Benefit and fun providing system details

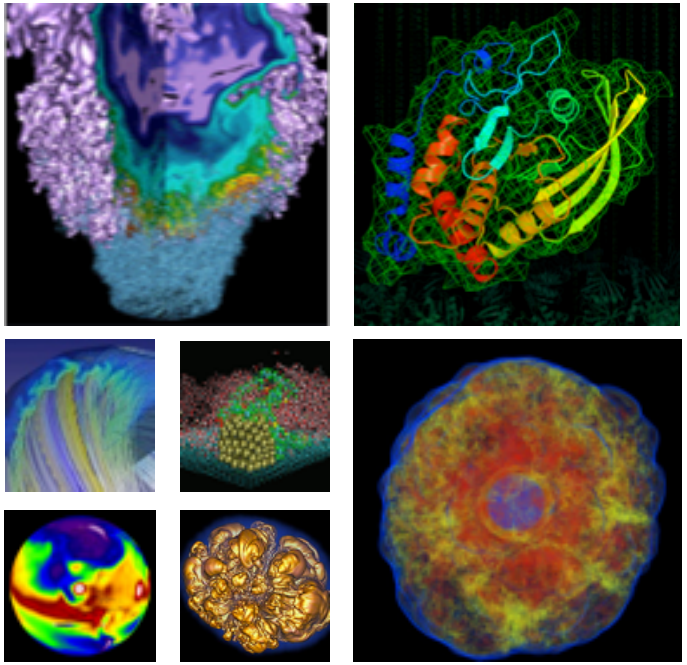

Glenn K. Lockwood, Ph.D.
Advanced Technologies Group

November 14, 2018

NERSC-9 and NERSC Center in 2020

- **CPU-only and CPU+GPU nodes**
- **30 PB, all-flash parallel file system**
- **200 PB, disk-based Community File System**
- **Cray *Slingshot* network directly connects to Ethernet network**

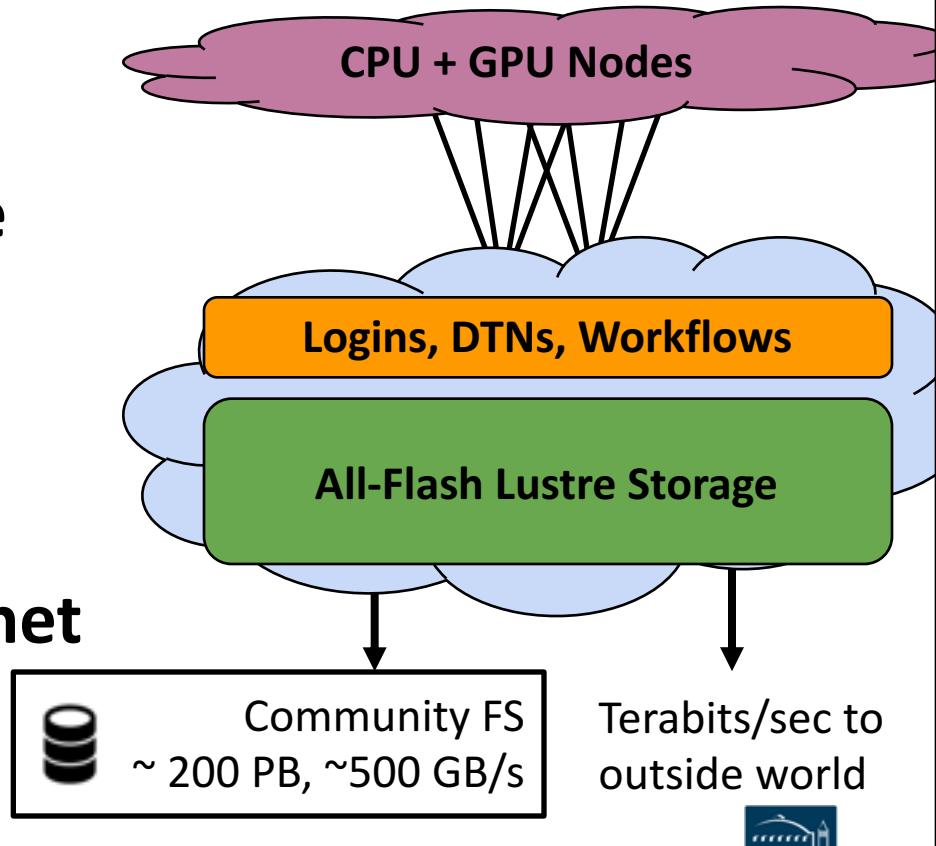

NERSC-9 and NERSC Center in 2020

- **But how do we rigorously describe this architecture?**
  - How are nodes connected to each other? Storage? External networks?
  - What are the data center facility requirements to host this all?
- **CDCL to the rescue!**

