IBM Systems

IBM BladeCenter Open Fabric Manager Installation and User's Guide

Version 3.1

IBM Systems

IBM BladeCenter Open Fabric Manager Installation and User's Guide

Version 3.1

Note

Before using this information and the product it supports, read the information in "Notices & Trademarks" on page 59.

This edition applies to version 3.1 of IBM BladeCenter Open Fabric Manager and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables v

Who should read this user's guide. . . vii

Conventions and terminology ix

Chapter 1. IBM BladeCenter Open Fabric Manager V3.1 1

Chapter 2. Overview of BladeCenter Open Fabric Manager 3

Accessibility features for BladeCenter Open Fabric Manager 3
License information 3
Hardware and software requirements 3
  Hardware requirements. 3
  Supported hardware 4
  Supported software 7
BladeCenter Open Fabric Manager components. . . 7
  Configuration file. 7
  AMM Web interface 13

Chapter 3. Planning for BladeCenter Open Fabric Manager 15

Preparing for OFM 15
  Upgrading firmware 15
  Setting up boot from SAN 19

Chapter 4. Installing BladeCenter Open Fabric Manager 21

Installing OFM 21
Installing OFM-Advanced upgrade 21
  Installing the IBM Systems Director plug-in version of OFM-Advanced upgrade for Windows 21
  Installing the IBM Systems Director plug-in version of OFM-Advanced upgrade for Linux . . 21
  Installing the OFM-Advanced upgrade stand-alone version for Windows 22
  Installing the OFM-Advanced upgrade stand-alone version for Linux 22
Migrating OFM-Advanced upgrade 23

Chapter 5. Configuring BladeCenter Open Fabric Manager 25

Creating a configuration file automatically 25
  Selecting domains 28
  Avoiding address duplication 28
Creating a requirements report 29
  Creating a requirements report from "The Configuration File Has Been Created" page in the AMM Web interface 29
  Creating a requirements report from the main Open Fabric Manager Configuration Management page on the AMM Web interface. . 29
Editing the configuration file manually 30
Applying a new configuration 30
Viewing the configuration in a local chassis . . . 31
Retrieving the current configuration 32

Chapter 6. Using OFM and OFM-Advanced upgrade 33

Using OFM 33
  Initial deployment 33
  Adding a new chassis to the domain 33
  Replacing a blade in the same slot. 34
  Swapping addresses between blades 34
  Replacing AMM IP addresses 34
  Replacing the AMM in a single AMM environment 34
Using OFM-Advanced upgrade. 35
  Creating a blade address manager configuration template 35
  Applying a blade address manager configuration template 39
  Creating a standby blade pool configuration template 39
  Deploying a standby blade pool configuration template manually 40
  Creating an event action plan 41

Chapter 7. Troubleshooting and support 43

Problems when you create or apply a standby blade pool 43
Incorrect OFM address usage 44
Configuration failure scenarios 44
Error messages 46

Chapter 8. Reference 49

Accessibility features for BladeCenter Open Fabric Manager 49
Session and credentials 49
Standby AMM 49
CLI command 50
Events 50
User interface. 51
  BladeCenter Open Fabric Manager summary . . 51
  Blade Address Manager Configuration wizard. . 52
  BladeCenter Standby Blade Pool Configuration wizard 56

Chapter 9. Publications and related information. 57

Notices & Trademarks 59

Notices 59
Trademarks 60

Glossary 63

Terms 63

Index 65

Tables

1. Instllatoin requirements for OFM-Advanced upgrade 3
2. Runtime requirements for OFM-Advanced upgrade stand-alone version 4
3. Expansion cards that support OFM 4
4. BladeCenter chassis' that support OFM. 5
5. BladeCenter servers that support OFM 5
6. Fibre channels switches that support OFM 5
7. Ethernet switches that support OFM 6
8. Pass-thru modules and switches that support OFM 7
9. Supported software - OFM-Advanced upgrade 7
10. 13
11. BladeScanner and ChassisUpdate log file locations 17
12. Blade address configuration creation and application problems 43
13. Standby blade pool creation and application problems 43
14. Parsing errors 47

Who should read this user's guide

This user's guide is for system administrators and operators using OFM from IBM® Director or the Advanced Management Module (AMM) to replace and recover blades in their environment.

Conventions and terminology

These notices are designed to highlight key information:

Note: These notices provide important tips, guidance or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.

CAUTION:
These notices indicate possible damage to programs, devices or data. An attention notice appears before the instruction or situation in which damage can occur.

Chapter 1. IBM BladeCenter Open Fabric Manager V3.1

IBM BladeCenter® Open Fabric Manager (OFM) is a solution that enables you to quickly replace and recover blades in your environment.

It does this by assigning Ethernet MAC, Fibre Channel WWN and SAS WWN addresses to the BladeCenter slots in such a way that any blades plugged into those slots will take on the assigned addresses. This enables the Ethernet and Fibre Channel infrastructure to be configured once and before any blades are connected to the BladeCenter chassis. BladeCenter Open Fabric Manager-Advanced upgrade adds capabilities to monitor blades for failure events and to take automatic action to failover from a faulty blade to a cold standby blade. The Ethernet MAC and Fibre Channel WWN addresses are moved from the faulty blade's slot to the standby blade's slot and the standby blade is automatically powered on. In a boot from SAN environment, the operating system and software that was once running on the faulty blade is now running on the standby blade.

There are two separate offerings of BladeCenter Open Fabric Manager. The main BladeCenter Open Fabric Manager function is provided in the OFM offering. Additional capabilities are available with the OFM-Advanced upgrade offering.

OFM

With BladeCenter Open Fabric Manager, you can pre-assign MAC and WWN addresses, as well as storage boot targets, for up to 100 chassis or 1400 blade servers. Using the management module Web interface, you can create addresses for blade servers, save the addresses to a configuration file, deploy the addresses to the blade slots in the same chassis or in up to 100 different chassis. This can be done without any blade servers installed in the chassis.

OFM-Advanced upgrade

With BladeCenter Open Fabric Manager-Advanced upgrade, you can monitor the health of blade servers and automatically - without user intervention - replace a failed blade from a designated pool of spare blades. After receiving a failure alert, OFM-Advanced upgrade attempts to power off the failing blade, read the BladeCenter Open Fabric Manager virtualized addresses and boot target parameters, apply these parameters to the next blade in the standby blade pool, and power on the standby blade.

You can also pre-assign MAC and WWN addresses, as well as storage boot targets, for up to 100 chassis or 1400 blade servers with BladeCenter Open Fabric Manager-Advanced upgrade. Using an enhanced graphical user interface, you can create addresses for blade servers, save the addresses profiles; deploy the addresses to the blade slots in the same chassis or in up to 100 different chassis. This can be done without any blade servers installed in the chassis. Additionally, you can create profiles for chassis that have not been installed in the environment by simply associating an IP address to the future chassis.

BladeCenter Open Fabric Manager-Advanced upgrade is available as a stand-alone offering or as an extension to IBM Systems Director. The stand-alone version includes an embedded version of IBM Systems Director.

Important: OFM is a prerequisite of OFM-Advanced upgrade

Chapter 2. Overview of BladeCenter Open Fabric Manager

This section provides an overview of BladeCenter Open Fabric Manager, including license information, supported hardware and software requirements, and a technical overview.

Accessibility features for BladeCenter Open Fabric Manager

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The BladeCenter Open Fabric Manager plug-in for IBM Systems Director supports the accessibility features that are supported in IBM Systems Director.

License information

When you download, install, and begin using BladeCenter Open Fabric Manager-Advanced upgrade plug-in for IBM Systems Director, you are granted a 60-day evaluation license. The stand-alone version comes with a 360-day evaluation license. Once the evaluation license expires, you must purchase a license in order to continue using BladeCenter Open Fabric Manager-Advanced upgrade.

The number of days remaining on the evaluation license is displayed on the BladeCenter Open Fabric Manager-Advanced upgrade summary page. When the evaluation license expires, you must install a permanent license key to continue using BladeCenter Open Fabric Manager-Advanced upgrade.

Hardware and software requirements

There are minimum hardware and software requirements the system must meet before you can install or use OFM and OFM-Advanced upgrade.

Hardware requirements

Review the processor, disk space, and memory requirements for the installation and use of OFM-Advanced upgrade, both stand-alone and plug-in versions.

Installation requirements

Stand-alone and plug-in

Table 1. Instllatoin requirements for OFM-Advanced upgrade

Operating system	Processor	Disk space	Memory
Linux®	1 Processor, 3 GHz Intel® Xeon®	6 GB[1]	1 GB
Windows®	1 Processor, 3 GHz Intel Xeon	6 GB[2]	1 GB

Note:

1. The stand-alone version for Linux requires at least 6 GB of free disk space in `/opt`, 3 GB of free disk space in `/tmp`, and at least 1.5 GB of RAM. The Director plug-in version for Linux requires 118, 167, 294 bytes of disk space on the installation target.

2. The stand-alone version for Windows requires at least 6 GB of free disk space on the target installation directory, 3 GB of free disk space in the system temp directory, and at least 1.5 GB of RAM. The Director plug-in version for Windows requires 118, 167, 294 bytes of disk space on the installation target.

Runtime requirements

Stand-alone

Table 2. Runtime requirements for OFM-Advanced upgrade stand-alone version

Disk space	Memory
3 GB temporary space, 2.5 GB on target	>1.5 GB

Plug-in

The hardware requirements for OFM-Advanced upgrade plug-in version are the same as those for IBM Systems Director 6.2.x. See Hardware requirements for IBM Systems Director (publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.plan.helps.doc/fqm0_r_hardware_requirements.html)

Note: OFM-Advanced upgrade 3.1 requires IBM Systems Director 6.2.x.

Supported hardware

The following hardware supports OFM and OFM-Advanced upgrade.

Expansion cards

Table 3. Expansion cards that support OFM

Expansion card	Firmware version	
	HS/LS blades	JS blades
QLogic 4Gb SFF Fibre Channel Expansion Card (26R0890)	Muti-boot v1.43 or later	Muti-boot v1.46 or later
QLogic 4Gb Fibre Channel Expansion Card (CFFv) for IBM BladeCenter (41Y8527)	1.43 or later	v1.46 or later
QLogic Ethernet and 4Gb Fibre Channel Expansion Card (CFFh) for IBM BladeCenter (39Y9306)	1.43 or later	v1.46 or later
QLogic 2-port 10Gb Converged Network Adapter (CFFh) for IBM BladeCenter		
Emulex 4Gb SFF Fibre Channel Expansion Card (39Y9186)	6.02a2 or later	6.02a7 or later
Emulex 4Gb Fibre Channel Expansion Card (CFFv) for IBM BladeCenter (43W6859)	6.02a2 or later	6.02a7 or later
IBM SFF Gb Ethernet Expansion Card (39R8624)	Any version	Any version
Ethernet Expansion Card (CFFv) for IBM BladeCenter (39Y9310)	Any version	Any version
NetXen 10 Gb Ethernet Expansion Card for IBM BladeCenter (39Y9271)	3.4.223	3.4.223

Note: JS22 cannot host NetXen 10 Gb Ethernet Expansion Card for IBM BladeCenter (39Y9271).

BladeCenter chassis'

Table 4. BladeCenter chassis' that support OFM

BladeCenter chassis	Firmware version	
	HS/LS blades	JS blades
BladeCenter E (8677)	AMM 4.0 (BPET42C) or later	AMM 4.0 (BPET42D) or later
BladeCenter H (8852)	AMM 4.0 (BPET42C) or later	AMM 4.0 (BPET42D) or later
BladeCenter T (8720/8730)	AMM 4.0 (BBET42C) or later	AMM 4.0 (BBET42C) or later
BladeCenter HT (8740/8750)	AMM 4.0 (BPET42C) or later	AMM 4.0 (BPET42D) or later
BladeCenter S (8886)	AMM 4.0 (BPET42C) or later	AMM 4.0 (BPET42D) or later

BladeCenter servers

Table 5. BladeCenter servers that support OFM

BladeCenter server	BIOS version	BMC version	NIC version
IBM BladeCenter HS21 (8853)	BCE128A or later	BCBT42BUS or later	See "Expansion cards" on page 4.
IBM BladeCenter HS21 XM (7995)	MJE119A or later	MJBT18AUS or later	See "Expansion cards" on page 4.
IBM BladeCenter LS21 (7971)	BAE139AUS or later	BABT42AUS or later	See "Expansion cards" on page 4.
IBM BladeCenter LS41 (7972)	BAE139AUS or later	BABT42AUS or later	See "Expansion cards" on page 4.
IBM BladeCenter JS22/JS12 (7998)	eFW 3.3 or later	eFW 3.3 or later	See "Expansion cards" on page 4.
IBM BladeCenter HS12 (8014)	N1E125A	N1BT07AUS	See "Expansion cards" on page 4.
IBM BladeCenter HS12 (8028)	N1E125A	N1BT07AUS	See "Expansion cards" on page 4.
IBM BladeCenter HS22 (7870) [1]	latest version	latest version	See "Expansion cards" on page 4.

[1] Does not include support for SAS. Support for SAS will be added in future release levels.

