NCAR Site Update

Spectrum Scale Users Group - US

16 Apr 19

NCAR

NSF

Big data is like teenage sex:

Everyone talks about it,

Nobody really knows how to do it,

Everyone thinks everyone else is

doing it, so everyone claims they are doing it…

- Dan Ariely

Storage Environment

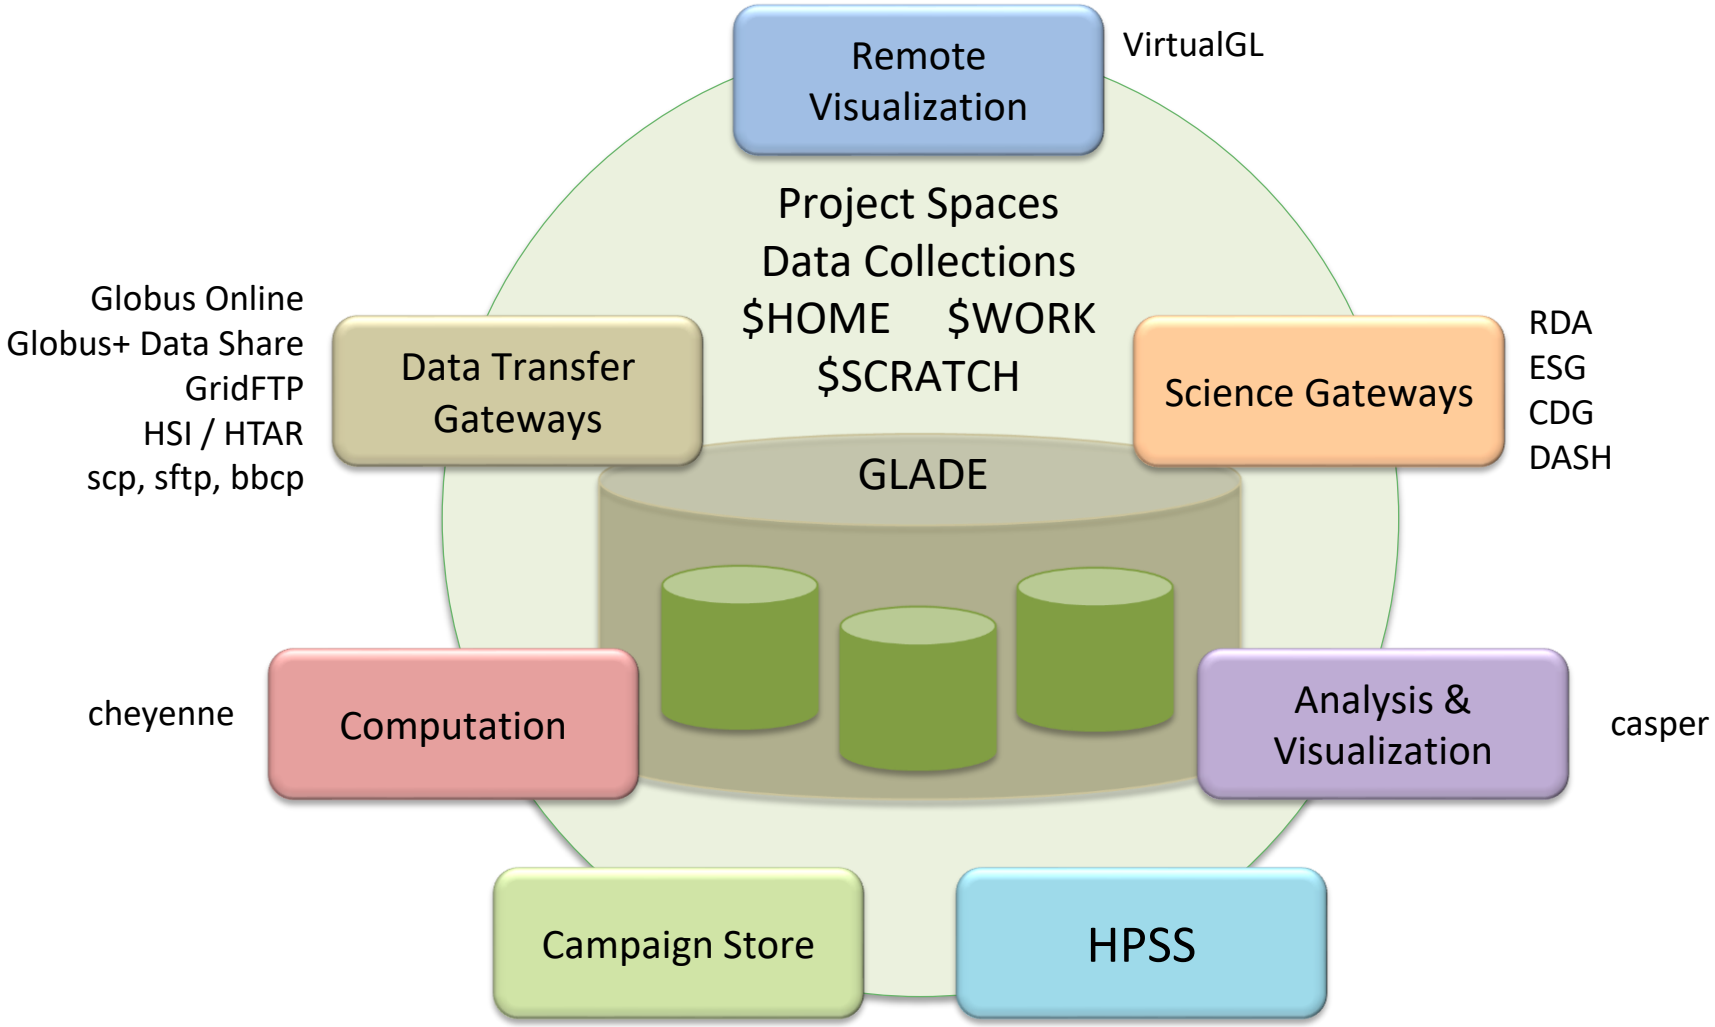

Storage Layer Cake

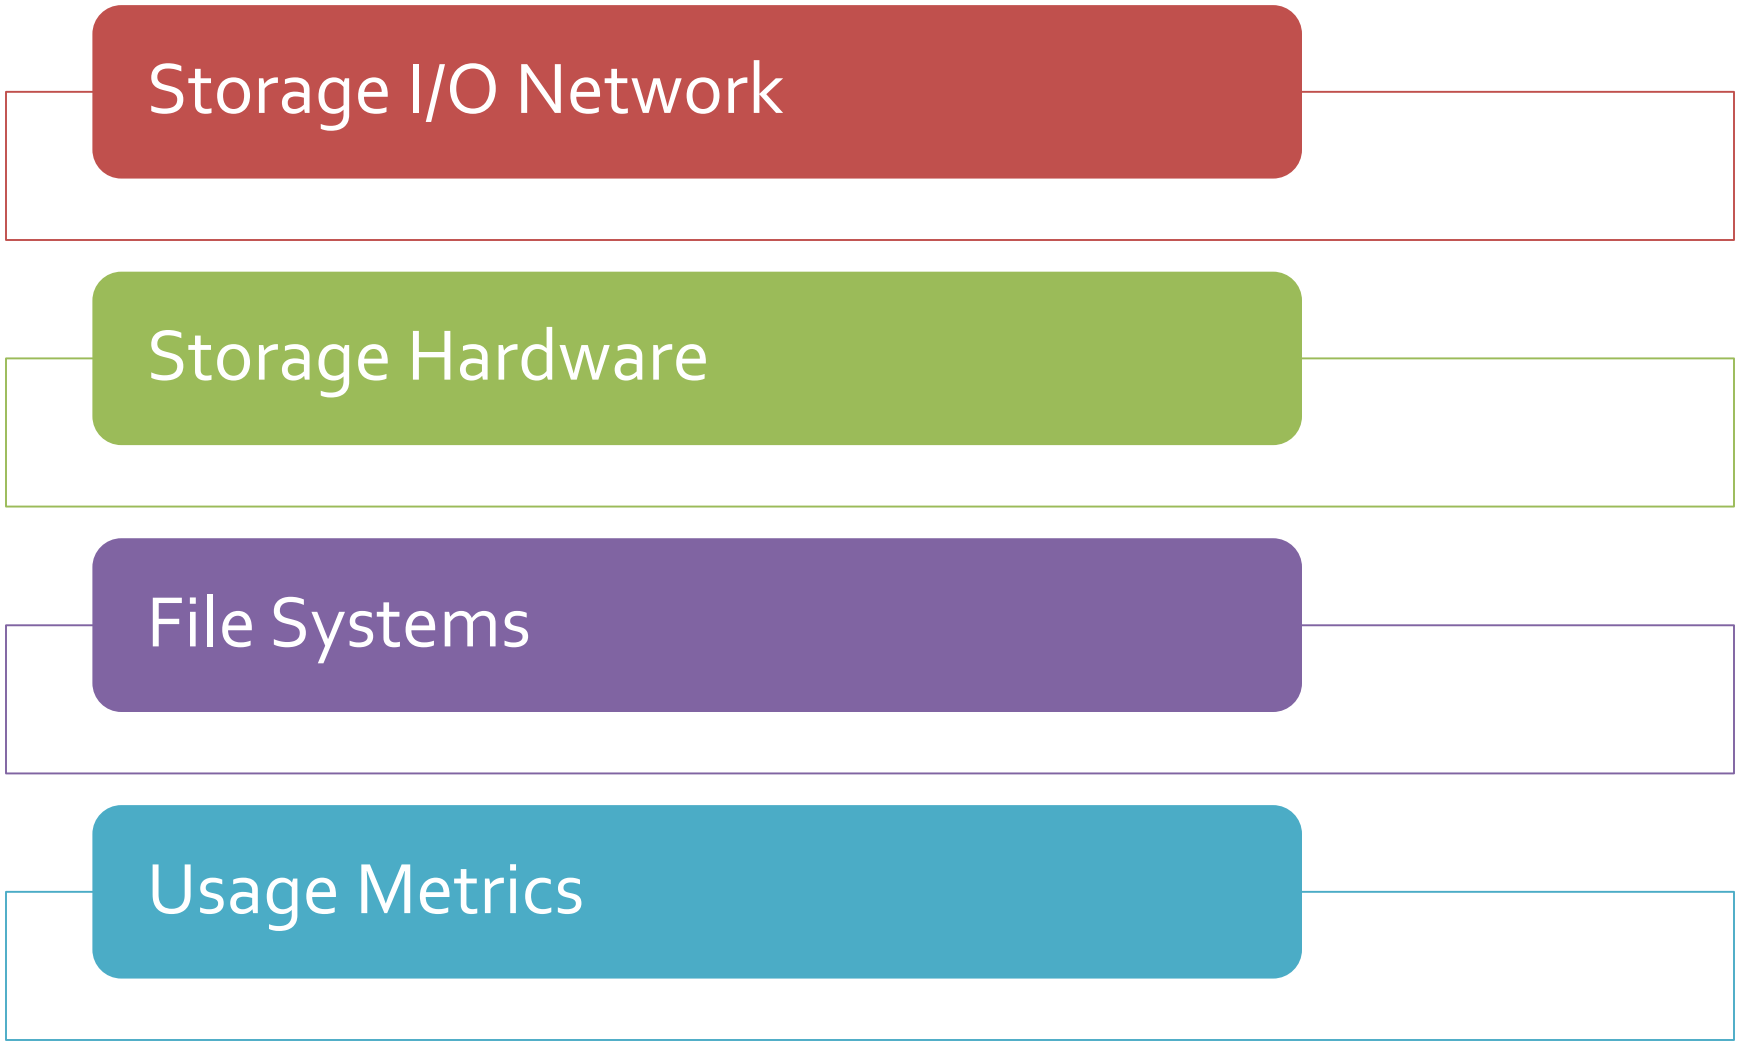

Storage I/O Network

- Network architecture providing global access to data storage from multiple HPC and Service resources
- Flexibility provided by support of multiple connectivity options and multiple compute network topologies
  - 10GBE, 40GbE, 100GbE, EDR
  - Full Fat Tree, Hypercube
- Scalability allows for addition of new HPC or storage resources
- Agnostic with respect to vendor and file system
- Can support multiple solutions simultaneously

