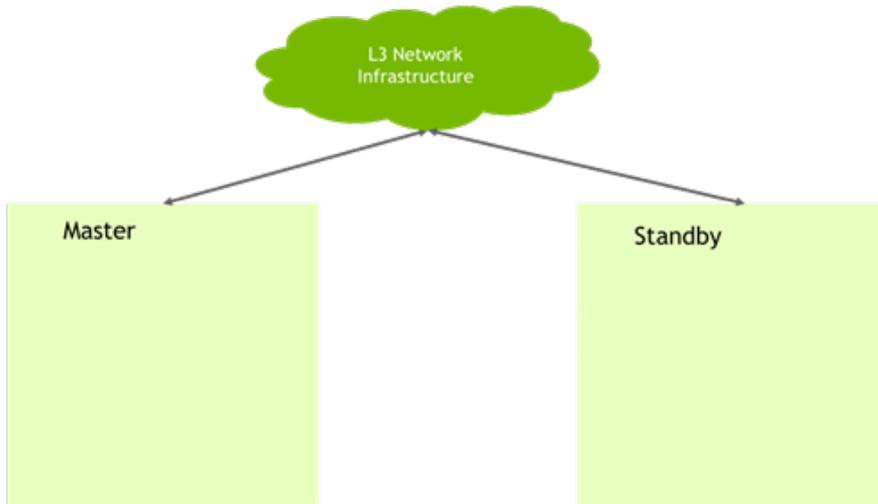
2.2 Connectivity Options

The master and standby nodes communicate with each other to establish and monitor a High-Availability solution. This connectivity is used by both the Pacemaker and DRBD. Below are connectivity options:

1. Cloud Connectivity. The following figure describes the external network infrastructure.

UFM ENTERPRISE HA CONNECTIVITY

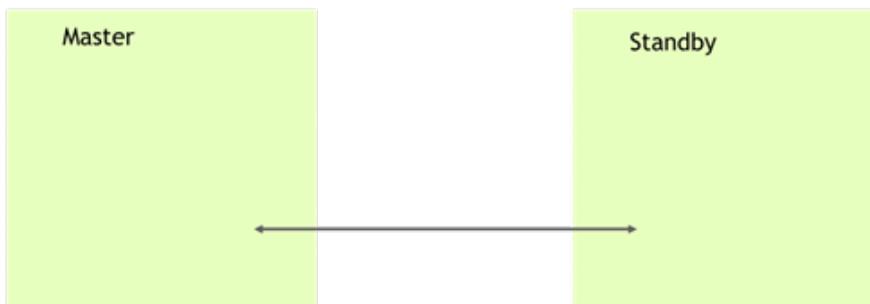
External Network Infrastructure



2. Back-to-back Connectivity, described in the following figure.

UFM ENTERPRISE HA CONNECTIVITY

Back to Back Connectivity



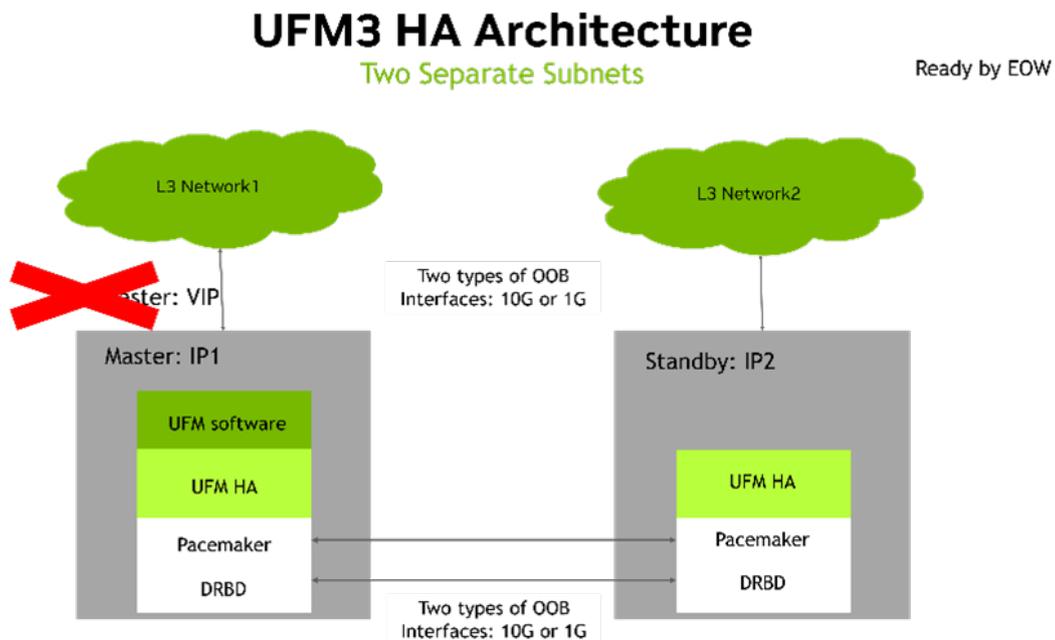
UFM-HA employs a dual-link configuration comprising primary and secondary connections to enhance system stability while mitigating the risk of connectivity challenges. It leverages two prioritized IP addresses, primary and secondary, which the Pacemaker utilizes to establish two connectivity links. Notably, DRBD utilizes the primary IP address to synchronize data. It is recommended to utilize this IP address for interfaces with high transfer rates such as InfiniBand interfaces for optimal