Fibre channels

Table 6. Fibre channels switches that support OFM

Fibre channel switch	Firmware version
Brocade Entry SAN Switch Module for IBM BladeCenter (26K5601)	Any version
Brocade Enterprise SAN Switch Module for IBM BladeCenter (90P0165)	Any version
Brocade 10-port SAN Switch Module for IBM eServer™ BladeCenter (32R1813)	Any version
Brocade 20-port SAN Switch Module for IBM eServer BladeCenter (32R1812)	Any version
Qlogic 6pt Fibre Channel Switch Module (26K6477)	Any version
QLogic 10-port Fibre Channel Switch Module for IBM eServer BladeCenter (32R1904)	Any version

Table 6. Fibre channels switches that support OFM (continued)

Fibre channel switch	Firmware version
QLogic 20-port Fibre Channel Switch Module for IBM eServer BladeCenter (26R0881)	Any version
QLogic 4Gb FC 10-port Full Fabric Switch (43W6724)	Any version
QLogic 4Gb FC 20-port Full Fabric Switch (43W6725)	Any version
QLogic 4Gb FC 20-port Pass-thru Switch (43W6723)	Any version
Cisco Systems 4Gb 10 port Fibre Channel Switch Module for IBM BladeCenter (39Y9284)	Any version
Cisco Systems 4Gb 20 port Fibre Channel Switch Module for IBM BladeCenter (39Y9280)	Any version

Ethernet switches

Table 7. Ethernet switches that support OFM

Network	Firmware version
BNT 1/10Gb Uplink Ethernet Switch Module for IBM	Any version
BNT 10-port 10Gb Ethernet Switch Module	Any version
Cisco Catalyst Switch Module 3012 for IBM BladeCenter (43W4395)	Any version
Cisco Catalyst Switch Module 3110G and 3110X for IBM BladeCenter (41Y8523 & 41Y8522)	Any version
Cisco Systems Gigabit Ethernet Switch Module for eServer BladeCenter (13N2281)	Any version
Cisco Systems Intelligent Gigabit Fiber Ethernet Switch Module for IBM eServer BladeCenter (26K6547)	Any version
Cisco Systems Intelligent Gigabit Ethernet Switch Module for IBM BladeCenter (32R1892)	Any version
Cisco Systems Fiber Intelligent Gigabit Ethernet Switch Module for IBM BladeCenter (32R1888)	Any version
Nortel Networks L2/L3 Copper Gigabit Ethernet Switch Module for IBM eServer BladeCenter (26K6530)	Any version
Nortel Networks L2/L3 Fiber Gigabit Ethernet Switch Module for IBM eServer BladeCenter (26K6531)	Any version
Nortel Layer 2/3 Copper Ethernet Switch Module for IBM BladeCenter (32R1860)	Any version
Nortel Layer 2/3 Fiber Ethernet Switch Module for IBM BladeCenter (32R1861)	Any version
Nortel Layer 2-7 Gigabit Ethernet Switch Module for IBM BladeCenter (32R1859)	Any version
Nortel Networks (TM) Layer 2 - 7 Gigabit Ethernet Switch Module for IBM eServer BladeCenter (73P9057)	Any version
Nortel 10 Gb Ethernet Switch Module for IBM eServer BladeCenter (32R1783)	Any version
IBM Server Switch Module (39Y9324)	Any version

Pass-thru modules and switches

Table 8. Pass-thru modules and switches that support OFM

Passthru	Firmware version
IBM BladeCenter Optical Pass-thru Module (02R9080)	Any version
IBM BladeCenter Optical Pass-thru Module (39Y9316)	Any version
Copper Pass-thru Module Option (39Y9320)	Any version
Copper Pass-thru Module Option (73P6100)	Any version
InfiniBand Passa-thru Module for IBM BladeCenter (43W4419)	Any version
Intelligent Copper Pass-Thru Module for IBM BladeCenter (44W4483)	Any version
QLogic 4Gb FC 20-port Pass-thru Switch (43W6723)	Any version

Supported software

OFM-Advanced upgrade is supported on selected Microsoft® Windows and Linux operating systems on x86 architecture. OFM is supported on any operating system supported by the blade.

Table 9. Supported software - OFM-Advanced upgrade

Operating System
Microsoft Virtual Server 2005
Microsoft Windows 2003 (SP1, SP2, R2)
Microsoft Windows 2000 (Up to SP4)
RHEL 3 (32-bit only; up to SP9)
RHEL 4 (up to SP6)
RHEL 5 with Xen 3.0 (up to SP1)
SLES 9 (up to SP4)
SLES 10 with Xen 3.0 (up to SP1)

BladeCenter Open Fabric Manager components

The OFM configuration file and AMM Web interface are essential for OFM functionality.

Configuration file

The OFM configuration file is the central tool for managing the OFM domain and contains the definitions that you need for a domain of up to 100 BladeCenters. You can generate it automatically, save it and edit it to conform to the needs of a specific domain and then apply it to the domain. You also have the option of creating your own configuration file.

The configuration file is a Comma Separated Value (CSV) file. Each non-blank and non-comment line defines a single entity within a domain of BladeCenters. The entities currently defined are **BladeCenters**, **Slots** and **Port Entries**, where a port is a single network device within a slot, which can have multiple definitions – one for each interface type.

The file is organized hierarchically by BladeCenters, Slots and Ports, with ample comment lines included to act as a guide to editing the file if needed. IBM recommends that you maintain the original structure as much as possible in order to retain the readability of the file. For certain purposes it might be appropriate to extract a smaller section of the domain into a new file so that you can update a particular BladeCenter or a particular blade individually.

To view an example of a configuration file, see “Example: Configuration file” on page 12.

Comments section

There are two types of comments: **line comments** and **field comments**.

```
// this is a line comment
localhost/# this is a field comment    ,bladecenter   ,apply
```

Line comments start with two slashes (//). The system ignores anything between this symbol and the end of the line. You can insert line comments anywhere in the file. However, inserting a comment inside a pair of quote marks results in a `No closing quote` error.

You can insert field comments in any field. Field comments start with a slash-hash (/#). The system ignores anything between this symbol and the end of the field. Inserting a field comment inside a pair of quotes does not cause an error.

Comments are included in the maximum line-length (currently 512 bytes including newlines). Very long comments should be broken over several lines to improve readability and to prevent exceeding the line-length limit. If you exceed the line-length limit, the system issues the following error: `Input line is too long`.

BladeCenter section

```
//BladeCenterIP       ,Type,        ,Mode
  bladecenter2.ibm.com  ,BladeCenter   ,apply
```

The BladeCenter section contains three parameters: the **IP address of the BladeCenter**, the entity **Type** (BladeCenter) and the **Mode**.

IP Address (required)
The IP Address can be any valid BladeCenter address in one of two formats:

- IPv4 dotted notation (`192.168.0.1`)
- Human-readable Internet addresses (`bladecenter2.ibm.com`)

  Note: Fully-qualified hostname is required for this format as indicated above (.ibm.com is required).

For larger sites, IBM strongly recommends using human-readable addresses only if the domain-name-server (DNS) is on the local network. If the DNS is not local, the lookup time can slow the parsing substantially, especially if there is an error and the name is not found.

Type (required)
The type is always `BladeCenter`. Any variation of upper and lowercase is acceptable.

Mode (required)
The mode is one of two options: `apply` or `ignore`. When ignore is selected, the system discards all slots and ports belonging to that BladeCenter. This allows

an entire BladeCenter to be commented-out without the need to modify each individual line, and without regard for the ordering of the file.

The BladeCenter section should always come before slots belonging to it, and only one BladeCenter section can exist for a particular BladeCenter. If no BladeCenter section exists, when a slot is defined, the system uses a reasonable default definition for the BladeCenter. The default definition is based on the IP address of the slot and its mode is `apply`.

You can define up to 100 BladeCenters with their blades and ports in a single file.

Slot section

```
//BladeCenterIP  ,Type   ,Slot,  ,Mode    ,Profile
  localhost      ,slot   ,1    ,enable   ,"TempProfile BC-1Slot-1"
```

The slot section represents a single slot within a BladeCenter. Its association to the BladeCenter is through the IP address of the BladeCenter. A slot entry is required before any port entries can be defined for that slot. Failure to define a slot before using it for a port results in an error. You can only define a slot once. Multiple definitions result in errors.

The slot section consists of five parameters: **IP Address**, **Type**, **Slot**, **Mode** and **Profile**.

IP Address (required)
The IP Address can be any valid BladeCenter address in one of two formats: IPv4 or human-readable internet addresses. If you have not previously defined a BladeCenter with that address, the system defines one using reasonable defaults.

Type (required)
Always contains the value `slot`. Any combination of upper and lowercase is acceptable.

Slot (required)
Indicates which BladeCenter slot you are referring to. This is a numeric value from 1 to 14. For a given type of BladeCenter the actual number of slots may be less than 14.

Mode (optional)
Can be one of three values: `enable`, `disable` or `ignore`.

- `enable` means that the AMM will push the OFM configuration to the blade.
- `disable` means that the AMM will clear the OFM configuration on the blade, so the blade goes to factory addresses.
- `ignore` means that the OFM configuration of this slot is unchanged by the configuration file.

The default value is `disable`.

Profile (optional)
This is a string value of up to 31 characters. You can use it to attach a human-friendly string to a particular blade. When you generate the configuration file, the system creates a value based on the cardinal position of the BladeCenter in the file and the slot number. You can edit this value at will, but care should be taken to stay within the 31 character limit. If you exceed this limit, the system truncates the string and issues a warning. If no profile is given, the system creates one based on the IP address.

Restriction: Quotes (") are not allowed. Commas can only be included if the string is quoted.

Port section

There are currently three types of port entries that the system understands: **Ethernet**, **FibreChannel** and **FibreChannel Targets**. The IP Address, Slot and Type parameters are common to all port types.

IP Address (required)
The IP Address can be any valid BladeCenter address in one of two formats: IPv4 or human-readable internet addresses. If no BladeCenter has previously been defined with that address, the system issues an `Attempt to use a blade slot that has not yet been defined` error.

Slot (required)
Which BladeCenter slot you are referring to. This is a numeric value from 1 to 14. For a given type of BladeCenter the actual number of slots may be less than 14. You can only define a port for a slot that has already been defined with a slot entry (see "Slot section" on page 9. Attempting to define a port for an undefined slot will result in an error.

Offset (optional)
This is a value between 0 and 3. For single-slot blades this value will always be 0. See "Mapping of devices to ports" on page 12 and "Multi-slot blades and the port offset parameter" on page 13 for more information.

Type (required)
This parameter defines which of the port types to use, and which fields are required. Any combination of upper and lowercase is acceptable. It can contain one of these values `Eth` for Ethernet; `FC` for FibreChannel; and `FCTarget` for FibreChannel Target.

- **Ethernet (Eth)**: In addition to the four common parameters, the Ethernet port entry also contains the following parameters.

  Port (required)
  The port to which the data is written. This is a value between 1 and 8, where 1 and 2 are reserved for the built-in on-board Ethernet cards. See "Mapping of devices to ports" on page 12 and "Multi-slot blades and the port offset parameter" on page 13for more information.

  MAC1 (required)
  This is the primary MAC address that is written to the Ethernet card attached to the port in question. It is a 48-bit EUI value represented in the field as six hexadecimal bytes (using values 0-9 and A to F, and not preceded by `0x`) separated by colons, for example, `12:34:56:78:90:AB`. Invalid addresses cause an error and the system ignores the line.

  VLAN1 (optional)
  This is the VLAN you use for this Ethernet connection. An empty field is equivalent to a value of zero. and the system informs the NIC that no VLAN was selected. Valid values for VLANs are 0 to 40. This field defaults to zero.

  This VLAN tag is used only by the BIOS for the PXE boot, if used. You must apply OS VLAN tags at the OS level.

  MAC2 (optional)
  This is the secondary MAC address for those Ethernet cards that support this option. If you do not supply a value for this field it is not applied.

VLAN2 (optional)
This is the secondary VLAN that the system uses for those cards that support a secondary MAC address. This field is applied only if MAC2 and VLAN2 contain a valid value. Valid values for VLANs are 0 to 40. A value of zero is equivalent to an empty field.

This VLAN tag is used only by the BIOS for the PXE boot, if used. You must apply OS VLAN tags at the OS level.

Here is an example of the Ethernet entry parameters:

```
//IP     ,Type  ,Slot  ,Offset ,Port  ,MAC_1              ,VLAN1 ,MAC_2 ,VLAN2
localhost  ,eth   ,1     ,0    ,1     ,25:00:c9:00:00:00
```

Note: In the generated file `Eth` is expanded to `Ethernet`, but this is not required.

- **FibreChannel (FC)**: In addition to the common parameters, the FibreChannel port entry also contains the following parameters.

Port (required)
The port to which the data is written. This is a value between 3 and 8 (1 and 2 are reserved for the built-in on-board Ethernet cards). See "Mapping of devices to ports" on page 12 and "Multi-slot blades and the port offset parameter" on page 13 for more information.

WWNN (optional)
This is the worldwide node name for the FibreChannel device attached to the given port. It is a 64-bit EUI value represented in the field as eight hexadecimal bytes (using values 0-9 and A to F, and not preceded by `0x`) separated by colons. Not all applications require this value, and some interface cards supply this value themselves by creating a number based on a transformation of the WWPN.

WWPN (required)
This is the worldwide port name for the FibreChannel device attached to the given port. It is a 64-bit EUI value represented in the field as eight hexadecimal bytes (using values 0-9 and A to F, and not preceded by `0x`) separated by colons. Invalid addresses result in an error and the line is ignored.

Boot-order (optional)
This defines which target the interface uses during the boot process. Valid values are `none`, `first`, `second`, or `both`. If the value is `first`, `second` or `both`, the boot process tries to use the equivalent target to boot the blade (targets can have a priority of first or second). This parameter defaults to `none`.

Here is an example of the FibreChannel entry parameters:

```
//IP      ,Type  ,Slot  ,Offset ,Port  ,WWPN,                    ,Boot-order
localhost ,fc   ,1     ,0     ,3    ,2f:fc:00:00:c9:00:00:00  ,none
```

- **FibreChannel Target (FCTarget)**: In addition to the common parameters, the FibreChannel Target port entry also contains the following parameters.

Priority (required)
The value of this parameter can be `first` or `second`. `first` denotes the primary target for the blade and `second` denotes the secondary target.

WWN (required)
This is the worldwide node name of the target.

LUN (required)
This is the LUN of the target. For numbers less than four bytes long this can be specified as a decimal or hexadecimal number, where hexadecimal numbers are preceded by 0x. For longer numbers, you must use the standard EUI notation (eight pairs of hexadecimal characters, divided by colons).

- Here is an example of the FibreChannel Target entry parameters:

```
//IP      ,Type       ,Slot  ,Priority  ,WWN                      ,LUN
localhost ,fctarget   ,1     ,first     ,00:00:00:00:00:00:00:00  ,0
```

Other format features of the configuration file

The configuration file also contains the following characteristics and requirements:

- **Case:** Characters in the configuration file are not case sensitive.
- **Whitespace:** The file is largely whitespace agnostic. In general, whitespace is stripped before the fields are parsed. To make editing easier for you when you choose not to edit in a spreadsheet program, whitespace is added to the end of fields in the generated files. This whitespace is entirely optional and you can remove it at your discretion.
- **Newlines:** The file supports spreadsheets that use the UNIX® line-feed only convention (OpenOffice Calc) and the DOS carriage-return/line-feed convention (Excel). It also supports line-feed only and carriage-return/line-feed text editors.
- **Line-length:** The maximum line length is 512 characters. This in the absolute length of the line. It includes comments, whitespace carriage-returns, line feeds and other hidden characters. Exceeding this length results in a line error and the line is discarded.

Example: Configuration file

This topic contains an example configuration file.

```
// GENERATED FILE STARTS

// Blade Center 192.168.0.1
//IP         ,Type (Center) ,Mode
192.168.0.1  ,bladecenter    ,apply

//IP        ,Type (Slot) ,Slot  ,Mode    ,Profile
192.168.0.1  ,slot        ,1     ,enable  ,"TempProfile BC-2Slot-1"

//IP        ,Type  ,Slot  ,Port ,MAC_1            ,VLAN1 ,MAC_2  ,VLAN2
192.168.0.1  ,eth  ,1    ,1    ,25:00:c9:00:00:70  ,0
192.168.0.1  ,eth  ,1    ,2    ,25:00:c9:00:00:71  ,0
192.168.0.1  ,eth  ,1    ,3    ,25:00:c9:00:00:72  ,0
192.168.0.1  ,eth  ,1    ,4    ,25:00:c9:00:00:73  ,0

//IP      ,Type  ,Slot  ,Port  ,WWPN,                    ,Boot-order
localhost ,fc   ,1      ,3    ,2f:fc:00:00:c9:00:00:00  ,none

//IP       ,Type      ,Slot  ,Priority  ,WWN                        ,LUN
localhost  ,fctarget  ,1     ,first     ,00:00:00:00:00:00:00:00  ,0
```

Mapping of devices to ports

You can map devices to ports.

