



OFED Server OEM Panel

SGI

Edward Mascarenhas
March 15, 2010

Agenda

- SGI and OpenFabrics
- SGI contributions to OFED
- Requirements, Directions

SGI Products and OpenFabrics



- All SGI Servers support InfiniBand and iWARP/10GigE with OFED
 - Altix® ICE Clusters with Integrated IB backplane
 - Highly scalable with Hypercube, Enhanced Hypercube, or fat-tree topology
 - Allows expansion of cluster while in production mode
 - Single or dual InfiniBand fabrics
 - Altix® XE clusters and Rackable Servers
 - Altix® NUMA systems - Altix® 4700/450 (IA64) and Altix® UV (x86_64)
 - CloudRack Servers and Octane III Personal Supercomputers
- Storage products

SGI Products and OpenFabrics



- SGI currently shipping OFED 1.4.1 as part of SGI Foundation Software
- Transition to use OFED from SLES/RHEL for most servers
- Some large SGI OFED/IB deployments
 - NASA Ames Pleides cluster – 544.3TF
 - NMCAC – 133.2TF
 - CINES, ARL, HLRN

SGI Contributions to OFED

- DOR routing support for Hypercube topology
- Weighted DOR routing (for I/O fabric)
- OpenSM: Sweep function toggle, Log subnet name
- Other major patches
 - NUMA fixes (e.g., coherent CQ)
 - Per device /proc/interrupts
 - IPoIB fixes
- Push patches upstream and also to distros

Requirements, Directions

- Continue high quality releases of OFED with feature enhancements – keep up the good work!
- Continue Distro OFED releases
 - Prefer Distro OFED to closely follow community OFED release versions in timing and content
 - Naming and packaging consistency
- Continue compatibility of API between OFED versions
- Sample code required for API users when new features/API are added to OFED
- Improve IB storage support – should be like FC

Requirements, Directions

- Scaling to 10,000s of nodes
 - OpenSM, ARP, IPoIB – will they scale?
- Improve ability to diagnose failures and performance issues
 - better tools
 - Making sense of log errors/messages
 - Localizing failures to specific devices in the fabric
- EDR and FDR – smooth transition from QDR
- Adaptive routing, congestion management