The network has become the most critical component of the environment

Storage Connectivity Options

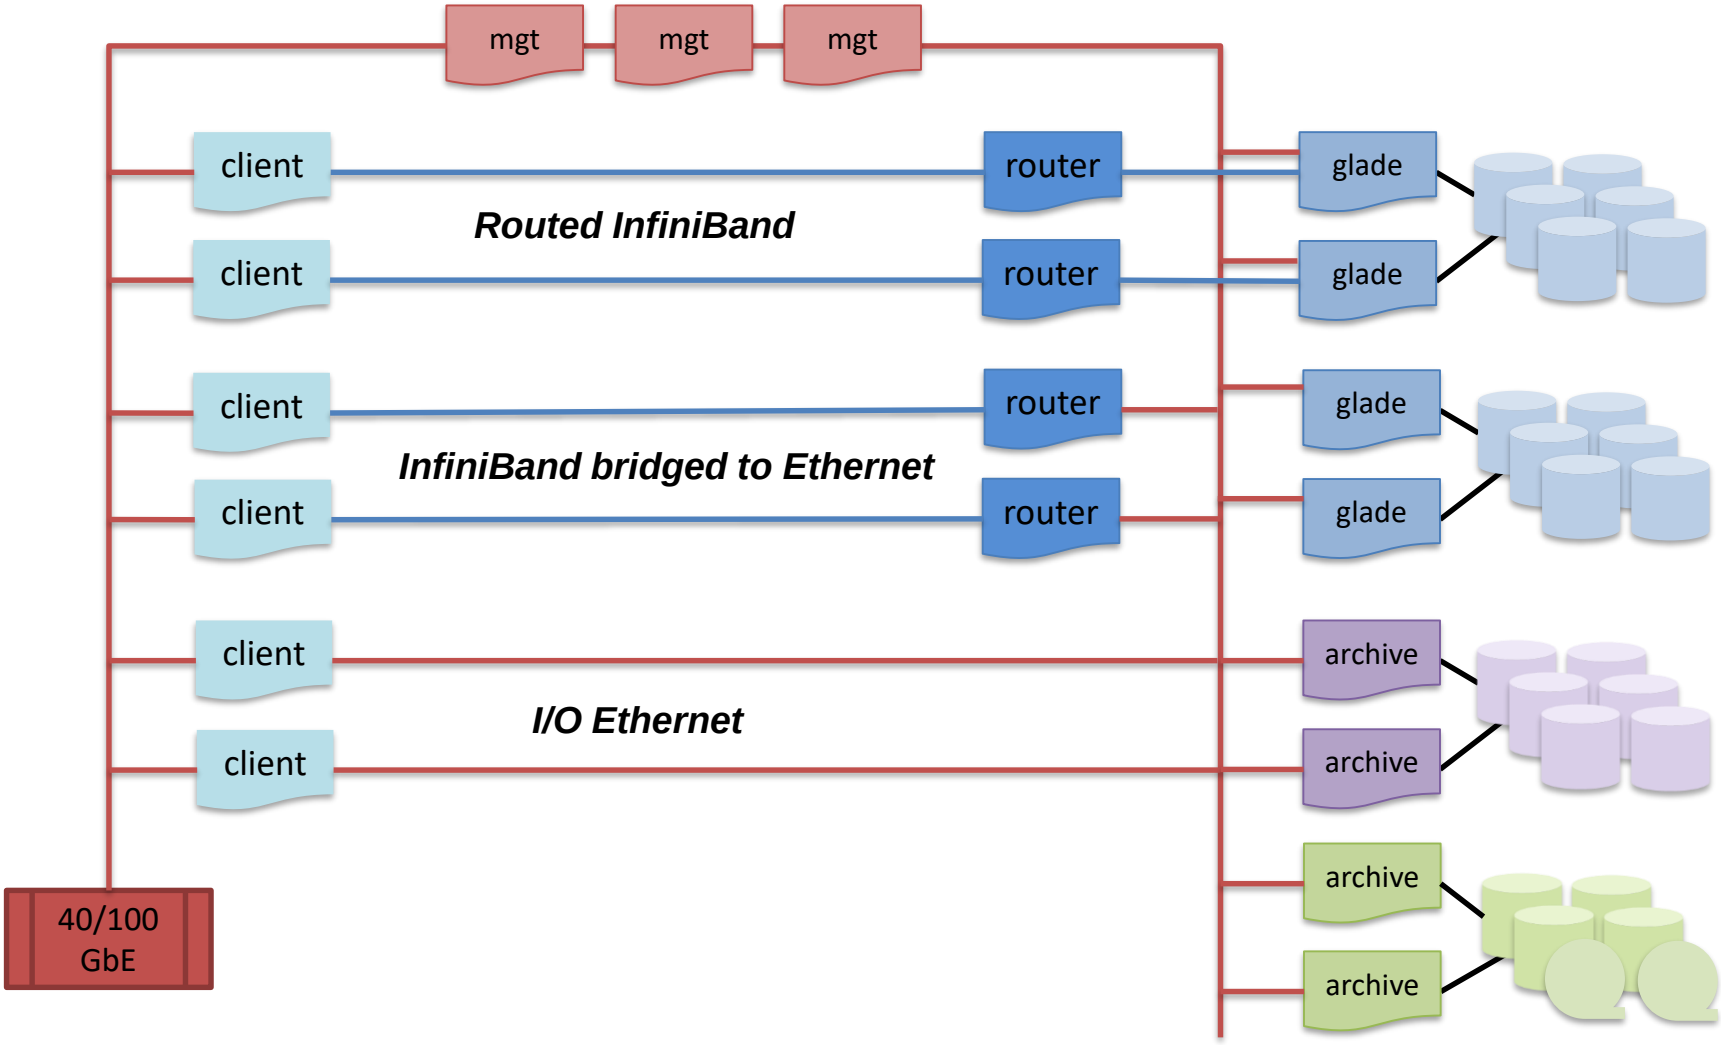

GLADE Routing

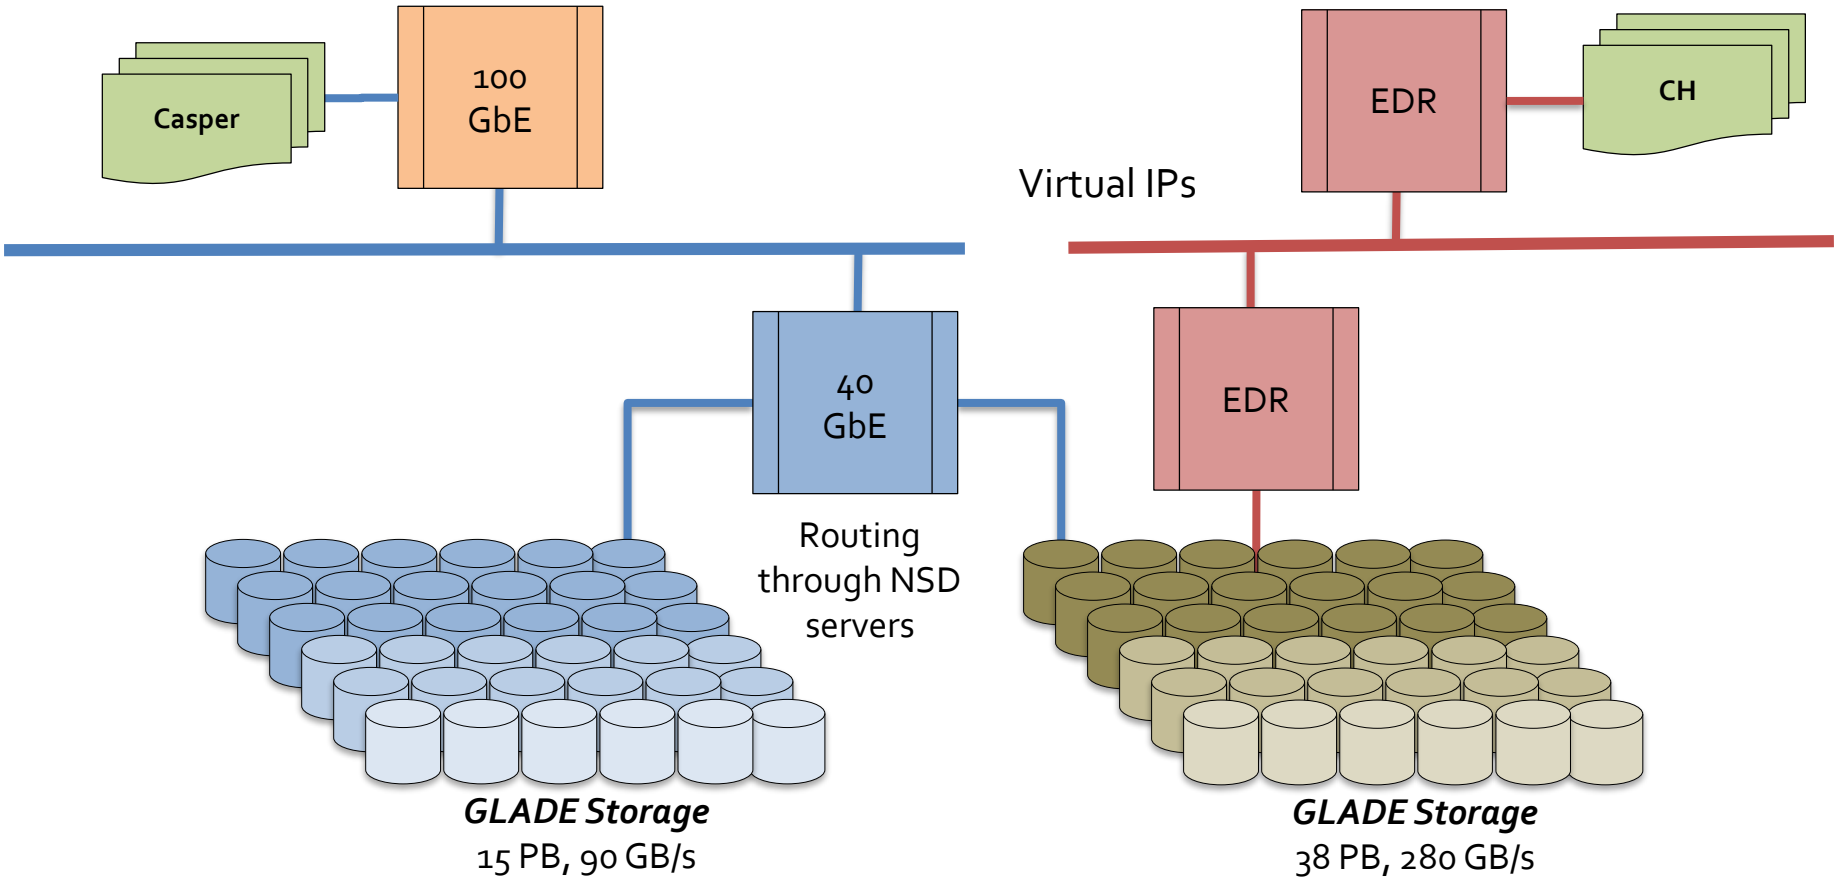

NCAR Storage Cluster

Storage Cluster

- 6 management servers
- 40/100 GbE main network

GLADE

- 300 GB/s bandwidth
- 38 PB useable capacity
- 8 DDN 14KXE (10 drawer)
- 6560 8TB drives
  - data only
- 96 800GB SSD drives
  - metadata
- 168 3.52 Flash drives
  - mixed data + metadata

Campaign Store

- 76 GB/s
- 26 PB useable capacity
- 2 DDN 14KXE (20 drawer)
- 3600 10TB drives
  - mixed data + metadata
- 8 GPFS NSD servers
- 40 GbE

