



Berkeley Storage & Backup Strategy

Common Solutions Group
May, 2010

Shel Waggener
(on behalf of the Berkeley team)



IST

Storage History @ Berkeley

2003: Every department on campus (including seven different departments within central IT organization runs own storage and backups.

2004: Data center move creates opportunity for common architecture for some central systems and backups

2006: Dedicated storage group formed. No further central storage purchases supported except through storage team.

2007: Hitachi wins bakeoff. 250 TB . Email team works with Storage group to move from direct attached to SAN

2010: Over 500 hosts backed up – 1.25 PB expanding to 3 PB this year.



IST

IST SAN – 2010 Technology

Serial Attachment SCSI (SAS) disk – Performance comparable to Fiber Channel disk at significantly lower cost

Virtual Tape Library (VTL) – Improved performance for UCBackup and support for Business Resumption

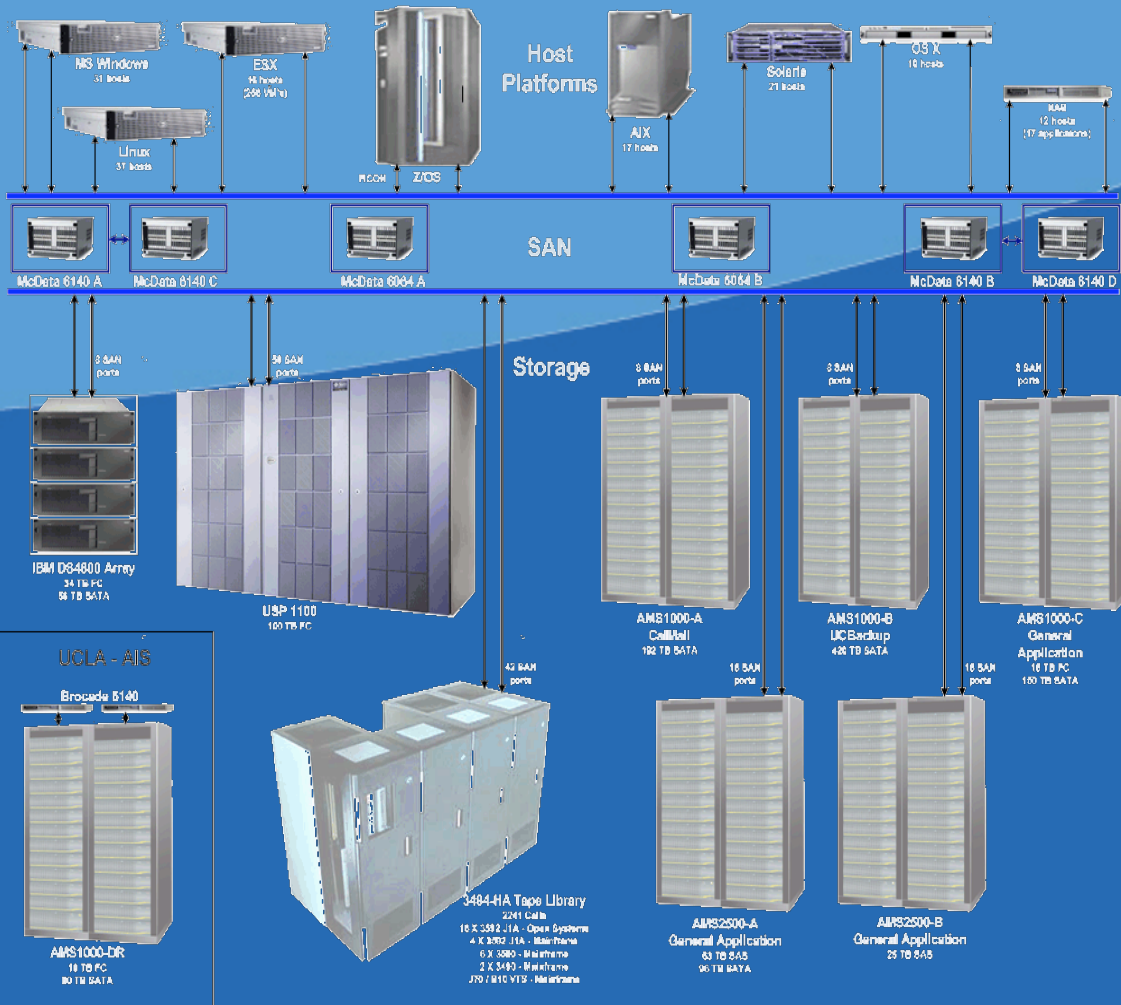
Clustered Network Attached Storage (NAS) - Consolidate standalone file servers while maintaining discrete, secure departmental environments

2 TB SATA – Potential to reduce Economy and Mass tier storage rates

Solid State Disk (SSD) - SAN solution for applications with extreme storage performance demands

