



How to Failover and Deploy a Power Based Blade Server with BladeCenter Open Fabric Manager

February 2011

Trina Bunting
Power Systems Advanced Technical Skills
© Copyright IBM Corporation, 2011

Revision History

1.0 - Initial Version 2/9/2011

1.1 - Made some minor corrections 7/18/2011

Special Thanks

A special thanks to Khalid M. Ansari from System x ATS and Kerry Anders from STG Lab Services for your input!

Notices:

This paper is intended to provide information regarding Choosing the Right Switch Module for the POWER processor-based blade environment. It discusses findings based on configurations that were created and tested under laboratory conditions. These findings may not be realized in all customer environments, and implementation in such environments may require additional steps, configurations, and performance analysis. The information herein is provided "AS IS" with no warranties, express or implied. This information does not constitute a specification or form part of the warranty for any IBM or non-IBM products.

Information in this document was developed in conjunction with the use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

The information contained in this document has not been submitted to any formal IBM test and is distributed **as is**. The use of this information or the implementation of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

IBM may not officially support techniques mentioned in this document. For questions regarding officially supported techniques, please refer to the product documentation, announcement letters, or contact the IBM Support Line at 1-800-IBM-SERV. This document makes reference to vendor-acquired applications or utilities. It is the customer's responsibility to obtain licenses of these utilities prior to their usage.

® Copyright International Business Machines Corporation 2011. All rights reserved. U.S. Government Users Restricted Rights

-Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Table of Contents

I.	Introduction	4
II.	Prerequisite	4
	A. BladeCenter Hardware/Software Configuration	4
	B. SAN Storage	5
III.	BOFM Overview	7
	A. Basic BOFM	8
	B. Advanced BOFM	8
	C. BOFM Licensing	9
IV.	Configuring BOFM	10
	A. Configuring the AMM	10
	B. BOFM GUI	11
	C. Discovering the Chassis	12
V.	Advanced BOFM Templates	14
	A. Address Manager Template	15
	B. Creating Standby Blade Pool	19
	C. Creating Failover Monitor	22
VI.	Testing Failover on the Blade	26
	A. Backup the Blade	26
	B. Manual Failover	27
	C. Automatic Failover	28
	D. Restoring the Blade Configuration	28
VII.	Configuring Rip & Replace from the AMM	30
	A. Creating the Profile	30

I. Introduction

In this whitepaper we discuss the steps I followed to failover and deploy a virtualized Power based Blade Server environment with IBM® BladeCenter® Open Fabric Manager (BOFM). We discuss the differences between BOFM Basic and BOFM Advanced as well as the caveats involved in using BOFM to deploy Power based Blade Servers. We learn how to configure event monitors, automatic failover and deployment of the new blade. We review the steps involved in restoring the LPAR Configuration and N_Port ID Virtualization (NPIV) mappings on the blade. We also discuss how to configure Rip & Replace from the Advanced Management Module (AMM). It is important to note the blades should be in a power off state before you configure BOFM. You can force the configuration while the blades are powered on but in some cases it can cause duplicate address conflicts.

This whitepaper assumes you are familiar with BOFM. For more detailed information on how to implement BOFM go to <http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101691>.

II. Prerequisite

The following prerequisites are required to implement BOFM on the POWER7 based Blades:

- A license key for BOFM Basic and Advanced is required and should be activated on the AMM under License Manager
- Advanced BOFM v4.0 Stand-alone version must be installed on a server running Windows or Linux on Power (requires 10Mb of disk space)
- Sun Java SE 1.6 or higher
- The POWER7 based Blades must be configured for SAN Boot
- The hardware used on the blades i.e. adapters, firmware levels, etc. must be compatible with BOFM v4.0

Note: Advanced BOFM software is NOT currently supported running on a server with AIX installed and is only supported with servers running windows or Linux on Power. BOFM can manage AIX servers but the Advanced BOFM software can not be installed on that operating system.

For more information on installing BOFM refer to the Installation and User's Guide at: ftp://ftp.software.ibm.com/systems/support/system_x_pdf/49y0230.pdf.

A. BladeCenter Hardware/Software Configuration

It is important to validate the hardware (i.e. adapters, blades, etc.), firmware and software levels to ensure support with BOFM. For more information on the supported BOFM

environment go to <http://www-03.ibm.com/systems/bladecenter/hardware/openfabric/openfabricmanager/index.html>.

Note: It is important to note the SAS adapter does not support BOFM on the POWER7 based Blades.

The following hardware configuration including firmware levels as well as software configuration was used in my lab for testing BOFM:

My Hardware/Firmware Levels

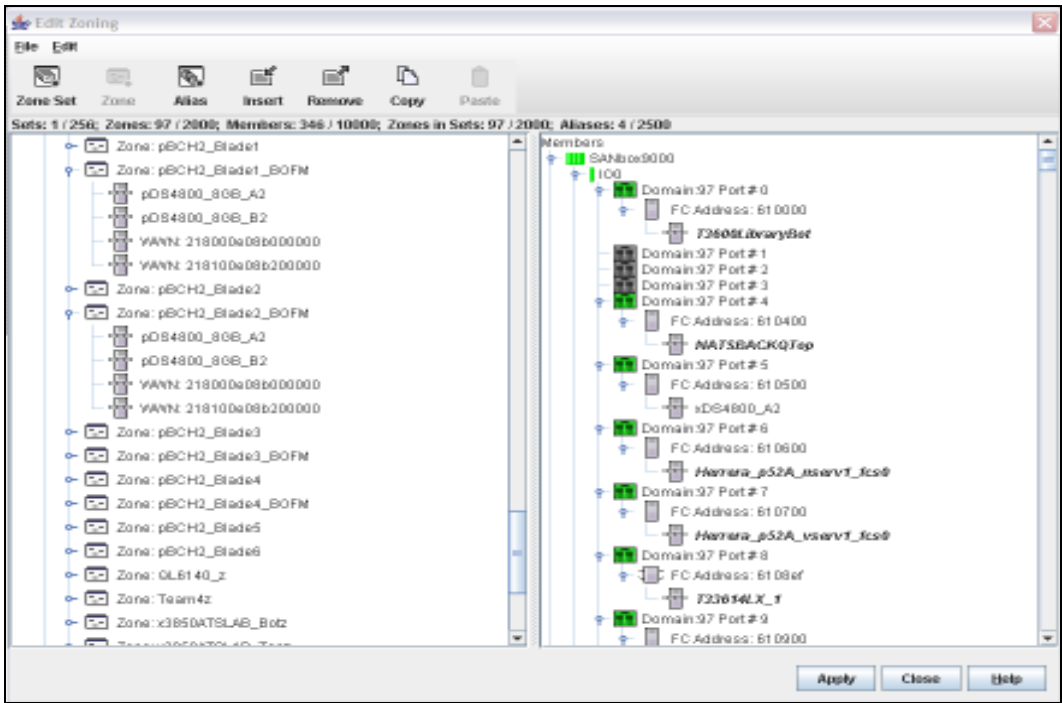
- Intel Blade – HS22
- Two PS700 Blades (Blade #1 is configured as the Primary blade and Blade #2 is configured as a backup)
 - Blade firmware level AA710-083
 - 8Gb QLogic Fibre Channel Adapter with firmware level 7710322577107601.039050201
- BladeCenter H Chassis
 - AMM Firmware level BPET54P
 - Two QLogic 8Gb Fibre Channel Switch Modules (Bays 3 and 4); firmware level BRFSM, Rev 7a14
- Both POWER7 Blades are configured to boot from SAN Storage - DS4800

My Software

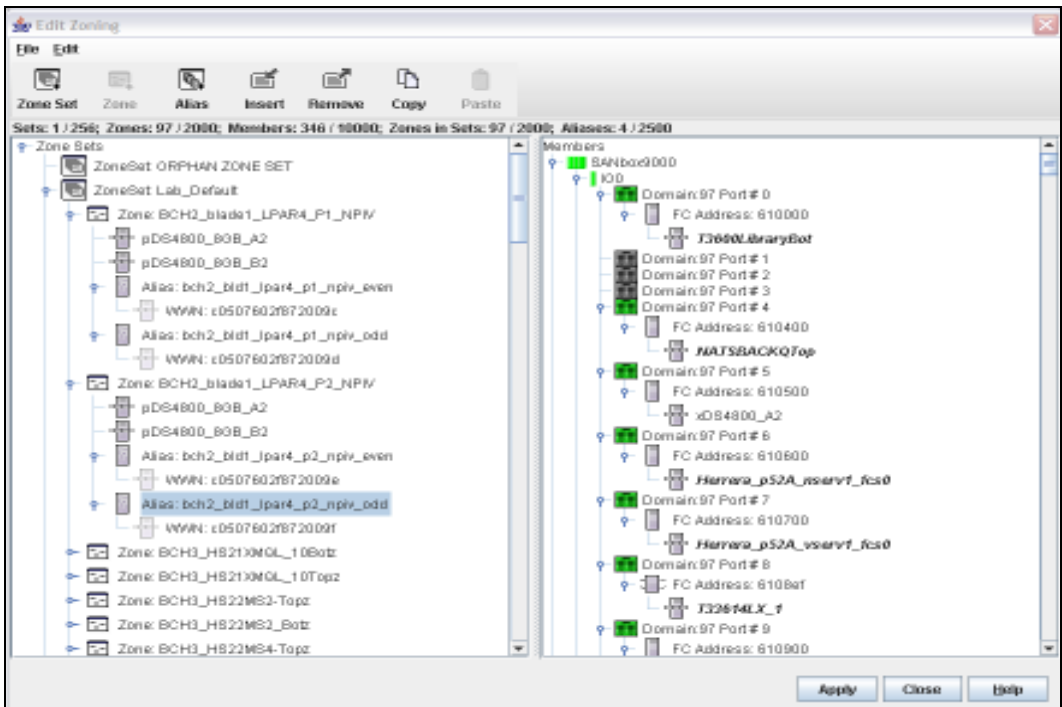
- Advanced BOFM v4.0 Server; Windows 2008 R2
- Integrated Virtualization Manager (IVM) 2.2 FP 24 installed
 - NPIV enabled LPAR with AIX 6.1 TL5 installed
 - vSCSI enabled LPAR with no Operating System

B. SAN Storage

When configuring BOFM on your blade the best practice is to use external storage to boot from. If a hardware failure is experienced on the blade, then only the failed component has to be replaced while the system image resides on the SAN attached storage. For my test environment I attached my blades to the DS4800 Storage and removed the internal disk drives. The following is an example of how I zoned my blade for BOFM. The WWNs in the “**p6BCH2_Blade1_BFM**” zone are for the 8Gb physical adapter.

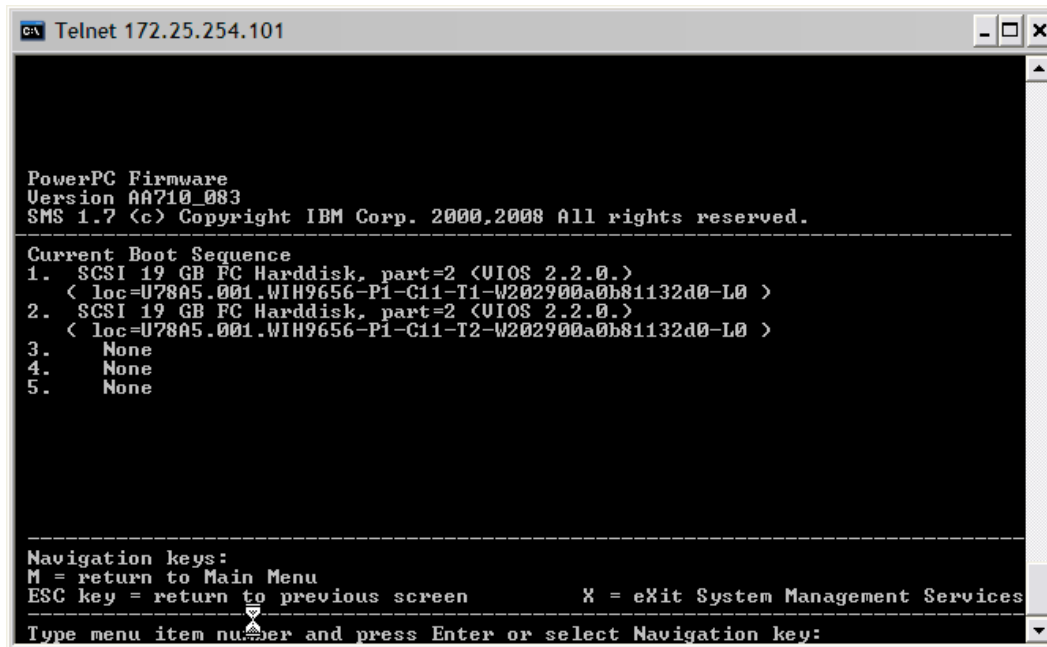


The “BCH2_blade1_LPAR4_P1_NPIV” and “BCH2_Blade1_LPAR4_P2_NPIV” zones (one for each port) were created for my NPIV enabled LPAR:



Once the zoning is completed the storage LUNs can be created. The blade should then be rebooted to the SMS Menu in preparation for installation. It is important to note that both

paths need to be enabled to boot the VIO server. From the SMS menu select both paths for the storage.



```
Telnet 172.25.254.101

PowerPC Firmware
Version AA710_083
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.

-----
Current Boot Sequence
1. SCSI 19 GB FC Harddisk, part=2 (UIOS 2.2.0.)
   < loc=U78A5.001.WIH9656-P1-C11-T1-W202900a0b81132d0-L0 >
2. SCSI 19 GB FC Harddisk, part=2 (UIOS 2.2.0.)
   < loc=U78A5.001.WIH9656-P1-C11-T2-W202900a0b81132d0-L0 >
3. None
4. None
5. None

-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen      X = eXit System Management Services

Type menu item number and press Enter or select Navigation key:
```

III. BOFM Overview

BladeCenter Open Fabric Manager is a tool used for simplifying deployment, failover and repurposing of blade servers. BOFM allows you to manage virtual addresses (MAC, WWN) and virtual storage (address and LUN of target storage devices) for up to 256 chassis and 3584 blades, as well as failover virtual I/O. BOFM does virtualization from open firmware on the blade and the AMM Firmware. The following versions of BOFM are currently supported:

- BOFM Basic
- BOFM for BladeCenter S (BCS)
- BOFM Advanced Director Plug-in v3.2
- BOFM 4.0 Stand-alone

Basic BOFM and BOFM for BCS functionality are built-in to the Advanced Management Module (AMM). BOFM 4.0 Stand-alone is a GUI based tool which should be installed on a stand-alone server running windows or Linux on Power. BOFM Advanced Director Plug-in v3.2 is currently still available but will be withdrawn this year and will no longer be updated. When planning to install BOFM Advanced the Best Practice is to install it on an external server instead of putting it on a blade that is the same chassis that is being managed by BOFM.

