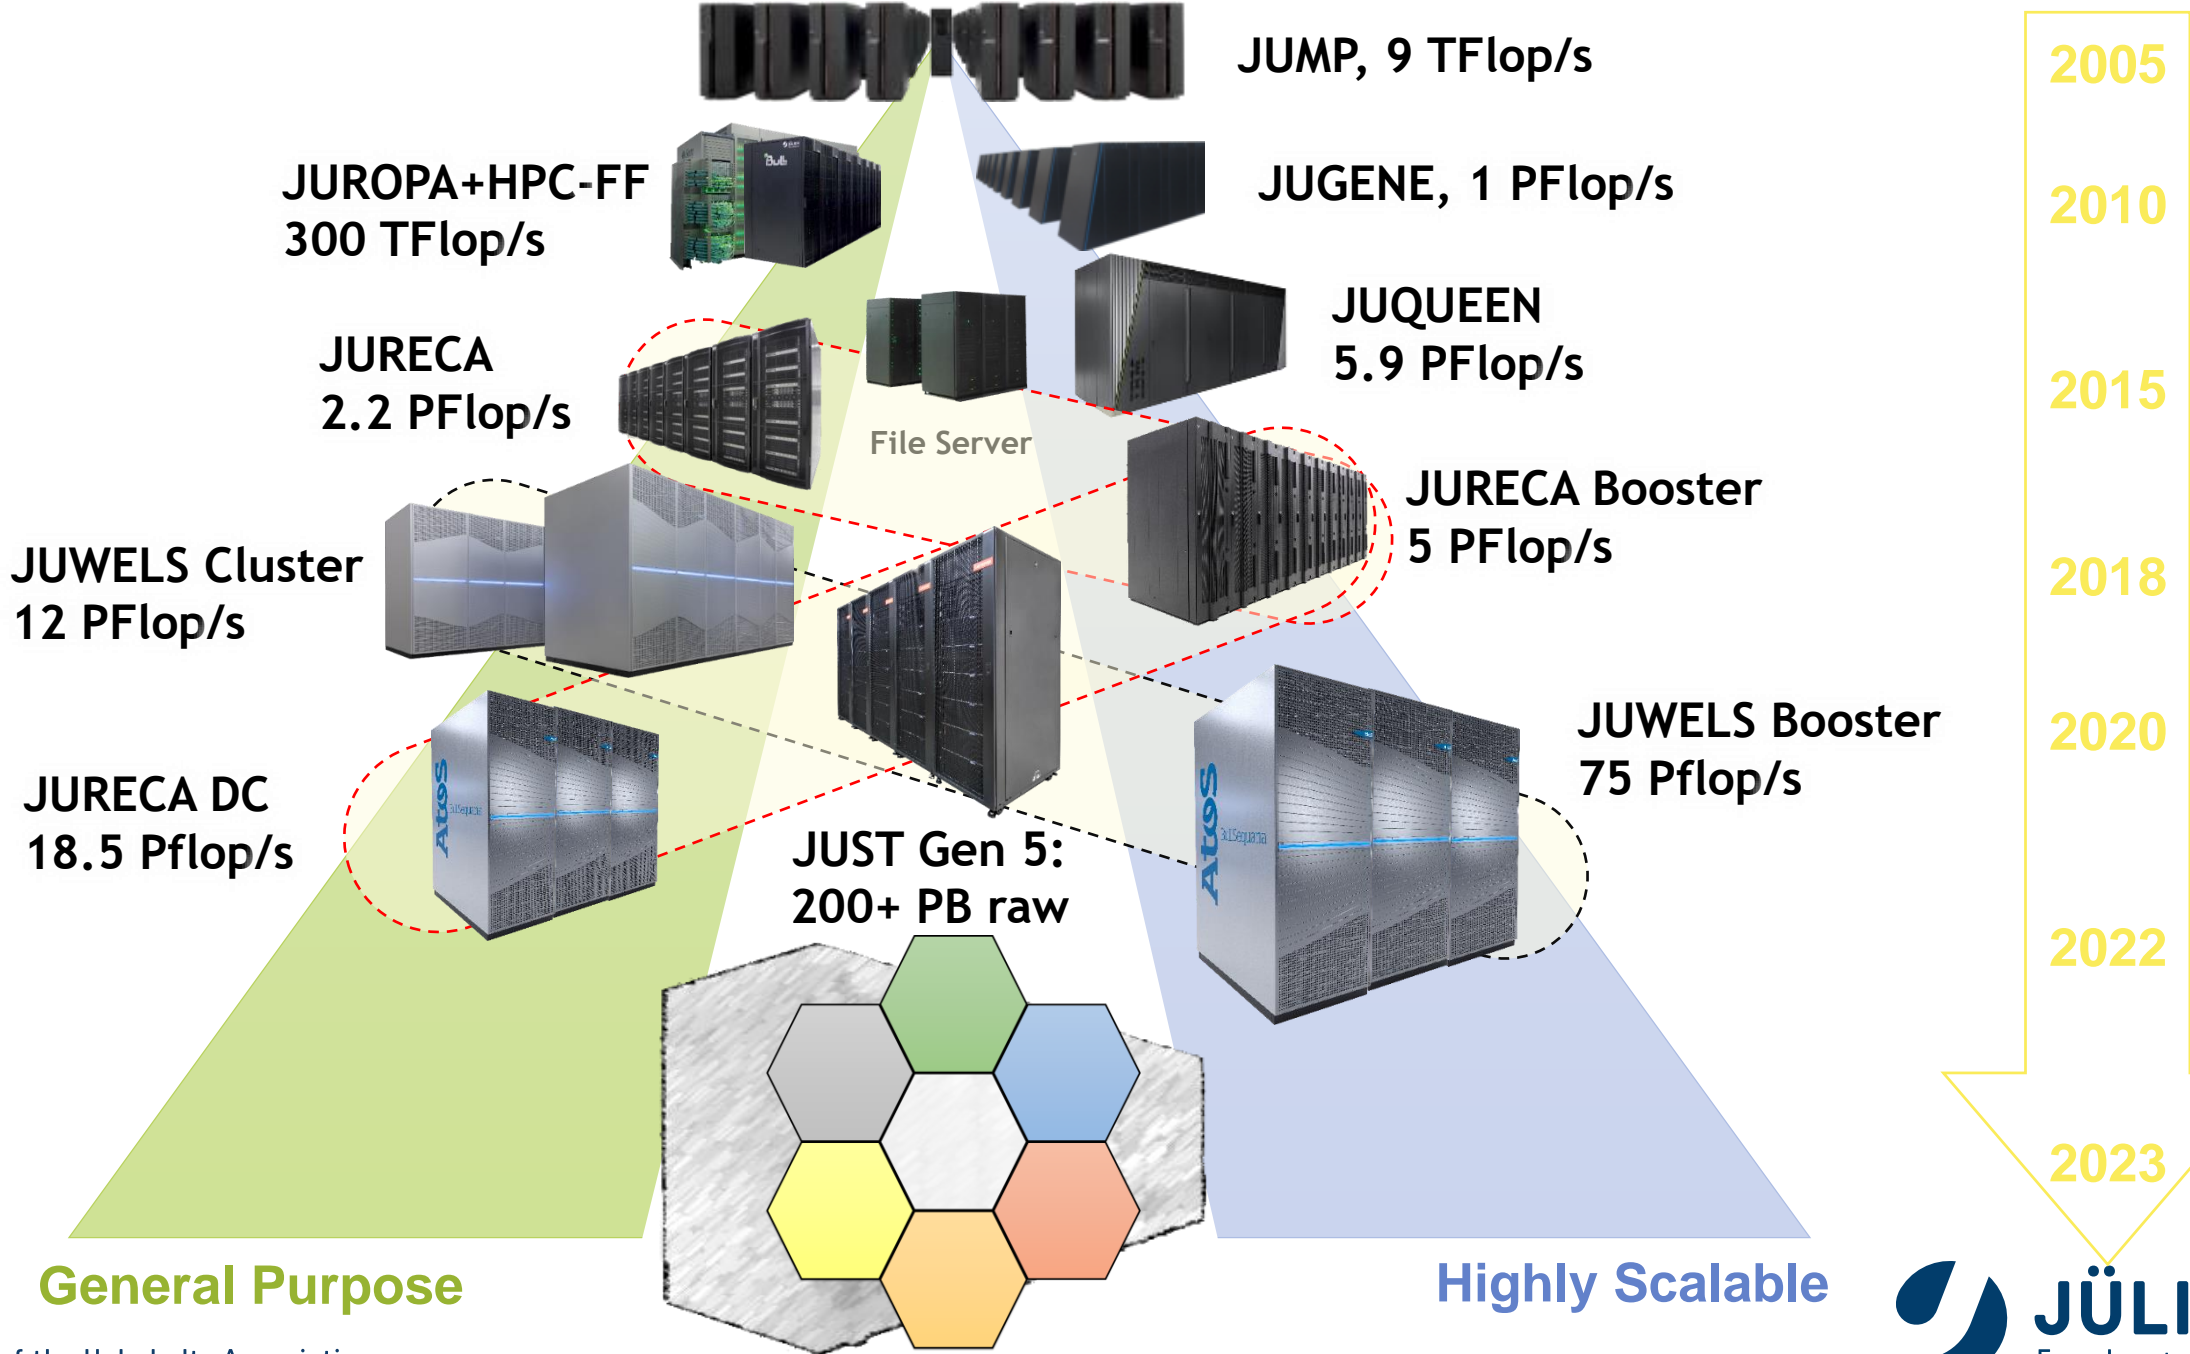



STORAGE SERVICES FOR HPC AND BEYOND JUELICH STORAGE CLUSTER JUST

22. SEPTEMBER 2023 | STEPHAN GRAF (JÜLICH SUPERCOMPUTING CENTRE)



Exascale Modular Supercomputer

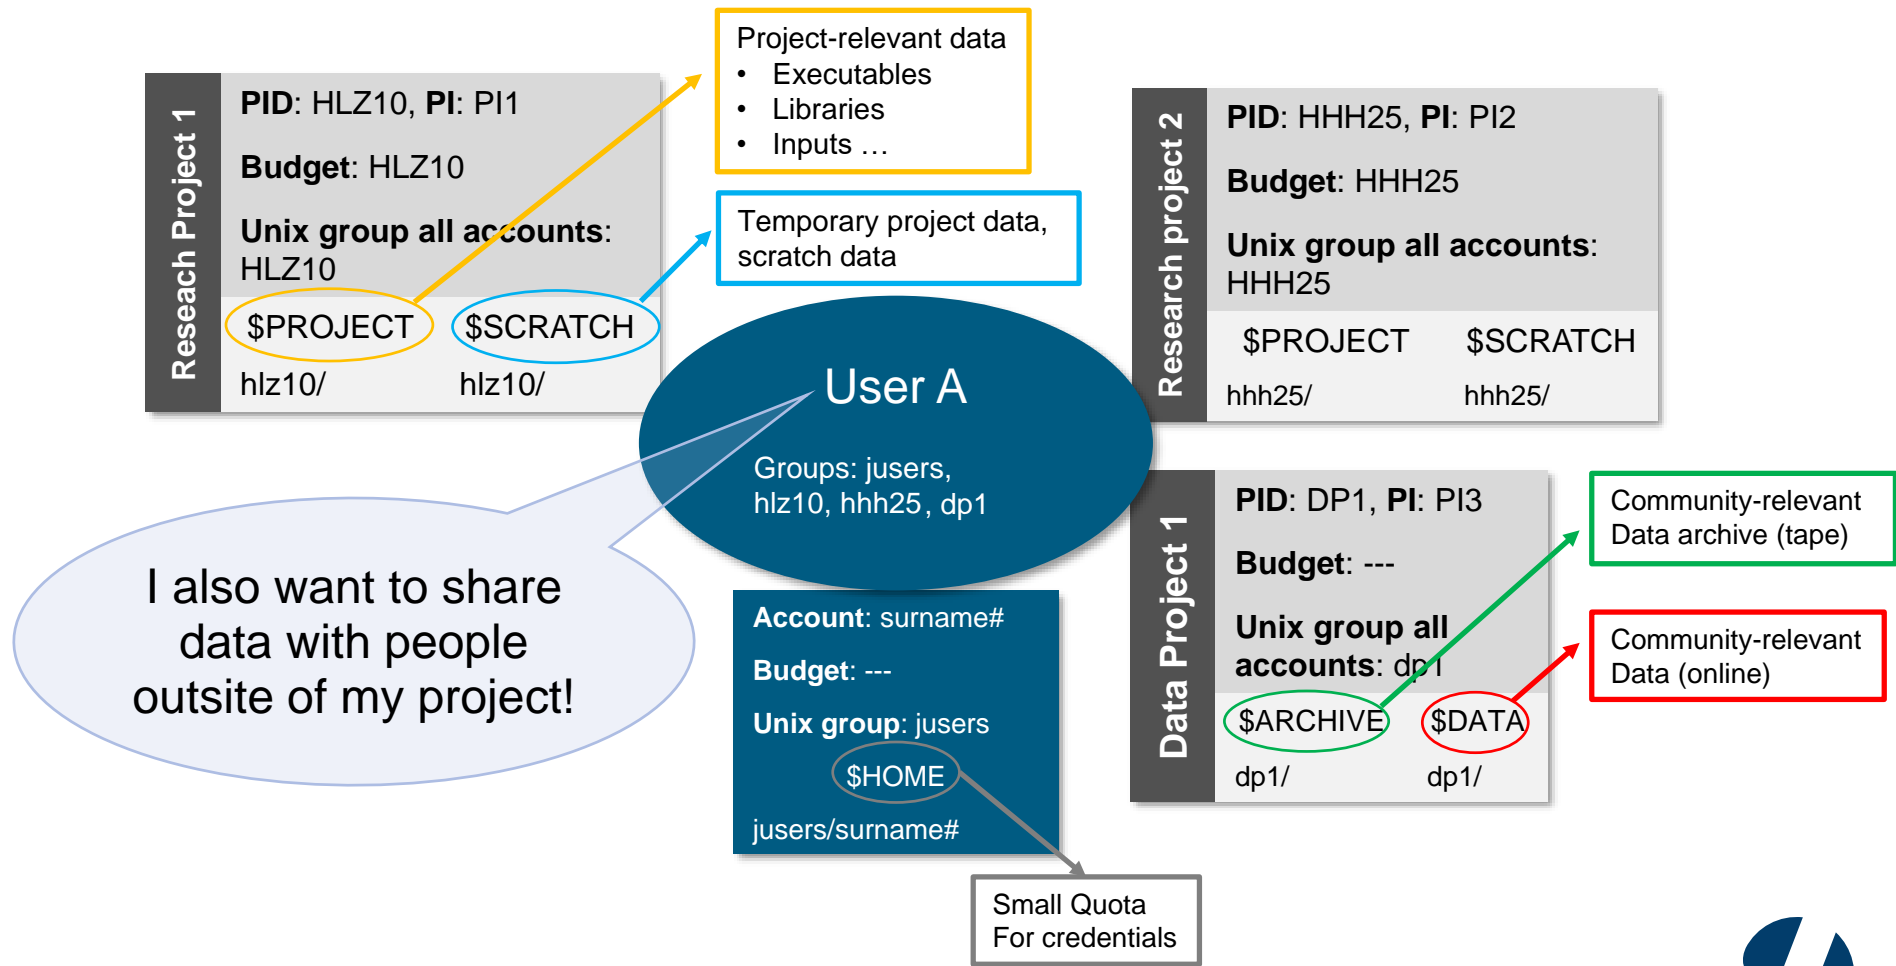
TIERED STORAGE OFFERING

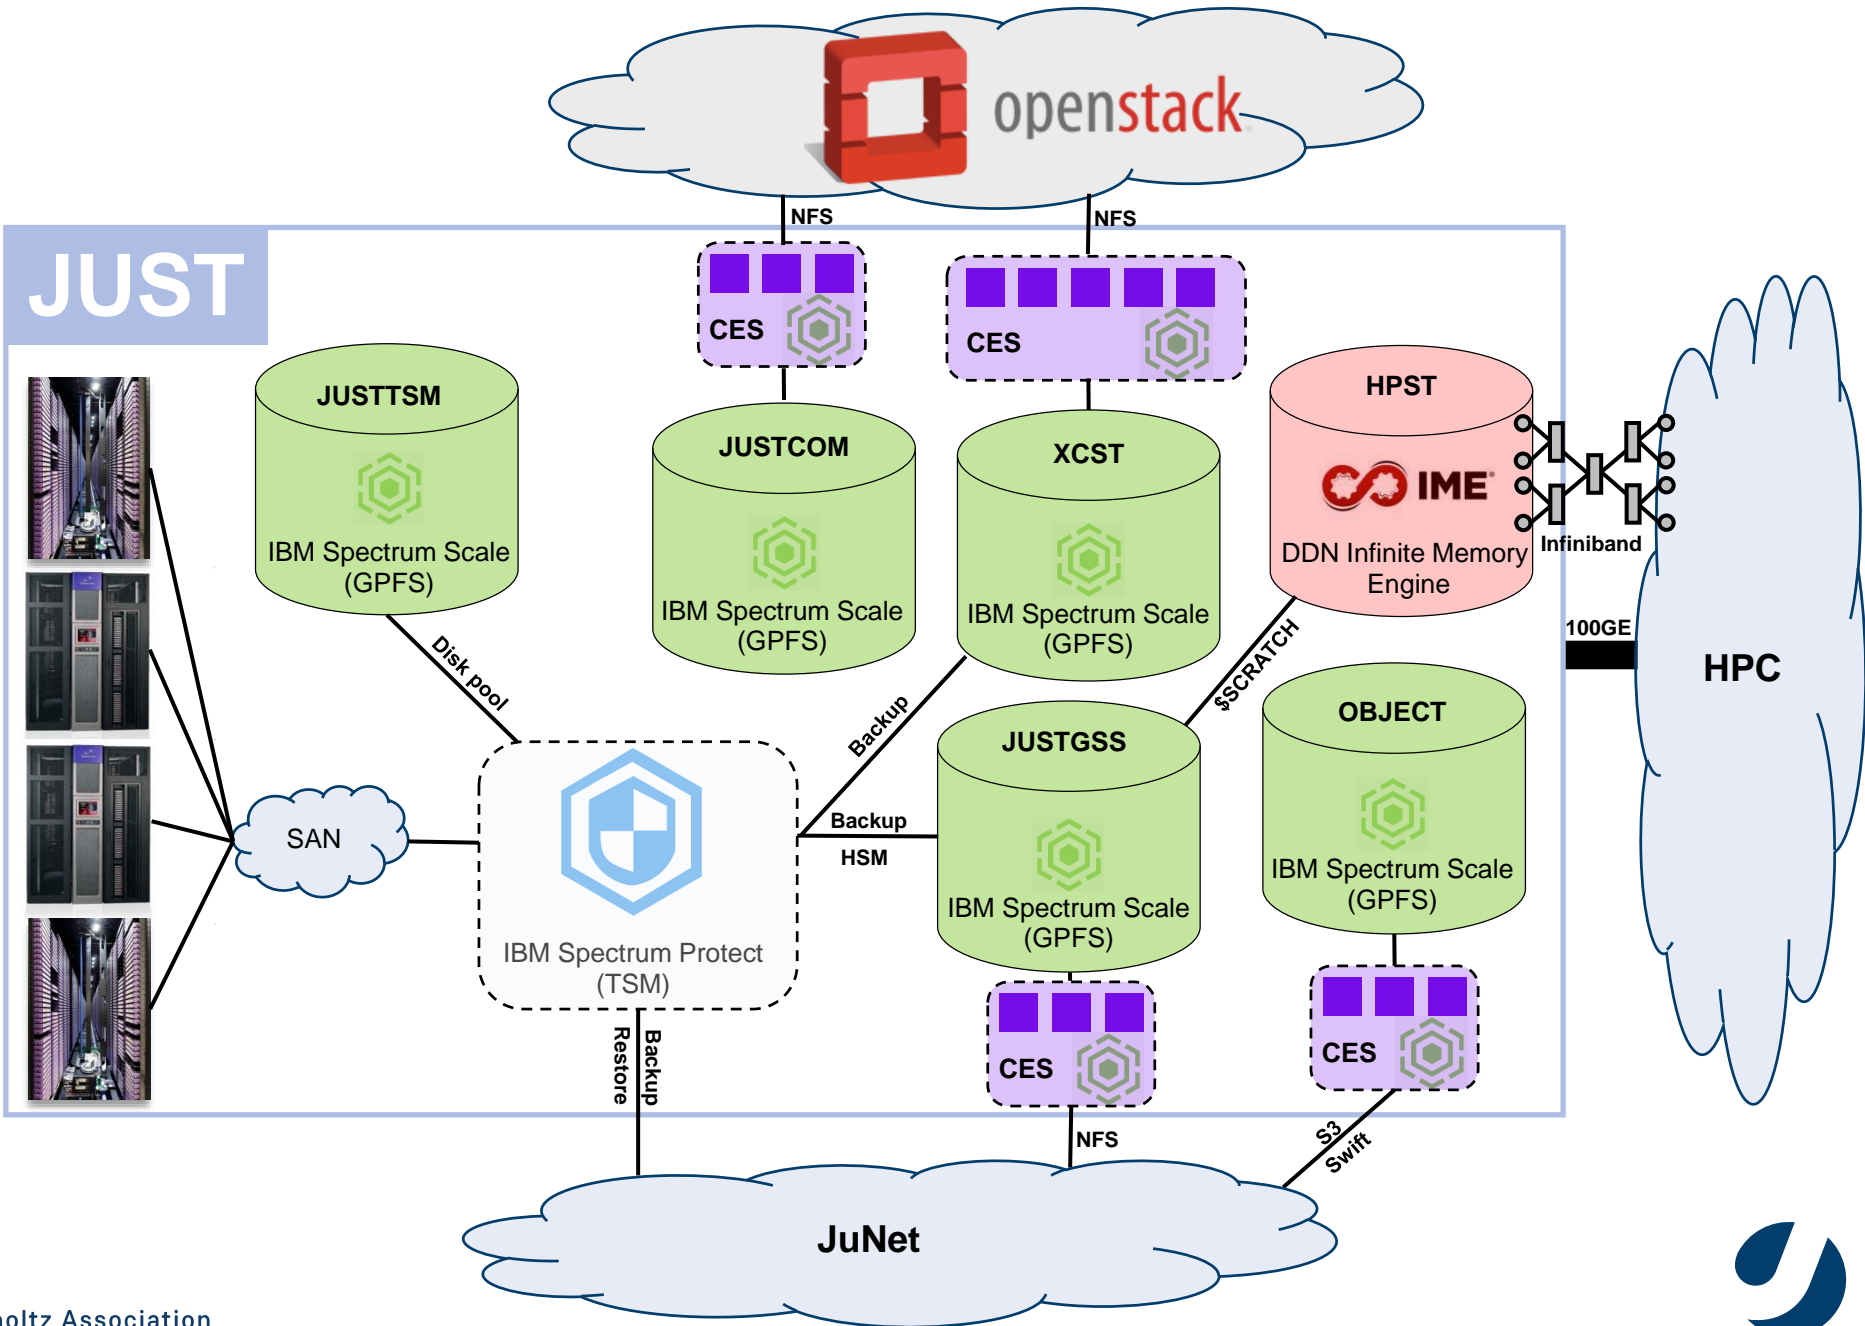



- High Performance Storage Tier (HPST): NVMe based Storage (low latency+high bandwidth)
- Large Capacity Storage Tier (LCST): Lenovo DSS Cluster (GNR, 5th Gen. of JUST, bandwidth optimized) → LCST
- Extended Capacity Storage Tier (XCST): GPFS Building Blocks (target: capacity) → XCST
- Archive: Tape storage (Backup + GPFS&TSM-HSM)

JÜLICH USAGE MODEL


User-centered







JUST



Disk pool



Backup
Restore



Backup

Backup
HSM

Backup/Archive Infrastructure:

- ~ 380 PB tape capacity
 - 4 tape libraries (2 buildings)
 - 110 tape drives
 - ~ 39.000 tapes
- Software (Client): IBM Spectrum Protect
- Backup/Archive HPC (core user)
- Backup/Archive FZJ (→ IT-Service)
- dCache backend (LOFAR)