The mapping of ports to the devices on the blade is as follows:

- Ports 1 and 2 are reserved for the on-board Ethernet devices.
- Ports 3 and 4 are reserved for standard expansion cards.
- Ports 5 to 8 are reserved for high-speed expansion cards.

The mapping between the OFM ports and the switch numbering on the chassis is dependant on the chassis. For example, the BCS chassis routes both on-board Ethernet devices to I/O Module bay 1.

The following table defines the mapping of the OFM ports to the switch numbering on the chassis.

Table 10.

	Chassis IO/M numbering				
OFM port	BC1	BCT	BCH	BCHT	BCS
1	1		1	1	1
2	2		2	2	1
3	3		3	3	3
4	4		4	4	4
5	n/a	n/a	7	7	2
6	n/a	n/a	8	8	n/a
7	n/a	n/a	9	9	2
8	n/a	n/a	10	10	n/a

Multi-slot blades and the port offset parameter

Some blades fill more than a single slot in the chassis. As a result, they can access more ports than a single-slot blade.

The maximum number of ports available to a single slot is 8. The maximum number of ports available to a single blade is 32 (a blade can fill four slots, which is the maximum number of slots any Blade can occupy at this time).

Ports beyond the eight ports of the first blade are referred to by using the port **Offset** parameter. `Port 1 Offset 0` refers to the first built-in Ethernet port of a single or multi-slot blade. `Port 1 Offset 1` refers to the third built-in Ethernet port of a blade that is double-width or more.

```
//IP          ,Type  ,Slot  ,Offset  ,Port ,MAC_1              ,VLAN1
localhost     ,eth   ,1     ,0       ,1    ,25:00:c9:00:00:00   ,1
localhost     ,eth   ,1     ,0       ,2    ,25:00:c9:00:00:01   ,2
localhost     ,eth   ,1     ,1       ,1    ,25:00:c9:00:00:00   ,1
localhost     ,eth   ,1     ,1       ,2    ,25:00:c9:00:00:01   ,2
```

The first two ports of each offset are reserved for the Ethernet attached card or built-in Ethernet (at offset = 0). Attempting to apply a FibreChannel (FC) port specification to `Port 1, Offset 2` causes an error. The parser reminds you that the port is reserved for Ethernet use only.

AMM Web interface

Some functions require you to access the management-module Web interface.

For detailed information related to the AMM Web interface, see the IBM BladeCenter: Management Module User's Guide (ftp://ftp.software.ibm.com/systems/support/system_x_pdf/42c4886.pdf).

Chapter 3. Planning for BladeCenter Open Fabric Manager

Planning involves understanding the hardware and software requirements for OFM and OFM-Advanced upgrade and preparing for OFM.

Preparing for OFM

In order to prepare your environment for OFM, you will need to upgrade the firmware of the AMMs and blades, including the BMC, BIOS, and additional expansion cards in your environment. In addition, optimum use of OFM requires that you setup your blade environment to boot from SAN.

Important: UXSPs simplify the updating of all of your firmware. However, before you can upgrade your firmware on Emulex and Qlogic drivers, you must ensure that these fiber channel cards are already installed and operating properly. If you are not using UXSPs, follow the instructions in "Steps to update firmware without an OS" on page 17.

Upgrading firmware

Before you can use OFM, you must first update the firmware of the BladeCenter and blades, including BMC, BIOS, and additional expansion cards in your environment.

About this task

For a list of the latest firmware, see: Software and device drivers - IBM BladeCenter

Steps to update AMM firmware using the AMM Web interface

Procedure

1. Login to the AMM Web interface and select **Firmware Update** on the left pane, under **MM Control**. The **Update MM Firmware** page will open in the right pane.
2. On the Update MM Firmware page, click **Browse** to find the AMM flash file.
3. A separate **Choose file** window will open. Select the AMM flash file and click **Open**. The AMM flash file will appear in the field next to the **Browse** button on the **Update MM Firmware** page.
4. Click **Update** button on the **Update MM Firmware** page and wait for the firmware to be uploaded to the AMM. If there is a standby MM installed, the firmware on the standby MM will automatically update to the same level.
5. Click **Continue** to actually perform the flash.
6. Once the flashing is complete, you must reboot the AMM.

What to do next

On reboot, the new firmware will be active and the standby AMM firmware will be automatically updated.

Steps to update AMM firmware using UpdateXpress for BladeCenter (UXBC)

Procedure

1. Download Python interpreter, version 2.3 or later at http://www.python.org. The UXBC uses the Python update scripts to update the firmware of the applicable systems. To run the Python scripts, you must install Python interpreter on the administrative system.

   Note: Python also comes with most Linux distributions.
2. Download the latest UpdateXpress CD2 at IBM System x Support Web site (www.ibm.com/servers/eserver/support/xseries/index.html).
3. Go to the Software and device drivers - IBM BladeCenter Web site (http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-63017&brandind=5000020) (www.ibm.com/servers/eserver/support/xseries/index.html). Select your blade type from the list. This action directs you to a page for your blade type. Click **Mangement Module**, which takes you to the Management Module section of the page. You can then select from a list of available AMM firmware updates. After you download the package, you must unzip the package and save it to the applicable server or network location for distribution to target systems.
4. To save the firmware updates for a BladeCenter management module, complete the following steps:
   a. Copy the firmware update package to a network directory that you can access from the administrative console.

      Note: Do not unzip the firmware update package in the \BladeCenterUpdates directory. Each firmware update package includes a readme file. If you unzip the package in the \BladeCenterUpdates directory, the UXBC readme file is replaced with the update package readme file.

   b. Make a note of the directory path. This information is used to create the response file when BladeScanner is run.

Steps to update AMM firmware using BladeScanner and ChassisUpdate: Procedure

1. Use Bladescanner to create a response file.

   Note: Running ChassisUpdate with the -s (scan) switch (with valid Management Module login credentials) will also create a default response file.
2. Make sure that you have the file name and directory location of the response file that you want to use.
3. In the MM section of the response file, set the paths of the AMM firmware updates and make sure mmdisable is set to false.

```
### MANDATORY Fields ###
# These fields must be specified.

# This is a mandatory field that specifies the dotted IP
# address of the BladeCenter Management Module.
mmipaddr=192.168.70.125

####  OPTIONAL Fields ###
# These fields may be necessary depending on the BladeCenter configuration.
# This is a mandatory field that specifies the username for the
# BladeCenter Management Module.
mmuser=UX2

# This is an optional field that contains the password of the specified
# username for the BladeCenter Management Module.  If not specified, an
# empty password will be supplied to the Management Module.
mmpass=nIKH7P!,

# This field is mandatory if you intend to update the Management Module.  If
# not overridden, the default paths are used.
mmFilename1=\\server\share\AMMFirmware\BPETXXX.TKT
# mmFilename2=FILE2
# mmFilename3=FILE3
# This is an optional field that disables the update of the BladeCenter
# Management Module.  If not specified, or specified as FALSE, the
# Management Module is updated.
mmdisable=FALSE
# This is an optional field that is used for informational purposes by
# BladeScanner. BladeScanner in scan mode detects the firmware revision of
# the MM and stores it in this field. BladeScanner in edit mode reads the MM
# firmware revision from this field and displays it on the UI.
# The update scripts ignore this field.
mmMainAppFirmwareRevision=BRET86L
mmMainAppRevisionNum=16
#mmBootRomFirmwareRevision=
#mmBootRomRevisionNum=
#mmRemoteControlFirmwareRevision=
#mmRemoteRevisionNum=
mmPS2toUSBFirmwareRevision=BREZ15
mmMMtoUSBFirmwareRevision=BRPI33
mmName=MM00096BCA2328
```

4. From a command-line prompt, change to the disk drive that contains the UpdateXpress for BladeCenter utilities.
5. Type the following command to run the ChassisUpdate utility:

   ```
   chassisupdate.py -r file
   ```

   where, *file* is the fully qualified file name of the response file that you want to use. The ChassisUpdate utility reads the parameter from the response file and updates the applicable systems.

Results

BladeScanner and ChassisUpdate record the transactions that they perform in a single log file. The log file is created in the following directories:

Table 11. BladeScanner and ChassisUpdate log file locations

Windows	**Linux**
%TEMP%\uxbc.log where %TEMP% is the temporary directory for the Windows operating system.	$HOME/uxbc.log

Information is appended to the uxbc.log file each time you run BladeScanner or ChassisUpdate. As a best practice, you should periodically delete this file.

Steps to update firmware without an OS

About this task

The use of UXSPs requires that you have an OS running on the blade. If you do not have an OS on your blade, you can follow these steps to upgrade your firmware.

Steps to update blade BMC firmware:
Procedure

1. Download the boot image (BMC update diskette) of the latest version of the BMC firmware for your system from the Software and device drivers - IBM BladeCenter Web site (http://www-304.ibm.com/jct01004c/systems/support/

supportsite.wss/docdisplay?lndocid=MIGR-63017&brandind=5000020). Select your blade type from the list. This action directs you to a page for your blade type. Click **Baseboard Management Controller**, which takes you to the Baseboard Management Controller section of the page. You can then select from a list of available boot images. The boot image (BMC update diskette) will have a file extension of `img`.

2. Boot the blade to the downloaded image by either creating a diskette from the image or using the remote drive feature of the AMM.
3. Follow the directions to update the BMC firmware.

Steps to update blade BIOS:
Procedure

1. Download the boot image of the latest version of the BIOS update diskette for your system from the Software and device drivers - IBM BladeCenter Web site (http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-63017&brandind=5000020). Select your blade type from the list. This action directs you to a page for your blade type. Click **BIOS**, which takes you to the BIOS section of the page. You can then select from a list of available boot images. The BIOS will have a file extension of `img`.
2. Boot the blade to the downloaded BIOS by either creating a diskette from the image or using the remote drive feature of the AMM.
3. Follow the directions to update the BIOS update diskette.

Steps to update Emulex HBA firmware for x86 architecture:
Procedure

1. Download the latest version of the Emulex HBA firmware for your system from the Software and device drivers - IBM BladeCenter Web site (http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-63017&brandind=5000020). Select your blade type from the list. This action directs you to a page for your blade type. Click **Fibre Channel**, which takes you to the Fibre Channel section of the page. You can then select from a list of available Emulex HBA firmware updates.
2. Create a bootable DOS diskette image containing the doslpcfg.exe flash tool and the <flash image name>.prg flash file.
3. Boot the blade to the downloaded image by either creating a diskette from the image or using the remote drive feature of the AMM.
4. Type the following commands:

```
> doslpcfg download n=1 i=<flash image name>.prg
> doslpcfg download n=2 i=<flash image name>.prg
```

Steps to update QLogic firmware for x86 architecture:
Procedure

1. Download the latest version of the QLogic HBA firmware for your system from the Software and device drivers - IBM BladeCenter Web site (http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-63017&brandind=5000020). Select your blade type from the list. This action directs you to a page for your blade type. Click **Fibre Channel**, which takes you to the Fibre Channel section of the page. You can then select from a list of available QLogic HBA firmware updates.
2. Create a bootable DOS diskette image containing the flashutil.exe flash tool and the <flash image name>.bin flash file.
3. Type the following command:

```
> flashutil /f /o<flash image name>.bin
```

Steps to update Emulex or QLogic HBA firmware for POWER PC architecture: Procedure

1. Pre-OS installations on POWER PC architecture systems can only be accomplished by use of the IBM Standalone Diagnostics CD-ROM. The CD-ROM can be ordered from the IBM Publications Center (http://www.elink.ibmlink.ibm.com/publications/servlet/pbi.wss?CTY=US&FNC=SRH&). Search for "**pSeries standalone**".
2. Obtain the latest firmware for Emulex or QLogic HBA from the IBM BladeCenter Support Web site (https://www-304.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000020). Select **JS22** or **JS12** blades in the **Product family** field. Refine the results by selecting **Fibre Channel firmware**, then select from a list of available Emulex and QLogic firmware updates.
3. Create an ISO image CD using the acquired image. **Note:** Some external Windows or AIX-based workstations must be used in this step to create the ISO image.
4. Load the Standalone Diagnostics CD from the blade containing the HBA requiring the firmware update. Use the SoL interface on the JS blade to access the Standalone Diagnostics CD. Follow the documentation included with the Standalone Diagnostics CD-ROM to download firmware to the Emulex or QLogic HBA. After starting the Standalone Diagnostics CD-ROM, the Diagnostics CD must be removed and replaced with the CD you created in the previous steps. The Standalone Diagnostics utility will use this new CD as the source of the firmware download.

Setting up boot from SAN

In order to take full advantage of the BladeCenter Open Fabric Manager solution, you will want to setup your blade environment to boot from SAN.

Before you begin

For more information, see:

- IBM BladeCenter 4Gb SAN Solution (http://www.redbooks.ibm.com/abstracts/sg247313.html?Open)
- Emulex's IBM BladeCenter HBA Installation and Management white paper (http://www.emulex.com/white/hba/IBMBlade.pdf)

Chapter 4. Installing BladeCenter Open Fabric Manager

This topic contains information on installing OFM and OFM-Advanced upgrade

Installing OFM

The use of OFM does not require any installation as OFM capabilities are accessible through the AMM software. However, there are prerequisite tasks that must be satisfied before you can begin using OFM.

Ensure you have completed the steps outlined in “Preparing for OFM” on page 15.

Installing OFM-Advanced upgrade

You can install OFM-Advanced upgrade as a stand-alone product and as an IBM Systems Director 6.1.1 plug-in.

Installing the IBM Systems Director plug-in version of OFM-Advanced upgrade for Windows

Follow these steps to install the IBM Systems Director plug-in version of OFM-Advanced upgrade for Windows.

Before you begin

The BladeCenter OFM-Advanced upgrade 3.0 plug-in for IBM Systems Director requires IBM Systems Director 6.1.1. Before you install OFM-Advanced upgrade 3.0, verify that the target system is running IBM Systems Director 6.1.1. If the system is running IBM Systems Director 6.1.0, use update manager to acquire and install the 6.1.1. update. Do not install the IBM Systems Director 6.1.1 after you install OFM-Advanced upgrade 3.0; IBM Systems Director 6.1.1 must be installed before you install OFM-Advanced upgrade 3.0.

Procedure

1. Using an account with either local or domain **Administrator** authority, log on to the operating system.
2. Double click on the **BladeCenter Open Fabric Manager Advanced Upgrade for IBM Systems Director** installation package. This will launch the installation wizard.
3. Click **Next**.
4. Accept the license.
5. Click **Next**.
6. Click **Install**.
7. Click **Yes** to restart the IBM Systems Director Server.
8. Click **Done**.

