



# High Speed Workstation Data Center Management Leverage Low Latency and Linux Systems

Subhendu Ghosh  
April 4th, 2011

**RED HAT®  
ENTERPRISE LINUX®**



- Uniform deployment and management
- Common high performance platform



# Latency and Red Hat Enterprise Linux

- Networking
  - RDMA, 10 GigE
- CPU/Memory ↔ Disk
  - Storage Architectures
- CPU ↔ Memory
  - NUMA aware, CPU Affinity
- Scheduler
  - Realtime





# Networking – New

- Multi-queue
- Tools to monitor dropped packets – tc, dropwatch.
- Read-Copy-Update (RCU) adoption in stack
- Multi-CPU receive to pull in from the wire faster.
- 10GbE driver improvements.
- Data center bridging in ixgbe driver.
- FcoE performance improvements throughout the stack.



# RHEL5.5 to RHEL6 AMQP TCP Latency



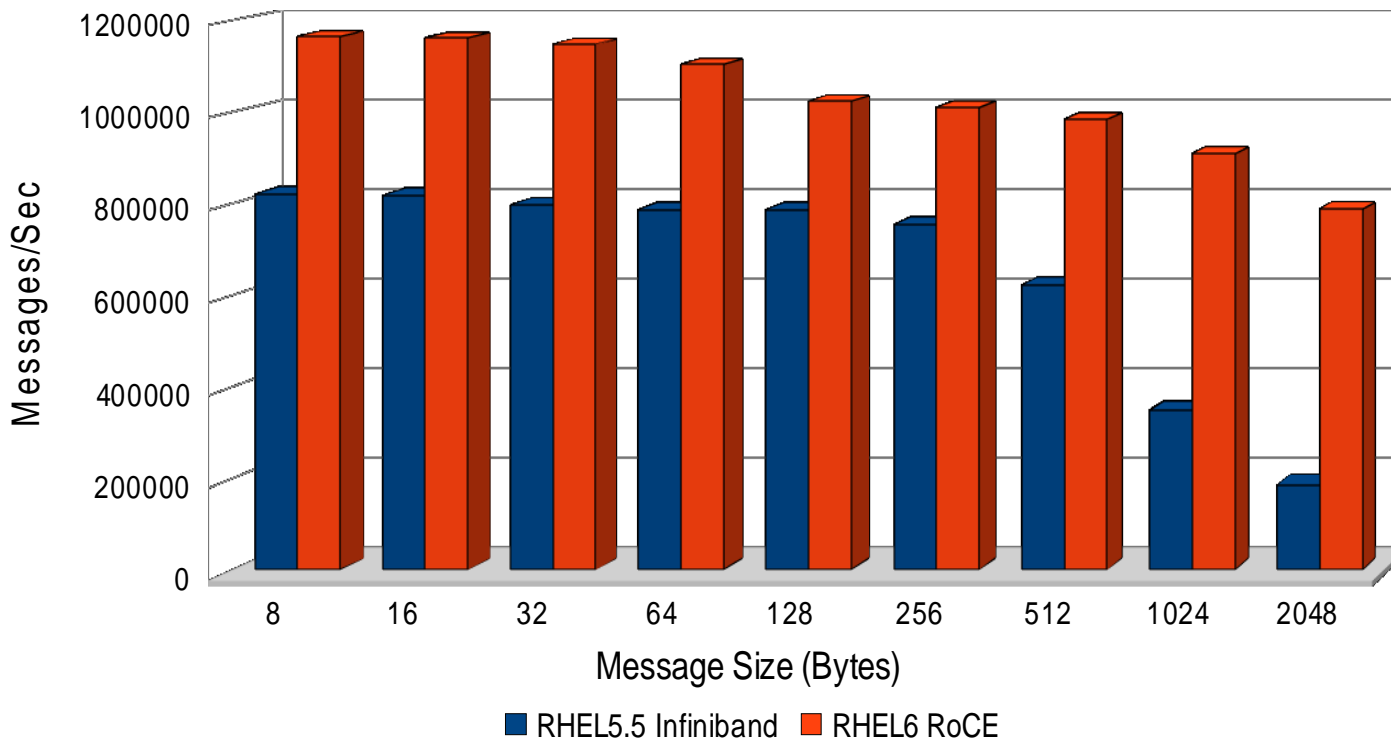
10 Gbit Ethernet (Mellanox)




