



Hierarchical Storage Management, year 2022

Brad Blasing
Principal Software Engineer, SAM-QFS
September 21, 2012

Agenda

- Simpler
- Scalability
- Data Integrity
- Data Management

Safe Harbor: The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



Simpler

It's complicated to set up

- GUI Configuration
- Disaster Recovery configurations
 - Currently complicated to set up
 - Variety of ways of doing the same thing
 - This should be, will be point & click
 - HA configurations likewise complicated
- HSM Appliance
 - Variety of NAS/SAN storage appliances exist
- Media Migration
 - Make easier to adopt next generation tape/disk
 - Cadence is still every 2-3 years



Scalability

Moore's law has not been overturned by the courts



- Disk
 - No particular end in sight in disk growth
 - Expect 2.5 doublings of disk in 10 years
 - 50's of petabyte disk systems
- Tape
 - Expect 5 doublings of tape in 10 years
 - Oracle T10000C – 5 TB @ ¼ GB/sec – fairly awesome amount of data
 - Horizontally scale tape I/O
 - Will make use of partitioning to pack/recycle data
- Files
 - 1B files was last years problem
 - Need administrative scalability too (dump/restore/scanning/...)

Scalability

- Hostless Operations
 - Push tasks out towards devices, libraries
 - Data integrity validation scanning
 - Tape-to-tape copy: 2nd copy of files, media migration



Data Integrity – Self Repair

We don't have time for manual repair

- End to end data integrity
 - Not just tape scanning
 - User → channel → disk → daemon → tape → 2nd copy → ...and back
- Automated Data Integrity Validation
 - Done without recalling data (today)
 - Automate by policy
- Self repair when problems found





Data Location & Management


It's not just disk and tape

- Multiple pools of different kind
 - LOC outlined 6-7 tiers
 - Fast disk, bulk disk, SSD, big tape, small tape, long-term archiving
 - Cloud
- Data placement on tiers based on general guidelines
 - High capacity/cheap, high performance, super available/redundant
 - Don't care where data is, just as long as goals met
- Automation
 - Automatically migrate between tiers
- Data import/export
 - Simple data exchange
 - LTFS, version n+2

Call to Action

Mountains of Data

- There will be 2.5 to 5 doublings of data in next 10 years
- Automate
- Virtualize
- Simplify
- Integrity checking / self healing



Hardware and Software

ORACLE®

Engineered to Work Together