Installing the IBM Systems Director plug-in version of OFM-Advanced upgrade for Linux

Follow these steps to install the IBM Systems Director plug-in version of OFM-Advanced upgrade for Linux.

Before you begin

The BladeCenter OFM-Advanced upgrade 3.0 plug-in for IBM Systems Director requires IBM Systems Director 6.1.1. Before you install OFM-Advanced upgrade 3.0, verify that the target system is running IBM Systems Director 6.1.1. If the system is running IBM Systems Director 6.1.0, use update manager to acquire and install the 6.1.1. update. Do not install the IBM Systems Director 6.1.1 after you install OFM-Advanced upgrade 3.0; IBM Systems Director 6.1.1 must be installed before you install OFM-Advanced upgrade 3.0.

Procedure

1. Download 46C3552GMAR.iso.
2. Using an account with **root** authority, log on to the operating system.
3. Mount the iso image to a mount point and change directory to the mount point directory as follows:

   ```
   > mkdir /mnt/iso
    > mount 46C3552GMAR.iso /mnt/iso -t iso9660 -o ro,loop=/dev/loop0
    > cd /mnt/iso
   ```
4. Run the script file as follows:

   ```
   > ./BOFMAUxx_DirectorExtension_Linux.sh
   ```

Installing the OFM-Advanced upgrade stand-alone version for Windows

The OFM-Advanced upgrade stand-alone version for Windows is downloaded as a single executable file. Launching this executable file starts the installation process.

About this task

Procedure

1. Download the standalone installation file (BOFMAUxx_StandAlone_Windows.exe).
2. Using an account with either local or domain **Administrator** authority, log on to the operating system.
3. Run the install program.

What to do next

Installing the OFM-Advanced upgrade 3.0 stand-alone version installs an embedded version of IBM Systems Director 6.1.1. No IBM Systems Director update is required to install the stand-alone version. Do not perform any IBM Systems Director Server updates after installing the OFM-Advanced upgrade 3.0 stand-alone version.

Installing the OFM-Advanced upgrade stand-alone version for Linux

Follow these steps to install the OFM-Advanced upgrade stand-alone version for Linux.

Procedure

1. Download 46C3551GMAR.iso.
2. Using an account with **root** authority, log on to the operating system.
3. Mount the iso image to a mount point and directory:

```
> mkdir /mnt/iso > mount 46C3551GMAR.iso /mnt/iso -t iso9660 > cd /mnt/iso
```

4. Run the script file:

   ```
   > ./BOFMAUxx_StandAlone_Linux.sh
   ```

What to do next

Installing the OFM-Advanced upgrade 3.0 stand-alone version installs an embedded version of IBM Systems Director 6.1.1. No IBM Systems Director update is required to install the stand-alone version. Do not perform any IBM Systems Director Server updates after installing the OFM-Advanced upgrade 3.0 stand-alone version.

Migrating OFM-Advanced upgrade

You can migrate configuration profiles from previous releases of OFM-Advanced upgrade. Configuration profiles can be migrated from OFM-Advanced upgrade V1.0 and V2.x to OFM-Advanced upgrade V3.0 using the IBM Systems Director Migration Tool. The migration tool migrates all existing OFM-Advanced upgrade configuration manager profiles and converts them to configuration templates and configuration plans.

IBM Systems Director Migration Tool

Use the IBM Systems Director Migration Tool to migrate OFM-Advanced upgrade data, for both the plug-in version and stand-alone version of OFM-Advanced upgrade. Because the stand-alone version of OFM-Advanced upgrade includes an embedded version of IBM Systems Director, the IBM Systems Director Migration Tool is used to migrate OFM-Advanced upgrade.

The IBM Systems Director Migration Tool allows you to migrate from systems running IBM Director 5.20 to systems running IBM Systems Director 6.1.

Important: You must install IBM Systems Director 6.1 on a different system than the system running OFM-Advanced upgrade, plug-in or stand-alone, on IBM Systems Director 5.20.

Using the IBM Systems Director Migration Tool, you can:

- Migrate from OFM-Advanced upgrade plug-in running on IBM Systems Director 5.20.x to IBM Systems Director 6.1.1.
- Migrate from OFM-Advanced upgrade 2.x stand-alone to OFM-Advanced upgrade stand-alone 3.x.
- Migrate from OFM-Advanced upgrade 2.x stand-alone to OFM-Advanced upgrade 3.x plug-in running on IBM Systems Director 6.1.1.
- Migrate from OFM-Advanced upgrade plug-in running on IBM Systems Director 5.20.x to OFM-Advanced upgrade stand-alone 3.x.

For more information on the IBM Systems Director Migration Tool, see Upgrading and migrating IBM Systems Director (publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.upgrade_6.1/fqm0_t_upgrading_and_migrating.html).

To migrate OFM-Advanced upgrade using the IBM Systems Director Migration Tool:

1. Download and install the IBM Systems Director Migration Tool.

See Obtaining and installing the IBM Systems Director Migration Tool (publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.upgrade_6.1/fqm0_t_obtaining_and_installing_migration_tool.html).

2. Export configuration profiles from the system running a previous release of OFM-Advanced upgrade.

   See Exporting data from IBM Director 5.20 (publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.upgrade_6.1/fqm0_t_exporting_5.20_data.html).
3. Import configuration profiles to the target system running OFM-Advanced upgrade 3.0.

   See Importing IBM Director 5.20 data into IBM Systems Director 6.1 (publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.upgrade_6.1/fqm0_t_importing_5.20_data_to_6.1.html).

Chapter 5. Configuring BladeCenter Open Fabric Manager

Configuring OFM involves a number of steps that include the creation of your configuration file and applying the new configuration, along with several other considerations.

Important: You must configure OFM through the AMM Web interface before you can use either the IBM Systems Director plug-in or stand-alone versions of OFM-Advanced upgrade.

Creating a configuration file automatically

When using OFM for the first time, you will need to create a configuration file in which you assign virtual addresses to each slot in each chassis.

About this task

The following example outlines the kind of steps you might follow when creating a configuration file automatically. It does not apply to all BladeCenter environments. These example steps assume that you have a single domain (no addresses are duplicated).

Procedure

1. Log in to the AMM Web interface and select **Open Fabric Manager** in the left pane, under **Blade Tasks**. The **Open Fabric Manager Configuration Management** page will open in the right pane.
2. Click **Create an Initial Configuration**. This will open the **Specify Virtual Addresses** page in the right pane.
3. For the **Ethernet Address Type**, use the pull-down under **Vendor** and select IBM.

   Note: Another option for Vendor is `User Defined`.
4. For the **FC Address Type**, use the pull-down under **Vendor** and select Emulex or QLogic.
5. For the **SAS Address Type**, use the pull-down under **Vendor** and select LSI or IBM range.
6. Click **Advanced option** and check the box next to **Generate an FC target place holder** or **SAS target place holder**.
7. In the **WWN** field, enter the storage system WWPN.

   Note: You also have an option here to specify a value in the **LUN** field.
8. Click **Next**. This will open the **Chassis to include** page in the right pane.
9. Click **Next**. This will allow you to optionally add to existing OFM configuration. This allows you to specify an existing OFM configuration file (CSV file) using the **Browse** button. This file will be prepended to the newly generated OFM configuration file that will contain addresses that continue those in the existing specified file. This extends the existing OFM domain. If you do not want the new OFM configuration prepended to an existing configuration, do not specify any file name. Instead, click **Next**.
10. On the **Chassis to include** page, there are two methods for providing the list of chassis to be included in the configuration file.

You can either create a file with the list of AMM IP addresses or use the chassis that were discovered by the AMM via SLP. When using the **Use AMM IP Addresses that were discovered on the AMM management network** button, you should first operate the Remote Chassis page using the SLP method to verify that all chassis that show up there are those you on which want to configure OFM. Otherwise explicitly specify the chassis address list in a file as explained in the next step.

If you elect to use an explicit list of AMM IP addresses instead, create a text file in which each line contains a single IP address or the hostname of a single chassis. If you use hostnames in this file, you will need to enable DNS and define at least one DNS server on the AMM Web interface **Network Protocols** page. When the text file is complete, you can then select the **Use AMM IP Addresses in a file that I specify** option on the **Chassis to include** page. Click **Browse** to locate the file that you created.

Note: You can also use a valid existing OFM configuration file to define the list of chassis.

11. Click **Next**. The AMM will now generate the configuration file and display **The Configuration File Has Been Created** page.
12. The browser will launch the **File Save** window allowing you to save the generated configuration file. If the **File Save** window does not appear, click **Download the configuration file manually** on **The Configuration File Has Been Created** page.

What to do next

It is recommended that you store the configuration file locally, and validate the new configuration. It is also important to store the configuration file in a safe location because this is your original copy of the OFM configuration. If an AMM has a hardware failure and you don't have a standby AMM, then this is your single source to reproduce the OFM configuration.

If you wish to apply the configuration directly or to create a **Requirements Report**, you can do it directly from this page.

Some considerations for the **Specify Virtual Addresses** page:

- Some applications will check the adapter type and vendor by looking at the address. You can select ranges that meet the type of adapter that you are using. Each vendor has allocated a special range for OFM outside of their normal range which guarantees that these addresses do not conflict with any previous or future burned-in addresses. Selecting a vendor automatically sets values within that range. You can also set it to `user defined` mode and select any range by editing the **From** and **To** fields.

  The default address ranges are as follows:

```
Ethernet: The range for IBM MAC range is:
00:1A:64:76:00:00 - 00:1A:64:76:FF:FF.

FC :
  Qlogic: WWPN odd port range : 21:80:00:E0:8B:0X:XX:XX
  WWPN even port range: 21:81:00:E0:8B:2X:XX:XX
  where : X = 0..F
  WWNN adresses are generated internally by Qlogic from the WWPN.

Emulex : WWNN odd port range: 2F:FE:00:00:C9:XX:XX:XX
  WWNN even port range: 2F:FF:00:00:C9:XX:XX:XX
  WWPN odd port range: 2F:FC:00:00:C9:XX:XX:XX
  WWPN even port range: 2F:FD:00:00:C9:XX:XX:XX
  where x = 0..F
```

```
SAS:
  IBM range for WWPN: 50:05:07:60:1A:80:00:02 to
  50:05:07:60:1A:BF:FF:FF
  LSI range for WWPN:
  50:00:62:B0:00:11:17:02 to 50:00:62:B0:00:12:16:ff
```

- For FC (FibreChannel), there are two ranges for each vendor: one for odd-numbered ports and one for even-numbered ports. The `non-OFM` default is that the system assigns an address from one of the ranges for each port of the device. When generating the file automatically, the system allocates the addresses for the even ports out of the first range and allocates the addresses for the odd ports out of the second range.

  Note: For some devices this might not be the appropriate allocation. For example, a Fibre Channel high speed adapter is connected to ports 6 and 8, which are both allocated out of the second range, and this might appear as two different devices rather then a single device. You can update the file manually to match your specific devices.
- For some vendors (such as QLogic), you do not need to define a WWPN, since it is automatically derived from the WWPN.
- By default, when you choose to assign MAC addresses they will be assigned for each of the ports 1 to 8, FC addresses will be assigned for each of the ports from 3 to 8, and SAS for ports 3 and 4 only. These configurations match a single slot blade. These are generic configurations, which contain a virtual address for each possible hardware type (Ethernet expansion card, FC expansion card, or SAS expansion card). As a result, when you change the type of blade or expansion card, you do not need to modify the OFM configuration. However, this option makes the configuration file bigger and error validation harder. If you want to generate a configuration that assigns an address to a subset of the ports, you can use the advanced options section. In this section, you can select which type of address to assign to each port.

  Note: You can not assign FC addresses to ports 1 and 2, because these are the on board Ethernet NICs.
- Some Ethernet expansion cards can have a range of MAC addresses per port. The range is defined by specifying two MAC addresses per port: MAC A and MAC B. To set up a range of MAC addresses for a port, click **Advanced option**, then click **Generate range of MAC addresses per port**. You will then be able to define the following values and make consideration for the JS and PS blade requirements noted below:
  - **List of ports to apply to (numbers in the range 1 to 8, comma or space separated)**: The default value is 1.
  - **Range size:** Enter a range size between 2 and 256. The default value is 16.
  - **Ethernet VLAN for the second MAC address:** The default is 0.

  Note: The above default values will allow the JS and PS blade onboard HEA Ethernet adapter to use the virtual addresses preassigned by the OS. To avoid conflicts, the MAC Address Step value (range size) must be between 16 and 256.
- For the NetXen 10 Gb Ethernet Expansion Card for IBM BladeCenter (39Y9271), do not use the **Advanced option** button to apply Fibre Channel (FC) addresses to FC ports that exist on the card. This will generate errors.
- Using the **Advanced option**, you can also assign addresses to multi-slot blades. To do this, select the slot offset that you wish to assign addresses to, and, in the table for that offset, select the ports and type of address that you would like to assign. The maximum number of addresses that you can assign to any single

blade is 32 addresses for a four-slot blade. For more information, see "Multi-slot blades and the port offset parameter" on page 13.

- In the **Advanced option**, you can also define the increment of the assigned addresses. The default is one (by default addresses are assigned sequentially). You can also define the VLAN tag for host-based VLAN tagging. Some Ethernet cards can have two MAC addresses per port. For these, you can elect to assign two MAC addresses by checking the **Generate range of MAC addresses per port** checkbox with a range of 2 addresses.
- In the advanced options for the FC section, you can also specify the step increment for the WWN and WWPN. In addition, you can select to create templates for the FC targets. After the file is created, you can edit the configuration file and fill in the correct target WWN and LUN for each slot.

Selecting domains

In some complex environments, you might need to create separate configuration files for multiple domains.

A single OFM configuration corresponds to a single OFM domain, in which duplicated addresses are not allowed. In general, unless you have a special reason to do so, it is recommended that you have one configuration file for the entire domain. This way the AMM will verify that there are no address duplications when you modify the OFM configuration.

However, in a more complex environment, it may be convenient to have multiple domains, where the same addresses can exist in different domains. In this case, you can generate a separate configuration file for each domain.

CAUTION:
We strongly advise against using multiple configurations on a single network, since this will override the protections against MAC or FC address duplication as discussed in "Avoiding address duplication."

Avoiding address duplication

When creating your OFM configuration file, you need to avoid address duplication.

Duplication of MAC addresses can cause serious issues with your network. Fibre Channel address duplication can lead to data corruption if more than one blade is trying to access the same volume at the same time.

To protect from duplicate addresses, the AMM performs as follows:

1. When a new configuration file is applied, the AMM verifies that it does not contain internal duplicate addresses.
2. Before applying the configuration, the AMM verifies that the blades that are about to be re-configured are powered off. This ensures that no addresses are currently in use.
3. Before writing the new OFM configuration to the blades, the AMM disables the OFM configuration on all the blades that are about to be re-configured. As a result, there can be no address duplication even if the operation has not completed.

