



# Site in a Box

## Improving the Tier 3 Experience

University of California as a prototype

across [UC Davis](#), [UC Santa Cruz](#), [UC Santa Barbara](#), [UC Irvine](#), [UC Riverside](#), [UC San Diego](#)

Jeff Dost (UCSD)



# Tier 3 Challenges

## System Administration

- Software maintenance
  - WN software stack
  - CA Certs / CRLs
  - Security patches
  - Experiment software maintenance
- Network debugging
- Performance Tuning
- Non-dedicated staff (Campus IT not HEP specialized)

## Workflow Administration

- Transporting data in / out of site
- Submitting to external resources
  - Q: Don't we have WMS (CRAB / PANDA) for that?
  - A: Not all workflows are WMS appropriate
  - Want to submit seamlessly between local resources and external (common interface)
- Letting external users in

For more details, see [Lightweight sites \(#165\)](#), M. Litmaath

# Solution

- Provide a physical host to the Tier 3 that is centrally maintained by dedicated Tier 2 staff
  - Software is installed, configured, and maintained remotely from a central Puppet instance
- Host provides necessary software to handle data transfers and send users out to the grid
  - Should integrate seamlessly with Tier 3 cluster
- No LHC or Grid knowledge required from local IT

# Hardware Specifications

(aka the “brick”)



