




Unstructured Data Accelerator (UDA)

Author: Motti Beck, Mellanox Technologies
Date: March 27, 2012

Market Trends Big Data

- Growing technology deployments are creating an exponential increase in the volume of data available
- Existing analytical techniques not adequate for business decision-making processes
- A successful approach of big data analytics will be a critical core competency
 - Delivering significant competitive advantage to organizations

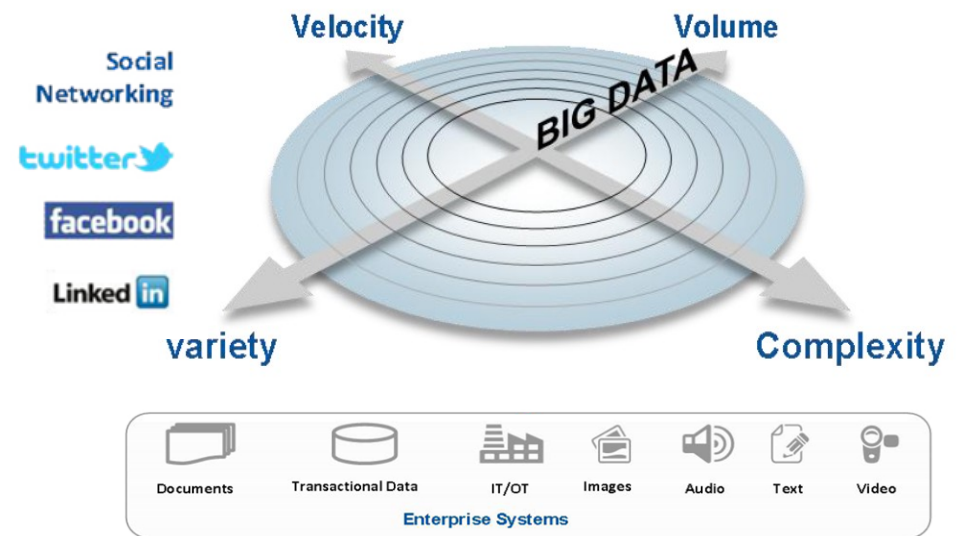


Nationwide Electricity Grid Analysis

Big Data – Not just a Volume Play

- Facebook statistics 2009

- 350 Million Named users
- 175 Million Active users in one day
- 35 Million Users updating status each day
- 2.5 Billion Photos / Month
- 1.6 Million Active pages
- Growth: 12 TB /day, 2 PB /year
- Global data volume : 8.7 PB




