



IBGranite Fabric Verification Suite: User's Guide

Rev 0.10

Mellanox Technologies

© Copyright 2005. Mellanox Technologies, Inc. All Rights Reserved.

IB Granite User's Guide

Document Number: 2196UG

Mellanox Technologies, Inc.
2900 Stender Way
Santa Clara, CA 95054
U.S.A.
www.Mellanox.com

Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies Ltd
Hermon Building
Yokneam 20692
Israel

Tel: +972-4-909-7200
Fax: +972-4-959-3245

Mellanox Technologies

Contents

1	Introduction	5
1.1	Supported Platforms and Operating Systems	5
2	Key Concepts And Terms	6
3	IBGFVS High Level Architecture	7
3.1	The Fabric Verification Process	7
3.2	IBGFVS Functionality and Steps	8
3.2.1	IB Granite Fabric Verification Suite (IBGFVS) Synopsis	8
4	Prerequisites and Dependencies	9
4.1	Prerequisites	9
4.2	Dependencies	9
5	IBGFVS Installation and Setup	10
5.1	Installation	10
5.2	Setup	10
6	IBGFVS Troubleshooting Guide	11
6.1	Troubleshooting General Errors	11
6.1.1	Locating a Problematic Device	11
6.2	Troubleshooting IBGFVS Steps	12
6.2.1	Troubleshooting Step 1: Subnet Manager Status	12
6.2.2	Troubleshooting Step 2: Subnet Analysis	13
6.2.3	Troubleshooting Step 3: Topology Matching	14
6.2.4	Troubleshooting Step 4: Hosts Info	15
6.2.5	Troubleshooting Step 5: Firmware Status	16
6.2.6	Troubleshooting Step 6: Fabric Stability	16
6.2.7	Troubleshooting Step 7: ULPs - IPoIB	17
6.2.8	Troubleshooting Step 8: ULPs - MPI	17
Appendix A	Topology File Description	19
Appendix B	Summary Report Example	21
Appendix C	Detailed Report Example (Output Logfile)	23
Appendix D	Fabric Design Specification	30

Mellanox Technologies

1 Introduction

Fabric design, bring up, operation, and maintenance stages comprise the life cycle of a fabric. The IBGranite product family offers solutions targeted at the different stages of the life cycle of an InfiniBand fabric. The IBGranite Fabric Verification Suite (IBGFVS) is one product in the IBGranite family, and is targeted at the design and bring up stages of the fabric life cycle.

The objective of IBGFVS is to enable end users to design, install and bring up an InfiniBand fabric in the most efficient and cost effective way. Once the processes defined by IBGFVS are executed and results pass the IBGFVS certification criteria, the fabric can be deemed ready for usage in the next stage of the fabric life cycle.

This document provides the user with IBGFVS related general concepts, high level architecture and dependencies, installation information, and detailed steps and procedures for executing recommended IBGFVS processes for verifying the design and bring up of the fabric.

1.1 Supported Platforms and Operating Systems

For the list of platforms and operating systems that IBGFVS supports, please refer to the *IBGFVS Release Notes* of the IBGFVS revision you are using.

Mellanox Technologies

2 Key Concepts And Terms

Throughout this manual, there are frequent references to certain IBGranite specific concepts and terms. These are described below:

- **IBGranite Master Node:** This is a dedicated node, such as an IT console, from which the user will launch IBGFVS. This node is required to be connected to the fabric, operational with IBGranite required software, and have special access to all other nodes in the fabric during the fabric bring up phase when the IBGFVS processes are executed. The requirements are discussed later in this document. There is no need for this node to be available in the fabric after the IBGFVS steps are completed with success.
- **IBGranite Agent Node:** All nodes in the InfiniBand fabric other than the IBGranite Master Node are called IBGranite Agent Nodes. This is because during the execution of the IBGFVS steps from the Master Node, IBGranite components - called Agents - are distributed and installed on all the nodes in the fabric. These Agents are uninstalled after the IBGFVS steps are completed.
- **Fabric Design Specification:** This is a design specification document that provides a checklist of items and design guidelines necessary for designing an InfiniBand fabric. Following the guidelines and checklist items can yield smoother and faster execution of IBGFVS processes.
- **Topology File:** This is a configuration file that feeds as input to IBGFVS. The Topology File describes the InfiniBand connectivity options and systems included in the fabric. The specified topology is verified by IBGFVS against the discovered one during the IBGFVS execution process.
- **Performance Threshold Configuration File:** IBGFVS verifies application level latency and bandwidth. The Performance Threshold Configuration File is a configuration file that feeds as input to IBGFVS. It can be used to specify defaults that a typical installation and configuration can use. Custom parameters can be inserted at installation time to override default parameters in this file. IBGFVS also verifies performance against error and other statistical counters supported by InfiniBand HCAs. The default thresholds for those counters are also stored in this file and they can be overridden by using the custom installation option.
- **In-Band:** A term assigned to communication between the Master Node and the Agent Nodes using the InfiniBand based connectivity within the fabric. IBGFVS operates through In-Band communication only.
- **Out-of-Band:** A term assigned to communication between the Master Node and the Agent Nodes using the non-InfiniBand based connectivity available within the fabric, for example Ethernet or I2C-compatible. Future IBGranite products designed for operation and maintenance stages of the fabric life cycle may use Out-of-Band communication.

The following section, [IBGFVS High Level Architecture \(page 7\)](#), further describes the application of the concepts presented in this section.

3 IBGFVS High Level Architecture

IBGFVS is based on multiple Software Agents running concurrently on every HCA node in the InfiniBand fabric. The Software Agents are designed and qualified to execute in parallel with Mellanox's IB Gold Distribution (IBGD) or any other InfiniBand HCA software stack, which includes drivers, access layers and ULPs chosen by the end user for use in the fabric. The software Agents are distributed and installed on the nodes during run time by the Master Node that runs the IBGFVS Master Software.

