

Performance Analysis of iSCSI & iSER in MPIO environment

**Seikh Basiruddin
NetApp India Pvt. Ltd.**

Agenda...

- ❑ iSCSI introduction.
- ❑ iSCSI implementation Vs today's Hardware
- ❑ RDMA introduction
- ❑ What is iSER (iSCSI Extension of RDMA)
- ❑ Hardware Platforms used
- ❑ Performance Data & Comparison
- ❑ Summary

- ❑ Internet Small Computer Systems Interface
- ❑ Emerging Storage Area Networking Technology
- ❑ Allows block-level access to storage devices, such as disks, over a computer network.
- ❑ Use TCP/IP Stack of the host OS
- ❑ Has two modules “initiator” & “target”

Advantages

- ❑ Leverages the existing LAN(Ethernet) infrastructure
- ❑ Less expensive, hence reduce the TCO (Total cost of ownership)

iSCSI implementation Vs today's Hardware

- ❑ Evolution of 10 gigabit Ethernet
- ❑ Availability of RNIC

Disadvantage of Today's Software iSCSI

- ❑ Can't utilize the 10gigabit bandwidth
- ❑ Data-copy overhead in standard TCP/IP layer
- ❑ All the Processing & interrupt handling done at host's CPU.

What is RDMA?


- ❑ **DMA:** "*Direct Memory Access.*" This is a hardware mechanism in modern computers that allows some I/O devices to directly transfer blocks of data into main memory without the intervention of the CPU

- ❑ **RDMA:** "*Remote Direct Memory Access.*" This technology takes DMA a step farther,
 - ❑ Allows a remote device/computer to perform a DMA transfer across a network
 - ❑ Avoid constant CPU intervention on either end of the transfer
 - ❑ Help to achieve faster data transfer

What is iSER (iSCSI Extension of RDMA)

- ❑ **New Standard defined by IETF**
- ❑ **Use RDMA technology to create session and transfer data**
- ❑ **It's a hardware assisted technology**
- ❑ **Follow most of the SCSI standard. Minimal change in SCSI command**
- ❑ **Can avoid the data-copy over head of the TCP/IP layer of the host.**
- ❑ **Take advantage of 10 gigabit Ethernet**

