A NEW VINTAGE OF COMPUTING CUG 2013

Napa Valley, California • May 6–9

SPIDER II

OLCF's next-generation Spider file system

20 Years of Excellence in Computational Science

OLCF

OAK RIDGE LEADERSHIP COMPUTING FACILITY

1992–2012

Dave Dillow
HPC Architect
Technology Integration Group
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

May , 2013

Spider at OLCF

- Deployed in 2008
  - 240 GB/s, 10 PB, proudly served more than 26,000 clients
  - 4 MDS, 192 OSS, and 1,344 OSTs
  - Center-wide, shared resource, scratch space for OLCF users/projects
  - Jaguar was the main consumer
    - >18,000 clients, ~2 PFLOPS, 300 TB main memory

Plans to upgrade Jaguar started early on, in-parallel with Spider II plans

Motivations for Spider

- Single shared storage pool
  - For all OLCF resources
- Aggregate performance and scalability
  - For all OLCF resources
- Resilience against system failures
  - internal to the storage system as well as failures of any computational resources
- Allow growth of the storage pool
  - independent of the computational platforms

Spider met all of these requirements!

All valid for next-generation Spider as well!

Titan at OLCF

- Jaguar upgrade
  - Still in acceptance
  - A magnitude higher in compute power compared to Jaguar
    - 200 cabinets, 18,688 nodes, 27 PFLOPs, 18,688 NVIDIA Kepler GPUs
  - Doubled in memory size
    - 600 TB scalar, 710 TB total (including GPUs)
  - Increased I/O requirements
    - Bandwidth and capacity

Replacing Spider I was more cost effective

Increased I/O requirements

Increasing maintenance costs

Cray XK7 Compute Node

XK7 Compute Node
AMD Opteron 6274 16 core processor @ 141 GF
Tesla K20x @ 1311 GF
Host Memory 32GB 1600 MHz DDR3
Tesla K20x Memory 6GB GDDR5
Gemini High Speed Interconnect

