



# Mellanox HPC-X™ Software Toolkit Release Notes

---

Rev 2.0

## NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "ASIS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies  
350 Oakmead Parkway Suite 100  
Sunnyvale, CA 94085  
U.S.A.  
www.mellanox.com  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2017. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

## Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>List Of Tables</b> .....	<b>4</b>
<b>Release Update History</b> .....	<b>5</b>
<b>Chapter 1 Overview</b> .....	<b>6</b>
1.1 HPC-X™ Requirements .....	6
1.2 Important Notes .....	6
<b>Chapter 2 Changes and New Features in This Release</b> .....	<b>7</b>
<b>Chapter 3 Known Issues</b> .....	<b>8</b>
<b>Chapter 4 Bug Fixes History</b> .....	<b>11</b>
<b>Chapter 5 Change Log History</b> .....	<b>13</b>
5.1 HPC-X Toolkit Change Log History .....	13
5.2 FCA Change Log History .....	16
5.3 MXM Change Log History .....	17
5.4 HPC-X™ Open MPI/OpenSHMEM Change Log History .....	19
5.5 HPC-X™ UPC Change Log History .....	20

## List Of Tables

Table 1:	Release Update History . . . . .	5
Table 2:	Changes and New Features . . . . .	7
Table 3:	Known Issues . . . . .	8
Table 4:	Bug Fixes History . . . . .	11
Table 5:	HPC-X Toolkit Change Log History . . . . .	13
Table 6:	FCA Change Log History . . . . .	16
Table 7:	MXM Change Log History . . . . .	17
Table 8:	HPC-X™ Open MPI/OpenSHMEM Change Log History . . . . .	19
Table 9:	HPC-X™ UPC Change Log History . . . . .	20

## Release Update History

*Table 1 - Release Update History*

Release	Date	Description
Rev 2.0	October 30, 2017	Initial version of this HPC-X version.

# 1 Overview

These are the release notes for the Mellanox HPC-X™ Rev 2.0. The Mellanox HPC-X™ Software Toolkit is a comprehensive software package that includes Open MPI, OpenSHMEM, PGAS, UPC, MXM, UCX, FCA tool suite for high performance computing environments. HPC-X provides enhancements to significantly increase the scalability and performance of message communications in the network. HPC-X™ enables you to rapidly deploy and deliver maximum application performance without the complexity and costs of licensed third-party tools and libraries.

## 1.1 HPC-X™ Requirements

The platform and requirements for HPC-X are detailed in the following table:

Platform	Drivers and HCAs
OFED / MLNX_OFED	<ul style="list-style-type: none"> <li>• OFED 1.5.3</li> <li>• MLNX_OFED 1.5.3-x.x.x, 3.3-x.x.x</li> </ul>
HCAs	<ul style="list-style-type: none"> <li>• ConnectX®-5 / ConnectX®-5 Ex</li> <li>• ConnectX®-4 / ConnectX®-4 Lx</li> <li>• ConnectX®-3 / ConnectX®-3 Pro</li> <li>• ConnectX®-2</li> <li>• Connect-IB®</li> </ul>

## 1.2 Important Notes

When HPC-X is launched in an environment without resource manager (slurm, pbs, ...) installed, or from a compute node, it will use Open MPI default rsh/ssh based launcher which does not propagate the library path to the compute nodes.

In such case, pass the `LD_LIBRARY_PATH` variable as following:

```
% mpirun -x LD_LIBRARY_PATH -np 2 $HPCX_MPI_TESTS_DIR/examples/hello_c
```

## 2 Changes and New Features in This Release

HPC-X™ Rev 2.0 provides the following changes and new features:

**Table 2 - Changes and New Features**

Category	Description
HPC-X Content	Updated the following communications libraries and acceleration packages versions: <ul style="list-style-type: none"> <li>• OpenMPI version 3.0.0</li> <li>• SHARP version 1.4</li> <li>• HCOLL version 3.9</li> <li>• UCX version 1.3</li> </ul>
UCX	<ul style="list-style-type: none"> <li>• UCX is now at GA level.</li> <li>• Added the following UCX features:                             <ul style="list-style-type: none"> <li>• <b>[ConnectX-5 only]</b> Added support for hardware Tag Matching with DC transport.</li> <li>• <b>[ConnectX-5 only]</b> Added support for Out-of-order RDMA RC and DC to support adaptive routing with true RDMA.</li> <li>• Hardware Tag Matching (See section <i>Hardware Tag Matching</i> in the User Manual)</li> <li>• SR-IOV Support (See section <i>SR-IOV Support</i> in the User Manual)</li> <li>• Adaptive Routing (AR) (See section <i>Adaptive Routing</i> in the User Manual)</li> <li>• Error Handling (See section <i>Error Handling</i> in the User Manual)</li> </ul> </li> </ul>
HCOLL	<ul style="list-style-type: none"> <li>• Added support for SHARP v1.4</li> <li>• Added support for NCCL on-host GPU based collectives.</li> <li>• Added support for Hierarchical GPU based allreduce using NCCL for scale-in and MXM/UCX for scale-out.</li> <li>• Improved shared memory performance for allreduce, barrier, and broadcast. Targeting high thread count systems, e.g. Power9.</li> <li>• Improved large message allreduce (multi-radix, zero-copy fragmentation, CPU vectorization.)</li> <li>• Added new and improved AlltoAllv algorithm - hybrid logarithmic pairwise exchange.</li> <li>• Added support for on-demand HCOLL memory. Improves HCOLL's memory footprint on high thread count system e.g. Power9.</li> <li>• Added a high performance multithreaded implementation to support MPI_THREAD_MULTIPLE applications. Designed specifically for high thread count systems, e.g. Power9.</li> <li>• HCOLL startup improvements.</li> </ul>
Open MPI / OpenSHMEM	<ul style="list-style-type: none"> <li>• Added support for Open MPI 3.0.0.</li> <li>• Added support for xpmem kernel module.</li> <li>• Added a high performance implementation of shmemptr() with UCX SPML.</li> <li>• Added a UCX allocator. The UCX allocator optimizes intra-node communication by allowing direct access to memories of processes on the same node. The UCX allocator can only be used with the UCX SPML.</li> <li>• Added a UCX one-sided component to support MPI RMA operations.</li> </ul>

### 3 Known Issues

The following is a list of general limitations and known issues of the various components of this HPC-X release.

**Table 3 - Known Issues (Sheet 1 of 3)**

Internal Ref.	Issue
-	<p><b>Description:</b> In OpenMPI 3.0.0, the MCA options <code>rmaps_dist_device</code> and <code>rmaps_base_mapping_policy</code> are not functional. A fix will be available in OpenMPI 3.1.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Process binding policy, NUMA/HCA locality</p> <p><b>Discovered in Version:</b> 2.0 (OpenMPI 3.0.0)</p>
-	<p><b>Description:</b> SHARP library is not available in HPC-X for the Community OFED and Inbox OFED.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> SHARP</p> <p><b>Discovered in Version:</b> 2.0</p>
1949	<p><b>Description:</b> Hardware Tag Matching may not function properly with UCX over DC transport. (Github issue: <a href="https://github.com/openucx/ucx/issues/1949">https://github.com/openucx/ucx/issues/1949</a>)</p> <p><b>Workaround:</b> Set the <code>mpi_add_procs_cutoff</code> parameter when running with OpenMPI, to a value higher than the total number of processes being run. * <code>"-mca mpi_add_procs_cutoff x"</code> where <code>x</code> is larger than the <code>np</code> (the total number of used processes) value.</p> <p><b>Keywords:</b> UCX, Hardware Tag Matching, DC transport</p> <p><b>Discovered in Version:</b> 2.0</p>
1162	<p><b>Description:</b> UCX currently does not support canceling send requests. (Github issue: <a href="https://github.com/openucx/ucx/issues/1162">https://github.com/openucx/ucx/issues/1162</a>)</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> UCX</p> <p><b>Discovered in Version:</b> 2.0</p>
1156724	<p><b>Description:</b> When using OpenSHMEM with MXM, a timeout on the UD transport might occur that will result in OpenSHMEM hanging during process finalization. The following error is displayed when this happens:  <code>ud_channel.c:768 Fatal: UD timeout sending to &lt;hostname&gt;</code></p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> MXM, OpenSHMEM, timeout</p> <p><b>Discovered in Version:</b> 2.0</p>



**Table 3 - Known Issues (Sheet 2 of 3)**

Internal Ref.	Issue
-	<p><b>Description:</b> MXM over Ethernet does not function for MTUs which are higher than 1024B when using firmware version 2.11.0500</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> MXM over Ethernet</p>
-	<p><b>Description:</b> While running, MXM may show excessive log message.</p> <p><b>Workaround:</b> To minimize the volume of log messages, use:  <code>-x MXM_LOG_LEVEL=fatal</code>  i.e. <code>% mpirun -x MXM_LOG_LEVEL=fatal ...</code></p> <p><b>Keywords:</b> Logs</p>
-	<p><b>Description:</b> A mixed configuration of active ports (one InfiniBand and the other Ethernet) on a single HCA is not supported.</p> <p><b>Workaround:</b> In such case, specify the port you would like to use with:  <code>"-x MXM_RDMA_PORTS"</code>  or  <code>"-x MXM_IB_PORTS"</code></p> <p><b>Keywords:</b> Port Configuration</p>
-	<p><b>Description:</b> When stack size is set to "unlimited", some application may suffer from performance degradation.</p> <p><b>Workaround:</b> Make sure that 'ulimit -s unlimited' is not set before running MXM.</p> <p><b>Keywords:</b> Performance</p>
-	<p><b>Description:</b> MXM v3.4 and v3.5 require that the <code>max_op_v1</code> value in OpenSM to be set as <code>&gt;=3</code>.</p> <p><b>Workaround:</b> Set the MXM environment parameter <code>MXM_OOB_FIRST_SL</code> to 0 from the command line:  <code>\$mpirun -x MXM_OOB_FIRST_SL=0 ...</code></p> <p><b>Keywords:</b> OpenSM Configuration</p>
-	<p><b>Description:</b> <code>MXM_IB_USE_GRH</code> must be set to "yes" when one of the following is used:  1. Socket Direct  2. Multi-Host  3. SR-IOV</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> EMXM parameters</p>
-	<p><b>Description:</b> Currently, the UPC Barrier does not utilize FCA Barrier, so while <code>GASNET_FCA_ENABLE_BARRIER</code> option that enables/disabled the FCA barrier does affect various UPC collectives, it does not affect UPC Barrier.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> UPC Barrier</p>

**Table 3 - Known Issues (Sheet 3 of 3)**

Internal Ref.	Issue
-	<b>Description:</b> UCX may not work properly with RoCE when running on a large scale.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> UCX
-	<b>Description:</b> Using UCX on ARM hosts may result in hangs due to a known issue in OMPI when running on ARM.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> UCX
-	<b>Description:</b> As UCX embedded in the HPC-X is compiled with AVX support, UCX cannot be run on hosts without AVX support. In case the AVX is not available, recompile the UCX that is available in the HPC-X with the option: <code>--with-avx=no</code>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> UCX

## 4 Bug Fixes History

Table 4 lists the bugs fixed in this release.

**Table 4 - Bug Fixes History (Sheet 1 of 2)**

Internal Ref.	Issue
-	<b>Description:</b> Fixed job data transfer from SD to libsharp.
	<b>Keywords:</b> SHARP
	<b>Discovered in Release:</b> 1.9
	<b>Fixed in Release:</b> 1.9.7
884482	<b>Description:</b> Fixed internal HCOLL datatype mapping.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
884508	<b>Description:</b> Fixed internal HCOLL datatype lower bound calculation.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
884490	<b>Description:</b> Fixed allgather unpacking issues.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
885009	<b>Description:</b> Fixed wrong answer in alltoallv.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
882193	<b>Description:</b> Fixed mcast group leak in HCOLL.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
-	<b>Description:</b> Added IN_PLACE support for alltoall, alltoallv, and allgatherv.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
-	<b>Description:</b> Fixed an issue related to multi-threaded MPI_Bcast.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406

**Table 4 - Bug Fixes History (Sheet 2 of 2)**

Internal Ref.	Issue
Salesforce: 316541	<b>Description:</b> Fixed a memory barrier issue in MPI_Barrier on Power PPC systems.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
Salesforce: 316547	<b>Description:</b> Fixed multi-threaded MPI_COMM_DUP and MPI_COMM_SPLIT hanging issues.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
894346	<b>Description:</b> Fixed Quantum Espresso hanging issues.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
898283	<b>Description:</b> Fixed an issue which caused CP2K applications to hang when HCOLL was enabled.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.7.405
	<b>Fixed in Release:</b> 1.7.406
906155	<b>Description:</b> Fixed an issue which caused VASP applications to hang in MPI_Allreduce.
	<b>Keywords:</b> HCOLL, FCA
	<b>Discovered in Release:</b> 1.6
	<b>Fixed in Release:</b> 1.7.406

