



Release Notes

# InfiniBand Administration Tools (IBADM)

Rev 1.7.0

Mellanox Technologies

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

## InfiniBand Administration Tools (IBADM) Release Notes

### Document Number:

Mellanox Technologies, Inc.

2900 Stender Way

Santa Clara, CA 95054

U.S.A.

[www.Mellanox.com](http://www.Mellanox.com)

Tel: (408) 970-3400

Fax: (408) 970-3403

Mellanox Technologies Ltd

PO Box 586 Hermon Building

Yokneam 20692

Israel

Tel: +972-4-909-7200

Fax: +972-4-959-3245

Mellanox Technologies

# 1 Overview

These are the release notes for Rev 1.7.0 of IBADM, a package of InfiniBand ADMINistration tools. The document includes the following sections:

- This “Overview” which describes: features and tools in this release, software and firmware dependencies, supported platforms and operating systems, and supported interfaces of the various systems.
- “Known Issues” on page 7
- “Bug Fixes” on page 9

## 1.1 Features and Tools

Detailed installation instructions along with complete descriptions of the various tools in the package can be found in the IBADM User’s Manual, Document no. 2130UM. The following is a list of the available tools in the package, together with a brief description of what each tool performs.

<b>ibtopogen</b>	A cluster design tool for fat-tree topologies
<b>ibdmchk</b>	Checks the topology and provides the means to diagnose the IB routing
<b>ibls</b>	Lists the available IB systems, devices, and cables in the IB Subnet
<b>ibmon</b>	Monitors all the devices and systems (for both IB and system related issues)
<b>ibfwmgr</b>	Loads and tracks the IB devices firmware
<b>ibcon</b>	Provides a view of some internal configuration registers of devices
<b>ibcert</b>	A cluster certification tool

## 1.2 Firmware Dependencies

Table 1 - Firmware Dependencies

Silicon Type	Required Firmware Version
MT23108 InfiniHost	3.2.0 or later
MT25208 InfiniHost III Ex (InfiniHost mode)	4.5.2 or later
MT25208 InfiniHost III Ex (MemFree mode)	5.0.1 or later
MT43132 InfiniScale	5.2.0 or later
MT47396 InfiniScale III	0.3.2 or later

## 1.3 Software Dependencies

Table 2 - Software Dependencies

Software Package	Required Version
Mellanox HCA Driver (thca)	3.2.2 or later
Transport Access Layer and IPoIB	1.6.1 or later
MST (Mellanox Software Tools)	3.3.7 or later
OpenSM	0.3.2 or later
Perl	5.6 or later
Expat	1.95 or later

## 1.4 Supported Platforms and Operating Systems

Table 3 on page 5 lists all supported platforms and operating systems by the tools included in this IBADM package.

Table 3 - Supported Platforms and Operating Systems

Platform	Operating System	Kernel
<b>X86</b>	Red Hat Enterprise Linux AS 3.0	2.4.21-20.ELsmp
	Red Hat Linux 9.0	kernel.org: 2.4.27 (smp)
	Red Hat Linux 9.0	2.4.20-8 (smp; bigmem)
	SuSe SLES 9.0	Update (2.6.5-7.111.xx-smp)
	SuSE Linux 9.1 Pro	2.6.9 / 2.6.10
	SuSE Linux 9.1 Pro	Update (2.6.5-7.111.xx-smp)
	Rocks 3.3.0	2.4.21-20.ELsmp
	Fedora Core 3	vanilla 2.6.9 (from kernel.org)
<b>IA-64</b>	Red Hat Enterprise Linux AS 3.0	2.4.21-15.EL
	SuSE SLES 9.0	2.6.5-7.97-default
<b>AMD64 (X86_64)</b>	Red Hat Enterprise Linux AS 3.0	2.4.21-20.ELsmp
	SuSE SLES 9.0	2.6.5-7.111.xx-smp
	SuSE 9.1 Pro	Update (2.6.5-7.111.xx-smp)
	SuSE 9.1 Pro	2.6.9/2.6.10
	Rocks 3.3.0	2.4.21-20.ELsmp
	Fedora Core 3	2.6.9-1.667smp
<b>Intel EM64T</b>	Red Hat Enterprise Linux AS 3.0	2.4.21-20.EL
	SuSE SLES 9.0 RC5	2.6.5-7.97-smp / 2.6.5-7.111.xx-smp
	SuSE 9.1 Pro	2.6.10
	SuSE 9.1 Pro	Update (2.6.5-7.111.xx-smp)
	Rocks 3.3.0	2.4.21-20.ELsmp
	Fedora Core 3	2.6.9-1.667smp