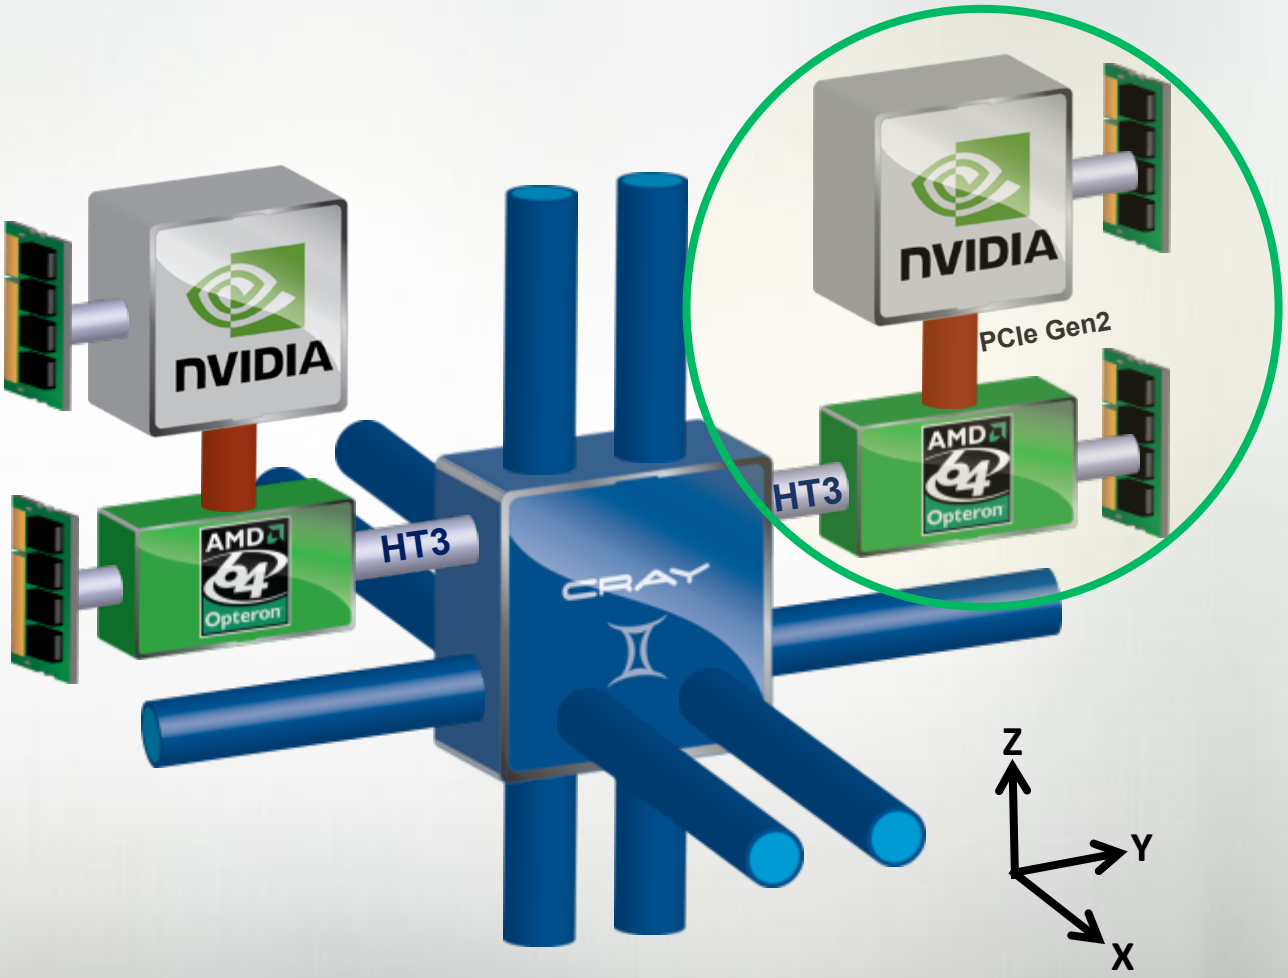

Slide courtesy of Cray, Inc.

Upgrading Spider

- Efforts started in late 2009
  - Right after Spider was deployed
  - It was a marathon, not a sprint
    - Perhaps decathlon would be a better term

Completed Efforts	On-going or Pending Efforts
Understanding Spider Evaluating storage technologies and solutions Writing and releasing the RFP Evaluation of the responses	Deployment and installation Acceptance Integration Commissioning Production

Understanding Spider – pre RFP

- Learned more about Spider and our production environment
  - "*Lessons Learned in Deploying the World's Largest Scale Lustre File System*," CUG'10
  - "*Workload characterization of a leadership class storage*," PDSW'10
  - "*Monitoring tools for large scale systems*," CUG' 10
  - "*I/O congestion avoidance via routing and object placement*," CUG'11
- Summarized findings and our comprehensive understanding
  - "*A Next-Generation Parallel File System Environment for the OLCF*, CUG'12

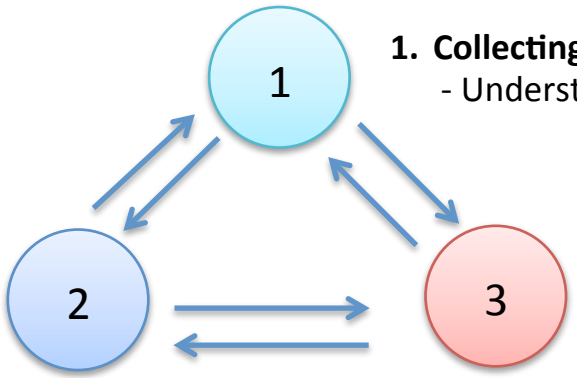

1. Collecting requirements
- Understand I/O demands

2. Design
- Architect and build storage system

3. Validation
Operation, maintenance (performance efficiency, capacity utilization)

Understanding Spider – pre RFP

Congestion is real and present!