A. Basic BOFM

The Basic version of BOFM is built-into the AMM and allows you to do the following:

- Add Virtualization (i.e. configure chassis, etc.)
- Configure SAN Boot Settings
- Rip & Replace

The Rip & Replace feature of BOFM allows you to configure slot based I/O Address assignment. For example when BOFM is configured on a blade in slot 1 of the chassis (blade #1) and the blade in that slot goes down, that blade can be removed from the chassis and a new blade can be inserted in that slot. When the new blade is booted it will have the same IP Address, WWN, etc as the old blade.

B. Advanced BOFM

The Advanced version of BOFMv4.0 is a GUI based stand-alone product that is not supported as a Systems Director plug-in. It is important to note that Basic BOFM is a prerequisite for using Advanced BOFM. Therefore a license for basic and advanced BOFM should be purchased for every chassis. Advanced BOFM can be installed on a server running windows or Linux on Power and it supports the following functions:

- Failover monitoring and deployment
- Manual Failover
- Automatic Failover

From Advanced BOFM you can discover the chassis as well as create profiles for failover. The Best Practice is to run Advanced BOFM v4.0 from a sever that is not being managed by BOFM, instead of running it from one of the blades in the same chassis which BOFM is managing. The following table compares the Advanced and Basic BOFM features:

Management Task	Basic BOFM	Advanced BOFM v4.0	Power based Blade Info
Slot-based I/O address assignment	Yes	Yes	
Manages Ethernet MAC addresses, Fibre Channel and SAS WWN numbers	Yes	Yes	No SAS Adapters support
Supports Fibre Channel and SAS boot targets	Yes	Yes	SAS & Fibre Channel boot targets not supported
Pre-assignment allows LAN/SAN configuration prior to blade installation	Yes	Yes	
Automatic re-assignment on blade swap	Yes	Yes	
Web-based GUI through AMM	Yes	No	
IBM Director-based GUI	No	No	
Blade Address Manager Configuration Wizard	No	Yes	
Standby Blade Pools	No	Yes	
Event Action Plans	No	No	
Provides I/O parameter and VLAN migration to standby blades in case of blade failure	No	Yes	

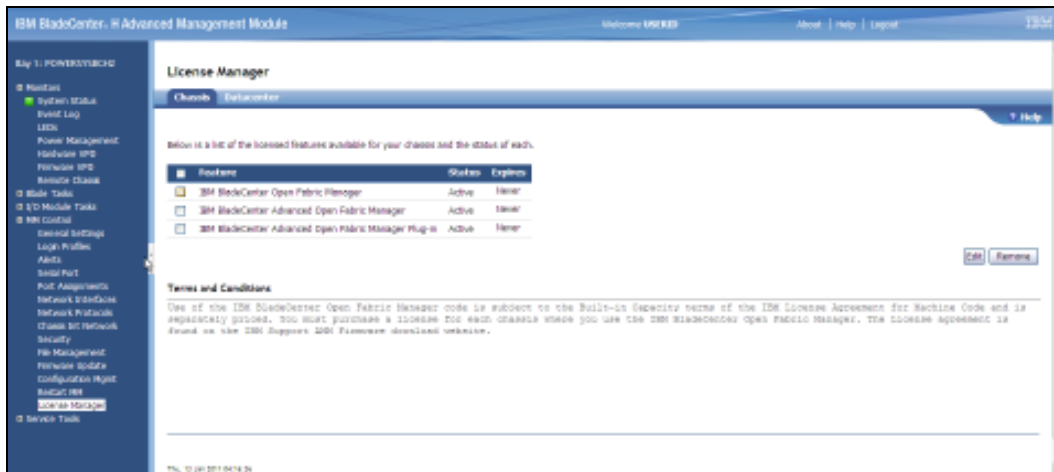
Management Task	Basic BOFM	Advanced BOFM v4.0	Power based Blade Info
Number of chassis supported by one BOFM console	100	256	
Number of manageable ports per blade	8	8	

For more information on BOFM go to <http://www-03.ibm.com/systems/bladecenter/hardware/openfabric/openfabricmanager.html>.

Note: Currently Advanced BOFM can not be installed on a server running AIX. Also none of the previous versions of BOFM v3.X (stand-alone) and System Director plug-in are supported with AIX as well. Support is planned for 3Q2011.

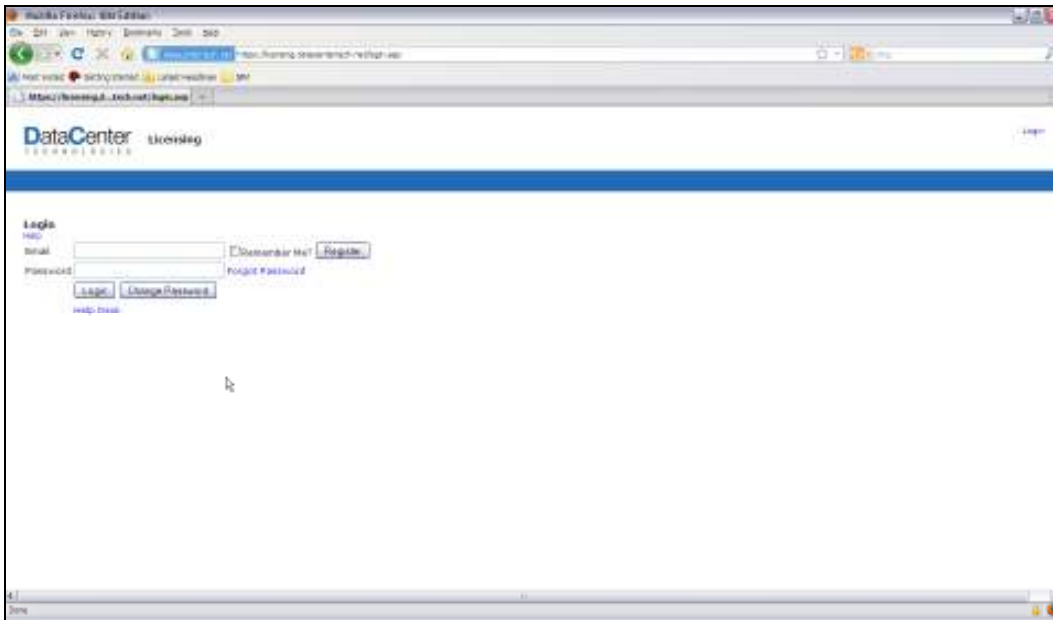
C. BOFM Licensing

Every chassis that will be managed by BOFM requires a license key. The license keys are activated from the AMM under License Manager. A license key is required for Basic BOFM and a license key is required for Advanced BOFM. If you are only interested in BOFM Basic then the BOFM Basic license key must be activated on the AMM. If you are interested in Advanced BOFM then you must have license keys activated for BOFM Basic and BOFM Advanced.



Note: The Basic BOFM and the Advanced BOFM licenses must be used together. The Plug-in Systems Director version will be sunset this year.

A 60 day trial license is available for BOFM Basic and Advanced. Once the trial period is up, BOFM will be disabled and you will have to purchase a license and will not be allowed to register for a second trial. For more information on BOFM Licensing go to <https://licensing.datacentertech.net/login.asp>.



If you have lost your BOFM license but you still have the order number send an email to ezevend@us.ibm.com with the order information and he will investigate and verify your account.

IV. Configuring BOFM

In this section we discuss the three main steps involved in configuring BOFM: configure the AMM, Chassis Discovery and authentication, and creating the profile. From the AMM some of the Network Interface parameters must be modified before you can use BOFM. The level of firmware supported on the AMM with BOFM v4.0 is BPET54p or higher. Also it is important to verify the firmware level on the blade and the expansion card. The profiles can be created from Basic and Advanced BOFM. If a profile is created from Basic BOFM it can be imported to Advanced BOFM. One profile can be applied to multiple blade slots or you can create one profile per blade slot (up to 14 depending on the chassis).

The best way to verify the BOFM options/configuration real time is from the AMM. If you have Advanced BOFM installed it will update to the AMM as well. Events from the Event Log on the AMM are automatically updated in BOFM which is how BOFM triggers failover events.

A. Configuring the AMM (Mandatory)

The following changes must be made to the Network Interface parameters on the AMM before running BOFM:

1. From the AMM under MM Control select “**Network protocol**”. Next select File Transfer Protocol (FTP) and make sure the FTP server is “**Enabled**”, then select “**Save**”.

File Transfer Protocol (FTP)

FTP server: Enabled

FTP idle timeout (seconds): 300

Save

2. Now go back to the SNMP Protocol link. The “**Access Type**” for Community Name “**trap**” should be changed to “**Set**” and the IP Address for the BOFM Server must be the first entry under Fully Qualified Hostnames or IP Addresses then select “**Save**”.

Simple Network Management Protocol (SNMP)

SNMP traps: Enabled

SNMPv3 agent: Enabled

Community Name: trap

Access Type: Set

Fully Qualified Hostnames or IP Addresses:

Community Name	Access Type	Fully Qualified Hostnames or IP Addresses
trap	Set	1. 172.23.0.99

Save

3. Now go back to the TCP Command Mode Protocol and change Command mode to “**10**” connections then select “**Save**”.

TCP Command Mode Protocol

Command mode: 10 connections

Secure command mode: 0 connections

Command mode inactivity timeout: 300 seconds

Save

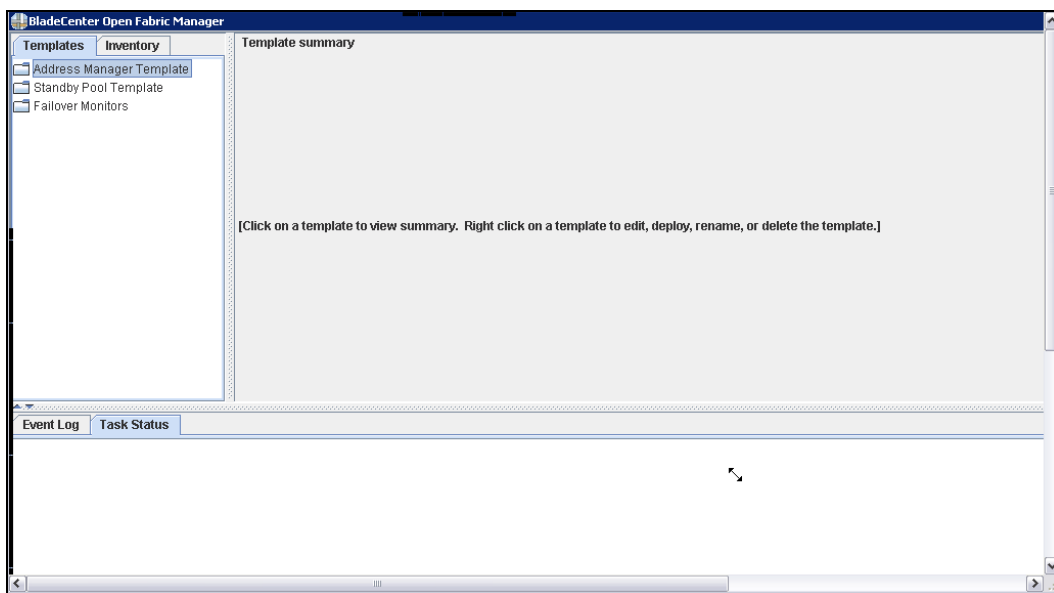
B. BOFM GUI

The BOFM GUI allows you to manage up to 256 chassis and up to 3584 blades from a single console. To start Advanced BOFM select the Advanced OFM Console Icon and login to the BOFM console. In the Pane on the left there are two tabs: Inventory and Templates. The Inventory tab is where you discover the chassis and collect inventory. The Template tab is where you create the Address Manager Profile, Create Standby Pool

and Failover Monitors. The Address Manager Template allows you to create/Update/Delete BOFM configuration templates; configure boot settings; per-port VLAN/boot priority customization; Apply/Deploy BAM configuration; and Import/Export configuration from/to CSV file. The Standby Pool Template allows you to manage resources for failover. The Failover Monitor Template allows you to manage failover criteria as well as dynamically Start/Pause/Resume monitoring for automatic failover. The bottom pane has two tabs: Event Log and Task Status.