HPSS Archive

- 10GB/s
- 2 PB front-end cache
- 48 T10KD drives

Storage Architecture

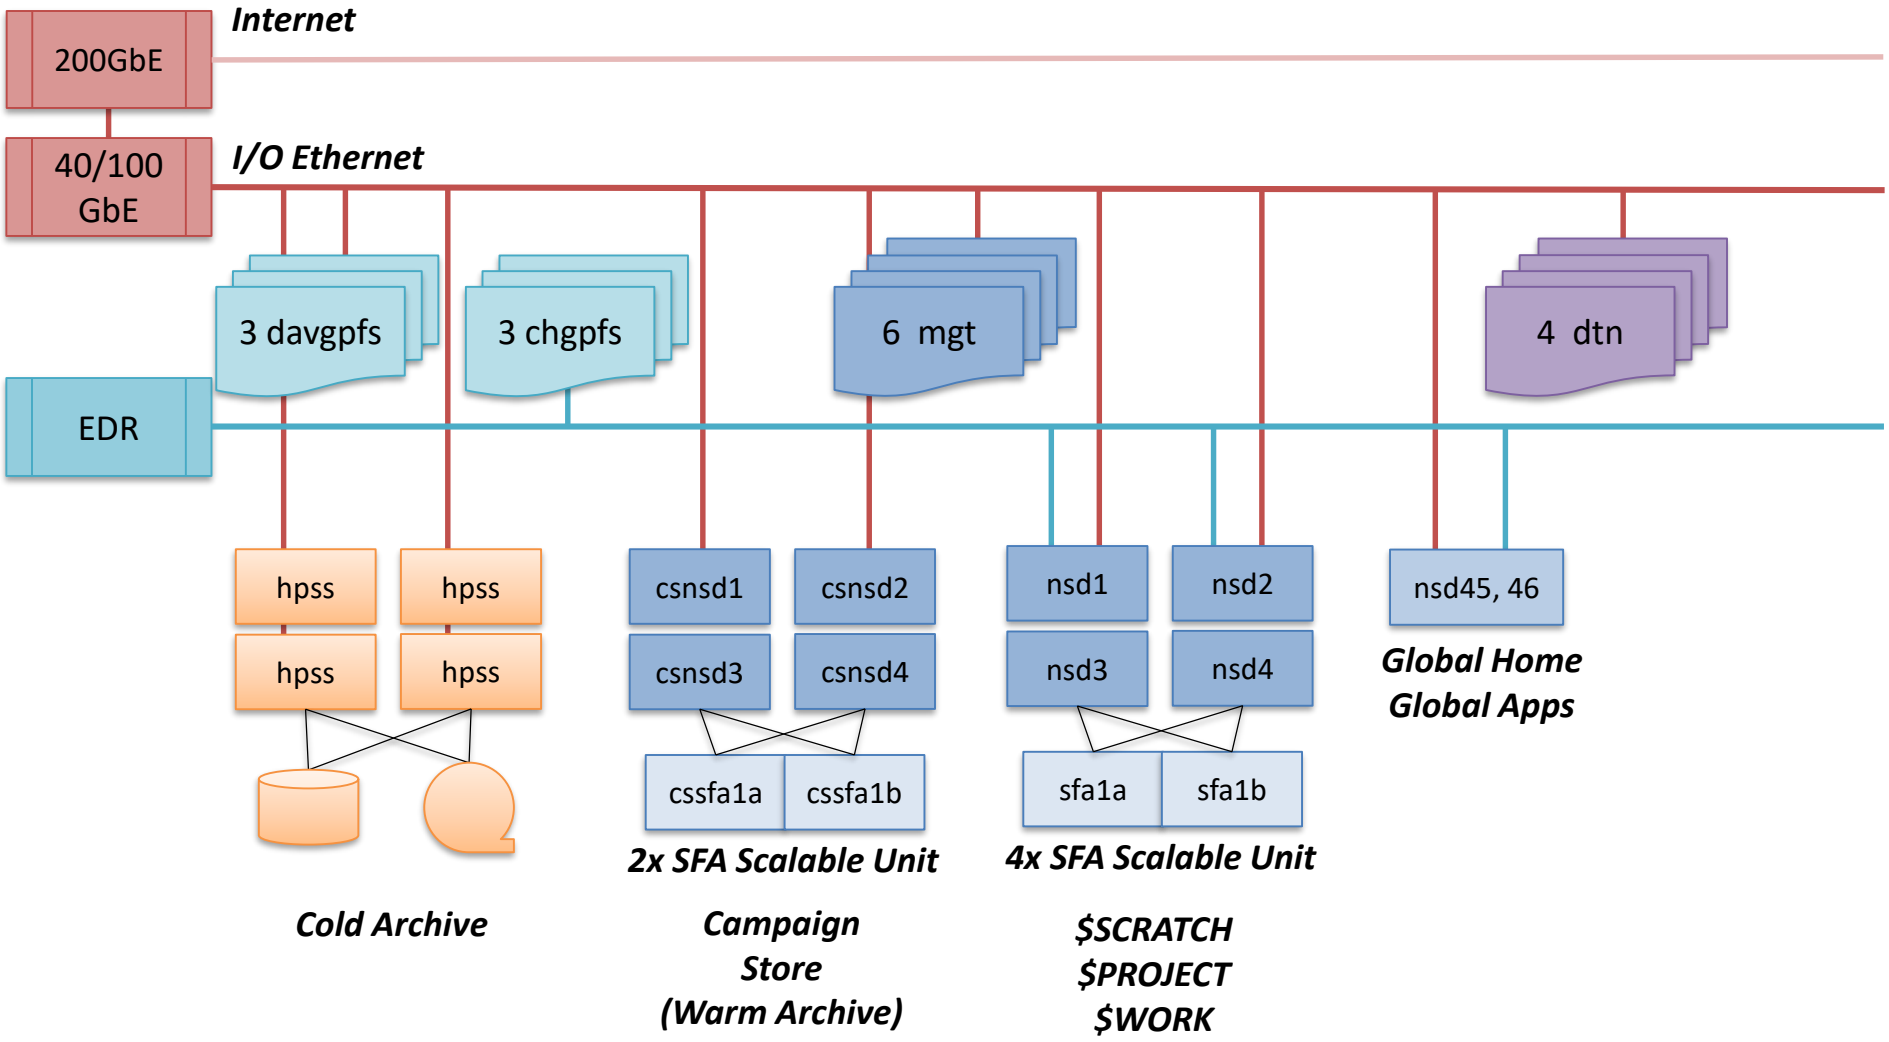

Spectrum Scale Multi-Cluster

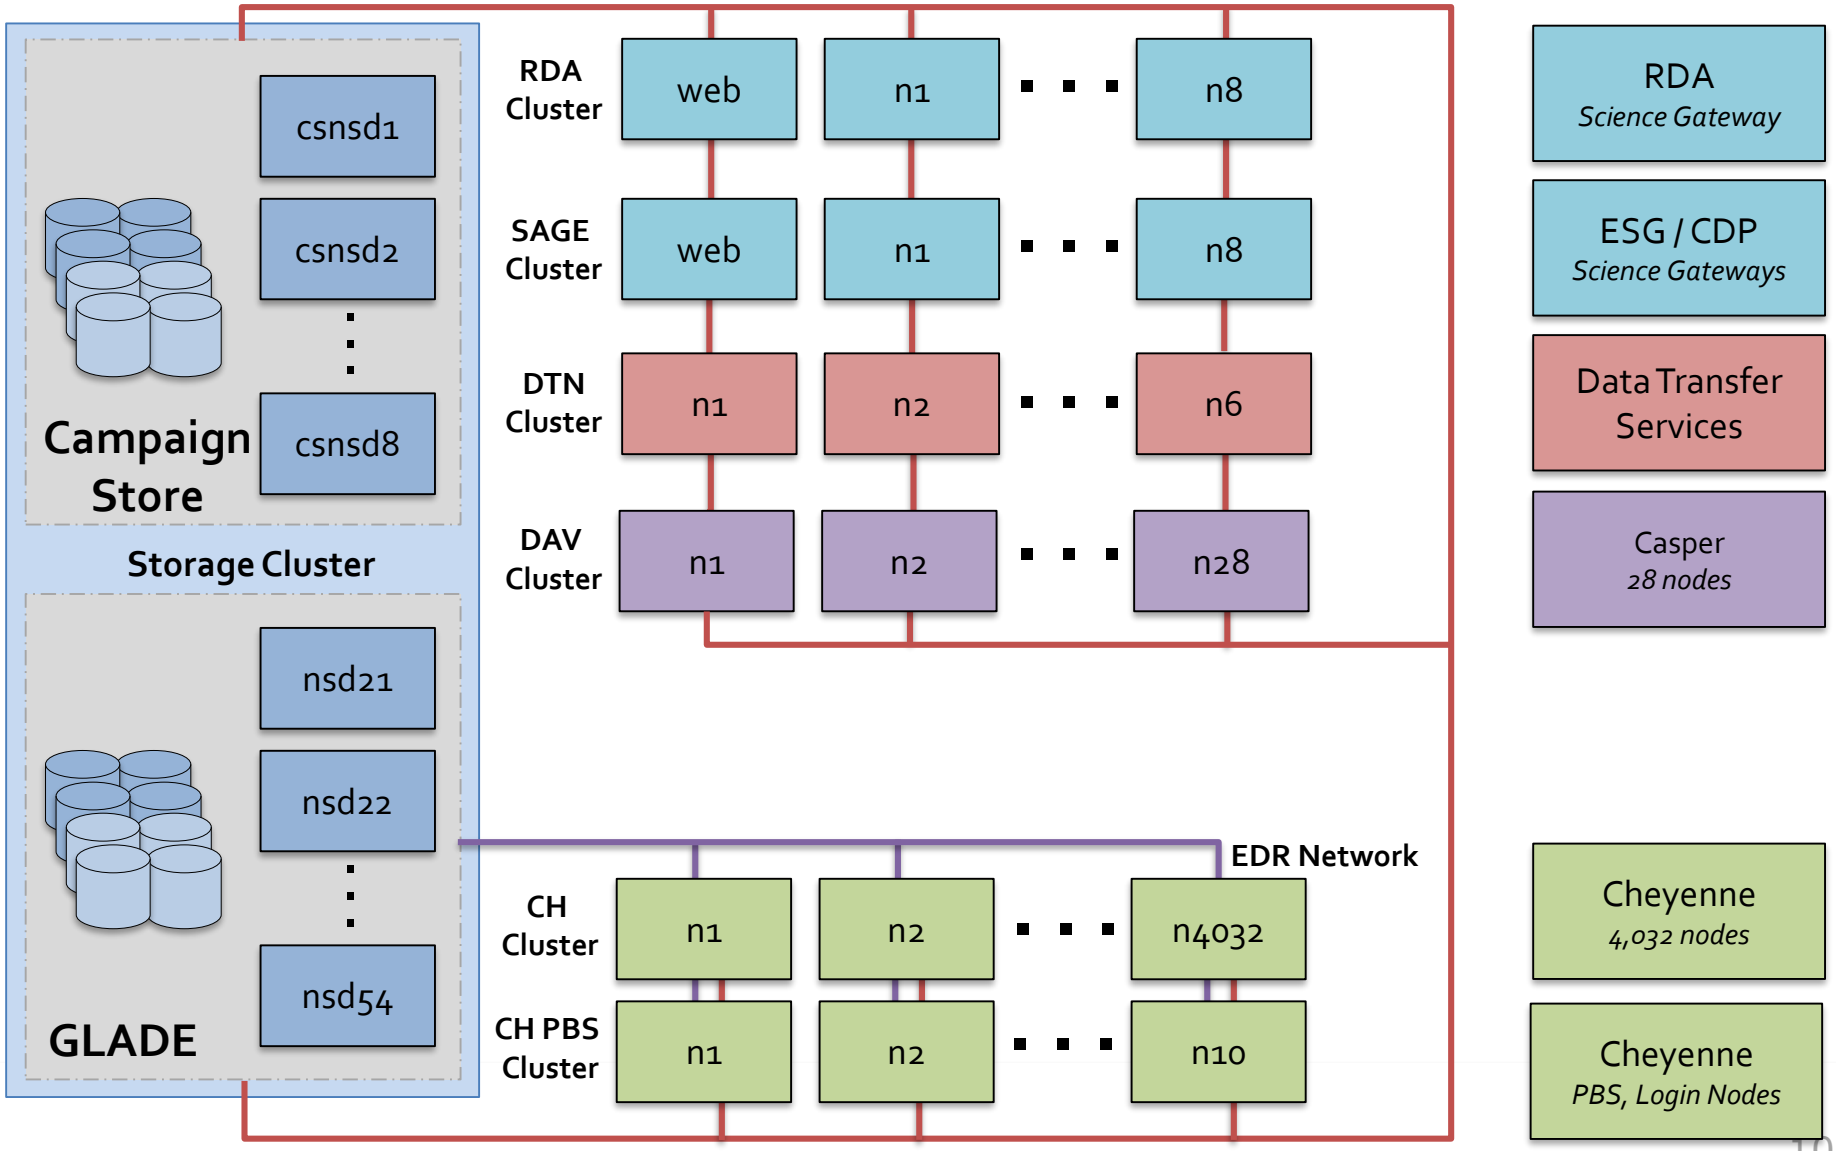

Storage System Configuration

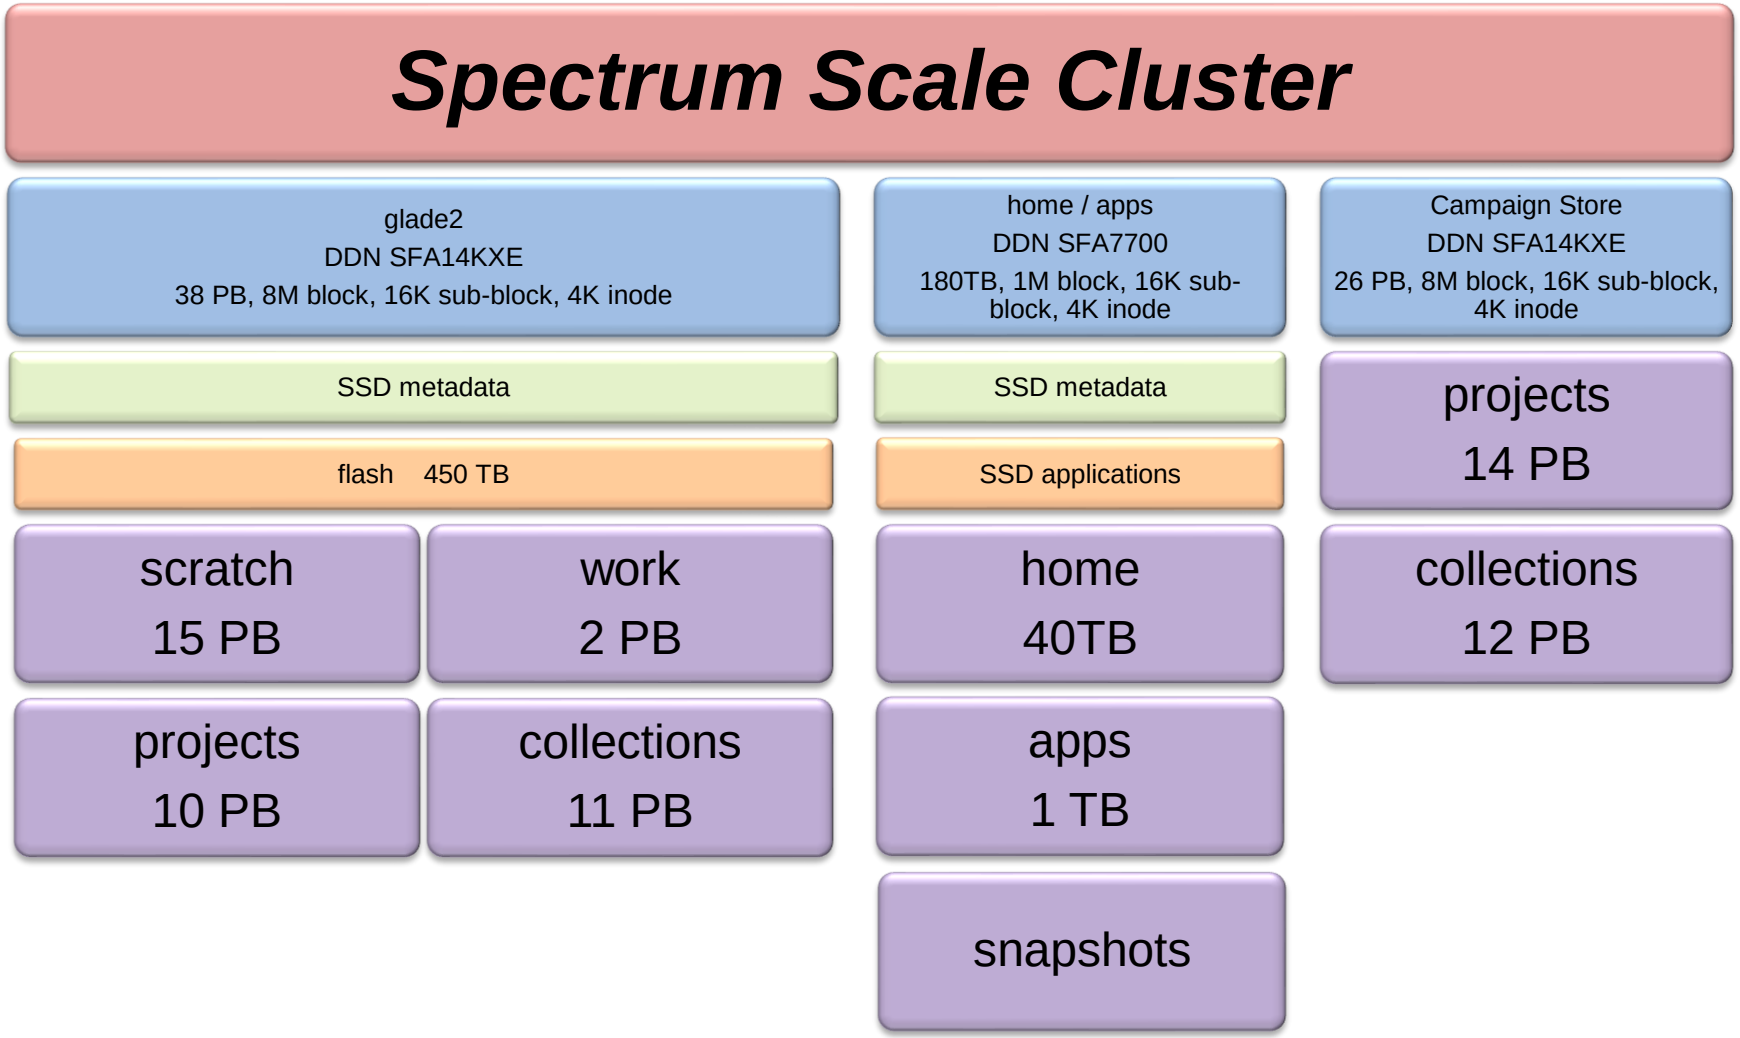

Allocation & Usage Policies

File Space	Peak Perf	Allocation	Quota	Purge	Backup
home	8 GB/s	by default	25 GB	No	Yes
scratch	300 GB/s	by default	10 TB	90 days	No
work	300 GB/s	by default	1 Tb	No	No
flash	300 GB/s	by request	None	2 weeks	No
project	300 GB/s	by allocation		1 year	No
datashare	300 GB/s	by request		45 days	No
Campaign Store	76 GB/s	by allocation		5 year	No
Tape (HPSS)	10 GB/s	by allocation		None	DR

Glade Space Usage

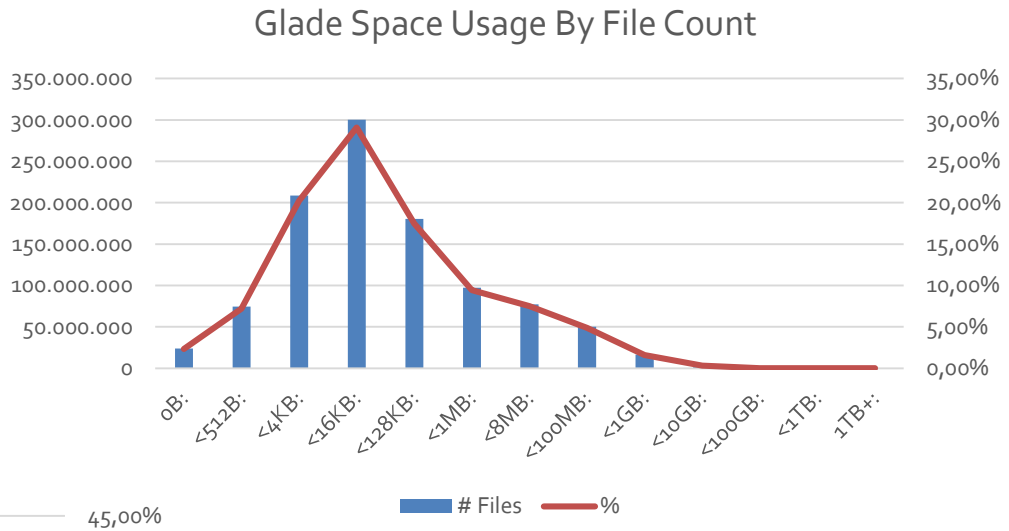

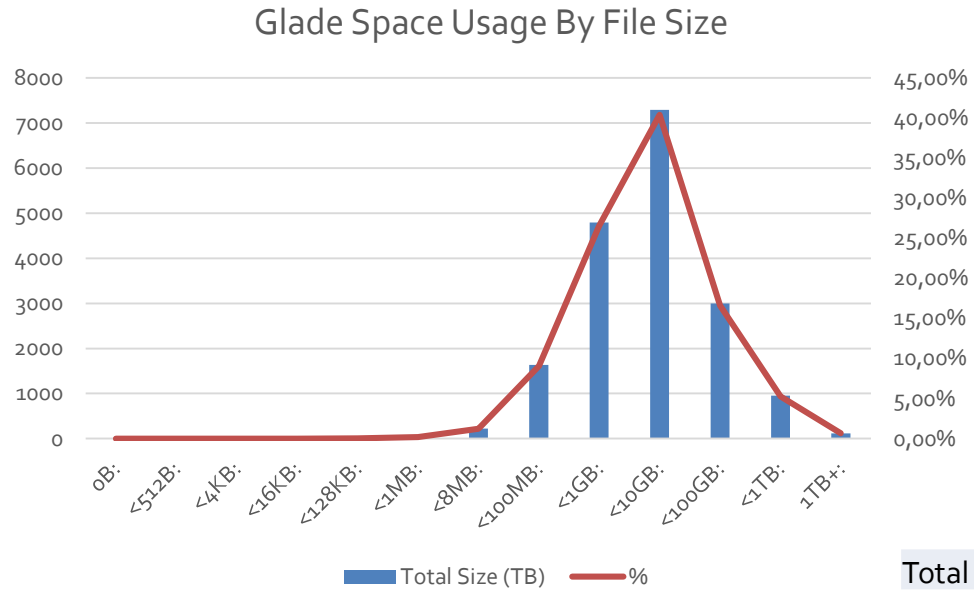

