Updates from the OLCF

Ashley Barker, Tjerk Straatsma, Tom Papatheodore, Jack Wells, Reuben Budiardja, Stephen McNally, Arjun Shankar and Justin Whitt

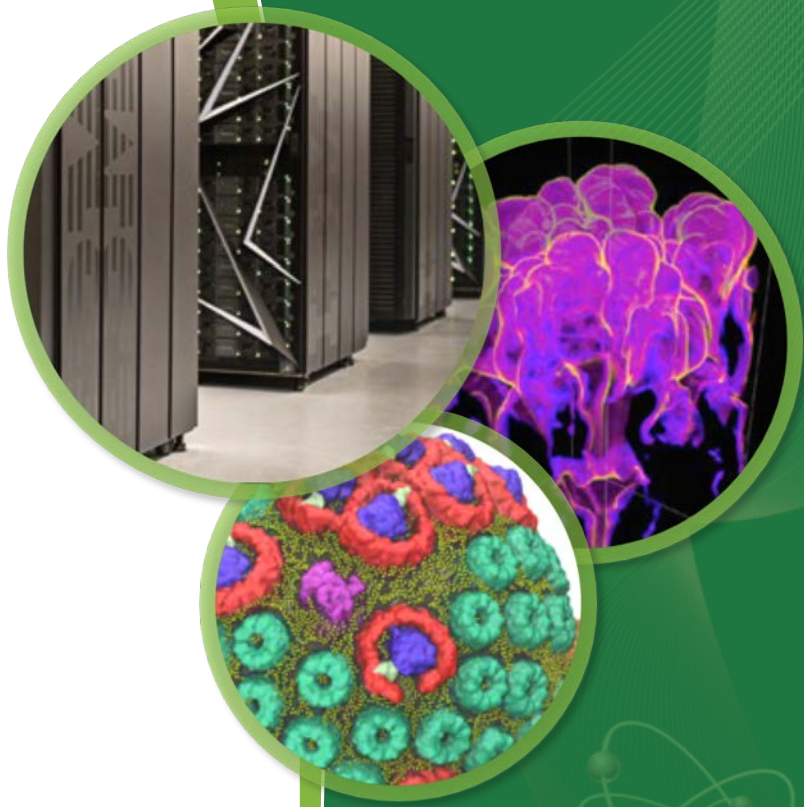

OLCF User Group Meeting

Oak Ridge

15 May 2018

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Some of the work presented here is from the TOTAL and Oak Ridge National Laboratory collaboration which is done under the CRADA agreement NFE-14-05227. Some of the experiments were supported by an allocation of advanced computing resources provided by the National Science Foundation. The computations were performed on Nautilus at the National Institute for Computational Sciences.

Overview of Today's Agenda

- OLCF Staff will give you a few Facility Updates
- Invited talks from OLCF Users Brant Robertson, Evan Schneider, Dan Jacobson, Wayne Joubert, P.K Yeung, and Matthew Clay.
- Two poster slam sessions at 10:30 am and 1:45 pm for our poster participants to briefly introduce their poster
- Lunch will be served next door, and we will have OLCF staff at each table to facilitate a discussion on a specific topic which will be noted on the table (outlined on next slide)
- We will take a Group Photo at 4:30
- There will be a voluntary tour of the OLCF and EVEREST facilities at 4:45 pm. For those interested, OLCF staff will escort you to the facilities.
- The poster session will take place out in the lobby from 5:15-6:30.

Round Table Lunch Topics

Tuesday Line-Up

- User Software Containers
- Tell your Story
- Continuous Integration
- OLCF Allocation Programs
- Constellation
- Employment Opportunities at the OLCF
- OLCF Website Feedback
- Performance Portability
- Deep Learning
- Center Management

Wednesday Line-Up

- Summit Early Science Program
- Data Initiatives/Programs
- Summit Burst Buffers
- OLCF Training
- JSRUN & LSF Tools on Summit
- Using Kokkos
- Software Engineering
- Machine Learning for Applications
- Analytics/Learning

Safety

- Emergency exits – There are two stair cases available in the case of an emergency.
- Please stop and listen if a safety announcement comes over the PA system.
- Please adhere to any instructions given over the PA system.
- As a DOE laboratory, we must maintain strict access control to the buildings. Therefore, please stay in this building unless you are escorted by an OLCF staff member.

Miscellaneous Announcements

- If you are interested in meeting with someone on the OLCF staff or elsewhere in the lab and need our help in arranging the meeting, please let us know.
- We have reserved several rooms within this building for additional meeting space. Please see Sherry Ray if you would like to reserve a meeting room while you are here.
- If you are a presenter and have not given us a copy of your slides, please see Ashley Barker or Sherry Ray at the next break so that we can get your presentation loaded on the system as soon as possible.
- We are broadcasting this meeting over BlueJeans for those users that could not join us in person.
- Restrooms and vending machines are located down the hall on the right-hand side.
- There is also a cafeteria located downstairs that is open for both breakfast and lunch.

2018 OLCF User Meeting Survey

- https://bit.ly/2Kqhste

What is a Leadership Computing Facility (LCF)?

- Collaborative DOE Office of Science user-facility program at ORNL and ANL
- Mission: Provide the computational and data resources required to solve the most challenging problems.
- 2-centers/2-architectures to address diverse and growing computational needs of the scientific community

- Highly competitive user allocation programs (INCITE, ALCC).
- Projects receive 10x to 100x more resource than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).

Origin of Leadership Computing Facility

Department of Energy High-End Computing Revitalization Act of 2004 (Public Law 108-423):

The Secretary of Energy, acting through the Office of Science, shall

- Establish and operate Leadership Systems Facilities.
- Provide access [to Leadership Systems Facilities] on a competitive, merit-reviewed basis to researchers in U.S. industry, institutions of higher education, national laboratories and other Federal agencies.

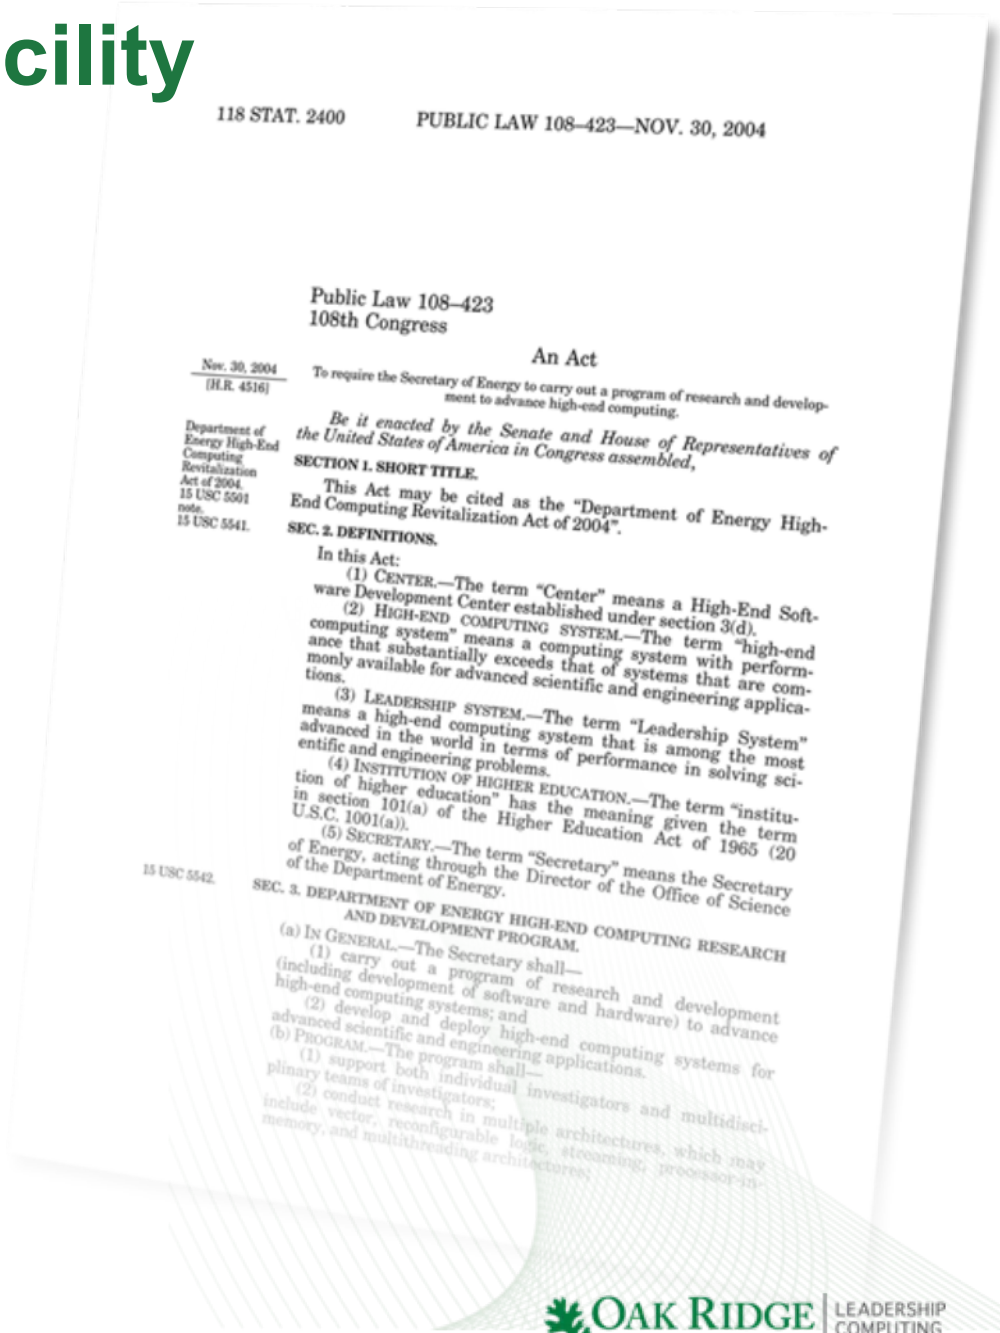

118 STAT. 2400 PUBLIC LAW 108–423—NOV. 30, 2004

Public Law 108–423
108th Congress

An Act

Nov. 30, 2004 [H.R. 4516]

To require the Secretary of Energy to carry out a program of research and development to advance high-end computing.