performance (IP over IB) and rapid DRBD synchronization. On the other hand, the secondary connectivity link may be effected via the management interface, typically an Ethernet interface.

DRBD and Pacemaker can use the same network interface or utilize different interfaces. For example, while the Pacemaker connectivity can be done through the management interface (usually an Ethernet interface), the DRBD synchronization could be done on an InfiniBand interface for better performance (IP over IB).

See below the configuration options for selecting a dif:

1. No VIP Connectivity Option



For some constrained network environments, the no VIP Connectivity option is supported. In this architecture, every UFM node has two physical IP addresses, primary and secondary. There is no VIP (floating) IP representing the whole cluster. This option allows two cluster nodes to lay in different subnets. In such a setup, clients who communicate with the UFM cluster should be aware of the active node status or constantly try to access both nodes.

2.3 HA Cluster Resources

The cluster software monitors the following HA cluster resources:

- UFM Enterprise
A systemd service runs and monitors all UFM Enterprise processes.
- UFM HA Watcher
The ufm-ha-watcher service monitors the health status of the UFM Enterprise and performs a failover in case the ufm-health process decides to perform a failover.
- Virtual IP
Also known as Cluster IP, a virtual IP is a unique IP resource allocated on the master node.

The virtual IP address should be reachable from any machine that uses it (REST API or UI). Virtual IP is not a mandatory configuration and can be omitted.

- DRBD and File System
DRBD needs its block device on each node. This can be a physical disk partition or a logical volume. The volume size planning should be done according to specific cluster sizing. The UFM-HA creates a DRBD resource and a filesystem resource with primary/secondary states based on the node if it is a master of a standby node.

2.4 Cluster Network Access

Cluster Network Access must consume UFM REST APIs and UI or performs management or monitoring tasks (ssh, scp, syslog, etc.).

For access to the UFM cluster, the below five IP addresses should be configured:

- Primary Physical IP1 - For the master node
- Secondary Physical IP1 - For the master node
- Primary Physical IP2 - For the standby node
- Secondary Physical IP2 - For the standby node
- Virtual (floating) IP (VIP)

Each two IP addresses of the same class should be configured in the same subnet and accessible (routable) by both cluster nodes. A virtual IP address should be in the subnet of one of the classes. The cluster manages the virtual IP address state. By default, the VIP is assigned to the master node. In case of failure of the master node, the VIP is moved by the cluster SW to the standby node. Network failures from the client to the UFM cluster are not monitored or handled by the HA cluster. Network infrastructure redundancy is out of the UFM HA solution scope. UFM HA cluster components utilize L3 and communication protocols (TCP/IP) for their internal communication and are agnostic to underlying L2 networking infrastructure.

2.5 Supported platforms

UFM HA is supported on the following Linux distributions:

1. Ubuntu 18.04, 20.04 and 22.04
2. CentOS7.7-9
3. CentOS8 Stream, RHEL8.5
4. CentOS9 Stream, RHEL9.X (2023)

3 Prerequisites

The following packages should be installed.

