

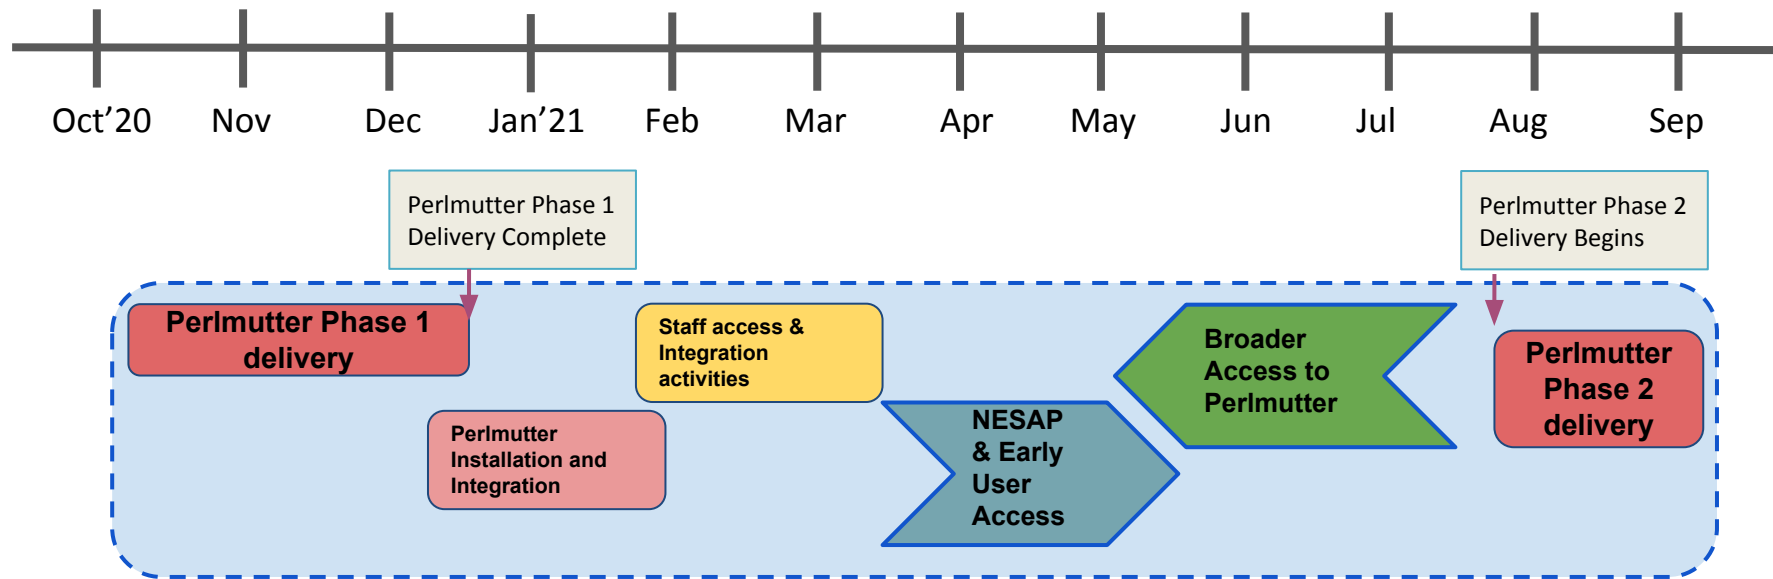
Perlmutter System Integration Overview



