Oracle's Converged Network

Richard Frank, Oracle

Based on OFA RDMA sockets

- RDS is an OFA verbs client, using IB today.
- rds-bcopy for all IPC messages (DLM,PQ,BCACHE,etc),
- rds-rdma for all data xfer > 8k size
- RDS is available on Linux (v3), AIX (v2), Solaris (v2), HPUX (v3).
- RDS v3 requires heterogeneous interoperability. HPUX,Linux interoperate today.
- Investigating Windows Port (WinOF).

RDMA for everyone via Ethernet

- Today we can only use RDMA with IB and this is fantastic… but we need RDMA to be available on Ethernet too.

- We want RDMA everywhere - ideally it would be on every motherboard and all converged network adapters – today.

ROEE - rdma in a box

- as defined in Paul Grun's “RDMA over Ethernet Proposal”
- leverages (flow control, congestion, rate control, etc)…CEE
- may translate to simpler RDMA over Ether silicon...
- may lead to more industry support with faster uptake into the market.

ROEE

- ROEE preserves the channel interface semantics we are using today. So we know RDS (and all other OFA verbs clients will continue to work).

- This really is a big deal for us..and no doubt others.

ROEE

- We know the IB interface and protocol layer work well, scale well, and are very well tested - they have nearly 10 years on the floor.

- We along with many of you have many years of development and testing demonstrating this..

ROEE

- For Oracle to use RDMA over Ether (any solution) - it takes a about a year of testing - as many of you can personally atest to.

- We really only have the resources (and desire) to test one RDMA over Ethernet solution - we think it should be ROEE...

ROEE

- We need an RDMA over Ether solution now..
- We encourage you to consider the ROEE proposal and move forward to building and deploying it ASAP