iSER Stack



Hardware

- ❑ Two IBMx3650 machines
- ❑ Two Chelsio RNIC
- ❑ Optic Fibre Cables

Software

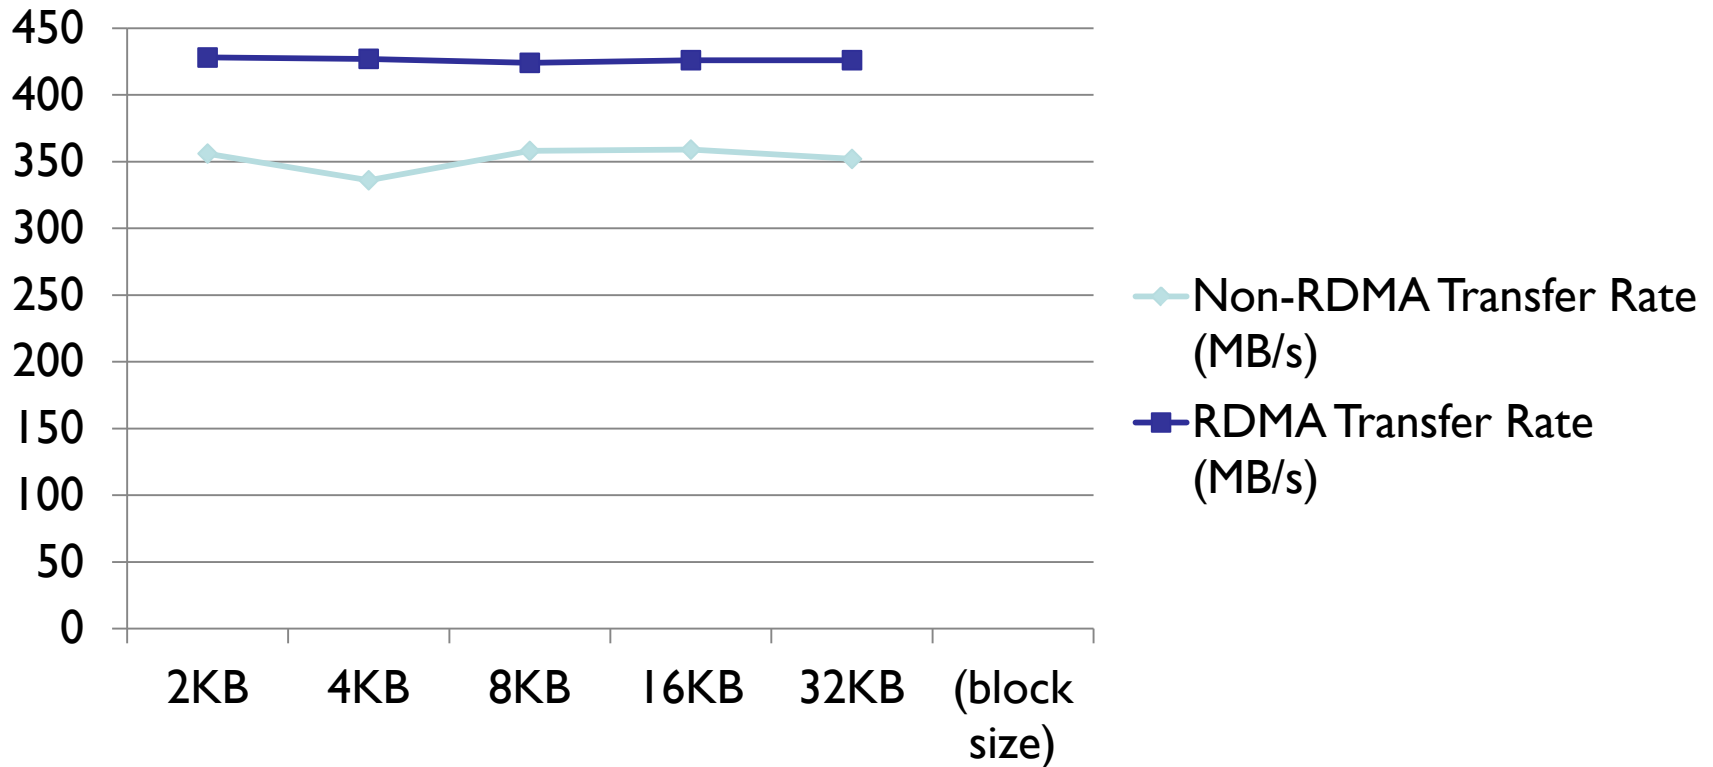
- ❑ Red Hat Enterprise Linux 5.3
- ❑ Chelsio Driver: 1.3.1.9 Firmware: 7.4.0
- ❑ OFED 1.4.1
- ❑ UNH-iSER2.2

Performance Data & Comparison

Trial	No of Threads	Block Size (KB)	Non-RDMA Transfer Rate (MB/s)	RDMA Transfer Rate (MB/s)	Impact on throughput
1	2	2	356	428	+20.3%
2	2	4	336	427	+27.1%
3	2	8	358	424	+18.4%
4	2	16	359	426	+18.7%
5	2	32	352	426	+21.0%

Approximately 20% performance improvements while using RDMA compared to non-RDMA transfer.

Performance Comparison Chart



Summary...

- ❑ Performance is the main criteria for selecting the storage solution
- ❑ With Evolution of 10 gigabit Ethernet, users want high performance & cost effective storage access.
- ❑ iSER technology can meet the expectation of industry.
- ❑ Our analysis shows over all 20% performance improvement while using RDMA compare to normal software iSCSI

Thank You!!!

Question???