\$SCRATCH


Backup

Backup
HSM




NFS

S3
Swift



NFS





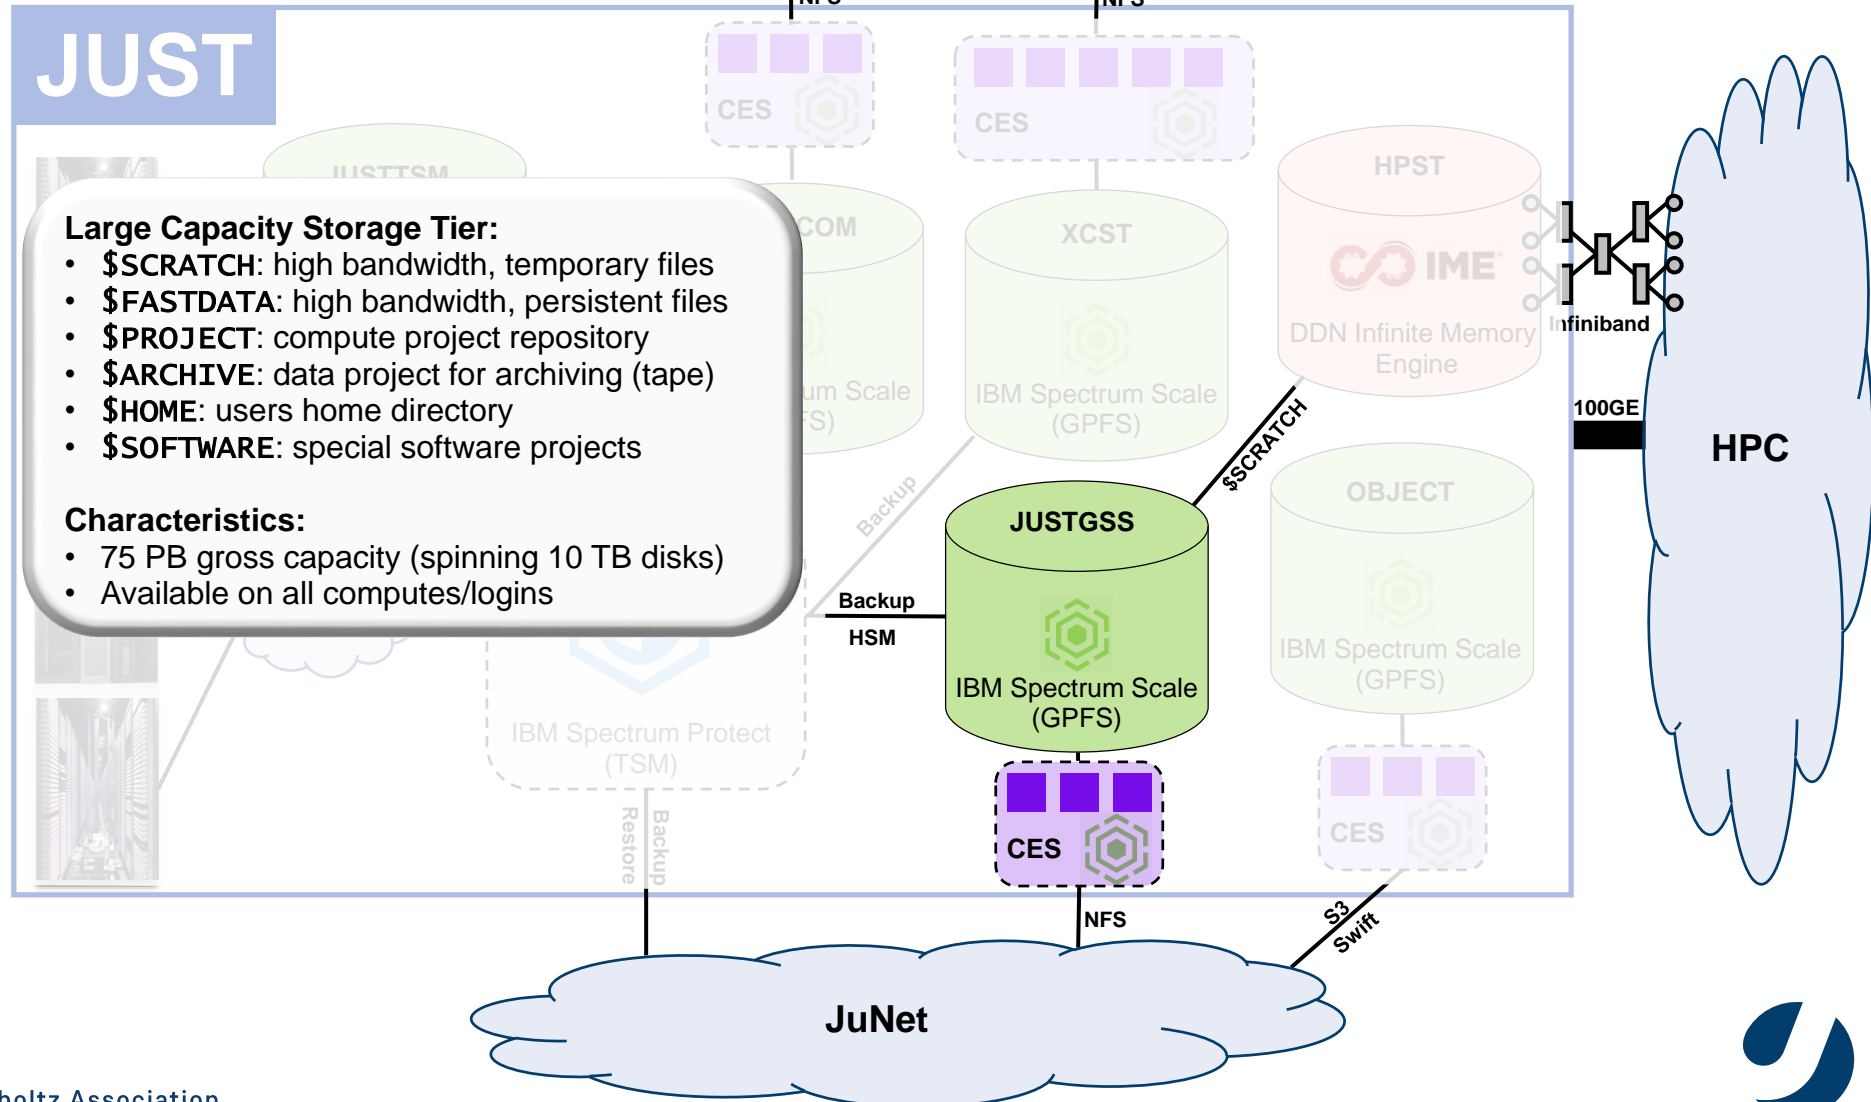
JUST

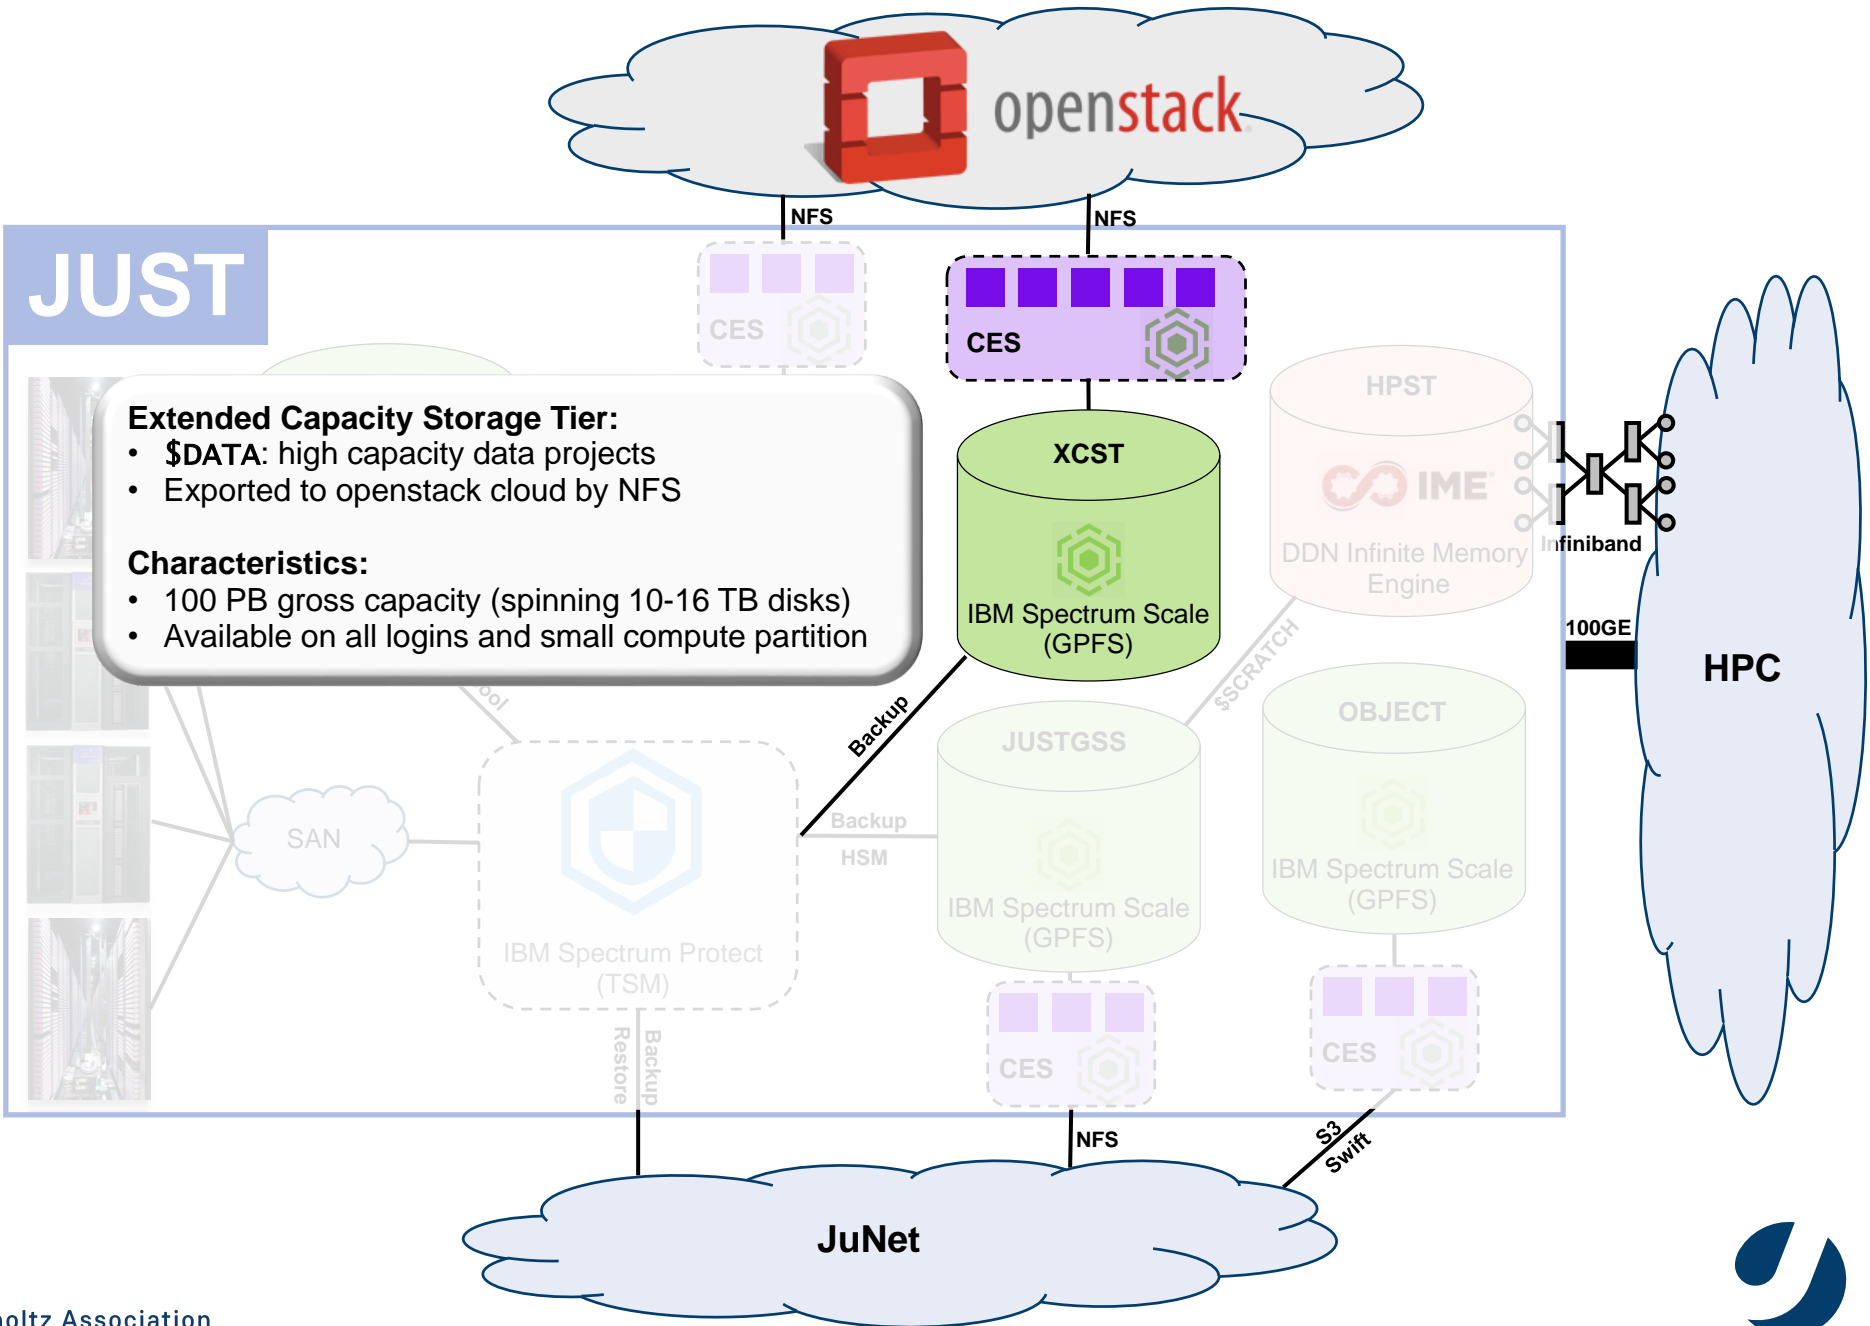
Large Capacity Storage Tier:

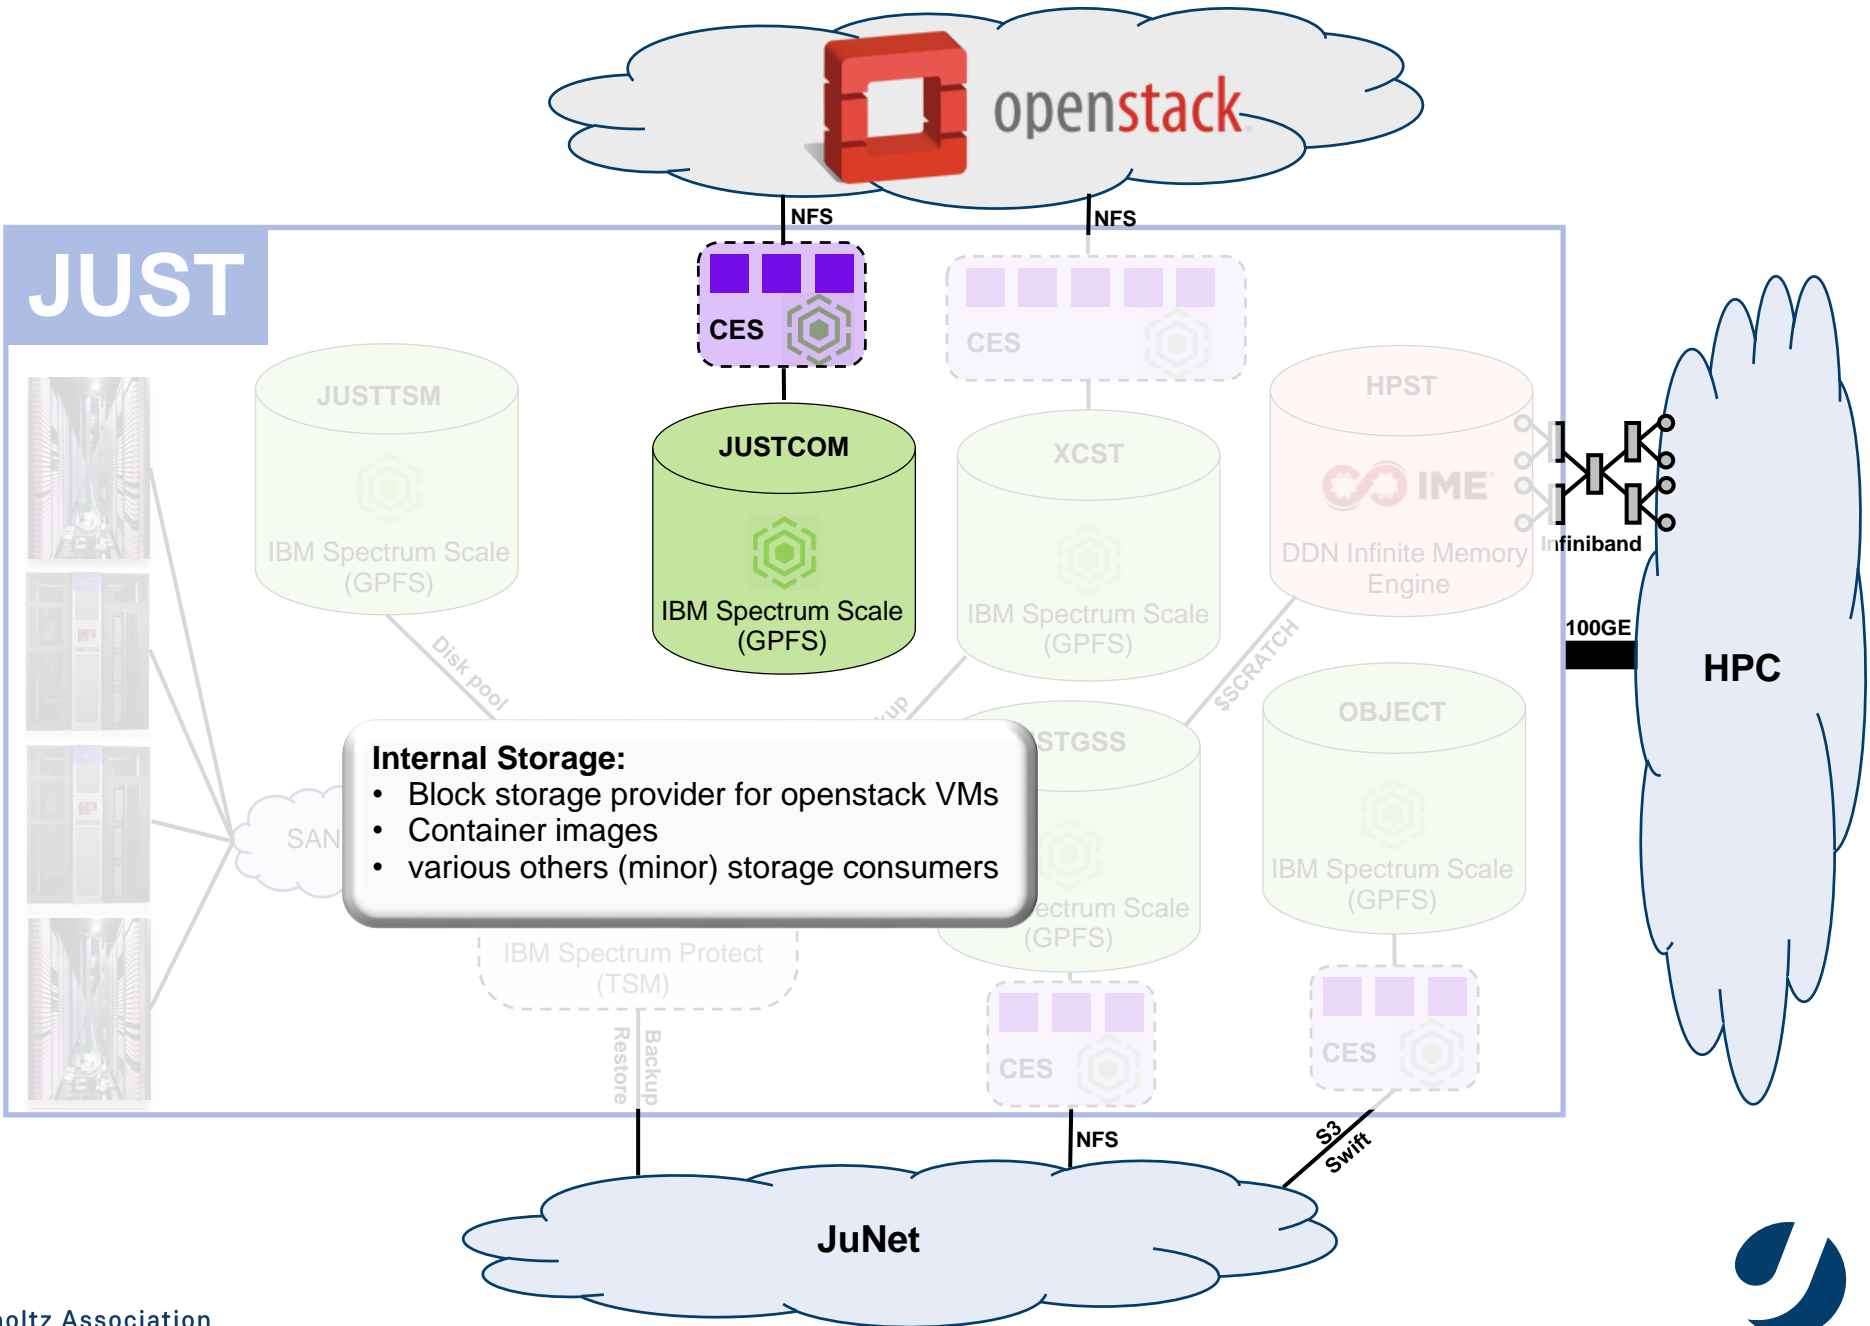
- \$SCRATCH: high bandwidth, temporary files
- \$FASTDATA: high bandwidth, persistent files
- \$PROJECT: compute project repository
- \$ARCHIVE: data project for archiving (tape)
- \$HOME: users home directory
- \$SOFTWARE: special software projects