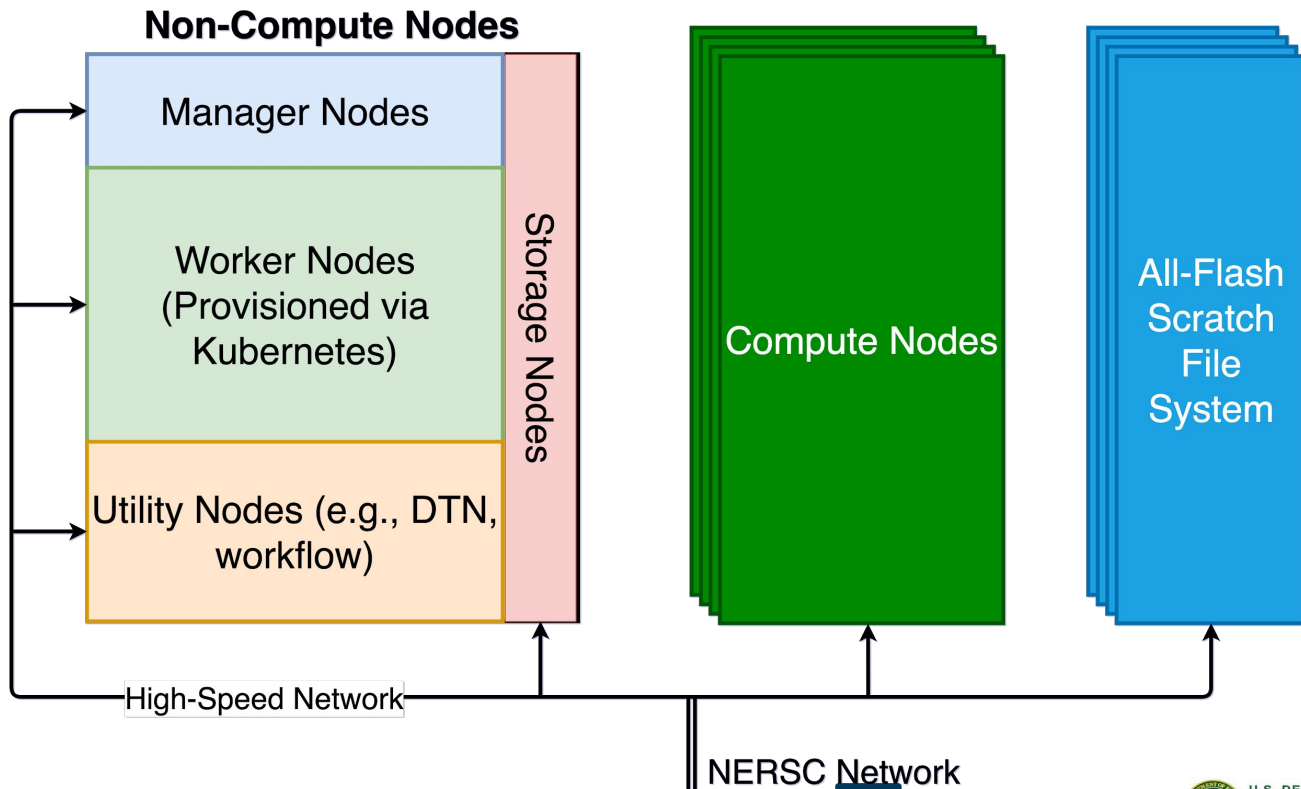
Douglas Jacobsen
Computational Systems Group