Department of Energy High-End Computing Revitalization Act of 2004. 15 USC 5501 note. 15 USC 5541.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. SHORT TITLE.

This Act may be cited as the "Department of Energy High-End Computing Revitalization Act of 2004".

SEC. 2. DEFINITIONS.

In this Act:

(1) CENTER.—The term "Center" means a High-End Software Development Center established under section 3(d).

(2) HIGH-END COMPUTING SYSTEM.—The term "high-end computing system" means a computing system with performance that substantially exceeds that of systems that are commonly available for advanced scientific and engineering applications.

(3) LEADERSHIP SYSTEM.—The term "Leadership System" means a high-end computing system that is among the most advanced in the world in terms of performance in solving scientific and engineering problems.

(4) INSTITUTION OF HIGHER EDUCATION.—The term "institution of higher education" has the meaning given the term in section 101(a) of the Higher Education Act of 1965 (20 U.S.C. 1001(a)).

(5) SECRETARY.—The term "Secretary" means the Secretary of Energy, acting through the Director of the Office of Science of the Department of Energy.

15 USC 5542.

SEC. 3. DEPARTMENT OF ENERGY HIGH-END COMPUTING RESEARCH AND DEVELOPMENT PROGRAM.

(a) IN GENERAL.—The Secretary shall—

(1) carry out a program of research and development (including development of software and hardware) to advance high-end computing systems; and

(2) develop and deploy high-end computing systems for advanced scientific and engineering applications.

(b) PROGRAM.—The program shall—

(1) support both individual investigators and multidisciplinary teams of investigators;

(2) conduct research in multiple architectures, which may include vector, reconfigurable logic, streaming, processor-in-memory, and multithreading architectures;

ORNL has systematically delivered a series of leadership-class systems

On scope • On budget • Within schedule

OLCF-1

OLCF-2

OLCF-3

18.5 TF
2004
Cray X1E
Phoenix

25 TF
2005
Cray XT3
Jaguar

54 TF
2006
Cray XT3
Jaguar

62 TF
2007
Cray XT4
Jaguar

263 TF
2008
Cray XT4
Jaguar

1 PF
2008
Cray XT5
Jaguar

2.5 PF
2009
Cray XT5
Jaguar

1000-fold improvement in 8 years

27 PF
2012
Cray XK7
Titan

Titan, five years old in October 2017, continues to deliver world-class science research in support of our user community. We will operate Titan through 2019 when it will be decommissioned.

