


# NERSC Data Ecosystem Overview

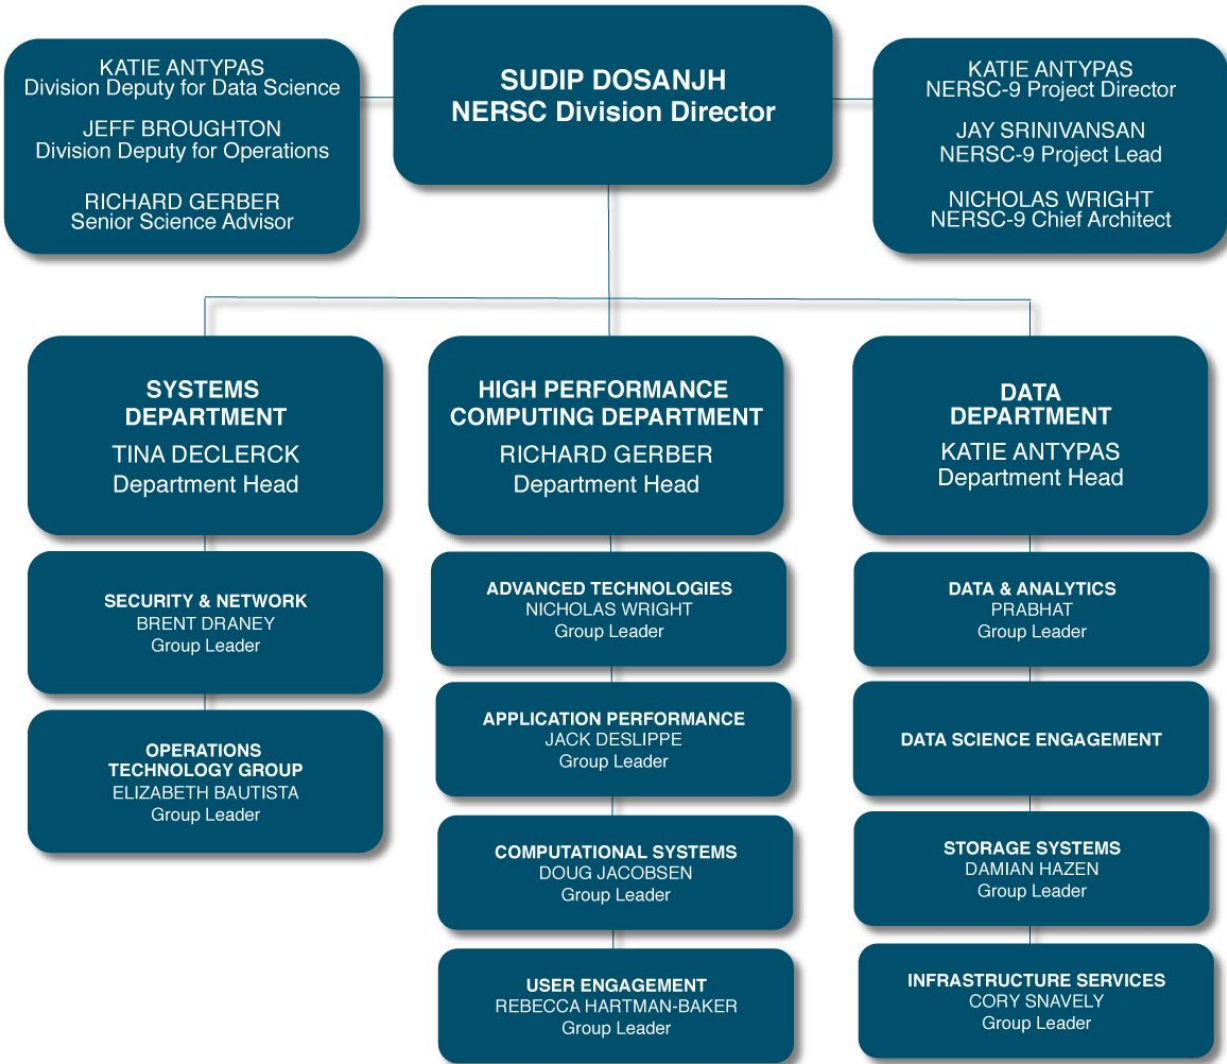


**Prabhat**  
**Data & Analytics Services Group**

# Afternoon Schedule



<b>1:40 pm</b>	<b>Data Transfer</b>	<b>Shreyas Cholia</b>
<b>2:10 pm</b>	<b>File Systems</b>	<b>Jialin Liu</b>
<b>2:30 pm</b>	<b>I/O Best Practices</b>	<b>Quincey Koziol</b>
<b>2:50 pm</b>	<b>Break</b>	
<b>3:10 pm</b>	<b>Python and Jupyter</b>	<b>Rollin Thomas</b>
<b>3:30 pm</b>	<b>Shifter</b>	<b>Rollin Thomas</b>
<b>3:50 pm</b>	<b>Deep Learning</b>	<b>Prabhat</b>
<b>4:10 pm</b>	<b>End</b>	



# Cori Brings HPC & Data Together

NERSC



Gerty Cori: Biochemist and first American woman to win a Nobel Prize in science

Phase I: 2388 x 32-core Intel Xeon “Haswell”

128 GB DDR4

Phase II: 9688 x 68-core Intel Xeon Phi “KNL”

96 GB DDR4 + 16 GB MCDRAM



U.S. DEPARTMENT OF  
**ENERGY**























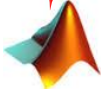




Office of  
Science



BERKELEY LAB  
Lawrence Berkeley National Laboratory

# Production Data Stack



Capabilities	Technologies
<b>Data Transfer + Access</b>	     
<b>Workflows</b>	  
<b>Data Management</b>	     
<b>Data Analytics</b>	         
<b>Data Visualization</b>	 

# Cori's Data-Friendly Features



12 32-core Haswell  
500 GB Login Nodes

 Jupyter  
Notebook  
Node

768 GB "bigmem"  
Haswell Nodes

Pipeline/Workflow  
Management Nodes

Serial Queue  
Shared-node Queue  
Transfer Queue

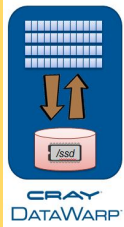
 Containerized  
Environments 

Real-Time Queues  
for Co-Scheduling  
w/Experiments

  
slurm  
workload manager

Streaming Data to  
Compute Nodes

External Network  
Access to/from  
Compute Nodes

Cray DataWarp:  
Burst Buffer for  
I/O acceleration 

Interactive Queue:  
64 Nodes x 4 Hours

- **Please engage with NERSC staff members**
  - Provide feedback, critique
  - Tell us about interesting science problems!

# Questions? Comments?

---

