

# **Voltaire InfiniBand GridStack Software Revision 4.3.0**

## **Release Notes**

**August 22, 2007**

## CONTENTS

1	Introduction.....	1
2	Release Description .....	1
3	Release Contents .....	1
4	Supported HCAs.....	1
5	Supported Operating Systems.....	2
6	Supported CPU Architectures.....	2
7	Performance Envelope .....	3
8	SW Components Implementations Status.....	4
9	Prerequisites.....	4
10	Installation of the Voltaire GridStack™ Package .....	4
11	Limitations .....	4
12	Patches Documentation .....	4
13	Known Issues .....	7

## 1 Introduction

This document provides the release notes for Voltaire InfiniBand Stack, Revision 4.3.0.

Please note some sections of this document are considered to be a technology preview or beta.

## 2 Release Description

This Voltaire GridStack™ package is an Open Fabrics EWG OFED based distribution designed to ease installation and management of the Voltaire InfiniBand Stack (or InfiniBand Host Stack).

This distribution is based on the OFED 1.2 GA release.

Please refer to the Voltaire HCA 4X0 User Manual for further detail.

## 3 Release Contents

Item	Description
Binary package compressed as a.tar file	4.3.0_11
Documentation	Linux GridStack™ for HCA 4X0 User Manual (399Z00121)

## 4 Supported HCAs

### 4.1 Voltaire HCAs

The following Voltaire HCAs are supported in this release:

VOLT Name	VOLT PN	MLNX PN
HCA400	501S12319	MHET2X-1TC
HCA400-EX	501S12317	MHEL-CF128-TC
HCA410-EX	501S00045	MHES18-XTC
HCA400-EX-D	501S00047	MHGA28-1TC
HCA410-Ex-D	501S00046	MHGS18-XTC
HCA400-EX-M	501S00078	MHEA28-XTC
HCA400-EX-D-M	501S00071	MHGA28-XTC

For the list of specific Mellanox names according to P/N  
[http://www.mellanox.com/support/firmware\\_table.php](http://www.mellanox.com/support/firmware_table.php)

Please note that Mellanox P/N are regularly updated and therefore may vary. Check the Mellanox website for the latest Part Numbers.

Note that the Voltaire firmware burning tool only supports the Voltaire HCAs.

## 4.2 Other HCAs

Voltaire support HCAs based on the following Mellanox HCAs:

- Single port InfiniHost III Lx PCI Express HCA cards (Cheetah / Tiger)
- Dual port InfiniHost III Ex PCI Express HCA cards (Lion Cub / Lion Mini)
- Dual port InfiniHost PCI-X HCA cards (Cougar Cub / Cougar)

## 4.3 Voltaire HCA Firmware Versions

1. HCA400 – 3.5.0
2. HCA410, HCA410-D – 1.2.0
3. HCA400-EX, HCA400-EX-D – 4.8.2
4. HCA400-EX-M, HCA400-EX-D-M – 5.2.0

## 5 Supported Operating Systems

1. RHEL 4 UP 5 – kernel 2.6.9.55-EL
2. RHEL 5 – kernel 2.6.18.8-EL5
3. Suse SLES 10 sp 1 – kernel 2.6.16.46-0.12

## 6 Supported CPU Architectures

1. x86\_64 (AMD & EM64t)
2. x86
3. ia64
4. ppc64

## 7 Performance Envelope

Measurement		Platform		
Test Name	Units	X86 PCI-X SDR (Tavor)	x86_64 PCI-E SDR (Arbel)	x86_64 PCI-E DDR (Arbel)
ib_write_lat	μSec	4.67		2.91
ib_rdma_bw	Mbits/Sec	4729		11268
ib_rdma_bw_bidirection	Mbits/Sec	6611		22013
ib_rdma_lat	μSec	4.67095		2.84679
ib_send_bw	Mbits/Sec	6535	7480	11238
ib_write_bw	Mbits/Sec	6578	7480	11279
ib_read_bw	Mbits/Sec	6462	7312	10873
ib_read_lat	μSec	10.53	6.6	5.82
ib_send_lat	μSec	6.56	4.34	3.89
iperf	Mbits/Sec	1090		2850
netperf/rds	Mbits/Sec	3041		6722
netperf/ipoib	Mbits/Sec	1522		2262
netperf/ipoib-cm	Mbits/Sec	3795		5987
netperf/sdp	Mbits/Sec	1286	3234	7414
voltaire_mpi_latency	μSec	6.0205	4	3.7895
voltaire_mpi_bandwidth	Mbits/Sec	6888	7800	11461
open_mpi_latency	μSec	6.508946	4.8	3.95
open_mpi_bandwidth	Mbits/Sec	6891	7600	11425

## 8 SW Components Implementations Status

1. Core IB support – GA
2. IPoIB-UD – GA, IPoIB-CM - beta
3. Voltaire MPI – GA
4. Open MPI – beta
5. RDS – beta
6. SDP – beta
7. Bonding driver – beta
8. GVD (GridVision Daemon) - GA
9. HIS (Host Identification Service) - GA

## 9 Prerequisites

1. One of the supported operating systems must be running on one of the supported CPU architectures.
2. For RHEL AS4, the `sysfs-utils` package must be already installed.
3. Before installing GridStack, any existing InfiniBand stack installation on your system must be first uninstalled.