Avoidable up to a certain degree
30% performance increase

42% read I/O workload!

Contrary to our beliefs

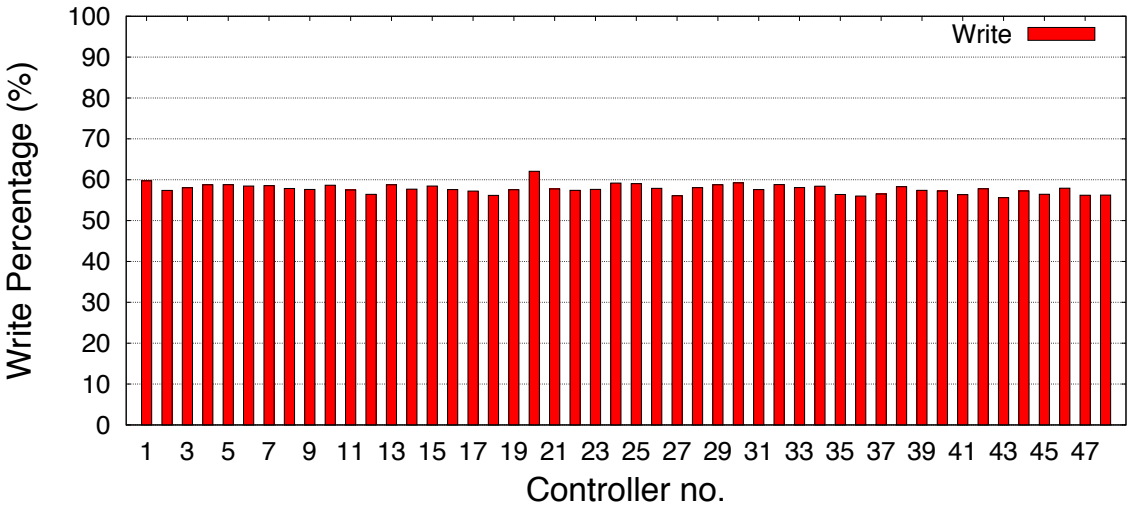

Understanding Spider – pre RFP

Substantial amount of small I/O requests

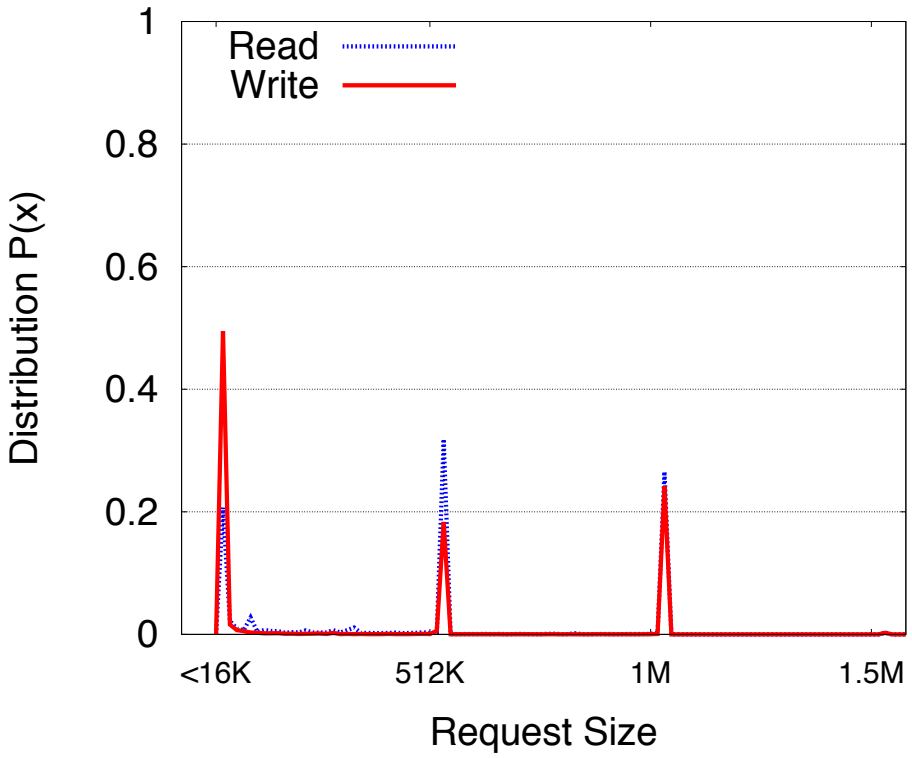

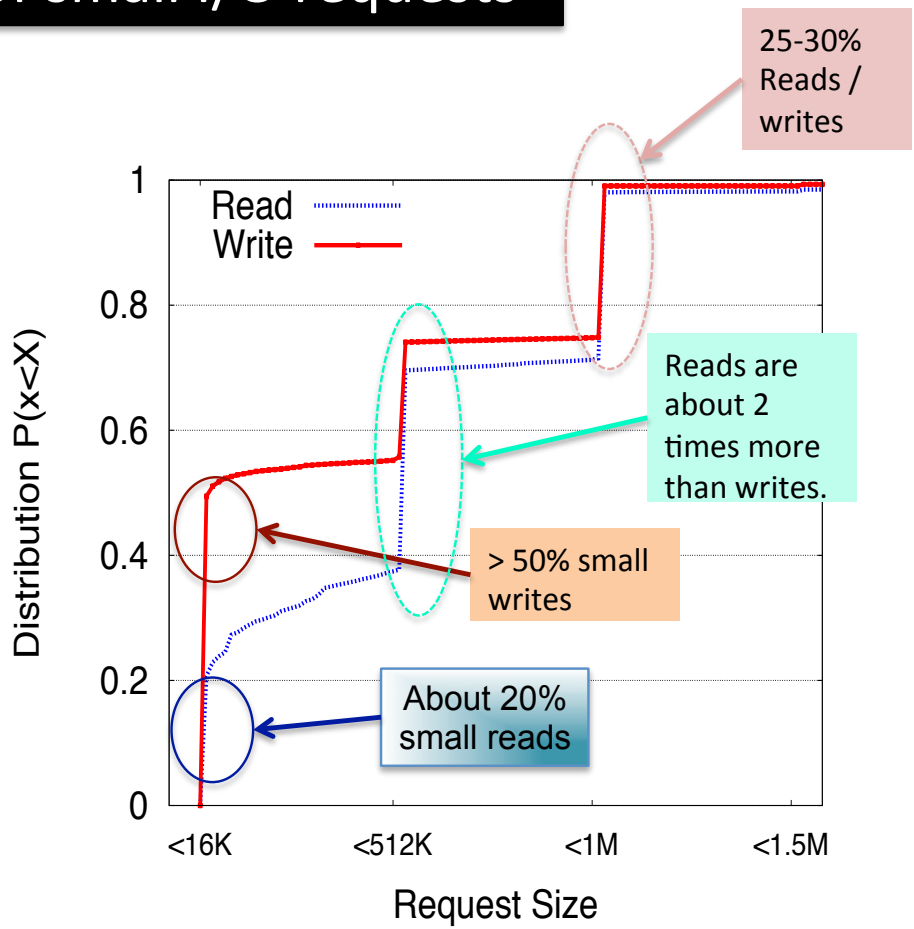

Evaluation Efforts – pre RFP

- A new benchmark suite developed
  - Block-level
    - Wrapper around fair-lio
    - Based on the *libaio* libraries
  - Lustre-level
    - Wrapper around obdfilter-survey
  - Catalogues and stores results, plots them with gnuplot
  - Released to public in 2010
  - Received positive feedback

Evaluation Efforts – pre RFP

- A new storage evaluation testbed was established
  - Testing solutions before they were GA
  - Evaluated
    - Embedded or integrated solutions
    - Block solutions
    - Host-to-storage network technologies
    - Host-to-compute network technologies