The Software Agents interface with various components in the HCA nodes to collect InfiniBand device specific information from the fabric. Each such Agent interfaces with HCA hardware using the PCI interface employing PCI transactions. The Software Agents can also interface with the driver access layer to extract information from the HCA driver. Finally, they can interface with ULPs using standard-based interfaces.

The Master Software running on the Master Node collects information from two sources, both via an In-Band connection. First, it collects information from the Software Agents via remote access (rsh or ssh). Second, it collects information about the state of the InfiniBand fabric from the local Subnet Manager by using queries based on standard InfiniBand management datagrams (MADs) or enhanced Mellanox Vendor-Specific MADs.

The Master Software and Agent Software have a set of dependencies that determine the extent of verification and results that can be provided by IBGFVS. See "[Prerequisites and Dependencies](#)" on [page 9](#).

3.1 The Fabric Verification Process

The IBGFVS process of verifying an InfiniBand fabric involves the following steps:

1. **Fabric Design Specification:** This is a fabric profile specification and checklist to be filled out as a first step. See the Appendix "Fabric Design Specification" on [page 30](#). Part of this process involves the definition of a Topology File (see the Appendix "Topology File Description" on [page 19](#)).
2. **Installation:** This step involves the installation of the Master Software on the Master Node and enabling access to the Agent Nodes as required by the Master Software (see "IBGFVS Installation and Setup" on [page 10](#)).
3. **Execution of the IBGFVS verification process:** This process is initiated from the Master Node and takes the Topology File and Performance Threshold Configuration files as input. The Master Node distributes and installs the Software Agents on all the HCA nodes in the fabric, and initiates communication with the Software Agents and Subnet Manager. Next, it ensures that the Software Agents have initialized successfully, and are capable of extracting management information based on requests received from the Master. Then it launches the process of executing fabric verification tests. The process of fabric verification involves the following steps:
 - Checking the Subnet Manager status
 - Performing subnet analysis
 - Topology matching
 - Getting hosts info
 - Checking firmware status
 - Testing fabric stability
 - Checking the status and performance of ULPs

4. **Inspection of Summary and Detailed Reports:** IBGFVS provides a summary report of the fabric status. It provides a status snapshot of each aspect of the fabric that has been verified, based on category groupings such as: Asset Management, Cluster Design Verification, Link Integrity and Performance Verification, Credit Loops Verification, and ULP functional and performance verification (see [Appendix B, “Summary Report Example” on page 21](#)). For items marked as ERROR or FAIL, the Detailed Report provides further details and error codes (see [Appendix C, “Detailed Report Example \(Output Logfile\)” on page 23](#)). In case no errors are reported, then the IBGFVS process is complete.
5. **Troubleshooting:** In case of errors in the Detailed Report, the user can refer to the [IBGFVS Troubleshooting Guide \(page 11\)](#) chapter in this document to find recommended ways for fixing the errors encountered by IBGFVS. The Troubleshooting Guide chapter offers recommendations for action per error code for easy reference.
6. **Rerun IBGFVS verification process:** After fixing all errors reported by the run, repeat steps 3 through 5 above until all errors are resolved.

3.2 IBGFVS Functionality and Steps

The **IBGFVS** tool provides a sequence of operations to be performed on a fabric in order to certify that the fabric is operational and error-free. This sequence is composed of the following steps:

- Step 1. **Subnet Manager Status:** Running the subnet manager and verifying that no critical errors occurred (e.g., duplicate GUIDs).
- Step 2. **Subnet Analysis:** Testing the cluster to verify that no credit loops exist in the topology and no other unpredicted/improper topology items exist in the cluster.
- Step 3. **Topology Matching:** Verifying that the physical topology matches the one specified in the topology file provided by the user.
- Step 4. **Host Info:** Querying cluster nodes for their system and driver information and versions.
- Step 5. **Firmware Status:** Querying cluster devices for their firmware versions.
- Step 6. **Fabric Stability:** Verifying that no bad cables exist in the cluster at a basic check level (without running data in the cluster).
- Step 7. **ULPs - IPoIB:** Verifying for each ULP its interfaces functionality and cluster cables condition under data stress.
- Step 8. **ULPs - MPI:** Verifying for each ULP if its performance meets the minimum requirements of performance (according to a pre-defined minimum bandwidth and lower than a pre-defined maximum latency).

Though passing all steps is required for a ‘healthy fabric’ certificate, the user can run **IBGFVS** steps in several ways: in single steps, in a consecutive subset of steps, or all steps in a row.

Some of the steps of IBGFVS must pass successfully, otherwise IBGFVS exits upon a failure. Steps 4 and 5 may fail and in that case IBGFVS logs warnings and failure information (unexpected results).

3.2.1 IB Granite Fabric Verification Suite (IBGFVS) Synopsis

```
ibgfvs [-h] [-v] [-V] [-o <num-of-step>] [-f <num-of-step>] [-t <num-of-step>]
```

The command line **options** are:

- h for help.
- v for IBGFVS version number.
- V for verbose mode (DEBUG | WARNING | INFORM).
- o perform a single step out of the **IBGFVS** steps above. This option overrides the ‘-f’ and ‘-t’ options.
- f perform **IBGFVS** steps starting from the specified step number. If ‘-f’ is not used, the first step is #1.
- t perform **IBGFVS** steps up to (including) the specified step number. If ‘-t’ is not used, IBGFVS runs up to the last step (inclusive).

4 Prerequisites and Dependencies

4.1 Prerequisites

The following conditions must be met for IBGFVS to operate:

- The fabric must include at least one device which is defined as a switch.
- On Master Node:
 - IB Gold package 1.8.0 is installed.
 - IBGranite package is applied on top of IB Gold.
 - Remote access enabled to all cluster nodes.
- On each Agent Node:
 - Remote access enabled for Master Node.
 - Root permissions for access without password.
- Fabric General (these prerequisites are for the Beta release only)
 - There should be no IB vendor-supplied Subnet Manager or Performance Manager applications running on any nodes in the fabric. The only Subnet Manager active and running in the fabric should be the IB Gold Subnet Manager (OpenSM) running on the Master Node.