August 17, 2020

Perlmutter Phased Timeline




Perlmutter Architecture: Conceptual Overview



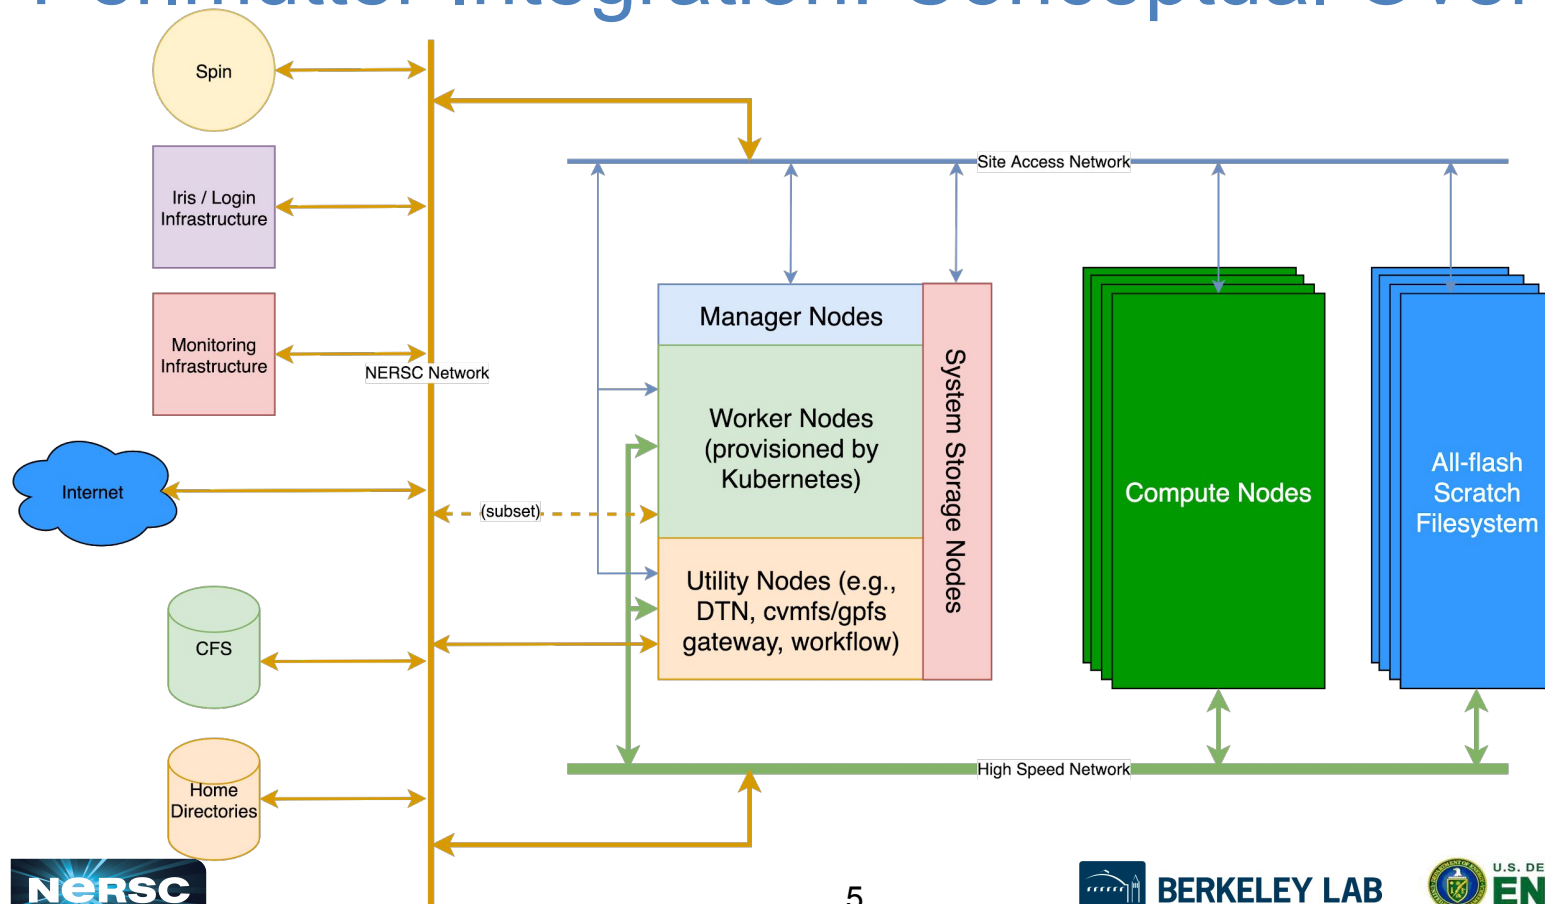
Perlmutter: Physical Integration

After moving cabinets into place, first step is physically connecting cabinets to power and water.

