Resource Provider Spotlight: Globus Storage & IT Services

GlobusWORLD 2012

David Champion – dgc@uchicago.edu
Jeremiah Stuppy – jstuppy@uhicago.edu
Anita Nikolich – anikolich@uchicago.edu

THE UNIVERSITY OF CHICAGO

IT Services

State of Affairs

- Working relationships in science divisions, Computation Institute
  - ATLAS
  - Ad hoc support to individual PIs
- "Research Storage"
  - In the wind (but not here yet): Institutional Archive
  - Enterprise NAS (EMC^2 Isilon)
  - Highly available, highly scalable, highly priced
  - Goal: make mid-scale, reliable storage available for critical data

Approach

- Globus Online plan
  - Globus Transfer opened 2011
  - Globus Storage in development
- Mutual interest in collaboration
  - Globus wish: a pattern of institutional/commercial provider cooperation
  - ITS wish: better integration with IT partners and resources
  - Shared technological experiences and interests

Approach

- Proposal
  - ITS provide Isilon for Globus Storage pilot & production
  - Upon launch, ITS becomes a campus provider of storage resources via Globus Online
  - Jointly develop “retail” model for access and growth
- Benefits
  - Globus: additional storage backend
  - Globus: foundation for developing service models
  - ITS: partnership with domain leaders
  - ITS: platform experience

Campus Provider Overview

- Incremental approach to service provisioning
- Simple deployment

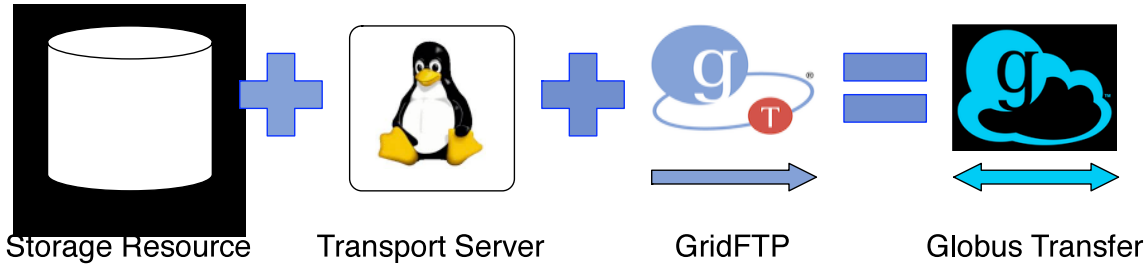

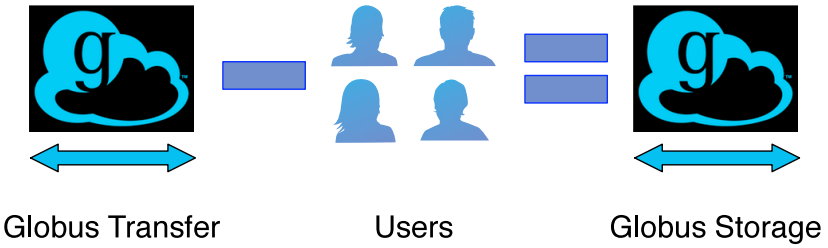

- Pure infrastructure — who can say no?