For flexibility, the user can override these checks and apply the configuration even if the AMM generates warnings that duplicate addresses might exist. In addition, if the user defines the same address, in two different configuration files, for two chassis that are on the same network, then the AMM can not protect against address duplication.

To avoid address duplication, it is best to use a single configuration file for a single network domain and not use the options to override the protection that the AMM provides.

Duplicate addresses

An error page results when the configuration file contains duplicate addresses.

If the configuration file contains duplicate addresses, an **Open Fabric Manager Configuration Failure** window will display in the AMM Web interface. The window contains a table. Each row of the table indicates an address that is duplicated and all of the lines in the configuration file on which the error appears. If you did not intend to have duplicate addresses, you must press **Cancel**, fix the configuration file, and then re-apply it.

CAUTION:
If you wish to ignore the duplicated addresses and apply the configuration anyway, you can click the Ignore button. We strongly recommend you do not do this unless you are absolutely sure you know what you are doing. Duplicate address can cause serious problem on your network. To avoid the address duplication check, you can also select the Ignore duplicate virtual addresses in the configuration file advanced option on the Apply a Configuration window Again, we strongly recommend you do not do this unless you are absolutely sure you know what you are doing.

Creating a requirements report

Before you apply the OFM configuration, IBM recommends that you create a requirements report.

About this task

The requirements report verifies the firmware level of the blades BMC, BIOS and adapters. It also goes through a dry run of the first part of applying the OFM configuration and checks to see if there are potential problems.

Creating a requirements report from "The Configuration File Has Been Created" page in the AMM Web interface

Procedure

1. Follow the steps outlined in “Creating a configuration file automatically” on page 25.
2. On **The Configuration File Has Been Created** page, click **creating a requirements report** .

Creating a requirements report from the main Open Fabric Manager Configuration Management page on the AMM Web interface

Procedure

1. Login to the AMM Web interface.

2. In the left pane, under **Blade Tasks**, select **Open Fabric Manager**. The **Open Fabric Manager Configuration Management** page will open in the right pane.
3. Select **Create a Requirements Report**. The **Create a Requirements Report** page will open in the right pane.
4. Click **Browse** to locate the configuration file for which you are creating the requirements report. This would typically be a configuration file you want to use to apply a new OFM configuration.
5. You can select **Advanced Options** to override certain checks. See “Applying a new configuration” for more information.
6. After you locate the configuration file, click **OK**.

What to do next

If no parsing errors are found in the configuration file, the **Requirements Report** page will be displayed.

For a description about how to read the requirements report, a list of typical requirements report errors, and how to take action on the errors, see “Error messages” on page 46.

Editing the configuration file manually

After the initial generation of the configuration file, you can edit it to make changes to accommodate the specific needs of your environment.

Applying a new configuration

You can apply a new OFM configuration from the main Open Fabric Manager page on the AMM Web interface.

Procedure

1. Login to the AMM Web interface.
2. In the left pane, under **Blade Tasks**, select **Open Fabric Manager**. The **Open Fabric Manager Configuration Management** page will open in the right pane.
3. Select **Apply a Configuration**. The **Apply a Configuration** page will open in the right pane.
4. Click **Browse** to locate the configuration file you want to apply.
5. You can click **Advanced Options** to override checks that protect against address duplication.
   a. If you select **Ignore duplicate virtual addresses in the configuration file** under **Advanced Options**, the AMM will not check for address duplication.
   b. If you select **Force the configuration to be applied to powered on blades** under **Advanced Options**, the AMM will not check the power state of the blades.

      Note: In general, you should not change the OFM configuration while the blade is powered on. Changing the OFM configuration while the blade is powered on can lead to duplicate addresses, and unexpected results.
   c. If you select **Continue on error** under Advanced Options, the AMM will continue on most errors. Ignoring some of these errors may lead to address duplication. For example if the AMM is unable to connect to one of the

chassis it will continue even though that chassis might be using an address that is now assigned to a different chassis.

6. After you locate your configuration file and select advanced options, click **OK**. The AMM will start to prepare for applying the new OFM configuration and any errors that occur during this process will be reported.

   Note: If the new configuration is identical to the current configuration, then a **No Configuration Change** page will display.

   Before the configuration is applied, the AMM displays the expected OFM changes. There is a single line displayed for each chassis that has changes. If the OFM configuration for a given slot is changed, the corresponding cell in the table contains an icon, otherwise the cell is blank. If OFM is enabled for that slot, the icon is blue. If OFM is disabled for that slot, the icon is gray.
7. If you approve the changes, click **Continue**.

What to do next

At this point the AMM starts to apply the OFM configuration. As explained in "Avoiding address duplication" on page 28, the first phase is to disable the OFM configuration for all the slots that are about to be reconfigured. The next phase is to apply the new configuration for all of the slots. Finally, OFM is re-enabled on the appropriate slots. When the process completes, **The configuration file was applied successfully** page will appear in the right pane of the AMM Web interface.

Viewing the configuration in a local chassis

After you change the OFM configuration, you can view it on any of the chassis.

Procedure

1. Login to the AMM Web interface.
2. In the left pane, under **Blade Tasks**, click **Configuration**. The **Configuration** page opens in the right pane. The Configuration page contains an **Open Fabric Manager Parameters** table. In this table, you can see the OFM overview displayed under these column headings:
   - **Bay:** The location of the blade.
   - **Blade Name:** The name of the blade at that bay location.
   - **OFM Mode:** Either `enabled` or `disabled`, as defined by the OFM configuration file.
   - **Profile:** As defined in the configuration file.
   - **System Mgmt Processor OFM Capable:** Displays `Yes` if the BMC (blade systems management processor) supports OFM and `No` otherwise.
   - **BIOS OFM Capable:** displays `Yes` if the blade BIOS supports OFM and `No` otherwise.

     Note: This value is accurate only after the blade boots for the first time after updating the BIOS firmware.
   - **OFM status:** If the BMC (blade systems management processor) is OFM-capable and the blade BIOS is OFM-capable, the status will be `N/A` (Not Applicable), until the blade is powered on for the first time after the OFM configuration has been enabled. The status (after blade BIOS boot is completed) can then be: `Normal`, `Error` or `Warning`. If this blade OFM mode is disabled or one of the above OFM capabilities is missing, then the status will stay N/A even after power on of the blade

3. If you click on a Blade Name, you can see the detailed OFM configuration page for that slot as was defined in the OFM configuration file.

What to do next

The Configuration page is divided into sections of device address type: `Ethernet`, `FC` or `FC target`. In addition, this page displays information about whether devices (attached to this slot) support OFM and if any address type and value was consumed by any of the slot attached devices. If a OFM address was consumed, then the **Address Status** column will say `Used` or `Error` , otherwise it will say `Not Used`.

The Configuration page optionally displays a table of non-OFM devices, provided there are such devices attached to this blade. This table has 3 columns: **Slot offset**, **Port**, and **Address Status**. If such a device (characterized by <Slot-offset , Port>) is referred to in the OFM configuration file under this blade, the **Address Status** column will be flagged as `Warning`. Otherwise it will be flagged as `N/A`.

In general, the addresses that you see on the AMM Hardware VPD page are the current actual addresses. After a change to the OFM configuration, the addresses will not be updated until the next power on of the blade. To view the new OFM addresses, click the **Reload MAC/Unique IDs** button on the hardware VPD page.

Retrieving the current configuration

You can retrieve the current OFM configuration by selecting **Retrieve the Current Configuration** on the main Open Fabric Manager page.

Procedure

1. Login to the AMM Web interface.
2. In the left pane, under **Blade Tasks**, click **Open Fabric Manager**. The **Open Fabric Manager Configuration Management** page will open in the right pane.
3. Click **Retrieve the Current Configuration**. The **Retrieve Current Configuration** page will open in the right pane.
4. You can select **Use the IP addresses that were discovered by the AMM** or **Use AMM IP addresses in a file that I specify**. For more information on these two options, see “Creating a configuration file automatically” on page 25.
5. Click **Retrieve**. The AMM will start to retrieve the information.

What to do next

On completion, your browser will launch the **File Save** window allowing you to save the generated configuration file. If the **File Save** window does not appear, click **Download the configuration file manually**.

Chapter 6. Using OFM and OFM-Advanced upgrade

This section contains information on using the two offerings: OFM and OFM-Advanced upgrade

Using OFM

This topic contains information on using OFM.

Initial deployment

Follow these steps for an initial deployment.

Procedure

1. Update the firmware of the AMM, blade service processor, blade BIOSes and adapters. See "Upgrading firmware" on page 15.
2. Create an initial configuration file. See "Creating a configuration file automatically" on page 25.
3. Download and store the new configuration file in a safe location. This file is the source of your OFM configuration.
4. Update the configuration file, if required, and save the resulting file.
5. Create a requirements report. Validate that you can apply the OFM configuration without errors. You can apply the configuration successfully even if no blades are present or if the blade's firmware does not support OFM. In this case, the AMM stores the configuration on the AMM and it will push it to the blade when it is inserted or when its firmware is updated. See "Creating a requirements report" on page 29.
6. Apply the OFM configuration. See "Applying a new configuration" on page 30.

Adding a new chassis to the domain

When you need to add a new chassis to the domain, you will need to add its OFM configuration to the OFM configuration file. You can either do it manually or use the generate feature to create a new configuration file. In that case, you will need to define the ranges so that they do not contain addresses that are already used by the current configuration.

About this task

Procedure

1. Create a file where each line contains the IP address or hostname of the new chassis that you wish to add to the domain.
2. Create a new configuration file using the list you just created.
3. After creating the new file, you can append it to the current file using a spreadsheet or text editor.
4. Apply the configuration.

What to do next

When you apply the combined file, only the changes to the new chassis are applied assuming that for the other chassis the configuration in the file matches the existing configuration.

Replacing a blade in the same slot

When you replace a blade in a slot that is OFM enabled, the OFM parameters are automatically applied to the new blade before it is given power permission. The boot sequence is not part of the OFM configuration and it is not pushed automatically.

Swapping addresses between blades

If you want to swap the addresses of one blade to a different blade in a different slot (potentially a different chassis), you can move the OFM configuration from the first slot to the second.

Before you begin

Procedure

1. Edit the OFM configuration and swap the configurations of the two slots.
2. Apply the new configuration. Only the configuration of these two slots will be updated.

What to do next

Note:

- You will have to manually swap the boot sequence of the two slots (if necessary).
- In cases where slot based configuration has been used on the switches, such as VLAN for Ethernet or Zoning for Fibre Channel, those configurations will need to be updated as well.

Replacing AMM IP addresses

The OFM configuration is defined per AMM IP address. If you change the IP of the AMM in a chassis and reapply the same OFM configuration file, then the OFM configuration changes. For example, assume that Chassis-A uses IP-A, and Chassis-B uses IP-B, and you apply a OFM configuration for these two chassis. If you swap the IPs so that Chassis-A uses IP-B, and Chassis-B uses IP-A, then re-apply the same OFM configuration file, the OFM configuration will be swapped between the two chassis.

Replacing the AMM in a single AMM environment

You can replace the AMM in a single AMM environment.

About this task

When you replace an AMM and you don’t have a standby AMM in the chassis, the OFM configuration is cleared. If the current AMM is functioning, the best approach is to insert the new AMM in the second slot as a standby AMM for a few minutes. This allows the primary AMM to synchronize the OFM configuration with the standby AMM before the primary AMM is removed. If this is not possible, you can reapply the OFM configuration after inserting the new AMM. By default, the AMM clears its OFM configuration when inserted into a new chassis. However, if any of the blades are already powered on when the AMM is inserted into a new chassis, the AMM takes the OFM configuration from those blades.

If the AMM was reinserted into the same chassis it was previously in, it continues to use the OFM configuration that is defined on the AMM. During the period that the AMM was not in the chassis, the OFM configuration might have changed. In this case, inconsistencies and address duplications can occur.

If the AMM configuration is reset to factory defaults, then the OFM configuration is cleared the same way as if a new AMM is inserted into a chassis. Also, the OFM configuration is not included in the AMM configuration file, when restoring the AMM configuration from a file the OFM configuration has not changed.

Note: You might have to wait longer if the AMM you insert in the standby slot does not have the same firmware level as the primary. In that case, the standby AMM will be flashed first and then the data will be synchronized.

Using OFM-Advanced upgrade

This topic contains information on using OFM-Advanced upgrade, including program launch and chassis discovery, creating a standby blade pool profile, creating an event action plan and manually applying a failover configuration profile.

Important: You must to configure the AMM to enable OFM-Advanced upgrade access. The following network protocols will need to be set on the AMM:

File Transfer Protocol (FTP)

1. From the AMM Web interface access the **MM Control** → **Network Protocols** page.
2. Click on the **File Transfer Protocol (FTP)** link and make sure that it is set to **Enabled**.

TCP Command Mode Protocol

1. From the AMM Web interface access the **MM Control** → **Network Protocols** page.
2. Click on the **TCP Command Mode Protocol** link and make sure that the Command mode is set to 20 connections.

Creating a blade address manager configuration template

You can create a blade address manager template using the Blade Address Manager Configuration wizard in IBM Systems Director.

About this task

To launch the Blade Address Manager Configuration wizard, complete the following steps:

1. In the IBM Systems Director navigation area, click **Welcome**.
2. On the Welcome page, click the Manage tab. A list of available summary pages is displayed.

3. On the Manage tab, scroll to the **BladeCenter Open Fabric Manager** section heading and click it. The BladeCenter Open Fabric Manager summary page is displayed.
4. Under the Configuration heading, click the **Create...** link located under the Blade Address Manager Templates section. The Blade Address Manager Configuration wizard is displayed.
5. Complete the following wizard pages to create a Blade Address Manager template:
   a. **Create template**

      Select the target type and specify a name for the template. You can also specify a description of the template.

   b. **Welcome**

      Click **Next**.

   c. **Import**

      Select whether you want to import an existing Blade Address Manager .csv configuration file. Importing an existing Blade Address Manager configuration file will populate this wizard with the settings and values from the imported configuration file. This can be used as a starting point to edit, then create a new configuration file based on the imported configuration.

   d. **Member**

      Select the BladeCenter chassis that will participate in the domain by clicking on the chassis in the **Available Chassis** box. Click **Add** to move the selected chassis to the **Domain Members** box. Chassis not listed in the **Available Chassis** box can be added to the domain by clicking **New...** and entering the IP address of the chassis.

   e. **Enable**

      Select the checkbox next to a chassis to enable the blade address configuration for the selected chassis. Unselect the checkbox next to a chassis to deactivate the blade address configuration for the selected chassis.

   f. **Address Pool**

      Complete the following fields:

      Maximum chassis in domain

      Enter the maximum number of chassis to include in the domain. The maximum size is 100. The software will pre-allocate address space for the maximum number of chassis chosen to ensure that there will be no address duplications in any of the templates.