AFS Deployment expanding

IST SAN - 2010

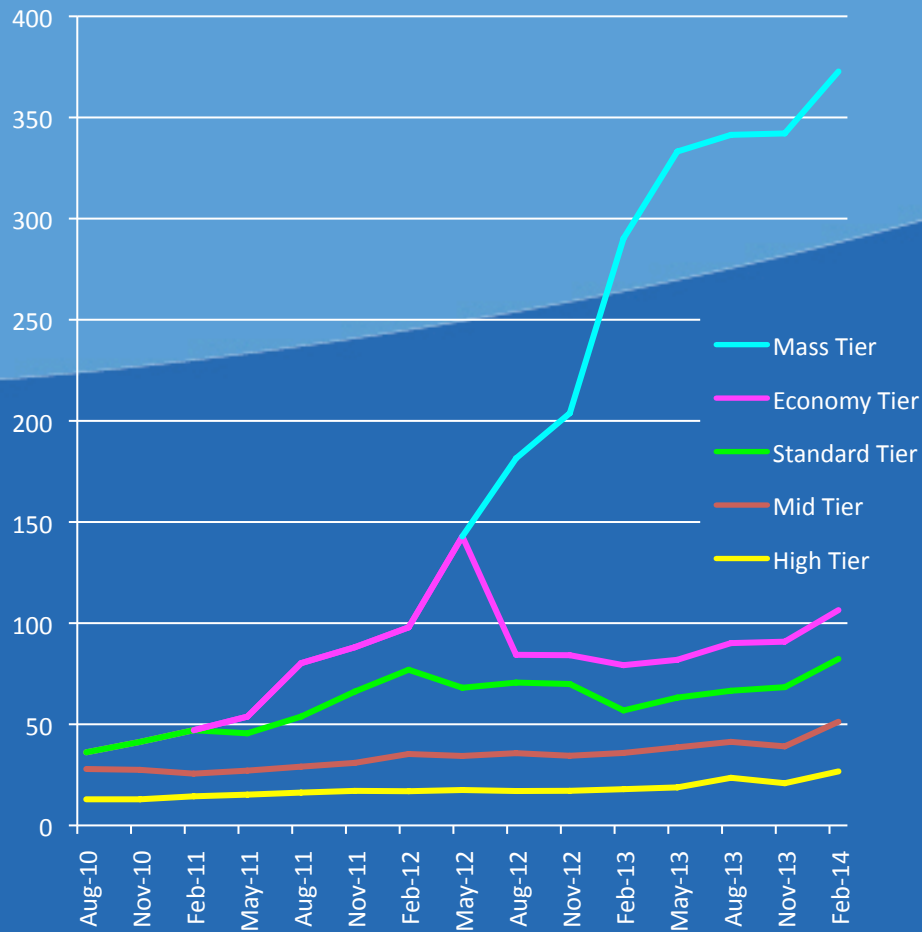


- High-availability enterprise arrays, tape library and FC network
- Five price/performance storage tiers
- 1.25 PB total raw disk capacity
- FC, SAS and SATA technology
- 768 Fiber Channel (FC) ports
- 170+ SAN hosts
- 7 Host OS platforms
- FC, iSCSI and NAS connectivity
- 350+ Virtual Machines
- Thousands of campus users
- High Availability Mirror to UCLA