Globus Storage Architecture

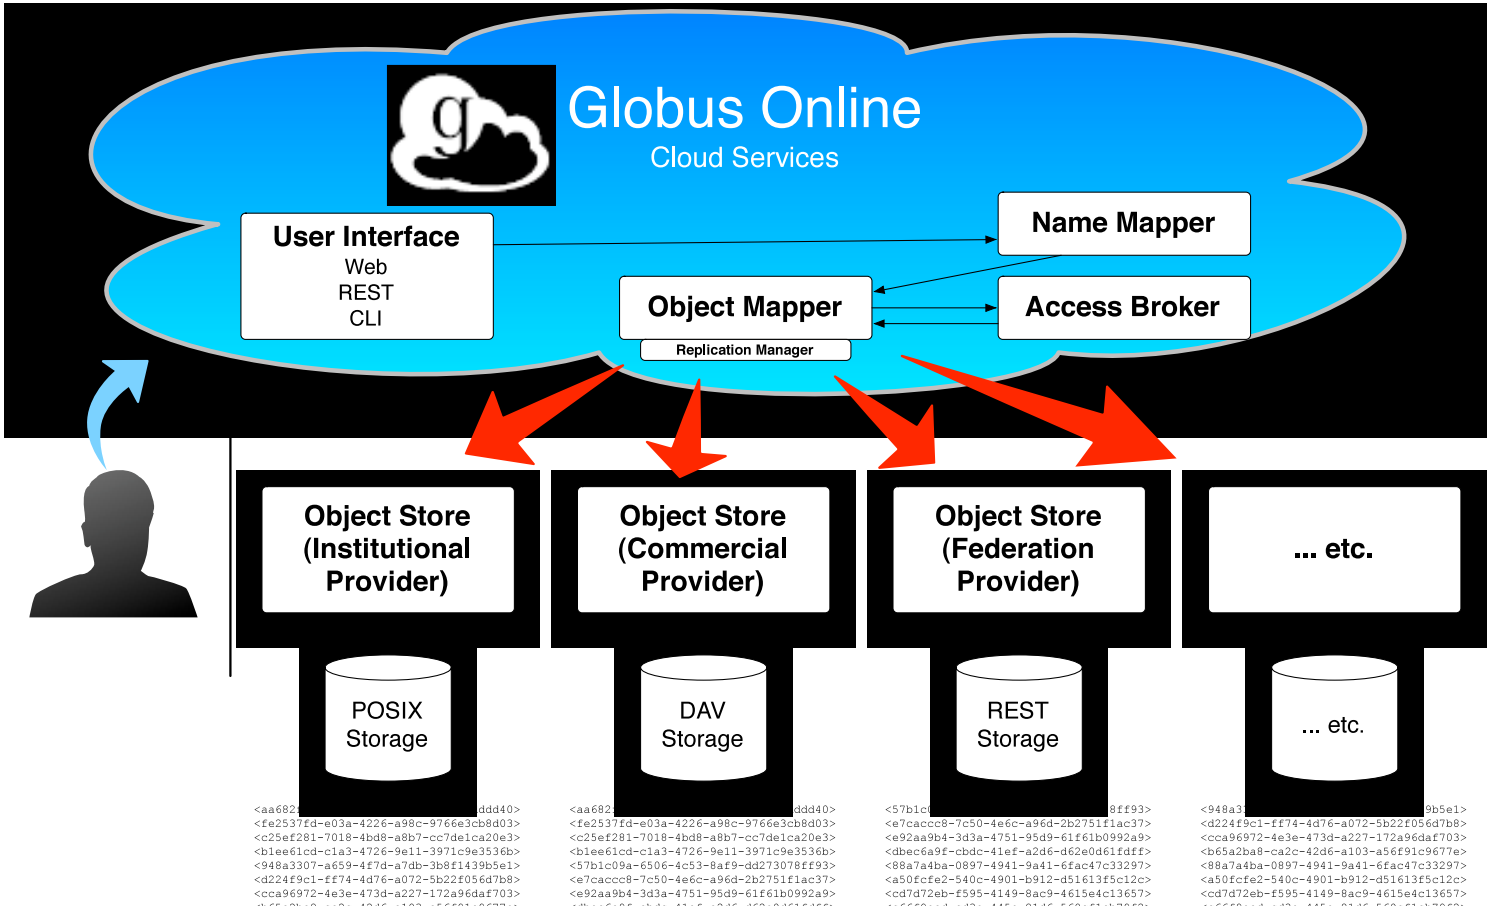

THE UNIVERSITY OF CHICAGO

ITServices

Globus Storage Architecture

- Filesystem Emulation
  - Access portal (UI) requests named files
  - Name Mapper maps named files to objects identified by UUID
  - Access Broker manages policy for requested object