## 5 Change Log History

### 5.1 HPC-X Toolkit Change Log History

*Table 5 - HPC-X Toolkit Change Log History*

Category	Description
<b>Rev 1.9.7</b>	
SHARP	Bug Fixes, see <a href="#">Section 4, “Bug Fixes History”</a> , on page 11
<b>Rev 1.9</b>	
HPC-X Content	Updated the following communications libraries and acceleration packages versions: <ul style="list-style-type: none"> <li>• OpenMPI version 2.1.2a1</li> <li>• SHARP version 1.3.1</li> <li>• HCOLL version 3.8.1652</li> <li>• MXM version 3.6.3103</li> <li>• UCX version 1.2.2947</li> </ul>
UCX	Point-to-point communication API, with tag matching, remote memory access, and atomic operations. This can be used to implement MPI, PGAS, and Big Data libraries and applications- IB transport
	A cleaner API with lower software overhead which provides better performance especially for small messages.
	Support for multitude of InfiniBand transports and Mellanox offloads to optimize data transfer performance: <ul style="list-style-type: none"> <li>• RDMA</li> <li>• DC</li> <li>• Out-of-order</li> <li>• HW tag matching offload</li> <li>• Registration cache</li> <li>• ODP</li> </ul>
	Shared memory communications for optimal intra-node data transfer: <ul style="list-style-type: none"> <li>• SysV</li> <li>• posix</li> <li>• knem</li> <li>• cma,</li> <li>• xpmem</li> </ul>
MXM	Enabled Adaptive Routing for all the transport layers (UD/RC/DC).
	Memory registration optimization.
SHARP	Improved the Out-of-the-box performance of SHARP.
Shared memory	Improved the intranode performance of allreduce and barrier.
Configuration	Changed many default parameter setting in order to achieve best out-of-the-box experience for several applications including - CP2K, miniDFT, VASP, DL-POLY, Amber, Fluent, GAMES-UK, and LS-DYNA.

**Table 5 - HPC-X Toolkit Change Log History**

Category	Description
FCA	As of HPC-X v1.9, FCA v2.5 is no longer included in the HPC-X package
	Improved AlltoAllv algorithm.
	Improved large data allreduce.
	Improved UCX BCOL.
OS architecture	Added support for ARM architecture.
<b>Rev 1.8.2</b>	
MXM	Updated MXM version to 3.6.2098 which includes memory registration optimization.
<b>Rev 1.8</b>	
Cross Channel (CC)	Added Cross Channel (CC) AlltoAllv
	Added CC zcpy Ring Bcast
SHARP	Added SHARP non-blocking collectives
Shared memory POWER	Added shared memory POWER optimizations for allreduce
	Added shared memory POWER optimizations for Barrier
Mixed data types	Added support for mixed data types
Non-contiguous Bcast	Added support for non-contiguous Bcast with UMR or SGE in CC
UMR	Added UMR support in CC bcol
Unified Communication - X Framework (UCX)	A new acceleration library, integrated into the Open MPI (as a pml layer) and available as part of HPC-X. It is an open source communication library designed to achieve the highest performance for HPC applications.
HPC-X Content	Updated the following communications libraries and acceleration packages versions: <ul style="list-style-type: none"> <li>• HCOLL updated to v3.7.</li> <li>Open MPI updated to v2.10</li> </ul>
FCA	FCA 2.x is no longer the default FCA used in HPC-X. As of HPC-X v1.8, FCA 3.x (HCOLL) is the default FCA used and it replaces FCA v2.x.
Bug Fixes	See <a href="#">Section 4, “Bug Fixes History”, on page 11</a>
<b>Rev 1.7</b>	
MXM	Updated MXM version to 3.6
FCA Collective	Added Cross-Channel based Allgather, Bcast, 8-byte Allreduce.
FCA	Added MPI datatype support.
	Added optimizations for PPC platforms.
	Added support for multiple SHARP leaders on a single host.
	Added support for collecting SHARP usage statistics.
	Exposed cross-channel non-blocking collectives to the MPI level.

