Emerging Storage Technology Panel

the MAID Perspective

Aloke Guha
CTO

22nd IEEE13th NASA Mass Storage Conference, Monterey, CA
April 14 2005

Different Data – Different Needs

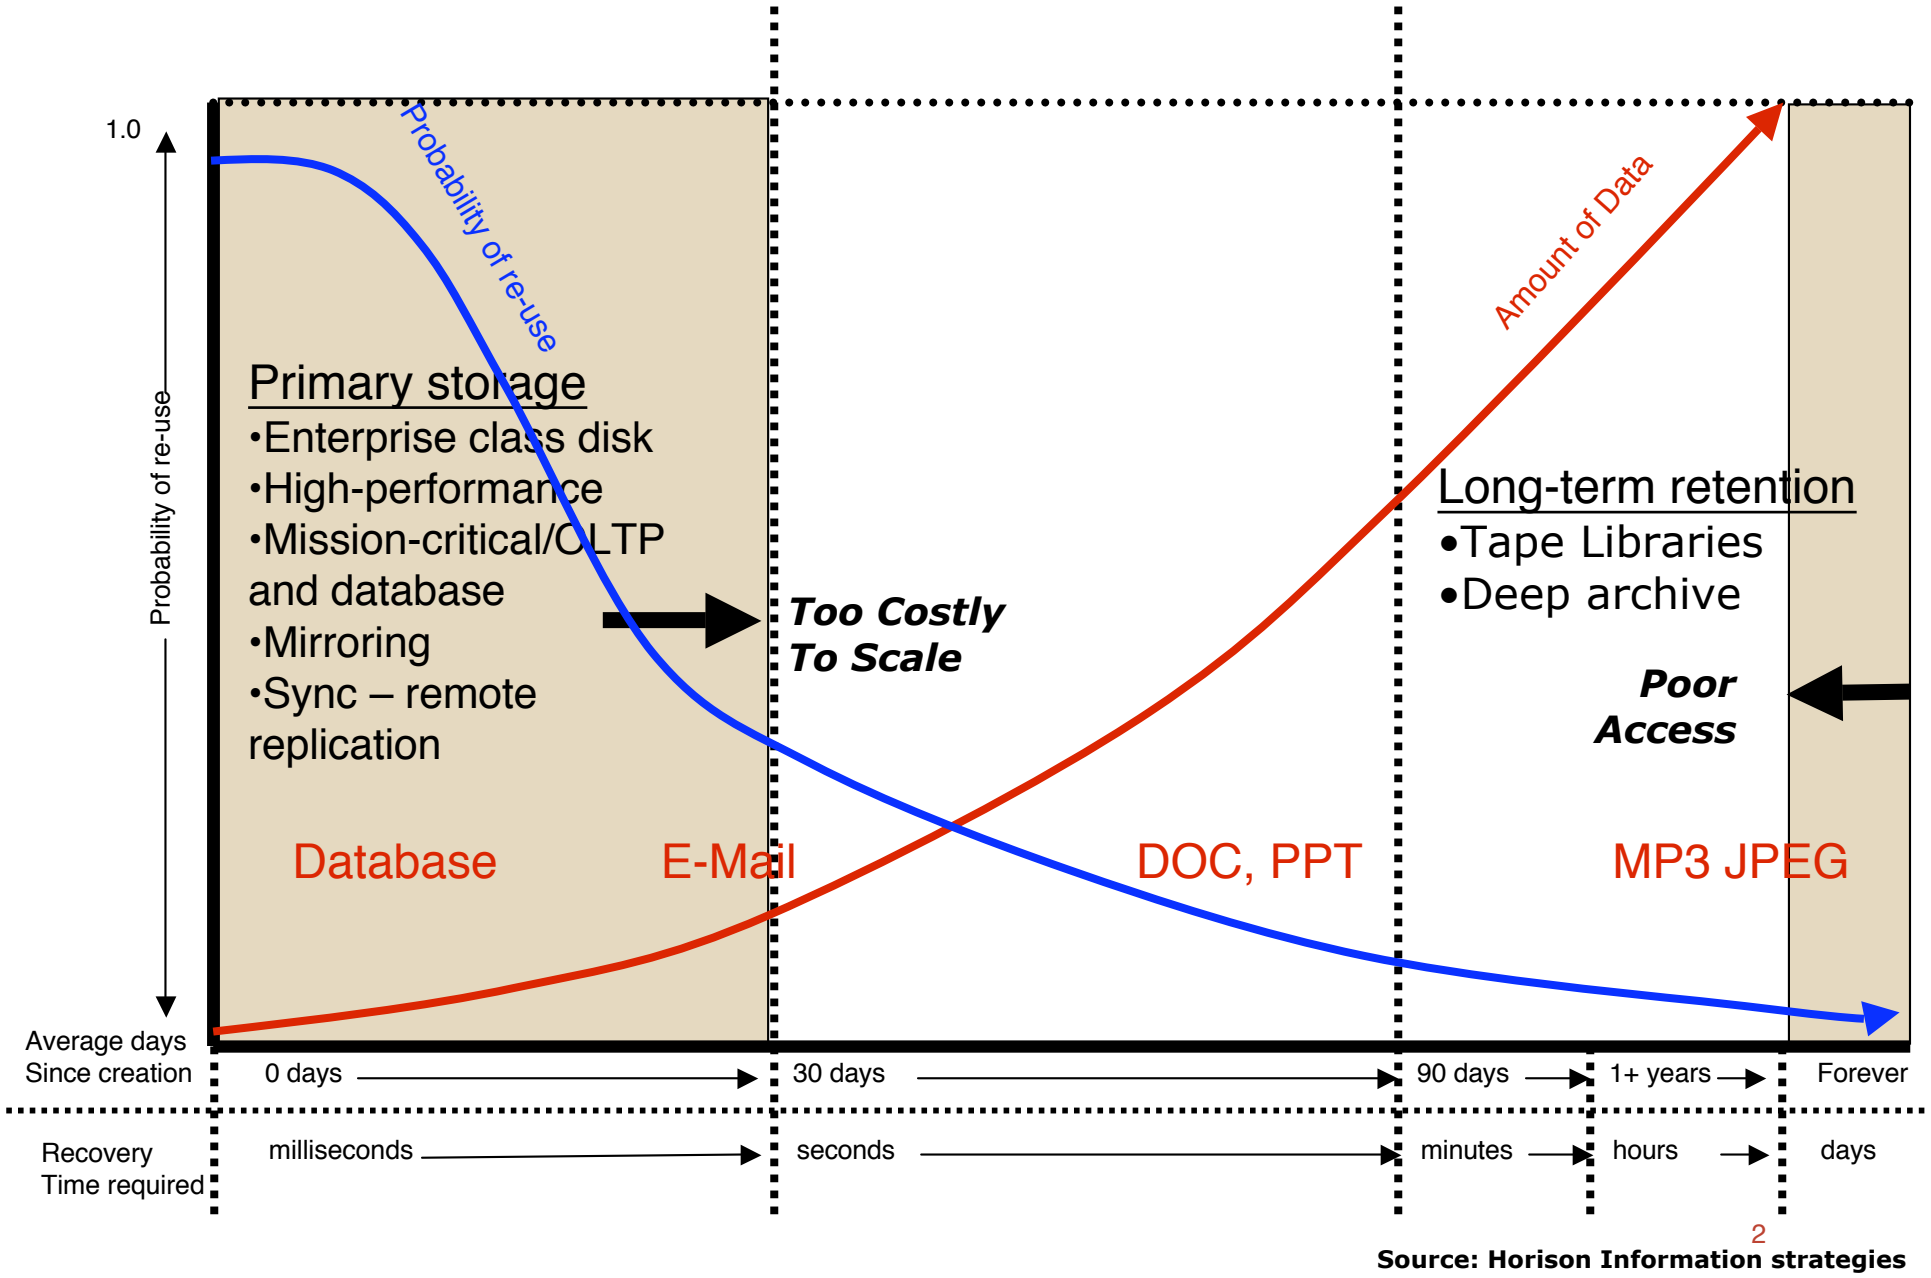

Source: Horison Information strategies

Tape Gets a Bad Rap

It is not as if tape storage became unreliable or that tape systems vendors vastly increased the cost of these solutions. . . . the pace with which the Information Age has emerged is having a ripple effect throughout all of IT…

In one sense, the requirements of today's "always on" business atmosphere are outpacing tape's ability to satisfy the needs of IT, and increasingly, the business.