## 10 Installation of the Voltaire GridStack™ Package

For information on how to install the Voltaire GridStack™, please refer to the Voltaire HCA 4X0 User Manual.

## 11 Limitations

Supports DHCP server with InfiniBand related changes [Internet Systems Consortium (ISC) DHCP Server V3.0.5] provided by Voltaire.

## 12 Patches Documentation

On top of the original OFED version 1.2, Voltaire added the following patches to this release:

Patch File Name (patches/Common/fixes/kernel/)	Patch Description
zzz_0010_cm_sidr_1.patch	<p><b>Subject:</b> Duplicate SIDR REQs.</p> <p><b>Bug:</b> The system sent reject messages if a duplicate was detected.</p> <p><b>Fix:</b> Duplicates are now simply discarded.</p>
zzz_0020_cm_sidr_2.patch	<p><b>Subject:</b> SIDR REQ not matching a listen.</p> <p><b>Bug:</b> Dropped through to the default case of status 2 (Rejected by Service Provider).</p> <p><b>Fix:</b> Replies with status value 1 (Service ID not Supported). This also fixes a bug where the cm_id_priv is removed from the remote_sidr_table twice.</p>
zzz_0030_vlt_cma_tavor_quirk.diff	<p><b>Subject:</b> For Tavor based HCAs.</p> <p><b>Goal:</b> To increase performance.</p> <p><b>Fix:</b> Override the MTU that is returned by the SM in an answer to a Path Query with 1024.</p>
zzz_0400_ib_find__partial_pkey.diff	<p><b>Subject:</b> P_Key lookup.</p> <p><b>Goal:</b> For matching full and partial membership keys of the same partition.</p> <p><b>Fix:</b> IPoIB sets the P_Key membership bit of limited membership P_Keys when creating a child interface. After that IPoIB looks for the full membership P_Key in the table to make the interface "RUNNING". This patch fixes the pkey lookup in order to match full and partial membership keys that belong to the same partition.</p>

<b>Patch Name</b> (patches/Common/fixes/user/)	<b>x86 with PCI-X</b>
zzz_0020_use_pf_sdp.diff	<b>Subject:</b> Fallback of SDP connection. <b>Goal:</b> Solve issue of libsdp not supporting a non-blocking connect. <b>Fix:</b> Since libsdp does not support a non-blocking connect in "both" mode, change the fallback of SDP connection from both to SDP.
zzz_0010_mstflint2.diff	<b>Subject:</b> mstflint screen output for PowerPC architectures. <b>Fix:</b> Fixed the screen output.



## 13 Known Issues

This section contains information on known limitations of the current version.

#	Subject	Description	Workaround
1.	Previous/Native InfiniBand modules backup	If a previous local installation of Open-IB is present on a host when installing the Voltaire GridStack™ package, it will be overridden by the new set of modules. When GridStack is removed, the original modules will become loadable again.	The command <code>modinfo &lt;module_name&gt;</code> tells which of the InfiniBand modules will be loaded when running <code>modprobe (ro insmod)</code> . Voltaire modules are installed under <code>/lib/modules/\$(uname -r) /updates/kernel/drivers/infiniband</code>
2.	Compiling 3rd party kernel modules that use InfiniBand symbols	The installation of GridStack does not update the file <code>Module.symvers</code> with the new signature of the InfiniBand symbols (the ones that come with the new InfiniBand kernel modules). The result is that kernel modules that use these symbols and compiled against the newer modules will not be usable since <code>modprobe/insmod</code> will fail loading them into the kernel.	install GridStack with the option <code>--fix-symvers</code> (Builder only)
3.	Using ib-bonding for Ethernet and InfiniBand together	When using bonding persistent configuration for InfiniBand slaves and having another bonding interface for Ethernet slaves, the operation of stopping GridStack cause all the bonding masters to go down, including the one that enslaves the Ethernet devices.	None

