

Taiga & Granite | The NCSA Research Data Platform

J.D. Maloney | Lead Storage Engineer | NCSA



**National Center for
Supercomputing Applications**

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Prepared for: LoC DSA Conference 2024

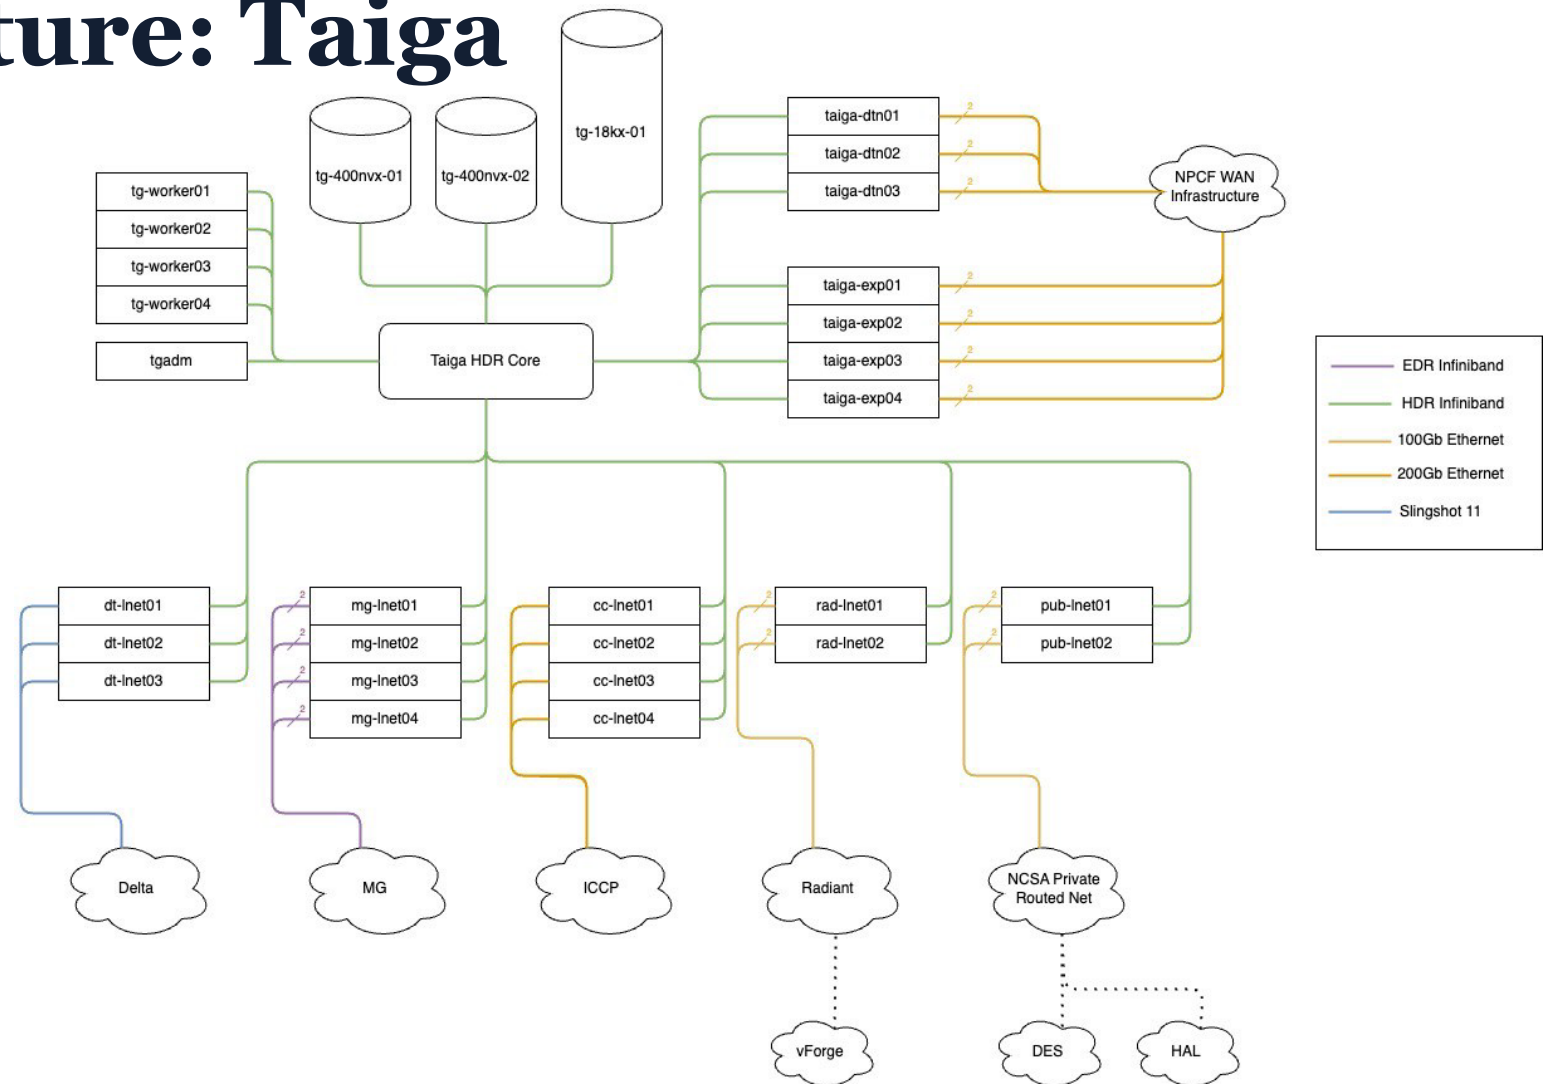
Design Goals: Taiga

- Easy to scale (both in capacity/performance) file system
 - Can add to the FS in reasonable chunks
- Hybrid NVME and HDD
 - More than just metadata on flash
- FS capable of peering with every compute environment natively on its HSN
- User/Group/Project Quotas
- Support for multiple Authentication Systems



System Architecture: Taiga

- Core network fabric is HDR Infiniband (200Gbps)
- File system is translated to every compute environment's high-speed fabric
- A set of worker nodes perform data packaging/migration work
- Each appliance contains 100TB usable NVME for holding small files (capturing ~70% of files)



