Communication & Mass Storage Futures

Topics: Network Challenges

Mass Storage Challenges

Cray Strategies

Paul Krueger

Disclaimer

The challenges, processes, and technologies discussed in this presentation are speculative in nature. This does not represent a commitment by Cray Inc. to implement, sell, or support any particular technology or anything else discussed.

HPC I/O Challenges

- Performance
- Configuration
- Management

Network Performance Challenges

- **The bandwidth challenge: How do I achieve maximum transmission rates for large data transfers or distributed applications?**
- **The latency challenge: How do I minimize response times for distributed applications or interactive users?**
- **The transaction rate challenge: How do I efficiently handle large numbers of network packets?**

Cray's Approach to Network Performance Challenges

- Offload network functions to special purpose processors
- Provide remote direct memory access (RDMA) to Cray hosts to enable offloaded functionality

Network Configuration Challenges

- The MTU mismatch challenge: How do I efficiently configure high and low-bandwidth network connections in a single network?
- The processor utilization challenge: How do I configure to conserve the use of relatively expensive mainframe processors?
- The network complexity challenge: How many different kinds of networks do I need?

Cray's Approach to Network Configuration Challenges

- Break end-to-end TCP connections to allow MTU conversion
- Offload TCP to avoid end-to-end TCP MTU matching requirement
- Track single-network approaches such as Infiniband

Network Management Challenges

- The reliability/availability/serviceability (RAS) challenge: How do I diagnose network problems, provide and make use of redundancy, and minimize downtime?
- The security challenge: How do I avoid unauthorized system access?
- The upgrade challenge: How do I migrate my environment to newer network technologies without throwing away previous investments?

Cray's Approach to Network Management Challenges

- Conform to network management standards such as SNMP
- Provide diagnostics for Cray-provided network components
- Provide bridges to previous-generation network technologies

Mass Storage Performance Challenges

- The bandwidth challenge: How can I rapidly stream large volumes of data to or from storage?
- The transaction rate challenge: How can I process many I/O transactions in a short period of time?
- The simultaneous data access challenge: How can I have many processes simultaneously writing to a single file with reasonable performance?
- The filesystem scaling challenge: What filesystem software efficiently supports many user processes?

Cray's Approach to Mass Storage Performance Challenges

- Keep paths to data as short as possible
- Minimize OS data handling (e.g. techniques such as page-flipping)
- Investigate software that offloads OS involvement (e.g. DAFS)
- Use fastest components available
- Partner to acquire technology we cannot develop alone

Mass Storage Configuration Challenges

- The network challenge: What's the least expensive network I can use to get the storage access performance that I need?
- The processor overhead challenge: How do I minimize the number of processors that are tied up managing storage access?
- The heterogeneity challenge: How do I cost-effectively get systems from several vendors to share data?

Cray's Approach to Mass Storage Configuration Challenges

- Evaluate all storage networks such as fibre channel, Infiniband, ethernet, etc.
- Use RAID devices
- Evaluate file sharing and processor offload technologies such as SANs and NAS

Mass Storage Management Challenges

- The backup challenge: How do I backup large data volumes without impact to operations.
- The storage availability challenge: How can I assure that storage is available when needed?
- The data archiving challenge: How do I ensure that data that isn't immediately needed doesn't consume primary storage, but is still easily available when needed?
- The security challenge: How do I allow data sharing when needed without risking unauthorized access?

Cray's Approach to Mass Storage Management Challenges

- Provide a variety of backup options such as tape, high-speed networks, interfaces to standard backup packages
- Hierarchical storage at slower speeds from third-party servers
- Evaluate third-party filesystems for security problems

Conclusions

- Rapid hardware evolution creates HPC I/O challenges
- Vendor solutions can be at odds with HPC requirements
- Cray will adopt solutions when possible and develop them when necessary