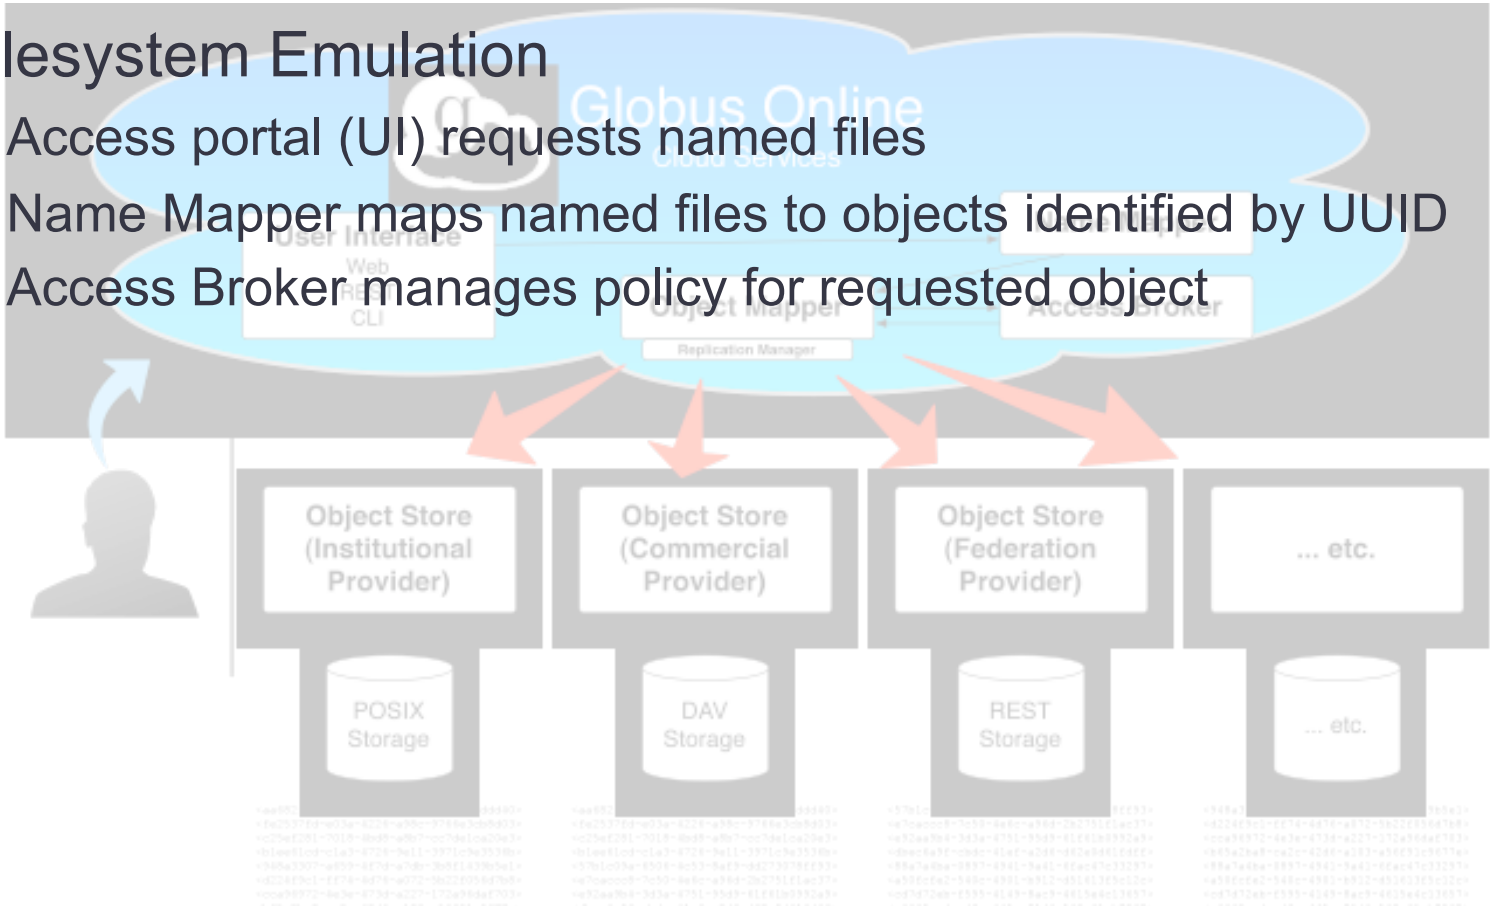

Globus Storage Architecture

- Object Management
  - Replication Manager may distribute multiple copies of object across disparate object stores
  - May choose object store based on network proximity or other factors
  - Replication and proximity policy selectable by user
- Object stores may have differing policies
  - User eligibility (institutional vs commercial)
  - Backup policy
  - Retention
  - Etc.