## Hardware:

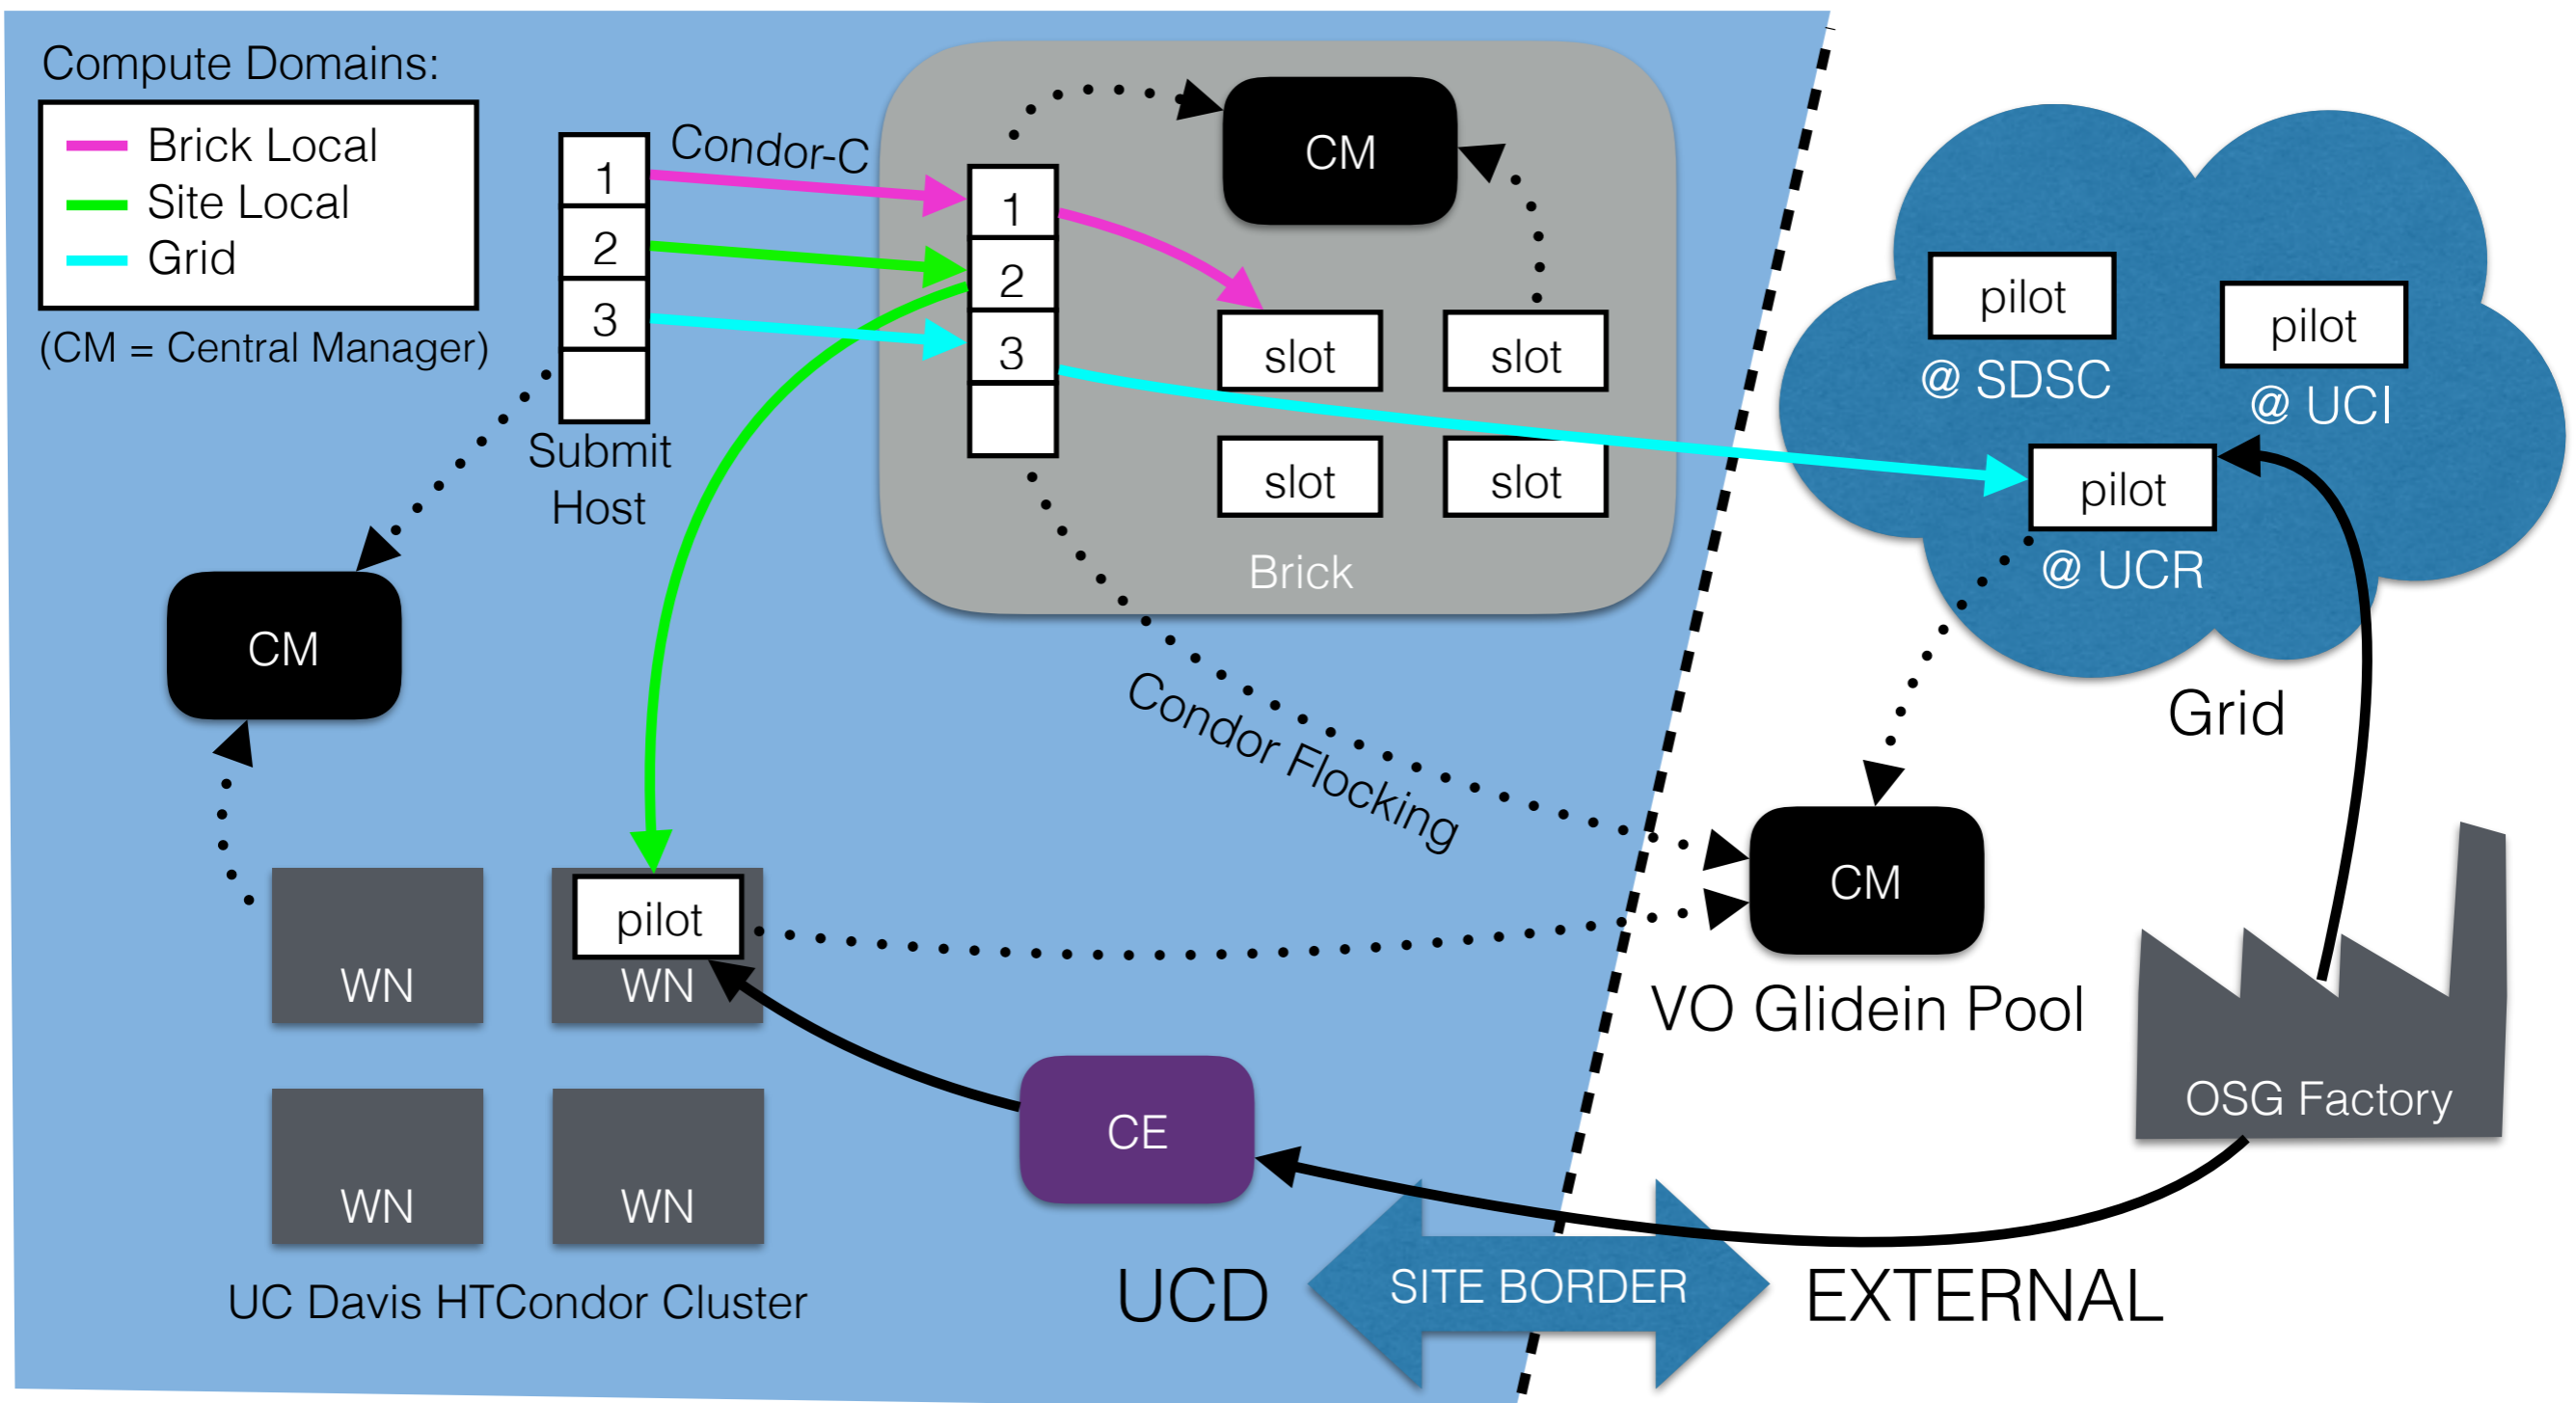
- 40 cores
- 12 x 4TB data disks (RAID6, XFS)
- 128 GB ram
- 2 x 10 gbit network interface

