



Fabric Collective Accelerator (FCA)

Release Notes

Version 2.2

www.mellanox.com

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” 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, Inc.
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies, Ltd.
Beit Mellanox
PO Box 586 Yokneam 20692
Israel
www.mellanox.com
Tel: +972 (0)4 909 7200 ; +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

© Copyright 2012. Mellanox Technologies, Inc. All Rights Reserved.

Mellanox Technologies. All rights reserved. Mellanox®, Mellanox logo®, BridgeX®, ConnectX®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, PhyX®, SwitchX®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

FabricIT™, MLNX-OS™, Unbreakable-Link™, UFM™ and Unified Fabric Manager™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Contents

1	Introduction	4
2	Overview of FCA	4
3	Key Features	4
4	System Requirements	5
5	Interoperability/Compatibility	6
6	Known Limitations	6
7	Product Documentation	6
	7.1 Related Documentation	6
8	Submitting a Service Request	7

1 Introduction

These release notes pertain to the Fabric Collective Accelerator™ (FCA™), software version 2.2.

2 Overview of FCA

Collective (or group) communications, which have a crucial impact on the application's scalability, are frequently used by scientific simulation codes like broadcasts for sending around initial input data, reductions for consolidating data from multiple sources and barriers for global synchronization. Any collective communication executes some global communication operation by coupling all processes in a given group. This behavior tends to have the most significant negative impact on the application's scalability. Additionally, the explicit and implicit communication coupling, used in high-performance implementations of collective algorithms, tends to magnify the effects of system-noise on application performance further hampering application scalability. Mellanox adapters and switches address the collective communication scalability problem by offloading the collective communications to the network. This solution provides the mechanism needed to support computation and communication overlap, allowing the communication to progress asynchronously in hardware while at the same time computations are processed by the CPU. It also provides a means to reduce the effects of system noise and application skew on application scalability. Mellanox Fabric Collective Accelerator (FCA) software utilizes the interconnect hardware engines in the Mellanox adapters (CORE-Direct technology) and selected switches to accelerate the collectives runtime, to increase the CPU availability to the application and to allow overlap of communications and computations with asynchronous collective operations.

3 Key Features

Key FCA features:

- Offloading collective function communication and computation from MPI process into Mellanox adapters and switches
- Efficient collective communication flow optimized to job and topology

New features in FCA version 2.2:

- Performance improvements
- Enabled dynamic offloading rules configuration based on the datatype and reduce operations
- Support for mixed MTU

New features in FCA version 2.1.1:

- Support for AMD/Interlagos CPUs

New features in FCA version 2.1:

- Support for Mellanox *Core-Direct*® technology (offloading collective operations onto HCA)
- Support for non-contiguous data layouts
- Support for PGI compilers

4 System Requirements

Before you begin be certain that:

1. InfiniBand Subnet Management is installed and running on a dedicated node in the fabric.
2. Mellanox OFED 1.5.3 or later is installed.

To download the latest MLNX_OFED version, go to:

[Mellanox OpenFabrics Enterprise Distribution for Linux \(MLNX_OFED\)](#)

3. Mellanox ConnectX®-2 or ConnectX®-3 HCA with firmware version 2.9.1000 or later.

To download the latest ConnectX® HCA firmware version, go to:

[Firmware Downloads](#)

The minimum system requirements for installing and running FCA are listed in the following table.



NOTE: Mellanox OFED 1.8 includes FCA 2.2, which is installed under `/opt/mellanox/fca`. If you have installed OFED 1.8, you do *not* need to download and install FCA.

Table 1: System Requirements

Item	Requirement
FCA 2.2	
Supported switches	Mellanox IB QDR/FDR switches
Linux distributions (<OS>)	RHEL 5.5, RHEL 5.6, RHEL 6.0, RHEL 6.1, RHEL 6.2, CentOS 5.5, CentOS 5.6, Centos 5.7, CentOS 6.0, SLES 10 SP4, SLES 11, SLES11 SP1
Supported HCAs	Mellanox ConnectX®-2 HCA with firmware version 2.9.1000 or later Mellanox ConnectX®-3 HCA with firmware version 2.10.0000 or later
Open Message Passing Interface (MPI) Project	Open MPI 1.4.3 or later.*
Open Fabrics Enterprise Distribution (OFED™)	1.5.3 or later
Root permission	The installer should have root permissions for post-installation tasks.
InfiniBand Subnet Management	All InfiniBand Subnet Management based software is supported in FCA version 2.2.



NOTE: MLNX_OFED v1.8 comes with OpenMPI 1.4.6 and 1.6 which is already compiled to support FCA v2.2.

5 Interoperability/Compatibility

FCA version 2.2 is fully interoperable and compatible with the following components:

- Version 3.6 or later of the following Mellanox's switches:
 - Mellanox Grid Director 4036/4036E/4200/4700
 - Mellanox IB QDR switches (IS5000 switch family)
 - Mellanox FDR switches (SX6000 switch family)

6 Known Limitations

This section lists the known limitations for FCA version 2.2.

Table 2: Product-Related Limitations

Description	Workaround
Support for all C data types, except for MPI_LONG_DOUBLE, for Reduce/AllReduce.	N/A
Congestion Control (ccm) must be disabled on all of the switches specified below when using FCA. It is disabled by default on switches shipped from the factory. <ul style="list-style-type: none"> • Mellanox Grid Director 4036/4036E/2004/2012/4200/4700/9024 	N/A
FCA only works on the default p-key.	N/A
When the fabric SM is running on a switch (embedded SM), the FCA module should be disabled in that specific switch.	N/A
FCA 2.1.1 cannot co-exist in the same fabric with previous versions of FCA.	Run a single version of FCA in the same fabric.

7 Product Documentation

This section lists the product documentation for FCA version 2.2.

Document Title	Description	Part Number
Fabric Collective Accelerator (FCA) User Manual	User Manual	DOC-00983

7.1 Related Documentation

For additional relevant information, please refer to the following documentation:

- Mellanox® OpenFabrics Enterprise Distribution (OFED) Software Version 1.5.3 Release Notes

- Mellanox® OFED Software Version 1.5.3 User Manual

8 Submitting a Service Request

The Mellanox Support Center is at your service. You may access Warranty Service through our Web Request Form by using the following link:

http://www.mellanox.com/content/pages.php?pg=support_index.