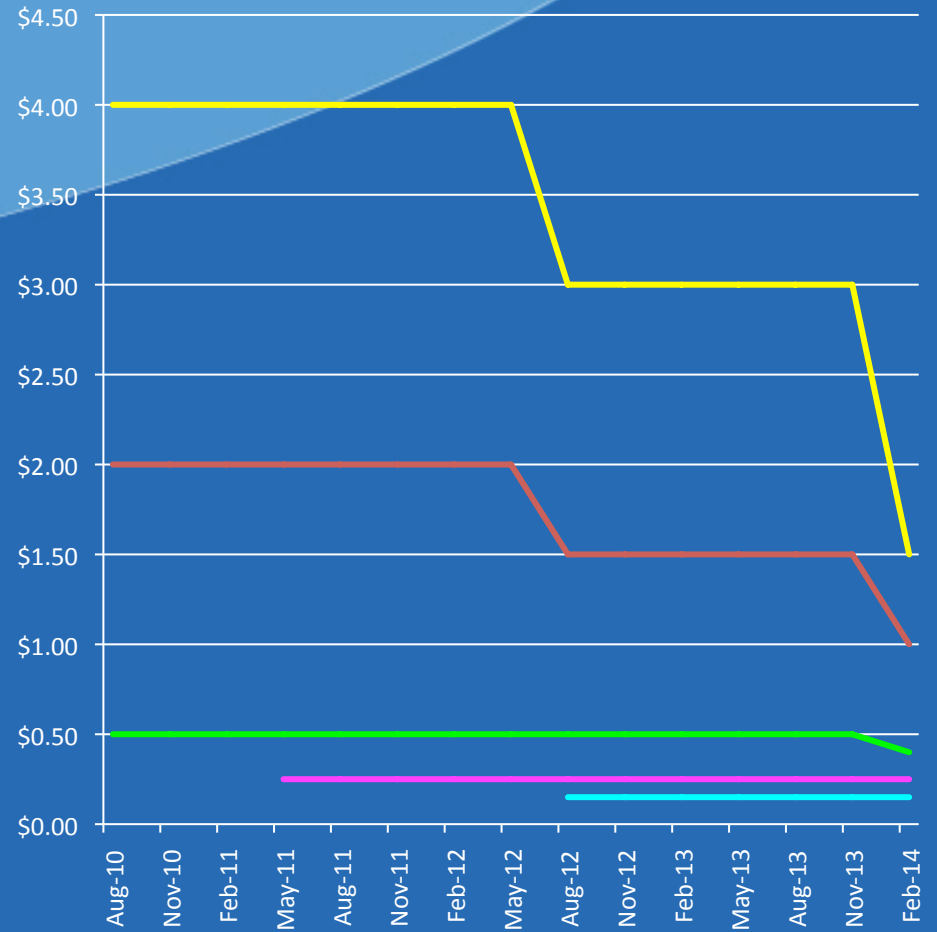


UC BERKELEY IST


Total Allocated SAN Capacity, Terabytes



Rates by Storage Tier, \$/GB/Month




SAN Utilization Central Uses





SAN Utilization - Other IST & Campus





Storage Futures

NAS Working Assumptions

- Scalability
 - Estimated 2 PB of shared, unstructured data across campus
- Security
 - Restricted Data management, offsite access
- Platform/Protocol Support and Standards Compliance
 - CIFS, NFS, HTTP, WebDAV, PNFS
- Minimum Cost



Storage Futures

NAS in the Marketplace

- Scale-Out NAS
 - Single managed environment supports thousands of concurrent users, consolidates “islands” of shared data
- Global Namespace
 - Facilitates inter-department data accessibility and security, reduces duplication of storage capacity
- Tiered Storage & Performance
 - Automated alignment of tier allocations with application requirements, from archives to HPCC
- Integration & Specialized Features
 - Gateway solutions; snapshots, replication, HSM



UC BERKELEY

IST

NAS Vendors and Products

Vendor	Product	Gateway	Cluster Type Max. Nodes	Namespace	Usable Capacity	Protocol Support
NetApp	DataONTAP 8 Cluster Mode	Yes	Scale Out 24 nodes	cluster-wide	14 PB	NFS CIFS
BlueArc	Mercury 100	Yes	N-active 8 nodes	cluster-wide	2 PB/256 TB	NFS CIFS HTTP iSCSI FTP
Isilon	IQ X	No	N-active 144 nodes	cluster-wide	45 PB	NFS CIFS HTTP FTP
IBM	SONAS	No	N-active 8 nodes	cluster-wide	14. PB	NFS CIFS HTTP SCP FTP
EMC	Celerra/MPFS	No	N+1 8 nodes	per node	896 TB	NFS CIFS iSCSI FTP
Oracle/Sun	?	No				



Application Example: Desktop Backups

How can we simplify backups and reduce costs for our mobile and client computing customers?



Desktop Backups on NAS

- IST Internal Desktop Backups
 - Logout script performs xcopy of selected directories and contents to CIFS share (“H:” drive)
 - 338 users, 1.3 TB currently
 - Users can check their folders on H: to verify backups, only most recent file versions retained
 - Entire CIFS share is backed up with TSM
- Campus departments could implement and administer this internally, with storage the only external cost



IST

DC Desktop Backup Systems

Vendor	Product	Client O/S Support	Session Connection Method(s)	Encryption
Code 42 Software	CrashPlanPro	MS Windows Linux MacOS Solaris	Schedule CDP	Yes
Symantec	Backup Exec Desktop Laptop Option	MS Windows	Schedule CDP (NTFS only)	Yes
IBM	Tivoli Continuous Data Protection	MS Windows	CDP	Yes
Open Source / GPL (sourceforge)	BackupPC	MS Windows Linux	No client SW	
Connected	Backup for PC	MS Windows MacOS	Schedule CDP	Yes



UC BERKELEY

IST

Storage Futures

Cloud Options

- Link Campus NAS into Cloud Providers?
 - Duraspace Integration?
 - UC Systemwide Storage Solution?
 - Direct with Commercial: ie: Amazon Simple Storage
- File Sharing through Cloud: Institutional Sharing?
 - Eliminated Xythos (done)
 - Common Contract with Dropbox?
- Student and Faculty Portfolios?
 - Alumni Offerings



IST Storage & Backup Group

Jack Shnell (jshnell@berkeley.edu), Supervisor

Joe Silva (silvaj@berkeley.edu), Senior Storage Administrator

Dennis Leong (spgdxl@berkeley.edu), Storage Administrator

Jim Neal (jrneal@berkeley.edu), Senior Backup Administrator

Tatiana Moll (tmoll@berkeley.edu), Backup Administrator

<http://ist.berkeley.edu/files/StorageForum20100315.pdf>

storageteam@lists.berkeley.edu

storage-ticket@lists.berkeley.edu

ucbackup@berkeley.edu

ucbackup-ticket@berkeley.edu