## Software:

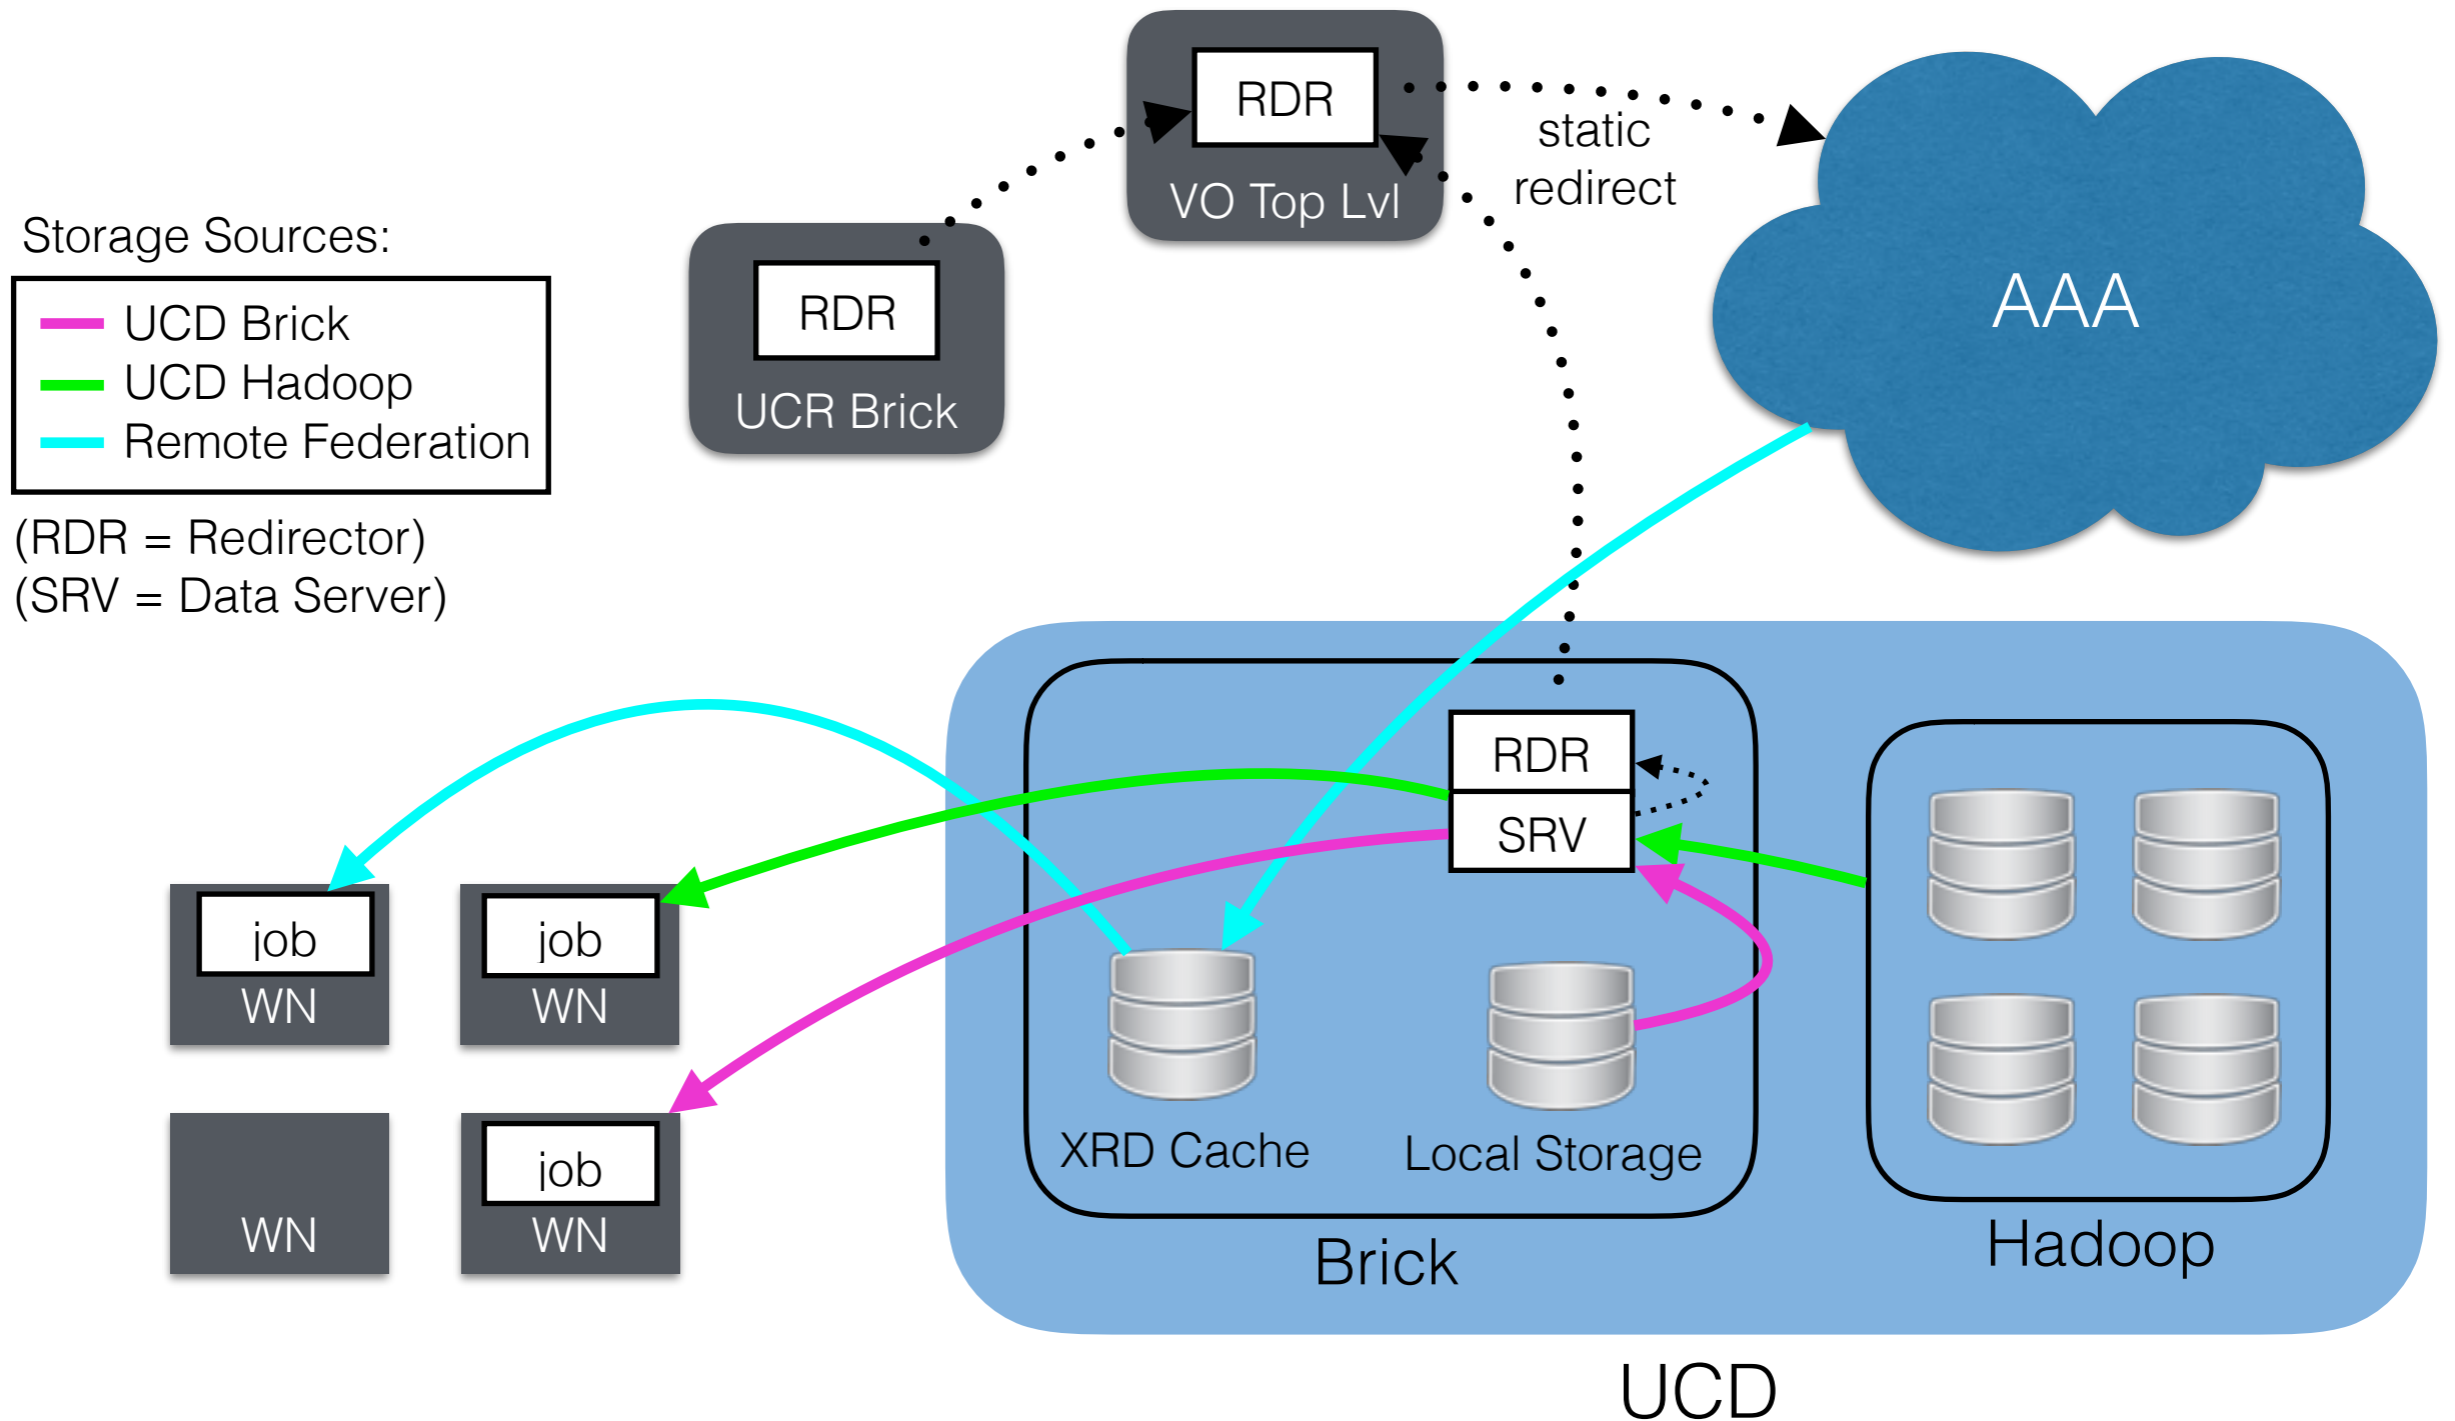
- Full HTCondor pool
- XRootD server, redirector, and proxy cache
- cvmfs w/ optional Squid