3.1 Pacemaker Packages

Pacemaker Package	Supported Versions
pacemaker	1.1.18 and 2.1.3
pcs	0.9.x, 0.10.x and 0.11.x
Corosync	2.4.3 and 3.1.5

3.2 DRBD

In the default sync mode, DRBD must be installed by the user. However, if NFS is chosen as the synchronization mechanism, DRBD is not mandatory.

DRBD	Supported Versions
DRBD utils	8.x.x, and 9.x.x

4 Installation and Configuration

4.1 Installation

The UFM HA package can be downloaded by running the following command:

```
wget http://www.mellanox.com/downloads/UFM/ufm_ha_6.1.1-2.tgz
```

The UFM HA package should be installed on both machines (Master and Standby) and the required UFM products. Installation order does not matter. To install the UFM-HA package:

- Untar the `ufm-ha` package:

```
tar xvfz ufm-ha-<version>.tgz
```

- Go to the directory you extracted and run the installation script. For example:

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

For NFS support, run the following installation script. For example:

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

Option	Description
<code>-l</code>	Sync Files Location. Must be always <code>/opt/ufm/files/</code>
<code>-d</code>	Disk name for DRBD. For example <code>/dev/sda5</code> (in case of using DRBD). Note that the <code>-d</code> option is not needed in case of NFS.
<code>-p</code>	Product Name. Must use “enterprise” to UFM Enterprise



In cases where you have a previous installation of `ufm_ha` and you want to upgrade to the newer version, run the following command:

```
./install.sh -u
```



UFM HA scripts are installed under `/usr/bin`.

4.2 Configuration

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.

4.2.1 Configure HA with SSH Trust

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 10.10.10.1 \  
--standby-primary-ip 10.10.10.2 \  
--master-secondary-ip 192.168.10.1 \  
--standby-secondary -ip 192.168.10.2 \  
--virtual-ip 10.10.10.5
```



The script `configure_ha_nodes.sh` is located under `/usr/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 ensure effective HA sync interface functionality for PCS version 0.9.X, employing back-to-back ports with local IP addresses, it is crucial to incorporate the relevant IP addresses and hostnames into the `/etc/hosts` file. This step is necessary to enable the HA configuration to accurately resolve hostnames based on the specific IP addresses in use.



`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



While configuring UFM HA on Oracle Linux, make sure the SELinux is disabled. You can check SELinux status with `sestatus`.
If it is enabled, follow the below steps to disable it:

- Run `vi /etc/selinux/config`
- Add `SELINUX=disabled`
- Reboot the machine
- Verify SELinux is disabled with the command `sestatus`.

Option	Description
<code>--cluster-password</code>	UFM HA cluster password for authentication by the pacemaker.
<code>--master-ip</code>	Master (main) server IP address
<code>--standby-ip</code>	Standby server IP address
<code>--virtual-ip</code> OR <code>--no-vip</code>	UFM HA cluster Virtual IP or configure HA without virtual IP

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

4.2.2 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:

Option		Description
-r	<code>--role <node role></code>	Node role (master or standby).
-e	<code>--peer-primary-ip <ip address></code>	Peer node primary IP address (mandatory).
-l	<code>--local-primary-ip <ip address></code>	Local node primary IP address (mandatory).
-E	<code>--peer-secondary-ip <ip address></code>	Peer node secondary IP address (mandatory).
-L	<code>--local-secondary-ip <ip address></code>	Local node primary IP address (mandatory).
-i	<code>--virtual-ip <virtual-ip></code>	Cluster virtual IP(v4).
	<code>--virtual-ip6 <virtual-ip></code>	Cluster virtual IP(v6).
-p	<code>--hacluster-pwd <pwd></code>	HA cluster user password.
-h	<code>--help</code>	Show this message
-N	<code>--no-vip</code>	Configure HA without virtual IP
-M	<code>--ignore-mgmt-failure</code>	Ignore management interface status if VIP is configured. Will not failover if master node's secondary IP is down.
	<code>--enable-single-link</code>	Enable single network interface mode.

To configure HA, follow the below instructions:



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

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

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

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