To login to the BOFM GUI complete the following steps:

- ___ 1. To start Advanced BOFM select the Advanced OFM Console Icon.
- ___ 2. Login to the BOFM Console. There are three templates available. Also in the bottom pane you notice the Event Log and Task Status panes.




C. Discovering the Chassis

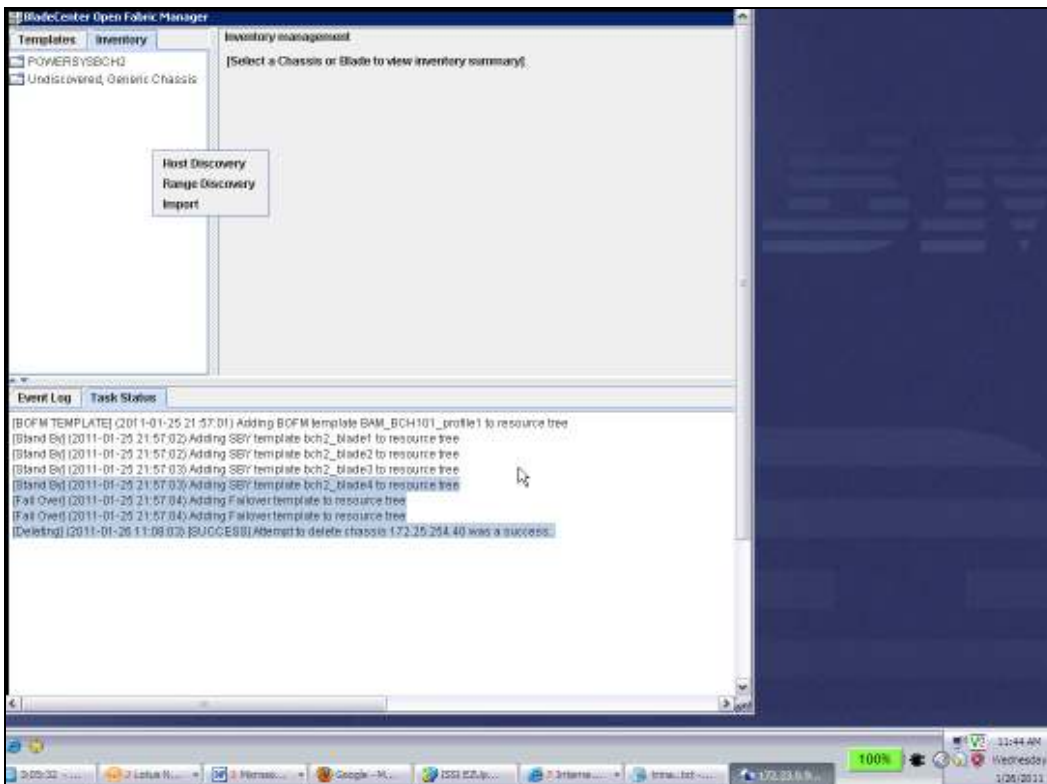
There are three ways to add Inventory in BOFM: Host Name, IP range discovery and chassis list import. When the inventory is collected you can choose to collect the chassis and blades, only the chassis inventory or put a placeholder for future chassis collection. The chassis can be discovered from Basic BOFM on the AMM or from Advanced BOFM. During the discovery process if a duplicate IP Address is encountered the discovery will skip the inventory import. Once the chassis is discovered by BOFM the chassis, blades, switches, etc. can be configured for management.

To add a chassis using the Host Name from Advanced BOFM v4.0, complete the following steps:

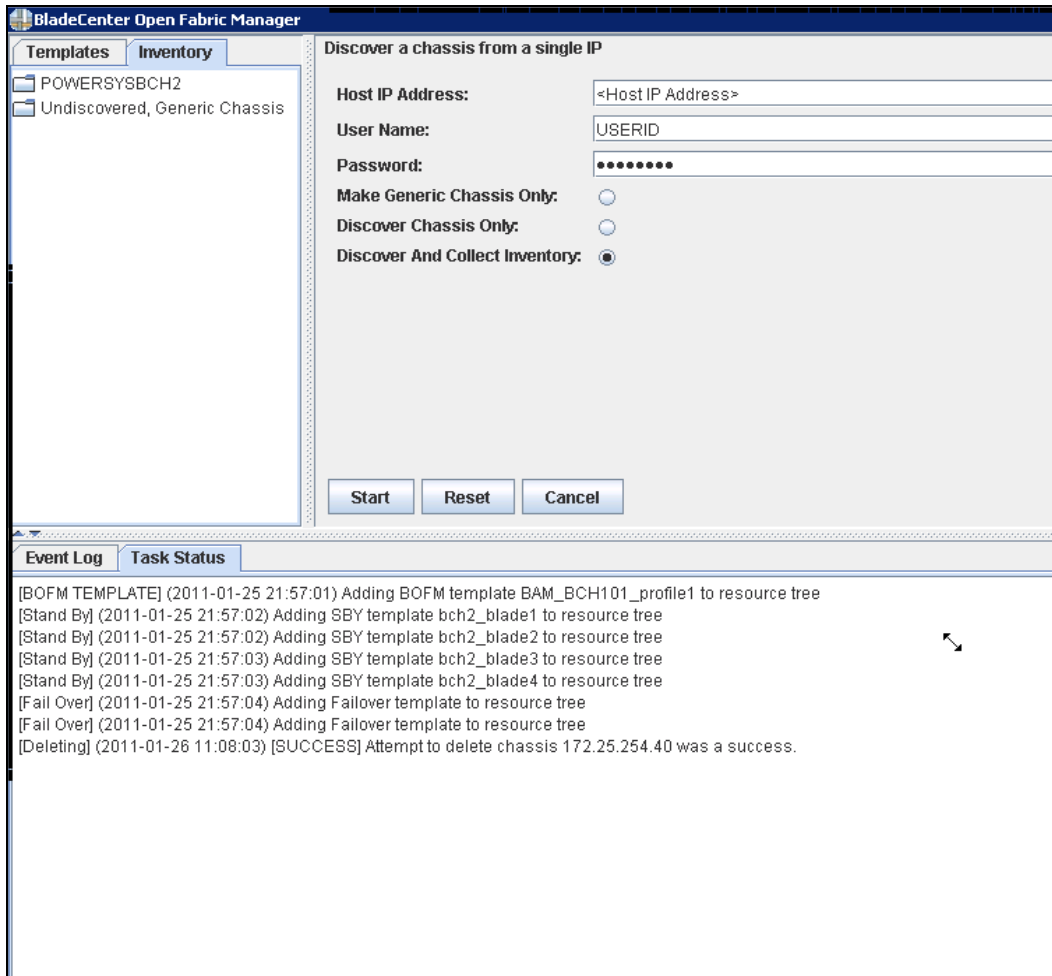
- ___ 1. Login to the BOFM Console. There are three templates available. Also in the bottom pane you notice the Event Log and Task Status panes.



2. From the Inventory tab and right click in the window, a pop-up window appears. Select **“Host Discovery”**.



3. Type in the “IP Address”, “User Name” and “Password” for the chassis you want to discover then select “Start”.



A summary screen is displayed. After the chassis has been authenticated it is added to the Inventory pane.

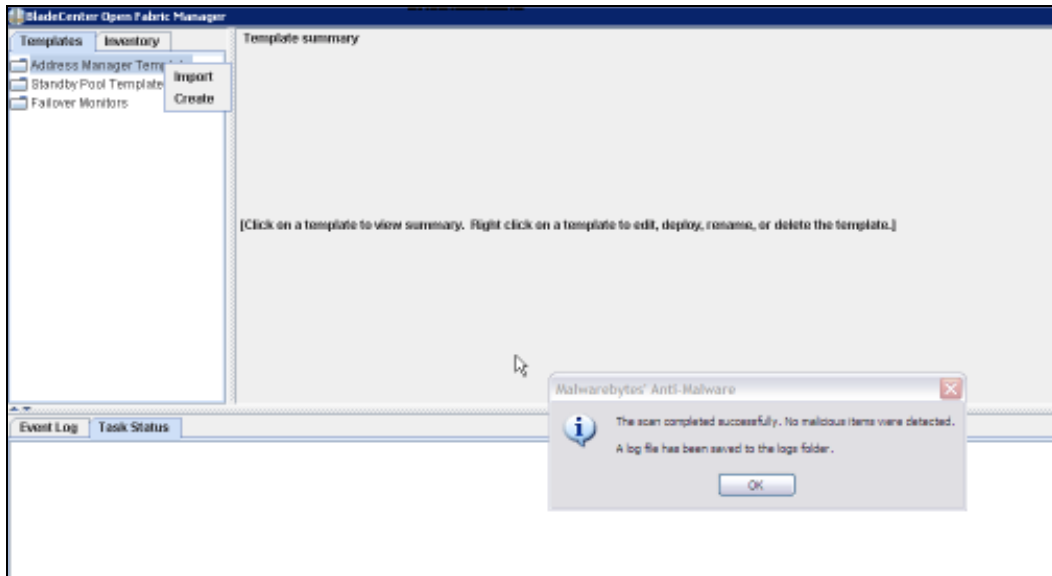
V. Advanced BOFM Templates

In this section we discuss how to create new templates from the Address Manager Template, create Standby Pool Templates which is used for backup blades and how to create event monitors for automatic failover. All of these features are configured from Advanced BOFM’s stand-alone GUI and can not be configured from the AMM.


A. Address Manager Template

The Address Manager allows you to create a new template and you can import a template. Once the template is created it can be deployed to a chassis. To create a new template, complete the following steps:


1. From the Template tab right click on “**Address Manager Template**“ then select “**Create**”.




2. Select the chassis you want to apply the new Template to and add it to the right pane.



- ___ 3. Select the chassis you want to apply the new Template to and add it to the right pane then select “Next”.
- ___ 4. Depending on the adapter you have installed on your blade select the appropriate tab. The Ethernet tab is displayed. The “Use a Range of MAC Addresses” option is required for Power based Blade Servers.




- ___ 5. Select the Fibre Channel tab. From the drop down list select the appropriate type of adapter installed on your blade. If you have both an Emulex and QLogic adapters on your blade then two profiles (one for Emulex and one for QLogic) must be created and applied to that blade.



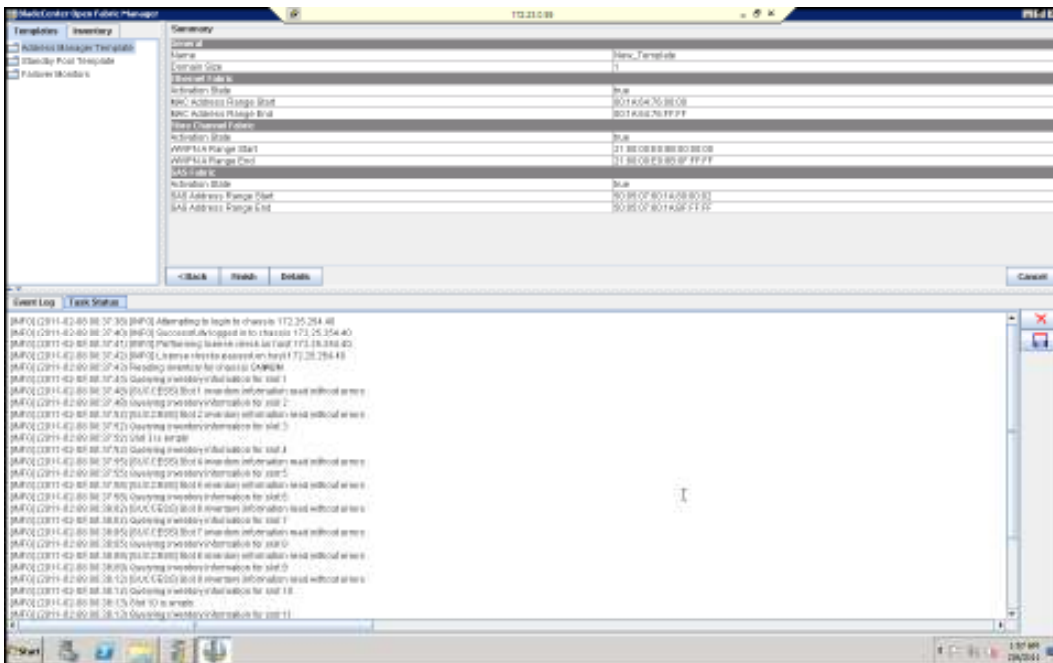
The type of adapter you have on your blade and the profile you create must match it. For example Emulex profiles are needed for Emulex adapters and QLogic profiles are needed for QLogic adapters. If your blade has both an Emulex and QLogic adapters then you have to deploy both profiles on that blade.