We are building on this record of success to enable exascale in 2021

OLCF Users' Publication Yield at All-Time High

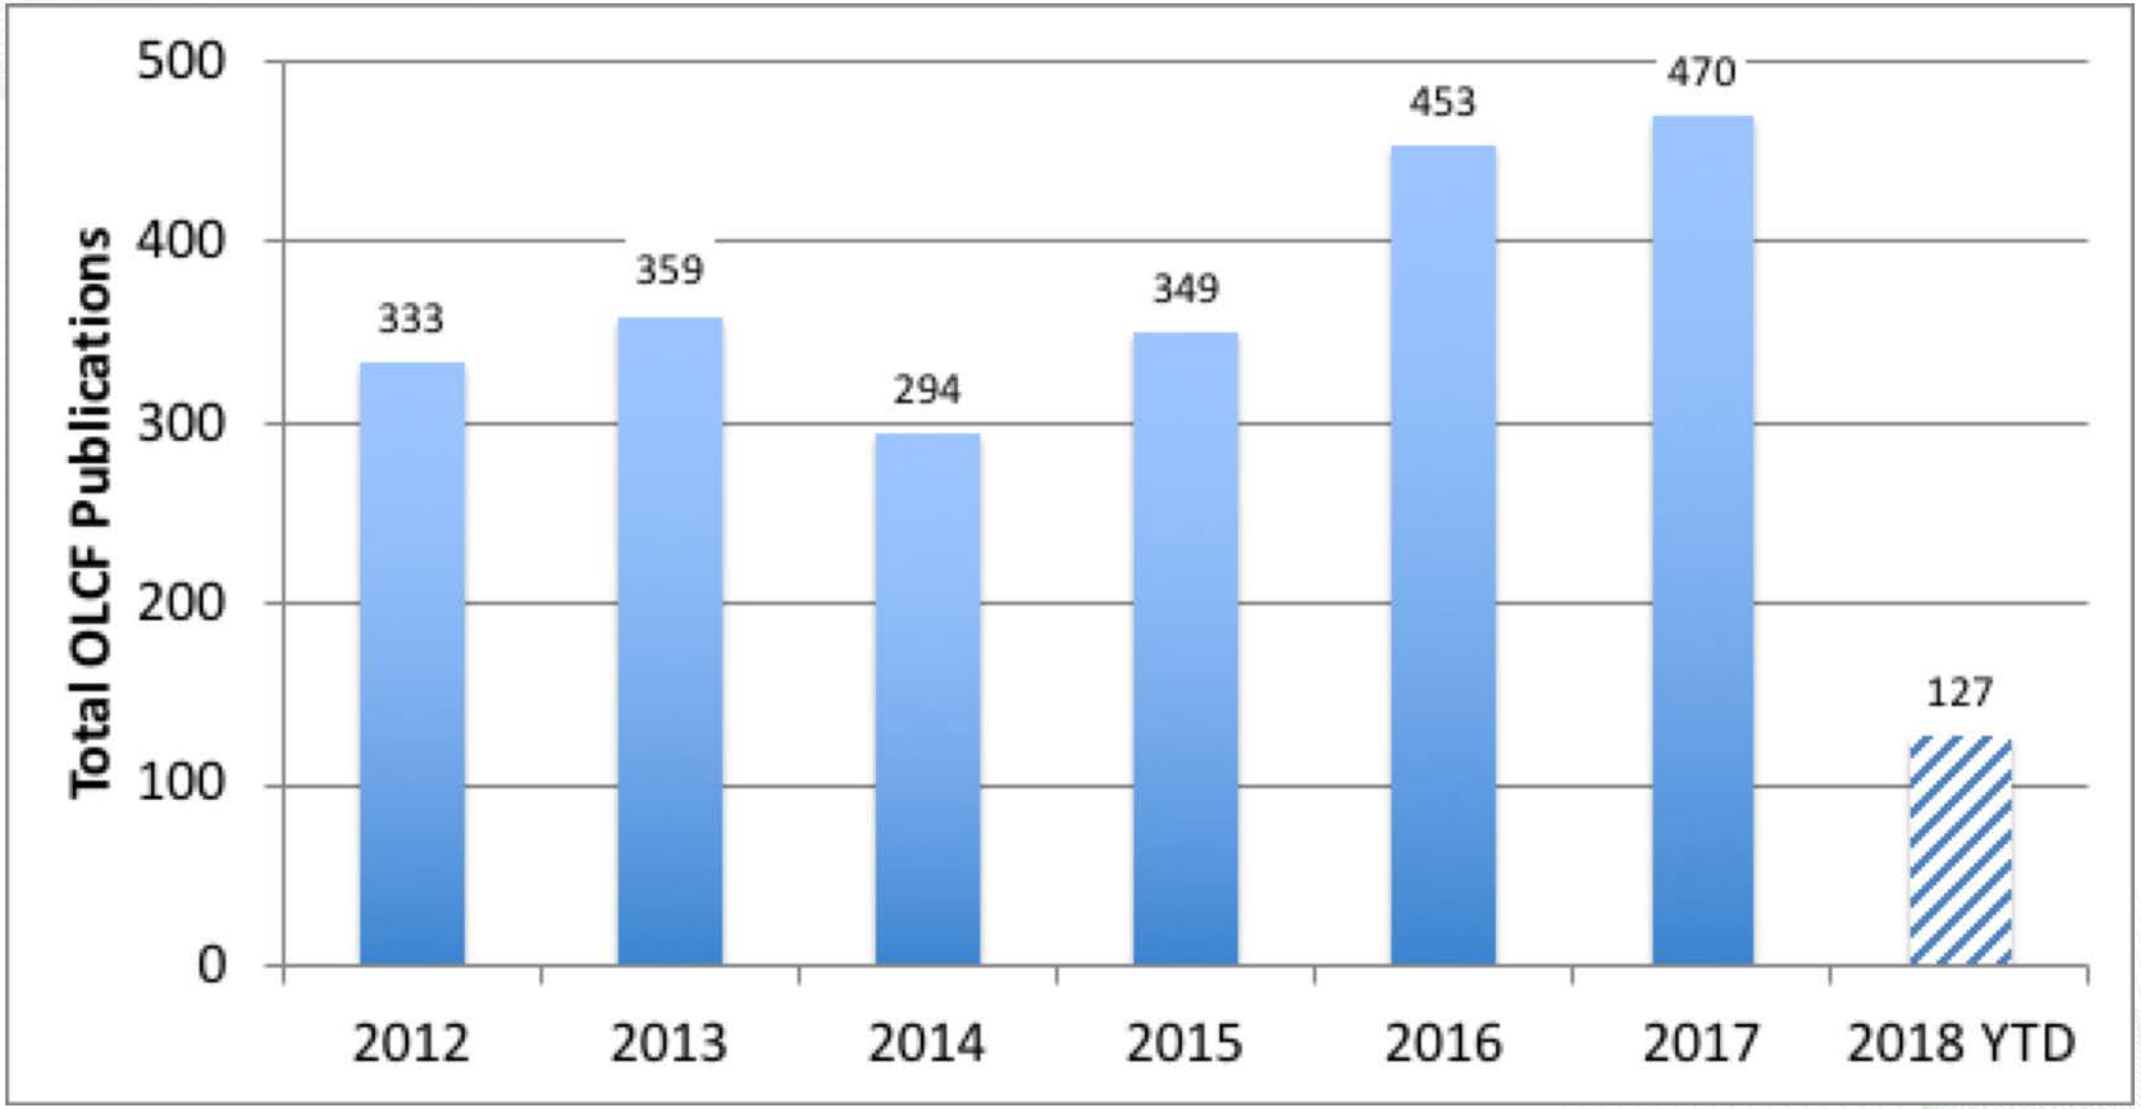

Data Snapshot: May 13, 2018

OLCF Users' Publications Increase in High-Impact Journals

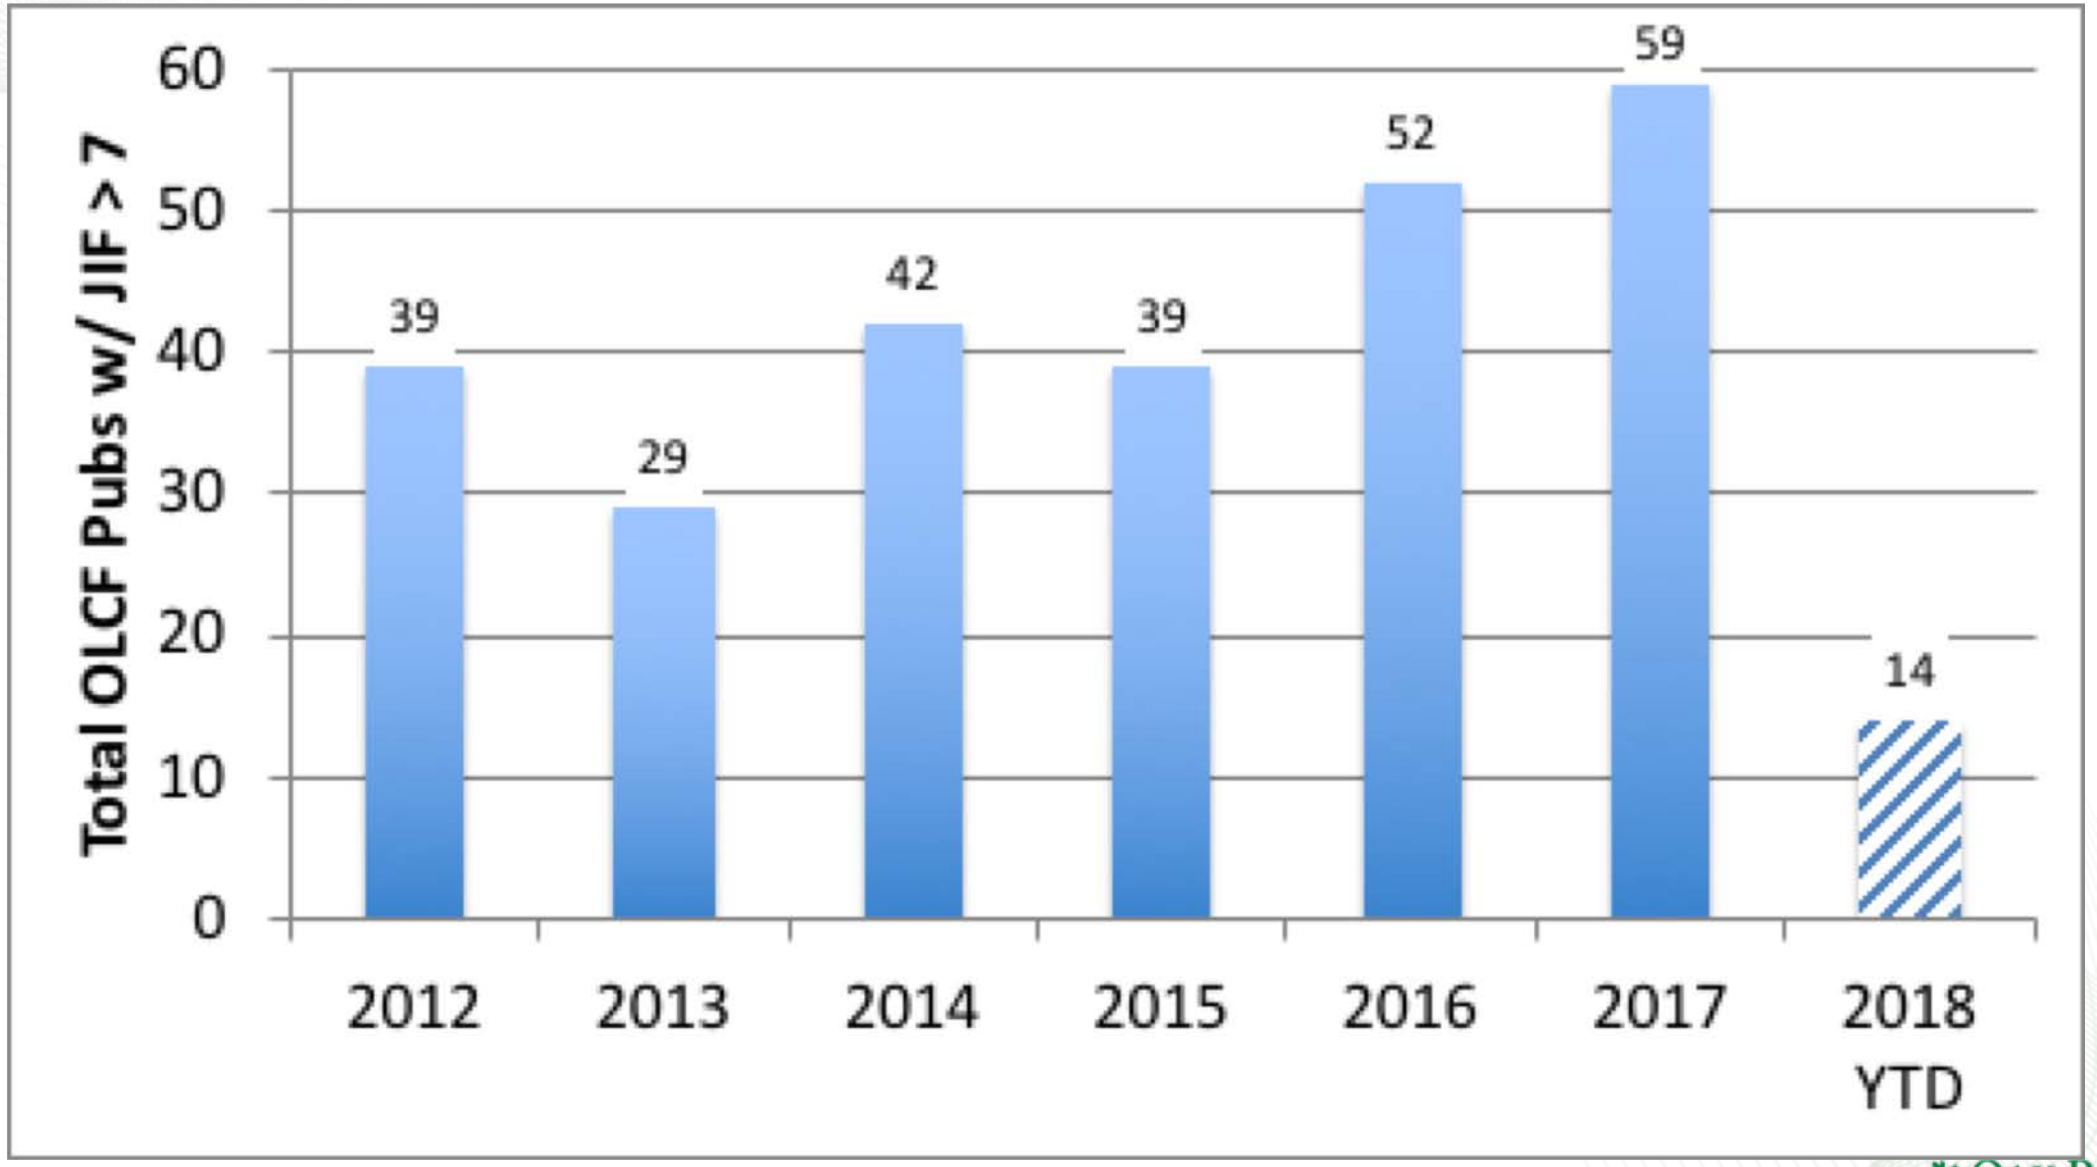

Data Snapshot: May 13, 2018

Setting Records

Measurement	2017 target	2017 actual	2018 YTD
Scheduled availability	95%	99.39%	99.54%
Overall availability	90%	**98.09%** ★	99.14%
Overall User Support Rating	3.5/5.0	4.6	N/A
User Problem Resolution Time	80% in 3 hrs	**93%** ★	97%
Total usage	NAM	**91%** ★	91%
Core-hours used	NAM	**4,389,163,123** ★	1,597,338,620
Core-hours available	NAM	**4,817,215,104** ★	1,746,954,240
INCITE Capability	NAM	**68.22%** ★	46.12%
Capability - All projects	35%	**59.81%** ★	47.56%