**The brick is effectively a site in a box for \$10k**

# Job Submission (UC Davis)



# XRootD Data Access





# Case Study: LHC @ UC

- Participating UCs:
  - ATLAS
  - CMS
- Science DMZs interconnected over 100gbps Pacific Research Platform network
- Resources currently provided:
  - Each UC
  - 50k core Comet cluster at SDSC
- Eventually:
  - Any other non-UC participating PRP site
  - Any OSG site beyond PRP
  - Other NSF XSEDE and DOE super computing centers
  - Cloud resources



# XRD Cache Scale Test

- 100 files ~1TB total stored at UCSC
- 10 clients running on UCSD brick fetch files at random simultaneously
- As cache fills, inbound network IO decreases and files begin to be read from disk cache
- Caching IO maxes at 5 Gbps (~640 MBytes/sec) during simultaneous disk reads and writes






# Conclusion

- The **site in a box** enhances Tier 3 sites by providing:
  - A unified way to submit locally and compute globally
  - The ability to decouple data placement from where the jobs run
- The central management of the services by dedicated Tier 2 admins allows the local Tier 3 users to worry less about infrastructure maintenance and focus more on getting science done

# XRD Scale Test No Cache



- Normalization scale test to show we can saturate at 10 Gbps (~1.3 GBytes/sec) with 10 parallel xrdcps copying straight to /dev/null and cache disabled

# Cross Site Data Access