Total Size:	1,154,681,318	17.6321 PB	17.6462 PB allocated
Files:	1,032,066,333	17.6318 PB	17.6460 PB allocated
Dirs:	46,572,220	291.80 GB	244.49 GB allocated
Links:	76,042,765	6.06 GB	0 b allocated

GLADE Performance – Past 6 Years

GLADE2 Performance – Past 2 Years

GLADE Futures

Challenges

Future Architectures

Data Storage Challenges

- Archival system ingest rates can't keep up with data production rates
  - Tape is best for long term storage of well tagged and cataloged data
  - For data that can't be reproduced
  - For preservation of published data collections
  - To facilitate better data discovery
- Are there better solutions for short term storage?
  - Data needed during computation only
    - Faster memory/storage layer available during job run
    - Better workflow methods to streamline data output
  - Data with well defined life time
    - What is the useful life span of the data set?
    - At what point is newer data actually more valuable

Data Management Challenges

- We can produce data faster than we can tag, catalog and publish
- How do you find the data you need?
  - Do you have effective tools to manage your data?
  - Do you have a good sense for the data you have archived?
  - Do you have methods for re-evaluating if data is still needed?
- How do you ensure data is deleted when no longer needed?
  - Storage costs require better data management techniques
- What is the effective life time of data collections?
  - At what point does newer data supersede older data sets?
  - At what point is it more effective to reproduce a data set than to store older data?