#	Subject	Description	Workaround
4.	Using bonding and GVD at the same time	If bonding uses child (P_Key) interfaces as slaves, then GVD auto-configuration of child interfaces might disturb the correct operation of ib-bonding.	Turn off GVD auto configuration by editing the file <b>/etc/gvd.conf</b>
5.	Bonding slaves and IP configuration	The bonding master interface should be configured with an IP address (via the ib-bond utility but not with a network script). So when using an IPoIB interface as a slave of a bonding interface, it <b>MUST NOT</b> have an IP address (i.e. configuration scripts with IP addresses).	None
6.	Limitations of memory pinning from user mode InfiniBand applications	Memory registration by user is limited according to the administrator settings.	Memory locking is managed by the kernel on a per-user basis. Regular users (as opposed to root users) have a limited number of pages that they may pin, where the limit is pre-set by the administrator. Registering memory for IB verbs requires pinning memory, thus an application cannot register more memory than it is allowed to pin. The user can change the system per-process memory lock limit by adding the following two lines to the file <code>/etc/security/limits.conf</code> : <ul style="list-style-type: none"> <li>* soft memlock &lt;number&gt;</li> <li>* hard memlock &lt;number&gt;</li> </ul> where <number> denotes the number of KBytes that may be locked by a user process.

#	Subject	Description	Workaround
			The above change to <b>/etc/security/limits.conf</b> will allow any user process in the system to lock up to <number> KBytes of memory. On some systems, it may be possible to use "unlimited" for the size to disable these limits entirely.
7.	ifconfig reports wrong HW address		<pre>linux:/root # ip address show dev ib0 5: ib0: &lt;BROADCAST,MULTICAST,UP&gt; mtu 2044 qdisc pfifo_fast qlen 128 link/infiniband 00:00:04:04:fe:80:00:00:00:00:00:00:0 8:f1:04:03:96:08:79 brd 00:ff:ff:ff:12:40:1b:ff:ff:00:00:00:00:00 :ff:ff:ff</pre> <pre>inet 193.168.70.16/24 brd 193.168.70.255 scope global ib0 inet6 fe80::208:f104:396:879/64 scope link valid_lft forever preferred_lft forever</pre>
8.	Using fork() with InfiniBand applications	<p>Using fork() in a program that uses InfiniBand is limited to the following conditions:</p> <ol style="list-style-type: none"> <li>1. Parent process may continue running without any limitations on memory access</li> <li>2. Child process gets a SISEGV signal (segmentation fault) when trying to access a memory that was registered by the parent</li> </ol>	Fork support from kernel 2.6.16 and above is available provided: that applications do not use threads. The fork() is supported as long as the parent process does not run before the child exits or calls exec(). The former can be achieved by calling wait(childpid), and the latter can be achieved by application specific means. The Posix system() call is supported.

#	Subject	Description	Workaround
9.	Host vs. VSM HIS interoperability	A host with more than one InfiniBand interface may appear on switch tables with the same IP address for its different ports.	This should be fixed in future switch software releases.
10.	ib1 gets the same configuration as ib0 on SLES10	Under SLES10, when configuring ib0, ib1 gets the same configuration. This happens even if /etc/sysconfig/network/ifcfg-ib1 does not exist.	Create a network configuration file for ib1 even if it is not connected to the network. In process of fixing that with Novell.
11.	HCA IRQ line may block on AMD with old BIOS installed	HCA may stop functioning and Kernel log (printed by dmesg) may show: "irq 169: nobody cared (try booting with the "irqpoll" option)". Also, most IB commands (like ibv_devinfo) may cause the shell to hang. This seems to happen only on AMD machines with old BIOS.	Upgrade the BIOS or add the noirqdebug option to the kernel boot line (in grub.conf).
12.	SLES10 YAST does not support IPoIB configuration.	On SuSE Linux distribution, the YAST setup tool does not recognize the InfiniBand interface. Therefore, it cannot be used to configure the interface.	Edit the network script manually or use the supplied utility: ib-config-as-eth. In process of fixing that with Novell.

#	Subject	Description	Workaround
13.	SLES10:Ping Broadcast/Multicast blocked by default	Ping broadcast might not work because the kernel might block the stack to reply to the ping request.	To enable ping broadcast reply, run the following command from your shell: echo 0 > /proc/sys/net/ipv4/icmp_echo_ignore_broadcasts or edit /etc/sysctl.conf file and add the line: net.ipv4.icmp_echo_ignore_broadcasts = 0
14.	On PowerPC: Removing RDS module after RDS traffic causes a kernel crash		Close applications using RDS before unloading the RDS module.
15.	uDAPL - Connection fail between client and server while changing the client data size and parameters on the next C/S session	The cause is that ib1 on the servers is also up and replies to ARPs from the client. Sometimes port 2 replies to the ARP, and the client tries to connect to this port rather than to port 1 (ib0). This causes the error.	Change this by sysctl net.ipv4.conf.all.arp_ignore=1.
16.	RDS support is missing in SLES9 up3	RDS is not supported on sles9sp3.	