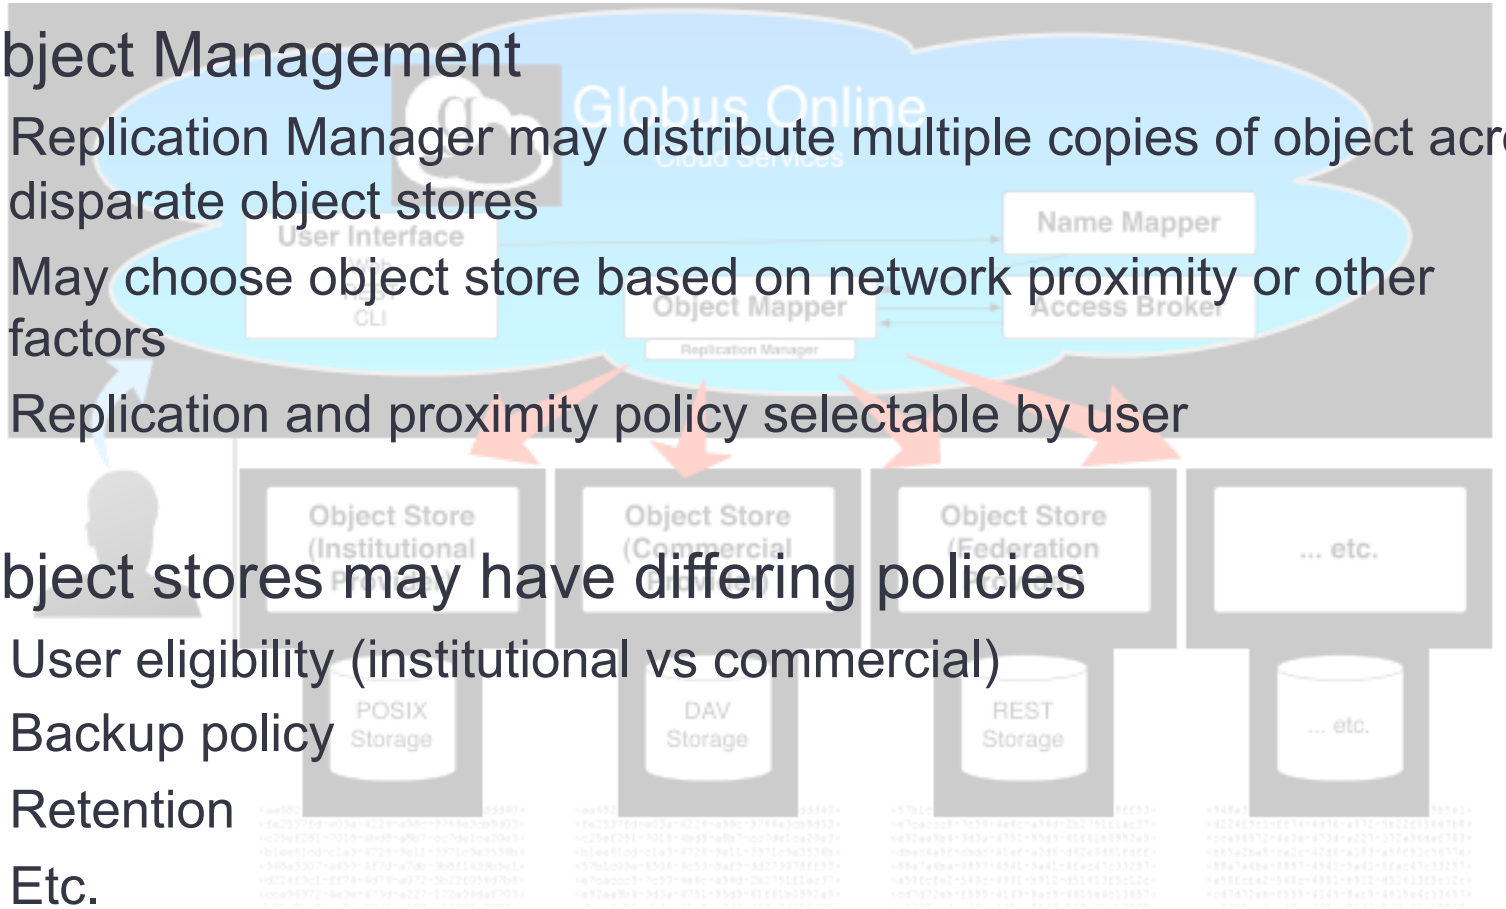

Globus Storage Partnership

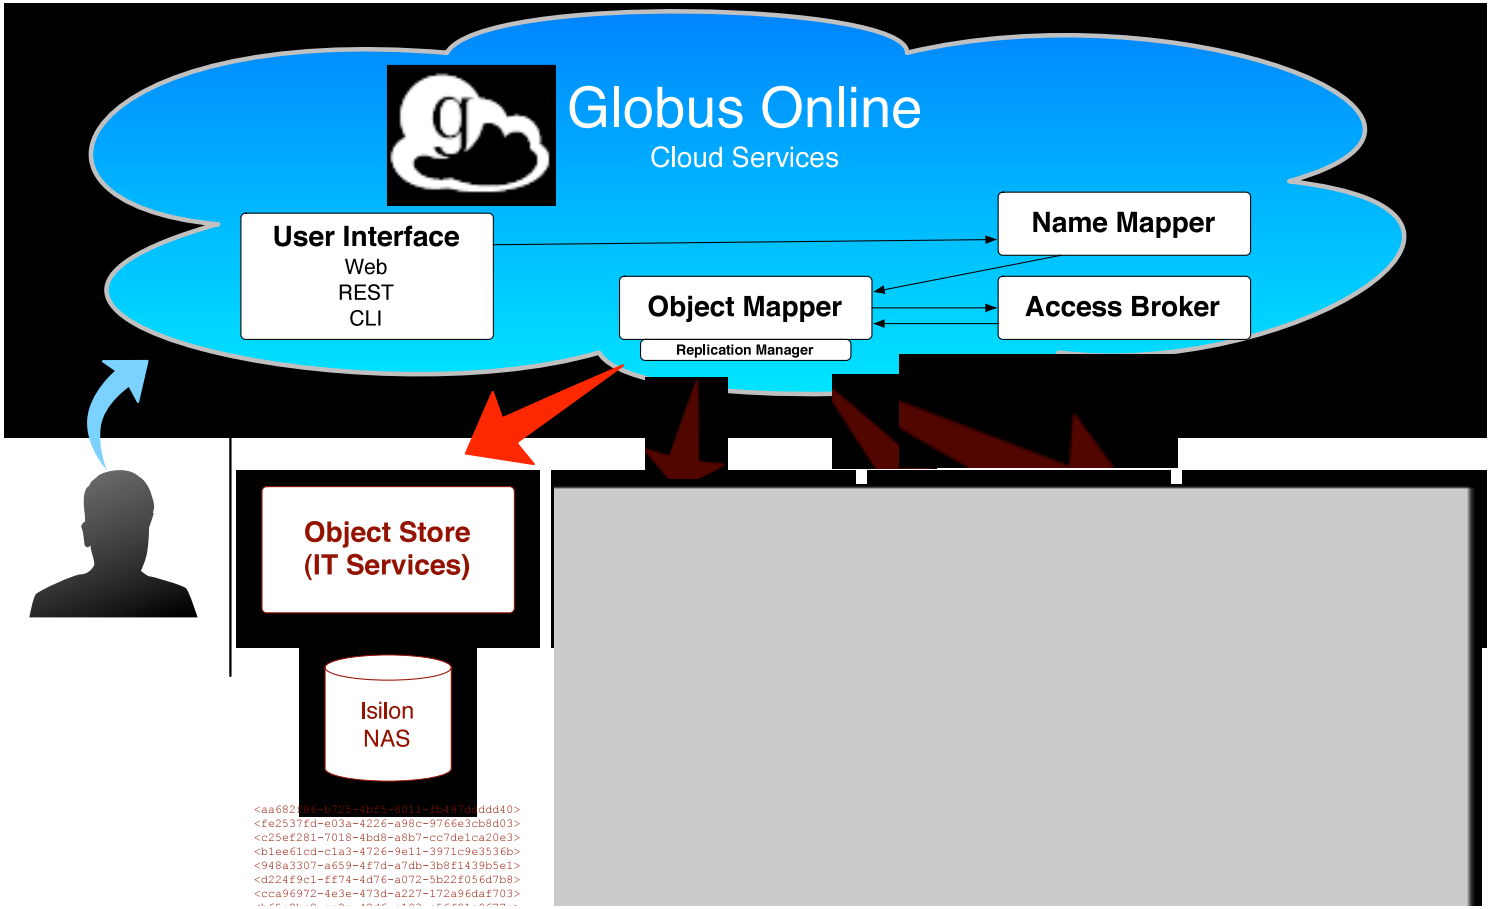

THE UNIVERSITY OF CHICAGO

ITServices

Globus Storage Provider Architecture

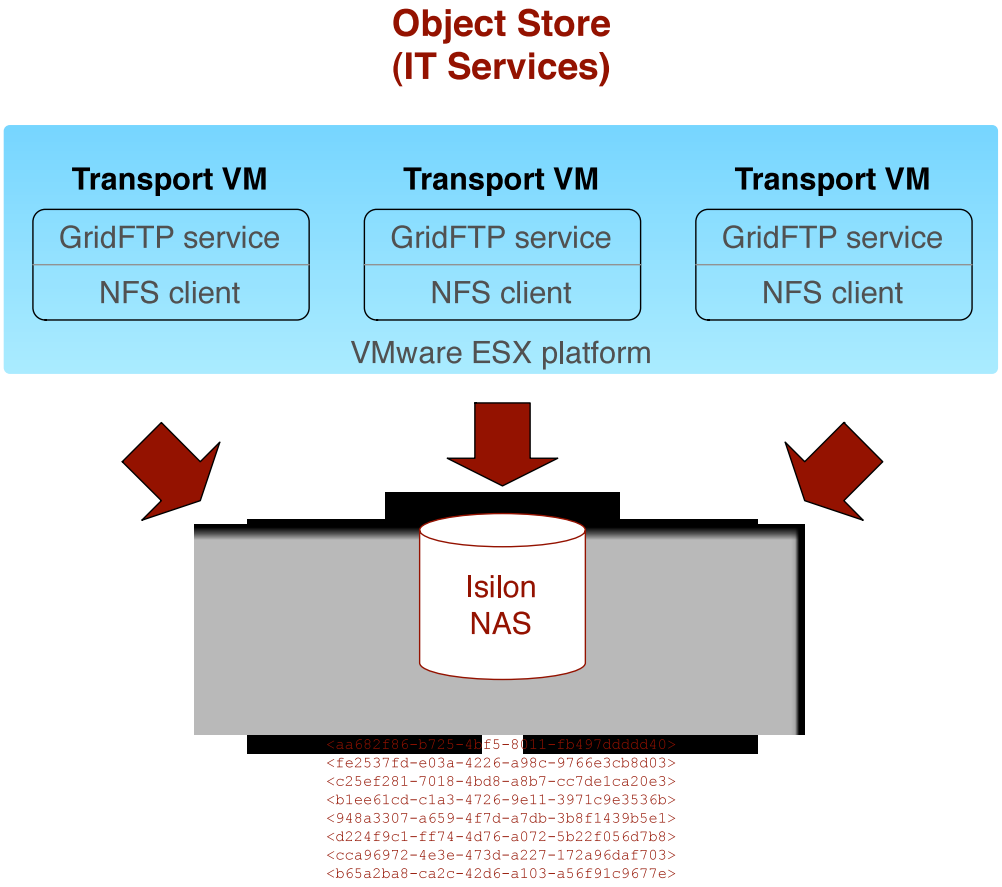

THE UNIVERSITY OF CHICAGO

ITServices

Globus Storage Provider Architecture

- Isilon Storage Array (EMC^2)