Note: The primary and/or secondary boot target option is not supported on Power based Blade Servers. Also, BOFM does not support SAS Adapters on the Power based Blade Servers.

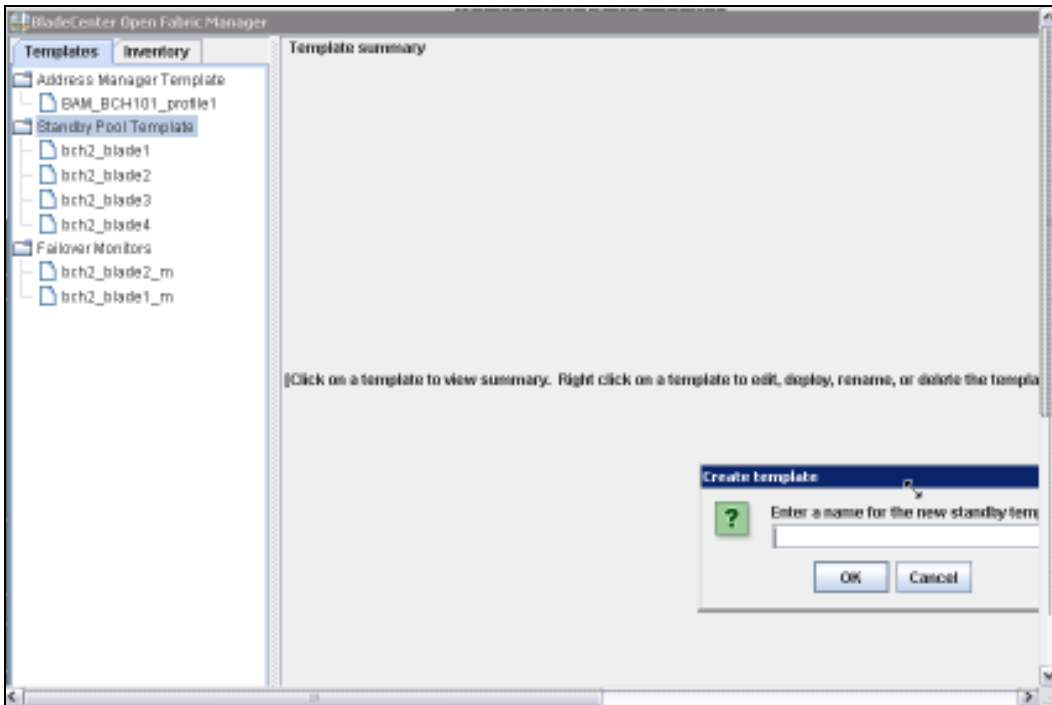
6. From the Advanced settings you can enable/disable/ignore individual blades; add additional blade offset for multi-wide blades, edit information for individual ports and edit boot target for individual blades. The Max Offset of 0 indicates Blade Slot 0 is a single wide blade (Max Offset 1 applies to a double-wide blade). Select “**Next**” to continue.



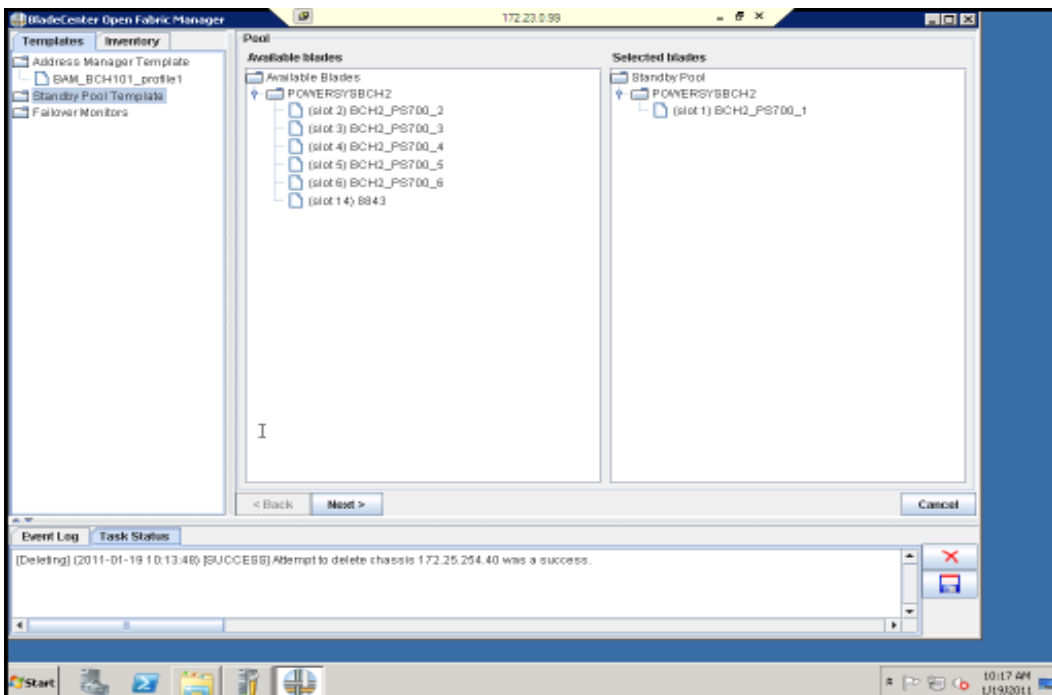
7. A Summary screen is displayed with your new template options. Select **“Finish”**, to complete the creation of the new template. The new template is now listed under the Address Manager Template.



8. Now select the new template and right click on it to list the options available for that template.

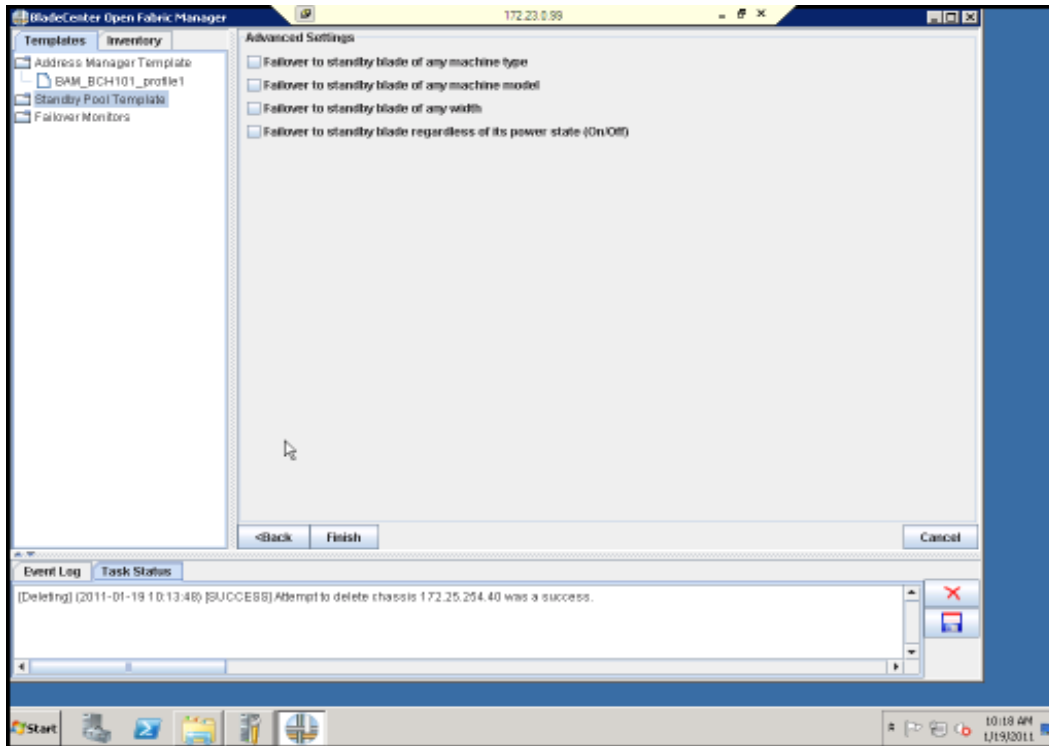


2. Select the twisty next to “*PowerSYSBCH2*”, which is the chassis we discovered to display the blade options available.

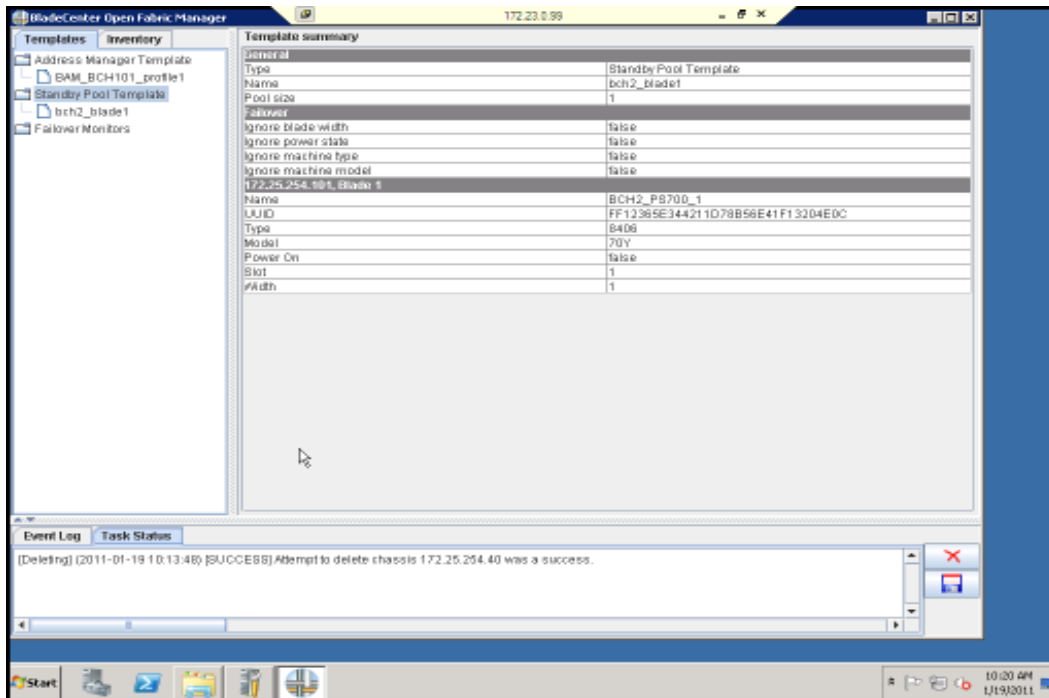


3. In our example Blade #2 is currently serving as the primary blade and Blade #1 is the Standby blade. Select Blade #1 from the list. It will appear in the Standby Pool section of the window now select “Next”. Select the Advanced

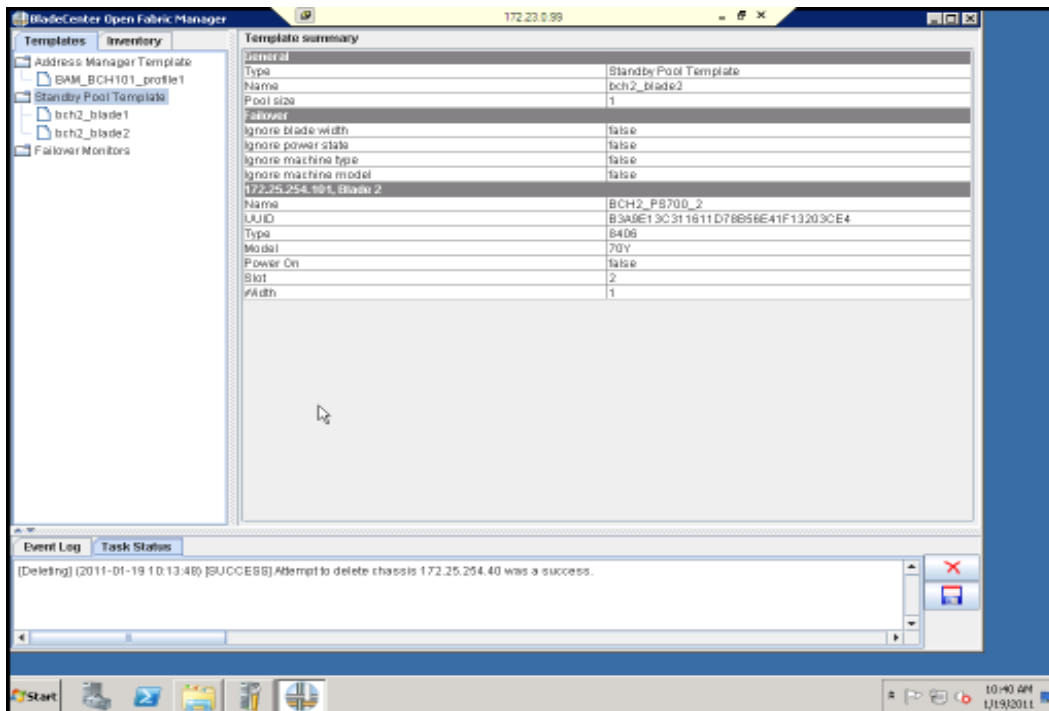
Settings to apply to the Standby Blade Pool. In my example we do not select any of these options. Now select “**Finish**”.