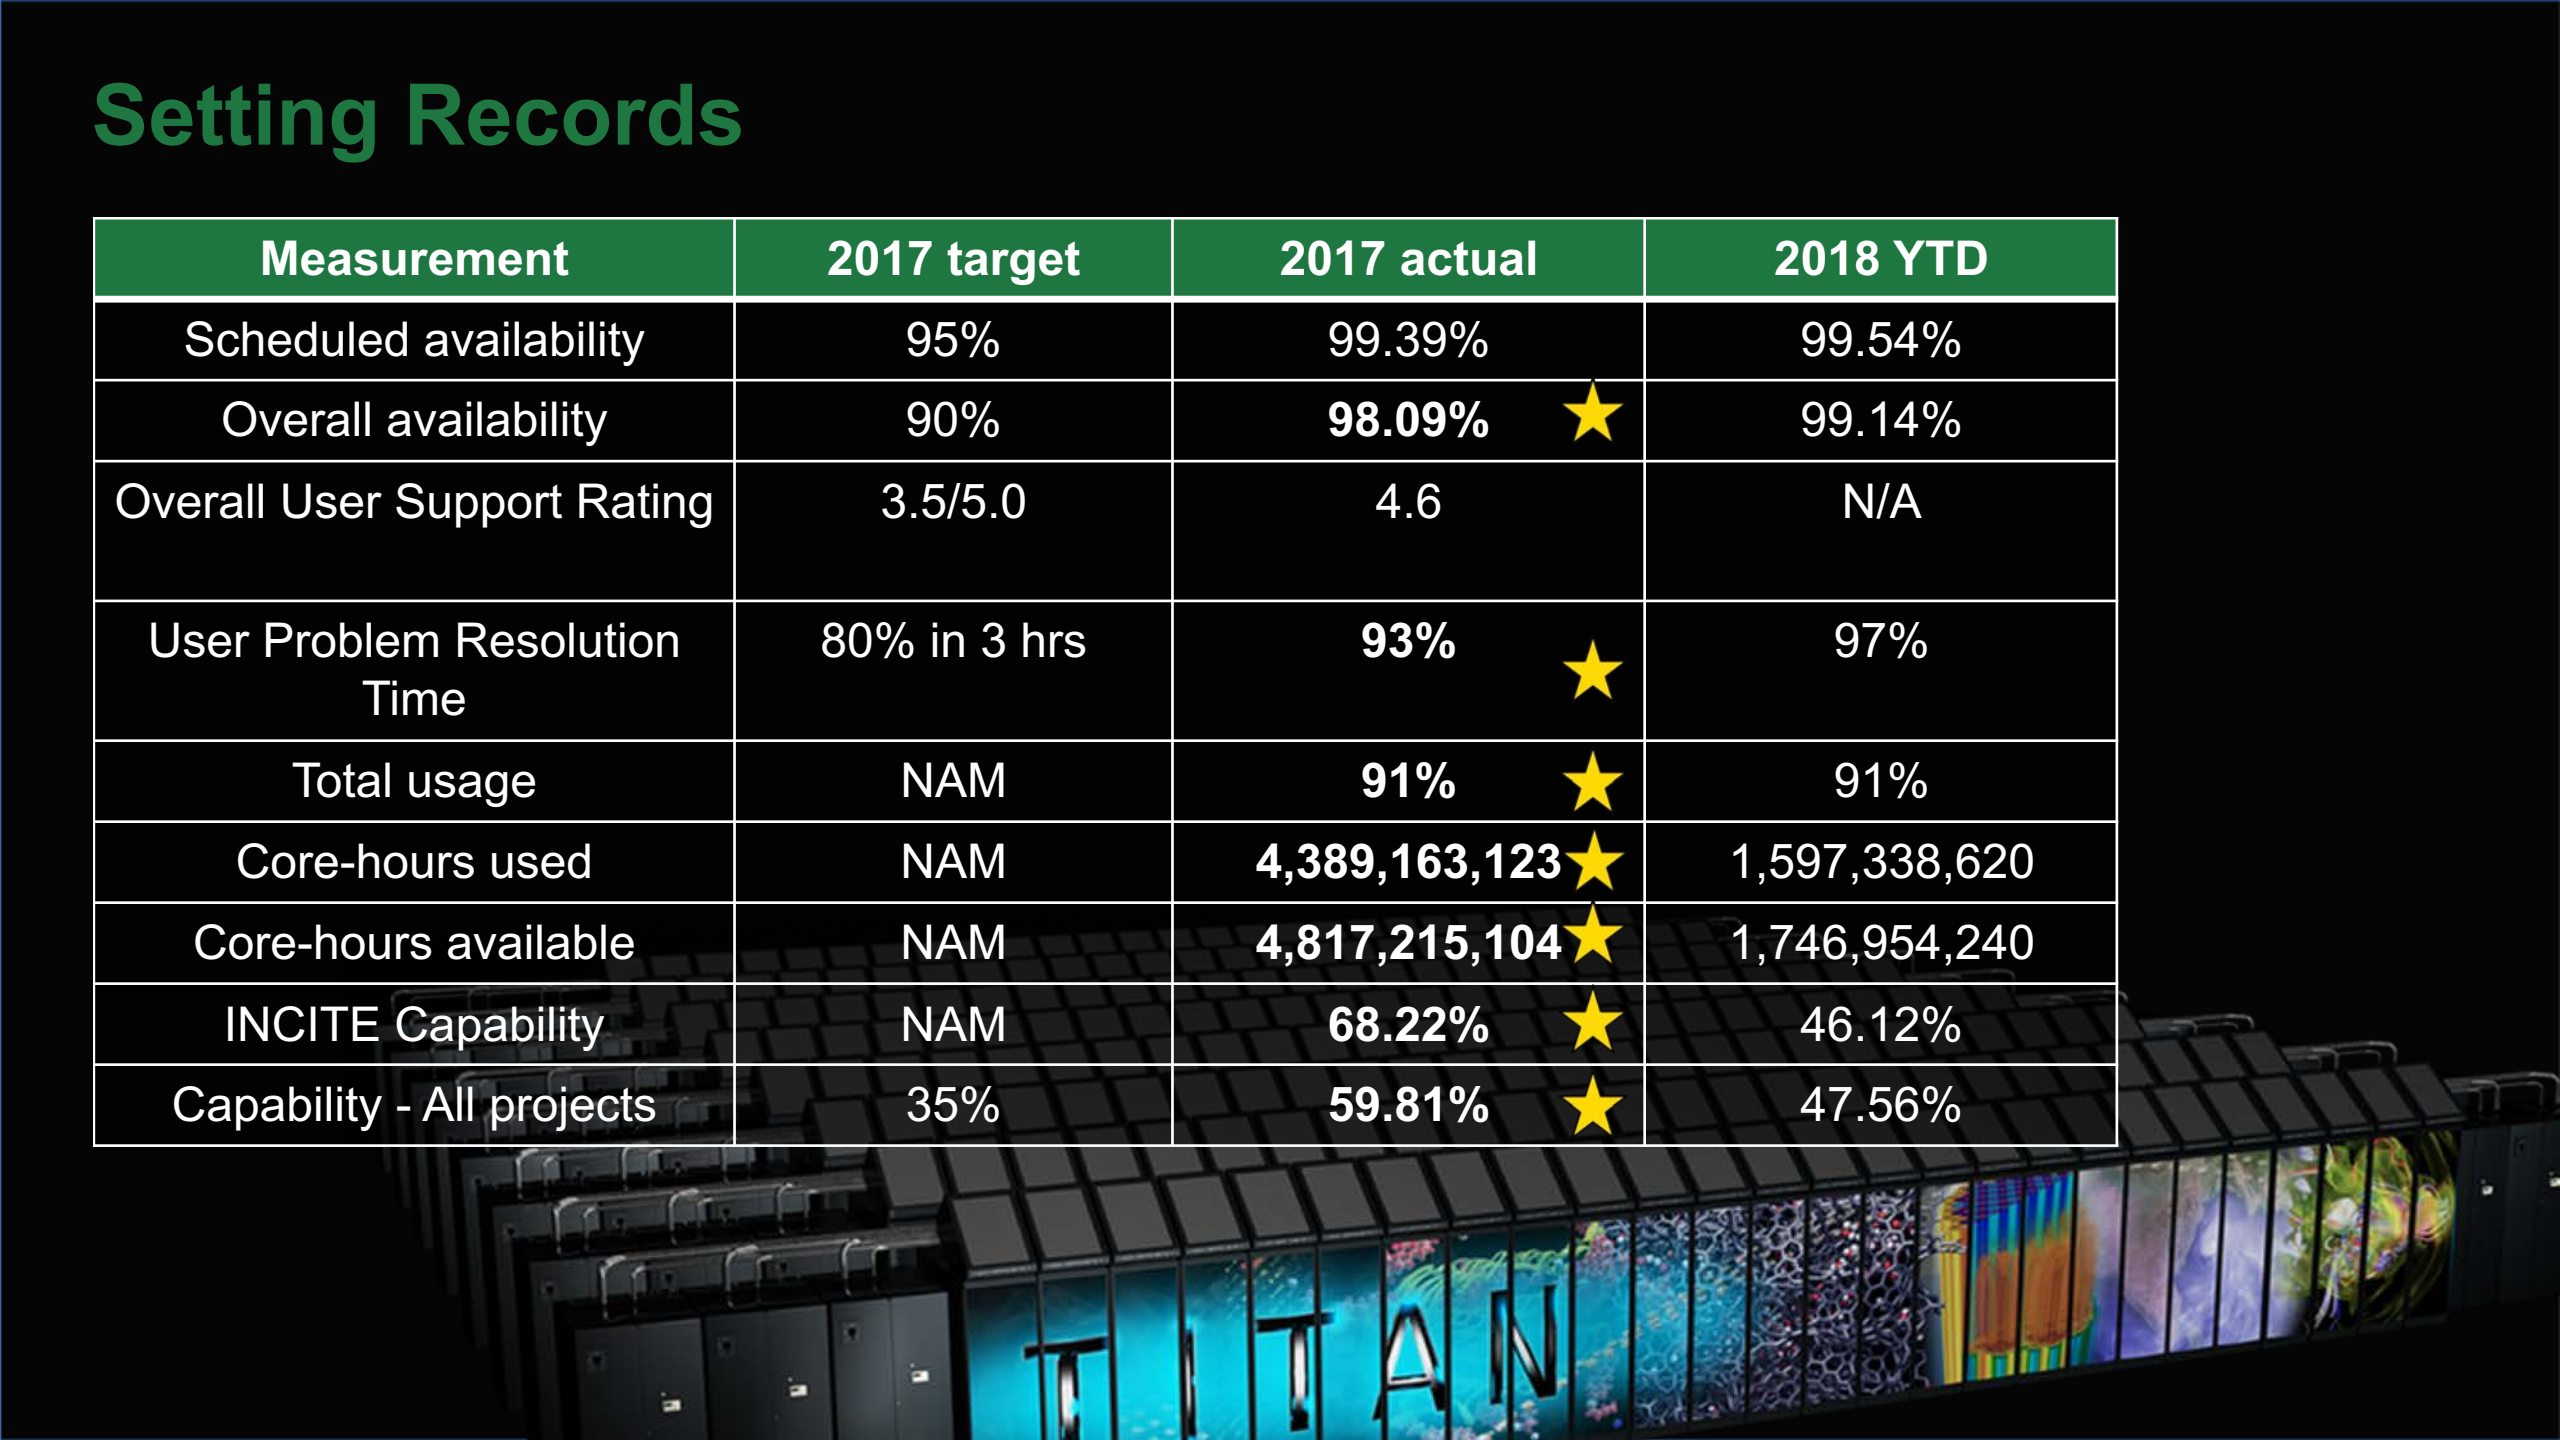

What about the Titan component failures?