  - A simple, file level, scalable NFS service
  - Node-based clustering: higher throughput using multiple pipelines
  - 10 TB space for Globus Storage pilot, exported to Transport VMs
  - Total platform capacity easily expanded
  - Easy to extend allocation to the Globus Storage project as needs change

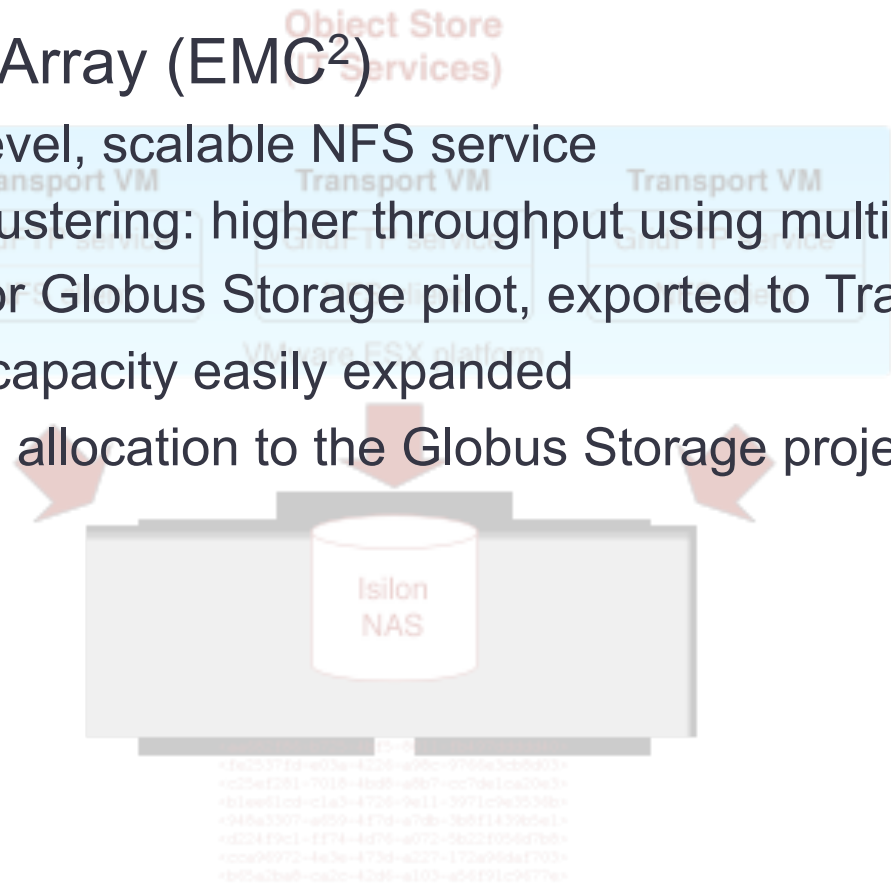

Globus Storage Provider Architecture

- Transport VMs
  - Simple Linux servers; no variation from enterprise profile
  - RHEL 5.7, because no special needs
  - Uses GridFTP for transport to Globus Online
  - Multiple instances to meet pipelining expectation of Isilon cluster due to node based architecture