**Table 5 - HPC-X Toolkit Change Log History**

Category	Description
<b>Rev 1.6</b>	
MXM v3.5	See <a href="#">Section 5.3, “MXM Change Log History”</a> , on page 17
IB-Router	Allows hosts that are located on different IB subnets to communicate with each other. This support is currently available when using the 'openib btl' in Open MPI. <b>Note:</b> When using 'openib btl', RoCE and IB router are mutually exclusive. The Open MPI inside HPC-X 1.6 is not compiled with ib-router support, therefore it supports RoCE out-of-the-box.
FCA v3.5	See <a href="#">Section 5.2, “FCA Change Log History”</a> , on page 16
<b>Rev 1.5</b>	
HPC-X Content	Updated the following communications libraries and acceleration packages versions: <ul style="list-style-type: none"> <li>• Open MPI updated to v1.10</li> <li>• UPC update to 2.22.0</li> <li>• MXM updated to v3.4.369</li> <li>• FCA updated to v3.4.799</li> </ul>
MXM v3.4.369	See <a href="#">Section 5.3, “MXM Change Log History”</a> , on page 17
FCA v3.4.799	See <a href="#">Section 5.2, “FCA Change Log History”</a> , on page 16
<b>Rev 1.4</b>	
FCA v3.3	See <a href="#">Section 5.2, “FCA Change Log History”</a> , on page 16
MXM v3.4	See <a href="#">Section 5.3, “MXM Change Log History”</a> , on page 17
<b>Rev 1.3</b>	
MLNX_OFED	Added support for OFED Inbox drivers
CPU Architecture	Added support for PPC architecture
LID Mask Control (LMC)	Added support for multiple LIDs usage when the LMC in the fabric is higher than zero. MXM will use multiple LIDs to distribute traffic across multiple links and achieve better resource utilization.
Performance	Performance improvements for all transport layers.
Adaptive Routing	Enhanced support for Adaptive Routing for the UD transport layer. For further information, please refer to the HPC-X User Manual section <i>“Adaptive Routing for UD Transport”</i> .
UD zero copy	UD zero copy support on receiver side to achieve better bandwidth utilization and reduce CPU usage.

## 5.2 FCA Change Log History

**Table 6 - FCA Change Log History**

Category	Description
<b>Rev 3.5</b>	
FCA Collective	Added MPI Allgatherv and MPI reduce
FCA	Added support for SHArP (including SHArP allreduce, reduce and barrier)
	Enhanced scalability for CORE-Direct based collectives
	Added support for complex data types
<b>Rev 3.4</b>	
General	UCX support
	Communicator caching scheme with eviction: improves jobstart and communicator creation time
Collectives	Collectives: Added Alltoallv and Alltoall small message algorithms.
<b>Rev 3.3</b>	
General	Ported to PowerPC
	Thread safety added
Collectives	Improved large message allreduce algorithm (Enabled by default)
	Beta version of network topology awareness (Enabled by default)
<b>Rev 3.0</b>	
Collectives	Offload collectives communication from MPI process onto Mellanox interconnect hardware
	Efficient collectives communication flow optimized to job and topology
MPI collectives	Significantly reduce MPI collectives runtime
MPI-3	Native support for MPI-3
Blocking and Non-blocking collectives	Support for blocking and nonblocking collectives
HCOLL	Supports hierarchical communication algorithms (HCOLL)
Collective algorithm	Supports multiple optimizations within a single collective algorithm
Performance	Increase CPU availability and efficiency for increased application performance
MPI libraries	Seamless integration with MPI libraries and job schedulers
<b>Rev 2.5</b>	
Multicast Group	Added MCG (Multicast Group) cleanup tool
Performance	Performance improvements
<b>Rev 2.2</b>	
Performance	Performance improvements



**Table 6 - FCA Change Log History**

Category	Description
Dynamic offloading rules	Enabled dynamic offloading rules configuration based on the data type and reduce operations
Mixed MTU	Added support for mixed MTU
<b>Rev 2.1.1</b>	
AMD/Interlagos CPUs	Added support for AMD/Interlagos CPUs
<b>Rev 2.1</b>	
Core-Direct®	Added support for Mellanox Core-Direct® technology (enables offloading collective operations to the HCA.)
Non-contiguous data layouts	Added support for non-contiguous data layouts
PGI compilers	Added support for PGI compilers