4. A summary window appears with your choices.



For testing purposes I added Blade #2 to the standby pool as well since I plan to failover and failback between my two blades. Here's a summary of the profile information for Blade #2:



C. Creating a Failover Monitor

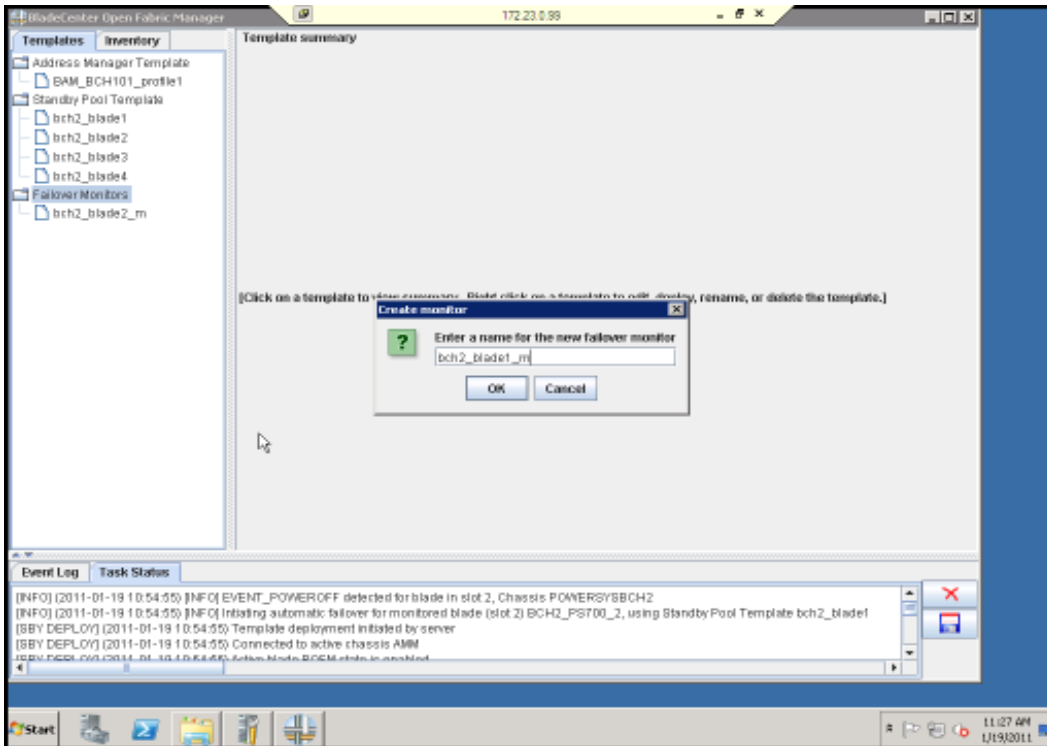
A Failover Monitor is created on the primary blade to monitor specific blade events. If the event is triggered it will cause the primary blade to automatically failover to the target blade. The following blade events are monitored by BOFM:

- CPU Failure
- Blade COMM Errors
- Blade Removal
- HDD Failure
- Blade Denied Power
- Memory Failure
- Voltage Warnings
- PFA Events

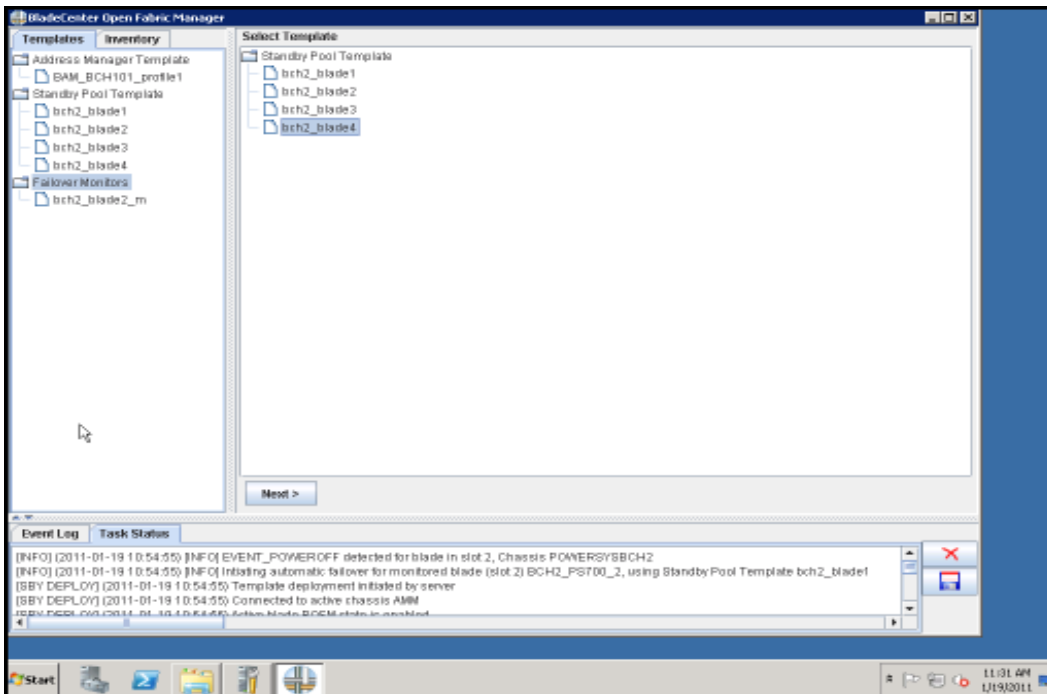
If one of these events occurs a blade from the Standby Pool will be powered on in the monitored blade's place.

To create a Failover Monitor, complete the following steps:

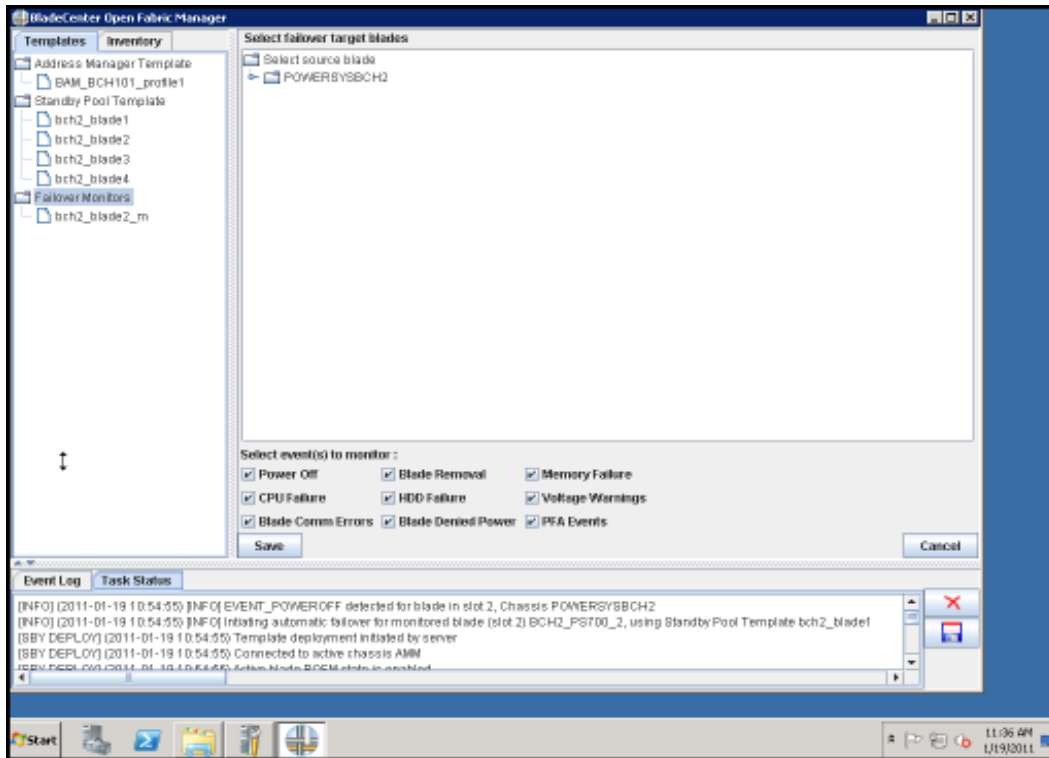
- ____ 1. From the Template tab right click on “**Failover Monitors**” and select “**Create**”. A pop-up window appears.



2. Enter the name for the new failover monitor and select “**Ok**”.
3. From the **Select Template** window select the blade you want to monitor then select “**Next**”.

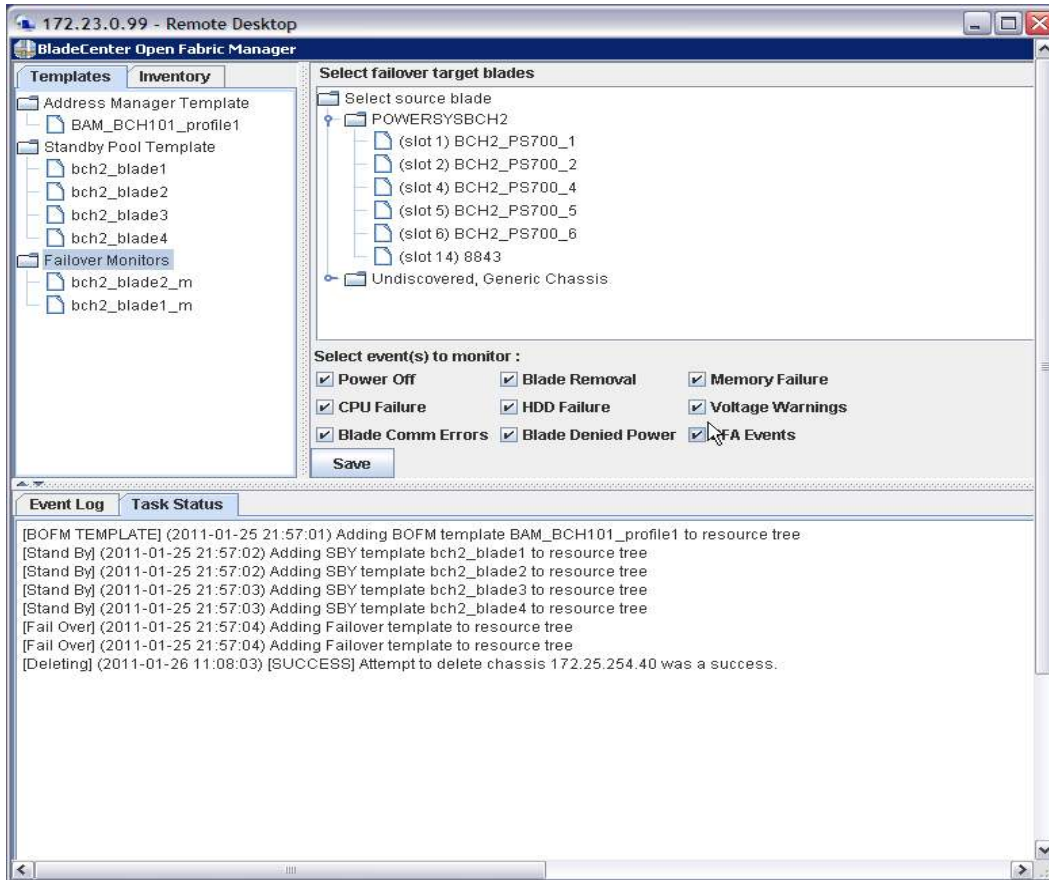


4. Towards the bottom of the window notice all of the monitored events are selected as the default. You can uncheck any event you do not want to monitor.