4.2 Dependencies

Dependency 1:

The Software Agent may not be successfully installed on an Agent Node due to one or more of the following reasons:

- No remote access available (neither rsh nor ssh)
- The operating system kernel is not supported by the IBGFVS Software Agent

In that case, the following steps in the IBGFVS process will not be executed:

- Gathering hosts information
- Gathering firmware information
- IBGFVS ULP steps will not be run

None of these (unexecuted) steps prevent the full run of the IBGFVS process. Warnings in this regard will be indicated in the Summary and Detailed Reports.

Dependency 2:

If an IB Vendor supplied Performance Manager continues to run or cannot be stopped on any node other than the Master Node, then IBGFVS step 6 will fail.

5 IBGFVS Installation and Setup

IBGFVS should be installed only after all prerequisites are met.

5.1 Installation

To install IBGFVS, run the following script from the installation directory: `./install.sh`

The following is an example of the output of the install script:

```
-----  
IBGranite Fabric Verification Suite installation - Build: 20050606-1112  
Copyright (C) June 2005, Mellanox Technologies Ltd. ALL RIGHTS RESERVED.  
Use of software subject to the terms and conditions detailed in the file:  
"<installation_directory>/LICENSE.txt".  
-----
```

```
IBGCVS installation done. Please run "ibgfvs setup".
```

5.2 Setup

After the installation of IBGFVS, run: `./ibgfvs setup`

6 IBGFVS Troubleshooting Guide

The following sections list the error and warning messages organized by the IB Granite step that reports them. The sections are:

- Troubleshooting General Errors (page 11)
- Troubleshooting IBGFVS Steps (page 12)

6.1 Troubleshooting General Errors

This section covers general errors which either have no clear relation to a specific step or occur during installation/ compilation steps. To fix some of the errors that may be detected, you may need to locate the problematic device.

Table 1 - Troubleshooting General Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E00001	Internal error		<ol style="list-style-type: none"> 1. See further information in log file 2. Contact Mellanox representative
E00002	No space left on device	The volume containing /tmp directory is full. IB Granite stores intermediate and log files in this directory	<ol style="list-style-type: none"> 1. Remove files from /tmp
E00003	Bad configuration	One of the configuration files which controls IBGFVS functionality may have been corrupted or has not been configured correctly	<ol style="list-style-type: none"> 1. Review log file 2. Review all IBGFVS configuration files 3. Re-run IBGFVS

6.1.1 Locating a Problematic Device

If an error is detected, the file /tmp/IB GraniteError.log will contain the full error data. To locate a device, follow the path given in this file starting from the node running IBGranite, through the switches on the path, to the required device. Devices are located using a 'direct route' scheme and the path provided is composed of a list of port numbers. These are output ports on the path to the device. Note that the input ports to which these output ports are connected are not listed.

6.2 Troubleshooting IBGFVS Steps

This section describes the following:

- Troubleshooting Step 1: Subnet Manager Status (page 12)
- Troubleshooting Step 2: Subnet Analysis (page 13)
- Troubleshooting Step 3: Topology Matching (page 14)
- Troubleshooting Step 4: Hosts Info (page 15)
- Troubleshooting Step 5: Firmware Status (page 16)
- Troubleshooting Step 6: Fabric Stability (page 16)
- Troubleshooting Step 7: ULPs - IPoIB (page 17)
- Troubleshooting Step 8: ULPs - MPI (page 17)

6.2.1 Troubleshooting Step 1: Subnet Manager Status

This step checks that the subnet manager is able to discover and configure the subnet. To fix some of the errors that may be detected, you may need to locate the problematic device. See [Section 6.1.1, “Locating a Problematic Device,” on page 11.](#)

Table 2 - Troubleshooting Step 1 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E10002	Bind Error	The SM did not bind to the local IB port, probably because the driver is not up	<ol style="list-style-type: none"> 1. Run <code>"/etc/init.d/openibd restart"</code> 2. Run <code>"vstat"</code> 3. Make sure that the connected IB port is not in the PORT_DOWN state
E10003	SM Port Down	The IB port that the SM is trying to use on the local machine is down	<ol style="list-style-type: none"> 1. OpenSM tries to query the subnet from the local IB port specified in file <code>/etc/osm.conf</code>. Make sure that this IB port is connected. 2. Check that the other end of the cable is connected 3. Check power on the other endnode is on 4. Check that the green LED on the local IB port is on
E10004	SM in Standby	Another SM with higher priority was detected in the fabric	<ol style="list-style-type: none"> 1. Locate the device running the other SM 2. Shut down the other SM
E10005	GUID 0 Detected	The SM detected a device with GUID value 0. GUID is a global unique identification of a node, port, or system in the InfiniBand fabric. Zero is not a valid value for GUIDs. This error is probably caused by a problem in the device firmware burning	<ol style="list-style-type: none"> 1. Locate the device 2. Burn the device firmware specifying the correct GUID value 3. Reboot the device
E10006	Duplicated GUID	GUID is a global unique identification of a node, port, or system in the IB fabric. This error is probably caused by a problem in the device firmware burning	<ol style="list-style-type: none"> 1. Locate one of the devices to be fixed 2. Burn the device firmware specifying the correct GUID value 3. Reboot the device
E10007	Irresponsive Port	The port is physically connected to the IB fabric, but does not respond to SM queries. Probably a hardware/firmware problem	Reboot the unresponsive node. For hosts with pciX AMD chipset 8131 - it may be caused by a known AMD bug described in doc errata 58.
E10008	Set Error	The SM failed setting attributes on a device. This may happen if the fabric is highly unstable	<ol style="list-style-type: none"> 1. Find the path to the problematic device 2. For each link in the path make sure that the green LED on the IB port is not blinking 3. If the LED is blinking, reconnect or change the cable

Table 2 - Troubleshooting Step 1 Errors (Continued)

Error Code	Error Title	Detailed Description	Suggested Action Steps
E11000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvcs.log for details

6.2.2 Troubleshooting Step 2: Subnet Analysis

This step checks for errors in the subnet routing. If an error is detected, the file /tmp/IB GraniteError.log will contain the full error data.