# RED HAT ENTERPRISE LINUX + High Performance Network

- RDMA over Converged Ethernet (RoCE) support
- Replaces Infiniband and TCP transport layers for RDMA applications

Comparing RHEL55 Mellanox Infiniband and RHEL6 Mellanox 10Gb with RoCE



# RHEL6 IB vs 10Gb RDMA IP vs Converged Ethernet (RoCE) (Latency msec - Smaller = Better)



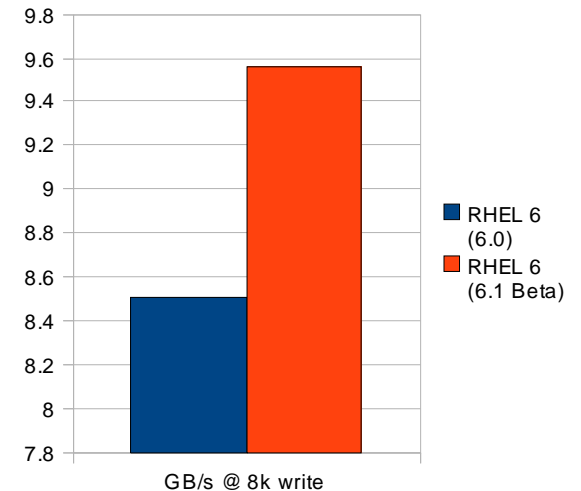
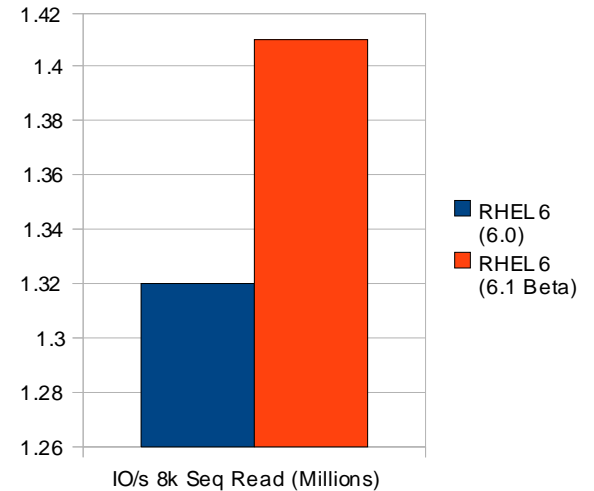
# Storage

- Storage architectures vary between workloads
  - Local storage: SAS, SATA, SSD, Flash
  - Traditional Storage Networks
    - Fibre Channel
  - Converged Storage Networks
    - FCoE
    - iSCSI



# Storage - IO/sec

- RHEL 6 and Fusion-io cards
  - 1.32 million IO/Sec @ 8kb Read
  - 12.5 GB/sec – peak write performance
  - Intel Boxboro-EX 32-core, 128 GB, 8x Fusion-io duo



| IO/Sec            | RHEL 6.1 Beta (2.6.32-94) | RHEL 6.0 GA (2.6.32-71) | % Gain ( 6.1 Beta / 6.0 ) |
|-------------------|---------------------------|-------------------------|---------------------------|
| IO/s 8k Seq Read  | 1406300                   | 1320188                 | <b>6.5%</b>               |
| IO/S 8k Seq Write | 1253313                   | 1114141                 | <b>12.5%</b>              |
| GB/s 8k Seq Read  | 10.73                     | 10.07                   | <b>6.5%</b>               |
| GB/s 8k Seq Write | 9.56                      | 8.5                     | <b>12.5%</b>              |

Disclaimer: RHEL 6.1 Beta results may not be indicative



# Bringing the Community, Vendors and Users Together