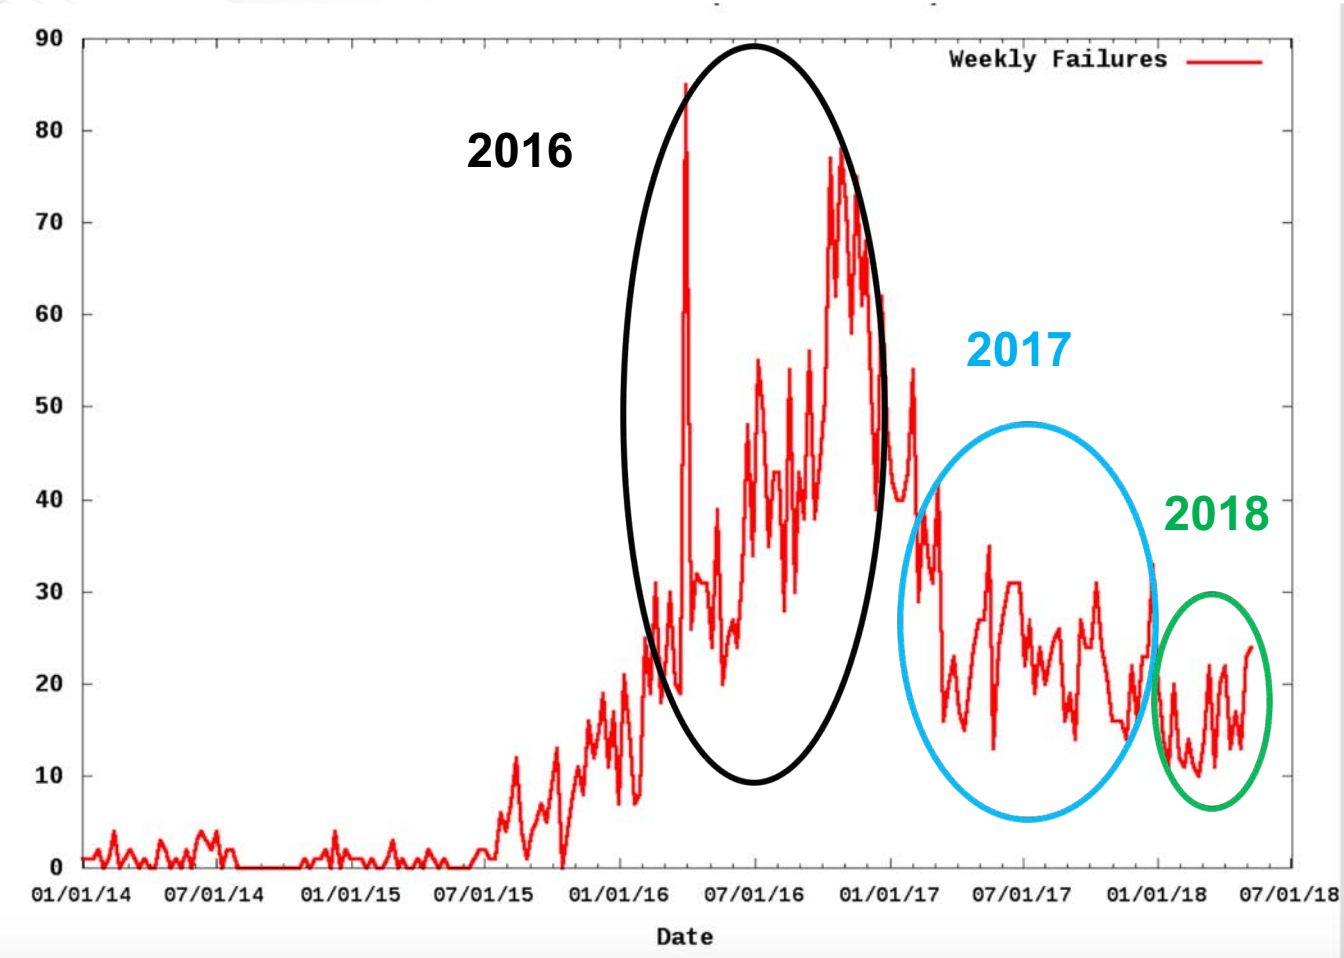

- **2016 – peak weekly failure rate**
- **2017 – completed 11,000 GPU SXM replacements**
- **2018 – numbers to date reflect policy enhancements that were enacted in late 2017**

What's next for Titan?

- CUDA 9.2 upgrade to Titan coming soon
  - Currently in testing and evaluation at the OLCF
- Scheduled end-of-life for September 2019 marking 7-years in production

Coming in 2018: Summit will replace Titan as the OLCF's leadership supercomputer

Summit, slated to be more powerful than any other existing supercomputer, is the Department of Energy's Oak Ridge National Laboratory's newest supercomputer for open science.

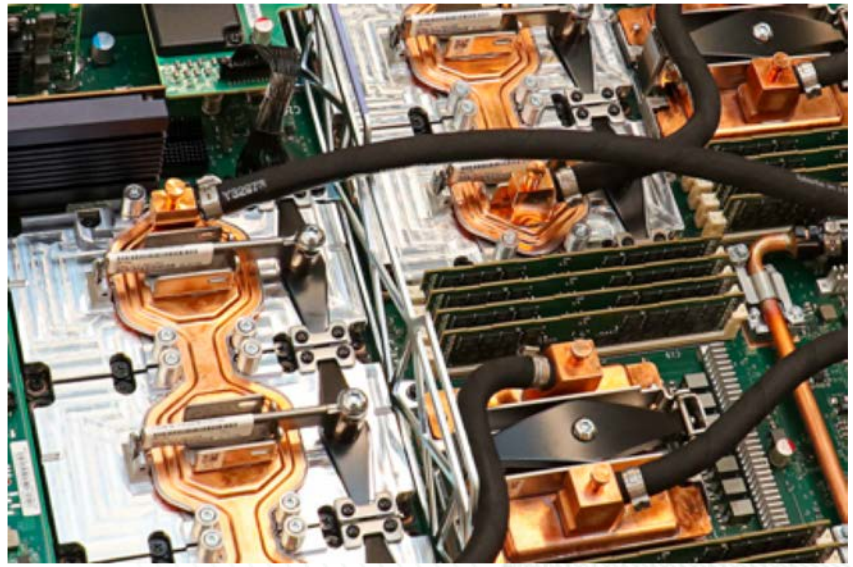

Coming in 2018: Summit will replace Titan as the OLCF's leadership supercomputer

- Many fewer nodes
- Much more powerful nodes
- Much more memory per node and total system memory
- Faster interconnect
- Much higher bandwidth between CPUs and GPUs
- Much larger and faster file system

Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	4,608
Node performance	1.4 TF	42 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM2
NV memory per Node	0	1600 GB
Total System Memory	710 TB	>10 PB DDR4 + HBM2 + Non-volatile
System Interconnect	Gemini (6.4 GB/s)	Dual Rail EDR-IB (25 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Bi-Section Bandwidth	15.6 TB/s	115.2 TB/s
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Power Consumption	9 MW	13 MW

Installation Nearing Completion

- Hardware installation completed in March
- Continuing to stabilize nodes, disks, and network
- In December, accepted 1,080 of 4,608 nodes to port codes
- OLCF is working with IBM, NVIDIA, Red Hat, and Mellanox to stabilize and debug system software

Summit is still under construction

- We expect to accept the machine in Summer of 2018, allow early users on this year, and allocate our first users through the INCITE program in January 2019.

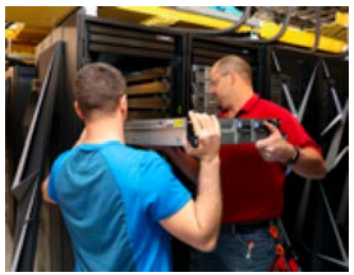

Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program for 2019

- Access to the most capable, most productive, fastest open science supercomputers in the nation
- Call for proposals submission window:
  - Apr 16 – Jun 22, 2018
- Applicable to a broad array of science, engineering, and computer science domains
- Proposals must be:
  - High-impact, computationally and/or data intensive campaigns
  - Must take advantage of unique HPC architectures
  - Research that cannot be performed anywhere else.
- INCITE Webinar will be held June 7th
- For more information visit http://www.doeleadershipcomputing.org/

How can you get access to Summit to prepare for INCITE and ALCC competitions?

- CAAR and Early Science Program (ESP) Letter of Intent (LOI) calls are closed;
  - 13 CAAR teams; 60+ ESP-LOI Teams have access to Summit Phase I
  - ESP Proposals due by 1 June 2018,
  - See T. Straatsma's presentation Wednesday Morning.
- INCITE CFP is open now, closes 22 June 2018.
  - Preparatory access via Titan, Summitdev, and ESP-LOI process on Summit
- DOE's ALCC program call is expected to open in Fall 2018, expected to close February 2019. For ALCC preparation access:
  - 2018: via Titan, Summitdev, and ESP-LOI process on Summit; Potential to access additional IBM AC922 cluster hardware for application development;
  - 2019: Director's Discretionary (DD) Research begins January 2019.

2018 Training Events: Preparing for Summit

- Workshops, webinars, screencasts
- GPU Hackathons
- Mini-GPU-Hackathons

Prepared and presented by OLCF staff, vendors, and partner organizations.