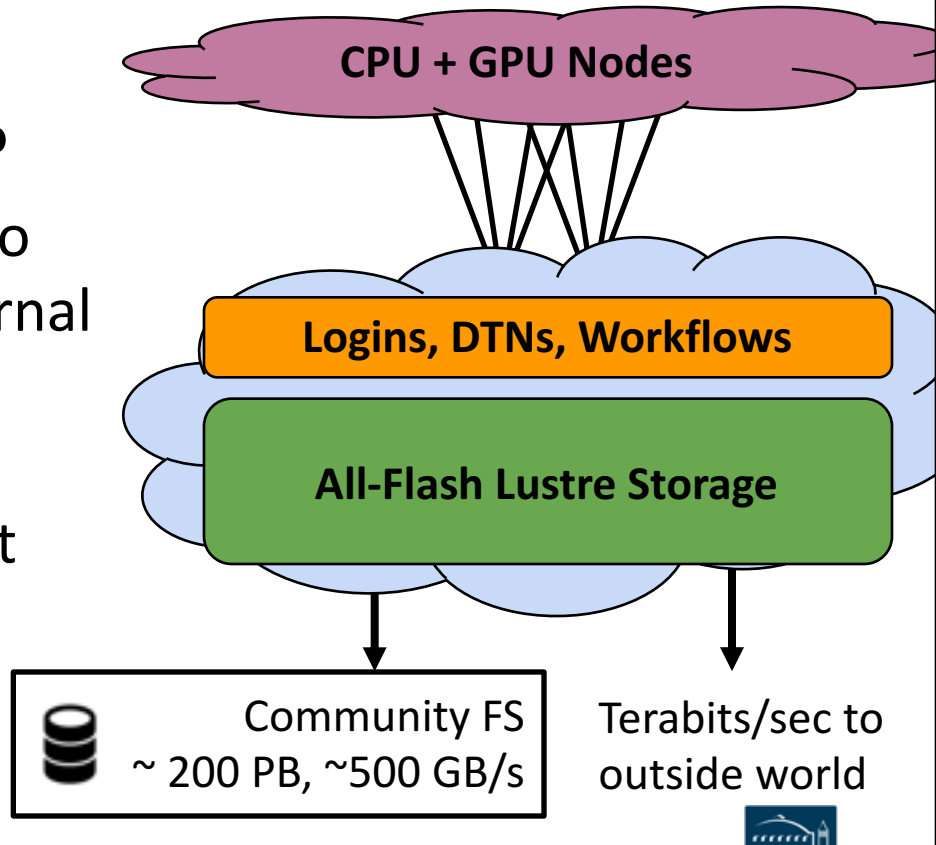

Comprehensive Data Center List (CDCL)

- **Like Top500 but without competitive angle**
- **Who should I ask about...**
  - scaling Lustre?
  - deploying tape?
  - tiering storage?
- **Your center may already be on it!**

The current list

#	site.institution	site.storage system.net capacity	site.supercomputer peak
		in PiB	*in PFLOPS*
1	Oak Ridge National Laboratory	250.04	
2	National Energy Research Scientific Computing Center	197.65	
3	Los Alamos National Laboratory	72.83	
4	German Climate Computing Center	52.00	
5	Lawrence Livermore National Laboratory	48.85	
6	RIKEN Advanced Institute for Computational Science	39.77	