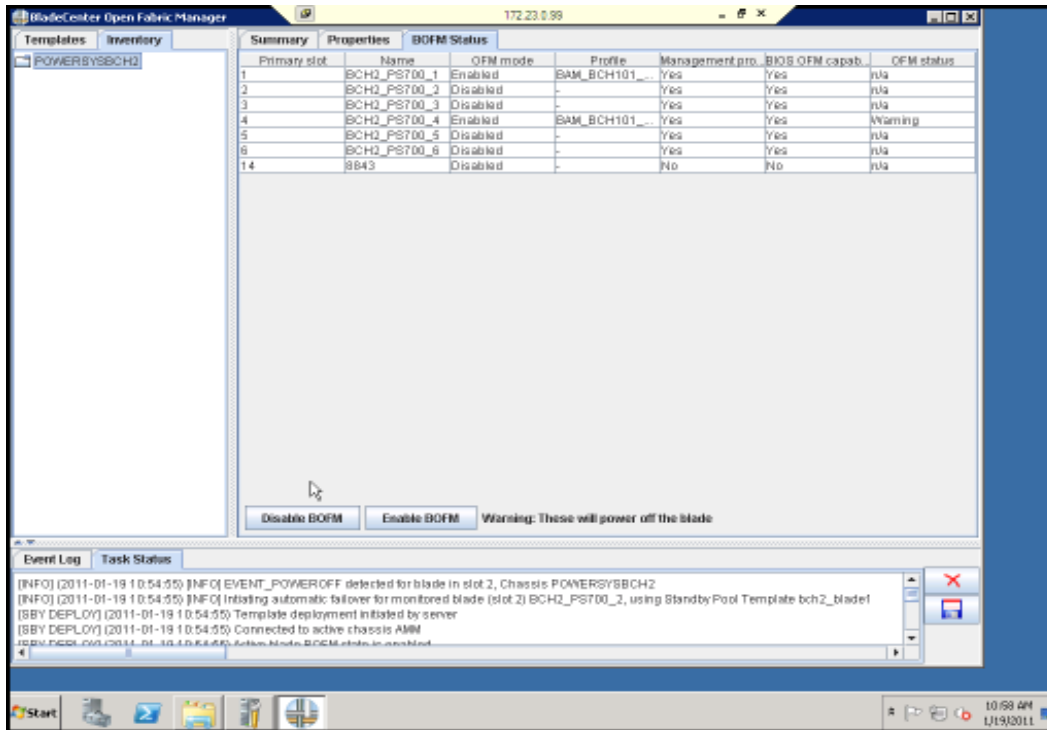


Note: Selecting the “**Power Off**” monitor as part of your failover event has to be carefully planned as any time the server is powered off (whether intentional or not) it will cause a failover.

5. Now select the key next to “**POWERSYSCH2**” which is the chassis we are using and select the blade you want to associate with the monitored events below then select “**Save**” to continue. In our example we configured blade #1 for monitoring therefore blade #2 is selected as the failover blade.



6. Notice all of the monitored events are selected as the default. Select the key next to “**POWERSYSCH2**”. Now select the blade you want to associate with the monitored events below then select “**Save**” to continue.
7. To view the BOFM Status go to the “**Inventory**” tab and select “**BOFM Status**”. Although the GUI provides status information for BOFM, it is recommended this information is viewed from the AMM since events are sent to the AMM before they get to BOFM so the AMM information is real time.



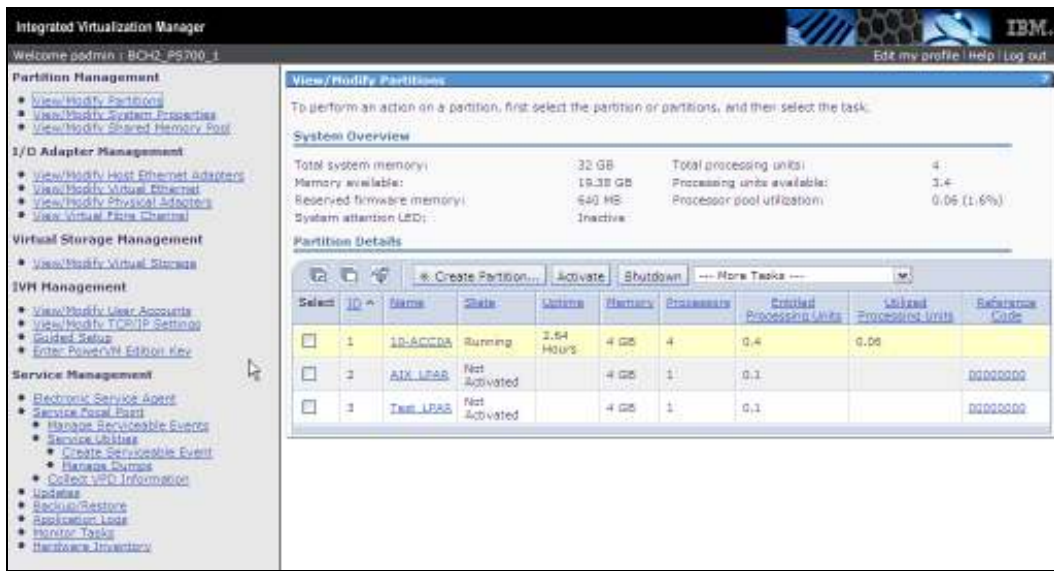
VI. Testing Failover on the Blades

Failover can be initiated from Advanced BOFM both manually or automatically if triggered by a specific event (i.e. CPU Failure, Power off, etc) that is monitored on the blade. With automatic failover you can have multiple blades configured for takeover. For example in a BladeCenter H chassis all 14 blades can be figured for BOFM, seven of these blades can be configured as the primary and seven as backup blades. If the blades are virtualized there will be several additional steps required before the blade can be fully operational.

It is important to note that Power based Blade Servers have serial numbers attached to the server and the LPARs and this information has to be modified to complete the failover process. Also there is information in NVRAM that is not transferred over to the target blade. As a result, there are several manual steps that have to be completed in order to complete the failover for Power based Blade Servers.

A. Backup the Blades

Before you initiate failover I recommend you backup the LPAR and mapping configuration as this information will be required to complete the failover on your blade (*Section D. Restoring the Blade Configuration, page 28*). During the failover process, all of the configuration information is not transferred to the new blade so this is an important step!



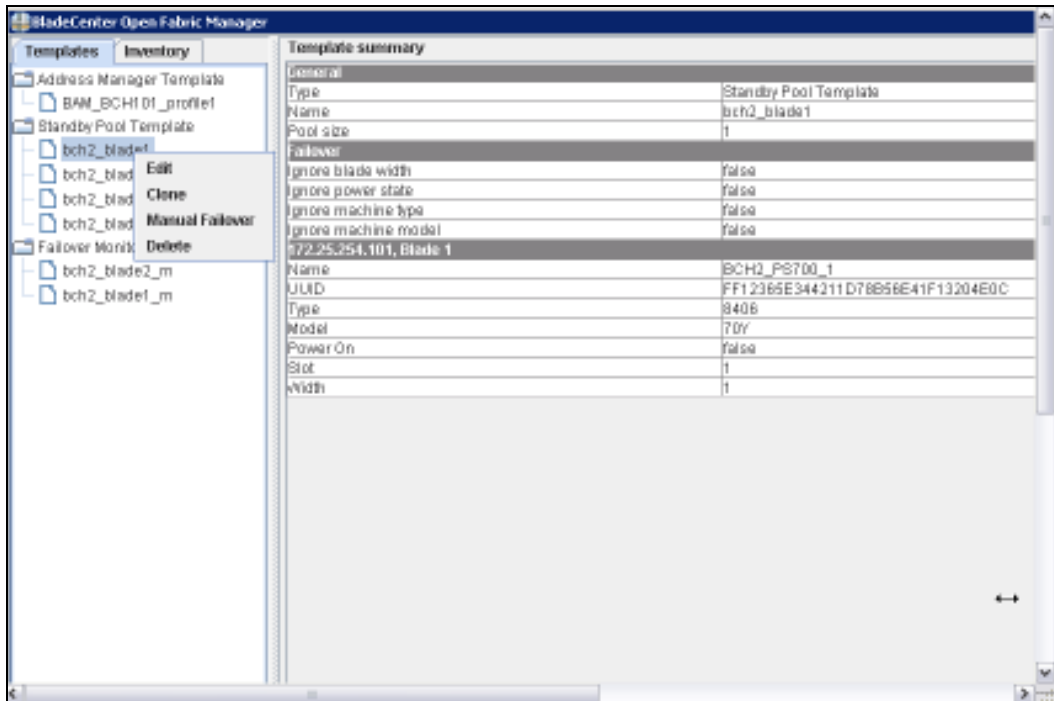
To backup the blade configuration files complete the following steps:

- ___ 1. To backup the VIOS type “**viosbr –backup –file filename**”. This is a good overall backup to take as it allows you to view your NPIV mapping information but you can not restore it using the “**viosbr**” command.
- ___ 2. Backup the profile.bak file by typing “**bkprofdata –o backup –f /home/padmin/profile.bak**” which contains the LPAR information.

B. Manual Failover

The blade can be configured to manually failover by completing the following steps:

- ___ 1. To manually failover a blade from the GUI select the blade from the Standby Pool Template and right click on it. Select “**Manual Failover**”.
- ___ 2. A message appears in the “Event Log” on the bottom pane. You can go to the AMM verify the failover blade is now powered on.



C. Automatic Failover

Automatic failover is triggered on a blade that has a failover monitor configured. The primary blade is powered off and a blade in the Standby Pool is powered on. When an event is triggered on the blade it is logged in the Event Log on the AMM and it is logged in BOFM.

To trigger an event complete the following steps:

- ___ 1. To test automatic failover login to your blade and shutdown the Operating System. Since this is one of our monitored events BOFM will automatically “**Power On**” the backup blade which will take on the new config (IP Address, WWN, etc.).
- ___ 2. A message should appear in the bottom pane of the BOFM GUI. Go back to the AMM and under Blade Task the target blade should be powered on.

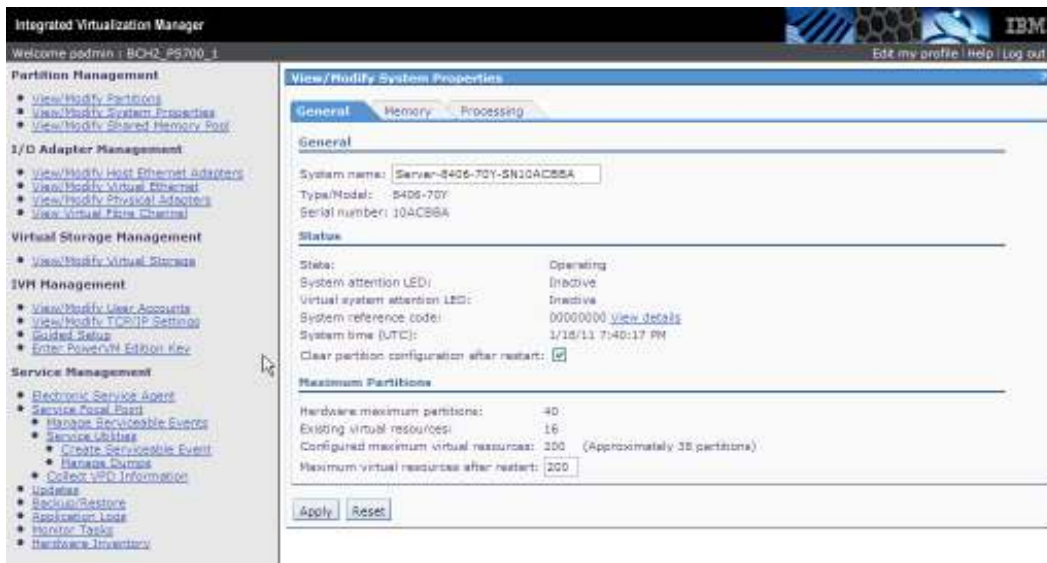
D. Restoring the Blade Configuration

When a failover is triggered and the target blade takes over for the primary blade, there are several manual steps required to restore the LPAR configuration to the target blade as some of the configuration information has not been transferred to the target blade. For example, the serial number of the blade is attached to the LPARs and there is information in NVRAM that is not transferred over to the target blade. If one of your LPARs has NPIV enabled then the serial number must be changed to match the new blade by editing and restoring the */home/padmin/profile.bak* file. This change will allow the virtual

WWPNs associated with the LPAR to be restored. Also, the NPIV mappings (vfchost to fcs) are not automatically transferred to the target blade. The NPIV mappings can not be restored with the “**viosbr**” command or from the “**profile.bak**” file.

To restore the LPAR definitions and NPIV mapping on your blade complete the following steps:

- ___ 1. Boot the new target blade to the SMS Menu to check the boot order as needed.
- ___ 2. The profile on the new target blade will need to be cleared depending on what was installed on the blade (i.e. any existing LPAR information). The profile can be cleared from the GUI under “**View/Modify System Properties**” and from the CLI.