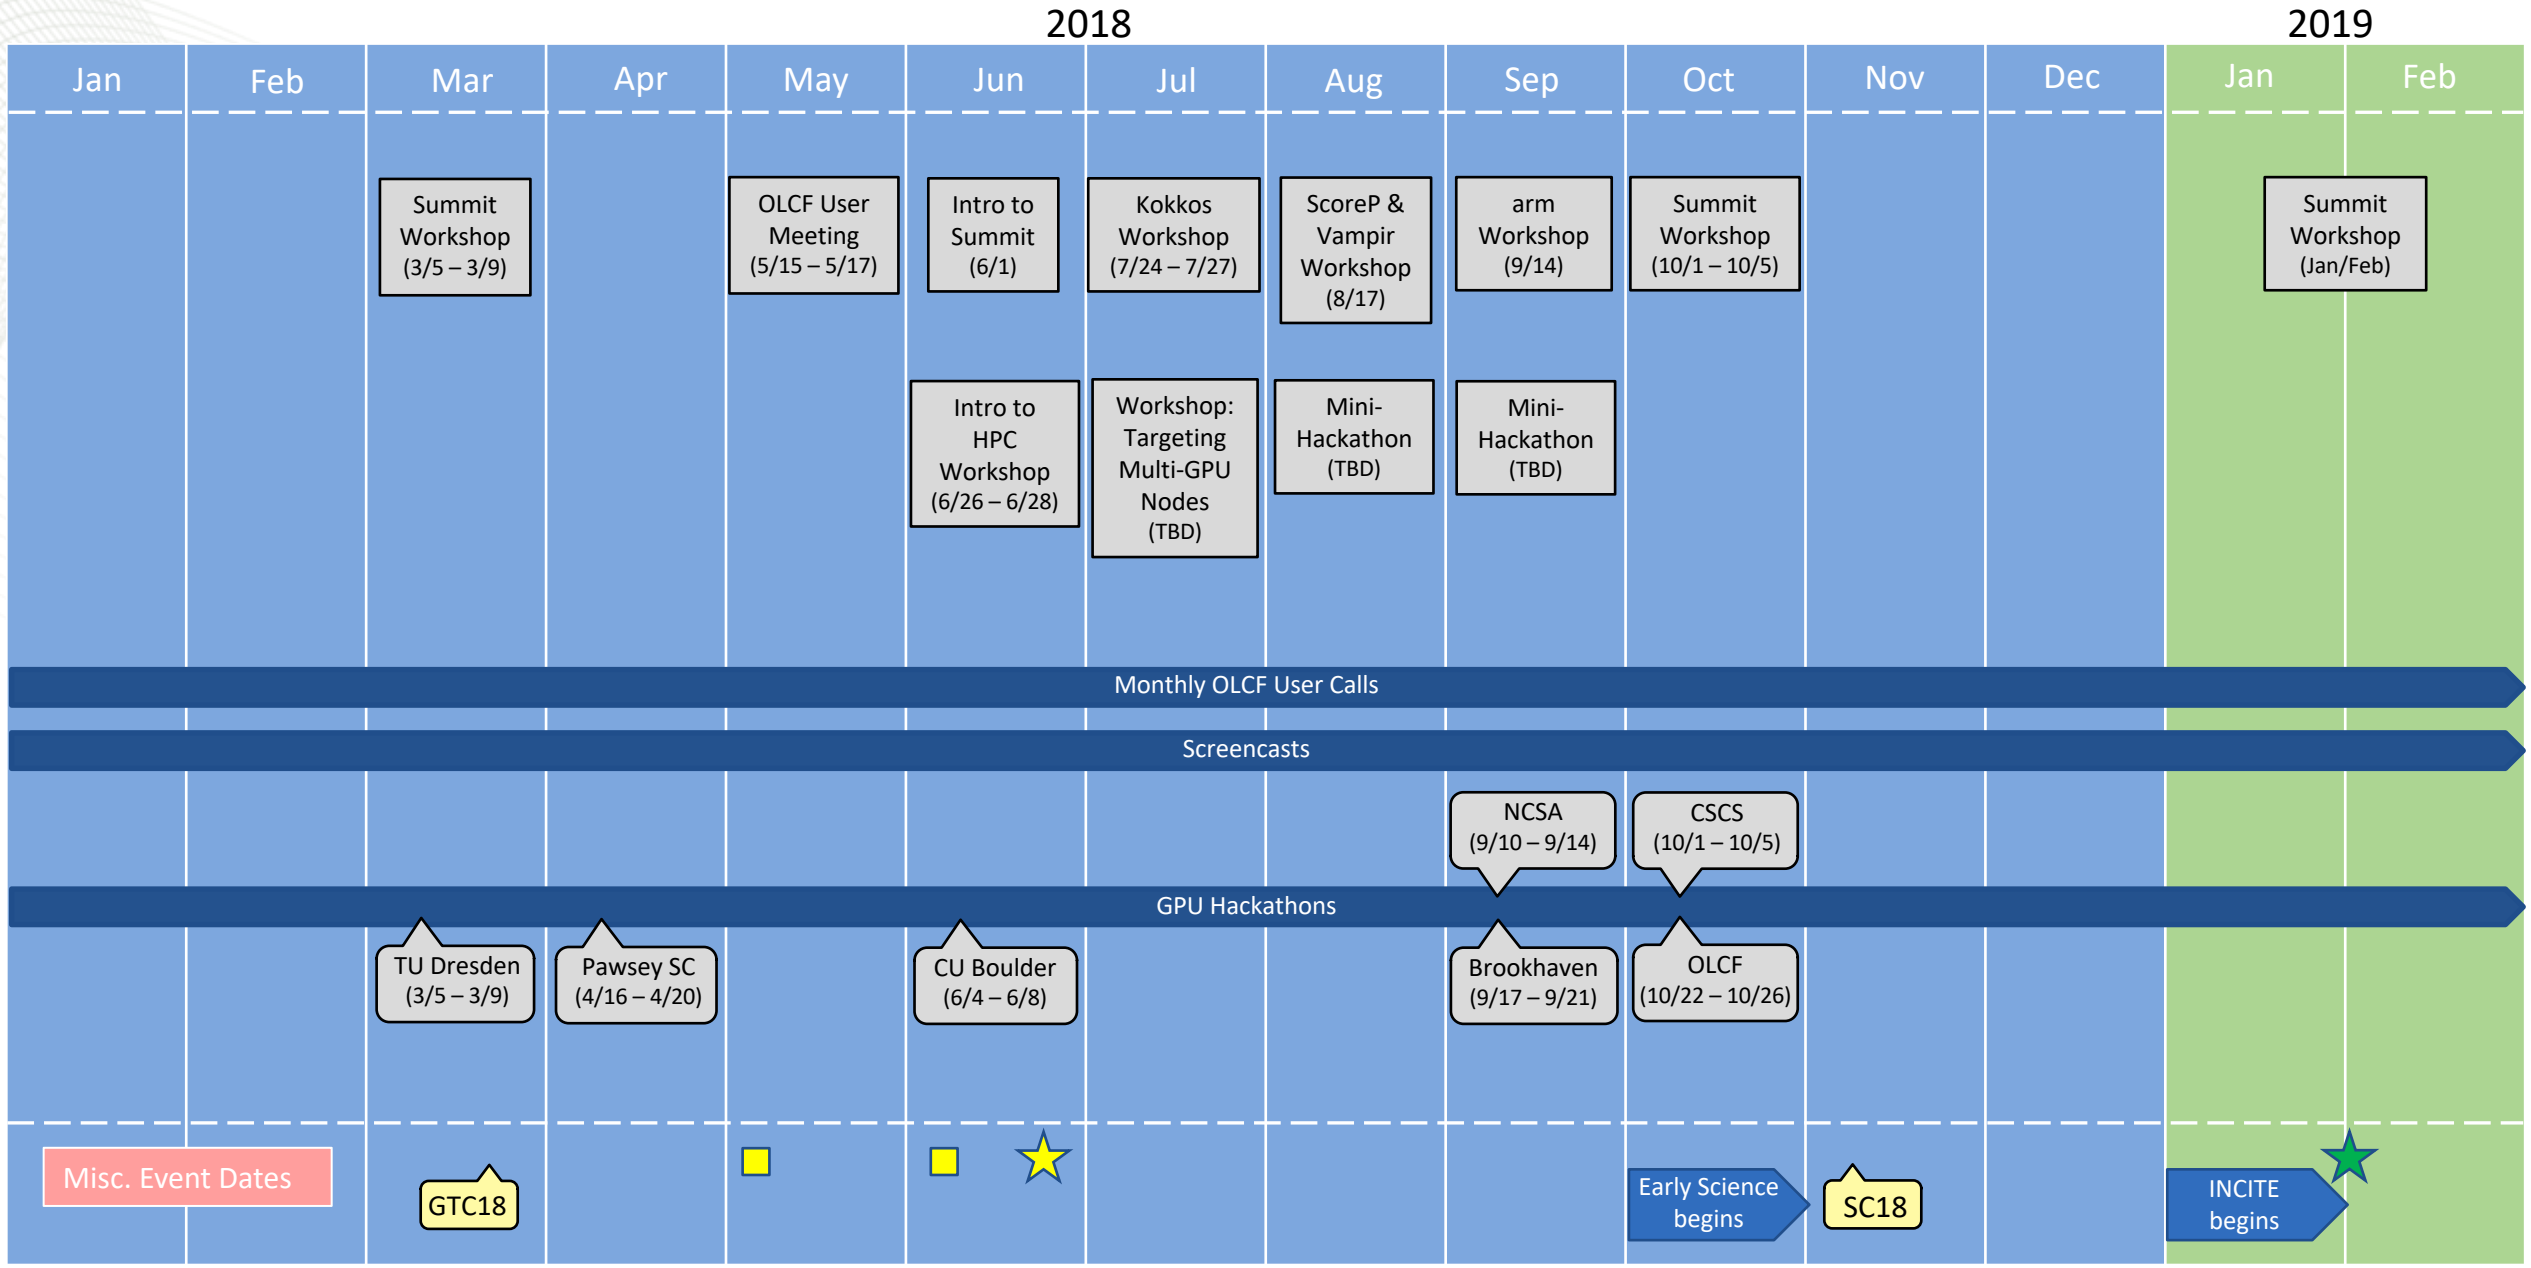

INCITE call closes on 6/22

ALCC call closes on 2/1

INCITE proposal webinars (5/2 and 6/7)

OLCF User Calls
Feb: jsrun (2/28)
Mar: Profiling with forge (3/28)
Apr: CUDA9.2 (4/25)
Jun+: TBD

Screencasts
Hello_jsrun Tool
Transferring Data between Atlas & AlpineTDS

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Summit will be the world's smartest supercomputer for open science

But what makes a supercomputer smart?

Summit provides unprecedented opportunities for the integration of artificial intelligence (AI) and scientific discovery. Here's why:

- **GPU Brawn:** Summit links more than 27,000 deep-learning optimized NVIDIA GPUs with the potential to deliver exascale-level performance (a billion-billion calculations per second) for AI applications.
- **High-speed Data Movement:** NVLink high-bandwidth technology built into all of Summit's processors supplies the next-generation "information superhighways" needed to train deep learning algorithms for challenging science problems quickly.
- **Memory Where it Matters**: Summit's sizable local memory gives AI researchers a convenient launching point for data-intensive tasks, an asset that allows for faster AI training and greater algorithmic accuracy.