      Automatically assign addresses

      Activate MAC

      To assign Ethernet MAC addresses to bays, click the **Activate MAC** check box. Then select **IBM** in the **Vendor:** dropdown list to allow the software to automatically assign addresses. Selecting **User Defined** in the Vendor dropdown list will allow you to edit the MAC address range to any range of your choosing.

      Some Ethernet expansion cards can have a range of MAC addresses per port. The range is defined by specifying two MAC addresses per port: MAC A and MAC B. For example, the JS22 and JS12 onboard HEA Ethernet adapter must use the default values shown below. To set up a range of MAC addresses for a port, select **User defined** from the **MAC addresses per port:** dropdown list. Then specify a value from 1 to 255 in the **Range:** field.

Note: The above default values will allow the JS22/JS12 onboard HEA Ethernet adapter to use virtual addresses. Even though the 16 MAC range for the onboard HEA adapter is declared via port 1, the actual routing of the 16 MACs is determined by configuration in the OS.

Activate FC

To assign Fibre Channel WWN addresses to bays, click the **Activate FC** check box. For **Vendor** select either **Qlogic** or **Emulex**, depending on the type of fibre channel cards in your environment, to allow the software to automatically assign addresses. Selecting **User Defined** will allow you to edit the Fibre Channel address range to any range of your choosing.

Valid address range for **Qlogic**:

- First port: 21:80:00:E0:8B:0X:XX:XX
- Second port: 21:81:00:E0:8B:2X:XX:XX

Valid address range for **Emulex**:

- WWNN: 2F:FE:00:00:C9:XX:XX:XX through 2F:FF:00:00:C9:XX:XX:XX
- WWPN: 2F:FC:00:00:C9:XX:XX:XX through 2F:FD:00:00:C9:XX:XX:XX

Activate SAS

To assign storage target WWPN addresses to bays, click the **Activate SAS** check box. Then select the **Vendor** type and enter the range of addresses.

Valid address range for **IBM**:

- 50:05:07:60:1A:80:00:00h to 50:05:07:60:1A:BF:FF:FFh

Valid address range for **LSI**:

- 50:00:62:B0:00:11:17:00h to 50:00:62:B0:00:12:16:FFh

g. **Boot Settings**

Select the primary and secondary ports of the storage system where the boot LUNs for the blades reside. Additionally, select the blade assigned LUN for the boot LUN. This will be the LUN that the blade will be booting from so it will need to have an operating system installed. These settings will be applied to every blade in the domain, but can be modified on a blade by blade basis on the Advanced Configuration panel of this wizard. If the storage system was discovered using Advanced Discovery, then the primary and secondary fields will show the WWPN information of the storage system.

h. **Advanced Options**

You can select to ignore three different types of errors when the template is applied. The types of errors are:

Ignore duplicate check when applying
Select this option to have BladeCenter Open Fabric Manager ignore errors occurring when a template containing duplicate addresses is applied to a domain of blades.

Apply the blades even if they are powered on
Select this option if you have blades in your domain that are currently powered on and you are changing their addresses. The new addresses will not be active until the next time that the blade is powered on. If this box is unchecked, then the template will fail when applied if there are blades in the domain which are currently powered on.

Continue to apply even if an error occurs
Select this option to have BladeCenter Open Fabric Manager ignore various errors, such as the inability to log into another chassis, and continue to apply addresses to blades and chassis that are accessible.

i. **Advanced Configuration**

You can select to:

- Assign addresses from the pool
- Manually override the default boot settings

.

j. **Advanced Settings**

If you selected to manually override the default boot settings, select the override options on the **Advanced Settings** page. You can also manually configure boot settings for each bay.

1) Select a chassis in the **Select Chassis** box by clicking on the IP address of the chassis.
2) Select a bay to edit in the **Select Bay** box by clicking on the bay number.
3) Under the **Select Chassis** box you can select to:

   Activate Bay
   Assigns addresses to the blade connected to that bay.

   Deactivate Bay
   Removes any addresses assigned to that bay and any blade connected to that bay will revert to their physical addresses.

   Ignore Bay
   Leaves the addresses in their previous settings. The addresses will remain either virtual or physical when this template is applied. The **Ignore Bay** setting is useful when two templates are created for a single chassis, such as when both Emulex and Qlogic adapter cards are in the same chassis.
4) Select the type of setting you would like to configure, then click **Edit...**.
5) In the **Fibre Channel Boot Settings** and **SAS Boot Settings**, you can specify the WWN and LUN number for the boot target. Then blade can then be booted from those addresses.

k. **Summary**

The Summary page displays the settings that you have selected. The hardware listed represents the chassis that you included in the domain members list. Ensure that the template options have been set correctly. To save the configuration template to a .csv file, click the **Export** button in the upper-right corner of the page.

Select **Deploy settings when finished** to apply the settings of this template to the hardware upon clicking **Finish**.

Applying a blade address manager configuration template

You can apply a blade address manager configuration template manually using IBM Systems Director.

Procedure

1. In the IBM Systems Director navigation area, click **Welcome**.
2. On the Welcome page, click the Manage tab. A list of available summary pages is displayed.
3. On the Manage tab, scroll to the **BladeCenter Open Fabric Manager** section heading and click it. The BladeCenter Open Fabric Manager summary page is displayed.
4. Under the Configuration heading, click the **Configuration templates** link located under the Common Tasks section.
5. Select a configuration template to deploy on one or more systems by selected the check box next to template in the Select column, then click the **Deploy** button.
6. On the Run page, select the target resource in the **Available** box, then click **Add >** to add the resource to the **Selected** box.
7. To deploy the template immediately, click **OK**. To schedule the deployment, click the **Schedule** tab at the top of the Run page.

Creating a standby blade pool configuration template

You can create a standby blade pool configuration template using the BladeCenter Standby Blade Pool Configuration wizard.

Procedure

1. In the IBM Systems Director navigation area, click **Welcome**.
2. On the Welcome page, click the Manage tab. A list of available summary pages is displayed.
3. On the Manage tab, scroll to the **BladeCenter Open Fabric Manager** section heading and click it. The BladeCenter Open Fabric Manager summary page is displayed.
4. Under the Configuration heading, click the **Create...** link located under the Standby Blade Pool Configuration Templates section.
5. On the Create template page, specify a name for the standby blade pool template. You can also specify a description of the template. Click **Continue**. The BladeCenter Standby Blade Pool Configuration wizard is displayed.
6. On the Welcome page of the wizard, click **Next**.
7. On the Blades Pool page, select the chassis bays to be used for failover in the **Available Bays** box, then click **Add >**. These bays will be tried in the order they appear in the **Selected Bays** box. Use the **Up** and **Down** buttons to specify the bay order in the **Selected Bays** box, then click **Next**.
8. On the Properties page, select the failover options to be used by the standby blade pool template, then click **Next**.
9. The Summary page displays the settings that you have selected. Ensure that the template options have been set correctly, then click **Finish**.

What to do next

Note: Once a standby blade pool has been created, the list of standby blades will remain in the template even after one or more of the blades has been removed from the BladeCenter chassis.

Deploying a standby blade pool configuration template manually

You can use IBM Systems Director to deploy a standby blade pool configuration template manually.

Procedure

1. In the IBM Systems Director navigation area, click **System Configuration** → **Configuration Templates**.
2. From the table of configuration templates, select the desired template, then click the **Deploy** button.
3. On the **Targets** tab, select the target systems to which you would like to apply the standby blade pool configuration template, then click **Add >**.
4. On the **Schedule** tab, you can select to deploy the template immediately, or select to schedule the deployment to run in the future.
5. Click **OK**.

What to do next

When the Standby Blade Pool is being applied is being applied to an active blade, the software will communicate with the active blade's corresponding chassis management device (the Advanced Management Module) to read the currently assigned addresses associated with that blade. These addresses will then be applied to the chassis management device associated with the selected Standby Blade's chassis slot.

During the applying of the Standby Blade Pool to an active blade, the Ethernet Switch Module's port based VLANs are migrated from the active blade to the standby blade. The software will connect to all of the Ethernet Switches in the active blade's chassis and read the VLAN port information associated with the active blade. The software will then connect to all of the Ethernet Switches in the standby blade's chassis and apply these port based VLANs to those Ethernet switches. This connection and configuration update is performed through the Ethernet Switch's Director plug-in which must be installed on the Director server otherwise the failover process will fail.

The OFM-Advanced upgrade software performs the following checks prior to implementing a failover from the source blade to the standby blade:

1. Checks for matching standby blade machine type
2. Check for matching standby blade model type
3. Check for standby blade initial power state to be off
4. Check for blade width
5. Skip migration of switch settings from source blade's switches to standby blade's switches

Creating an event action plan

Refer to the IBM Systems Director documentation for instructions to create event action plans.

About this task

For detailed instructions, see Managing event action plans (http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_t_managing_event_action_plans.html) in the IBM Systems Director documentation.

For a list of typical OFM events, see “Events” on page 50.

Chapter 7. Troubleshooting and support

This topic covers known problems and workarounds, limitations, and error messages related to OFM. It also contains information to get support and troubleshoot problems.

Problems when you create or apply a standby blade pool

When creating a standby blade pool or applying the standby blade pool, you may run into a number of error conditions. This section describes typical problems and how to resolve these issues.

Table 12. Blade address configuration creation and application problems

Problem	Solution
Template failed to be applied.	• Check that blades are not powered on. • Check that you have discovered all of the chassis in the domain. • Check that you do not have address duplications in the domain.
No chassis show up in the Available chassis box	Make sure that the chassis have been completely discovered.

Table 13. Standby blade pool creation and application problems

Problem	Solution
No blades show up in the available blades section.	Make sure that the blades and the BladeCenter chassis have been completely detected.
Only the blades from one chassis appear in the available blades section.	When creating a standby pool through a targeted action only the blades from that chassis will appear in the available blades section. To see all of the blades that have been detected, start BladeCenter Configuration Manager through an un-targeted action, such as clicking on the task.
Standby Blade Pool failed after manually applying to a blade.	• Check that the blades in the standby pool have the same model and type as the source blade. • Check that the standby blade pool does not contain only the source blade (a failover attempt to the same blade will always fail). • Check that the blades in the standby blade are powered off. • Make sure that the Director plug-in for your network switches has been installed.

Table 13. Standby blade pool creation and application problems (continued)

Problem	Solution
Standby Blade Pool failed to be applied after creating an event action plan and applying the event action plan to a blade.	• Check the logs to make sure the event that you are filtering on was actually triggered. • Make sure that the event action plan was applied to the correct blade. • Check that the event was sent to the blade object and not just to the BladeCenter chassis.

Related tasks

"Creating a standby blade pool configuration template" on page 39
You can create a standby blade pool configuration template using the BladeCenter Standby Blade Pool Configuration wizard.

Incorrect OFM address usage

If the blade is not using the OFM addresses as you expect, use the following steps to find the cause.

- Check the AMM Web page:
  - Select **Blade tasks** -> **Configuration** -> **Open Fabric Manager**.
  - Click on the name of the blade.
  - Verify that the OFM configuration is enabled.
  - Verify that the blade supports OFM.
  - Verify that the blade system management processor supports OFM.
    - If not, then upgrade the firmware of the Blade System Management processor.
  - Verify that the blade BIOS supports OFM.
    - If not, then upgrade the BIOS firmware.
- The OFM status for the blade is `N/A`:
  - If the blade is off:
    1. Power it on.
    2. Wait for the OS to reboot.
    3. Check the OFM status again.
  - If the OS is running:
    1. Reboot the blade.
    2. Check the OFM status again.
- Check the address status:
  - If the address status for the adapter is `Used` and the addresses are not actually used, replace the adapter.
  - If the address status for the adapter is `Error`, replace the adapter.
  - If the address status for the adapter is `Not Used`:
    - Verify that the address type matches the adapter type.

Configuration failure scenarios

You may encounter errors when applying your OFM configuration.

Parsing failures

If the configuration file contains errors while being parsed, the AMM Web interface displays an **Open Fabric Manager Configuration Failure**. The page contains a table that shows the line number of each error and a brief description of the error. You must fix all of the errors in the configuration file before continuing.

For a detailed description of these errors, see "Error messages" on page 46.

Connection failures

If the AMM is unable to connect to any of the chassis' that are listed in the configuration file, the Web interface displays an **Open Fabric Manager Configuration Failure**. The page contains a table that shows the IP address of the chassis it is unable to connect to. If you want to ignore this chassis and apply the configuration anyway, you can press **Ignore**. Ignoring a chassis might result in duplicate addresses, because the AMM cannot compare the addresses against the ones that are being used by the unreachable chassis. You can also suppress this warning by selecting the **Continue on error** advanced option on the **Apply a Configuration** page.

Login failures

The AMM uses the user name and password that you used to login to the current AMM to attempt to login to the other AMMs. If the AMM fails to login to the other AMMs using these credentials, the Web interface displays an **Open Fabric Manager Autentication Failure** page. This page gives you the option to provide an alternate user name and password for the chassis. If you check the box next to **use this user and password for the rest of the AMMs**, this user name and password will be attempted on all other chassis the AMM fails to connect to.

After you fill in the user name and password, click **Retry**. If you wish to ignore this chassis and continue to apply the configuration anyway click **Ignore**. If you wish to ignore all the chassis to which the AMM fails to login click **Ignore All**. Ignoring a chassis might result in duplicate addresses since the AMM cannot compare the addresses against the ones that are being used by the unreachable chassis. You can also suppress this warning by selecting the **Continue on error** advanced option on the **Apply a Configuration** page.

Important: You can get a **login failed** message in cases where all the TCP Command Mode connections to the AMM are in use. If you get the login failed message when using a valid user name and password, make sure that the AMM connection limit of the AMM you failed to login to is more than the number of connections in use (if, for example, Director is using one connection). This is controlled through the AMM Web interface on MM Control→Network Protocols→TCP Command Mode Protocol.

Retrieve failure

If the AMM fails to retrieve the OFM configuration from any of the chassis, the Web interface displays an **Open Fabric Manager Retrieve Failure** page. Failure to retrieve the OFM configuration usually occurs when the other chassis' do not support OFM or if there are communication errors.

If you want to ignore this chassis and apply the configuration anyway, you can press **Ignore**. Ignoring a chassis might result in duplicate addresses, because the AMM cannot compare the addresses against the ones that are being used by the unreachable chassis. You can also suppress this warning by selecting the **Continue on error** advanced option on the **Apply a Configuration** page.

Blade power state failure

Before applying a configuration, the AMM validates that all the blades that are about to be re-configured are powered off. If any of these blades are powered on, the Web interface displays an **Open Fabric Manager Configuration Error - Found Powered On Blades** page. The page contains a table that shows the blades that are currently powered on (indicated by a red X icon).

