

Data Analysis Challenge 2
HEC FSIO Workshop
Aug 2, 2010

Jim Finlayson

Department of Defense

james.m.finlayson@us.army.mil

State of I/O - Historical

- Presented at SGPFS Workshop, Santa Fe – Sept 23-24, 1999
 - Manage 1 trillion files over 5 year periods (>6300 files / second sustained for five years)
 - Seamless Hierarchical Storage Management (HSM) integration for disk, optical disk, and tape
 - Manage file sizes from bytes to terabytes
 - Seamless backup and recovery integration
 - 7x24x365 operation

State of I/O - Historical

- Available to variety of clients from desktops to supercomputers
- Provide disaster recovery hook for multiple copies, exporting to salt mine, etc
- Efficient addition and deletion of storage assets to/from the file system
- Ease of migration from legacy storage devices to modern storage devices

Stage of I/O - Historical

- Open/close/read/write interface with hints or ioctl calls to activate advanced features
- Proper dashboarding for lights out operation
- Very large number of clients

The Coming I/O Scenario

- 1 trillion files per file system
- 10 billion files per directory
- 100PB system capacity
- 1PB single file size
- >30K member nodes sharing the file system
- 100K open files concurrently
- End to end data integrity (T10-PI) or equivalent
- 100 hour “file system check”
- No discernable performance impact during RAID parity rebuild
- 40,000 file creates/second from a single node
- 10,000 directory listings/second
- 30GB/sec single client streaming
- 1.5TB/sec aggregated I/O across the system
- 1 IOP/GB (100PB = 100 million IOPS!!!!)

Questions?