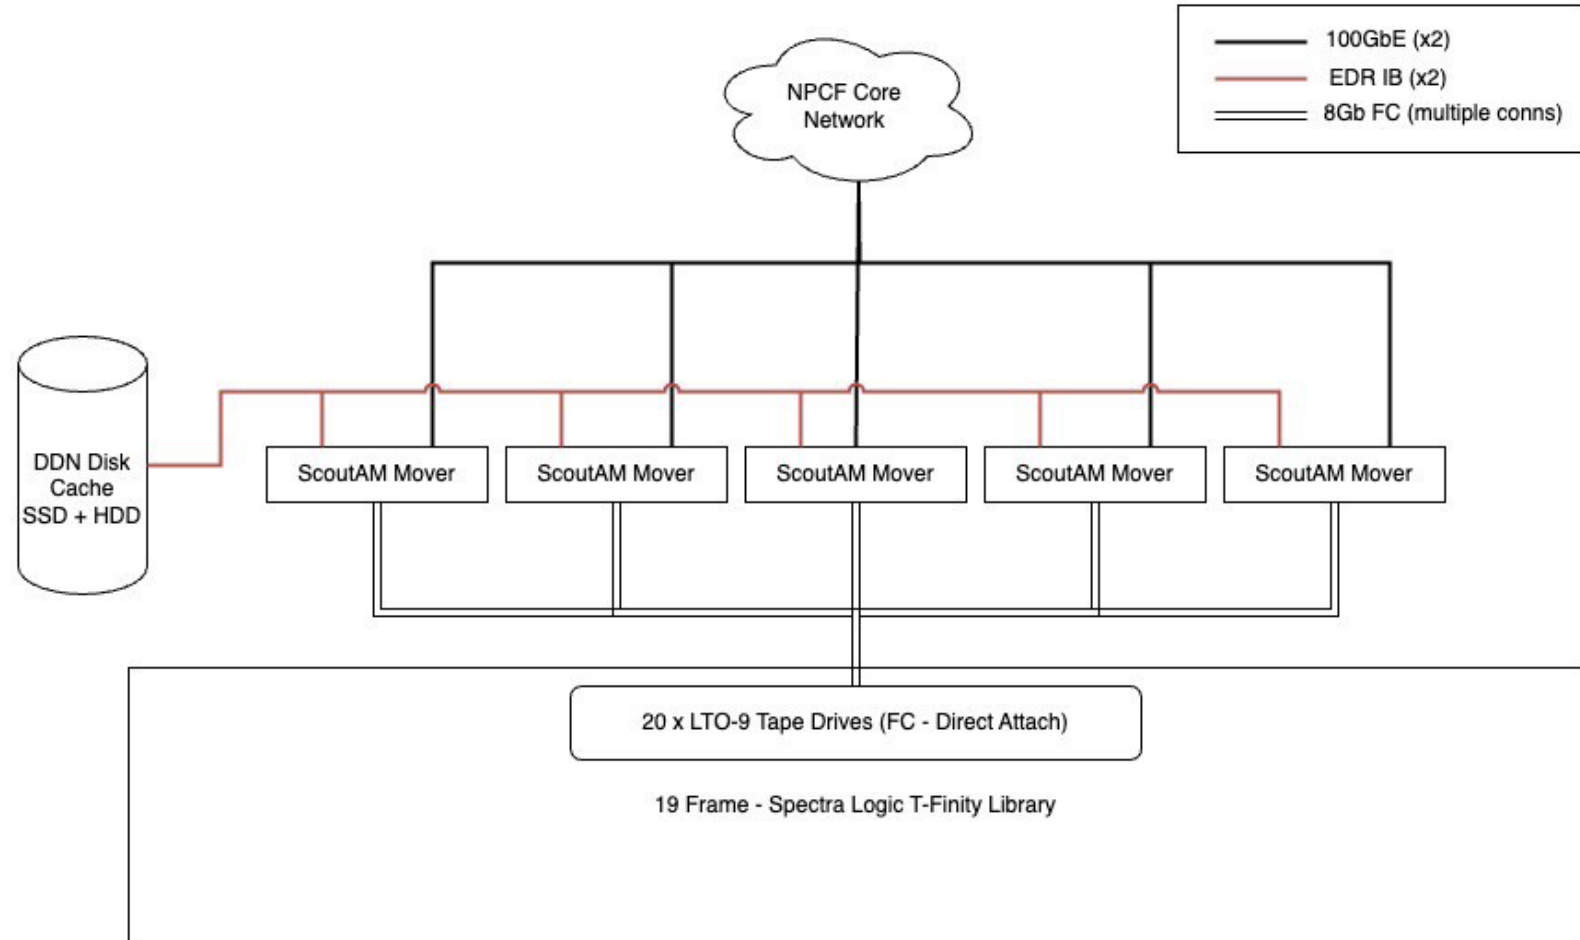
Design Goals: Granite

- Support for a large-scale (up to an Exabyte) archive across many thousands of tapes
- Fine grained ability to route specific datasets on to specific tape pools
- Scalable performance by adding mover nodes and tape drives (up to 60+ drives)
- Allows for cost effective design
- Ability to access data via POSIX (NFS) and via S3



System Architecture: Granite

- Tape Drives and RIMs direct connect to ScoutAM movers
- Disk Cache max performance of ~20GB/s
 - 14 x SSD for metadata
 - 600 x HDD for data
- Uses IB-SRP to present disk-cache targets to ScoutAM movers
- Dual 100GbE per ScoutAM mover node for client-facing data movement



Capabilities Unlocked

- Researchers can compute on their datasets across a variety of batch, high-throughput, and VM/cloud computing resources simultaneously
 - Greatly reduced cross-system data transfer needs
- Supports the full data life cycle from storing data directly after capture to the preservation of the data on tape
- Provides users with multiple ways to interact with their datasets (Globus, Lustre, NFS, SMB, S3, etc.)
- Both are cost recovered services, rates are available to use in grant proposals



The Future of Taiga

- System upgrade slated for 2026, seriously exploring a move to an all NVME platform
 - QLC media + 400/800Gb Ethernet
- Much improved policy engine and data manipulation tools and capabilities (both for admins & users)
- Tenant-based performance QoS capability
- Even better tooling for moving data between active and archive systems
- Automation of dataset sharing/access (FAIR)
- Partial hydration in AWS/Azure/GCP



The Future of Granite

- Move to a smaller, but all-NVME disk cache
 - Better performance for data movement
 - Movers are refreshed here as well to new hardware
- Finish upgrade to LTO-9 & Library software upgrade
- Rolling out more robust S3 access capability
- Better tooling/integration with primary tier
- Partnering with Illinois Data Bank for hosting large curated datasets
- Opening system to ACCESS users; 20% capacity for national datasets





Questions?

Email: malone12@illinois.edu



**National Center for
Supercomputing Applications**

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN