

Globus and the Centralized Research Data Infrastructure at CU Boulder

Daniel Milroy, daniel.milroy@colorado.edu

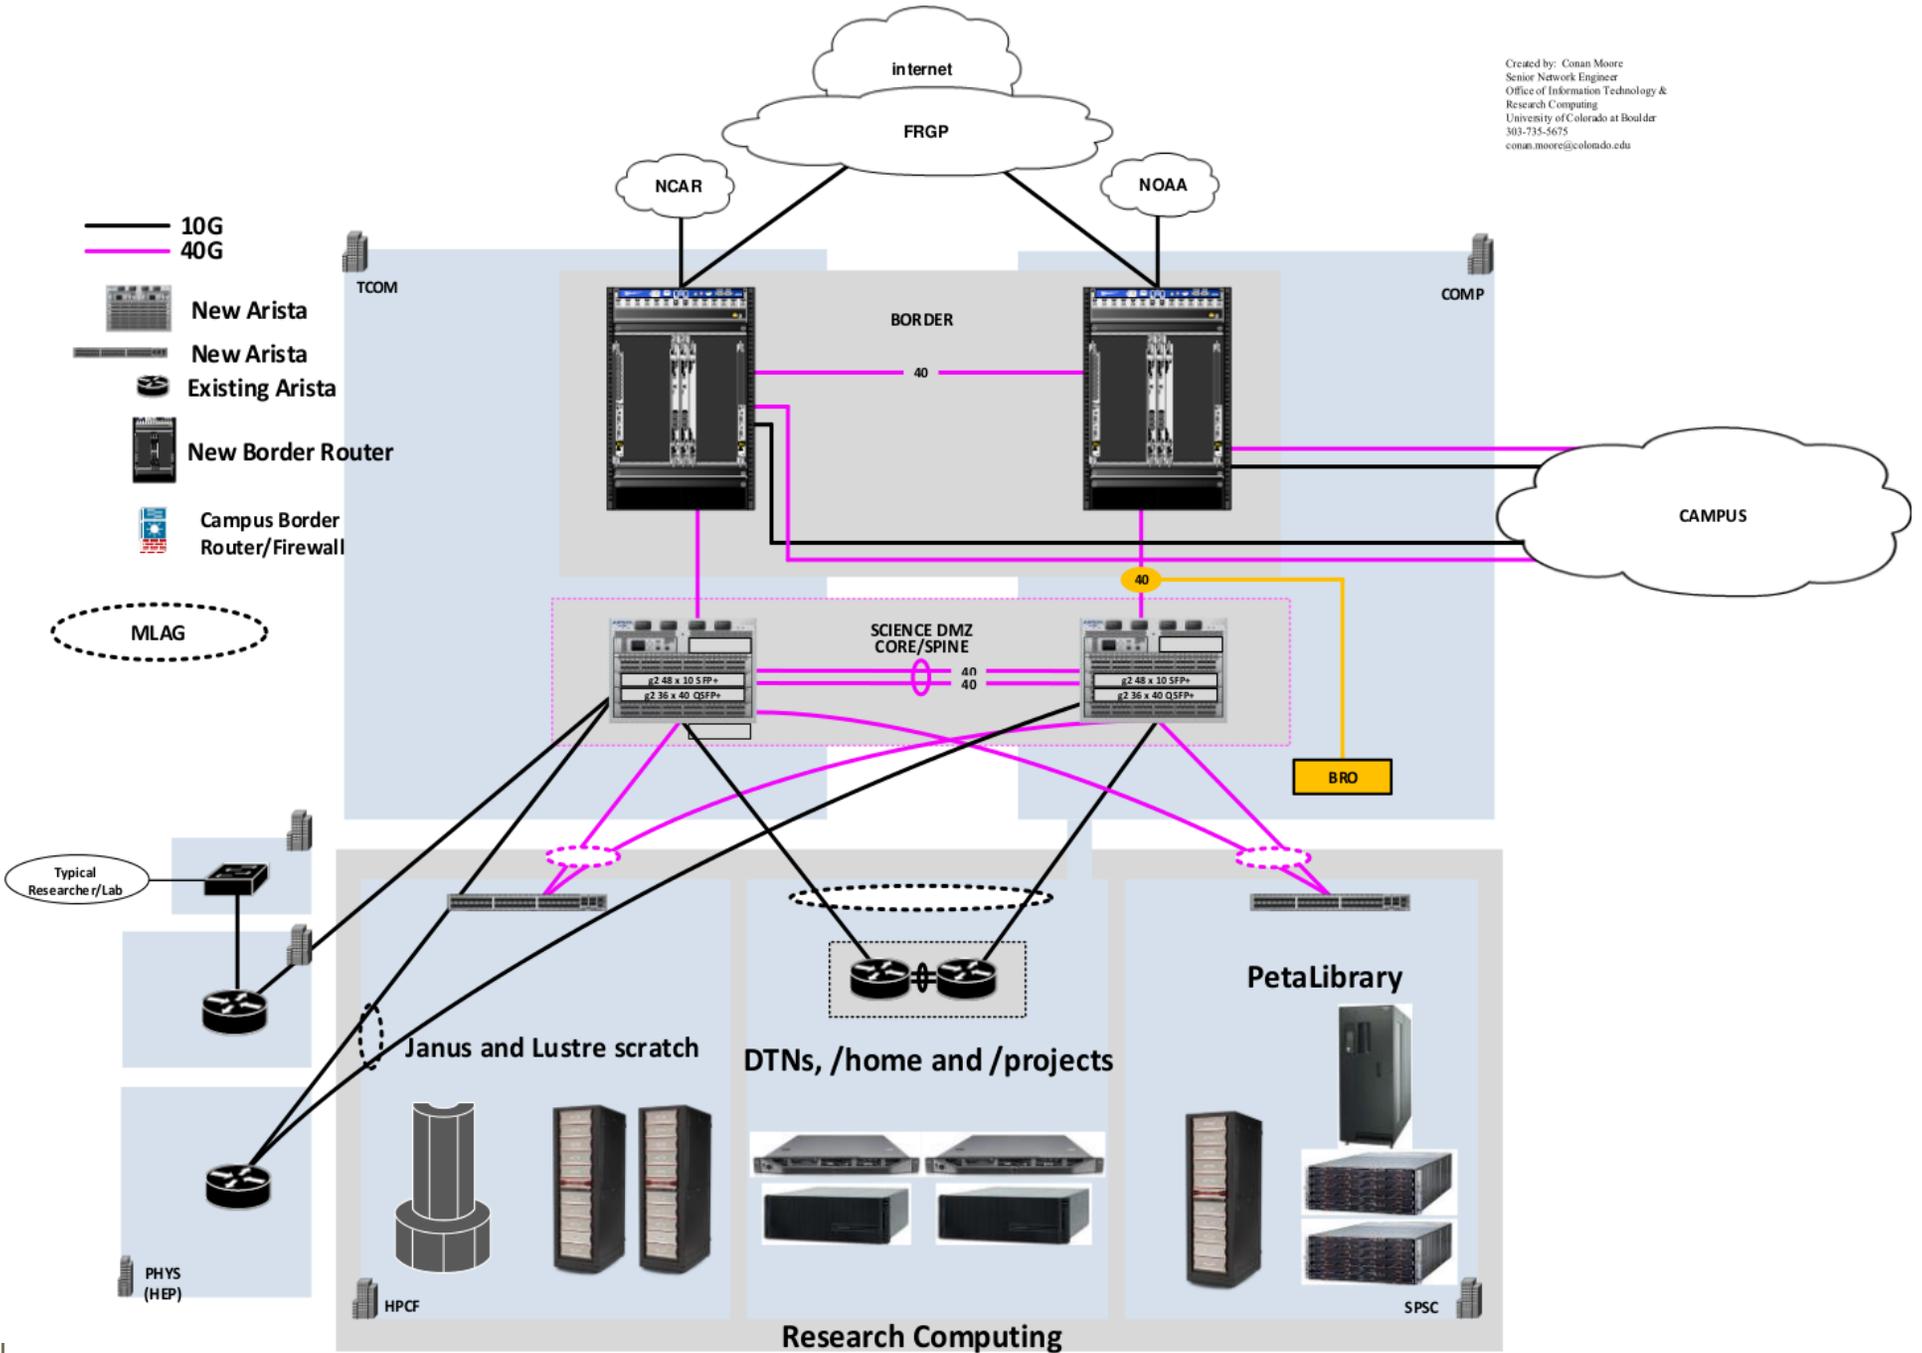
Conan Moore, conan.moore@colorado.edu

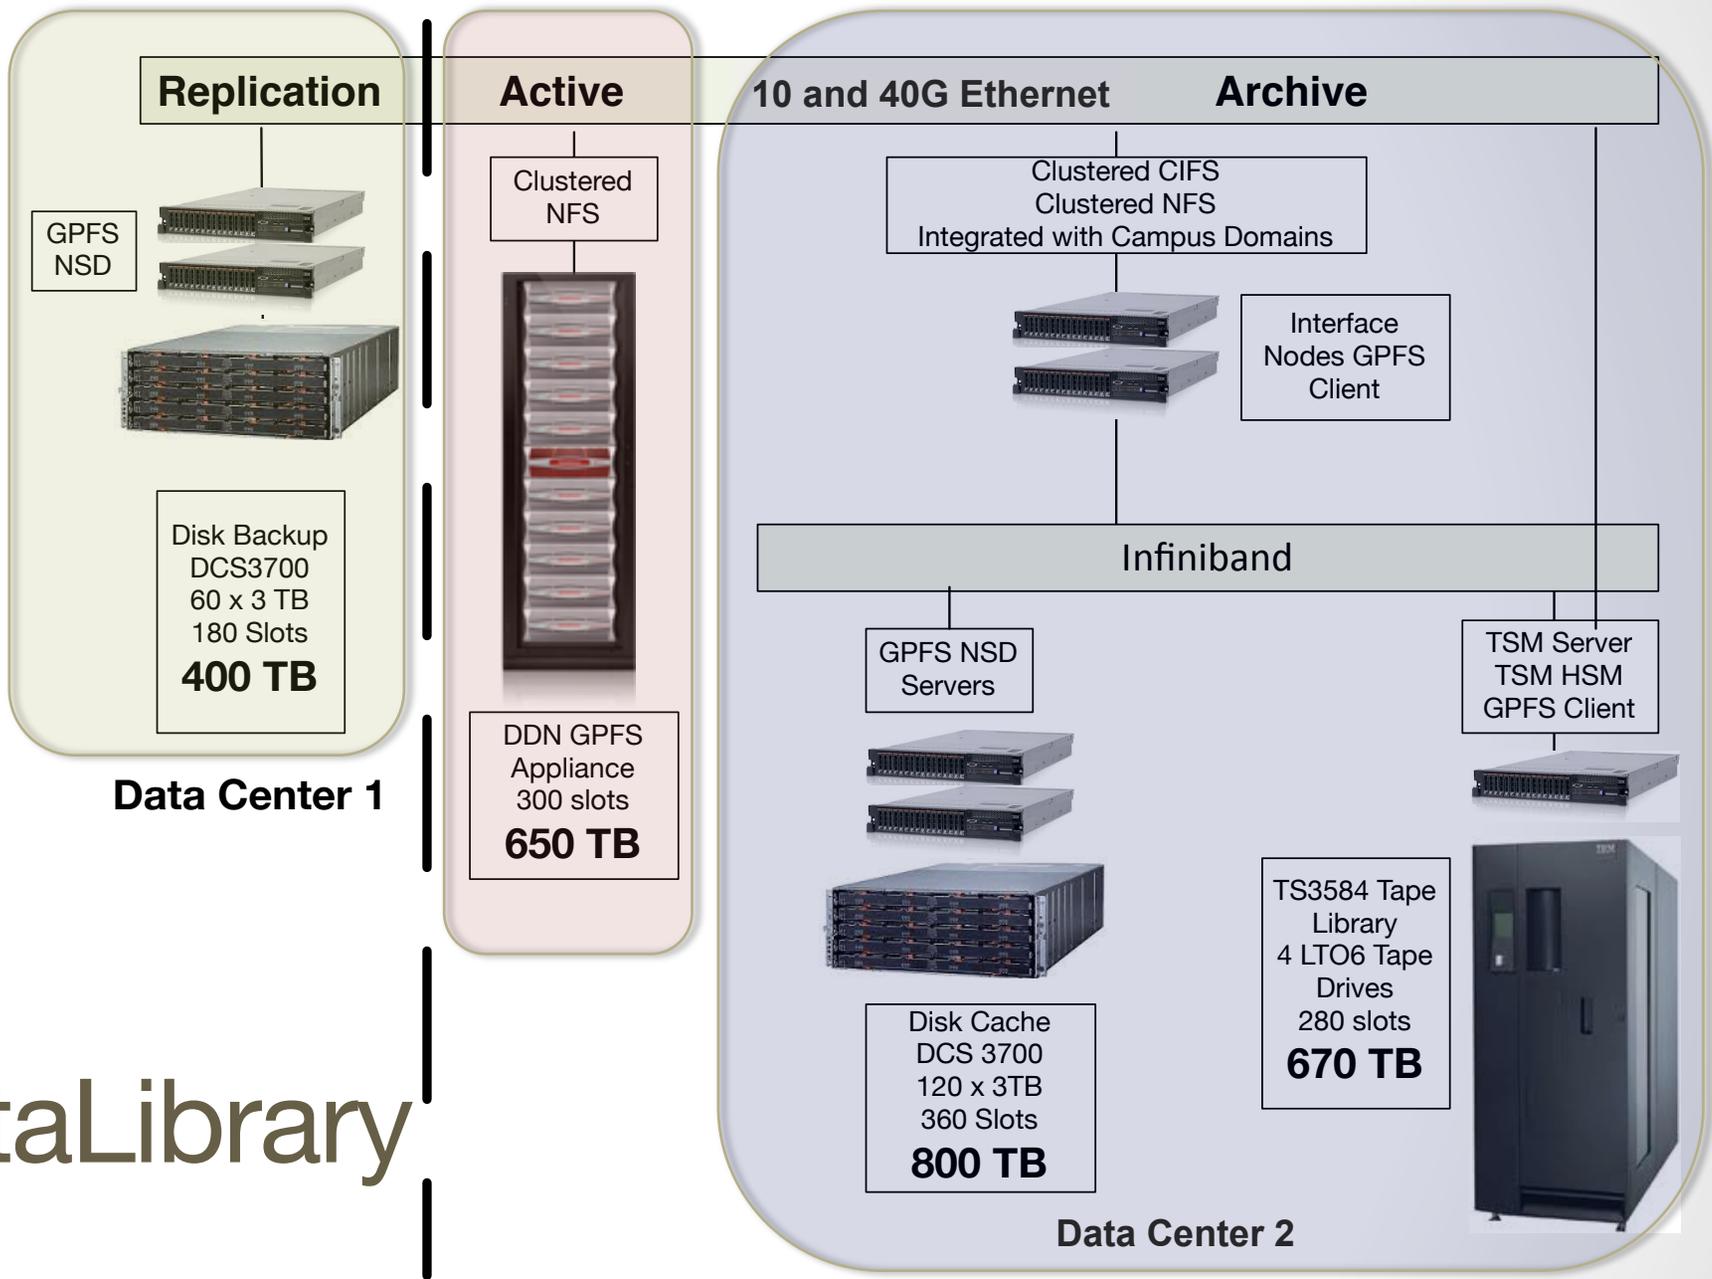
Thomas Hauser, thomas.hauser@colorado.edu

Peter Ruprecht, peter.ruprecht@colorado.edu

What CU researchers want

1. Cheap, robust storage
 - Snapshotted, replicated /home and /projects
 - PetaLibrary: paid storage options for active, archive and tape tiers under GPFS and HSM/TSM
2. Fast scratch storage for parallel IO
 - Petascale Lustre filesystem attached via IB to Janus: 1368 node compute cluster
3. Easy and rapid access to data
 - 80Gb core Science DMZ with redundant paths
 - Globus Online





PetaLibrary

Data Transfer Nodes

Access portal to Research Computing Storage Infrastructure

- Dell R320 servers
 - Single socket Xeon E5-1410 with 16GB memory
 - Bonded NICs for increased bandwidth and resiliency
 - Node 1: active-active (bonding mode 0) 10Gb links to switches MLAGed to Science DMZ core
 - Node 2: 40Gb uplink to the core as primary, dual 10Gb NICs secondary (bonding mode 1)
 - Multi-homed for VLANs
 - NFS automounts for /home, /projects, PetaLibrary
 - Lustre TCP client for DDN Exascaler

Data Transfer Nodes

Globus Connect Server configuration

- MyProxy CA running on one host (ID server)
 - Copy certs from CA server to `/var/lib/globus-connect-server/grid-security/certificates/` on other nodes
- Four instances: `colorado#gridftp` (public data interface), `colorado#jila`, `colorado#nsidc`, and `colorado#ibg`
 - Logical endpoints load-balance across physical hosts
 - To specify unique ports for control channel, and set `data_interface` edit, e.g.:
 - `/var/lib/globus-connect-server/gridftp.d/globus-connect-server-jila`
 - `/etc/gridftp-jila.conf`
 - Add symlinks in `/etc/gridftp-jila.d/`
 - Create entries in `/etc/init.d/` accordingly (PIDs, confs, etc...)

Data Transfer Nodes

Globus Connect Server configuration (contd.)

- Globus Sharing enabled for PetaLibrary paid storage
 - One complimentary Sharing account per contract
 - Server “perms” set to 0770 in /etc/gridftp.conf
 - Not process umask, file and directory creation permissions
 - Default is (0)644
 - Doesn't clobber setgid bit

Globus Use Cases

What problems does it solve?

- Characteristic examples: speed, automatic restart and performance optimization, fire-and-forget, ease of use for non-traditional HPC users
 - InCommon login to Globus
- RC employs One Time Password authentication
 - Globus' credential lifetimes facilitate long transfers
 - Potential to automate transfers upon compute job completion, not possible without `cli.globusonline.org` access, or interface to Globus' REST API
- Collaboration via Globus Sharing
 - Killer app for PetaLibrary customers

Globus Use Cases

- Data transfer within RC storage infrastructure
- Libraries
 - PetaLibrary is primary storage
 - Peace and Justice, Labor, Western Americana, etc...
 - Globus Sharing with ACLs for unified access control and user management
- Museum of Natural History
 - Large digitization effort of existing collections: Botany, Entomology, Paleo Vertebrates
 - Links to geographic, genomics data
- Ongoing Development
 - Neuroscience research group testing Globus Sharing for international collaboration
 - Will utilize multiple Sharing Groups with different permissions and compartmentalized access
 - Execute transfer on job completion

Globus Statistics

- External transfers from Nov 2013 to March 2014
 - 1.4 Gb/s to NCAR
 - 1.3 Gb/s to Stampede
- Transfers from March 2013 to March 2014
 - 58 TB in, 112 TB out
 - 335 M files
- Greater than 10 checksum verification failures in April 2014
 - TCP with 16 bit checksum translates to approximately 1 in 300 M packets silently corrupted (Vern Paxson, 1999)

Thank you!

Questions?