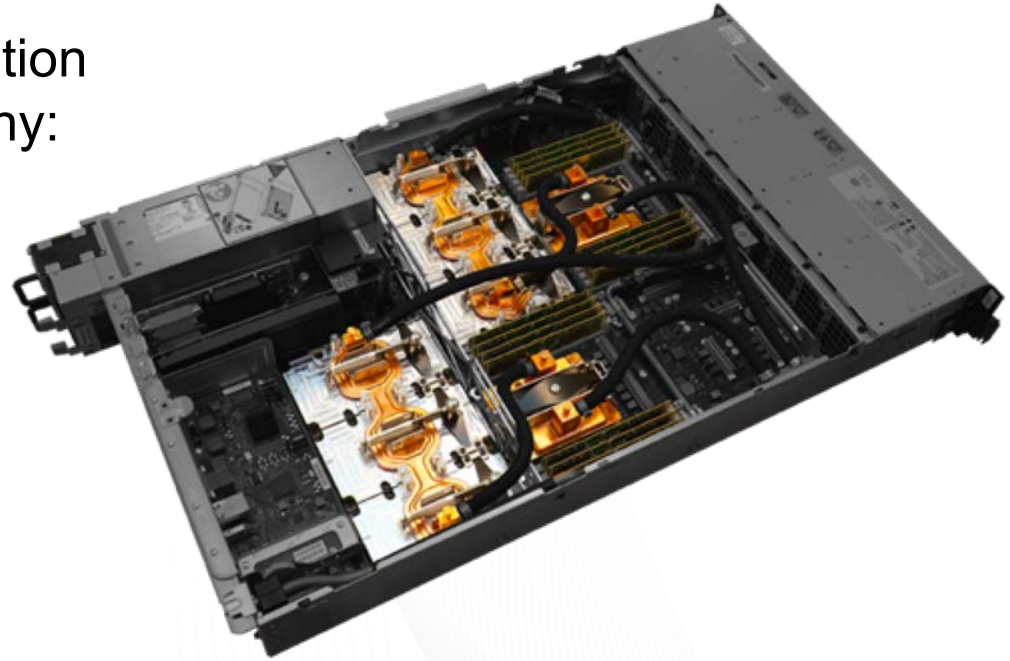

One of Summit's 4,600 IBM AC922 nodes. Each node contains six NVIDIA Volta GPUs and two IBM Power9 CPUs, giving scientists new opportunities to automate, accelerate and drive understanding using artificial intelligence techniques.

Summit will be the world's smartest supercomputer for open science

But what can a smart supercomputer do?

Science challenges for a smart supercomputer:

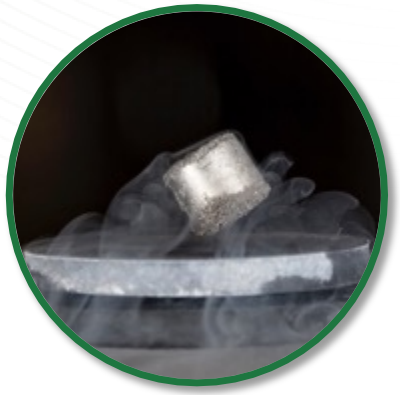

Identifying Next-generation Materials
By training AI algorithms to predict material properties from experimental data, longstanding questions about material behavior at atomic scales could be answered for better batteries, more resilient building materials, and more efficient semiconductors.

Predicting Fusion Energy
Predictive AI software is already helping scientists anticipate disruptions to the volatile plasmas inside experimental reactors. Summit's arrival allows researchers to take this work to the next level and further integrate AI with fusion technology.

Deciphering High-energy Physics Data
With AI supercomputing, physicists can lean on machines to identify important pieces of information—data that's too massive for any single human to handle and that could change our understanding of the universe.

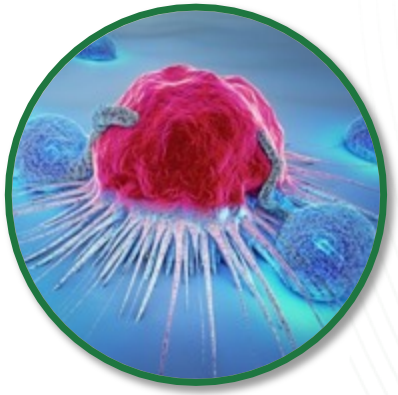

Combating Cancer
Through the development of scalable deep neural networks, scientists at the US Department of Energy and the National Cancer Institute are making strides in improving cancer diagnosis and treatment.

Emerging Science Activities:
Selected Machine Learning Projects on Titan: 2016-2017

Program	PI	PI Employer	Project Name	Allocation (Titan core-hrs)
ALCC	Robert Patton	ORNL	Discovering Optimal Deep Learning and Neuromorphic Network Structures using Evolutionary Approaches on High Performance Computers	75,000,000
ALCC	Gabriel Perdue	FNAL	Large scale deep neural network optimization for neutrino physics	58,000,000
ALCC	Gregory Laskowski	GE	High-Fidelity Simulations of Gas Turbine Stages for Model Development using Machine Learning	30,000,000
ALCC	Efthimions Kaxiras	Harvard U.	High-Throughput Screening and Machine Learning for Predicting Catalyst Structure and Designing Effective Catalysts	17,500,000
ALCC	Georgia Tourassi	ORNL	CANDLE Treatment Strategy Challenge for Deep Learning Enabled Cancer Surveillance	10,000,000
DD	Abhinav Vishnu	PNNL	Machine Learning on Extreme Scale GPU systems	3,500,000
DD	J. Travis Johnston	ORNL	Surrogate Based Modeling for Deep Learning Hyper-parameter Optimization	3,500,000
DD	Robert Patton	ORNL	Scalable Deep Learning Systems for Exascale Data Analysis	6,500,000
DD	William M. Tang	PPPL	Big Data Machine Learning for Fusion Energy Applications	3,000,000
DD	Catherine Schuman	ORNL	Scalable Neuromorphic Simulators: High and Low Level	5,000,000
DD	Boram Yoon	LANL	Artificial Intelligence for Collider Physics	2,000,000
DD	Jean-Roch Vlimant	Caltech	HEP DeepLearning	2,000,000
DD	Arvind Ramanathan	ORNL	ECP Cancer Distributed Learning Environment	1,500,000
DD	John Cavazos	U. Delaware	Large-Scale Distributed and Deep Learning of Structured Graph Data for Real-Time Program Analysis	1,000,000
DD	Abhinav Vishnu	PNNL	Machine Learning on Extreme Scale GPU systems	1,000,000
DD	Gabriel Perdue	FNAL	MACHINE Learning for MINERvA	1,000,000
		TOTAL		**220,500,000**

OLCF Accelerating ML/AI Deployment for Science

1. Establishing early science and significant applications
   - Summit/Dev Early Science Applications (e.g., CANDLE)
   - INCITE projects (e.g., Co-evolutionary Networks: From Genome to 3D Proteome, Jacobson, et al.)
   - Directors Discretion (e.g., Fusion RNN, MiNerva)

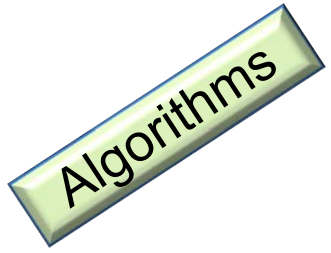

2. Creating leadership class analytics capabilities
   - Leadership analytics (e.g., Frameworks: pbdR, TensorFlow + Horovod)
   - Large-scale targeted algorithms (e.g., non-negative matrix factorization)

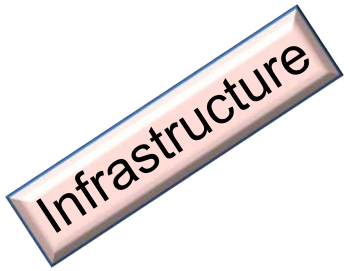

3. Enabling infrastructure for analytics/AI
   - Workflows to include data from observations for analysis within OLCF
   - Analytics enabling technologies (e.g., container deployments for rapidly changing DL/ML frameworks, analytic notebooks, etc.)

Enabling Leadership Science: Data, Simulation, and Analytics at Scale

- Cross-Facility Workflows
- Platform as a service (e.g., OpenShift)
- Container Deployments
- Machine Learning/Deep Learning at Scale