Table 3 - Troubleshooting Step 2 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E20001	Dead-end in CA to CA path	The routing tables are not setup correctly to allow a path from the source HCA to the destination HCA	<ol style="list-style-type: none"> 1. Examine the dead-end list in the log file 2. If all dead-ends are on the same switch then reboot it 3. Restart the SM 4. Inspect the SM log file for errors: if it still contains some errors go back to action step 1 5. If there are no errors in the SM log and still there are CA to CA paths, contact your SM vendor
E20002	Credit loop potential deadlock found	The routing tables are setup such that a credit loop exist. Such credit loops can cause deadlock under high data stress	<ol style="list-style-type: none"> 1. If you are using OpenSM modify the /etc/opensm.conf to use -u flag (up/down algorithm) 2. Contact your SM vendor
E20003	Multicast group is not fully connected	Some switches in the multicast group are disconnected from the rest	<ol style="list-style-type: none"> 1. Double check the SM log file for errors 2. Contact the SM vendor
E20004	Multicast group routing has credit loops	If the multicast group root is not chosen correctly, then data sent on the multicast group may cause a loop deadlock or some nodes may be unreachable	Contact your SM vendor
E21000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvcs.log for details

6.2.3 Troubleshooting Step 3: Topology Matching

The topology matching process compares the actual discovered subnet topology with the one specified in the given topology file. (By default, the file `/etc/ibadm.topo` is used.)

For details about the topology file syntax, please refer to the *IBADM User's Manual (IBADM UM)*, Document no. 2130UM.

Table 4 - Troubleshooting Step 3 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E30001	Failed to parse topology file	Probably dues to a bad topology syntax	<ol style="list-style-type: none"> 1. See error message in the log file to find the error type and line number 2. Fix the error according to the file syntax described in IBADM UM
E30002	Unknown system type in topology file	Topology file declares a system type which is unknown to the topology checker	<ol style="list-style-type: none"> 1. See error message in the log file to find the error type and line number 2. Check the topology file for a possible typo in the system type field
E30003	Root (master) node not found in the topology file	Master node name is either not assigned or not indicated correctly in the topology	<ol style="list-style-type: none"> 1. Check whether the <code>/etc/ibadm.conf</code> <code>IBBS_NAME</code> parameter is assigned correctly 2. Check if in the topology file the master node matches the <code>IBBS_NAME</code> in <code>/etc/ibadm.conf</code>
E30010	Missing system(s)	A system specified in the topology file was not found in the actual subnet	<ol style="list-style-type: none"> 1. Make sure that the system is connected to the right IB port and that it is powered on 2. Make sure that the physical link LED (green LED) of the missing system port is on; if not, try reconnecting the cable. 3. If it is OK for the system to be missing remove it from the topology file
E30011	Extra system(s) found	A system found in the actual subnet is missing from the topology file	<ol style="list-style-type: none"> 1. See error(s) description in IBGranite error log file to get the system name or its GUID, and which port of the system is connected 2. Remove the system from the subnet or add it to the topology file
E30012	Wrong link speed detected	One or more links went up incorrectly operating at SDR rather than DDR	<ol style="list-style-type: none"> 1. Verify firmware version matches the required configuration on both ends of the cable. If not, re-burn firmware and reboot the device(s) 2. Replace the cable connected to the port
E30013	Wrong link width detected	One or more links went up incorrectly as 4X instead of a 12X	<ol style="list-style-type: none"> 1. See error(s) description in IBGranite error log file to get the erroneous link 2. Try reconnecting the cable on both ends
E30014	Too many topology errors	Introduced topology is far from actual cluster topology	<ol style="list-style-type: none"> 1. Review <code>/etc/ibadm.topo</code> file and verify it is the right one for this cluster 2. Rebuild the <code>/etc/ibadm.topo</code> file
E30015	Internal system error(s) detected	A topology mismatch was detected inside a switch system	<ol style="list-style-type: none"> 1. For modular switches make sure all the switch leafs and spines exist indeed and are correctly installed 2. Reboot the system
E30100	Topology mismatch detected	A topology mismatch was detected in the cluster and no current mapping of the error is assigned in this manual	Follow <code>ibgcvs.log</code> errors to fix topology

Table 4 - Troubleshooting Step 3 Errors (Continued)

Error Code	Error Title	Detailed Description	Suggested Action Steps
E31000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

6.2.4 Troubleshooting Step 4: Hosts Info

This step verifies the homogeneity of architectures, operating systems, kernels, etc., and that installed components are fully qualified by Mellanox QA on cluster host nodes.