Characteristics:

- 75 PB gross capacity (spinning 10 TB disks)
- Available on all computes/logins



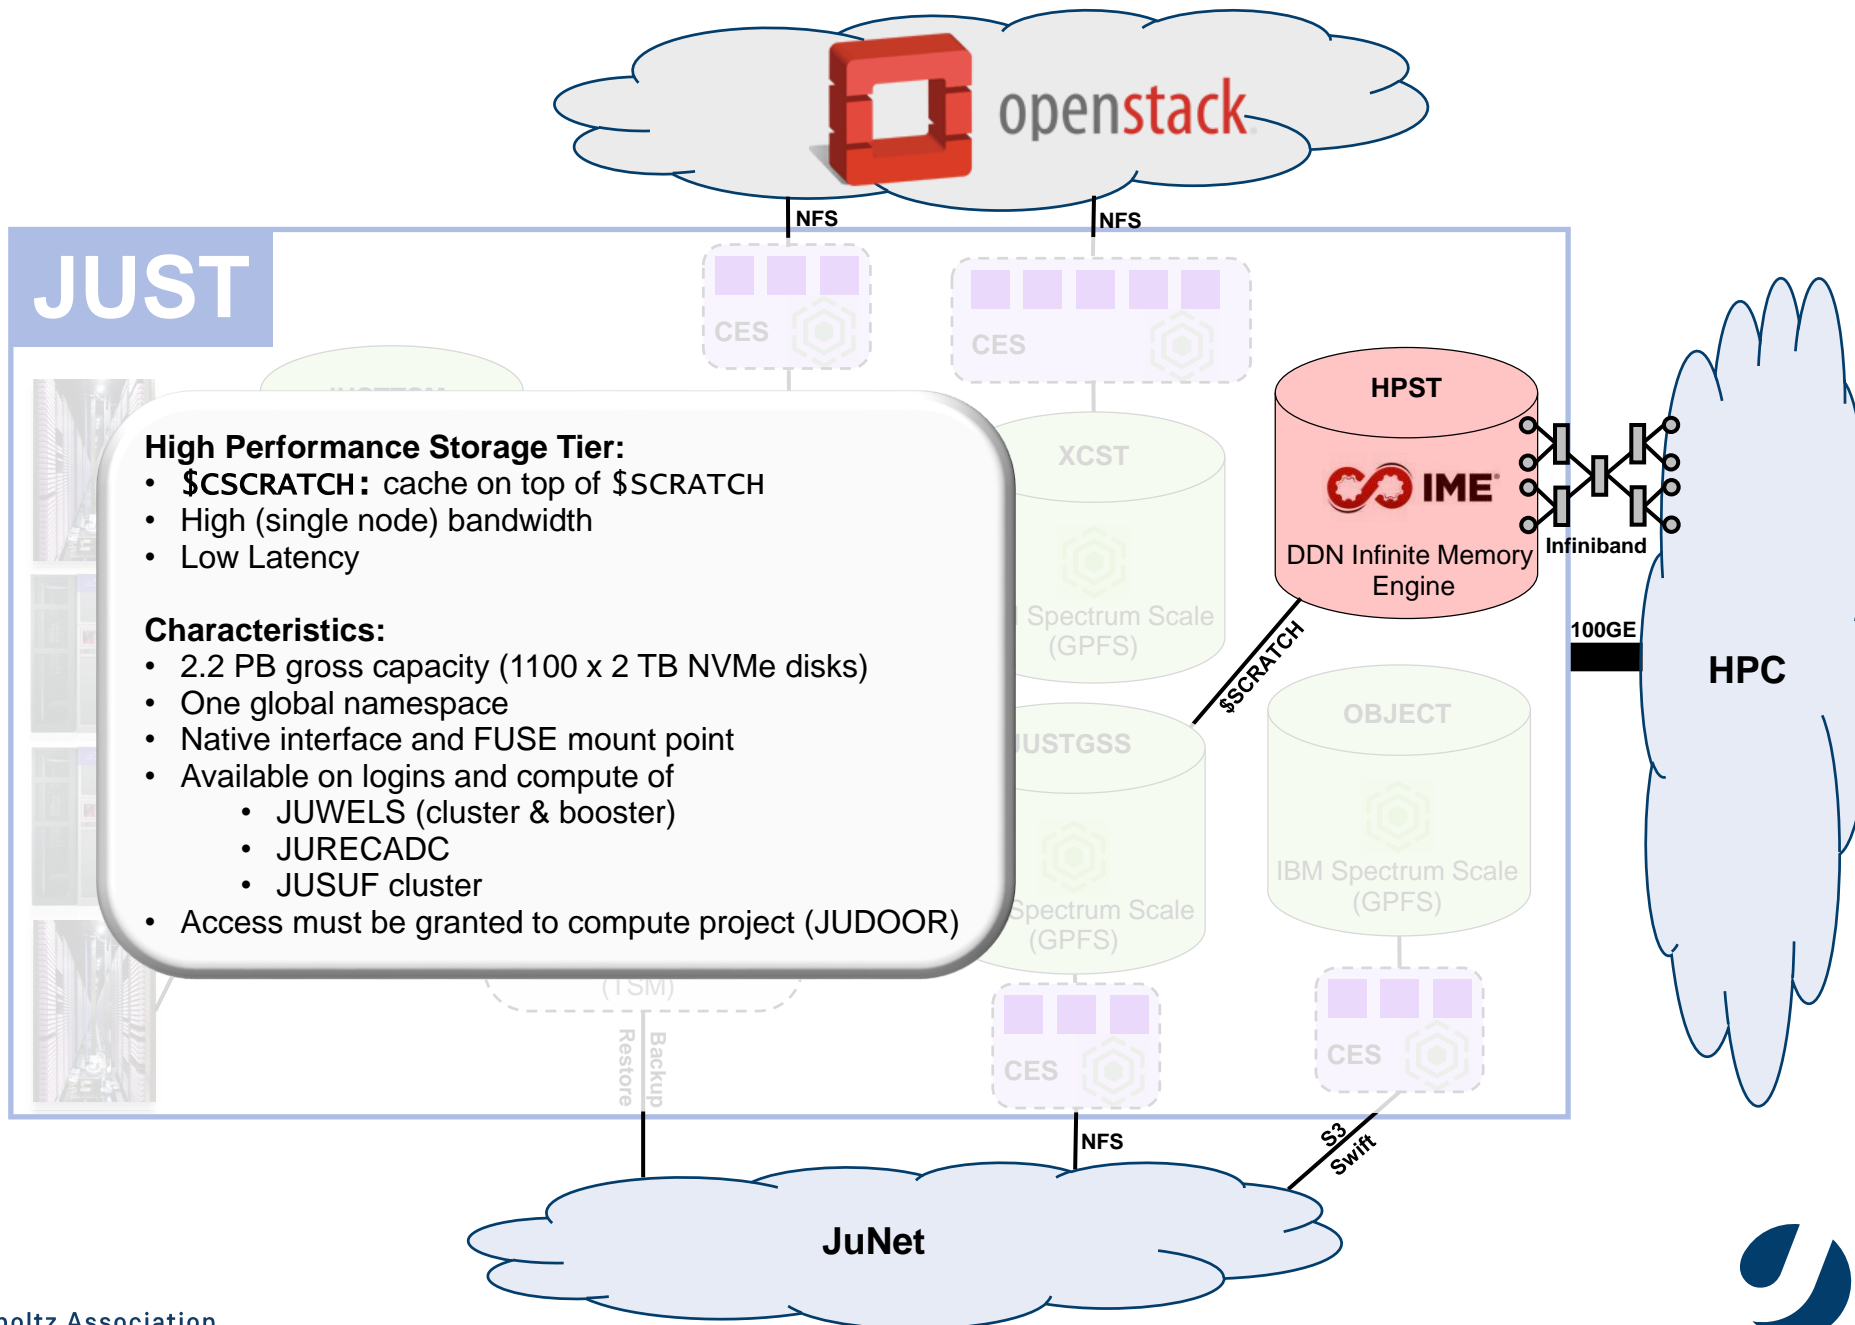




JUST


Internal Storage:

- Block storage provider for openstack VMs
- Container images
- various others (minor) storage consumers