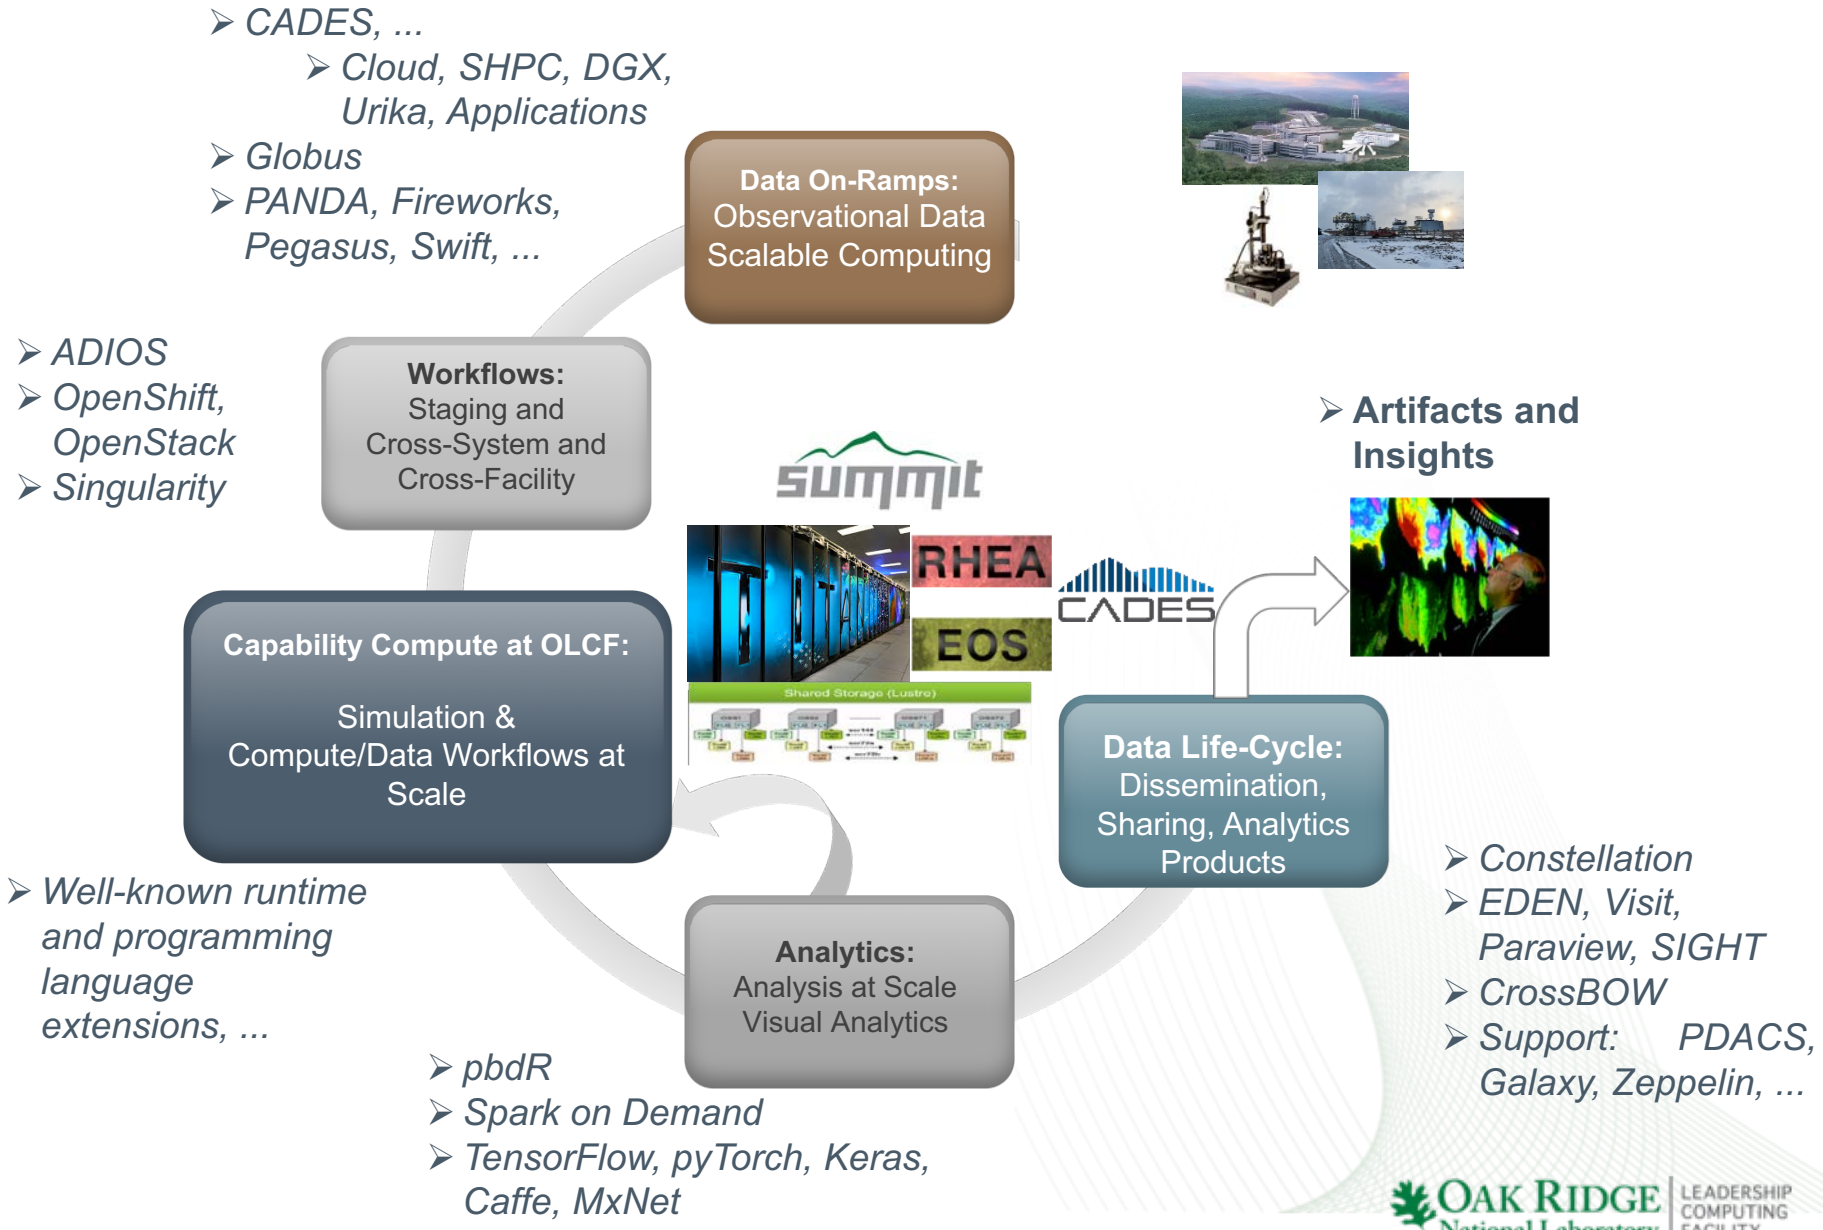

CORAL COLLABORATION
OAK RIDGE ARGONNE LIVERMORE

In April 2018, DOE Secretary Rick Perry announced an open, competitive solicitation for exascale computers to be delivered in 2021 and 2022 known as the CORAL-2 Request For Proposals

http://procurement.ornl.gov/rfp/CORAL2/

- The designs for these computers will target a **50x increase in scientific application performance** over current DOE leadership computers.
- CORAL will award up to 3 computer systems
  - beginning with a 2021 installation at Oak Ridge National Laboratory, with the requirement that the system must be diverse from the ALCF 2021 system
  - followed by a 2022 installation at Lawrence Livermore National Laboratory
  - with the potential for a 3rd system at Argonne National Laboratory
- Responses to this solicitation are due on May 24, 2018

OLCF Path to Exascale

Competitive procurement asking for:

- 50–100× application performance of Titan
- Support for traditional modeling and simulation, high-performance data analysis, and artificial intelligence applications
- Smooth transition for existing and future applications

> "The Exascale Computing Project has emphasized that Exascale is a measure of application performance, and this RFP reflects that, asking for nominally 50× improvement over Sequoia and Titan."
>
> -- Design Reviewer

2008 — Jaguar: 2.3 PF
World's Fastest

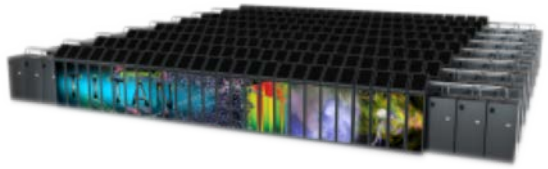

2012 — Titan: 27 PF
Accelerated Computing
World's Fastest

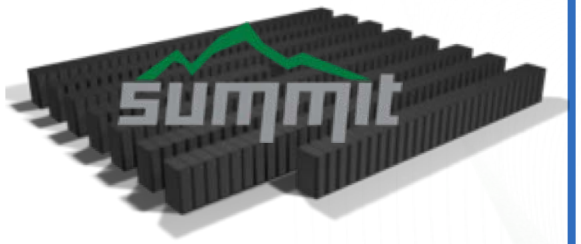

2017 — Summit: 200 PF
Accelerated Computing
5–10× Titan Performance

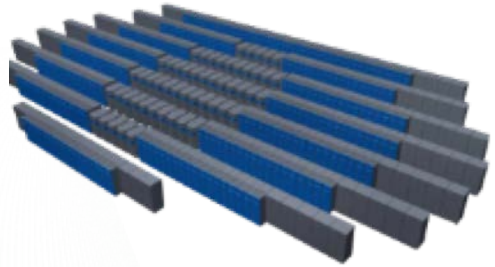

2021 — Frontier: >1000 PF
Competitive Procurement
5-10× Summit Performance

FRONTIER

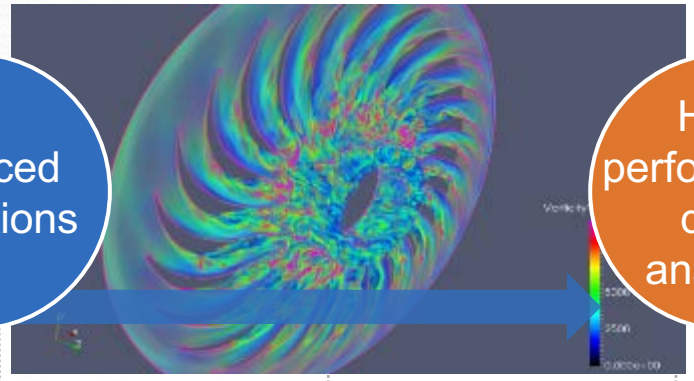

Advanced simulations

High-performance data analytics

Artificial intelligence

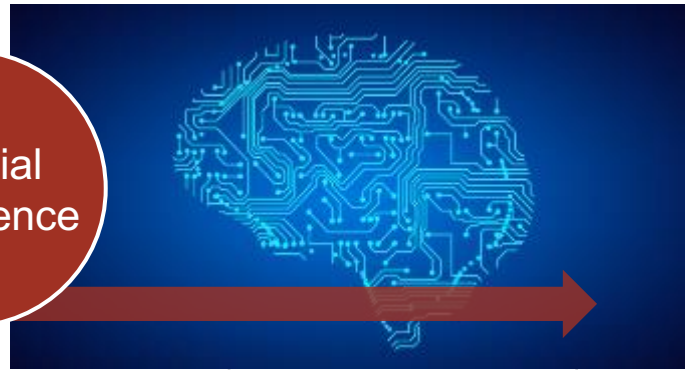

Design review

Solicit bids for computer build

Award contract

Go/No-go

2021 delivery planned with >1000 PF performance (50–100× the scientific productivity of current Titan system)

Frontier NRE: hardware and software engineering

Prepare power, cooling, and space

Installation

Prepare scientific software applications

Frontier available

Summit available: 200 PF production system

Titan available: 27 PF

CY 2017 CY 2018 CY 2019 CY 2020 CY 2021 CY 2022 CY 2023

The OLCF User Group (OUG)

https://www.olcf.ornl.gov/about-olcf/oug/

- "*The purpose of OUG is to provide advice and feedback to the OLCF on the current and future state of OLCF operations and services.*"
- The OUG Executive Board represents OUG to the OLCF and serves as advocates for the OLCF user community.
- Six users are running for three open positions on the board this year: https://tinyurl.com/OUGElection2018
- Vote until Thursday, May 17th 9AM EDT: https://tinyurl.com/OUGVote2018
- Results will be announced at the conclusion of the User Meeting

2018 OLCF User Meeting Survey

- https://bit.ly/2Kqhste