Enterprise Strategy Group

"It's the workload, ___!"

MAID: New Tier in the Storage Hierarchy

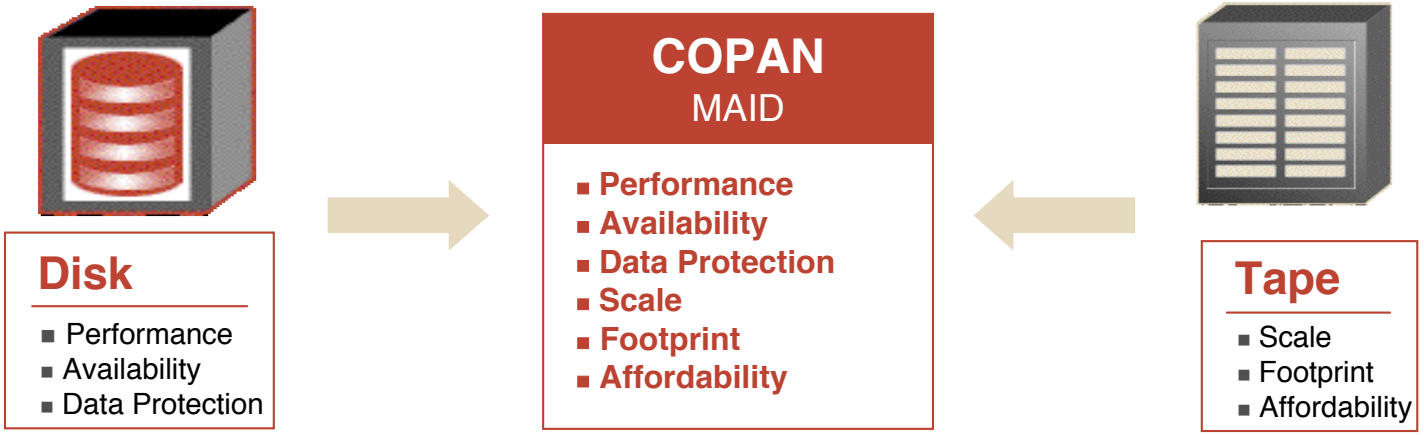

- *System solution to active data: accessible* large-scale data on power-managed disks
  - Disk-like performance
  - Cost in range of tape
- *Not* a replacement for inactive archive data on Tape
- *Not* a replacement for transactional data on Disk

Approaches in building MAID-based storage

- Storage/File Approaches
  - Controller w/ VTL software
  - File System on power-managed LUNs
  - Non-Traditional: purpose-built
- Purpose-Built (RISC): VTL, File or Disk
  - Optimize
    - Performance: WORO, mostly seq. access, MBs/sec, ≤ sec latency
    - Cost, Reliability, Serviceability, Data protection, . . .
  - Architectural implications
    - No need for large RAM cache
    - No need for non-blocking
    - No need to access all data all time