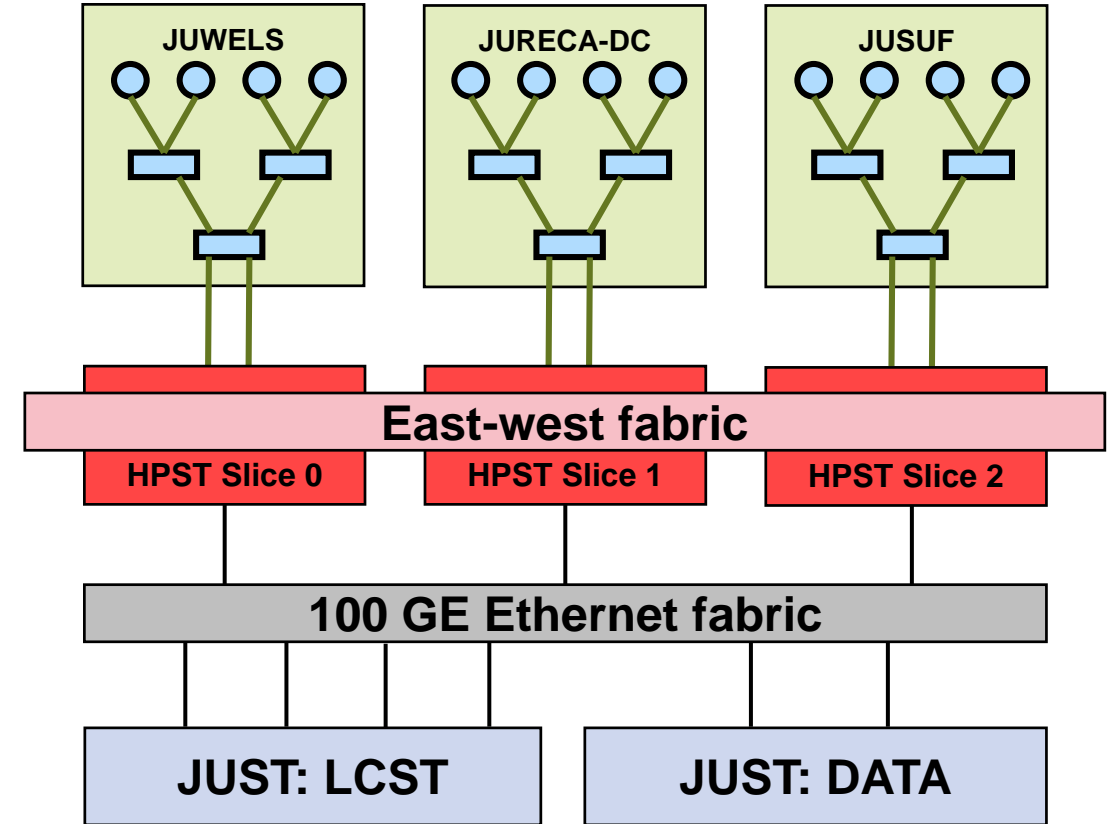



HIGH PERFORMANCE STORAGE TIER

Usage scenarios




1. Read file using HPST

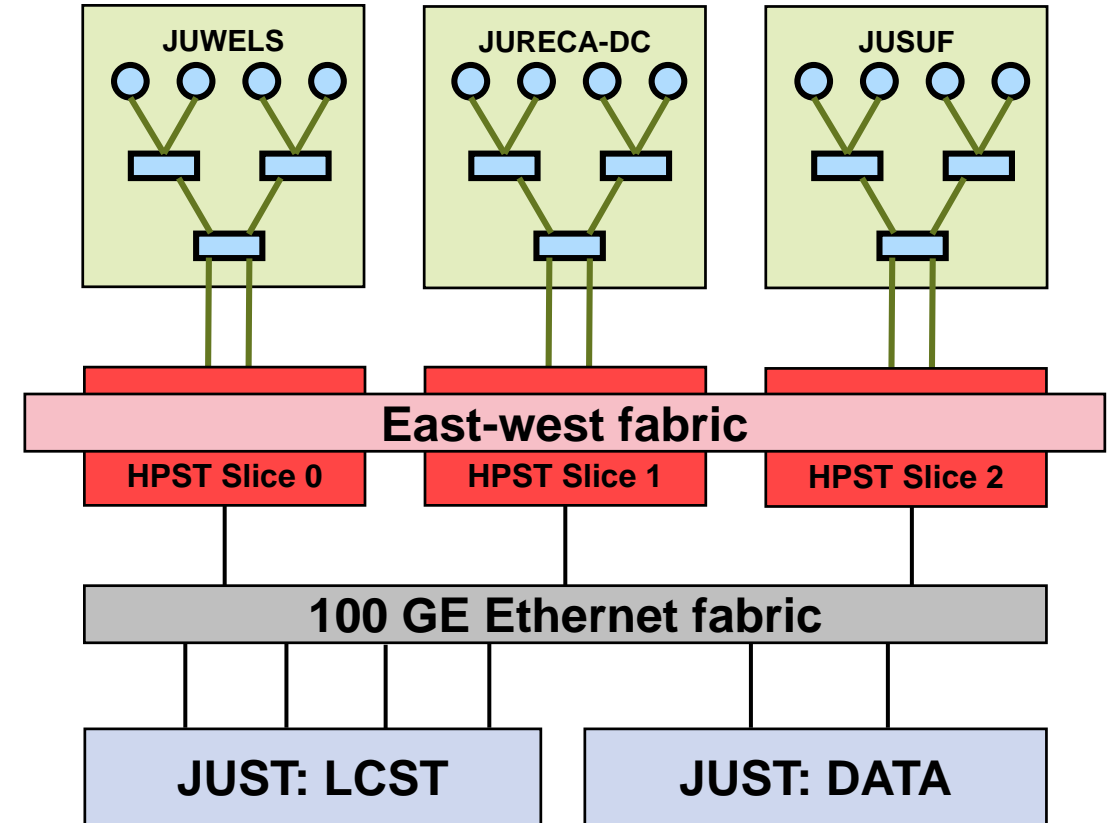



HIGH PERFORMANCE STORAGE TIER

Usage scenarios




2. Write new file to HPST

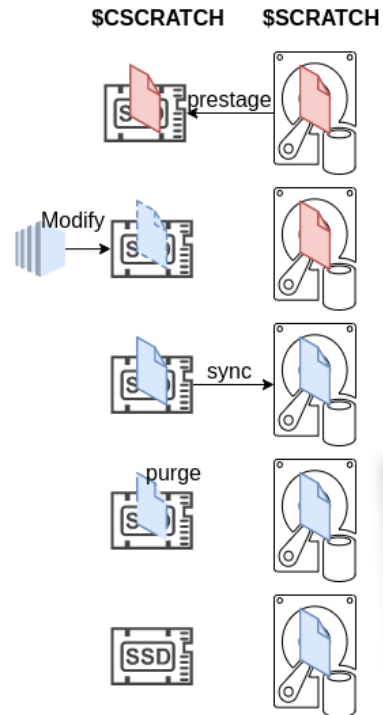


HIGH PERFORMANCE STORAGE TIER

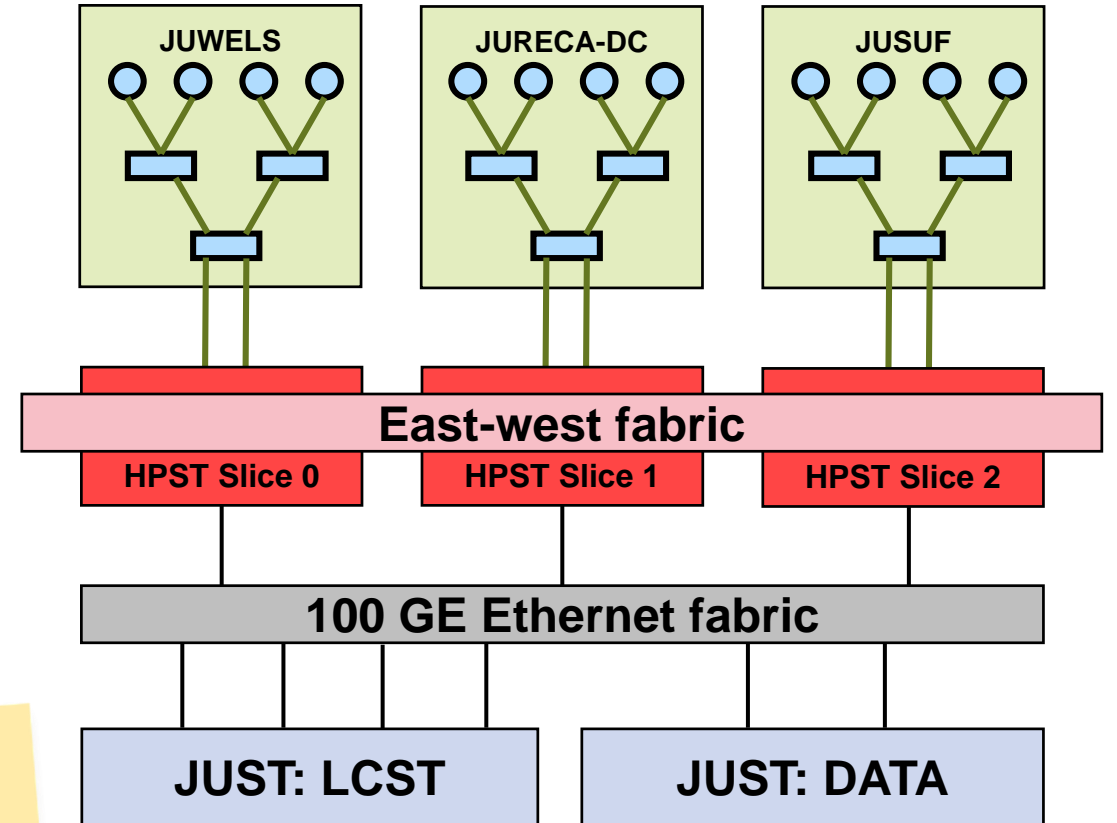
Usage scenarios



3. Read and modify file using HPST