Table 5 - Troubleshooting Step 4 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E40001	Cannot access host(s)	Access to host(s) listed in log file using rsh or ssh is not enabled	<ol style="list-style-type: none"> 1. Ping listed host(s) from IB Granite Master. If there is no ping response then reset the Ethernet switch 2. Check whether rsh/ssh is enabled on listed host(s)
E40002	Remote agent installation failed	Installation failed	Review install log file as specified in IB Granite log file on remote log file.
E40003	Remote agent activation failed	Activation failed	<ol style="list-style-type: none"> 1. Re-install the listed node(s) 2. Review log file on remote host(s)
E40004	Access to IBGFVS shared directory failed	IBGFVS requires a shared directory in order to distribute the agents. If it is not provided or cannot be accessed then agents will not be operational	<ol style="list-style-type: none"> 1. If running in "export local directory" mode, try running setup again 2. If running on a network directory check the mount to this directory on hosts
W40010	Unqualified machine type	Unqualified configuration according to the Mellanox QA matrix	<ol style="list-style-type: none"> 1. No action is mandatory - this is an informative note, however, it is recommended to migrate to supported architecture. 2. Migrate to supported Architecture.
W40011	Unqualified OS type	Unqualified configuration according to the Mellanox QA matrix	<ol style="list-style-type: none"> 1. No action is mandatory - this is an informative note, however, it is recommended to migrate to supported Operating system (OS). 2. Migrate to supported OS
W40012	Non-homogeneous driver version	Multiple driver versions are installed in the cluster	<ol style="list-style-type: none"> 1. No action is mandatory - this is an informative note, however, it is recommended to have the cluster with the same driver version in all the cluster 2. Update the driver on components with the old version according to the log file list
E40020	Unsupported Architecture/OS	Cluster node(s) has a configuration which will not be able to use any of the existing software components	<ol style="list-style-type: none"> 1. Replace problematic component 2. Contact your SW vendor for assistance
E41000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

6.2.5 Troubleshooting Step 5: Firmware Status

This step verifies the firmware status on the cluster components to for being up-to-date with latest versions as introduced by the IB Granite content.

This will be noted as an alert and not as an error.

Table 6 - Troubleshooting Step 5 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E50001	PSID is not assigned	Configuration file for the system PSID was not found	Burn a firmware image to system with PSID (see IBADM UM)
E50002	Firmware version too updated	Firmware version on cluster node(s) is newer than the firmware directory	Re-install a new /etc/ibfw directory
W50010	Non-homogeneous firmware version	Cluster has multiple firmware versions installed on similar components	<ol style="list-style-type: none"> 1. No action is mandatory - this is an informative note. However, it is recommended to have the cluster with the same firmware version for similar components 2. Update firmware on devices with old versions according to the log file list
W50011	Firmware is not up-to-date	Cluster is installed by old firmware versions	<ol style="list-style-type: none"> 1. No action is must - this is an informative note, it is recommended to have the cluster with the newest version of firmware on all the cluster systems. 2. Update the firmware versions to newer ones according to the latest
E51000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

6.2.6 Troubleshooting Step 6: Fabric Stability

This step will verify in several methods that the cluster connectivity has no faulty cables or ports. Failing this stage terminates the certification process until its issues are resolved.

Table 7 - Troubleshooting Step 6 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E60001	Error on external link connections	An error occurred in one of the counters of the ports	<ol style="list-style-type: none"> 1. Disconnect both ends of cable and reconnect them 2. [Optional] Test cable with Eye-Opener tool 3. Replace cable 4. Reboot the two systems at cable ends 5. Test the link end ports with other ports and locate problematic edge
E60002	Error on internal link connections	An error occurred in one of the counters of the ports	Replace the hardware component
E61000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

6.2.7 Troubleshooting Step 7: ULPs - IPoIB

This step verifies three aspects of the IP-over-IB standard interface: (1) Availability, and (2) functionality of the MPI interface and (3) performance results are better than the required minimum.

Table 8 - Troubleshooting Step 7 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E70001	No active IPoIB interface found	IPoIB interface is not active and cannot be accessed	<ol style="list-style-type: none"> 1. Verify that driver is loaded on the indicated node as specified in the log file 2. Upload the driver
E70002	Ping failed	One of the interfaces does not function correctly	<ol style="list-style-type: none"> 1. Locate problematic node according to log file and output 2. Restart driver on problematic node 3. Reboot the host
E70010	IPoIB ping latency is above limit	Performance between nodes in the cluster does not meet expected limits	<ol style="list-style-type: none"> 1. Locate the problematic path 2. Verify there is no over-subscription on ports in systems 3. Verify that endnode hosts function correctly 4. Reboot all devices on problematic path
E70020	IPoIB ping bandwidth is below limit	Performance between nodes in the cluster does not meet expected limits	<ol style="list-style-type: none"> 1. Locate the problematic path 2. Verify there is no over-subscription on ports in systems 3. Verify that endnode hosts function correctly 4. Reboot all devices on problematic path
E71000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

6.2.8 Troubleshooting Step 8: ULPs - MPI

This step verifies three aspects of the MPI standard interface: (1) Availability, and (2) functionality of the MPI interface and (3) performance results are better than the required minimum.

Table 9 - Troubleshooting Step 8 Errors

Error Code	Error Title	Detailed Description	Suggested Action Steps
E80001	MPI test completion failed	Completion of MPI test failed due to some unknown reason	<ol style="list-style-type: none"> 1. Locate problematic node according to log file and output 2. Restart driver on problematic node 3. Reboot the host
E80002	MPI test run failed	MPI run failed due to a general reason	<ol style="list-style-type: none"> 1. Locate problematic node according to log file and output 2. Verify MPI is installed and running correctly on endnodes 3. Restart driver on problematic node 4. Reboot the host
E80010	MPI test exceeded maximum latency	Performance between nodes in the cluster does not meet expected limits	<ol style="list-style-type: none"> 1. Locate the problematic path 2. Verify there is no over-subscription on ports in systems 3. Verify that endnode hosts function correctly 4. Reboot all devices on problematic path

Table 9 - Troubleshooting Step 8 Errors (Continued)

Error Code	Error Title	Detailed Description	Suggested Action Steps
E80020	MPI test exceeded minimum bandwidth	Performance between nodes in the cluster does not meet expected limits	<ol style="list-style-type: none">1. Locate the problematic path2. Verify there is no over-subscription on ports in systems3. Verify that endnode hosts function correctly4. Reboot all devices on problematic path
E81000	An error occurred while running this step	An error occurred which was not identified as one of the errors above. This may be an error issued by the IBGCVS tool itself, or an error which is related to the specific check	See file ibgcvs.log for details

Mellanox Technologies

Appendix A Topology File Description

The topology file describes the IB connectivity and systems included in the cluster. The default file name is /etc/ibadm.topo. It serves two purposes:

1. Support for arbitrary system names to be later used in every report that IBGFVS generates.
2. Connectivity verification: The specified fabric topology is verified against the discovered one. Mismatch errors resulting from missing cables and/or wrong connections are reported by IBGFVS.

The topology file is composed of “system” sections. Each such section describes the connectivity of one system to other systems in the cluster. The first line of each section is a declaration of the system composed of a system-type, its local system-name section, and optional configuration details. The lines to follow until the next empty line describe the connections between the local system ports to the other systems. The following is a formal definition of a system section syntax. An example is listed afterwards.

```
SYSTEM-TYPE LOCAL-SYSTEM-NAME [CFG: <board Name>=R|Removed|4X|12X, [<board Name>=R|Removed|4X|12X]...]
    LOCAL-PORT-PANEL-NAME -> REM-SYS-TYPE REM-SYS-NAME REM-PORT-PANEL-NAME
    LOCAL-PORT-PANEL-NAME -> REM-SYS-TYPE REM-SYS-NAME REM-PORT-PANEL-NAME
    LOCAL-PORT-PANEL-NAME -> REM-SYS-TYPE REM-SYS-NAME REM-PORT-PANEL-NAME
    ...

SYSTEM-TYPE LOCAL-SYSTEM-NAME [CFG: <board Name>=R|Removed|4X|12X, [<board Name>=R|Removed|4X|12X]...]
    LOCAL-PORT-PANEL-NAME -> REM-SYS-TYPE REM-SYS-NAME REM-PORT-PANEL-NAME
    LOCAL-PORT-PANEL-NAME -> REM-SYS-TYPE REM-SYS-NAME REM-PORT-PANEL-NAME
    ...
...
```

Where:

- Board Name can be something like leafN or spineN (N is a number).
- SYSTEM-TYPE = REM-SYS-TYPE = MTEK43132|MTS2400-12T4|MTS2400-24|MTS2400|MTS9600|MTS14400|MTPB23108|MHX-CEXXX-T|MHXL-CEXXX-T|MTLP23108|MHEL-CFXXX-T|MHEAXX-XT
- LOCAL-SYSTEM-NAME = the name of the system described in this topology file section.
- LOCAL-PORT-PANEL-NAME = a name of the local system port. The numbers printed on the front panel are used together with L<N> for Leaf no. N or S<N> for Spine no. N.
- REM-SYS-NAME = the name of the system connected to the local port.
- REM-PORT-PANEL-NAME = a name of the remote system port. We use the numbers as printed on the front panel and L<N> for Leaf number N or S<N> for Spine number N.

The optional “CFG:” section in the system declaration line describes the special customization of each board of the system. The format of the CFG field for a system that supports N leafs and M spines is:

```
CFG: leaf1=R,leaf2=12X,...leafN=4X, spine1=R, spine2=R, ... spineM=12X.
```

That is the CFG string is a set of comma-separated sub-fields. Each sub-field (if exists) describes some special configuration of a corresponding system board (starting with leaf1 and ending with the last spine).

All switch systems that have plug-in cards support the following syntax for the leaf-cfg and spine-cfg fields:

- A 'D' or space or empty string between the commas stands for the default configuration.
- 'R' or 'Remove' stands for "remove". That is the board is not to be installed in the system.
- 12X stands for a 12X-port board configuration.

Further system specific specialization options for the system board are provided in the following table:

Table 1 - Board Configuration Options And Supported Systems

Option	Mnemonic	MTS9600	MTS14400
Remove the Board	R	Yes	Yes
Use 12X ports	12X	No	Also for spines
Add 4X ports	4X	No	Only for spines

Example. The following is an example of a definition-line in a topology file of the MTS9600 switch system. This switch system can have up to eight leafs and four spines. This example of the MTS9600 lacks ('R') leafs no.6,7 and 8, and lacks spines no. 3 and 4.

MTS9600 PartialGz1 CFG: leaf3=R,leaf5=R,leaf7=R,spine1=R

Mellanox Technologies

Appendix B Summary Report Example

The following is an example of a summary report produced by IBGranite.

IBGranite Cluster Verification Suite - Summary Report

Test conducted: Mon Jun 06 04:25:16 PM IDT 2005
Cluster Name: MyCluster

Asset Management Summary: =====

Hosts OK: 100% (2/2) *
Switches OK: [NA]
Subnet Manager: Failed

Firmware Version: Not Checked
Host OS Version: [NA]
Driver Version: [NA]

ULP Status:
- IPoIB: [NA]
- MPI: [NA]

Credit Loop Verification Summary: =====

Credit Loops: Not Checked
Routing: Not Checked

Cluster Design Verification Summary: =====

Topology specification: Not Checked
Topology matching: Not Checked

Link Stability Verification Summary: =====

Stability: Not Checked
Width: [NA]
Speed: [NA]

ULP Performance Verification Summary:

=====

IPoIB Latency: [NA]

MPI Bandwidth: [NA]

MPI Latency: [NA]

- E- The following error(s) were detected:
- E- E10004: SM in Standby
- E-
- E- Please see log file /tmp/ibgcvs.log for details.

Mellanox Technologies

Appendix C Detailed Report Example (Output Logfile)

```
+-----+
| CLUSTER CERTIFICATION LOG FILE |
+-----+
```

```
+-----+
| Step 1: Subnet Manager status. |
+-----+
| Step 1: Subnet Manager status. - Completed successfully. |
+-----+
```

```
+-----+
| Step 2: Subnet analysis. |
+-----+
| Step 2: Subnet analysis. - Completed successfully. |
+-----+
```

```
+-----+
| Step 3: Topology matching. |
+-----+
| Step 3: Topology matching. - Completed successfully. |
+-----+
```

```
+-----+
| Step 4: Hosts Info. |
+-----+
```

Asset Logging Report

=====

HostId	CpuNum	CpuModel	CpuMhz	MemSize	KernelRev
10.4.3.12	2	Intel(R) XEON(TM)	CPU 2.00GHz	1994.315	2074472 2.6.9-5.ELsmp

Intel Corp. E7500 Memory Controller Hub (rev 03) Red Hat Enterprise Linux AS
release 4 (Nahant)

10.4.3.11 2 Intel(R) XEON(TM) CPU 2.00GHz 2000.135 2074472 2.6.9-5.ELsmp
Intel Corp. E7500 Memory Controller Hub (rev 03) Red Hat Enterprise Linux AS
release 4 (Nahant)

10.4.3.10 2 Intel(R) XEON(TM) CPU 2.00GHz 1994.380 2074472 2.6.9-5.ELsmp
Intel Corp. E7500 Memory Controller Hub (rev 03) Red Hat Enterprise Linux AS
release 4 (Nahant)

10.4.3.9 2 Intel(R) XEON(TM) CPU 2.00GHz 1994.354 2074472 2.6.9-5.ELsmp
Intel Corp. E7500 Memory Controller Hub (rev 03) Red Hat Enterprise Linux AS
release 4 (Nahant)

```
+-----+
| Step 4: Hosts Info. - Completed successfully. |
+-----+
```

```
+-----+
| Step 5: Firmware status. |
+-----+
```

-W50011- Firmware is not up to date

The following devices need firmware update:

rhino/leaf12/U1 rhino/leaf1/U1 rhino/leaf2/U1 rhino/spine2/U2 rhino/
leaf7/U1 rhino/leaf10/U1 rhino/leaf9/U1
rhino/leaf4/U1 rhino/leaf6/U1 rhino/leaf3/U1 rhino/leaf5/U1
mtlm_reindeer/U1 rhino/spine2/U3 rhino/leaf11/U1
rhino/leaf8/U1 rhino/spine2/U1

-E50002- FW version on a device is newer than provided FW

The following devices FW is newer than the provided FW:

mtlmd12/U1 mtlmd11/U1 mtlmd09/U1 mtlmd10/U1

-W50010- Non-homogeneous firmware version

The following device type(s) have non homogeneous firmware versions:

MT47396: 45.57.38 44.57.38

Firmware query report:

=====

GUID	DEVICE	NAME	VER	FW-CONF
PSID	STATUS	REQ-VER	FW-FILE	
0x0002c901097c83c0	MT47396	mtlm_reindeer/U1	45.57.38	Reindeer4X
MT_0060000001	NEEDS_UPGRADE	62.79.64	/etc/ibfw/fw-47396/IS3FW.BIN	

```

0x0002c9000100d887 MT23108 mtlmd09/U1      16.17.11 MHX-CE128-T
MT_0000000001 UP_TO_DATE      5.25.15 /etc/ibfw/fw-23108/fw-23108-a1-rel.mlx
0x0002c901097642f0 MT23108 mtlmd10/U1      16.17.11 MHX-CE128-T
MT_0000000001 UP_TO_DATE      5.25.15 /etc/ibfw/fw-23108/fw-23108-a1-rel.mlx
0x0002c901093dad50 MT23108 mtlmd11/U1      16.17.11 MHX-CE128-T
MT_0000000001 UP_TO_DATE      5.25.15 /etc/ibfw/fw-23108/fw-23108-a1-rel.mlx
0x0002c9000100d050 MT23108 mtlmd12/U1      16.17.11 MHX-CE128-T
MT_0000000001 UP_TO_DATE      5.25.15 /etc/ibfw/fw-23108/fw-23108-a1-rel.mlx
0x0002c9010ab1f0a0 MT47396 rhino/leaf1/U1    44.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1ebd0 MT47396 rhino/leaf10/U1   44.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1d7b0 MT47396 rhino/leaf11/U1   44.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1cc50 MT47396 rhino/leaf12/U1   45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010able5b0 MT47396 rhino/leaf2/U1    44.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1d4a0 MT47396 rhino/leaf3/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010b0242b0 MT47396 rhino/leaf4/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1cfd0 MT47396 rhino/leaf5/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1d190 MT47396 rhino/leaf6/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010b024400 MT47396 rhino/leaf7/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010able540 MT47396 rhino/leaf8/U1    44.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab1da50 MT47396 rhino/leaf9/U1    45.57.38 Rhino4XLeaf
MT_0070000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab8a720 MT47396 rhino/spine2/U1   45.57.38 Rhino4XSpine
MT_0080000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab8a728 MT47396 rhino/spine2/U2   44.57.38 Rhino4XSpine
MT_0080000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN
0x0002c9010ab8a730 MT47396 rhino/spine2/U3   45.57.38 Rhino4XSpine
MT_0080000001 NEEDS_UPGRADE 62.79.64 /etc/ibfw/fw-47396/IS3FW.BIN

```

```

+-----+
| Step 6: Fabric stability. |

```

```

+-----+
-E60001- Error in link connections
        See link stability report below.

```

Link stability report:

```
=====
```

```
-E- 17:01:57 Threshold Exceeded on Cable: mtlm_reindeer/P9 <= mtlmd09/P1
```

```

+-----+
| Step 7: ULPs - IPoIB. |
+-----+

```

```

-E70010- Ping exceeded maximum latency
        The following hosts failed ping:
        From                To

```

IPOIB Ping Report (results in uSeconds):

```
=====
```

Pairs checked:	16
Pairs failed ping:	0
Pairs exceeded minimum ping latency (1000)	4
Results min / avg / max:	60.0 / 488.6875 / 1531.5
Minimum ping pair:	10.4.3.11 -> 10.4.3.12
Maximum ping pair:	10.4.3.11 -> 10.4.3.10

```

+-----+
| Step 8: ULPs - MPI. |
+-----+

```

-E80020- MPI test exceeded minimum bandwidth

MPI Test minimum bandwidth (550.433466) is below 600 MB/Second.

