



# **IBM Unified Fabric Manager Custom – Sourced From Mellanox (Part Number 8867-FM2) Customization README**

---

**Rev 1.4**

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies  
350 Oakmead Parkway Suite 100  
Sunnyvale, CA 94085  
U.S.A.  
[www.mellanox.com](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2017. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

# Table of Contents

<b>Document Revision History</b> .....	<b>5</b>
<b>Document Introduction</b> .....	<b>6</b>
<b>1 Defaults to Modify Immediately</b> .....	<b>7</b>
<b>2 Standalone UFM Customization</b> .....	<b>8</b>
2.1 Standalone UFM: Single Network Setup.....	8
2.1.1 Linux OS License.....	8
2.1.2 UFM Preparatory Setup.....	8
2.1.3 Routing Customization.....	8
2.1.4 Start UFM.....	8
2.2 Standalone UFM: Multi-Network Setup.....	8
2.2.1 Linux OS License.....	8
2.2.2 UFM Preparatory Setup.....	8
2.2.3 Unique subnet_prefix for each Network .....	8
2.2.4 Routing Customization.....	9
2.2.5 Start UFM.....	9
<b>3 Advanced Customization</b> .....	<b>10</b>
3.1 Miscellaneous Customization Considerations .....	10
3.2 HA UFM Overview .....	10
3.2.1 Flow For Master / Slave Servers .....	10
3.2.2 Master Server specific setup .....	11

# List of Tables

Table 1: Document Revision History ..... 5

## Document Revision History

*Table 1: Document Revision History*

Revision	Date	Description
Rev. 1.4	Match 01, 2017	Initial version of this document.

## Document Introduction

This document describes the typical customization that may need to be performed for the IBM Unified Fabric Manager Custom (PN 8867-FM2). The hardware will be shipped pre-installed with the Linux distribution and the Mellanox Unified Fabric Management software. The UFM User Manual should be referenced to understand the full capabilities of UFM.

### Documentation and Files

The following Mellanox documentation should be available prior to customization.

- Mellanox UFM User Manual
- Mellanox UFM Fabric Best Practice
- Mellanox UFM SDK User Manual
- Mellanox OFED User Manual

# 1 Defaults to Modify Immediately

- Server Password

The server comes with default OS password of “passw0rd”

- Ethernet Connection

The server Ethernet ip address should be modified to your specifications at  
`/etc/sysconfig/network-scripts/ifcfg-eno1`

- InfiniBand HCA IPoIB

The server IPoIB address should be modified to your specifications at  
`/etc/sysconfig/network-scripts/ifcfg-ib0`.

## 2 Standalone UFM Customization

### 2.1 Standalone UFM: Single Network Setup

#### 2.1.1 Linux OS License

Before proceeding with customization ensure your OS distribution is properly licensed and registered.

#### 2.1.2 UFM Preparatory Setup

1. connect the UFM server to the InfiniBand switch
2. copy your UFM license to `/opt/ufm/files/licenses/`

#### 2.1.3 Routing Customization

For a large HPC InfiniBand topology you may want to modify the default MINHOP routing engine. If this is the case then:

1. modify the routing engine stanza in `/opt/ufm/files/conf/opensm/opensm.conf` to reflect the preferred routing engine
2. you may also want to modify the `/opt/ufm/files/conf/opensm/root_guid.conf` to contain the guid of all the spine modules

#### 2.1.4 Start UFM

1. Reload the mlx5 driver modules by executing “service openibd restart”
2. Start UFM by executing “service ufmd start”

### 2.2 Standalone UFM: Multi-Network Setup

If this cluster is configured as a multi-network topology then each network will require a UFM server. For the Standalone Multi-network setup, follow these steps:

#### 2.2.1 Linux OS License

Before proceeding with customization ensure your OS distribution is properly licensed and registered for each UFM server

#### 2.2.2 UFM Preparatory Setup

On each UFM server, follow the steps below:

1. connect the UFM server to the InfiniBand switch
2. copy your UFM license to `/opt/ufm/files/licenses/`

#### 2.2.3 Unique subnet\_prefix for each Network

Since this is a multi-network topology then each UFM server must have a unique `subnet_prefix` for the driver to send messages to the correct network.

