IXPUG 2018 SPRING MEETING

Argonne NATIONAL LABORATORY

ALCF UPDATE AND EARLY SCIENCE PROGRAM

KATHERINE RILEY
Director of Science
Argonne Leadership Computing Facility

MARK FAHEY
Director of Operations
Argonne Leadership Computing Facility

DAVID MARTIN
Manager, Industry Partnerships and Outreach
Argonne Leadership Computing Facility

THETA – ALCF'S KNL MACHINE

	Per node	18 cabinets June 2016	20 cabinets Jan 2017	24 cabinets Dec 2018
Cores	64	207,360	231,936	281,088
DDR4	192 GB	622,080 GB	695,808 GB	843,264 GB
HPL			5.88 PF	6.92 PF
HPL % peak			61.1%	59.2%
MCDRAM	16 GB	51,840 GB	57,984 GB	70,272 GB
Nodes		3,240	3,624	4,392
Peak		8.626 PF	9.648 PF	11.693 PF
SSD	128 GB	414,720 GB	463,872 GB	562,176 GB

NEW: NODE LOCAL SSDS AVAILABLE

- ALCF has enabled node local SSDs to provide temporary scratch storage for caching and processing workloads during job runs.
  - If you'd like to use this feature, project PIs can request access by contacting support@alcf.anl.gov. A use case will need to be provided.
  - For more information, visit: https://www.alcf.anl.gov/user-guides/xc40-file-systems and https://www.alcf.anl.gov/user-guides/running-jobs-xc40#requesting-local-ssd-requirements

NEW: SINGULARITY AVAILABLE

- Singularity 2.4.2, a virtualization framework, has been installed on Theta.
  - With Singularity, users can take their complete application environment and reproducibly run it anywhere Singularity is installed.
  - For information on using Singularity at the ALCF, visit: https://www.alcf.anl.gov/user-guides/singularity
- ALCF will create some known base images to help users
- Working with users to build more configurations

ALCF APPLICATIONS READINESS: EARLY SCIENCE PROGRAM

- Ten Years of Early Science have given us
  - Breakthrough science
  - Technical reports on code porting & tuning
  - Open community workshops (science & technology)
  - Synergy with Tools & Libraries project
  - Stable production platform (problems shaken out)
  - Persisting culture of apps readiness for next generation
  - Success stories for postdocs

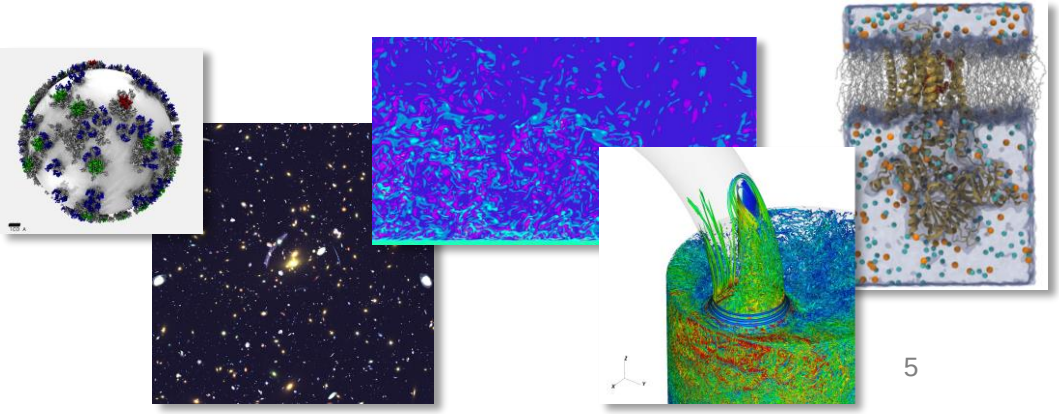

ALCF EARLY SCIENCE PROGRAM (ESP)

Applications Readiness

- Prepare applications for next-gen system:
  - Architecture
  - Scale
- ~Two year lead time

Proposals

- Ambitious targeted science calculation
- Parallel performance
- Development needed
- Team

Support

PEOPLE

- Funded ALCF postdoc
- Catalyst staff member support
- Vendor experts

TRAINING

- Training on HW and programming
- Community workshop to share lessons learned

COMPUTE RESOURCES

- Current ALCF systems
- Early next-gen hardware & simulators
- 3 months dedicated Early Science access
  - Pre-production (post-acceptance)
  - Large time allocation
  - Continued access for rest of year

CALL FOR PROPOSALS: A21 ESP DATA, LEARNING PROJECTS

- **CFP 10 January 2018**
  - Deadline 8 April 2018
- Selections June 2018
  - 5 Data projects
  - 5 Learning projects
- Two-year funded ALCF postdoc
- Cross-cutting proposals targeting the convergence of simulation, data and learning are very much encouraged.

DATA

- Experimental/observational data
  - Image analysis
  - Multidimensional structure discovery
- Complex and interactive workflows
- On-demand HPC
- Persistent data techniques
  - Object store
  - Databases
- Streaming/real-time data
- Uncertainty quantification
- Statistical methods
- Graph analytics

LEARNING

- Deep learning
- Machine learning steering simulations
  - Parameter scans
  - Materials design
  - Observational signatures
- Data-driven models and refinement for science using ML/DL
- Hyperparameter optimization
- Pattern recognition
- Reduced model derivation
- Bridging gaps in theory

AURORA 2021 ESP TIMELINE (NOTIONAL)