- Visited another site for more experience on a missing technology

Writing and releasing the RFP

- RFP process started in Fall of 2010
  - Gathering requirements, document writing
- Encountered some setbacks
  - Thailand flooding in July 2011 caused disk prices to spike
  - Had to wait until prices settled down
  - Budget sensitivity and continued disk price elevation near end of FY2012 caused delay into FY2013

RFP released November 2012 and responses were due December 2012

Writing and releasing the RFP

- Requirements
  - Both block-level and Integrated Lustre appliances were allowed
  - 1.2 TB/s block-level performance
    - Sequential writes and reads
  - 1.0 TB/s Lustre-level performance
    - Nice and well aligned writes and reads
  - 240 GB/s block-level random writes and reads
  - Minimum of 18 PB storage (after RAID)
  - SAS or IB FDR host-to-storage connectivity
  - Parity check on read
  - Performance under rebuild, etc

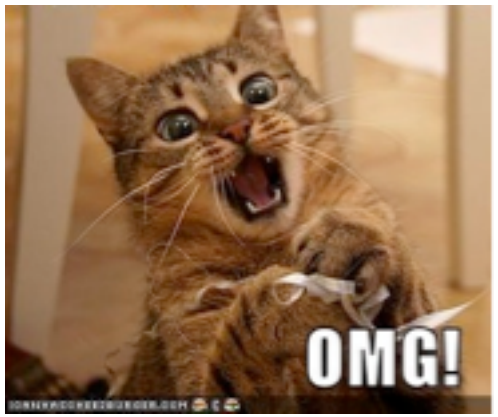

Evaluation of the Responses

- Data Direct Networks' proposal was selected
- Final Negotiated system

Scalable Storage System
36 SFA12K40 Infiniband FDR 10 60-disk enclosures per couplet 560 2 TB NL SAS drives per couplet 20,160 drives 32 PB capacity (after RAID) > 1 TB/s aggregate performance
Test and Development System
1 SFA12K40 Infiniband FDR 5 60-disk enclosure 280 2 TB NL SAS drives

DataDirect NETWORKS

Facts
32 PB capacity (after RAID)
> 1 TB/s aggregate performance
288 Lustre OSS total
8 OSS per couplet
4 MDS and 2 MGS
Configured in 4 rows
2x 108-port Core FDR IB switches
36x 36-port FDR IB switches
432 Lustre Titan LNET routers

