Storage needs for Life Science Informatics

Geek Cred

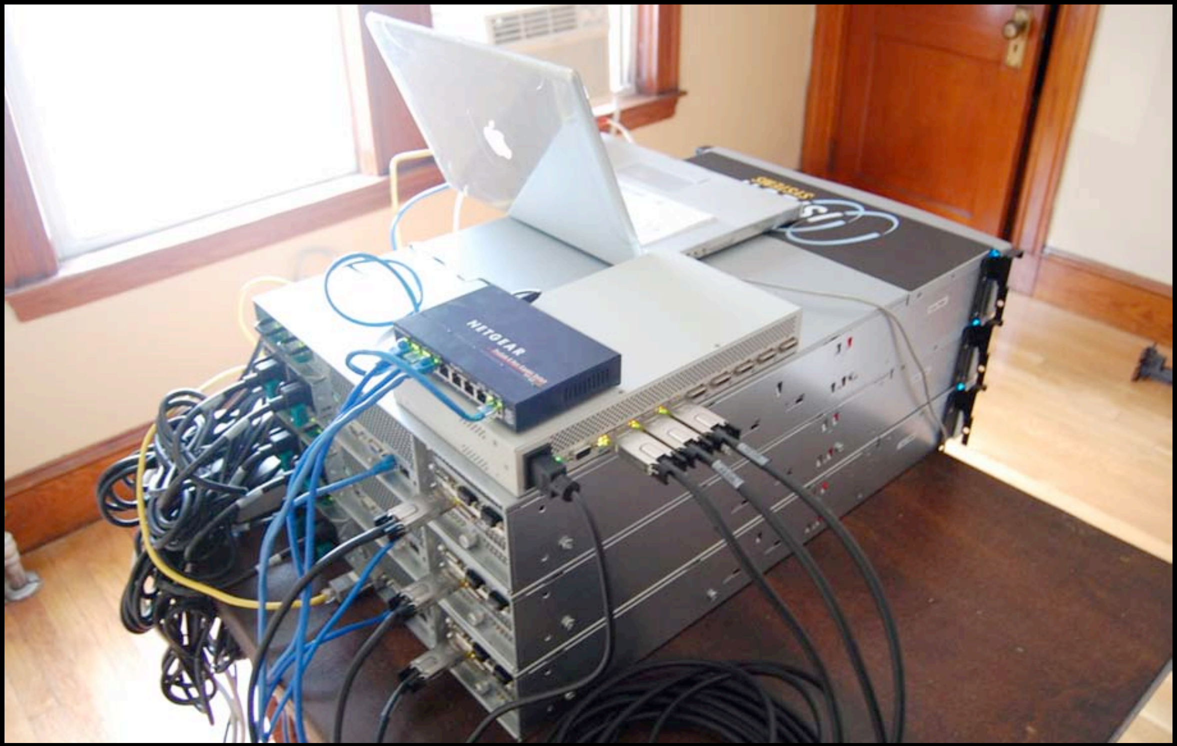

- Picture taken 7/30
- Isilon Evaluation
  - Symmetrically clustered storage
  - Infiniband interconnected storage nodes
    - Each “brick” is 3TB
    - Bricks autodiscover and join the storage cluster
  - Cluster NFS/NAS
    - Distributed NFS via standard ethernet

Requirements for Science

1. High capacity & high scaling headroom
2. Variable file types and access patterns
3. Multi-protocol access options
4. Concurrent read/write access

chris@bioteam.net

Capacity needs

Capacity Needs

- Life Science needs lots of storage
  - Already a cliché in 2006:
    - “Data Deluge”, “Data Tsunami”
- What changed in 2007:
  - Terabyte-scale laboratory instruments
    - Confocal microscopy
    - Imaging (fMRI, etc.)
    - “Next Generation” DNA Sequencers

