



Red Hat Opinions and Musings on this brave new world

Author: Doug Ledford <dledford@redhat.com>

Date: March 27, 2012

OFED stuns us

- OFED commits to basing releases on upstream kernel releases
- There was much dancing and rejoicing in the streets (the offices were too confining for the level of dancing and rejoicing)
- UNH IOL decides to start testing distro packages
- We tried to dance and rejoice some more but we were still groggy from the last go around
- But all is not complete rainbows and unicorns, the UNH IOL testing looked pretty abysmal :-)

What tripped up the IOL tests?

- The "out of the box" criteria for one thing
- Differences in the OFED stack relative to our stack
- Lack of knowledge on the part of the UNH people about how to configure a RHEL InfiniBand offering

What now?

- Work with OFED to fine tune methods of collaboration
- Work with UNH IOL to define the differences between OFED and Red Hat in terms of configuration so that the test results for future logo runs will be more realistic of the true state of the stack
- Work with UNH IOL to possibly modify the test plans to work better with distros (we aren't on an OFED release schedule, so twice a year tests prior to OFED releases doesn't jive with us well)

Impacts on RHEL5

- RHEL5 used the old OFED
- Red Hat is extending the lifetime of RHEL5 and later to 10 years
- There will be one more OFED update on RHEL5 to sync to a later version of various drivers. This will not be a full stack update, but will provide updated hardware support and bug fixes

To read more...

- The press release that goes along with this change in lifecycle:
<http://www.redhat.com/about/news/press-archive/2012/1/red-hat-enterprise-linux-stability-drives-demand-for-more-flexibility-in-long-term-operating-system-deployments>
- The Red Hat Enterprise Linux life cycle web page:
<https://access.redhat.com/support/policy/updates/errata/>
- An FAQ on the Red Hat Customer Portal:
<https://access.redhat.com/kb/docs/DOC-69647>

RHEL5.5 IB to RHEL6 AMQP w/ RoCE mess/sec and MB/sec (Bigger=Better)

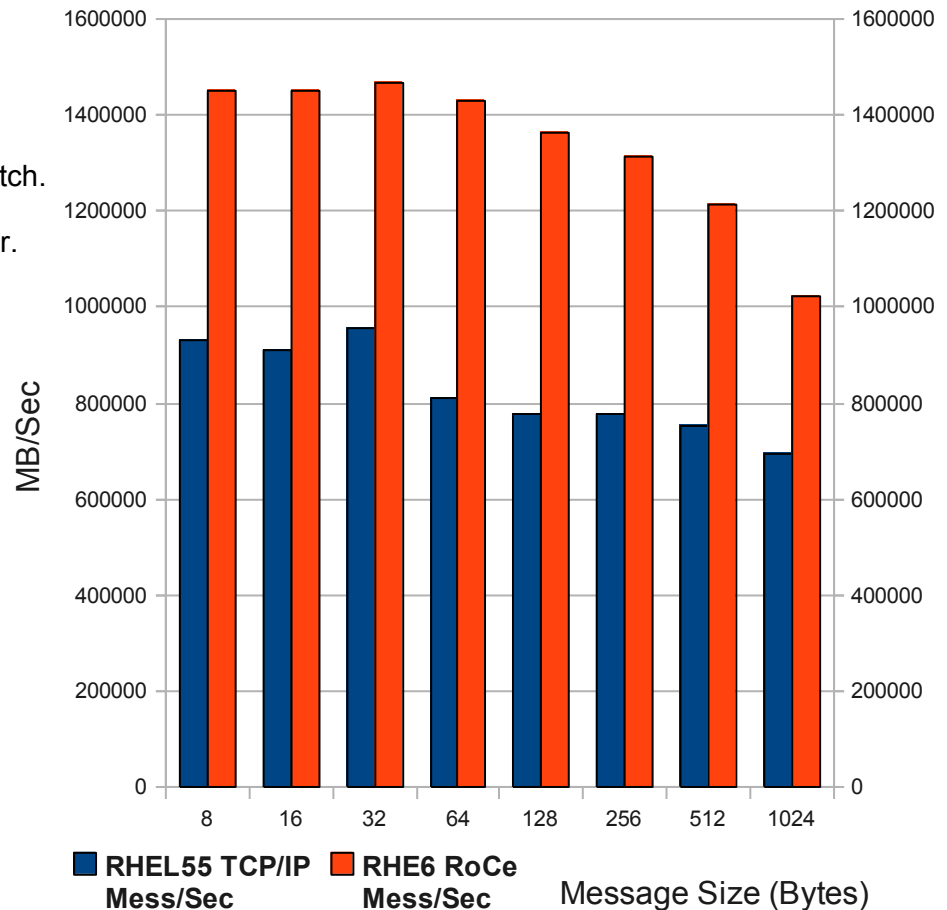


Mellanox 10Gb RHEL55/RHEL6 RC1 MRG1.3 Comp:

Network Enhancements

- Multi-queue send/recieve
- Tools to monitor dropped packets – tc, dropwatch.
- RCU adoption in stack
- Multi-CPU receive to pull in from the wire faster.
- 10GbE driver improvements.
- Data center bridging in ixgbe driver.

Reliable over Converged 10GB Ethernet RoCE
Chart -
20% kernel gain
20% w/ RDMA over 10Gbit



10 Gbit Ethernet (Mellanox)