



Hadoop Benchmark with Fujitsu Servers

~Mellanox UDA Accelerates the Hadoop Performance~

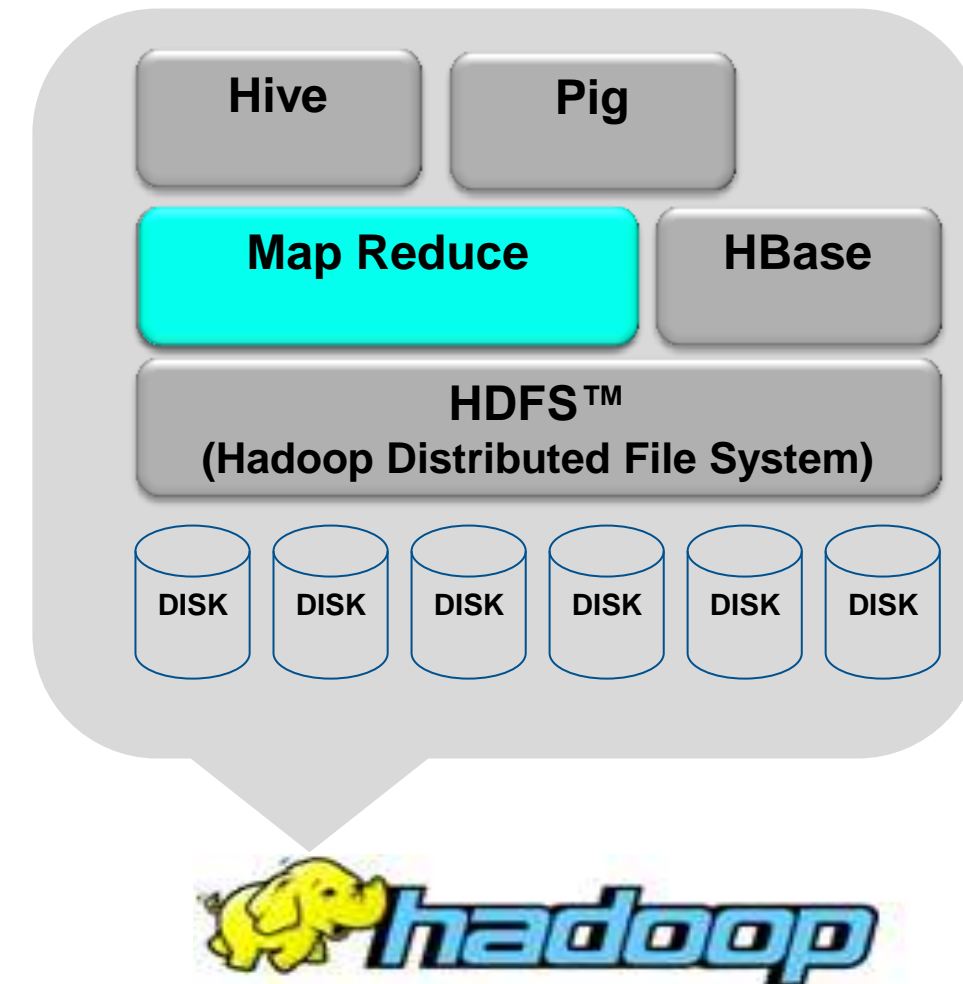
August 2013

 **Mellanox**
TECHNOLOGIES
Connect. Accelerate. Outperform.™

Unstructured Data Accelerator - UDA



- Plug-in architecture
 - Open-source, latest GA version 3.1 (6/10/2013)
 - Google code repository at: <https://code.google.com/p/uda-plugin/>
- Accelerates Map Reduce Jobs
 - Accelerated merge sort
- Efficient Shuffle Provider
 - Data transfer over RDMA
 - Supports InfiniBand and Ethernet
- Supported Hadoop Distributions
 - **Apache 3.0 – In the main trunk!**
 - **Apache 2.0.3 – In the main trunk*!**
 - Apache Hadoop 1.0.x ; 1.1.x
 - Cloudera Distribution Hadoop 3 update 4 (CDH3u4)
 - Cloudera Distribution Hadoop 4 (CDH4)
 - Hortonworks HDP 1.1
- Supported Hardware
 - ConnectX®-3 VPI
 - SwitchX-2 based systems