You can optionally force the configuration changes to these powered on blades by clicking **Force**. This will keep the blades in their powered on state and cause the configuration to be applied to those blades during their next reboot. Forcing the configuration might result in duplicate addresses. You can also force the configuration changes by selecting the **Force the configuration to be applied to powered on blades** advanced option on the **Apply a Configuration** page.

Disable OFM failure

As described in "Avoiding address duplication" on page 28, before applying the new configuration, the AMM disables the OFM configuration for the slots that are about to be re-configured. If there is a failure disabling the OFM configuration for any of the blades, the Web interface displays an **Open Fabric Manager Configuration Failure** page. The page contains a table that shows the blades that could not be disabled (indicated by a red X icon). A green check mark icon indicates blades that were properly disabled.

Important: This is a critical error so you must resolve this issue before continuing. If it is not resolved immediately, blades in your OFM environment can be disabled when that was not your intent.

Apply OFM failure

If there is a failure applying the OFM configuration to any of the slots, the AMM will continue to apply the OFM configuration to all the other slots. After the process completes, the Web interface displays a **The configuration file could not be successfully applied to one or more blades** page. The page contains a table that shows the blades to which the configuration could not be applied (indicated by a red X icon). If you apply the configuration again, the AMM will try to configure only the slots that it failed to configure on the previous attempt.

Error messages

This section describes error messages you might encounter when working with OFM.

Requirements reports errors

- AMM firmware needs to be upgraded in order to support Open Fabric Manager.
- Baseboard Management Controller (BMC) needs to be upgraded in order to support OFM.
- Blade BIOS needs to be upgraded in order to support OFM.

- Communication error to the blade or user doesn't have permission for OFM on a blade.
- HBA firmware needs to be upgraded to support OFM.

Parsing errors

This table contains a list of the parsing errors that can be generated when processing the OFM configuration file, with an explanation of the cause of each error.

Table 14. Parsing errors

Error message	Description
Assuming boot type `None`	The “Boot Type” field has been left empty – the system will assume a value of “none” and continue processing the line.
Assuming mode `Ignore`	The “Mode” field has been left empty – the system will assume a value of “Ignore” and continue processing.
Attempt to add multiple boot targets to a blade with the same priority	A single priority value for a target can only be used once for a given slot.
Attempt to redefine slot	There are multiple lines defining a slot. A slot must be defined before it is used and cannot be defined more than once.
Attempt to use a slot that has not yet been defined	A slot must be defined before it used. The OFM configuration file contains a reference to a slot that has not been defined when referenced.
Attempt to write non-Ethernet data to Ethernet port	The OFM configuration file includes an attempt to define a non-Ethernet port to an Ethernet-only port number (1 or 2).
Bad boot type	The “Boot type” string cannot be parsed.
Bad entry type	The entry type (One of: BladeCenter, Slot, Eth, FC, FCTarget) was not recognized.
Bad IP address	The IP address was incorrectly formed or, in the case of a hostname, the DNS lookup failed.
Bad offset number	Cannot parse the offset number, or the number is out of range. (Offsets must be in the range 0-3)
Bad port number	Cannot parse the port number, or the number is out of range. (Ports must be in the range 1-8 for Ethernet ports or 3-8 for all other types)
Bad slot number	Cannot correctly parse the slot number or it is out of range. The exact range depends on the BladeCenter type (BCH, BCS, BCHT), but is never less than 1 or more than 14. This error can also occur when the slot field does not have a closing comma.
Bad target priority - must be first or second	Cannot parse the target priority, or it was out of range (first/second). Targets can only be set as the first or second priority.

Table 14. Parsing errors (continued)

Error message	**Description**
Bad value in mac/wwpn/wnn	The value in the EUI field is correctly formed but contains characters that cannot be interpreted as a hexadecimal number. EUI format is a 16-character hexadecimal value with the leading prefix of *eui*, for example: eui:*0123456790ABCDEF*.
Bad vlan number	Cannot parse the VLAN number (not a number or similar problem).
Could not interpret the LUN value	The LUN value in the FCTarget was not correct.
Failure opening the configuration file	The OFM configuration file could not be opened: probably a bad file name or path, or a problem with file permissions.
Failure reading the configuration file	The OFM configuration file could not be read until the end.
Incomplete Line	This line is missing required fields.
Input line is too long	The maximum length for a single line in the OFM configuration file is 512 characters. Lines longer than this will be discarded.
Insufficient addresses in the range for chassis	The range of addresses defined for this OFM configuration file is not sufficient for the number of chassis required.
Invalid integer	Cannot parse an integer number.
mac or wwnn field is too short or too long	Too few or too many bytes in an EUI field.
Maximum value is 0xfffffffe - for longer values use EUI notation	LUN values can be specified in decimal, hexadecimal or EUI formats. Values above 4294967295 (0xffffffff) must be specified as an EUI.
No closing quote	One of the fields on this line is missing a closing quote. To prevent problems the line will be ignored.
Profile is too long and has been truncated	The profile string cannot be longer than 32 characters.
Second reference to a specific port or target	A port or target has been reused with the same type (for example, Eth, FC, etc.). Ports and targets can only be defined once for a particular type.
Too many BladeCenters	The maximum number of BladeCenters that can be processed from a single OFM configuration file is 100.

Chapter 8. Reference

This section provides reference information for OFM, including accessibility, the CLI command and events.

Accessibility features for BladeCenter Open Fabric Manager

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The BladeCenter Open Fabric Manager plug-in for IBM Systems Director supports the accessibility features that are supported in IBM Systems Director.

Session and credentials

The AMM displays an error page when another user attempts to login with the same credentials.

For most of the OFM operations, the AMM needs to interact with other chassis. By default, it will use the user name and password that you used to login to the current AMM. If the AMM cannot login to other chassis, it will prompt you to provide an alternative user name and password. The alternative user name and password are stored for the local session and are not cleared until you log out. Thus, while you are using OFM, another user cannot. If another user tries to open the OFM page while it is in use by you, the Web interface displays an **Open Fabric Manager is in Use By Another Session** page. If the new user clicks **Continue**, the your operation will be aborted.

Note: This page is displayed only in the case where a new user starts OFM while you have started a OFM operation but did not complete it. If you complete your last operation, this page is not displayed and the credentials are cleared automatically.

Standby AMM

Because an AMM failure results in a configuration loss, IBM recommends that you install a Standby AMM when using OFM.

The OFM configuration is not included in the AMM configuration backup. Such a backup would allow OFM configuration updates from a file that might be outdated and would cause inconsistencies with the configuration of other chassis.

The OFM configuration is chassis-based and will not transfer with the physical AMM. When an AMM is moved to a new chassis, it clears out its OFM configuration, and the OFM configuration must be reapplied on the new AMM.

Note: When an AMM is moved to a new chassis, it will clear the configuration for all slots except for slots that have blades powered-on with a valid OFM configuration that is already in use. For these slots, the AMM will use the OFM configuration as defined on the blades.

If the AMM configuration is reset to factory defaults, the OFM configuration is handled as if a new AMM was inserted to the chassis.

If the primary AMM fails, the standby AMM will contain the OFM configuration and will take over. As an additional precaution, you should always save your OFM configuration file or files via the OFM interface whenever changes are made or a new configuration is applied.

If you only have one AMM installed in the chassis and must replace it, before you do the replacement, insert the new AMM as a standby unit, let it power on and wait for a few minutes. This will allow the OFM configuration to be transferred to the new unit. The old unit can then be removed and you will not have to reapply the OFM configuration.

Note: You might have to wait longer than a few minutes if the AMM you insert in the standby slot does not have the same firmware level as the primary.

CLI command

You can apply a new configuration using the **bofm** CLI command. To use the command, you must have a tftp server available for uploading the configuration file.

Format

bofm -d [on|off] | -i <ip> |-l <file name> | -p [on|off] | -v

Options

-d: check duplicate [on|off]

-i: ip of tftp server

-l: configuration file name

-p: check blade power state [on|off]

-v: verbose mode

Usage

You will need to supply the tftp server IP address in the –i option, and the configuration file name in the –l (similar to the **update** CLI command). Adding the **–d off** option will cause OFM to ignore duplicate addresses in the configuration file. Adding the **–p off** option causes the OFM to apply the configuration even on powered-on blades. These parameters default to "on".

Events

There are a few OFM-related events:

1. If the OFM configuration changes for a slot, then an information event is generated for that slot.
2. If a slot has OFM configuration enabled and the system management processor of the blade in that slot does not support OFM, then a warning event is generated.
3. If a slot has OFM configuration enabled and the blade's BIOS in that slot does not support OFM, then a warning event is generated.
4. If a slot has OFM configuration enabled and any device in that slot does not support OFM, thean a warning event is generated.
5. If any of the adapters reported a OFM error, then a warning event is generated.

6. If the AMM was unable to apply a OFM configuration to a blade for which the system management processor supports OFM, then the blade is not given power permission and an error event is generated for that blade.
7. If the AMM cleared the OFM configuration after moving to a new chassis, then a warning alert is generated.
8. If the AMM discovered that a OFM configuration is in use by a blade that is powered on and it is different from the configuration defined on the AMM, then a warning event is generated.

User interface

This section describes the pages in the BladeCenter Open Fabric Manager interface.

BladeCenter Open Fabric Manager summary

Use the BladeCenter Open Fabric Manager summary page to configure and manage the Standby Blade Pool Configuration and Blade Address Manager.

Management

This section displays the status of the Standby Blade Pool Configuration and Blade Address Manager.

- In the **Common Tasks** area, the following links are provided:

  Event action plans

  Click to access the Event Automation Plan Wizard, which allows you to create and edit event automation plans. You can use event automation plans to designate events within your systems-management environment for which one or more actions are executed.

  System Discovery

  Click to access the System Discovery task, which allows you to discover resources in your environment.

  Advanced System Discovery

  Click to access the Advanced System Discovery task, which allows you to discover specific types of resources using discovery profiles.

Configuration

This section displays the number of the Standby Blade Pool Configuration templates and Blade Address Manager templates. Click the **Create...** link to use a wizard to create a template.

- In the **Common Tasks** area, the following links are provided:

  Configuration templates

  Click to view and manage existing configuration templates. You can use configuration templates to deploy settings on one or more systems.

  Configuration plans

  Click to view and manage existing configuration plans. You can use configuration plans to deploy a sequence of configuration templates across one or more systems.

  Current configuration

  Click to view the configuration settings for selected resources.

Import permanent license

Once you install BladeCenter Open Fabric Manager, you are granted a

60-day evaluation license. When the evaluation license expires, you must purchase a permanent license to continue using BladeCenter Open Fabric Manager. To install a permanent license key, click the **Import permanent license** link at the top of the BladeCenter Open Fabric Manager summary page.

Note: The 60-day evaluation period begins once BladeCenter Open Fabric Manager has been installed and IBM Systems Director Server has been restarted. The day the evaluation period will expire, as well as the number of days remaining on the license, is displayed at the top of the BladeCenter Open Fabric Manager summary page.

Blade Address Manager Configuration wizard

Use the Blade Address Manager Configuration wizard to create a Blade Address Manager template. A Blade Address Manager template can be used to start the BladeCenter Address Manager.

Complete the following wizard pages to create a Blade Address Manager template:

1. **Create template**

   Select the target type and specify a name for the template. You can also specify a description of the template.
2. **Welcome**

   Click **Next**.
3. **Import**

   Select whether you want to import an existing Blade Address Manager .csv configuration file. Importing an existing Blade Address Manager configuration file will populate this wizard with the settings and values from the imported configuration file. This can be used as a starting point to edit, then create a new configuration file based on the imported configuration.
4. **Member**

   Select the BladeCenter chassis that will participate in the domain by clicking on the chassis in the **Available Chassis** box. Click **Add** to move the selected chassis to the **Domain Members** box. Chassis not listed in the **Available Chassis** box can be added to the domain by clicking **New...** and entering the IP address of the chassis.
5. **Enable**

   Select the checkbox next to a chassis to enable the blade address configuration for the selected chassis. Unselect the checkbox next to a chassis to deactivate the blade address configuration for the selected chassis.
6. **Address Pool**

   Complete the following fields:

   Maximum chassis in domain

   Enter the maximum number of chassis to include in the domain. The maximum size is 100. The software will pre-allocate address space for the maximum number of chassis chosen to ensure that there will be no address duplications in any of the templates.

   Automatically assign addresses

   Activate MAC

   To assign Ethernet MAC addresses to bays, click the **Activate MAC** check box. Then select **IBM** in the **Vendor:** dropdown list to allow the software to automatically assign addresses.

Selecting **User Defined** in the Vendor dropdown list will allow you to edit the MAC address range to any range of your choosing.

Some Ethernet expansion cards can have a range of MAC addresses per port. The range is defined by specifying two MAC addresses per port: MAC A and MAC B. For example, the JS22 and JS12 onboard HEA Ethernet adapter must use the default values shown below. To set up a range of MAC addresses for a port, select **User defined** from the **MAC addresses per port:** dropdown list. Then specify a value from 1 to 255 in the **Range:** field.

Note: The above default values will allow the JS22/JS12 onboard HEA Ethernet adapter to use virtual addresses. Even though the 16 MAC range for the onboard HEA adapter is declared via port 1, the actual routing of the 16 MACs is determined by configuration in the OS.

Activate FC

To assign Fibre Channel WWN addresses to bays, click the **Activate FC** check box. For **Vendor** select either **Qlogic** or **Emulex**, depending on the type of fibre channel cards in your environment, to allow the software to automatically assign addresses. Selecting **User Defined** will allow you to edit the Fibre Channel address range to any range of your choosing.

Valid address range for **Qlogic**:

- First port: 21:80:00:E0:8B:0X:XX:XX
- Second port: 21:81:00:E0:8B:2X:XX:XX

Valid address range for **Emulex**:

- WWNN: 2F:FE:00:00:C9:XX:XX:XX through 2F:FF:00:00:C9:XX:XX:XX
- WWPN: 2F:FC:00:00:C9:XX:XX:XX through 2F:FD:00:00:C9:XX:XX:XX

Activate SAS

To assign storage target WWPN addresses to bays, click the **Activate SAS** check box. Then select the **Vendor** type and enter the range of addresses.

Valid address range for **IBM**:

- 50:05:07:60:1A:80:00:00h to 50:05:07:60:1A:BF:FF:FFh

Valid address range for **LSI**:

- 50:00:62:B0:00:11:17:00h to 50:00:62:B0:00:12:16:FFh

7. **Boot Settings**

Select the primary and secondary ports of the storage system where the boot LUNs for the blades reside. Additionally, select the blade assigned LUN for the boot LUN. This will be the LUN that the blade will be booting from so it will need to have an operating system installed. These settings will be applied to every blade in the domain, but can be modified on a blade by blade basis on the Advanced Configuration panel of this wizard. If the storage system was discovered using Advanced Discovery, then the primary and secondary fields will show the WWPN information of the storage system.

