

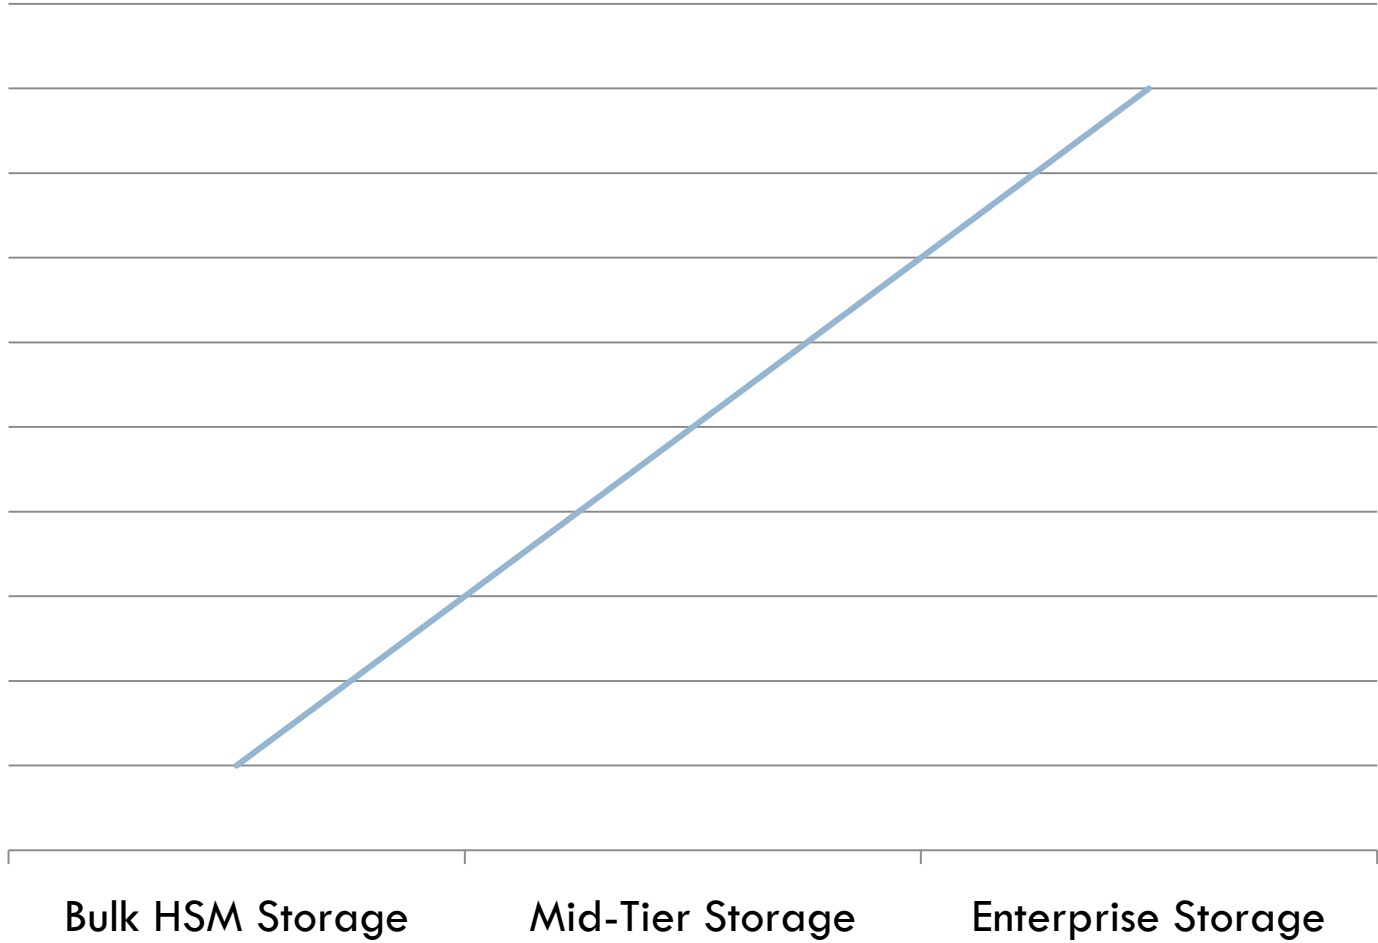
STORAGE SERVICES AT VIRGINIA

CSG Meeting
May 2010

Jim Jokl
University of Virginia

Storage Landscape @ Virginia

Cost and Performance



Hierarchical Storage Managers

- Services
 - ▣ Storage for long-term research data
 - Centrally funded but not well marketed
 - ▣ Library materials
 - Funding via Library contributions to infrastructure
 - ▣ RESSCU
 - Off-campus service for departmental Disaster Recovery backups
 - Centrally funded
- Technology #1 – Quantum StorNext
 - ▣ Quantum/ADIC Tape Robots in two buildings
 - ▣ Disk Cache accessible via NFS
 - ▣ ~94 TB in use for research data and library materials
- Technology #2 – Tivoli HSM for backups
 - ▣ On and off campus (@ Virginia Tech) tape robots