```
ufm_ha_cluster config -r master -e <peer primary ip address> -l <local primary ip address> -E <peer secondary ip address> -L <local secondary ip address> -p -i <virtual ip address>
```

4.3 DRBD Configuration

The DRBD is used for syncing File System between the two nodes. DRBD is the default sync method, unless stated otherwise in the configuration file. (see section: "Using File Configuration" below). The DRBD disk that was assigned during installation phase will be mounted as a File System directory, the default option of this mount is "data=ordered", however, it can be override in the configuration file in the "DRBD" section, in order to set the data option to "journal" which offers the highest level of data integrity, but it can impact write performance.

4.4 NFS File Sharing

NFS synchronization mechanism can be used instead of DRBD. Multi-Nodes Support can be used with NFS synchronization mechanism only, as described in the following section. To activate this functionality, users must define the following parameters:

- Mode: NFS
- NFS Server
- Shared Folder

Ensure that the NFS version supports nfs4. It is recommended that the NFS server is not one of the UFM-HA nodes. Refer to the section below for details on configuring the file.

4.5 Multi-Nodes Support

The UFM-HA cluster can comprise of more than two nodes. Among these nodes, one will serve as the master, while the others will operate in standby mode.

To configure multiple nodes, users must populate the configuration file '/etc/ufm_ha/ha_nodes.cfg' on all nodes (ensuring that the file is identical across all nodes).

This file contains details about each participating node, including:

- Role: Master/Standby
- Primary IP address
- Secondary IP address

4.5.1 Using File Configuration

The '/etc/ufm_ha/ha_nodes.cfg' file contains all the necessary information for HA configuration and can serve as a replacement for command-line configuration. The only configuration not saved in the file is the password for security reasons.

To configure, use the following command (should be executed after setting the configuration):

```
ufm_ha_cluster config -p <password>
```



The standby nodes must be configured at first, with the last node being set as the master node.

4.5.2 Configuration File

The sample configuration file includes up to three sections for nodes, but users can add additional sections as needed.

```
[General]
# Connection mode
# in case dual_link is true, each node must have primary and secondary IPs
dual_link = true
# automatic failure cleanup interval (in hours)
# will perform cleanup of failures to enable automatic failover
automatic_failure_cleanup_interval = 24

[Node.1]
# valid role options: master/standby
role = master
# Mandatory
primary_ip =
# Mandatory if dual_link = true
secondary_ip =

[Node.2]
role = standby
primary_ip =
secondary_ip =

[Node.3]
role = standby
primary_ip =
secondary_ip =

# Add other Node.x sections if needed.

[Virtual]
# If virtual IP should not be added, set `no_vip = true`
no_vip =

virtual_ip =

virtual_ip6 =

ignore_mgmt_failure = false
# when using BGP virtual IP, you must use the loopback interface, set `interface = lo`
# in other cases we let the pcs to decide on the relevant network interface.
interface =

[FileSync]
# valid options are: drbd/nfs
mode = nfs

[DRBD]
# fill in case the FileSync.mode is drbd
# drbd data mode. options are: ordered/journal (default is ordered)
# data=journal offers the highest level of data integrity,
# but it can impact write performance.
data = ordered

[NFS]
# fill in case the FileSync.mode is nfs
nfs_server =
shared_folder =
```

4.6 UFM HA Cluster Operations

4.6.1 Show UFM HA version

Run the following command to show UFM HA version:

```
ufm_ha_cluster version
```

4.6.2 Starting UFM HA Cluster



Before starting the UFM cluster, ensure that the DRBD sync is completed.

To start UFM HA cluster:

```
ufm_ha_cluster start
```

4.6.3 Checking UFM Cluster Status

To check UFM HA cluster status:

```
ufm_ha_cluster status
```

4.6.4 Stopping UFM HA Cluster

To stop UFM HA cluster:

```
ufm_ha_cluster stop
```

4.6.5 Takeover Services

The takeover command can be executed on the standby machine so that it will be the master.

```
ufm_ha_cluster takeover
```

4.6.6 Master Failover

The failover command can be executed on the master machine so that it will be the standby.