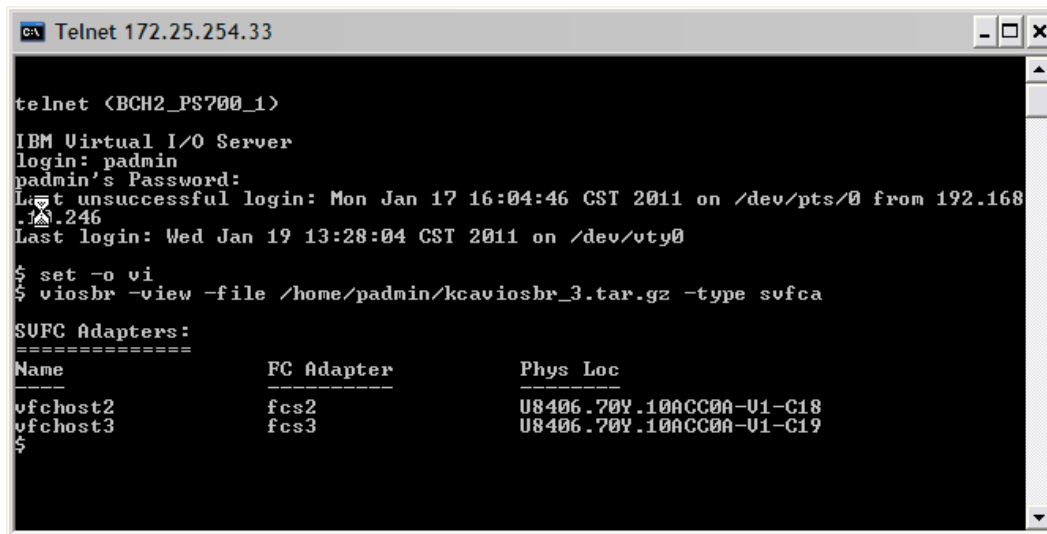


To clear from the command line type “**lpcfgop -o clear**”. Validate you want the LPAR information cleared.

- ___ 3. Reboot the new target blade
- ___ 4. From SOL console edit the “**/home/padmin/profile.bak**” file if NPIV is enabled on the LPAR and change the serial number to the current blade. If you are NOT using NPIV then this profile does NOT have to be changed.
- ___ 5. If your LPAR has NPIV enabled or NPIV and Virtual SCSI then type the following command to restore the profile: “**rstprofdata -l 1 -f filename**”.
- ___ 6. If you are using Virtual SCSI only then type this command to restore the profile: “**rstprofile -l 1 -f /home/padmin/profile.bak --ignoremtms**”.
- ___ 7. Reboot the blade
- ___ 8. You will need to remap the vfchosts to the physical fcs devices for every LPAR so type the following:
“**vfemap -vadapter vfchost2 -fcp fcs2**”
“**vfemap -vadapter vfchost3 -fcp fcs3**”.

Note: If your blade has more than one expansion card installed on it and if there are multiple LPARs then you need to write down what vfchost number is assigned to what LPAR ID to ensure the mappings are correct.

9. To view the original mappings for vfchost device and the physical fcs type the following:
“viosbr -view -file filename -type svfca”



```
telnet <BCH2_PS700_1>
IBM Virtual I/O Server
login: padmin
padmin's Password:
Last unsuccessful login: Mon Jan 17 16:04:46 CST 2011 on /dev/pts/0 from 192.168.1.246
Last login: Wed Jan 19 13:28:04 CST 2011 on /dev/tty0
$ set -o vi
$ viosbr -view -file /home/padmin/kcaviosbr_3.tar.gz -type svfca
SUFC Adapters:
=====
Name                FC Adapter          Phys Loc
-----
vfchost2            fcs2                U8406.70Y.10ACC0A-U1-C18
vfchost3            fcs3                U8406.70Y.10ACC0A-U1-C19
$
```

Note: The vfc mappings do not come back with the “viosbr” command (different physical locations because of serial number change). Also if the virtual SCSI mappings did not restore then you can use the “viosbr” command to recover them.

VII. Configuring Rip & Replace from the AMM

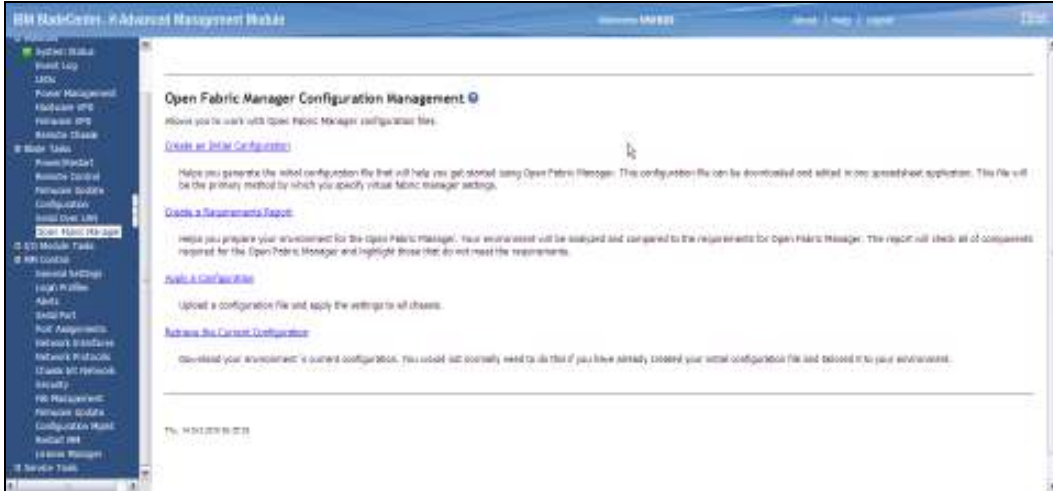
As previously discussed Rip and Replace allows you to configure slot based I/O Address assignment. It is a feature of Basic BOFM and is configured from the AMM. Rip & replace allows you to assign a profile to the Blade slot and if that blade fails the next blade inserted into that same slot will inherit its IP Address, WWN information, etc.

A. Creating the Profile

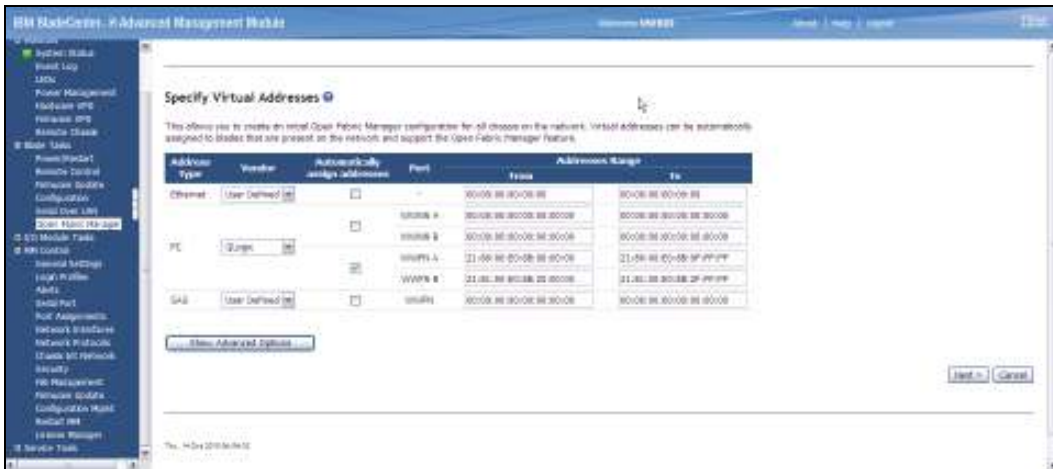
A profile is created for each blade slot in the chassis. The profile created on the AMM can be imported (csv file) to Advanced BOFM. One profile can be created per blade slot or multiple profiles can be created and assigned to the same blade slot. A profile created on the AMM can be imported to Advanced BOFM.

The following steps are required to create a profile from the AMM:

1. From the AMM under “**Blade Task**” select “**Open Fabric Manager**”. Select “**The requirements Report**” to ensure your environment is prepared for Open Fabric Manager. Your environment will be analyzed and a report will be produced.



2. From the drop-down box select the type of adapters you have installed on the blade. Select “**Advanced Options**”. Now select “**Next**” to continue.



Note: The type of adapter you have on your blade and the profile you create must match it. For example Emulex profiles are needed for Emulex adapters and QLogic profiles are needed for QLogic adapters. If your blade has both an Emulex and a QLogic adapters then you have to deploy both profiles on that blade.

3. Select “**Show Advanced Options**”. The “**Enable Blade Offsets**” should be “**Enable Offset 0**” for a single wide blade (a double wide blade would be offset 2). For Power based Blade Servers must use the “**Generate range of MAC addresses per port**” for Ethernet adapters.



Note: The Power based Blade Servers do not support Fibre Channel/SAS boot targets. Also BOFM does not support SAS adapters on Power based Blade Servers.

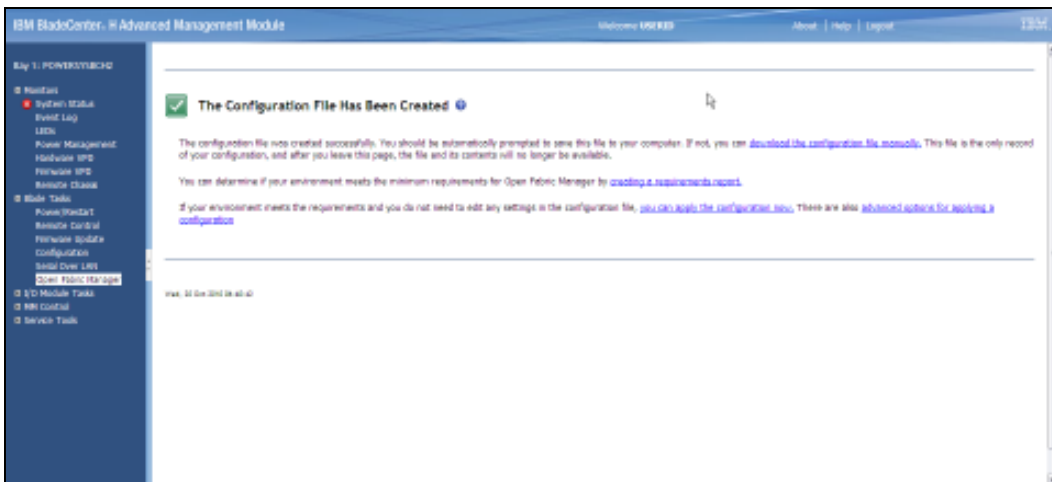
4. Select “**Next**” to have the AMM discover additional chassis on the network.



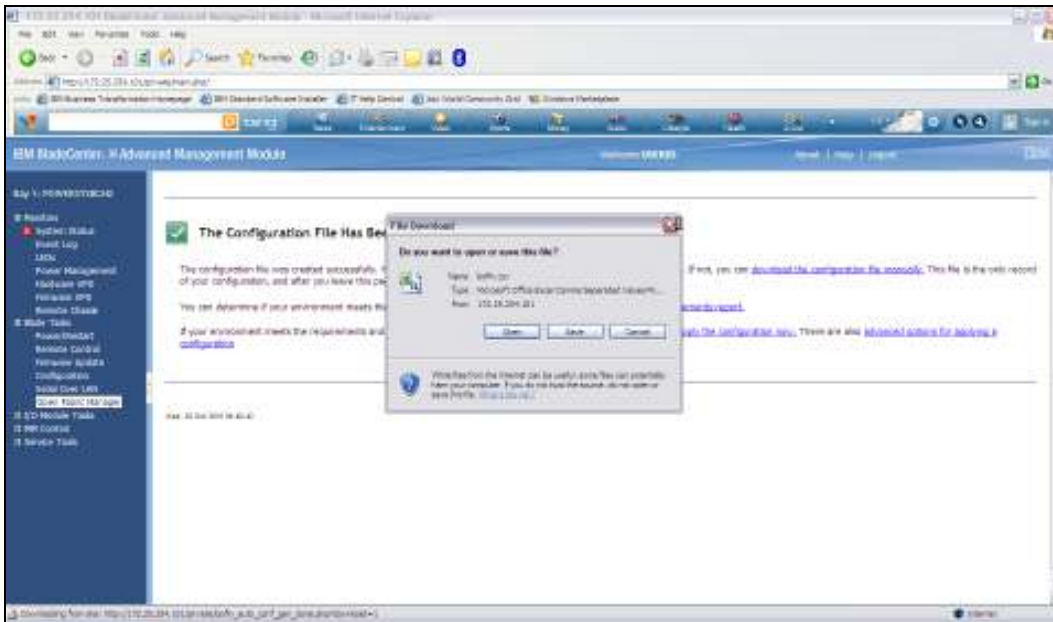
5. If you have an existing configuration file you can add the new config to an existing file.



6. The configuration file has been saved. Select the “**download the configuration file manually**” link to edit the file. The BOFM configuration file (*BOFM.csv*) file can be downloaded from the AMM to the Management System.