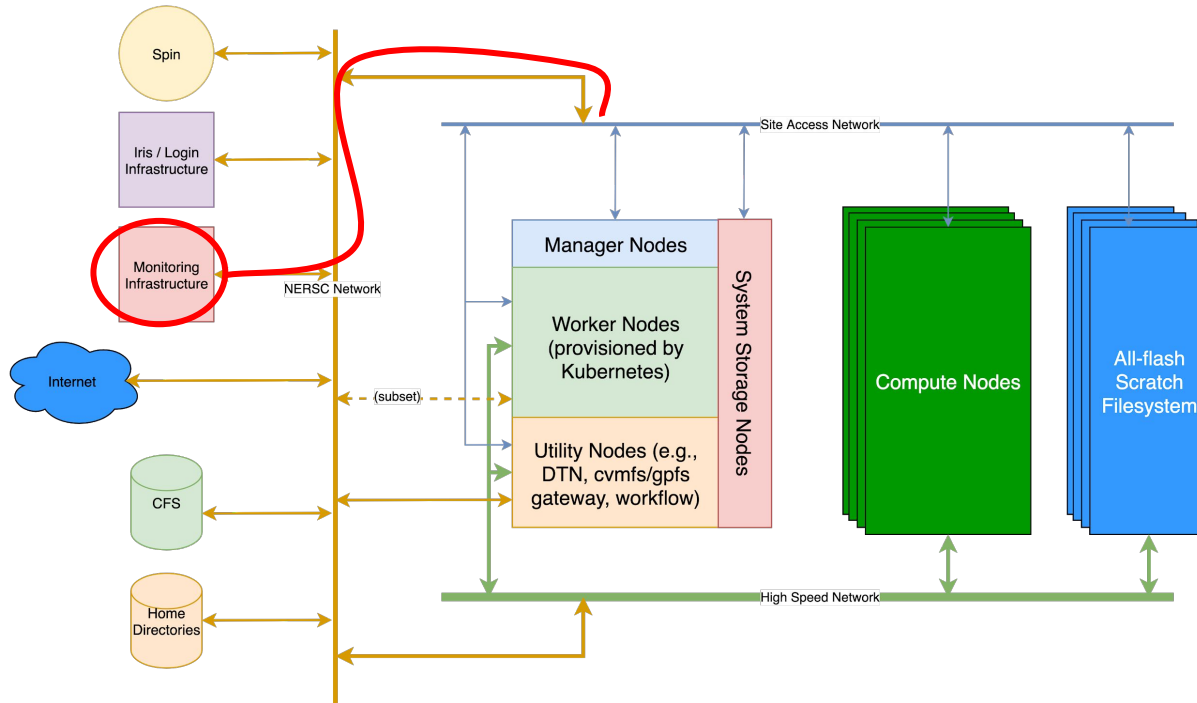
Then connections to the many networks that the system and NERSC need!



Perlmutter Integration: Conceptual Overview



Monitoring Integration



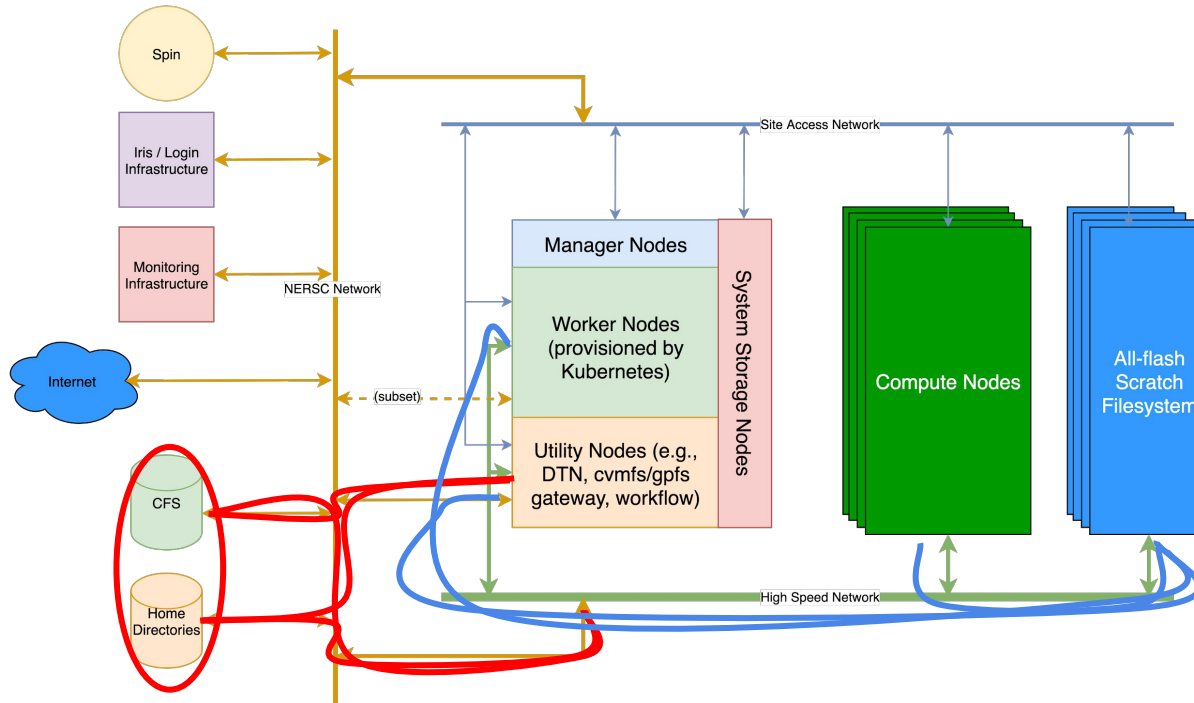
Work is starting now on our development system!