Source: Gartner

Market Trends – Serial to Parallel Computing

Big Data Analytics is Parallel Processing

Single Core



Multi cores

Single Computer




Clustered
Computing

Traditional Data-Processing Pipeline



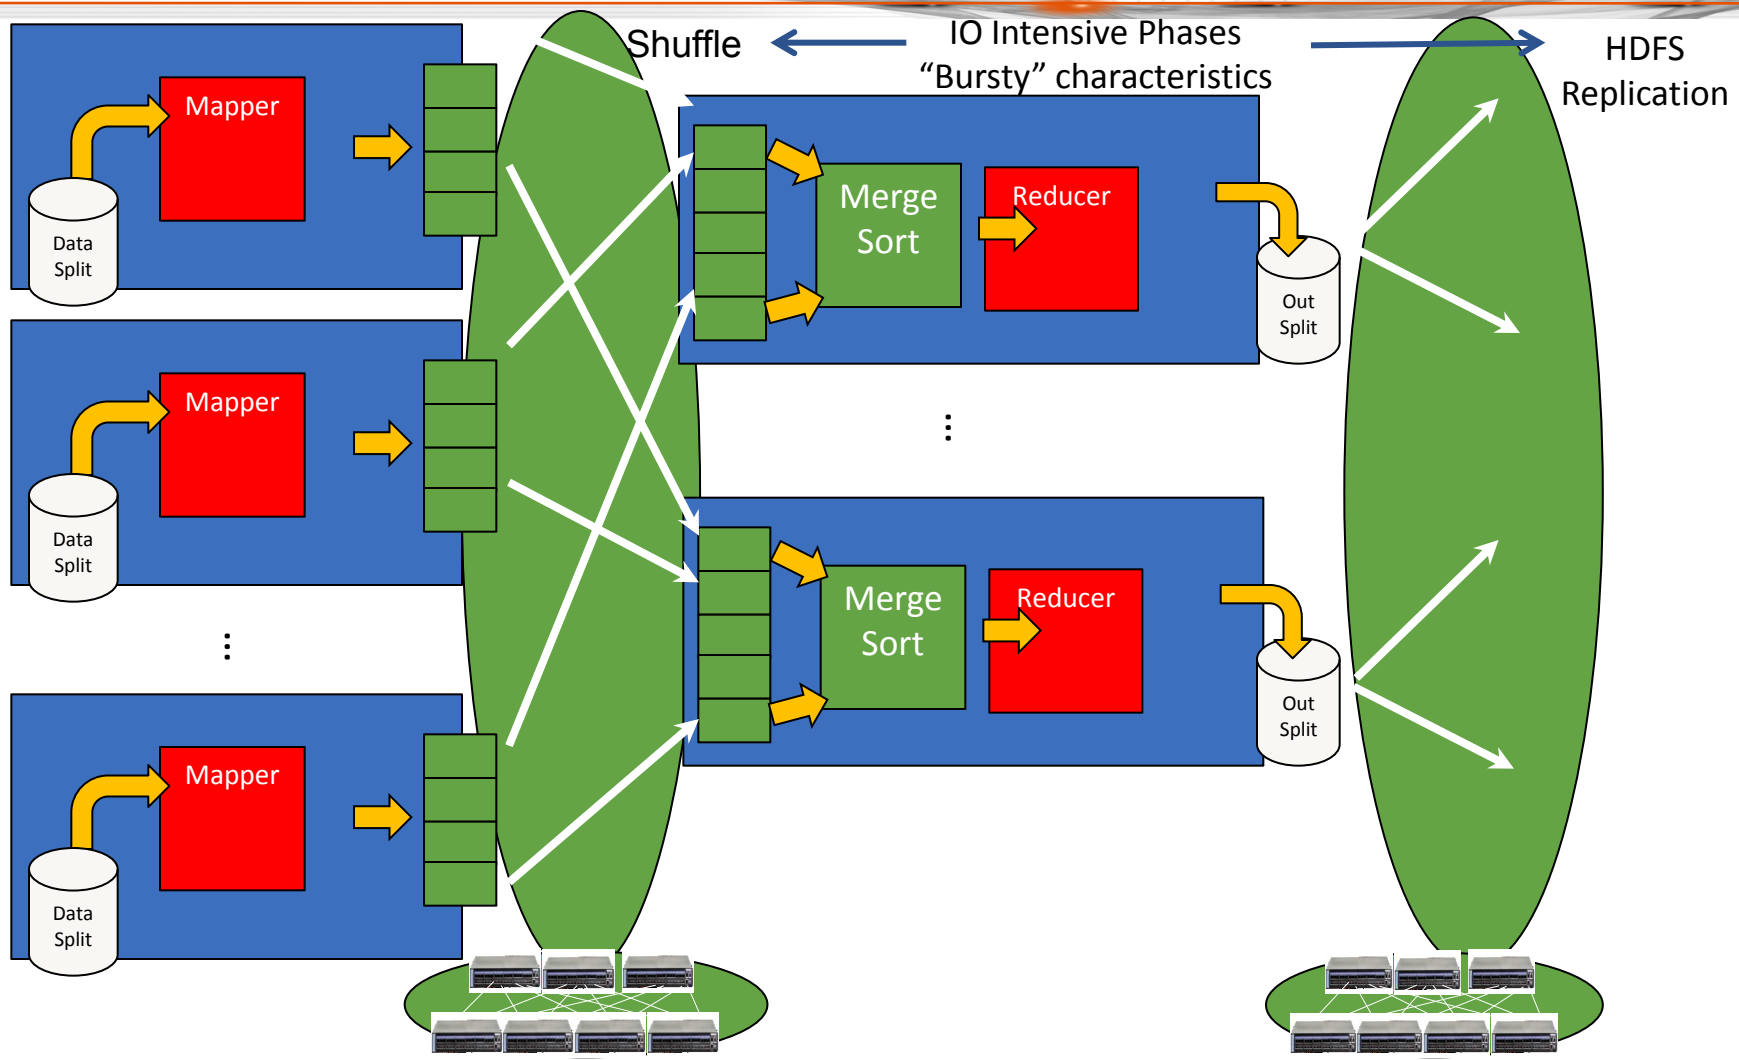
© 2011 Gartner, Inc. and/or its affiliates. All rights reserved.

