



Connect. Accelerate. Outperform.™

Mellanox Unstructured Data Acceleration (UDA) Release Notes

Rev 3.4.1

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 NONINFRINGEMENT 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 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, CloudX logo, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

Accelio™, CyPU™, FPGADirect™, HPC-X™, InfiniBridge™, LinkX™, Mellanox Care™, Mellanox CloudX™, Mellanox Multi-Host™, Mellanox NEO™, Mellanox PeerDirect™, Mellanox Socket Direct™, Mellanox Spectrum™, NVMeDirect™, StPU™, Spectrum logo, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
List Of Tables	4
Chapter 1 Overview	5
1.1 Content of Mellanox Unstructured Data Accelerator (UDA)	5
1.2 Hardware and Software Dependencies	5
Chapter 2 Changes and New Features in This Release	6
Chapter 3 Bug Fixes History	7
Chapter 4 Change Log History	8

List Of Tables

Table 1:	New Features, Changes and Fixes	6
Table 2:	Fixed Bugs List	7
Table 3:	Change Log History	8

1 Overview

These are the Release Notes for Mellanox Unstructured Data Accelerator (UDA) Rev 3.4.1.

The UDA plug-in software package provides a novel shuffle approach for Hadoop's MapReduce framework. RDMA based networks, with its low latency and high bandwidth features, build the most efficient shuffle provider for MapReduce. Compared to a 1GbE network, benchmark results show nearly double performance of Hadoop® clusters using UDA with 10GbE networks and quadruple the performance using FDR InfiniBand. UDA is an Open Source software package, available under Apache 2.0 License.

1.1 Content of Mellanox Unstructured Data Accelerator (UDA)

Mellanox UDA software contains the following:

- README
- journal.txt
- libuda.so - RDMA handling
- uda-hadoop-1.x.jar - Java side plugin for Hadoop-1
- uda-hadoop-2.x.jar - Java side plugin for Hadoop-2 and above

1.2 Hardware and Software Dependencies

- Linux OSs
 - RHEL 6.0 and above
 - Ubuntu 12.04
- MLNX_OFED-1.5.3 or MLNX_OFED-2.x
- Apache Hadoop versions:
 - Community Hadoop: 1
 - Community Hadoop: 2.2 and above
 - Hadoop Data Platform (HDP) 1.1
 - Greenplum Hadoop (GPHD) 1.2
 - Cloudera's Distribution including Apache Hadoop (CDH) 4.1.2
 - CDH 4.2.1
 - CDH 4.3.1
 - CDH 4.4.0

Other Hadoop versions and distributions may also be supported. Please contact Mellanox support for your specific Hadoop version.

- Hadoop configurations:
 - Disabled speculative reduce tasks
 - Disable systems swap ("swapoff -a" or unmount from /etc/fstab)
 - Use heuristic (default) overcommit VM scheme ("echo 0 > /proc/sys/vm/overcommit_memory")

2 Changes and New Features in This Release

UDA Rev 3.4.1 provides the following new features:

Table 1 - New Features, Changes and Fixes

Category	Description
Optimized Memory Access	Contiguous pages are used by default
HBase libs	Added support for ImmutableBytesWritable compare function

3 Bug Fixes History

Table 2 lists the bugs fixed in this release.

Table 2 - Fixed Bugs List

Issue	Description	Discovered in Release	Fixed in Release
UDA reducers	<ul style="list-style-type: none"> Fixed the hanging of reducers when there are a lot of MOFs with compression. 	3.2.2	3.4.1
libsnapy.so missing from search path	<ul style="list-style-type: none"> Aligned the search path of libsnapy.so in UDA with the path of libsnapy.so in Hadoop 	3.1.11	3.2.2
RDMA verbs operations	<ul style="list-style-type: none"> Fixed a fork() support issue for RDMA verbs operations 	3.0	3.1.11
Variable Length Text sort/compare function	<ul style="list-style-type: none"> Fixed Variable Length Text sort/compare function used by TestDFSIO and Word Count 	3.0	3.1.11
UDA reducers	<ul style="list-style-type: none"> Fixed leftover UDA reducers 	3.0	3.1.11
RDMA connection	<ul style="list-style-type: none"> Fixed several RDMA connection related errors 	3.0	3.1.11

4 Change Log History

Table 3 - Change Log History

Release	Description
Rev. 3.4.0	Added support for MLNX_OFED 2.x
	Added a new configuration parameter (mapred.rdma.mem.use.contig.pages), to allow contiguous pages allocation (supported as of MLNX_OFED 2.1)
	Added a new configuration parameter (mapred.rdma.developer.mode), to prevent UDA from falling back to default built-in shuffle in case of failure. Note that this should not be used in production environment.
	Added compare function support for BytesWritable keys, used by Sort example program.
Rev 3.3.3	Added support for Hadoop-2.2.x (Hadoop GA version)
	Unified UDA v2, UDA v3 and UDA CDH plugins into a single uda-hadoop-1.x.jar which supports all of them
Rev 3.2.2	Added support for Hadoop 2.x.y
	Added support for Ubuntu 12.04 including .deb packaging
	Added a new configuration parameter, mapred.rdma.shuffle.total.size to control UDA memory
Rev 3.1.11	Added snappy compression support for intermediate MOF file
	Added resilience fallback to default built-in shuffle in cases of: <ul style="list-style-type: none"> • Unsupported user defined compare function • Unsupported configurations like some compression modes • Not enough memory to complete job with RDMA shuffle • General UDA failure
	Enabled UDA log integration into Task Tracker and Reduce Task log system
	Added additional Apache Hadoop and 3rd party Hadoop distributions support