## 1.5 Interfaces Used for Tool Access

The following table lists the interfaces used for tool access for Mellanox boards and systems. The table uses the following abbreviations:

<b>IB</b>	(In-Band)	Application runs over InfiniBand ( <b>ibfwmgr</b> may use the IPoverIB functionality)
<b>IP</b>	(Out-of-Band)	Application runs over Ethernet
<b>I2C</b>	(Out-of-Band)	Application runs using I2C bus

Table 4 - Interfaces for Tools Access on Mellanox Boards and Systems

Tool	Board or System Support				
	HCA Cards: MHX-CExxx-T (Cougar) MHXL-CFxxx-T (Cougar Cub) MHEL-CFxxx-T (Lion) MHEA28-XT (Lion Mini)	MTS9600 (Gazelle) Switch System	MTS2400 (Reindeer) Switch System	MTS2400M (Reindeer- Managed) Switch System	MTS14400 (Rhino) Switch System
ibls	IB / IP <sup>1</sup>	IB / IP	IB / I2C	IB / IP	IB / IP
ibmon / pm <sup>2</sup>	IB / IP	IB / IP	IB / I2C	IB / IP	IB / IP
ibmon / bm <sup>3</sup>	NA	IB / IP	IB / I2C	IB / IP	IB / IP
ibfwmgr	IP	IP	IB / I2C	IB / IP	IB / IP
ibcon	IB / IP	IB <sup>4</sup>	IB / I2C	IB / IP	IB
ibcert	IB / IP	IB / IP	IB / I2C	IB / IP	IB / IP

1. IP may be native IP or IPoverIB
2. Performance check.
3. Baseboard check.
4. **ibcon** enables accessing an individual InfiniBand device (e.g., an MT43132 InfiniScale or MT47396 InfiniScale III device) within the system by its GUID, but not accessing the whole switch system.

## 2 New Features and Tools

### 2.1 New Tools

- **ibdmchk** - Checks the topology and provides the means to diagnose the IB routing. See the *IBADM User's Manual, Document no. 2130UM*, for details.
- **ibcert** - Certifies that a cluster is working properly by running various checks in a specific order. See the *IBADM User's Manual* for details.
- **mlxburn** - Intended for single node burning and firmware image creation. See the *IBADM User's Manual* for details.
- An IBADM Configuration Wizard was added to assist in the automatic configuration of IBADM for different usages. See the *IBADM User's Manual* for details.

### 2.2 New Features

- **ibfwmgr** now displays burn progress messages for each device.

- The MST (Mellanox Software Tools) package is now installed as part of the IBADM installation. It is no longer necessary to install it separately.

## 3 Known Issues

The following table provides a list of known bugs and limitations in regards to this release of IBADM.