```
ufm_ha_cluster failover
```

4.6.7 Automatic cleanup of failed actions

When an action failed in one of the HA nodes, for example DRBD failure, service failure or any other HA resources failure, the failed node will no longer be a candidate of automatic failover until these failed actions are cleaned up. To manually cleanup failed action, the user can run the following command:

```
pcs resource cleanup
```

The UFM-HA performs automatic cleanup of failed actions every 24 hours. This period is configurable and can be changed in the `General` section in the `ha_nodes.cfg` configuration file. See section "Configuration File" above.

4.6.8 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
```

- Config the new standby node; please refer to [Configuration](#).
- 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 <cluster_password>
```

4.6.9 Uninstalling UFM HA

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

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

5 Monitoring and Troubleshooting

Check UFM Status	Run the below command on the master node: <pre>systemctl status ufm-enterprise.service</pre>
Check HA Status	Run the below command: <pre>ufm_ha_cluster status pcs status</pre>
Check DRBD Status	Run the below command: <pre>ufm_ha_cluster status</pre>
Show DRBD Resource	Run the below command: <pre>drbdadm sh-resources</pre>
Show DRBD Disk State	Run the below command: <pre>drbdadm dstate ha_data</pre>
Show DRBD Role	Run the below command: <pre>drbdadm role ha_data</pre>
Show DRBD Connectivity	Run the below command: <pre>drbdadm cstate ha_data</pre>

<p>Check UFM Status</p>	<p>Run the below command on the master node:</p> <pre>systemctl status ufm-enterprise.service</pre>
<p>Split-Brain Recovery</p>	<p>For automated HA solution, is it recommended to configure STONITH agents to kill (power-off) a peer node.</p> <p>Step 1: Manually choose a node which 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.</p> <pre>ufm_ha_cluster enable-maintain</pre> <p>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:</p> <pre>victim# ufm_ha_cluster reset standby</pre> <p>Resync starts automatically if the survivor is in a WfConnection network state. If the split-brain survivor is still in a Standalone connection state, reconnect it:</p> <pre>survivor# ufm_ha_cluster reset master</pre> <p>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.</p>
<p>Communication Timeout during HA Configuration</p>	<p>During the configuration phase of high availability, if you encounter errors regarding connectivity, such as 'Error: Unable to communicate with <master/standby IP>' or connection timeouts—even when server connectivity appears fine, consider checking the ybind service, as it may be affecting communication. Stop the ybind service on the master and standby and configure HA. After the configuration succeeds, enable the ybind service again.</p> <pre>systemctl stop ybind # configure HA systemctl start ybind</pre>

6 Standby Node Replacement

UFM Standby Node Replacement automates the process of replacing a failed standby node in a UFM High Availability (HA) cluster, minimizing downtime and reducing the need for manual operations.

This capability streamlines the preparation, configuration, and reintegration of a replacement node into the existing UFM-HA cluster.

The `ufm_replace_standby` CLI tool manages the entire workflow, hiding the underlying complexity and providing a simple, user-friendly interface to prepare, execute, and validate the standby node replacement.

6.1 Objectives

In the event of a standby node failure, this feature enables smooth integration of a replacement standby node into the UFM cluster, ensuring minimal service interruption and minimal manual effort.

6.2 Prerequisites

Before starting the replacement procedure, ensure the following conditions are met:

- The failed standby node has been completely removed from the cluster and disconnected from the network.
- The replacement standby node is physically connected to the cluster network using the correct interfaces – both InfiniBand (IB) and Management links – in a Back-to-Back topology.
- The replacement standby node is running the same UFM appliance version as the master node.
- UFM-HA version 6.0.0-6 (with Seamless RMA support) is installed on both the master and standby nodes.

6.3 Replacement Process

The `ufm_replace_standby` tool automates the replacement of a UFM-HA standby node. It manages the complete workflow – from preparing and validating the new standby node to fully integrating it into the cluster.

To run the full replacement procedure, use:

```
ufm_replace_standby run <Options>
```

Command Options:

Option	Description
<code>--standby-primary-ip=<ip></code>	Primary IP address of the new standby node.

Option	Description
<code>--standby-secondary-ip=<ip></code>	Secondary IP address of the new standby node.
<code>--hacluster-password=<pwd></code>	Optional – Password for the hacluster user. If the <code>--reset_hacluster-password</code> option is also specified, this value will be set as the new hacluster password. If not provided, the user will be prompted to enter the password interactively during the replacement process.
<code>--reset_hacluster-password</code>	Optional – Resets the existing <code>hacluster</code> password. The default value is <code>false</code> .
<code>--preserve-ssh-trust</code>	Optional – Preserves the established SSH trust after the replacement process completes (whether successful or failed). If not set, SSH trust will be removed at the end of a successful run – unless it existed prior to execution.
<code>--enable-single-link</code>	Enable single network interface mode.
<code>--verbose</code>	Optional – Enables extended debug output in the console.

6.4 Run Command – Full Replacement Workflow

The `ufm_replace_standby run` command executes the complete standby node replacement sequence, which includes:

1. Establish SSH Trust
 - Checks if SSH trust is already in place.
 - If not, generates SSH keys and prompts the user to securely copy them to enable passwordless access.
2. Validate the New Standby Node
 - Confirms readiness for integration by:
 - Verifying required network interfaces.
 - Ensuring UFM-HA and HA stack versions match the master node.
 - Checking that `ufm_ha` is installed and supports standby replacement.
 - All validation steps are performed remotely over SSH.
3. Transfer Files
 - Copies `ha_nodes.cfg` from the master node to `/etc/ufm_ha/ha_nodes.cfg` on the standby node to maintain consistent HA configuration.
4. Detach Old Standby
 - Removes the old standby node from HA configuration (PCS and DRBD), even if it is already disconnected.
 - Ensures the cluster is ready to accept the new standby.
 - Performed on the master node.
5. Configure the New Standby Node
 - Prepares the standby node by:
 - Initializing the DRBD disk (if applicable).
 - Setting the `hacluster` user password.
 - Performing required pre-join setup tasks.
 - Executed remotely on the standby node via SSH.

6. Attach the New Standby

- Adds the standby node to the UFM-HA cluster by:
 - Authenticating with the `hacluster` user.
 - Updating PCS configuration.
 - Adjusting DRBD settings to include the new standby.
- Actions performed from the master node.

7. Clean-up

- Removes temporary SSH trust files (e.g., encrypted password, secret file).
- If `--preserve-ssh-trust` is not set, removes SSH trust created during the process.

8. Cluster Validation

- Runs automated checks to confirm:
 - The new standby is successfully integrated into the cluster.
 - HA daemons and UFM services are running.
 - DRBD connectivity and disk state (if applicable) are healthy.

6.5 Resuming After a Failure

If the replacement process fails at any stage, `ufm_replace_standby` can resume from the failure point using its built-in progress tracking.

After fixing the cause of the failure, resume without starting over:

```
ufm_replace_standby run --resume [--hacluster-password=<pwd>]
```

Short form:

```
ufm_replace_standby --resume [--hacluster-password=<pwd>]
```



Note: If the process fails, the `hacluster` password is not retained, even if it was provided in the original run. You must either re-enter it using `--hacluster-password` or provide it interactively.

6.6 Aborting the Previous Run

To abort an incomplete process and start fresh:

```
ufm_replace_standby run --abort
```

Short form:

```
ufm_replace_standby --abort
```

6.7 Cluster Validation

Once `ufm_replace_standby` completes all replacement steps, it automatically runs a series of validation checks to confirm successful integration of the new standby node.

The validation process includes:

- Confirming that the new standby has joined the UFM-HA cluster.
- Verifying that HA stack daemons and UFM services are running.
- Checking disk state and DRBD connectivity (if applicable).

Each validation step is subject to a predefined timeout. For example, the tool will wait up to 10 seconds for DRBD to reach a connected and active state. If a timeout is exceeded, the validation process stops.

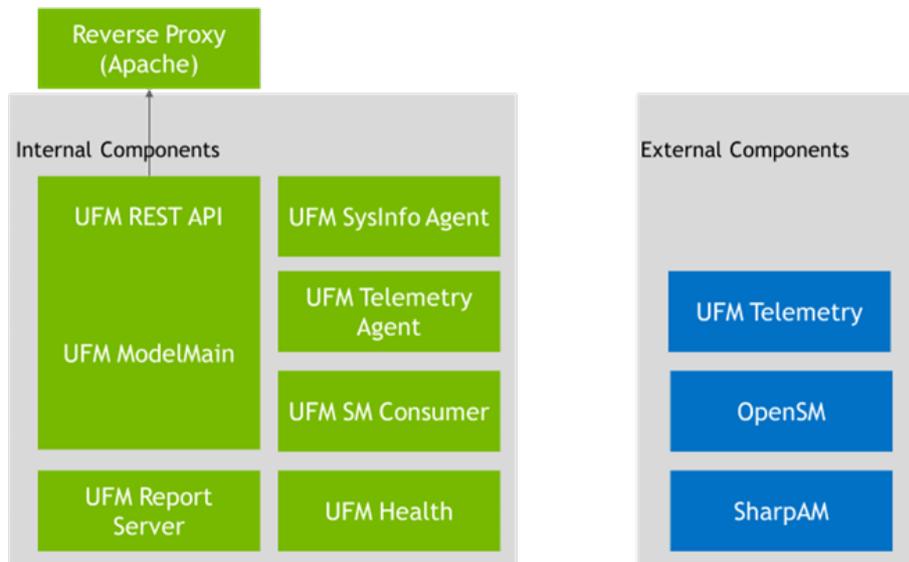


Note: A validation timeout does not necessarily indicate a cluster failure. In such cases, manually verify the cluster status using `ufm_ha_cluster` and other HA monitoring tools to ensure the cluster is functioning properly.

7 UFM High-Level Architecture

The below figure illustrates the UFM high-level architecture.

UFM HIGH LEVEL ARCHITECTURE



7.1 FR#1

Support of Active-Standby HA approach. UFM is not designed to run with multiple instances (active-active mode). There are several constraints:

1. Single SM
2. Single SharpAM
3. Single UFM Telemetry
4. UFM is stateful and manages its internal state (cluster topology model) in RAM

7.2 FR#2

Persistent storage usage is required for the following:

1. Configuration files (UFM, SM, SharpAM, UFM Telemetry, Apache)
2. DB (SQLite) - history telemetry + configuration + app state
3. Operation history - logs, events, alarms

7.3 Solution Options

7.3.1 FR#1

Develop “ufm operator” examples, refer to:

- active-standby-controller

- Implementing leader election for Kubernetes pods
- kubernetes-active-passive
- ActiveStandbySingletonPod

7.3.2 FR2#

1. KVS DB (etcd), Config Maps
2. 3rd party Cache\DB with load-balancing HA built-in (Redis, MongoDB, etc)

8 Document Revision History

Date	Description of Changes
Nov 10, 2025	<ul style="list-style-type: none"> • Changes and New Features • Bug Fixes in This Release • Updated Configure HA without SSH Trust - Added the " <code>--enable-single-link</code>" option • Updated Standby Node Replacement
Sep 5, 2025	Updated: <ul style="list-style-type: none"> • Changes and New Features • Bug Fixes in This Release
Aug 7, 2025	Updated: <ul style="list-style-type: none"> • Changes and New Features • Bug Fixes in This Release Added: <ul style="list-style-type: none"> • Standby Node Replacement
May 5, 2025	Updated: <ul style="list-style-type: none"> • Installation and Configuration • Changes and New Features
Feb 10, 2025	Added Release Notes
Nov 7, 2024	Updated Monitoring and Troubleshooting
Aug 12, 2024	Added a note to Configure HA with SSH Trust
May 7, 2024	<ul style="list-style-type: none"> • Updated HA package installation link across the document • Updated Installation and Configuration
Feb 8, 2024	Updated Installation and Configuration
Dec 12, 2023	Updated the following sections: <ul style="list-style-type: none"> • Prerequisites • Installation and Configuration • Using File Configuration • NFS File Sharing • Using File Configuration
Nov 30, 2023	<ul style="list-style-type: none"> • Updated DRBD • Updated Installation and Configuration
Nov 5, 2023	<ul style="list-style-type: none"> • Updated the UFM HA package link across the document • Added Multi-Nodes Support
Aug 14, 2023	Updated installation command.
May 10, 2023	Updated the following sections: <ul style="list-style-type: none"> • Overview • Prerequisites • Installation and Configuration • Monitoring and Troubleshooting
Feb 6, 2023	First Release

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 NON-INFRINGEMENT, 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 its



affiliates 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

© 2025 NVIDIA Corporation & affiliates. All Rights Reserved.