Big Data parallel processing over Hadoop





© 2011 Gartner, Inc. and/or its affiliates. All rights reserved.

Hadoop Cluster over High Performance Networking



I/O Bottlenecks: from Disk to Network


- Trends
 - Use multiple SATA drives
 - SSD is around the corner
- Matching network speed
 - 10Gigabit Ethernet and beyond



Source: "SSD Summary", Gartner, May 2011

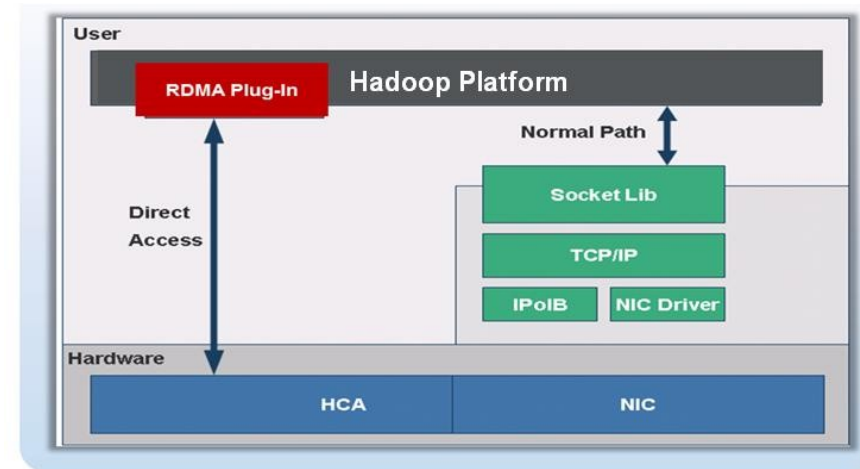
Efficient Fabric Services

- Scalable and non-blocking
 - Matches extensive data exchange at peak rates
 - Non blocking east-west traffic
- High Capacity
 - Match all I/O bandwidth of the server
- Losslessness
 - Efficiency and avoidance of retransmission
- Offload
 - Full transport offload – reduces CPU utilization
 - RDMA zero copy operations