#	Subject	Description	Workaround
17.	Warning messages in kernel log after loading the bonding	<p>After loading the bonding module the kernel log (viewed with dmesg) shows the following:</p> <p style="color: red;">bonding: Warning: either miimon or arp_interval and arp_ip_target module parameters must be specified, otherwise bonding will not detect link failures! See bonding.txt for details .</p> <p style="color: red;">bonding: bond0: Warning: failed to get speed and duplex from ib0, assumed to be 100Mb/sec and Full.</p> <p style="color: red;">bonding: bond0: Warning: The first slave device you specified does not support setting the MAC address. This bond MAC address would be that of the active slave.</p>	These messages should be ignored.
18.	ib-config and ib-config-as-eth script does not check correctness of their input parameters	It is possible to pass illegal values to the scripts that will be accepted by them without warning or rejection.	

#	Subject	Description	Workaround
19.	Output warning during the running of some IB utilities	<p>The warning looks as follows:</p> <p><b>libibverbs: Warning: RLIMIT_MEMLOCK is 32768 bytes .This will severely limit memory registrations</b></p> <p><b>May appear for non root users when running some InfiniBand utilities that use libibverbs.</b></p>	<p>Increase the maximum size that may be locked into memory by <code>ulimit -l &lt;value&gt;</code>. This operation requires root permissions.</p>
20.	Warning messages when burning firmware	<p>When burning firmware, it is possible that the following message will appear during the operation:</p> <p>You are about to replace current PSID in the image file - "XXXXXX" with a different PSID - "YYYYYY."</p> <p>Note: It is highly recommended not to change the image PSID.</p> <p>And after that:</p> <p>You are about to replace current PSID on flash - "YYYYYY" with a different PSID - "HP_XXXXXX".</p> <p>Note: It is highly recommended not to change the PSID.</p>	<p>Ignore the recommendation and answer with 'y' to complete the operation.</p>

#	Subject	Description	Workaround
21.	Burning firmware to an HCA that was not purchased from Voltaire	<b>ib-burn</b> identifies the type of HCA by reading an identification string from the HCA flash memory. If the HCA was not purchased from Voltaire then <b>ib-burn</b> will not be able to identify the HCA type and decide which firmware to burn.	The exact type of HCA should be passed to <b>ib-burn</b> (use option -i). Alternatively, firmware image file may be given to <b>ib-burn</b> (using option -a).
22.	IPoIB Interface order might change after adding an additional HCA	It is not guaranteed that a specific InfiniBand port will have the same name (ib0, ib1, etc.) after adding an HCA. This may cause, for example, a hardware address change for an IPoIB interface.	None
23.	Bonding module for Ethernet	The bonding module that comes with GridStack replaces the original bonding driver for Ethernet that comes with the OS.  It is possible to use the new bonding driver for Ethernet interfaces but it was not tested to work with it.	None
24.	On SLES10SP1 ip multicast not being sent through ib1	If a sles10sp1 host tries to route traffic of some IP multicast subnet through interface ib1, the traffic will not be sent.	Bring ib0 down manually (ifdown ib0), or manually give it a different IP than the IP of ib1.
25.	Voltaire MPI does not support multiple HCA configuration		None
26.	Open iSCSI over TCP does not support "data digest" option	Scatterlist length for data digest is not calculated correctly.	None
27.	Performance	Performance: TFTP and iSCSI data transport speed can be improved	Fixed version was provided after the QA start
28.	Open MPI and ibutils are not supported under ppc64 with SLES10 sp1		None



#	Subject	Description	Workaround
29.	Open MPI does not support gen1 VAPI driver	Although the <i>mvapi</i> component exists in Open MPI, it is not being maintained for more than a year now.	None
30.	Open MPI CPU affinity and NUMA awareness	Open MPI does not support CPU affinity and NUMA awareness for memory allocation. We are currently working to add this to the Open MPI components.	None
31.	Open MPI threads safe	Current Open MPI 1.2 does not fully support threads safe.	None
32.	Open MPI fault tolerance	Open MPI in version 1.2 does not support "fault tolerance" or "check-point restart"	None
33.	Open MPI multi HCAs or multi ports	The support for nodes with multi HCAs (or ports) has only limited flexibility and we are currently working to add more features to the next version of Open MPI.	None
34.	Open MPI TCP and IB interconnect working together	<b>Do not</b> run Open MPI using both TCP and IB communication (in IB fabric use TCP for debug and problem isolation only).	None
35.	Open MPI progress thread	Open MPI currently does not support the use of progress thread for overlapping of computation and communication.	None
36.	Open MPI RDMA in collective operations	Open MPI currently does not use RDMA capabilities in collective operations.	None