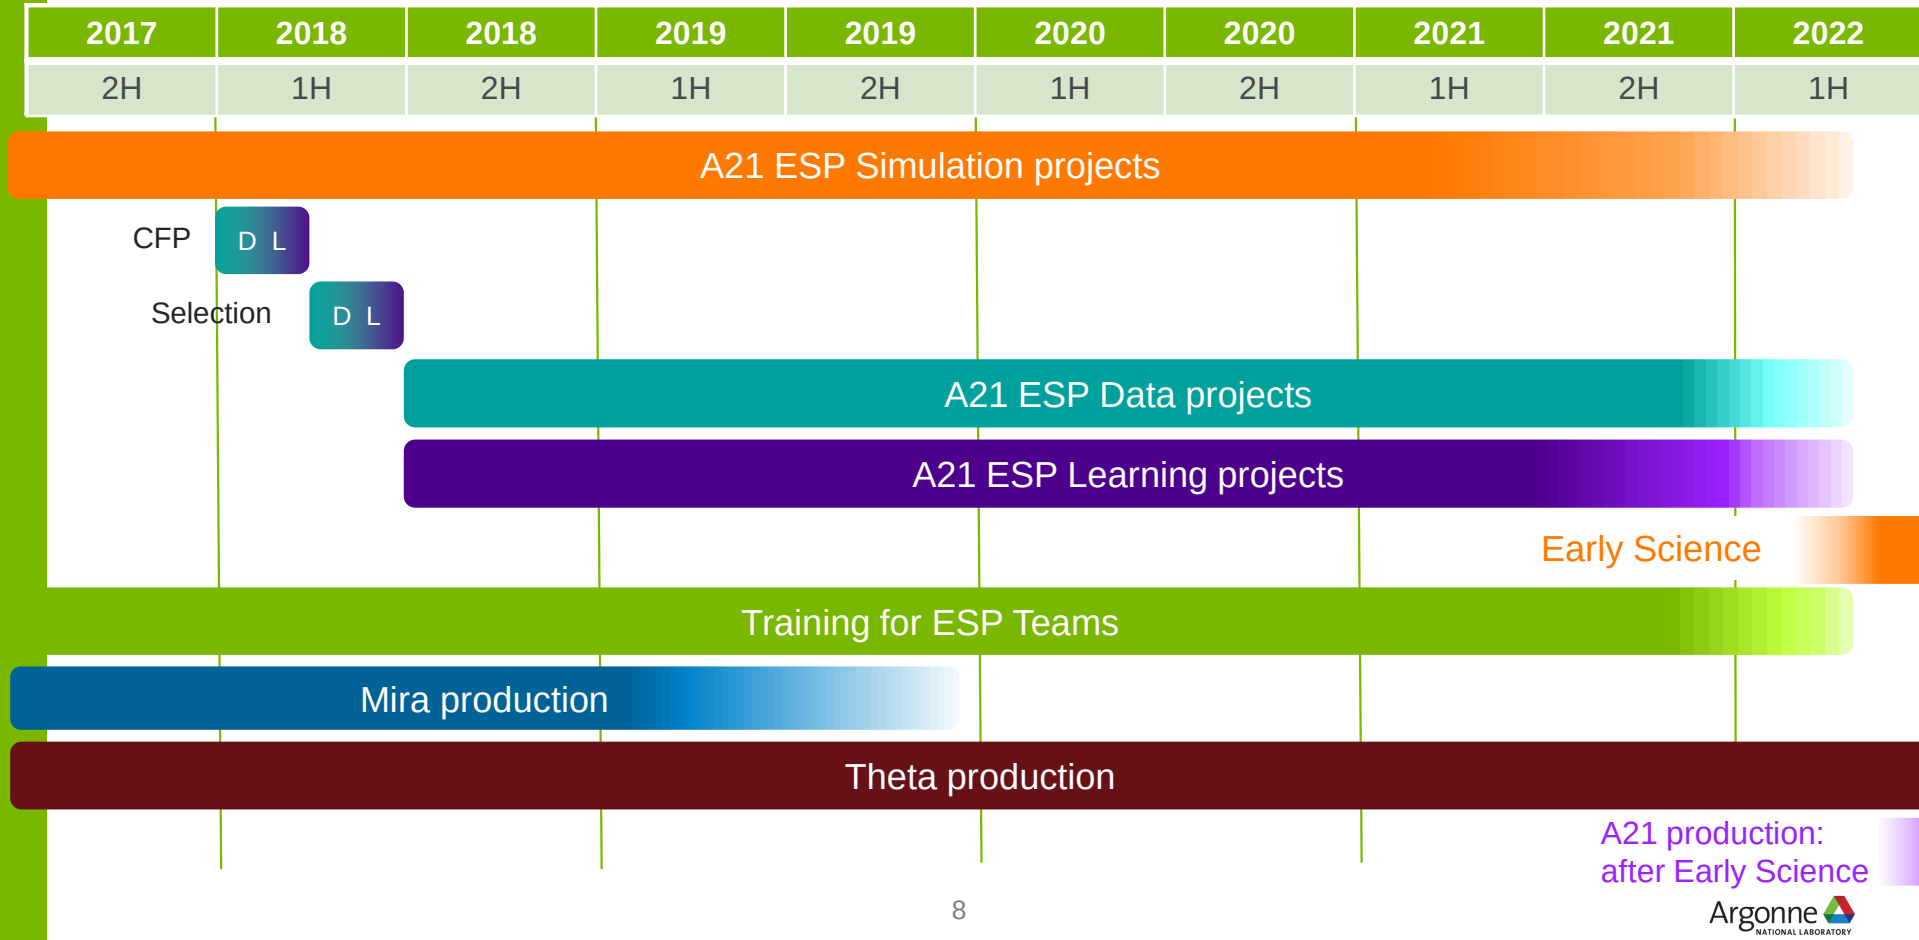

Argonne NATIONAL LABORATORY

THANK YOU