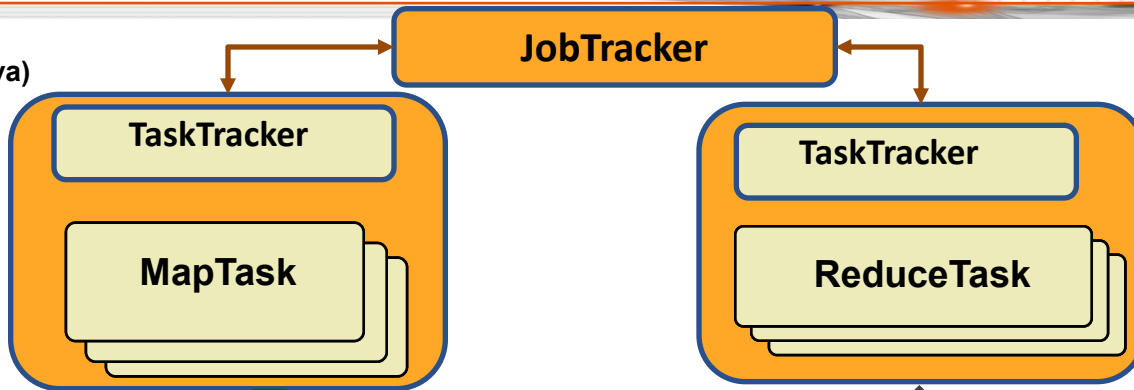
UDA Plug In Architecture

- Plug-in architecture
 - Hadoop applications are unmodified
 - Plug-in to Apache Hadoop
 - Enabled via xml configuration
- Efficient Map Reduce
 - Data communication over RDMA
 - Using RDMA for In-Memory processing
 - Enables to start the Reduce operation in parallel to the Shuffle operation
 - Reduce disk IO operation
 - Supports InfiniBand and Ethernet
 - Zero copy, transport offload, kernel bypass




Software Architecture

Hadoop (Java)



UDA Plugin

(C++)