8. **Advanced Options**

You can select to ignore three different types of errors when the template is applied. The types of errors are:

Ignore duplicate check when applying
Select this option to have BladeCenter Open Fabric Manager ignore errors occurring when a template containing duplicate addresses is applied to a domain of blades.

Apply the blades even if they are powered on
Select this option if you have blades in your domain that are currently powered on and you are changing their addresses. The new addresses will not be active until the next time that the blade is powered on. If this box is unchecked, then the template will fail when applied if there are blades in the domain which are currently powered on.

Continue to apply even if an error occurs
Select this option to have BladeCenter Open Fabric Manager ignore various errors, such as the inability to log into another chassis, and continue to apply addresses to blades and chassis that are accessible.

9. **Advanced Configuration**

   You can select to:

   - Assign addresses from the pool
   - Manually override the default boot settings

   .

10. **Advanced Settings**

    If you selected to manually override the default boot settings, select the override options on the **Advanced Settings** page. You can also manually configure boot settings for each bay.

    a. Select a chassis in the **Select Chassis** box by clicking on the IP address of the chassis.

    b. Select a bay to edit in the **Select Bay** box by clicking on the bay number.

    c. Under the **Select Chassis** box you can select to:

       Activate Bay
       Assigns addresses to the blade connected to that bay.

       Deactivate Bay
       Removes any addresses assigned to that bay and any blade connected to that bay will revert to their physical addresses.

       Ignore Bay
       Leaves the addresses in their previous settings. The addresses will remain either virtual or physical when this template is applied. The **Ignore Bay** setting is useful when two templates are created for a single chassis, such as when both Emulex and Qlogic adapter cards are in the same chassis.

    d. Select the type of setting you would like to configure, then click **Edit...**.

    e. In the **Fibre Channel Boot Settings** and **SAS Boot Settings**, you can specify the WWN and LUN number for the boot target. Then blade can then be booted from those addresses.

11. **Summary**

    The Summary page displays the settings that you have selected. The hardware listed represents the chassis that you included in the domain members list.

Ensure that the template options have been set correctly. To save the configuration template to a .csv file, click the **Export** button in the upper-right corner of the page.

Select **Deploy settings when finished** to apply the settings of this template to the hardware upon clicking **Finish**.

Ethernet Settings

Use the **Ethernet Settings** page to enable the virtualization of the MAC addresses.

To enable the virtualization of Ethernet MAC addresses:

1. Select the **Activate** checkbox.
2. Select the Ethernet port. To remove a port, select the port to be removed and click **Remove**. To add a port, enter the port number and click **Add**. The valid range of MAC ports is from 1 to 8.
3. In the **Blade MAC1** field, enter the MAC address in the form: xx:xx:xx:xx:xx:xx. Each byte must be a valid hex value.

   Note: Some Ethernet expansion cards can have a range of MAC addresses per port. The range is defined by specifying two MAC addresses per port in both the **Blade MAC2** and **Blade MAC2** fields.
4. You can also configure VLANs for the Ethernet ports. These VLAN settings are used during pre-booting to an operating system, such as when the blades are booting through PXE. The default is 0. The value can be entered as a decimal, octal, or hexadecimal number. Values starting with a 0 (for example, 077) are assumed to be octal. Values starting with a 0x (for example, 0xa) are assumed to be hexadecimal. The VLAN's maximum value is 65535.

Fibre Channel Settings

Use the **Fibre Channel Settings** page to enable the virtualization of Fibre Channel addresses.

To enable the virtualization of Fibre Channel addresses:

1. Select the **Activate** checkbox.
2. Select the Fibre Channel port. To remove a port, select the port to be removed and click **Remove**. To add a port, enter the port number and click **Add**. The valid range of Fibre Channel ports is from 3 to 8.
3. In the **Blade WWNN** field, enter the World Wide Node Name (WWNN). This value must conform to the EUI standard for 64-bit values (hexadecimal number): xx:xx:xx:xx:xx:xx:xx:xx. The first two bytes are either hex 10:00 or 2x:xx. Or the first half-byte(x) is either hex 5 or 6.
4. In the **Blade WWPN** field, enter the World Wide Port Name (WWPN). This value must conform to the EUI standard for 64-bit values (hexadecimal number): xx:xx:xx:xx:xx:xx:xx:xx. The first two bytes are either hex 10:00 or 2x:xx. Or the first half-byte(x) is either hex 5 or 6.
5. In the **Priority** list, select the boot priority.

SAS Settings

Use the **SAS Settings** page to enable the virtualization of Serial Attached SCSI (SAS) addresses.

To enable the virtualization of SAS addresses:

1. Select the **Activate** checkbox.

2. Select the SAS port. To remove a port, select the port to be removed and click **Remove**. To add a port, enter the port number and click **Add**. The valid range of SAS ports is from 3 to 4.
3. In the **Blade WWNN** field, enter the World Wide Node Name (WWNN). This field must conform to the EUI standard for 64-bit values (hexadecimal number): xx:xx:xx:xx:xx:xx:xx:xx.
4. In the **Priority** list, select the boot priority.

BladeCenter Standby Blade Pool Configuration wizard

Use the BladeCenter Standby Blade Pool Configuration wizard to create a Standby Blade Pool Configuration template.

Blades Pool

On the **Blades Pool** page, select the chassis bays to be used for failover in the **Available Bays** box, then click **Add >**. These bays will be tried in the order they appear in the **Selected Bays** box. Use the **Up** and **Down** buttons to specify the bay order in the **Selected Bays** box, then click **Next**.

Properties

On the **Properties** page, select the failover options to be used by the standby blade pool profile, then click **Next**.

Summary

The **Summary** page displays the settings that you have selected. Ensure that the profile options have been set correctly. Select **Deploy settings when finished** to apply the settings of this template to the hardware upon clicking **Finish**.

Note: Once a standby blade pool has been created, the list of standby blades will remain in the profile even after one or more of the blades has been removed from the BladeCenter chassis.

Chapter 9. Publications and related information

You can view the same BladeCenter Open Fabric Manager content that resides in the Information Center in a PDF document. To view a PDF file, you need Adobe Acrobat Reader, which can be downloaded for free from the Adobe Web site at www.adobe.com/products/acrobat/readstep.html.

BladeCenter Open Fabric Manager resources on the World Wide Web

- **BladeCenter Open Fabric Manager Web site**

  www.ibm.com/systems/bladecenter/hardware/openfabric/openfabricmanager.html

  Obtain an overview of BladeCenter Open Fabric Manager and links to download the product and user's guide.
- **IBM Systems Director Web site**

  www.ibm.com/systems/software/director

  Get overview information, demos, and downloads for the IBM Systems Director product, as well as IBM Systems Director plug-ins.
- **IBM Systems and servers: Technical support page**

  www.ibm.com/systems/support/

  Locate support for IBM hardware and systems-management software.
- **IBM ServerProven compatibility**

  www.ibm.com/systems/info/x86servers/serverproven/compat/us/

  Obtain compatibility information about IBM System x® and IBM BladeCenter.

Notices & Trademarks

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following are trademarks of International Business Machines Corporation in the United States, other countries, or both:

BladeCenter

Netfinity®
eServer
IBM
IntelliStation®
ServerProven®
xSeries®

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

Terms

This topic provides definitions of terms that are used in the OFM product.

This glossary defines technical terms and abbreviations used in this IBM PowerExecutive™ document. If you do not find the term you are looking for, view the IBM Terminology Web site at: http://www.ibm.com/ibm/terminology.

Selection of Terms: A term is a word or group of words to be defined. In this glossary, the singular form of the noun and the infinitive form of the verb are the terms most often selected to be defined. If the term may be abbreviated, the abbreviation is indicated. The abbreviation is also defined in its proper place in the glossary.

A

available blades
This is a collection of blades that can you can add to a standby blade pool. This section will either contain all the blades in a chassis or all the blades that have been discovered. To see all the blades that have been discovered, run BladeCenter Configuration Manager un-targeted by double clicking on the task. To see only the blades in a chassis run BladeCenter Configuration Manager targeted by dragging the task onto a BladeCenter chassis.

B

blade

blade server
An IBM BladeCenter server. A high-throughput server on a card that supports symmetric multiprocessors (SMP).

BladeCenter chassis
A BladeCenter unit that acts as an enclosure. It enables the individual blade servers to share resources such as the management, switch, power and blower modules.

BladeCenter Open Fabric Manager
An IBM solution that enables you to quickly deploy, replace and recover blades in your environment.

C

chassis
The metal frame in which various electronic components are mounted.

D

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved.

E

event
An occurrence of significance to a task or system, such as the completion or failure of an operation. There are two types of events: alert and resolution.

event action
The action that IBM Systems Director takes in response to a specific event or events

event-action plan
A user-defined plan that determines how IBM Systems Director manages certain events. An event action plan comprises one or more event filters and one or more customized event actions.

I

IBM Systems Director environment
The complex, heterogeneous environment managed by IBM Systems Director. It includes systems, BladeCenter chassis, software, Simple Network Management Protocol (SNMP) devices.

IBM Systems Director plug-in
A tool that extends the functionality of IBM Systems Director. Some of the IBM Systems Director extensions are Active Energy Manager, BladeCenter Open Fabric Manager, BladeCenter Management, Capacity Manager,

ServeRAID Manager, and Remote Deployment Manager.

IBM Systems Director Server
The main component of IBM Systems Director software. When installed on the management server, it provides basic functions such as discovery of managed systems, persistent storage of configuration and management data, an inventory database, event listening, security and authentication, management console support, and administrative tasks.

M

managed group
A group of systems or objects managed by IBM Systems Director.

managed object
An item managed by IBM Systems Director. In IBM Systems Director Console, a managed object is represented by an icon that shows its type (such as chassis, cluster, system, or scalable system).

managed object ID
A unique identifier for each managed object. It is the key value used by IBM Systems Director database tables.

managed system
A system that is being controlled by a given system management application, for example, a system managed by IBM Systems Director.

management module
The BladeCenter component that handles system-management functions. It configures the chassis and switch modules, communicates with the blade servers and all I/O modules, multiplexes the keyboard/video/mouse (KVM), and monitors critical information about the chassis and blade servers.

management server
The server on which IBM Systems Director Server is installed.

O

OFM See *BladeCenter Open Fabric Manager*.

P

plug-in
A software module that adds function to an existing program or application.

S

source blade
The blade that is currently running an operating system and is currently active. It is the source blade that will be associated with an event action plan.

standby blade
This is a cold spare blade that will be powered on in the case of a failure of the source blade and will take the place of the source blade.

standby blade pool
A collection of designated cold spare blades. These blades will have the same connectivity to the infrastructure as the source blades, but they will remain off until they are called upon to take over for a source blade.

standby blade pool template
A task that is created by the user and located under the BladeCenter Configuration Manager task. It contains the Standby Blade Pool and is applied to a source blade to trigger a failover.

system
The computer and its associated devices and programs.

T

target system
A managed system on which an IBM Systems Director task is performed.

W

World Wide Name (WWN)
A unique identifier in a Fibre Channel or Serial Attached SCSI storage network.

Index

A

accessibility
- features 3, 49
- keyboard 3, 49
- shortcut keys 3, 49

adding a new chassis to the domain 33
AMM Web interface 7
- overview 13

applying a new configuration 30
avoiding address duplication 28

B

Blade Address Manager Configuration template
- applying 39

Blade Address Manager template
- creating 35, 52

BladeCenter section of the configuration file 7
BladeCenter Standby Blade Pool Configuration wizard 39
BladeScanner 16
bofm CLI command 50
boot from SAN 19

C

ChassisUpdate 16
CLI 50
comments section of the configuration file 7
configuration file 7, 45
- BladeCenter section 7
- comments section 7
- create automatically 25
- other format considerations 7
- Ports section 7
- selecting domains 28
- Slots section 7

creating a configuration file automatically 25
creating a requirements report 29
Creating an events action plan 41
CSV file example 12

D

deploy
- standby blade pool configuration template 40

disability 3, 49
duplicate addresses 29

E

editing the configuration file 30
error messages 46
events 50
example
- configuration file example 12
- CSV file 12

F

failover from one slot to another 34
failure scenarios 45
features, accessibility 3, 49

G

glossary 63

H

hardware
- BladeCenter chassis 4
- BladeCenter servers 4
- expansion cards 4
- fibre channel 4
- network 4
- passthru 4
- requirements 3

hardware requirements 3

I

incorrect OFM address usage
- troubleshooting
  - OFM address usage 44

initial deployment 33
installation requirements
- disk space 3
- memory 3
- requirements
  - disk space 3
  - memory 3

installing
- OFM standalone version on Windows 22

installing IBM Systems Director plug-in version of OFM
- on Windows 21

installing Systems IBM Director plug-in version of OFM
- on Linux 22

installing the OFM standalone version for Linux 22

K

keyboard 3, 49

L

legal notices 59
licensing
- requirements 3

M

mapping devices to ports 12
multi-slot blades 13

O

offset parameter 13
OFM components 7
OFM configuration
- failure scenarios 45

OFM session and credentials 49
OFM-related events 50

P

parse errors 46
Planning for OFM 15
Ports entries section of the configuration file 7
Preparing for OFM 15

R

reference 49
replacing a blade in the same slot 34
replacing AMM IP addresses 34
replacing the AMM in a single AMM environment 34
requirements
- hardware 3
- licensing 3
- software 3

requirements report 29
retrieving the current configuration 32

S

scenario
- adding a new chassis to the domain 33
- failover from one slot to another 34
- initial deployment 33
- replacing a blade in the same slot 34
- replacing the AMM in a single AMM environment 34

scenarios
- replacing AMM IP addresses 34

selecting domains 28
session and credentials 49
shortcut keys 3, 49
Slots section of the configuration file 7
software
- requirements 3

Standby AMM 49

standby blade pool configuration template
  deploying 40
Standby Blade Pool Configuration wizard 56
standby blade pool problems 43
standby blade pool template
  creating 39
summary page 51
supported hardware 4
supported software 7

T

template
  Blade Address Manager 35, 52
trademarks 60
troubleshooting and support 43
troubleshooting OFM-Advanced upgrade 43

U

UpdateXpress for BladeCenter 16
updating AMM firmware 15
updating AMM firmware using BladeScanner 16
updating AMM firmware using ChassisUpdate 16
updating AMM firmware using UXBC 16
updating blade BIOS 18
updating blade BMC firmware 17
updating Emulex firmware for POWER PC 19
updating Emulex HBA firmware for x86 18
updating firmware 15
updating QLogic firmware for POWER PC 19
updating QLogic firmware for x86 18
UXBC 16

V

viewing the configuration
  in a local chassis 31

Readers’ Comments — We'd Like to Hear from You

IBM Systems
IBM BladeCenter Open Fabric Manager
Installation and User's Guide
Version 3.1

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name	Address
Company or Organization	
Phone No.	E-mail address

Readers' Comments — We'd Like to Hear from You

Cut or Fold
Along Line

Fold and Tape

Please do not staple

Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Dept. CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990

Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold
Along Line

Printed in USA