■ Hardware configuration

	Node type	Server	CPU	MEM	Disk	PCIe
Server	Name node	1x Fujitsu RX300 S7 Host name : UDA1	2x Intel Xeon E5-2690(8cores) 2.90GHz	64GB	3x 200GB	Gen3
	Data node	2x Fujitsu RX200 S7 Host name : UDA4,UDA5	2x Intel Xeon E5-2670(8cores) 2.60GHz	32GB	3x 600GB	Gen3
		2x Fujitsu RX300 S7 Host name : UDA2,UDA3	2x Intel Xeon E5-2660(8cores) 2.20GHz	32GB	3x 400GB	Gen3
InfiniBand Switch	Mellanox SX6036 (MLNX OS : v3.3.3500)					
InfiniBand HCA	Mellanox ConnectX-3 VPI, FW Rev 2.11.0500					

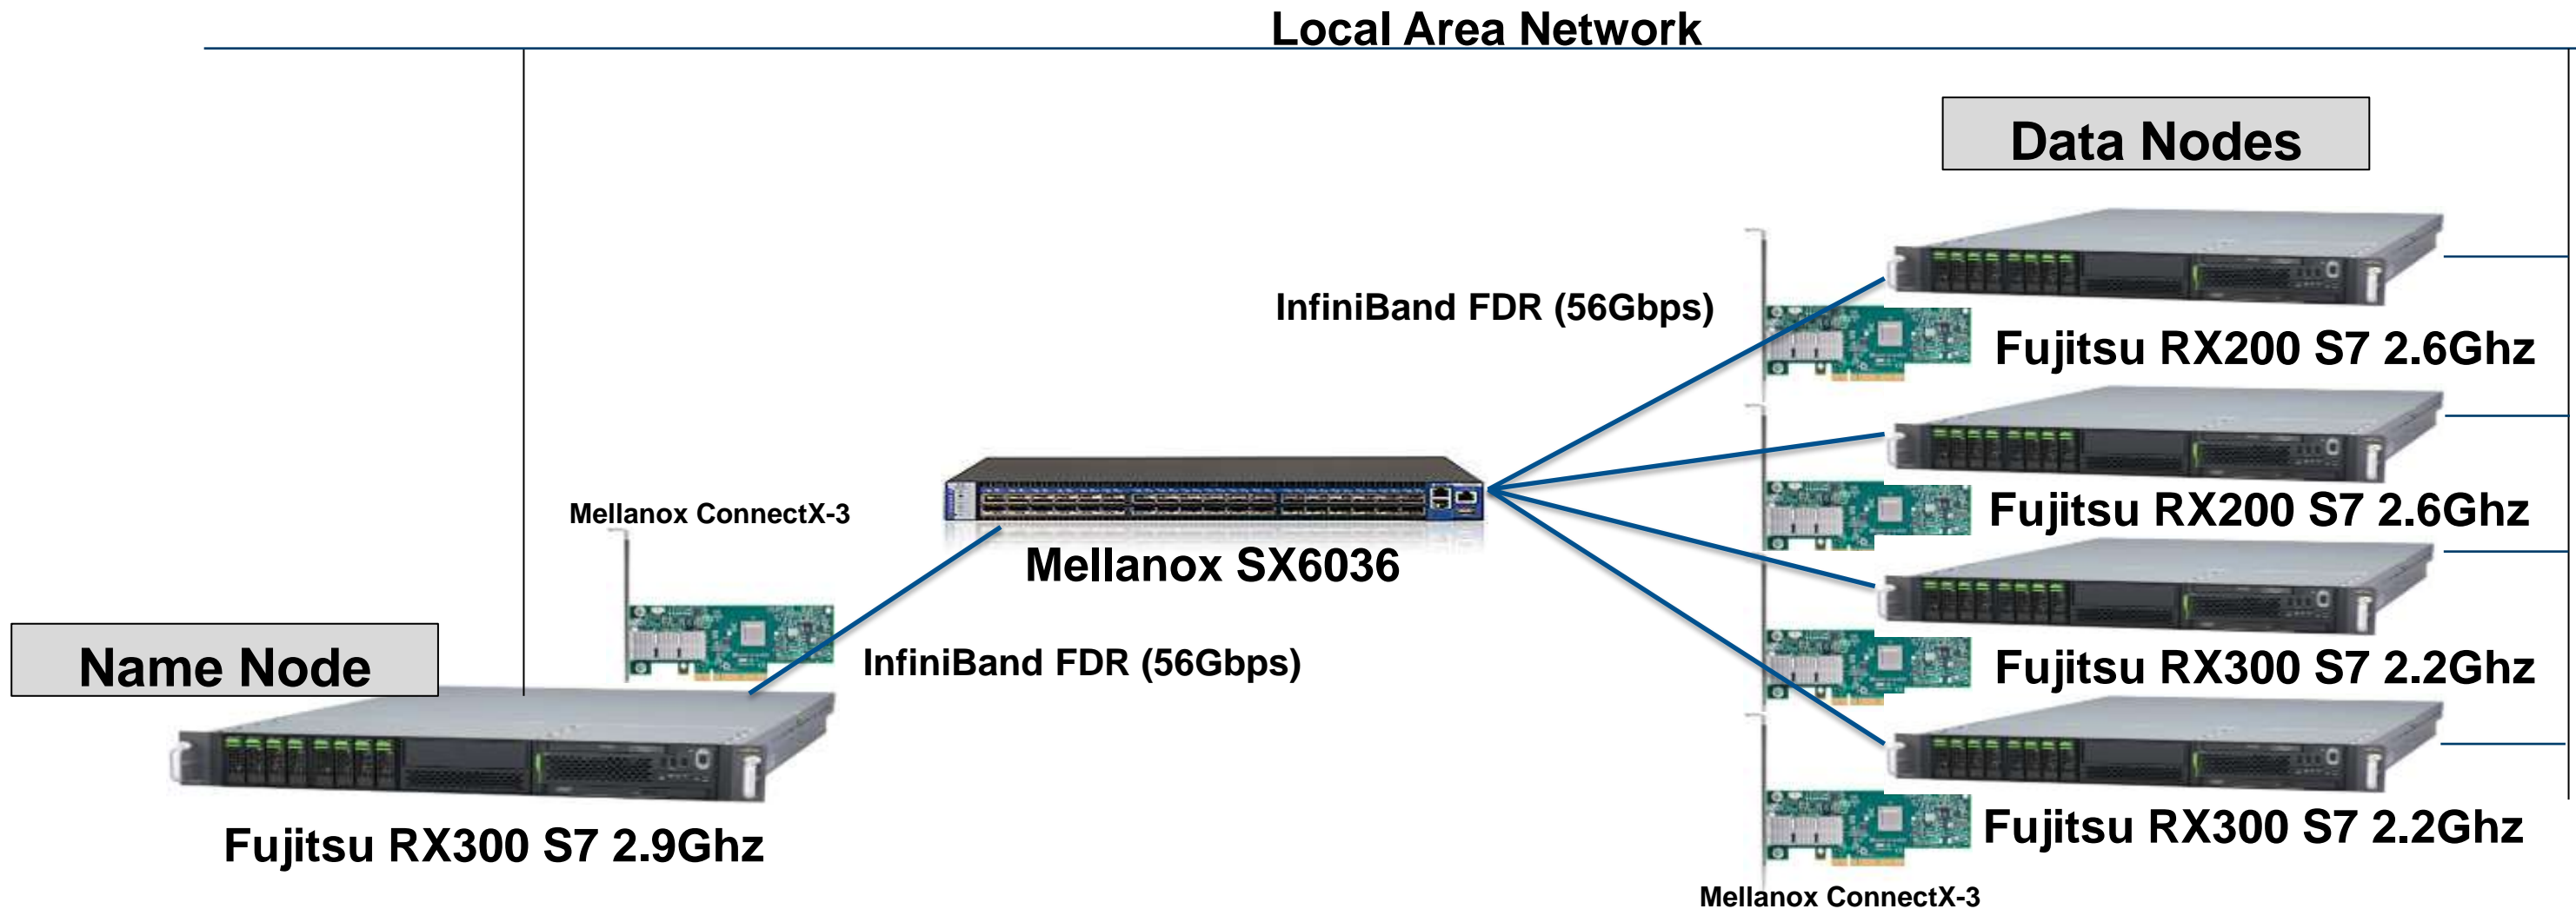
■ Software configuration

- OS : RHEL6.3 Kernel : 2.6.32-279.el6
- Driver : MLNX OFED 2.0-2.0.5
- UDA : ver3.1.11
- Hadoop : Apache Hadoop ver 1.1.2

*This test was conducted by Altima/Mellanox. (July-1~July-12, 2013)


■ Hadoop Configuration

- 1 Name Node
- 4 Data Nodes



*This test was conducted by Altima/Mellanox. (July-1~July-12, 2013)

Result – Terasort Benchmark : IB FDR(56Gbps) IPoIB vs UDA



Mellanox UDA accelerated Hadoop performance by 2.5X faster

*This test was conducted by Altima/Mellanox. (July-1~July-12, 2013)

Thank You