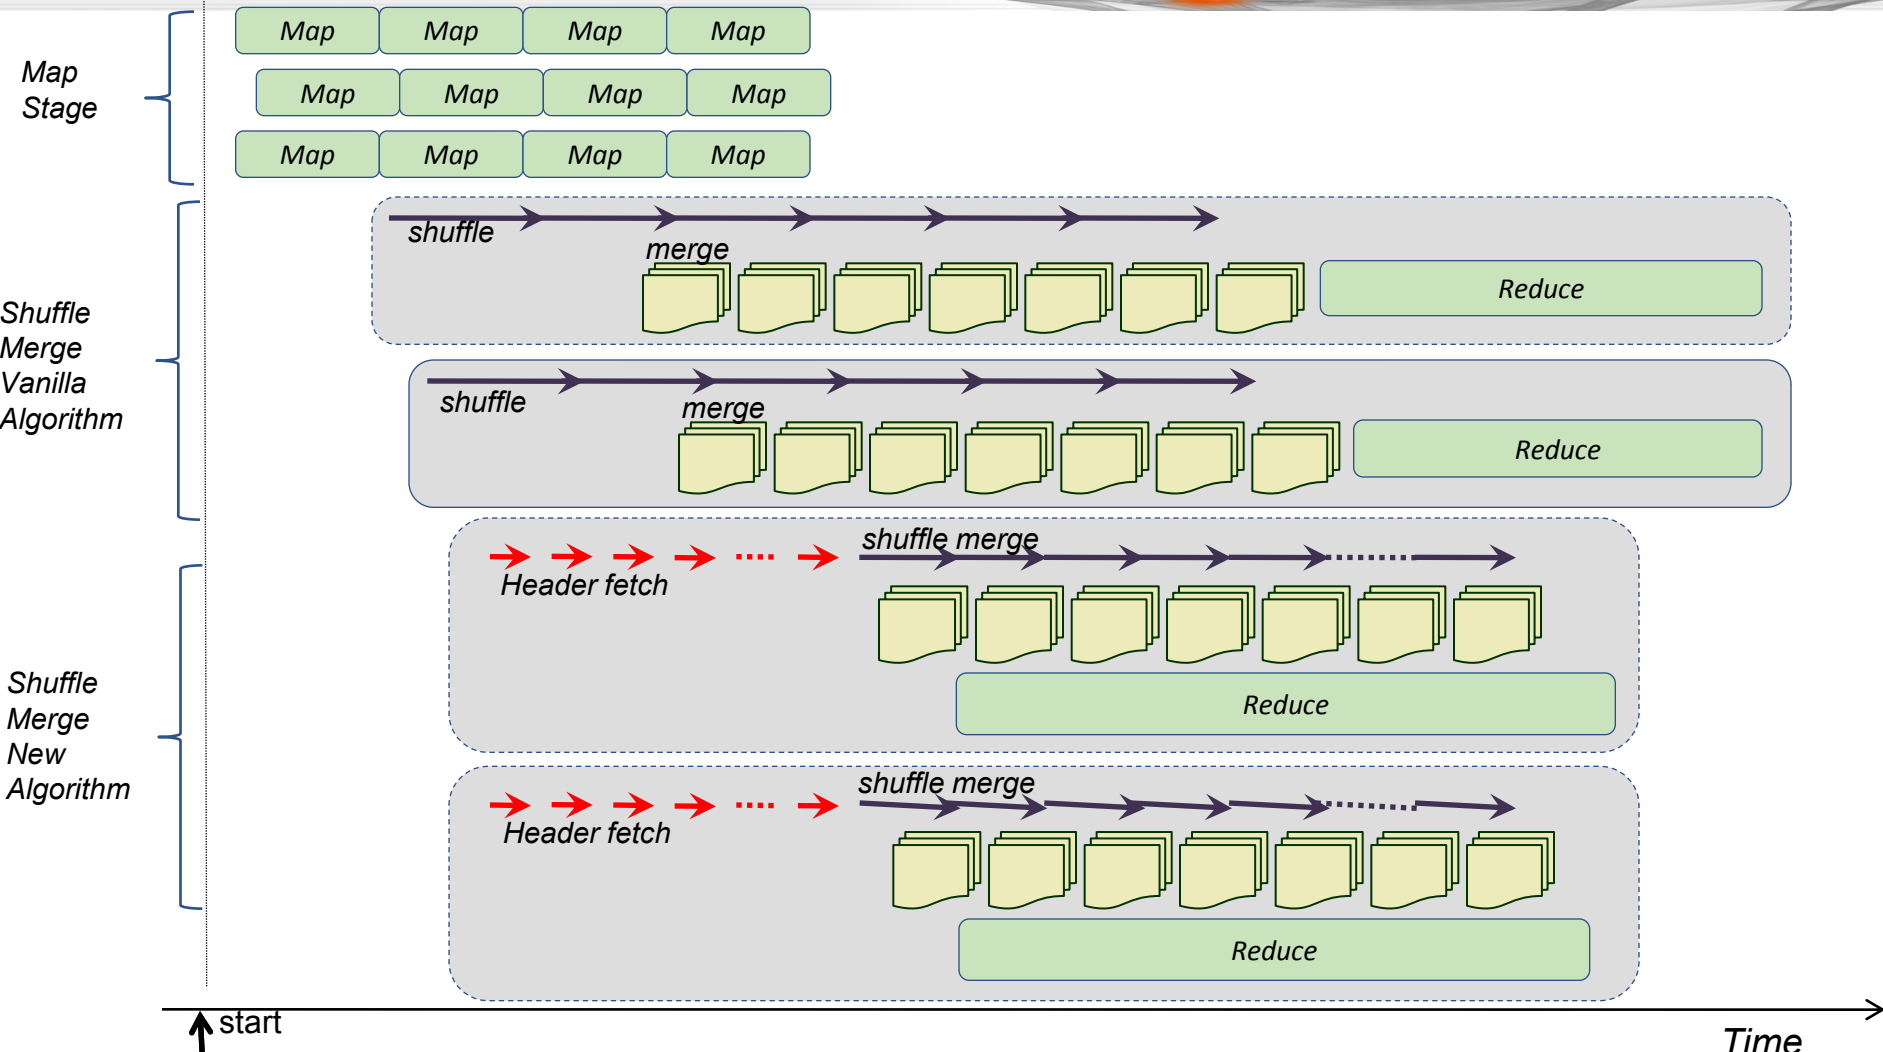


RDMA NIC / HCA

Plug-in Benefits:

- Zcopy datapath
- Transport offload
- Improved merge algorithm

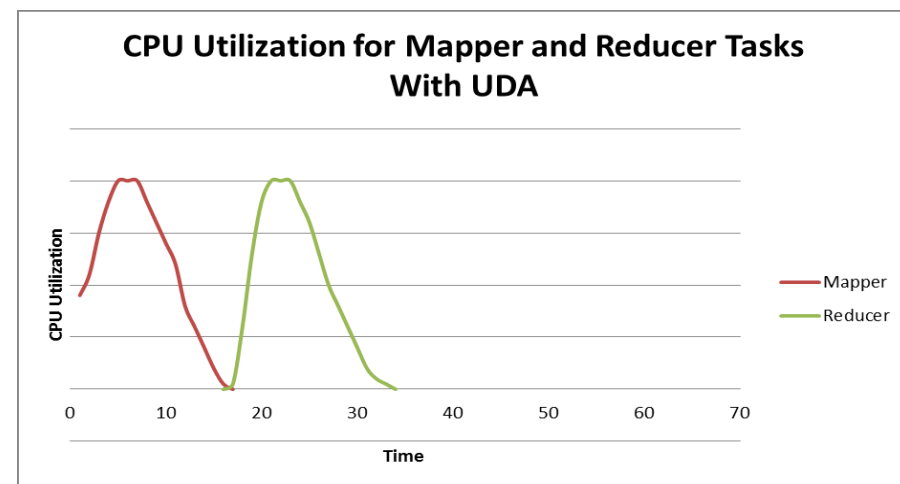
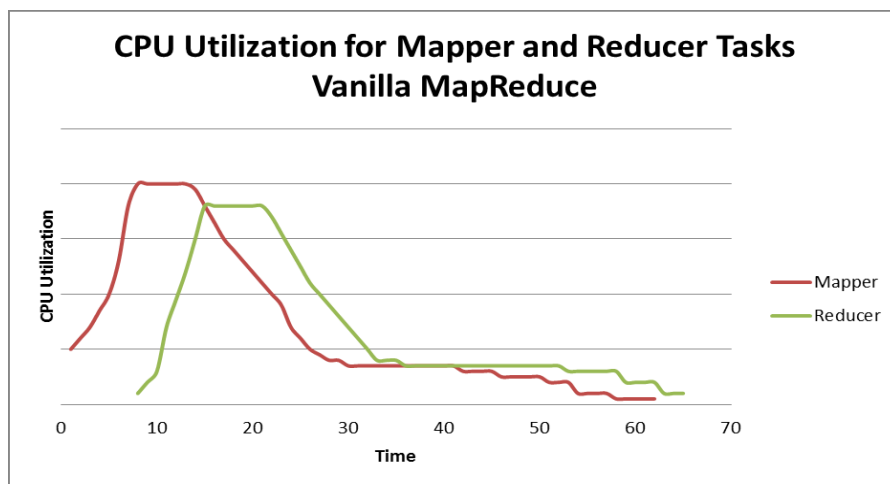
New Pipelined Data Flow



UDA MapReduce for Hadoop™ 1.X




- Shuffle portion executed in-memory
 - Eliminating time consuming HDD read/writes
- UDA reads Map Output Files (MOF) from mappers
 - Predefined portions of the MOFs, default 1KB
 - Reduce tasks begins only with all mapper tasks completion



UDA Enables Higher Performance

Terasort Benchmark*

(20GB file size, 16GB data per node, 8 Mappers, 4 reducers, 4 Disks)



*TeraSort is a popular benchmark used to measure the performance of Hadoop cluster

Accelerating Big Data Analytics



- EMC 1000-Node Analytic Platform
- Accelerates Industry's Hadoop Development
- 24 PetaByte of physical storage
 - Half of every written word since inception of mankind
- Mellanox VPI Solutions



Hadoop
Acceleration

2X Faster Hadoop Job Run-Time

High Throughput, Low Latency, RDMA Critical for ROI



Thank You



Email: motti@mellanox.com