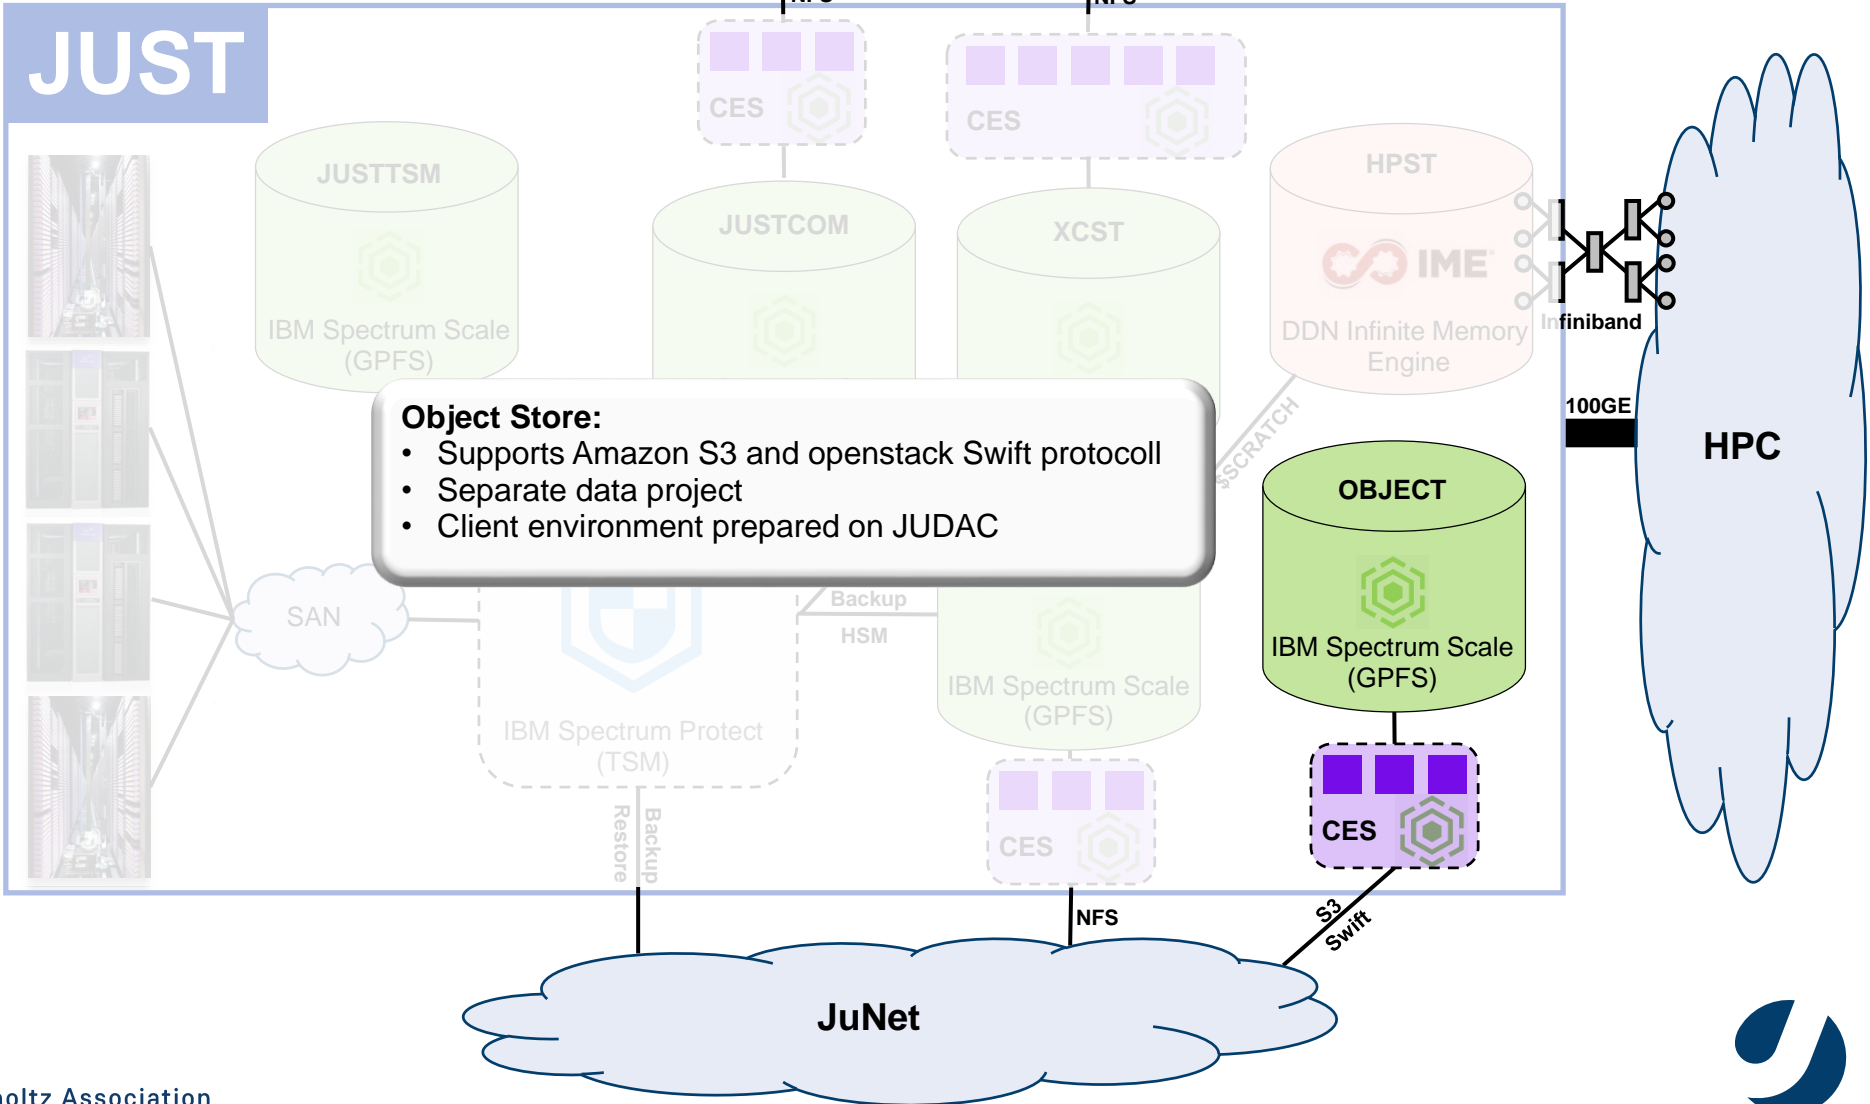


**In preparation:
IO statistics
included in job reporting**

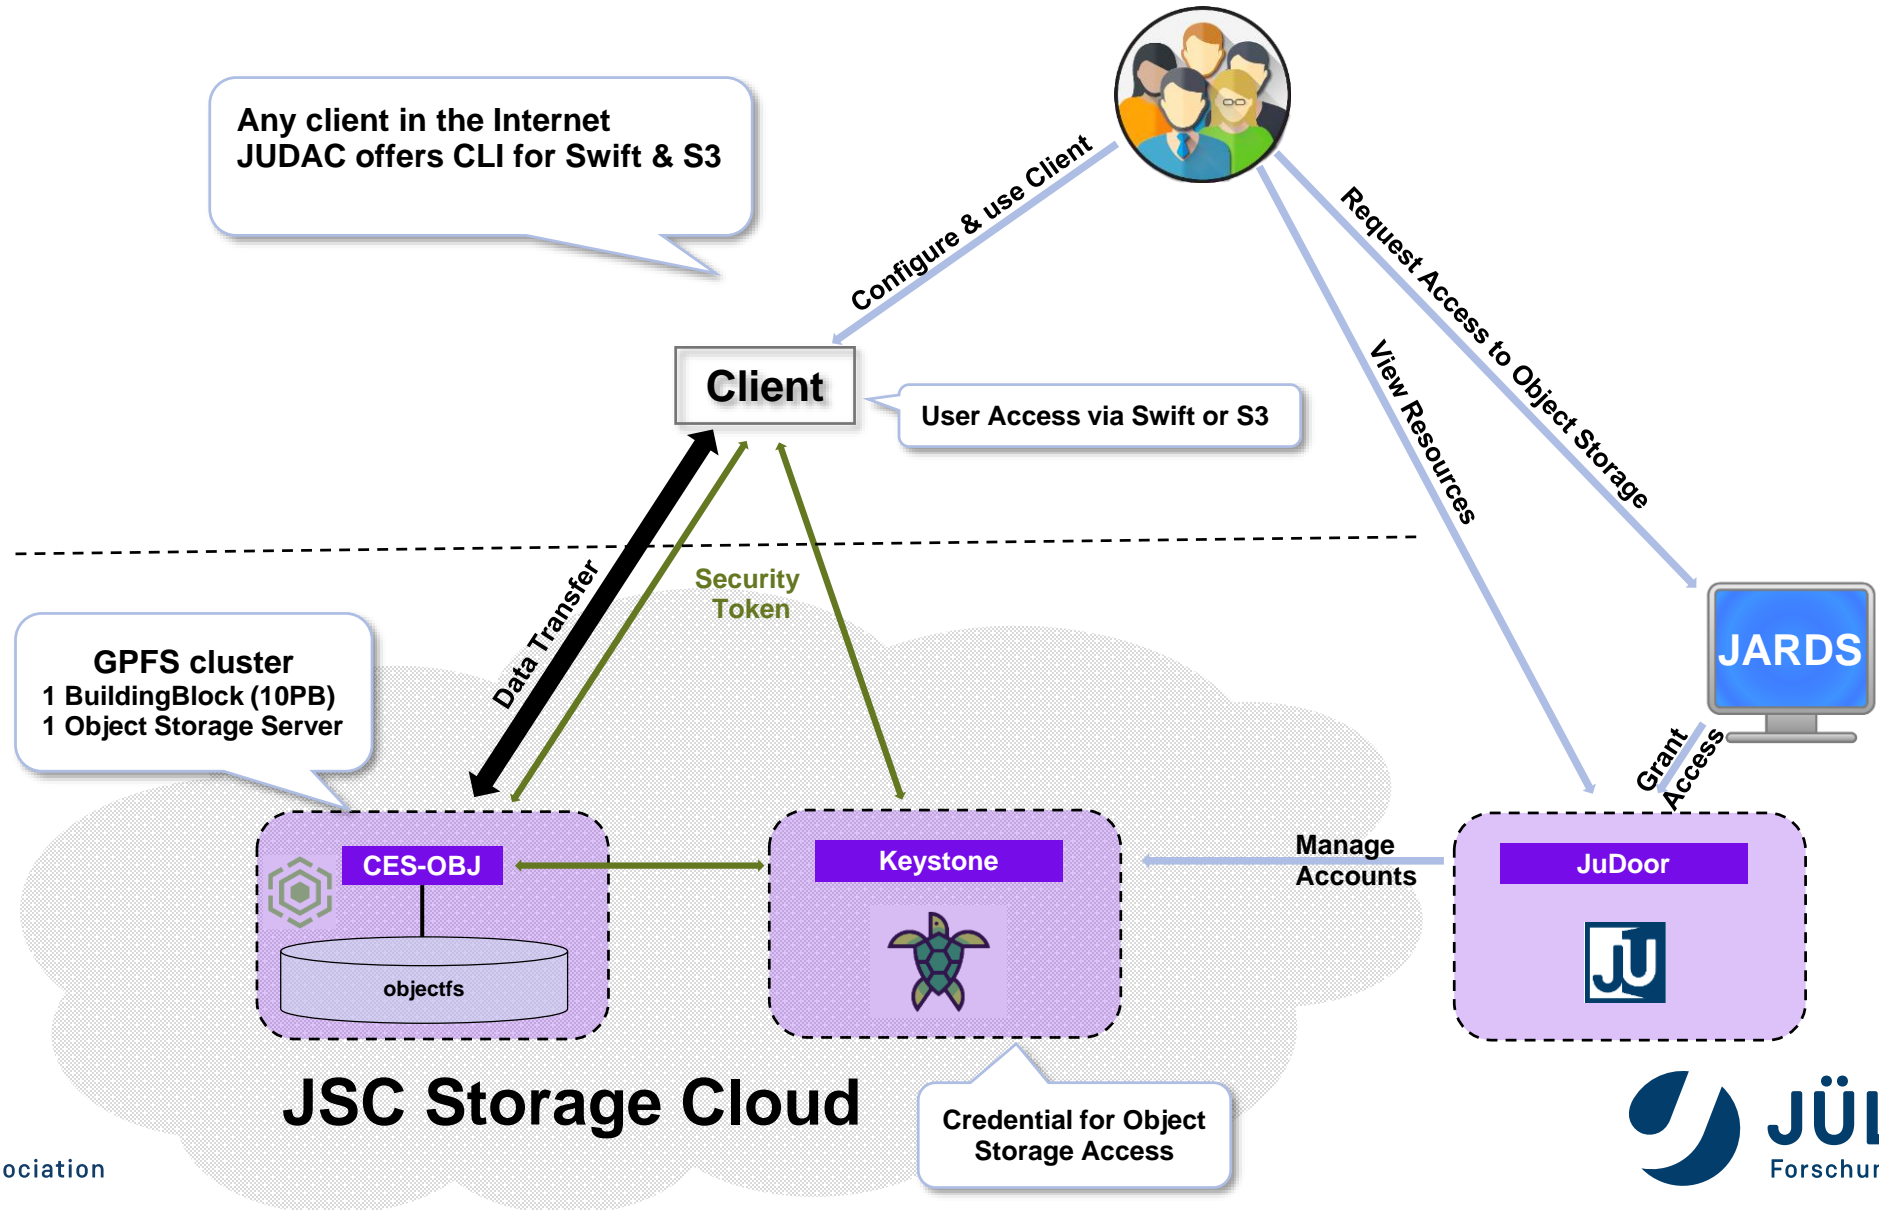




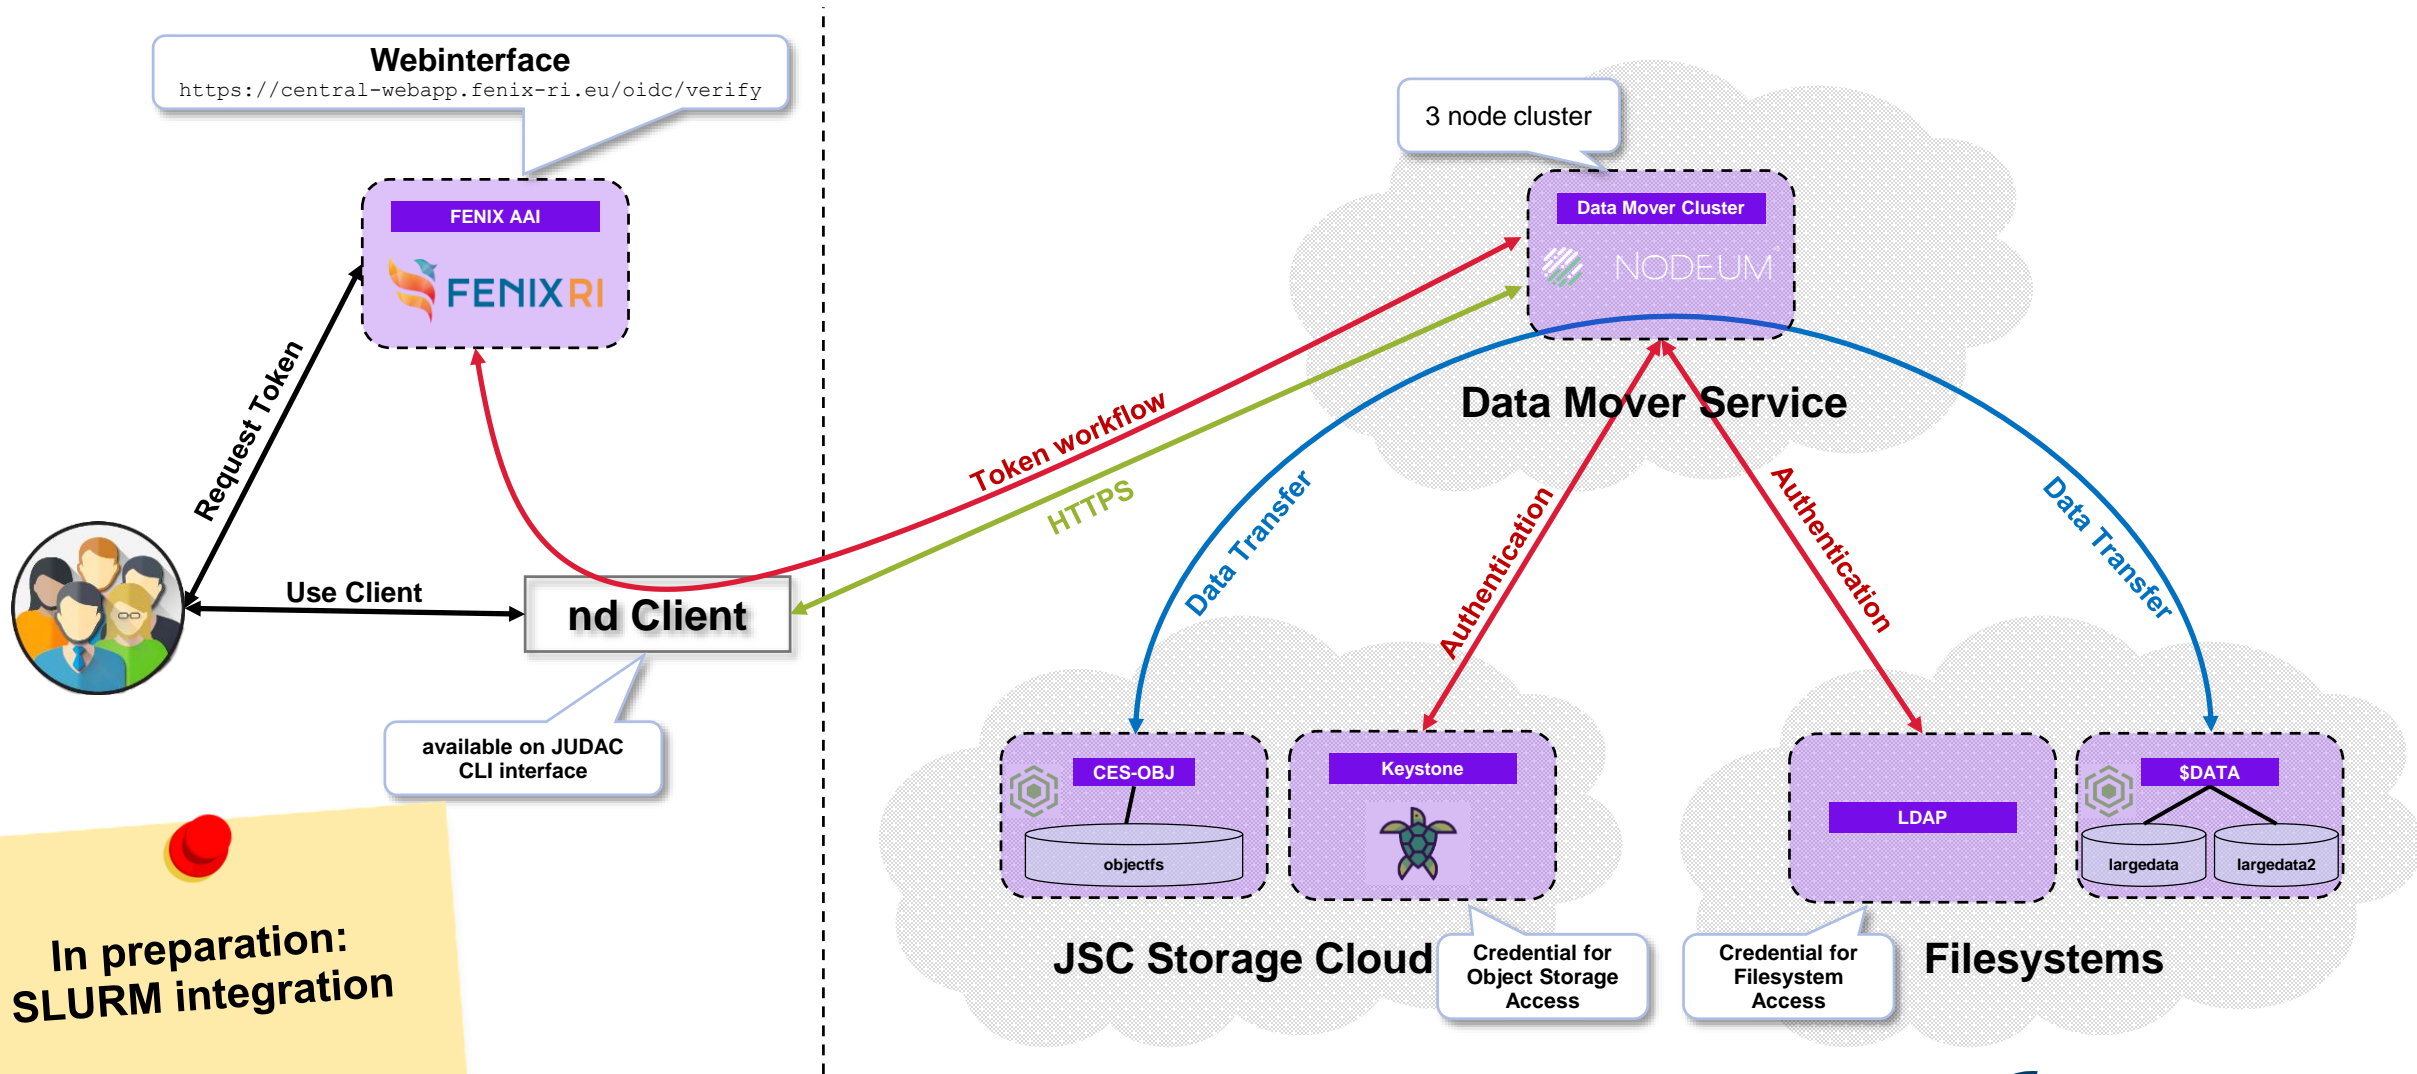
JUST



OBJECT STORAGE SERVICE



JSC DATA MOVER SERVICE




**In preparation:
SLURM integration**

OUTLOOK

- Exascale System JUPITER
 - „Joint Undertaking Pioneer for Innovative and Transformative Exascale Research“
 - Procured by EuroHPC, hosting site Jülich
 - Modular Supercomputer Architecture (MSA)
 - Includes FLASH module as \$SCRATCH
 - Contract signed?
- JUST & Exastore procurement started
 - JUST 6 as the successor of JUST 5
 - Exastore is in-kind contribution of Jülich to JUPITER
 - „One Exabyte for one Exascale“ (spinning disk & Tape)





THANK YOU FOR YOUR ATTENTION

[HTTPS://GO.FZJ.DE/JUDAC](https://go.fzj.de/judac)