https://www.vi4io.org/hpsl/start

NERSC on the CDCL

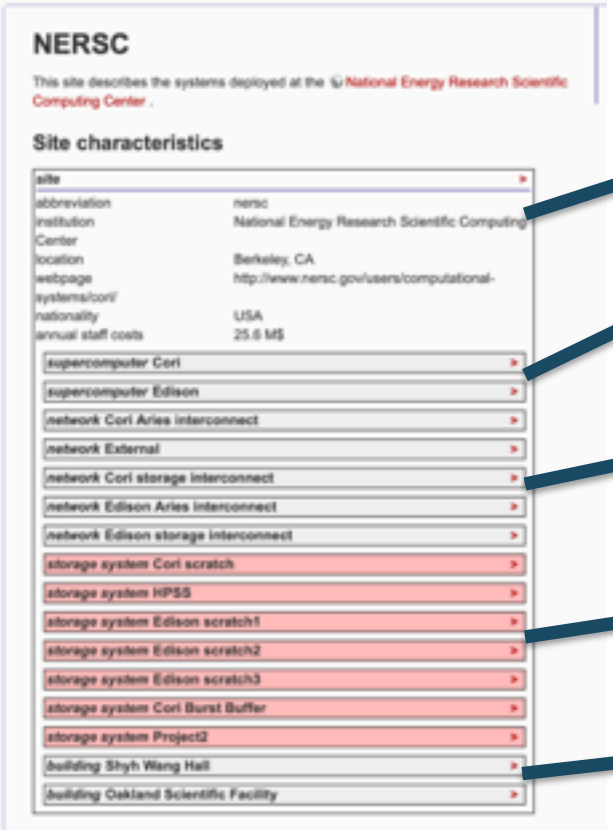

Center-wide metadata

Two large HPC systems

Four major networks

Seven major storage systems

Two buildings

Drill-down into details

NERSC occupies two data centers

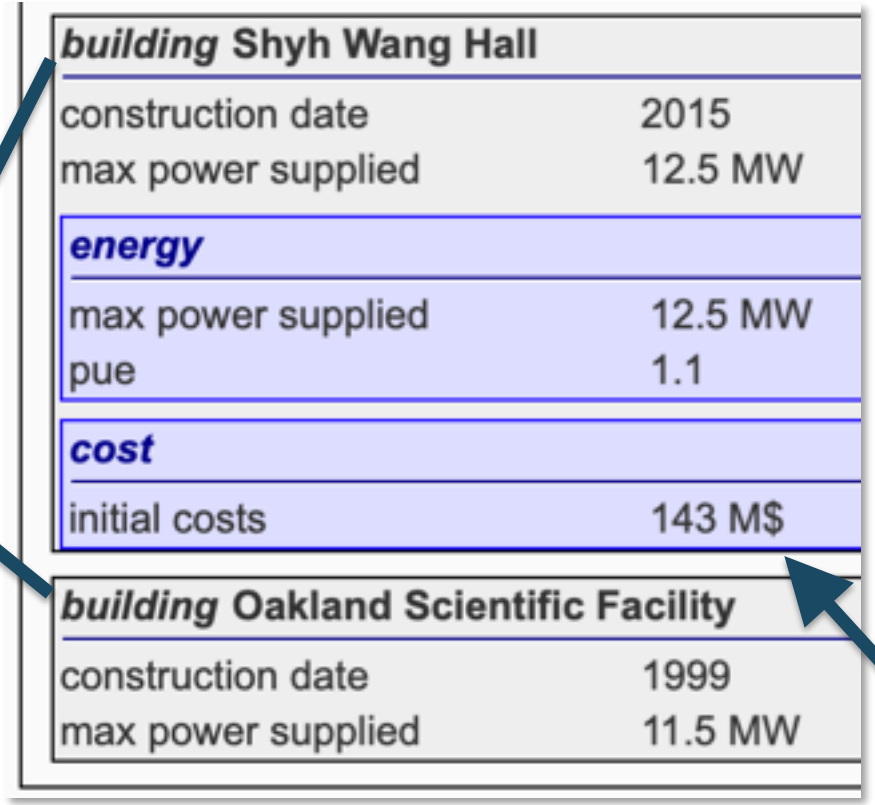

building Shyh Wang Hall	
construction date	2015
max power supplied	12.5 MW
energy	
max power supplied	12.5 MW
pue	1.1
cost	
initial costs	143 M$

building Oakland Scientific Facility	
construction date	1999
max power supplied	11.5 MW

- Green data center
- On a hillside
- Over a fault zone
- Not cheap!