7. Select “**Open**” to view/modify the BOFM configuration file.

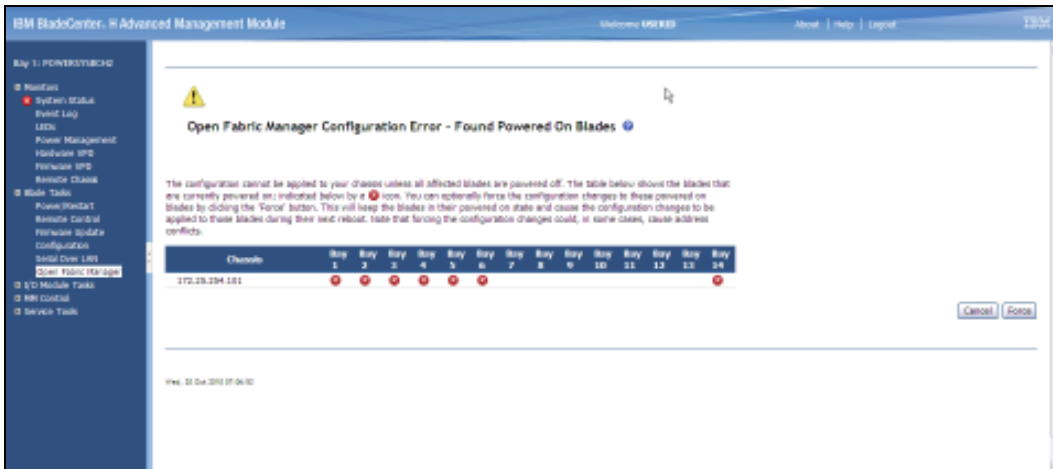


8. This spreadsheet can be modified and viewed. The Mode Column shows “**Enable**” for each Blade Slot. The blade can be *Enabled* or *Disabled*. Also the boot option must be selected if booting from SAN. After all changes save the spreadsheet.

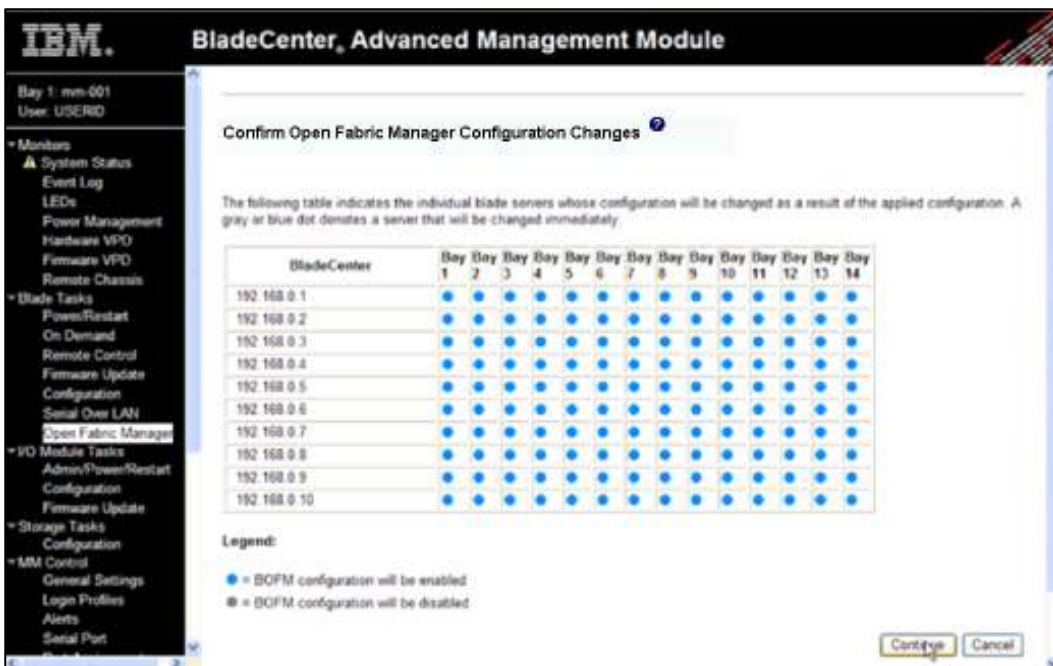
Blade Slot	Type	Mode	Profile
1	Blade Center 172 25 254 101		
2	//IP Type (Slot Slot)		
3	1172 25 254 BladeCent.appl		
4	//IP Type (Slot Slot)	1 enable	TempProfile BC-101e-1
5	1172 25 254 Slot		
6	//IP Type (Slot Slot)	2 enable	TempProfile BC-101e-2
7	1172 25 254 Slot		
8	//IP Type (Slot Slot)	3 enable	TempProfile BC-101e-3
9	1172 25 254 Slot		
10	//IP Type (Slot Slot)	4 enable	TempProfile BC-101e-4
11	1172 25 254 Slot		
12	//IP Type (Slot Slot)	5 enable	TempProfile BC-1Slot-5
13	1172 25 254 Slot		
14	//IP Type (Slot Slot)	6 enable	TempProfile BC-1Slot-6
15	1172 25 254 Slot		
16	//IP Type (Slot Slot)	7 enable	TempProfile BC-1Slot-7
17	1172 25 254 Slot		
18	//IP Type (Slot Slot)	8 enable	TempProfile BC-1Slot-8
19	1172 25 254 Slot		
20	//IP Type (Slot Slot)	9 enable	TempProfile BC-1Slot-9
21	1172 25 254 Slot		
22	//IP Type (Slot Slot)	10 enable	TempProfile BC-1Slot-10
23	1172 25 254 Slot		
24	//IP Type (Slot Slot)	11 enable	TempProfile BC-1Slot-11
25	1172 25 254 Slot		
26	//IP Type (Slot Slot)	enable	TempProfile BC-1Slot-12
27	1172 25 254 Slot		

Note: The “csv” file created in BOFM Basic can be imported to BOFM Advanced.

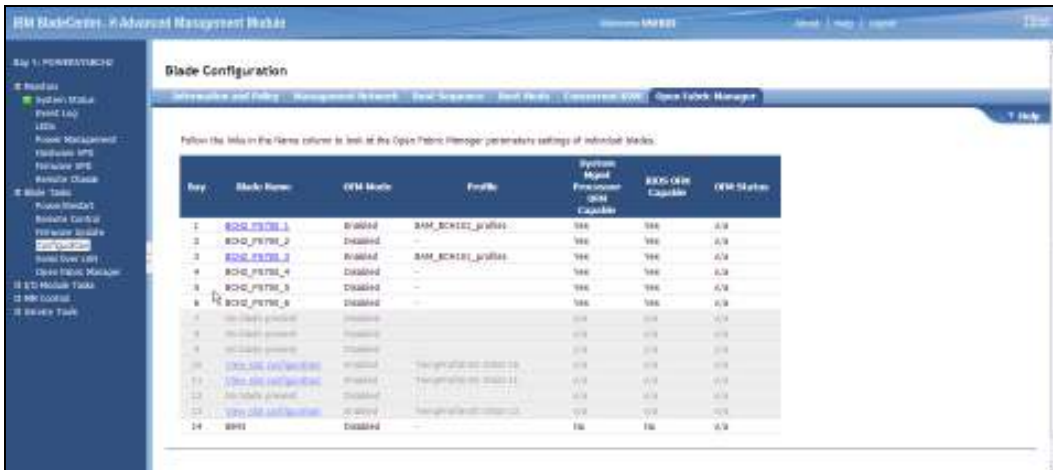
9. The next step is to create a requirements report which checks for the compatible firmware for the Management Module, BMC and BIOS versions installed on the blades. Select the “creating a requirements report link”. If the blades are powered on you will receive the following error:



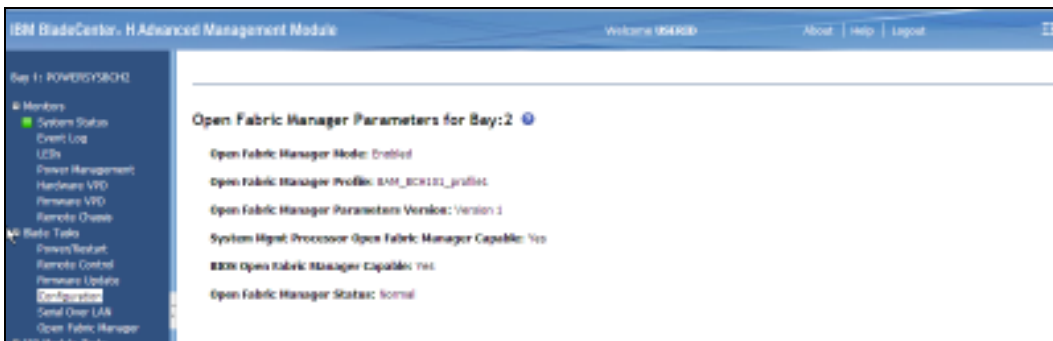
10. Forcing this operation could possibly cause duplicate Addresses so it is not recommended. The Blades should be powered off first then the configuration can be completed. The next step would be to confirm the configuration.



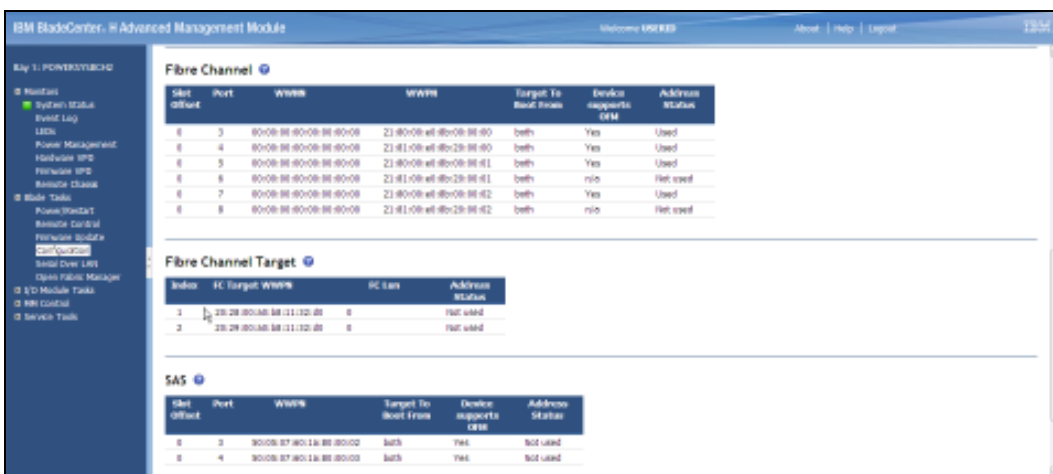
11. The configuration is complete and the blade must be rebooted so the new config can take affect.
12. To view the BOFM configuration on the blade from the AMM under Blade Task select “**Configuration**” then go to the “**Open Fabric Manager**” tab.



13. Now select one of the links to review the BOFM configuration on that blade.



Now scroll down to the bottom of the screen to review the Fibre Channel / SAS configuration. The Power based Blade Servers do not use the MAC addresses for Fibre Channel Target. Also SAS Expansion card is not currently supported for in BOFM on Power based Blade Servers therefore those MAC addresses are not used.



At this point if a blade is removed from a slot, a new blade can be inserted into the same slot and the original IP Address, WWN, zoning, etc for the blade you removed will be

available on the new target blade. If your blade has NPIV enabled LPARS, you will need to backup, modify and restore the “**profile.bak**” file and the NPIV mappings, as this information is not transferred to the new blade (refer to Section VI of this whitepaper).

This completes my whitepaper!

References

What is New in BOFM v4.0

[https://w3-](https://w3-03.sso.ibm.com/sales/support/ShowDoc.wss?docid=T671709Z18878F92&infotype=SK&infosubtype=W0&node=doctype.W0%7Cdoctype.MUM%7Cdoctype.WBC%7Cbrands.B5000%7Cbrands.B5G00%7Cgeography.AMR&appname=CC_CFSS)

[03.sso.ibm.com/sales/support/ShowDoc.wss?docid=T671709Z18878F92&infotype=SK&infosubtype=W0&node=doctype.W0%7Cdoctype.MUM%7Cdoctype.WBC%7Cbrands.B5000%7Cbrands.B5G00%7Cgeography.AMR&appname=CC_CFSS](https://w3-03.sso.ibm.com/sales/support/ShowDoc.wss?docid=T671709Z18878F92&infotype=SK&infosubtype=W0&node=doctype.W0%7Cdoctype.MUM%7Cdoctype.WBC%7Cbrands.B5000%7Cbrands.B5G00%7Cgeography.AMR&appname=CC_CFSS)

BladeCenter Open Fabric Manager v4.0 User's Guide

http://ftp.software.ibm.com/systems/support/system_x_pdf/49v0230.pdf

BladeCenter Open Fabric Manager

<http://www.redbooks.ibm.com/abstracts/TIPS0713.html?Open>

Troubleshooting BOFM

http://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp?topic=/com.ibm.blade_center.advmgtmod.doc/jr1bs_bofm_tp_advanced.html

Implementing BladeCenter Open Fabric Manager v3.1

<http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101691>

BladeCenter Open Fabric Manager Hardware and Software Support

<http://www-03.ibm.com/systems/bladecenter/hardware/openfabric/openfabricmanager/index.html>

Power based Blade Server Wiki

<https://www.ibm.com/developerworks/wikis/display/WikiPtype/POWER+Blades>