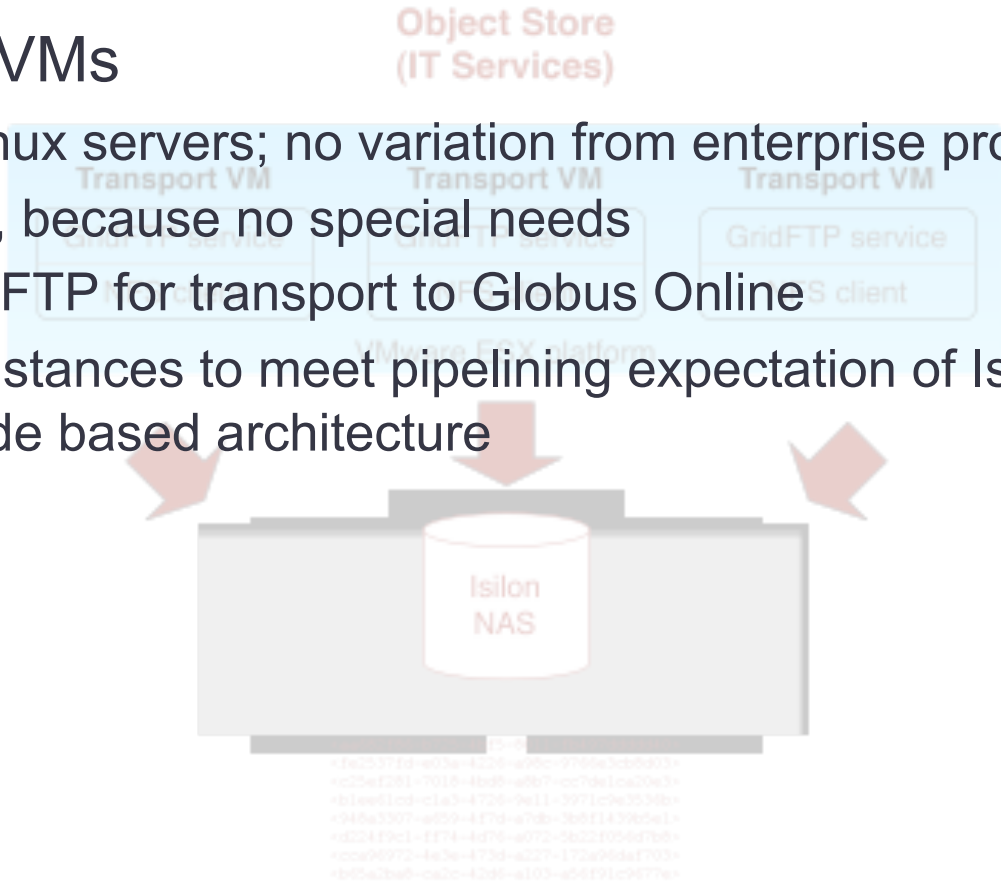

Globus Storage Provider Architecture

- Performance
  - Ideally, one transport VM per Isilon node
  - ESX hosts are Dell server blades
  - Blade chassis has 10G direct to data center core
  - Virtualization allows best match between Isilon node and I/O share in server chassis