- Failure alerting
- LDMS metrics (e.g., job / usage statistics)
- Environmental Integration

This is a critical requirement for bringing perlmutter as a production-ready resource.

Excellent monitoring needed for system acceptance and on-call support.

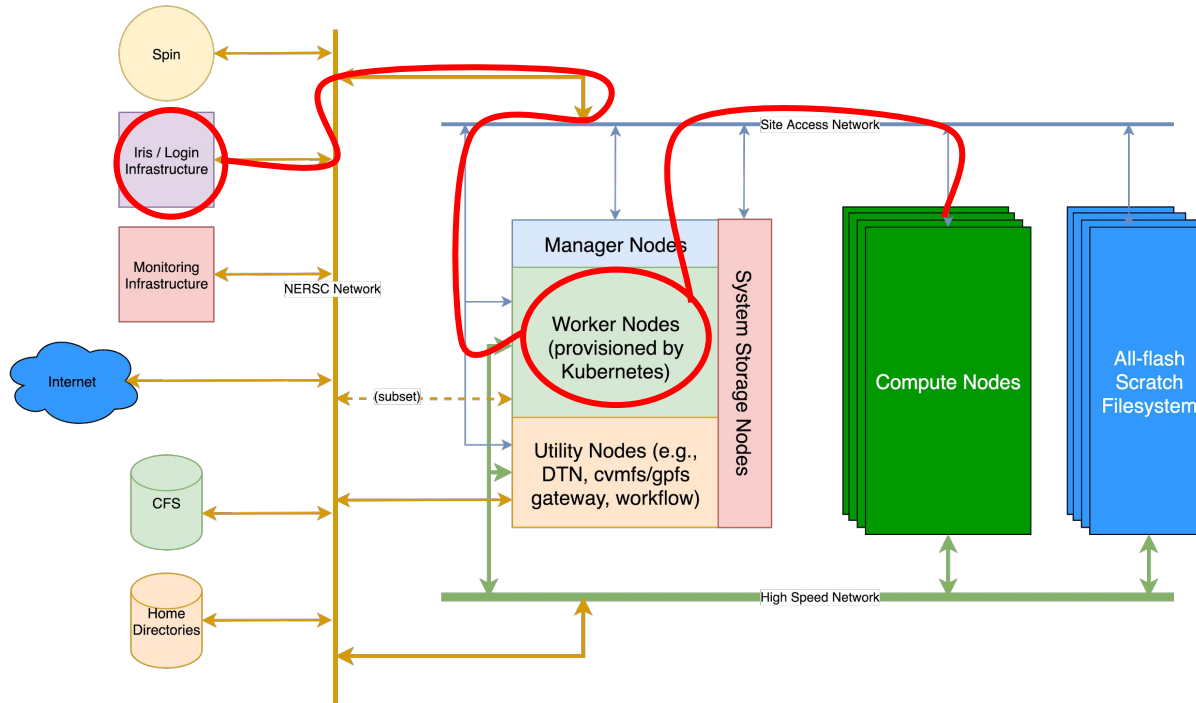
Filesystem Integration



Filesystems

- Lustre scratch is directly attached to the system via HSN
- GPFS (CFS/Homedirs) have multiple possibilities for integration
- Will be able to test and integrate the different options, even change over time!

Slurm Integration



Slurm controller needs access to:

- user login information
- your project limits
- account information
- node status

Slurm controller runs in Kubernetes

Slurm node management daemons and job daemons run as usual.

Conclusions

The perlmutter system integration involves bringing all aspects of NERSC onto the machine.

- The updated network and system software introduce new challenges and new opportunities
 - Expect more options to deliver integration solutions optimized for both performance and reliability
 - Integrating features will be a phased approach, in some ways similar to the physical system integration



Questions?