Capacity Needs - continued

- Terabyte storage issues were “lab” or “workgroup” problems in the past …
- Now:
  - Individual researchers & individual lab instruments can generate a terabyte of data ***per-experiment***
  - Real world example:
    - 40TB storage for each Solexa instrument in small labs
    - If IT groups do not step in …
      - Expect to see Sun Thumper boxes crammed under benches in your wet labs

chris@bioteam.net

Capacity - HMS Dilemma

- "Data triage" easier to implement in enterprise environments
  - New trend:
    - Primary data around only for QC/QA; then deleted
    - Keep only the derived or distilled data online
- Very hard in academic environments to tell staff that they must delete primary data
    - May be unavoidable ...

chris@bioteam.net

File types & access patterns

File types & access patterns

- Many storage products are optimized only for certain use-cases and file types
  - Life Science requires them all:
    - Many small files vs. fewer large files
    - Text vs. Binary data
    - Sequential access vs. random access
    - Concurrent reads against large files

Protocol Requirements

Multi-protocol storage needs

- Storage Area Networks (SANs) are not the best storage platform for discovery research environments
- The *overwhelming* researcher requirement is for *shared* access to *common* filesystems
  - Lab instrument, cluster nodes & desktop workstation can all see the same data
- Shared storage in a SAN world is non trivial to implement

Storage - Protocol requirements

- Simultaneous access via:
  - NFS
    - Standard method of filesharing between Unix hosts
  - CIFS/SMB
    - Mount shared storage on Windows desktop
    - Ideally with authentication & ACLs coming from Active Directory
  - FTP & HTTP
    - Distribute datasets and large files among collaborators
- IP-SAN and/or FC-SAN abilities
  - We still need block storage LUNs for Oracle, etc.

Concurrent Access

Concurrent Access

- Ideally we want read/write access to files from:
  - Researcher desktops
  - HPC / Cluster systems
  - Lab instruments
- If we don’t have this:
  - Lots of time & core network bandwidth consumed with data movement
  - Lots of data stored in multiple locations
  - Harder to secure; harder to backup (if at all …)
  - 30TB NAS arrays start showing up under desks and in nearby telco closets

Real world example

- From audit interviews conducted this week
  - A Children's Hospital Lab:
    - Team develops applications & data on local workstations
    - Everything then replicated to Orchestra for analysis
    - Results data must be replicated back
- Quote: "*Each one of my staff spends more than an hour per day replicating data via rsync to and from Orchestra*"

General Observations

Observations

- Storage is “cheap” & getting cheaper
- Operational costs seem to be remaining the same
- Backup and data continuity costs are exploding
  - I’m in awe of the backup people who still manage to stay afloat in the age of 1TB SATA disks

chris@bioteam.net

Observations continued

- End users have no clue regarding the true costs of keeping data online & available
  - *"I can get a terabyte from Costco for $220!"*
  - Significant outreach is needed
- Storage as a centralized resource makes sense:
  - Data continuity & security
  - Avoid 30TB islands scattered across every lab
  - Access to enterprise-class features:
    - Management software/tools that actually work
    - Thin provisioning & spindle virtualization
    - "Green" power management
    - Tiered storage & data deduplication options
    - Play nicely with VTLs and other tech that makes backup tasks easier

Final thoughts

- HMS needs a forward looking research storage roadmap
  - The rise of “terabyte” instruments is already having a major effect on storage environments
    - We see individual labs deploying 100TB storage systems
    - If a lab needs 100TB; what does “HMS Community” need?
  - Petabyte-scale needs will appear within the decade
    - Faster if disruptive technology (Church Lab …) appears

Final thoughts continued …

- The successful storage projects
  - Are closely aligned with researcher requirements
- Backup will continue be the hardest problem
- Personally: I think data triage is unavoidable
  - Simply not possible to give researchers unlimited storage; given costs of backup & management

chris@bioteam.net

Conclusion

- It may make sense for HMS to distinguish between “research” and “enterprise” storage platforms
- Future storage approaches should be evaluated based on:
  1. Compatibility with existing backup methods
  2. Ability to satisfy end-user requirements
  3. Lowest possible operational burden
  4. Sufficient scaling headroom

chris@bioteam.net