### 5.3 MXM Change Log History

**Table 7 - MXM Change Log History**

Category	Description
<b>Rev 3.6</b>	
General	Updated MXM version to 3.6
<b>Rev 3.5</b>	
Performance	Performance improvements
<b>Rev 3.4.369</b>	
Initialization	Job startup performance optimization
Supported Transports	DC enhancements and startup optimizations
<b>Rev 3.4</b>	
Supported Transports	Optimizations for the DC transport at scale
<b>Rev 3.3</b>	
LID Mask Control (LMC)	Added support for multiple LIDs usage when the LMC in the fabric is higher than zero. MXM will use multiple LIDs to distribute traffic across multiple links and achieve better resource utilization.
Adaptive Routing	Enhanced support for Adaptive Routing for the UD transport layer.
UD zero copy	UD zero copy support on receiver side to achieve better bandwidth utilization and reduce CPU usage.
<b>Rev 3.2</b>	

**Table 7 - MXM Change Log History**

Category	Description
Atomic Operations	Added hardware atomic operations support in the RC and DC transport layers for 32bit and 64bit operands. This feature is set to ON by default. To disable it run: <code>oshrun -x MXM_CIB_USE_HW_ATOMICS=n ...</code> <b>Note:</b> If hardware atomic operations are disabled, the software atomic is used instead.
MXM API	Added two additional functions ( <code>mxm_ep_wireup()</code> and <code>mxm_ep_power-down()</code> ) to the MXM API to allow pre-connection establishment for MXM (rather than on-demand). For further information, please refer to the HPC-X User Manual section “MXM Performance Tuning”.
Event Interrupt	Added solicited event interrupt for the rendezvous protocol for potential performance improvement. For further information, please refer to the HPC-X User Manual section “MXM Performance Tuning”.
Performance	Performance improvements for the DC transport layer.
Adaptive Routing	Added Adaptive Routing for the UD transport layer. For further information, please refer to the HPC-X User Manual section “Adaptive Routing for UD Transport”.
<b>Rev 3.0</b>	
Service Level	Service Level support (at Alpha level)
Adaptive Routing	Adaptive Routing support in UD transport layers
Supported Transports	Dynamically Connected Transport (DC) (at GA level)
Performance	Performance optimizations
<b>Rev 2.1</b>	
Supported Transports	Dynamically Connected Transport (DC) (at Beta level)
	RC is currently fully supported
	KNEM support for Intra-node communication
Performance	Performance optimizations
<b>Rev 2.0</b>	
Reliable Connected	Added Reliable Connection (RC) support (at beta level)
MXM Binding	MXM process can be pinned to a specific HCA port. MXM supports the following binding policies: <ul style="list-style-type: none"> <li>static - user can specify process-to-port map</li> <li>cpu affinity based - HCA port will be bound automatically based on process affinity</li> </ul>
On-demand connection establishment	Added on-demand connection establishment between the processes
Performance	Performance tuning improvements

**Table 7 - MXM Change Log History**

Category	Description
<b>Rev 1.5</b>	
MXM over Ethernet	Added Ethernet support
Multi-Rail	Added Multi-Rail support

## 5.4 HPC-X™ Open MPI/OpenSHMEM Change Log History

**Table 8 - HPC-X™ Open MPI/OpenSHMEM Change Log History**

Category	Description
<b>Rev 1.8.2</b>	
Acceleration Packages	Added support for new MXM, FCA, HCOLL versions
Job start optimization	Added job start optimization
Performance	Performance improvements
<b>Rev 2.2</b>	
Performance	Added Sandy Bridge performance optimizations.
memheap	Allocated memheap using contiguous memory provided by the HCA.
ptmalloc allocator	Replaced the buddy memheap by the ptmalloc allocator.
multiple pSync arrays	Added the option of using multiple pSync arrays instead of barrier synchronization between collective routines (fcollect, reduction routines)
spml yoda	Optimized small size puts
Performance	Performance optimization
Memory footprint optimizations	Added memory footprint optimizations

## 5.5 HPC-X™ UPC Change Log History

*Table 9 - HPC-X™ UPC Change Log History*

Category	Description
<b>Rev 2.18.0</b>	
Acceleration Packages	Added support for new MXM, FCA, HCOLL versions
PMI2 support	Added job start PMI2 support
<b>Rev 2.2</b>	
FCA library	Linking with FCA library instead of using dlopen at runtime.
MPI	Fixed an issue using some of MPIs as job spawner (e.g. MPICH2) Use MPI_BYTE rather than MPI_CHAR, and use MPI_IN_PLACE.