Data Lifecycle Management

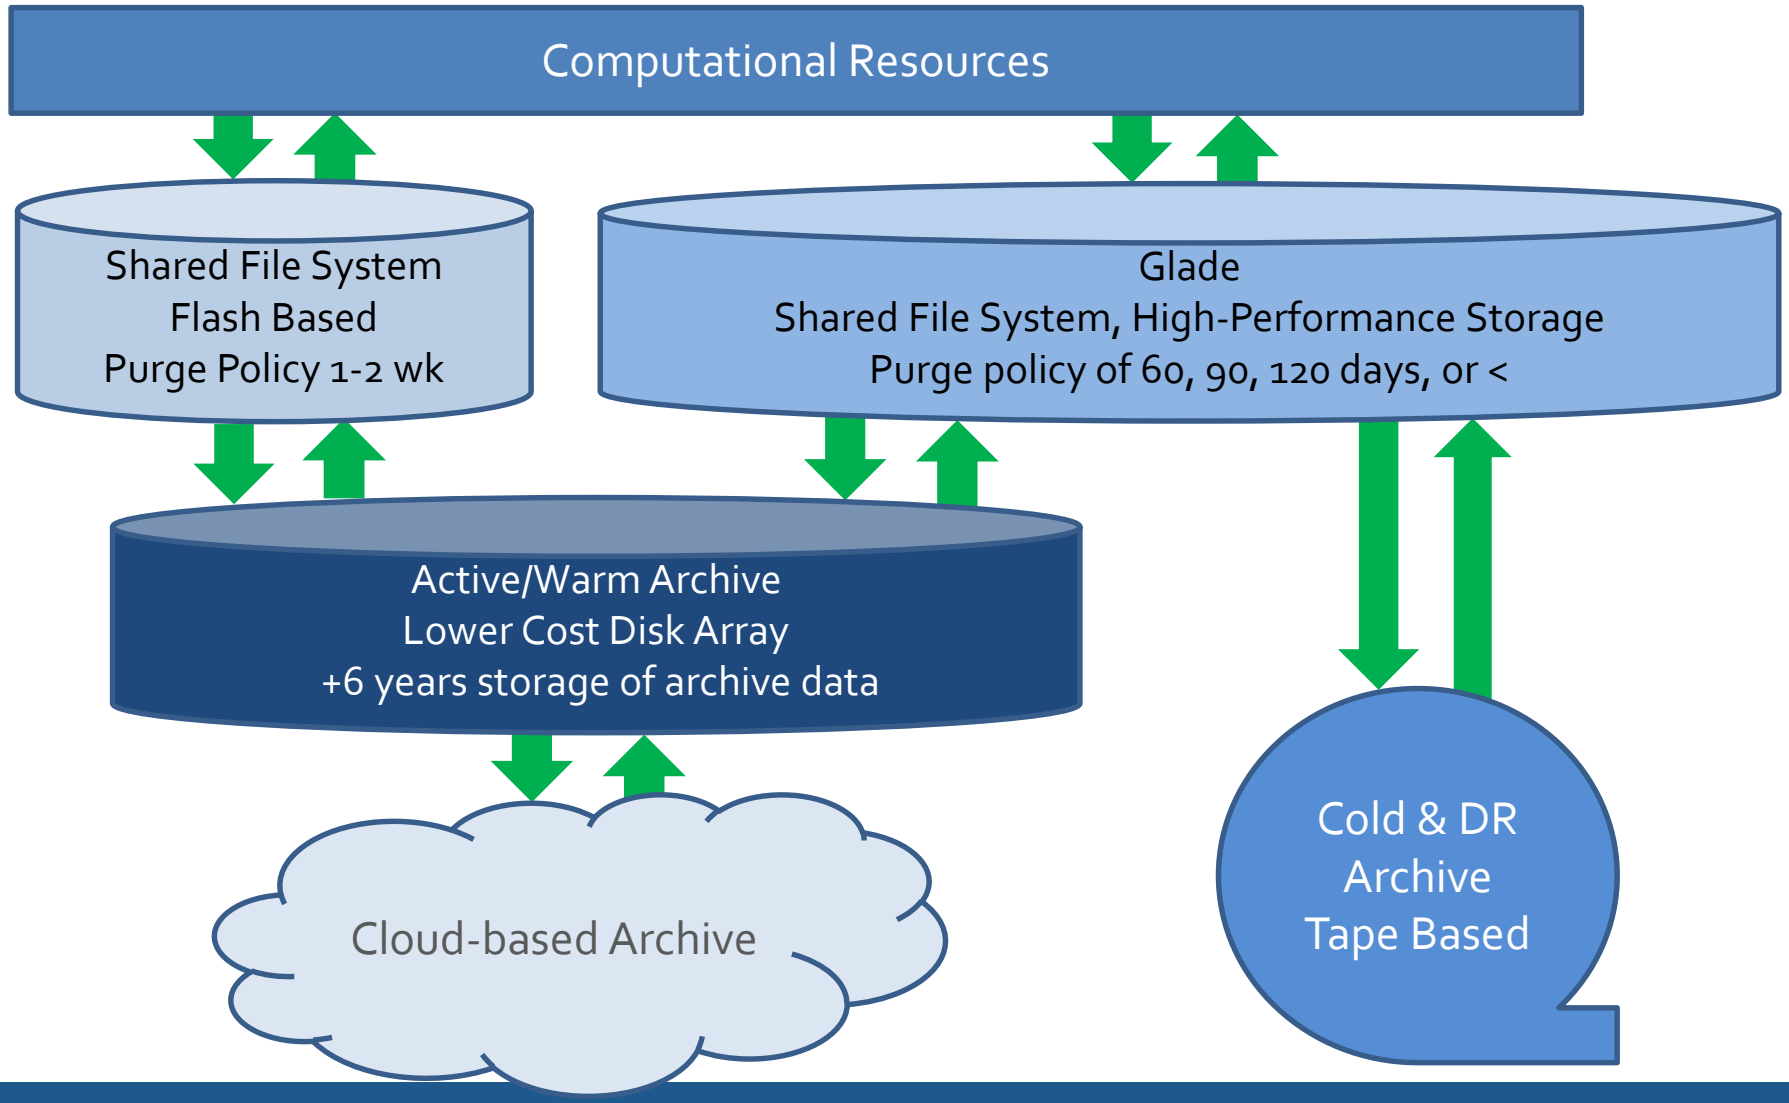

Storage Hierarchy

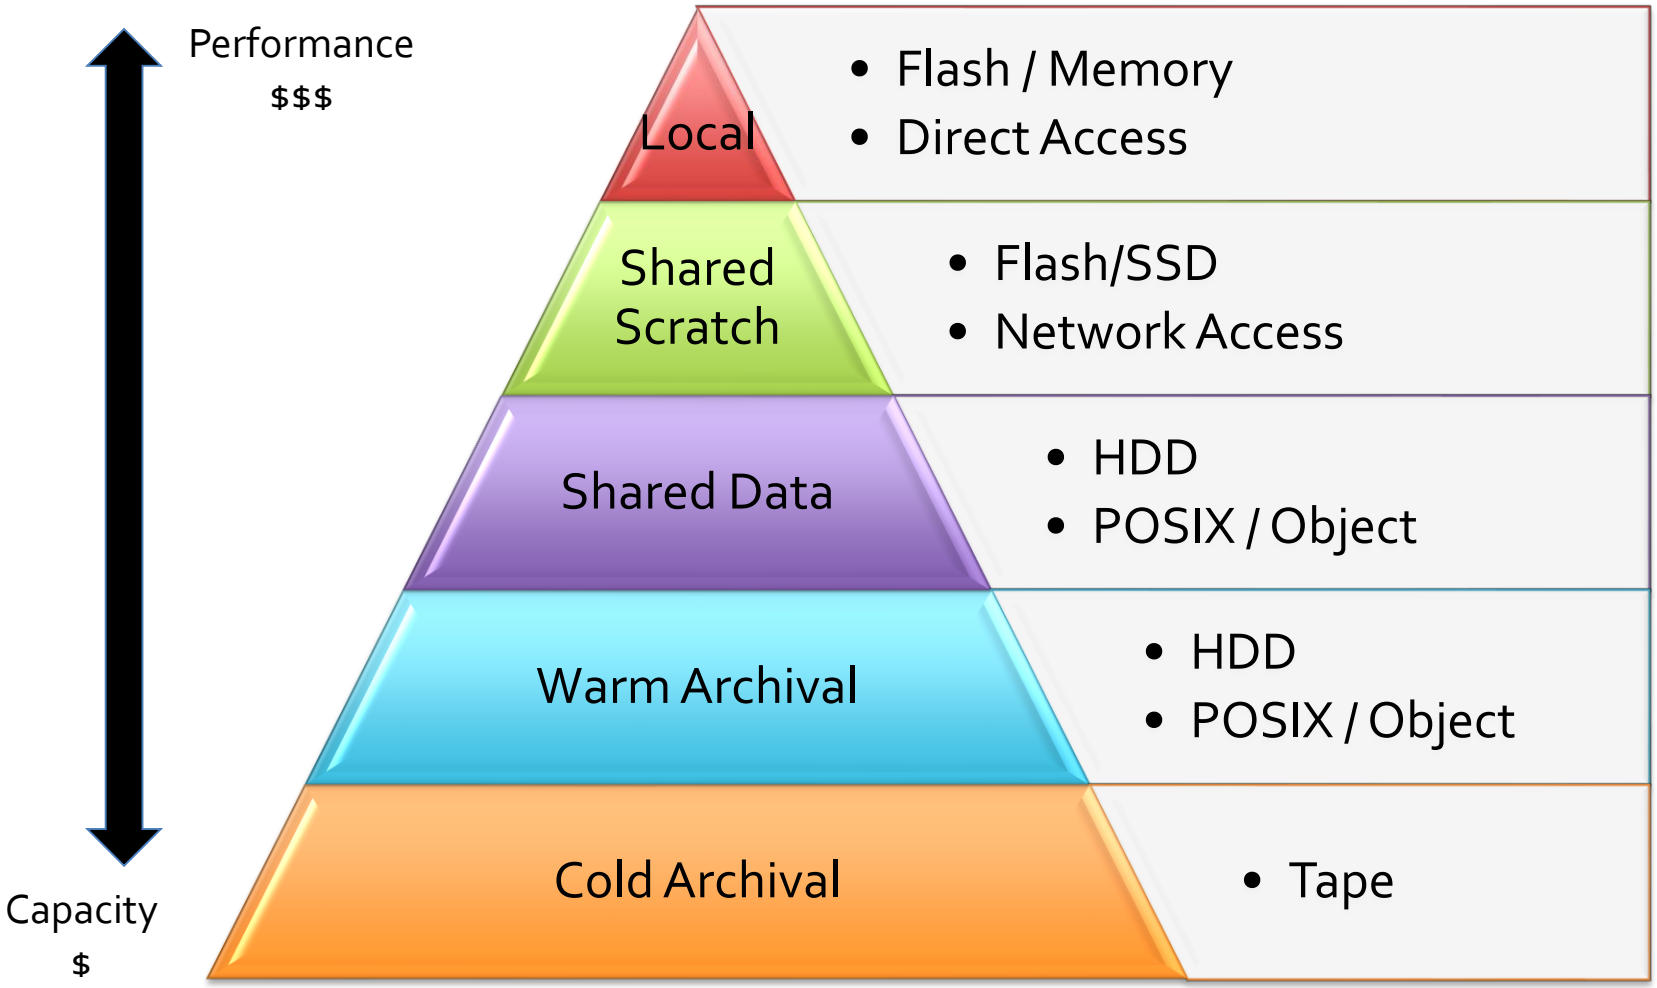

Data Hierarchy

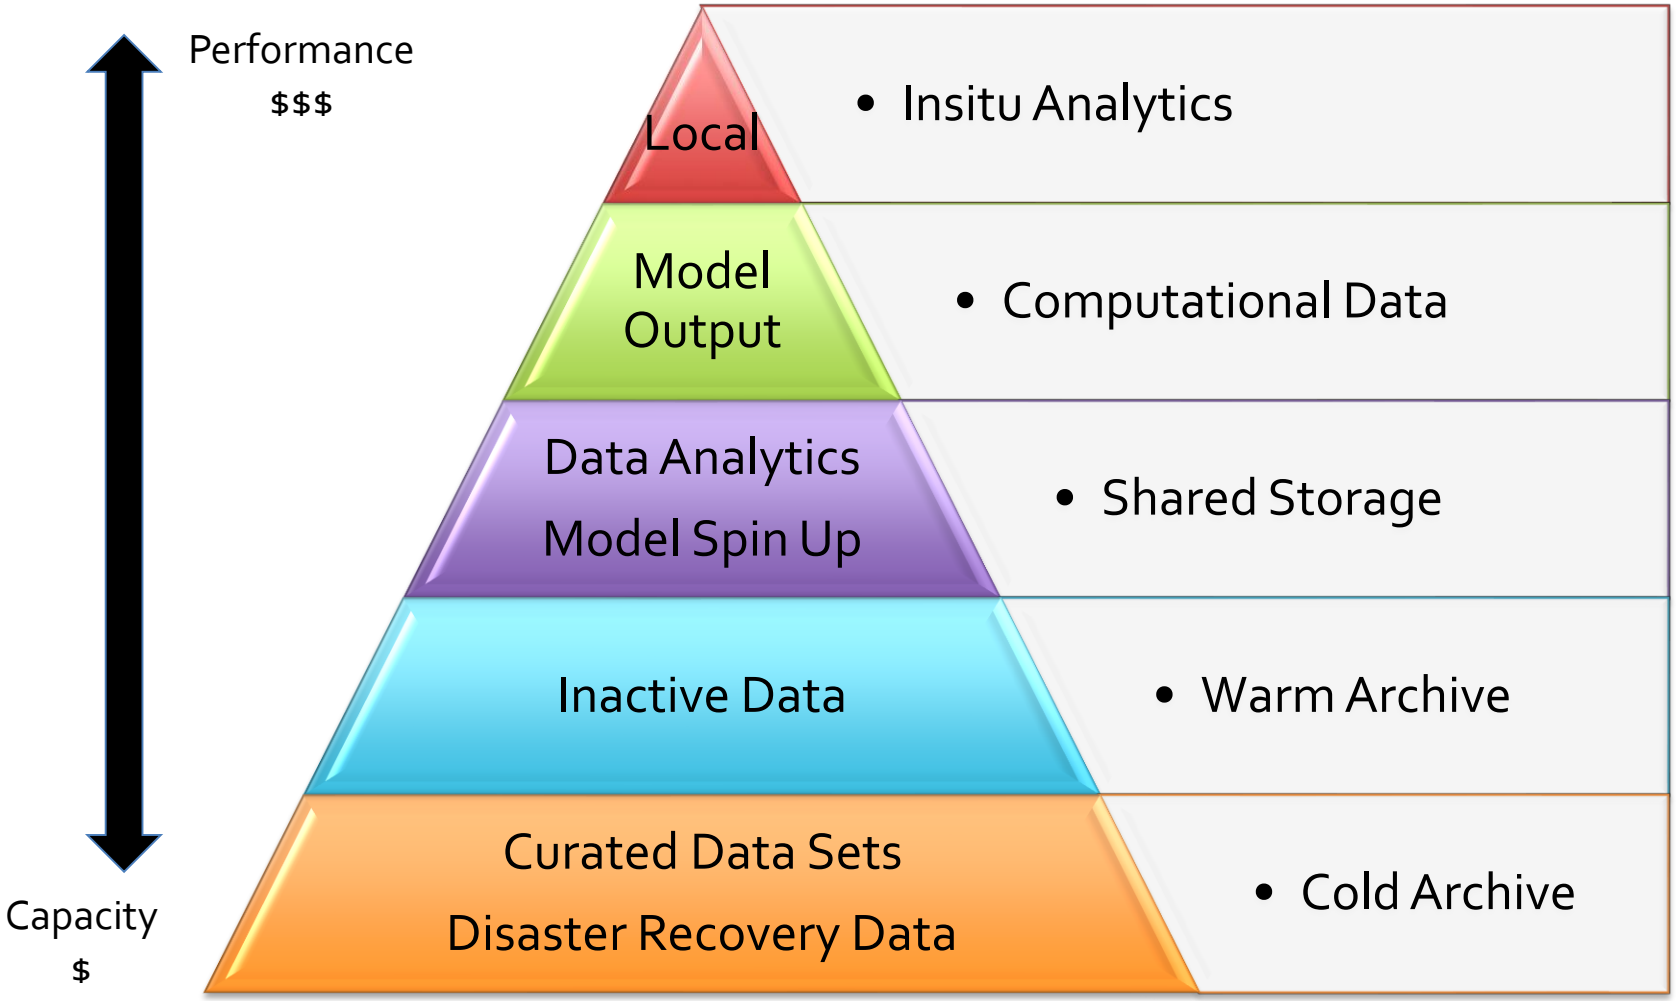

HPC Data Infrastructure Group
NCAR / CISL / HSS / HDIG

Data Services

- High-Performance File Systems
  - GLADE Spectrum Scale Storage Cluster
- Data Archival Resources
  - Campaign Store – Warm Archive
  - HPSS – Cold Archive
- Data Transfer Protocols
  - Globus
  - GridFTP
  - SCP/SFTP
  - HSI/HTAR
- Science Gateway Support
  - RDA/ESG/CDP

Innovative I/O Solutions

- New network technologies
- New storage technologies
- Innovative memory architectures

Staff

- Pamela Hill
- Marc Genty
- Bill Anderson
- Chris Hoffman
- Joey Mendoza
- Zach Mance

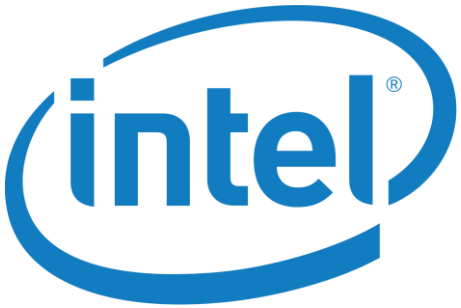

IBM
Spectrum
Scale

QUESTIONS?