Table 5 - Known Bugs and Limitations

	Description	Details
1.	The -I and -L options of <b>ibls</b> do not display any output	<b>ibls -I</b> and <b>ibls -L</b> do not produce any output if Timeout value in <code>/etc/ibadm.conf</code> is too small. The user can raise the Timeout value to get output.
2.	<b>ibfwmgr</b> does not support the reset option	To be fixed in 3.0.0 release.
3.	After HCA firmware burning using <b>ibfwmgr</b> or <b>mlxburn</b> , the <b>infiniburn</b> tool cannot read the firmware image from the device	<b>ibfwmgr/mlxburn</b> replaces the <b>infiniburn</b> Mellanox tool provided as part of the MST package. It is recommended to stop using <b>infiniburn</b> and to use <b>ibfwmgr</b> or <b>mlxburn</b> instead.
4.	<b>ibfwmgr</b> may hang if a device stops responding during burn	If the IBADM server running on the device undergoing a burn operation stops responding (due to an IBADM server crash, device reset or network problems), <b>ibfwmgr</b> may hang. HCAs will hang forever. Switches burnt In-Band will recover after a few minutes. This scenario can be detected if the burn process does not progress for more than one minute.  Workaround: Kill <b>ibfwmgr</b> , fix the problem, and re-initiate the burn.
5.	<b>ibfwmgr</b> may report incorrect percentile progress while burning multiple MT47396 InfiniScale III switch devices	This may occur if a previous FW burn session was prematurely aborted. Note that this erroneous percentile reporting of progress does not have any real impact on the burning process itself. However, it may double the duration of burning.
6.	LID re-assignment during an In-Band device burn may lead to firmware image corruption on devices in the cluster	If the SM re-assigns port LIDs while <b>ibfwmgr</b> is burning a device via InfiniBand (In-Band), the burn directives will be routed to the wrong device leading to the corruption of its firmware image.  Workaround: Do not re-assign LIDs while <b>ibfwmgr</b> is running.
7.	Limited Baseboard-Management-Agent (BMA) access and support	Using <b>ibmon</b> to get BMA data is restricted to Out-Of-Band and for MTS2400 or MTS14400 switch systems. It will be supported in the 2.0.0 release.
8.	Switch SMA Port 0 is not reported in <b>ibmon</b> logs	The <b>ibmon</b> option for querying Switch SMA Port 0 is not functioning
9.	Accessing a switch device undergoing an In-Band burn via the management port may cause the burn to fail	This scenario may cause a contention on the internal I2C-compatible bus. In some cases it may require the switch system to reboot.  Workaround: Do not access a switch system via its management port when one of its devices is undergoing an In-Band burn operation

	Description	Details
10.	<b>ibcert</b> scalability with data transfer	<b>ibcert</b> steps 6 and 7 (see <i>IBADM User's Manual</i> ) involve data transfers between all nodes of a cluster. On a cluster with more than 16 nodes, these steps will take a significant time. In the mean time, it is recommended to run 'ibcert -f 1 -t 5'. Will be fixed in the 2.0.0 release.
11.	<b>ibcert</b> requires a cluster to include at least one switch device	
12.	<b>mlxburn</b> supports burning HCA devices only	
13.	<b>mlxburn</b> supports burning Mellanox boards only	For a board of another vendors, please consult the board vendor
14.	IBADM cannot run on a machine with the tclreadline package installed	To workaround the problem, uninstall the tclreadline package

Mellanox Technologies

## 4 Bug Fixes

The following table lists the bugs that were fixed in this Rev 1.7.0 of IBADM.

Table 6 - Bug Fixes

	Description	Details
1.	<b>ibadm</b> scalability for more than 256 nodes problems	The problem was encountered with a number of open sockets in-flight
2.	Internal topology error in MTS14400 when matching topology	The internal topology had an error
3.	<b>ibls</b> response time improved	Topology matching response time improved dramatically
4.	<b>ibfwmgr</b> freezes while burning a cluster with over 32 nodes	An attempt to burn a cluster with over 32 nodes did not end
5.	Error handling methods fixed to log and catch all errors	
6.	<b>ibfwmgr -s</b> flag was not functional	Further verification on file format is now available and error messages can be listed to the user if requested
7.	<b>ibfwmgr -q</b> halts if any PSID error on any device occurs	<b>ibfwmgr -q</b> was not working properly when PSID was wrong for one of the systems. <b>ibfwmgr</b> query/burn is now fixed to continue as long as there are error-free targets. An error log is printed in the end.
8.	Servers failing to start and/or recover from errors	Servers could not recover from a failure such as getting an empty list in response to a query

Mellanox Tech

Mellanox Technologies