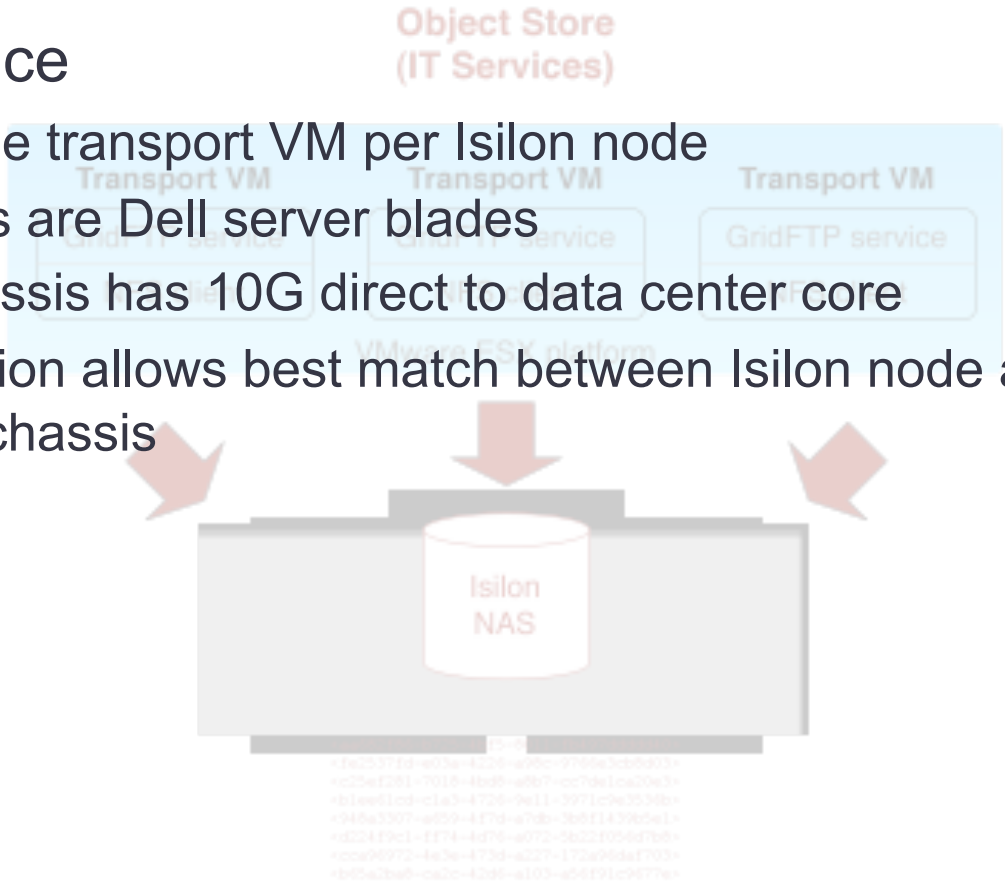

Globus Storage Provider Architecture

- Globus components from EPEL
  - Fedora project, Extra Packages for Enterprise Linux
  - Standard software distribution point for our environment; no extra requirements
  - Straightforward GridFTP installation with single local user mapped by `/etc/grid-security/grid-mapfile`
  - Easy to set up, but does require a trusted certificate
  - Now even easier: Globus Connect Multi-User (**GCMU**). `https://www.globusonline.org/forhpcowners/`

Caveats

Certain responsibilities are delegated to Globus Online:

- Only GO has visibility into per-user resource utilization
  - Campus Provider has aggregate view
  - Affects **chargeback** — easiest for GO to proxy billing
- Only GO has control of object storage
  - Individual objects have no metadata properties at provider end
  - Affects **access management** and **quotas/allocations**

Campus Provider Benefits

- Lowers Provider's cost for additional storage capacity
- Single relationship to manage
  - Provider works with Globus Online, not with users — even though they are our own users
- Single charge point
  - Globus Online pays Provider for aggregate use, rebills for individual use

Lessons Learned

- More testing needed
  - Sizing of transport servers is not tuned to workload
  - How many GridFTP servers are needed for optimum bandwidth?
  - How does resource consumption of VMs correspond to physical hosts under full workload?
  - Is virtualization the right approach?
  - Impact to enterprise workload: must scale these transport servers to have predictable maximum effect on the infrastructure as a whole

Future Plans

- Science DMZ placement
  - High throughput
  - Less restrictive access
  - Colocated with related applications, data
- Separate physical server infrastructure
  - Dedicated to science data flows
  - Easier management
  - Lower potential impact to enterprise
- Improved throughput to storage
  - Second Isilon array
  - Use Isilon replication for enhanced data protection

Future Plans

- Integration of Globus metadata with campus IDM
  - Automatic availability
  - Means of asserting metadata to Globus (allocation size, eligibility, groups)
- Closer proximity of Globus Storage software to array
  - Meeting with EMC^2 at SC11
  - Build Globus Storage/Globus Transfer target embedded within Isilon shell

Suggested Enhancements (Wishlist)

- Capacity and Performance Planning, Insight, Reporting
  - Users choose where data goes at Globus Online; no mechanism to check whether provider is supplying adequate capacity and per-workload performance
  - Local analysis tools can only discern aggregate utilization
  - ➢ Need a resource administrator interface to the metadata services behind Globus Storage

- Data Management
  - Legal restrictions, HIPAA, etc.
  - May have data that should not be backed up to our central systems
  - May have data that should not be stored extra-institutionally
  - Provider may have accountability to authorities for data provenance and residence
  - ➢ Need provider interface to classify data and flag it for specific policies

THE UNIVERSITY OF CHICAGO

IT Services

Try This at Home

- IT: fulfill service mission without talking to users
  - IT departments are good at providing core service
  - IT departments are less good at matching service to need
- Globus: compelling product but limited resources
  - Globus Online depends on external resource providers
  - Providers with existing commitment to mutual customers are cheaper to work with
  - By lowering cost to provider, provider is easier to work with
- Research: single entity managing data
  - No IT required

THE UNIVERSITY OF CHICAGO

ITServices

Questions?

THE UNIVERSITY OF CHICAGO

IT Services