CDCL Network Graph

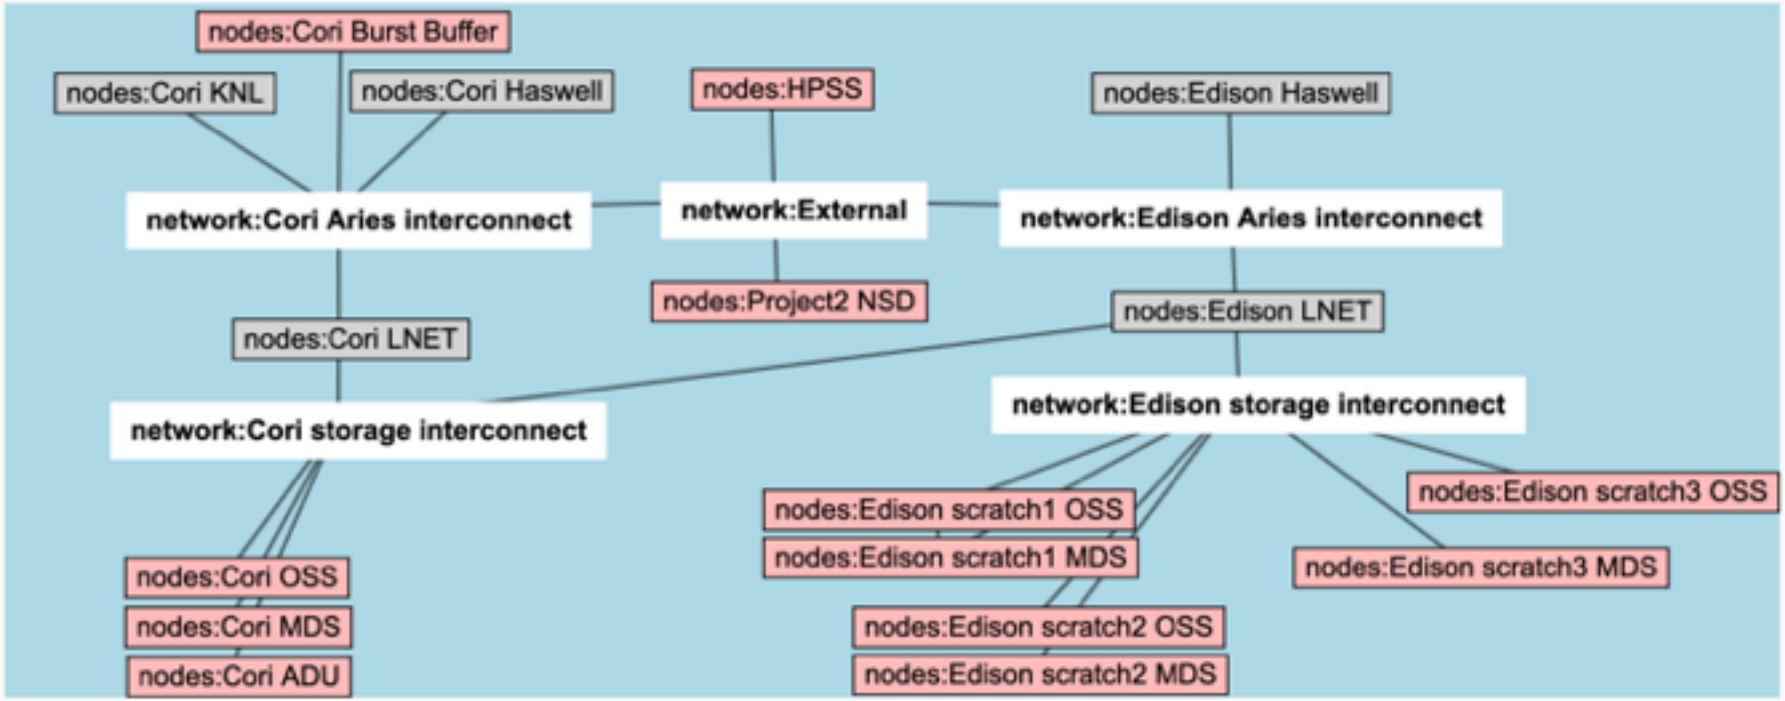

CDCL Building Graph

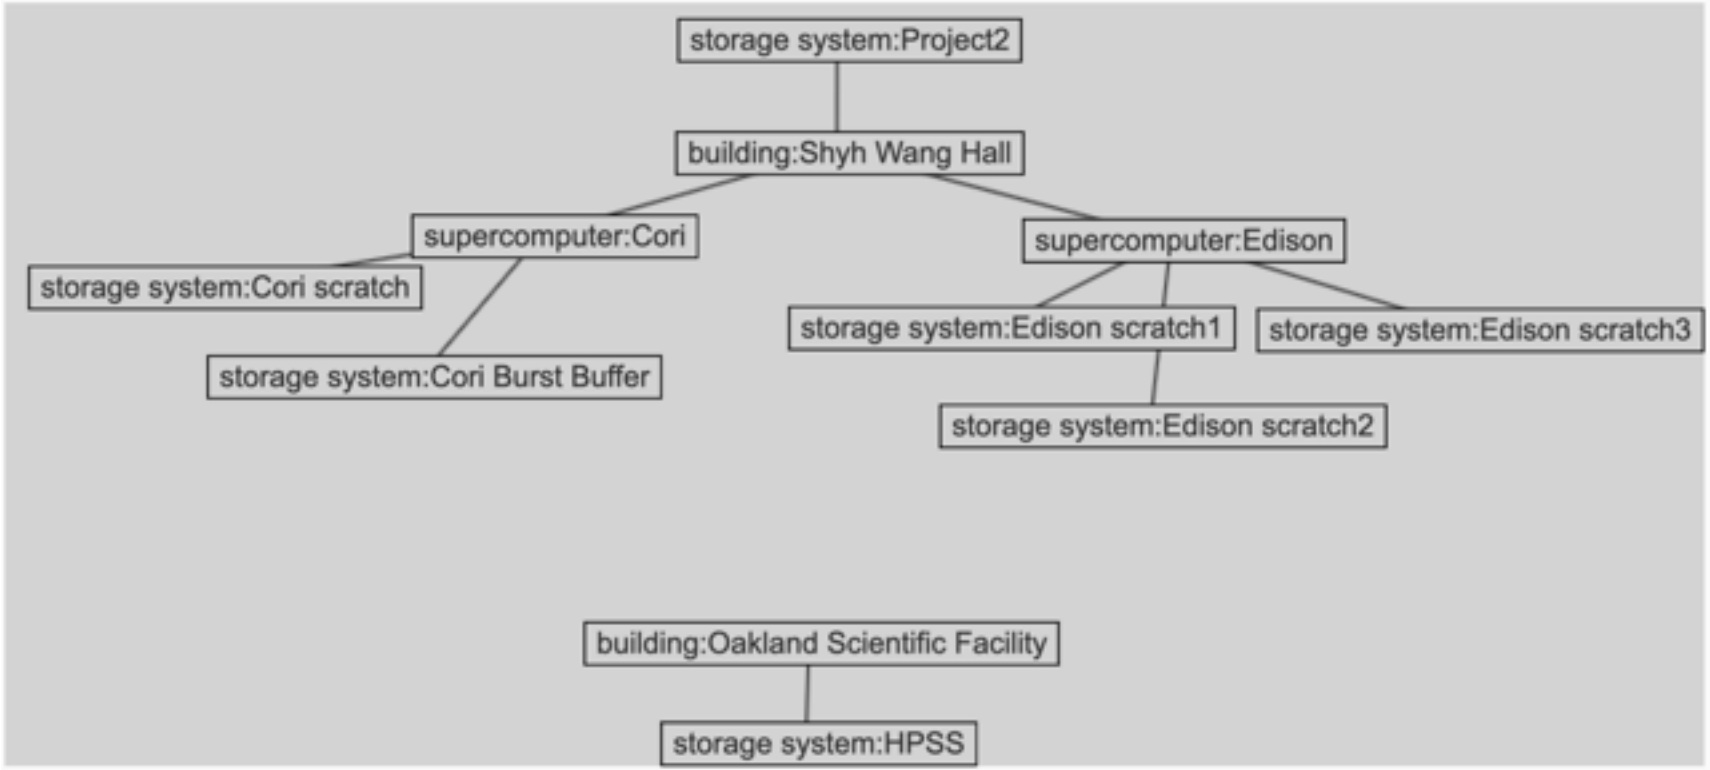

CDCL

- **Point-n-click interface to describe your data center**
- **JSON-based schema under the hood**
- **Complete description data center features: HPC, storage, infrastructure**
- **Identify centers facing the same scaling issues**
  - Connect with peers using the same technologies as you
  - or experts who have already hit the issues coming for you