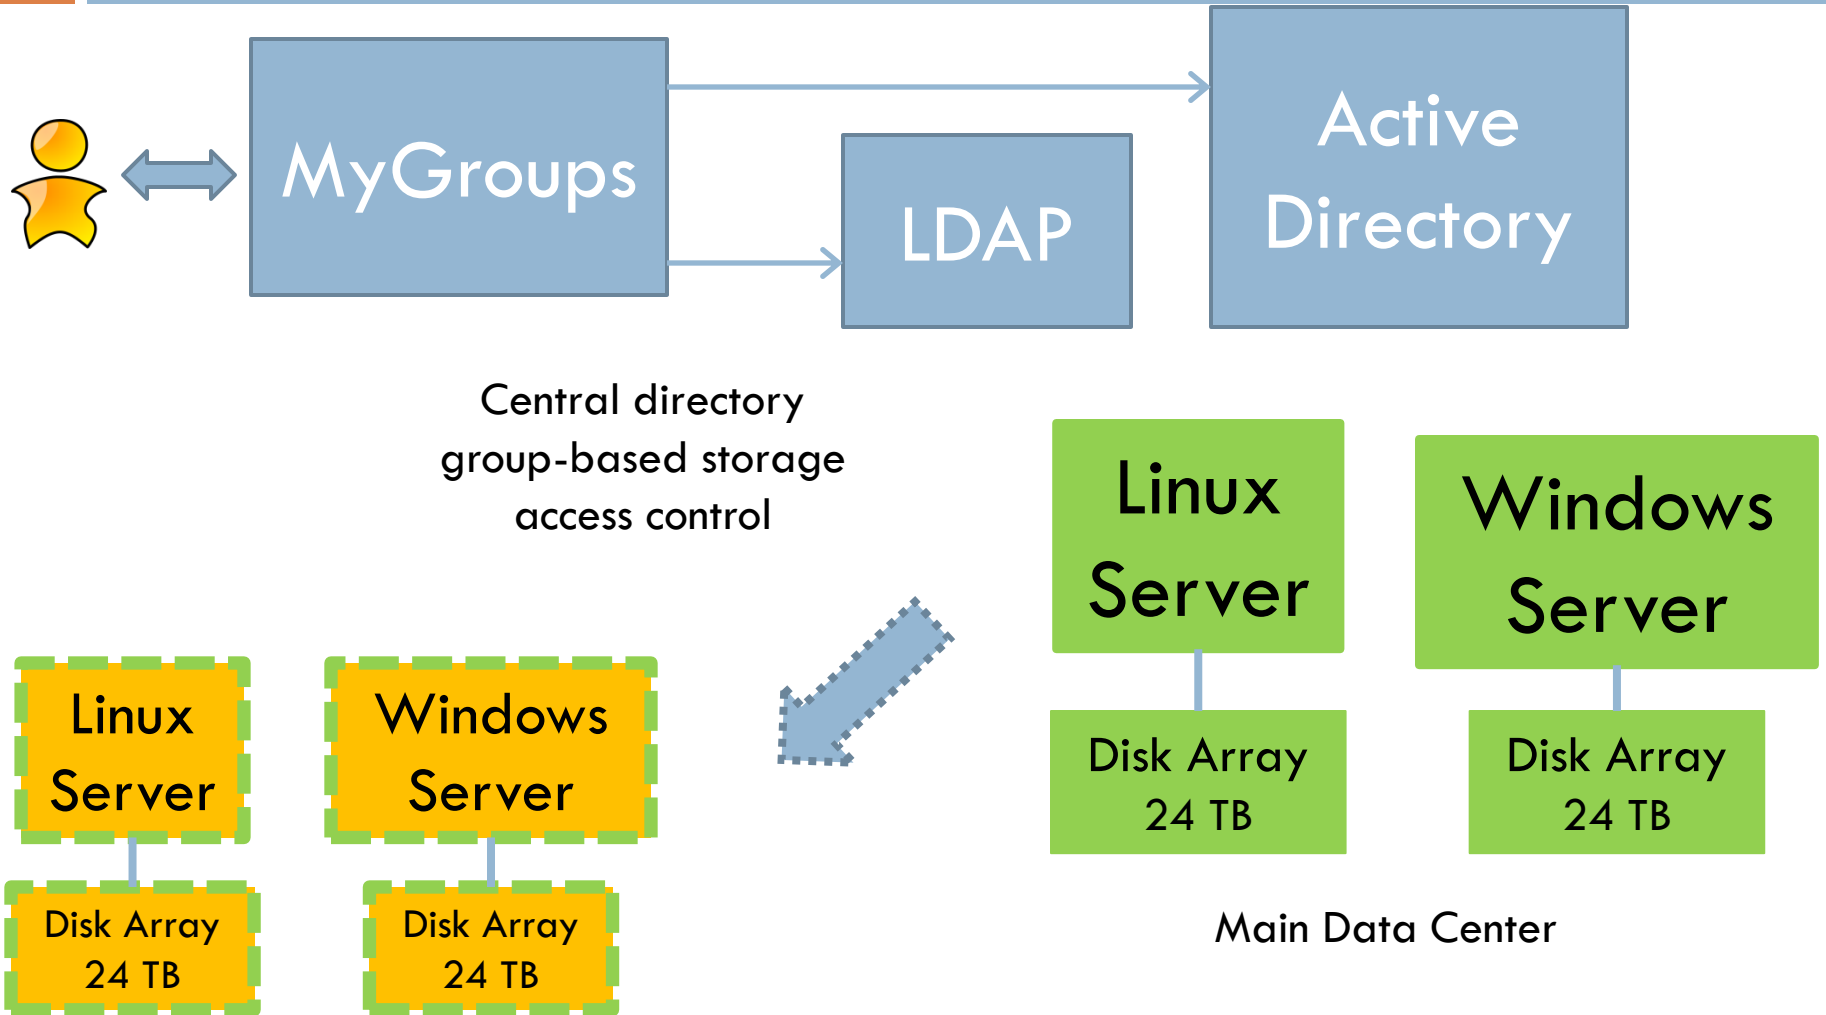
Enterprise Storage

- Based on Network Appliance file clusters
 - ▣ NFS, CIFS, iSCSI, SAN
 - ▣ The technology that we use to operate central services
 - ▣ Highly reliable, high performance, much cross site replication, etc
- Centrally Funded Services
 - ▣ 2 GB available on most central services
 - ▣ Services such as email, home directory, etc, etc
- Cost Recovery Services
 - ▣ Fee-based expansion to central service
 - Larger email quotas, home directory expansion, etc
 - ▣ Lease the type of storage that we use for large central applications
 - CIFS and NFS only; CIFS authn via central AD
 - ~12 TB leased
 - Annual rates match the type of hardware used to build the service
 - 3.20/GB to \$3.50/GB per year without backups
 - \$5.40/GB to \$6.10/GB per year with off-site backups

New Mid-Tier Storage Service

- Project Focus
 - ▣ A reasonable *and affordable* storage service for departments and researchers
- Requirements
 - ▣ Reliable, low cost, low overhead, self service
 - ▣ Unbundled services
 - Optional remote site replication
 - Optional backups
 - ▣ Access via NFS and CIFS
 - ▣ Snapshots; users deal with their own restores
 - ▣ Self-service storage access control

Mid-Tier Storage System



Remote Data Center Replication: optional

Mid-Tier Storage Status

- Build storage the way that a department would do it
 - ▣ Focus on automation to keep costs down
- Cost Model
 - ▣ Setup plus 1 hour plus \$0.34/GB per year
 - ▣ With remote replication \$0.67/GB per year
 - ▣ Next year we expect rates to drop by around 50%
- Status: after about one year in service
 - ▣ ~22 TB leased – NFS service
 - ▣ Marginal Windows use to date
 - Central AD adoption prerequisite for Windows/CIFS slows transitions
 - ▣ Lots of energy/discussions re: transitions, timing, when to adopt
 - ▣ We no longer hear complaints about the cost of central storage

Some References

- Disk Services
 - <http://itc.virginia.edu/datastorage/>
- Off-site Disaster Recovery Storage
 - Remote Electronic Storage for Safeguarding the Continuity of the University
 - <https://www.itc.virginia.edu/unixsys/resscu/>
- Research Data HSM
 - <http://itc.virginia.edu/research/hsm.html>
- Group-based Access Control (MyGroups)
 - <https://mygroups.virginia.edu/>