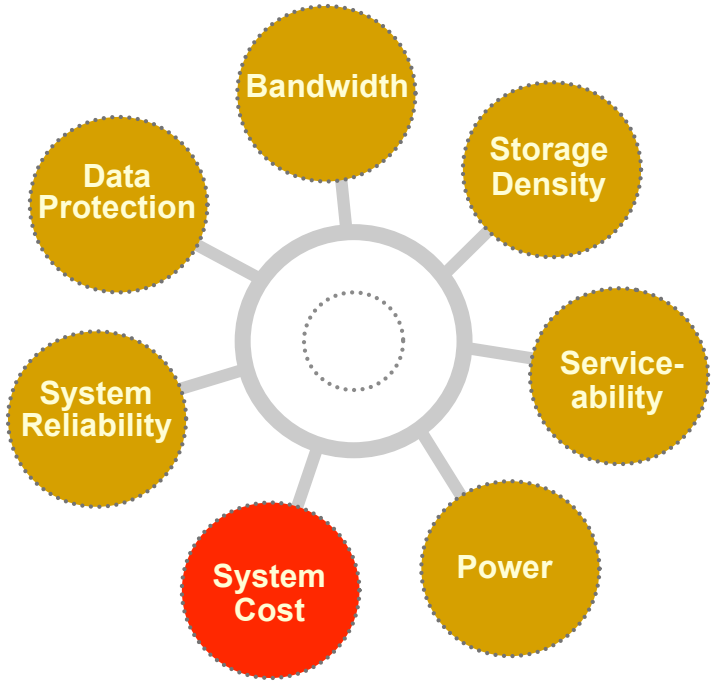

COPAN Approach: 3-Tier MAID Architecture

Three levels of processing separates functionality, simplifies management, scales performance w/ capacity

- Layer 2 – Storage Personality
  - **Physical Domain (Rack Controller)**
  - Storage Network Protocols
  - Logical Volume/Block Management
  - Performance and Load Balancing
- Layer 1 – Data Protection
  - **Physical Domain (Shelf Controller)**
  - RAID Support and Caching
  - Power Management
  - Device Management
- Layer 0 – Data Path Routing
  - **Physical Domain (Canister Controller)**
  - Protocol Router
  - Monitoring of environmental parameters

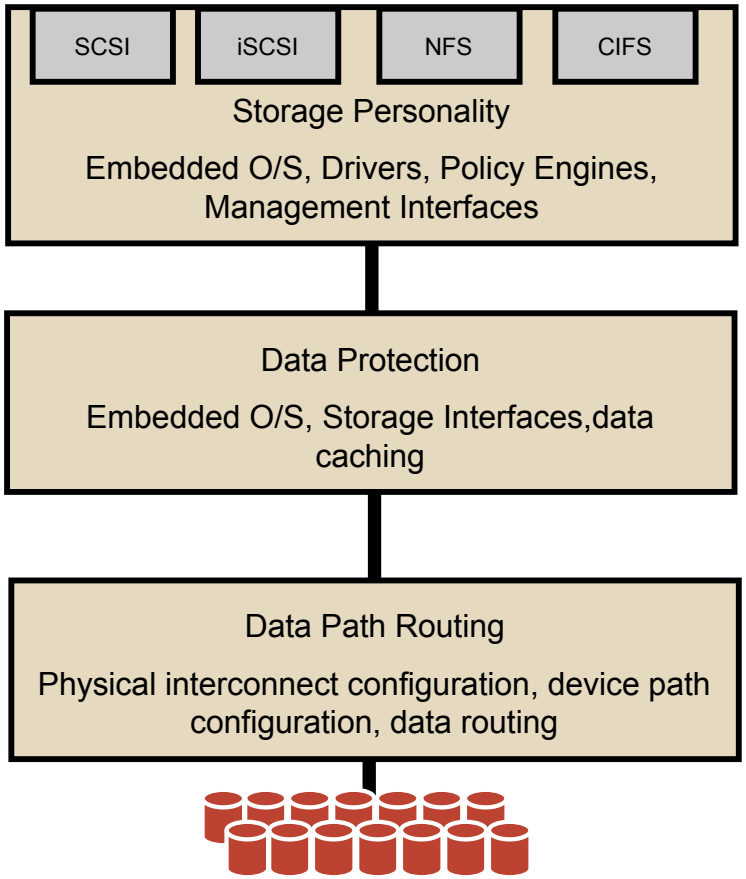

New Tiered Storage

Tier I Enterprise Disk

Critical and Replicated Data

Tier II Modular Array

Lower Priority, Non-Replicated Data

Tier III-IV MAID
(Massive Array of Idle Disks)

File Reference Information Regulatory Data

On Line Recovery

896 250GB drives
= 224 TB
(uncompressed)

22TB/sq.ft.
4 – 2Gb FC
2.4TB / HR

VTL

Tier V DR Off-Site Deep Archive

Disaster Recovery, Off-Site Data