MPI Test Report:

=====

 Latency-Bandwidth-Benchmark R1.3 (c) HLRS, University of Stuttgart
 Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Details - level 2

Message Length: 8

Latency	min / avg / max:	0.006625 /	0.006625 /	0.006625 msecs
Bandwidth	min / avg / max:	1.208 /	1.208 /	1.208 MByte/s

message size:

max time :	50.000000 sec
latency for msg:	0.006625 msec
estimation for ping pong:	1.192500 msec
max number of ping pong pairs	= 41928
max client pings = max server pongs	= 204
stride for latency	= 1

Message Length: 8

Latency	min / avg / max:	0.006500 /	0.006771 /	0.006937 msecs
Bandwidth	min / avg / max:	1.153 /	1.182 /	1.231 MByte/s

Message Length: 2000000

Latency	min / avg / max:	3.471500 /	3.471500 /	3.471500 msecs
Bandwidth	min / avg / max:	576.120 /	576.120 /	576.120 MByte/s

message size:

max time :	150.000000 sec
latency for msg:	3.471500 msec
estimation for ping pong:	41.658000 msec
max number of ping pong pairs	= 3600
max client pings = max server pongs	= 60
stride for latency	= 1

Message Length: 2000000
 Latency min / avg / max: 3.335000 / 3.531083 / 3.633500 msecs
 Bandwidth min / avg / max: 550.433 / 567.243 / 599.700 MByte/s

Message Size: 2000000 Byte
 Natural Order Bandwidth: 341.028681 MB/s
 Avg Random Order Bandwidth: 340.974601 MB/s

Execution time (wall clock) = 69.590 sec on 3 processes
 - for cross ping_pong latency = 0.010 sec
 - for cross ping_pong bandwidth = 0.340 sec
 - for ring bandwidth = 69.240 sec

 Latency-Bandwidth-Benchmark R1.2 (c) HLRS, University of Stuttgart
 Written by Rolf Rabenseifner, Gerrit Schulz, and Michael Speck, Germany

Major Benchmark results:

Max Ping Pong Latency: 0.006937 msecs
 Min Ping Pong Bandwidth: 550.433466 MB/s
 Naturally Ordered Ring Bandwidth: 341.028681 MB/s
 Randomly Ordered Ring Bandwidth: 340.974601 MB/s

 Detailed benchmark results:

Ping Pong:

Latency min / avg / max: 0.006500 / 0.006771 / 0.006937 msecs
 Bandwidth min / avg / max: 550.433 / 567.243 / 599.700 MByte/s

Ring:

Benchmark conditions:

The latency measurements were done with 8 bytes
 The bandwidth measurements were done with 2000000 bytes
 The ring communication was done in both directions on 3 processes
 The Ping Pong measurements were done on
 - 6 pairs of processes for latency benchmarking, and

- 6 pairs of processes for bandwidth benchmarking,
out of $3 \times (3-1) = 6$ possible combinations on 3 processes.
(1 MB/s = 10^6 byte/sec)

Mellanox Technologies

Appendix D Fabric Design Specification

Attached is the 'Fabric Profile' specification. For a soft copy of this document, please contact your local FAE.

Mellanox Technologies



Fabric Profile

Fabric ID:

Rev 1.0



Dear Customer,

The purpose of this “*Fabric Profile*” document is to ensure proper and timely InfiniBand-based fabric bring-up and support. Specifically, it aims to help Mellanox’s Technical Support team to:

- Identify in advance potential problems relating to fabric setup, equipment, and/or implementation schedule.
- Identify non-standard installation aspects.
- Prepare and act quickly for the bring-up and/or maintenance procedures.

Please fill in the requested information accurately. Mellanox’s team will be using this fabric profile to prepare in the best manner for your fabric installation.

1. Contact Information

Customer Information:

Location	Name	Title	Contact Info.

Mellanox Team Information:

Location	Name	Title	Contact Info.



2. Fabric Purpose and Required Benchmarks

Examples:

- HPC (Linpack, HPCC, others)
- Industrial HPC (LSDyna, Fluent, MM5, others)
- Other

Fabric Purpose	Required Benchmarks

3. Schedule

Item	Due Date
Fabric room preparation:	
Power	
Cooling	
Racks arrival	
Machine arrival	
HCA arrival	
Switch arrival	
Cables arrival	
HW installation	
SW installation	
Integration & Tuning	



4. Fabric Information

Fabric location			
Fabric size			
Target performance			
HW Profile	Computer Info		
	Brand		
	Number and type of CPU		
	Mem Type & Size		
	Chipset		
	BIOS version		
	Storage system		
	PCI bus type and speed		
	IB HCA		
	Type		
	Supplier		
	Version		
	Firmware version		
	IB Switch		
	Type		
	Supplier		
	Version		
	Firmware version		
	IB Cables		
	Type		
	Length/s		
	Cabling Schema		
	Cabling conduits		
	IB link		
	Width (4x, 12x)		
	Speed (2.5/5 GB/sec)		
	Racks		
	Size		
	Type		
	Cooling		
	Power		
	SW Profile	Operating System	



	Distribution	
	Kernel version	
	gcc version	
	Fortran compiler type and version	
	Installed Packages profile (full/minimal etc)	
	IB SW	
	Vendor	
	Type	
	Version	
	ULPs	
	MPI Type and version	
	SDP	
	SRP	
	DAPL	
	Fabric Use Static IP/DHCP	
	Management IBADM Other	

5. Tasks Prior to Physical Installation:

Item for checking	Done?
IBADM configuration files: Topofile Host file	
Single node performance check Linpack	
Two nodes performance (perf_main) test: BW Latency	



6. Additional Information