- UFM1: Network1  
modify the subnet-prefix in `/opt/ufm/files/conf/opensm/opensm.conf`.



Example:

```
subnet-prefix: 0xfec0000000000001
```

- UFM2: Network2

modify the subnet-prefix in `/opt/ufm/files/conf/opensm/opensm.conf`. Example:

```
subnet-prefix: 0xfec0000000000002
```

## 2.2.4 Routing Customization

For a large HPC InfiniBand topology you may want to modify the default MINHOP routing engine. If this is the case then:

1. modify the routing engine stanza in `/opt/ufm/files/conf/opensm/opensm.conf` to reflect the preferred routing engine
2. you may also want to modify the `/opt/ufm/files/conf/opensm/root_guid.conf` to contain the guid of all the spine modules

## 2.2.5 Start UFM

On each UFM server, start the UFM service

1. Reload the `mlx5` driver modules by executing “`service openibd restart`”
2. Start UFM by executing “`service ufmd start`”

## 3 Advanced Customization

For advanced customizations, you may want to consider using the services of a qualified InfiniBand service provider.

### 3.1 Miscellaneous Customization Considerations

There are additional customizations that a customer may wish to investigate to comply with an overall site IT infrastructure policy. These are outside the scope of this quick start setup and it is recommended that the UFM User Manual be referenced. Some examples:

- QoS (quality of service)
- Event thresholds/severity tuning
- Remote event logging
- Fabric monitoring

### 3.2 HA UFM Overview

The HA UFM customization requires additional setup beyond the basic standalone customization. As such, it is best to follow the instructions in the UFM User Manual. There are considerations, such as, creating an IP alias, defining identical DRBD partitions on each server, special command syntax required for installing and configuring the HA UFM, and understanding the proper procedure for starting, stopping, and/or failing over the UFM.

The following is a high-level flow as an introduction only. Please follow the instructions in the Mellanox UFM User Manual.

#### 3.2.1 Flow For Master / Slave Servers

- The Master & Slave server HCA need to be connected to different physical switches in the Infiniband network to avoid the possibility of a switch component failure impacting the UFM HA.
- Identify an IP address to be used as the "common" address for launching the UFM GUI. This IP address will be configured by UFM HA as a VIP address. This "common" address should be used when connecting to the UFM server since it will only be configured by UFM for the active Master UFM server.
- An identical disk partition must be created on each server. The partition must not have a filesystem on it and it must not be mounted. If you plan to use a partition created during install time then unmount it, remove the mount point, and remove the entry in `/etc/fstab`. However the partition is created, it must be zeroed out before starting the ufm installation. Refer to the Mellanox UFM User Manual for recommended partition sizes.
- Since the drbd partition will contain the contents of what is traditionally in `/opt/ufm` (including log files), it's important that the clocks on the two servers are synchronized with a common clock source.
- The UFM HA installation is launched from what will be the master server. To accomplish the installation, as well as, subsequent UFM HA operation, a passwordless ssh key must be created for each server and then installed on the other server.

- Refer to subnet\_prefix and routing discussion in the Standalone section of document after the Master and Slave have been installed.

### 3.2.2 Master Server specific setup

IMPORTANT: the following applies only to the server that is initially designated as the master.

- The UFM package is located in /UFM directory
- Copy the license to the designated master server
- unzip the UFM package and gunzip the respective ufm linux distribution (.tgz)

You're now ready to install UFM on the master. Installation of the stand-by is controlled and performed by the master server!!

Example:

```
./install.sh -o ib -m 192.1.1.30:/dev/master_drbd -s  
192.1.1.31:/dev/slave_drbd -c 192.1.1.32 -n 255.255.255.0 -b 192.1.1.255
```

where:

- -m = master IP address for accessing the server the drbd partition
- -s = standby IP address for accessing the server the drbd partition
- -c = the common IP address (VIP) which will be used when launching the UFM GUI and when accessing the current Master
- -n = the network netmask for the IP subnet
- -b = the broadcast address