Spider II Architecture

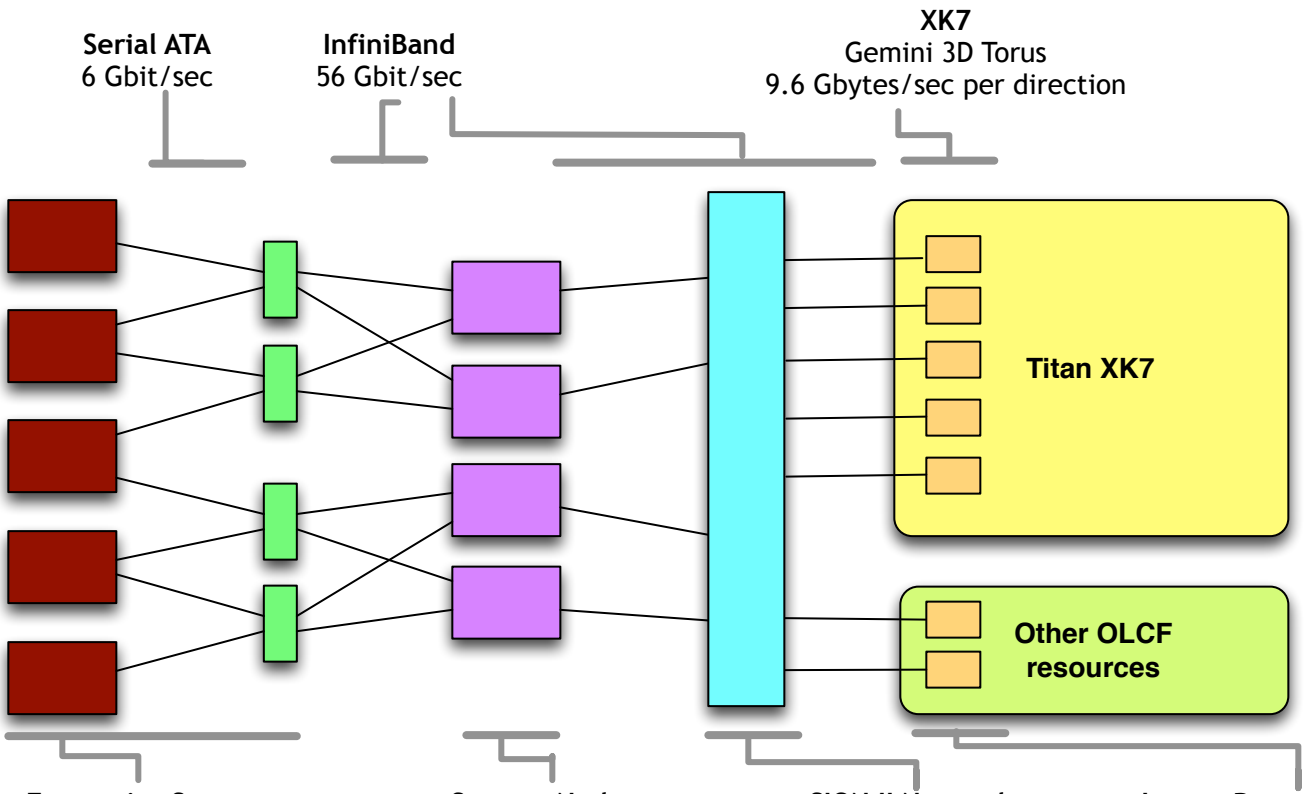

Enterprise Storage
controllers and large racks of disks are connected via InfiniBand.

36 DataDirect SFA12K-40 controller pairs with 2 Tbyte NL- SAS drives and 8 InifiniBand FDR connections per pair

Storage Nodes
run parallel file system software and manage incoming FS traffic.

288 Dell servers with 64 GB of RAM each

SION II Network
provides connectivity between OLCF resources and primarily carries storage traffic.

1600 ports, 56 Gbit/sec InfiniBand switch complex

Lustre Router Nodes
run parallel file system client software and forward I/O operations from HPC clients.

432 XK7 XIO nodes configured as Lustre routers on Titan

What are we delivering to users?

- > 1 TB/s Lustre scratch space
  - Based on Lustre 2.4
    - Latest maintenance branch
    - Includes features we want and require
      - Large stripe count
      - Distributed Namespace (DNE)
      - Metadata performance improvements
      - Imperative recovery (IR)
  - Will not be using DNE to start
    - Planning ahead to allow for this feature in the future

Integration efforts

- Lustre 2.4 testing
  - Small-scale
    - Round the clock testing for stability, regression, and performance on a single cabinet Cray XK7 (Arthur)
    - Home built Cray Lustre 2.4 client as well as servers
    - Early detection and correction of problems and bugs
  - Large-scale
    - Monthly testing of small-scale tested code drops on Titan
    - Did three tests and four more to go
    - Identified some number of problems at scale
    - Partnership with Cray
- IB FDR testing on Cray
  - Cray and Mellanox

Schedule

- System infrastructure delivery
  - Completed
- Block storage delivery
  - Started in late March and will end in early May
- Release an RFP for Lustre Support by end of April
  - Level 1, 2, and 3 support. Contract will be awarded soon
- Block acceptance
  - Starts after storage and infrastructure are installed
  - To be completed by May 31
- Complete file system integration by late August
- Commission the system by September


Networks
DataDirect


Lustre File System – Spider


DataDirect
Networks

SPIDER II

Questions?

dillowda@ornl.gov

Want to join our team?
ORNL is hiring. Contact us at
http://jobs.ornl.gov

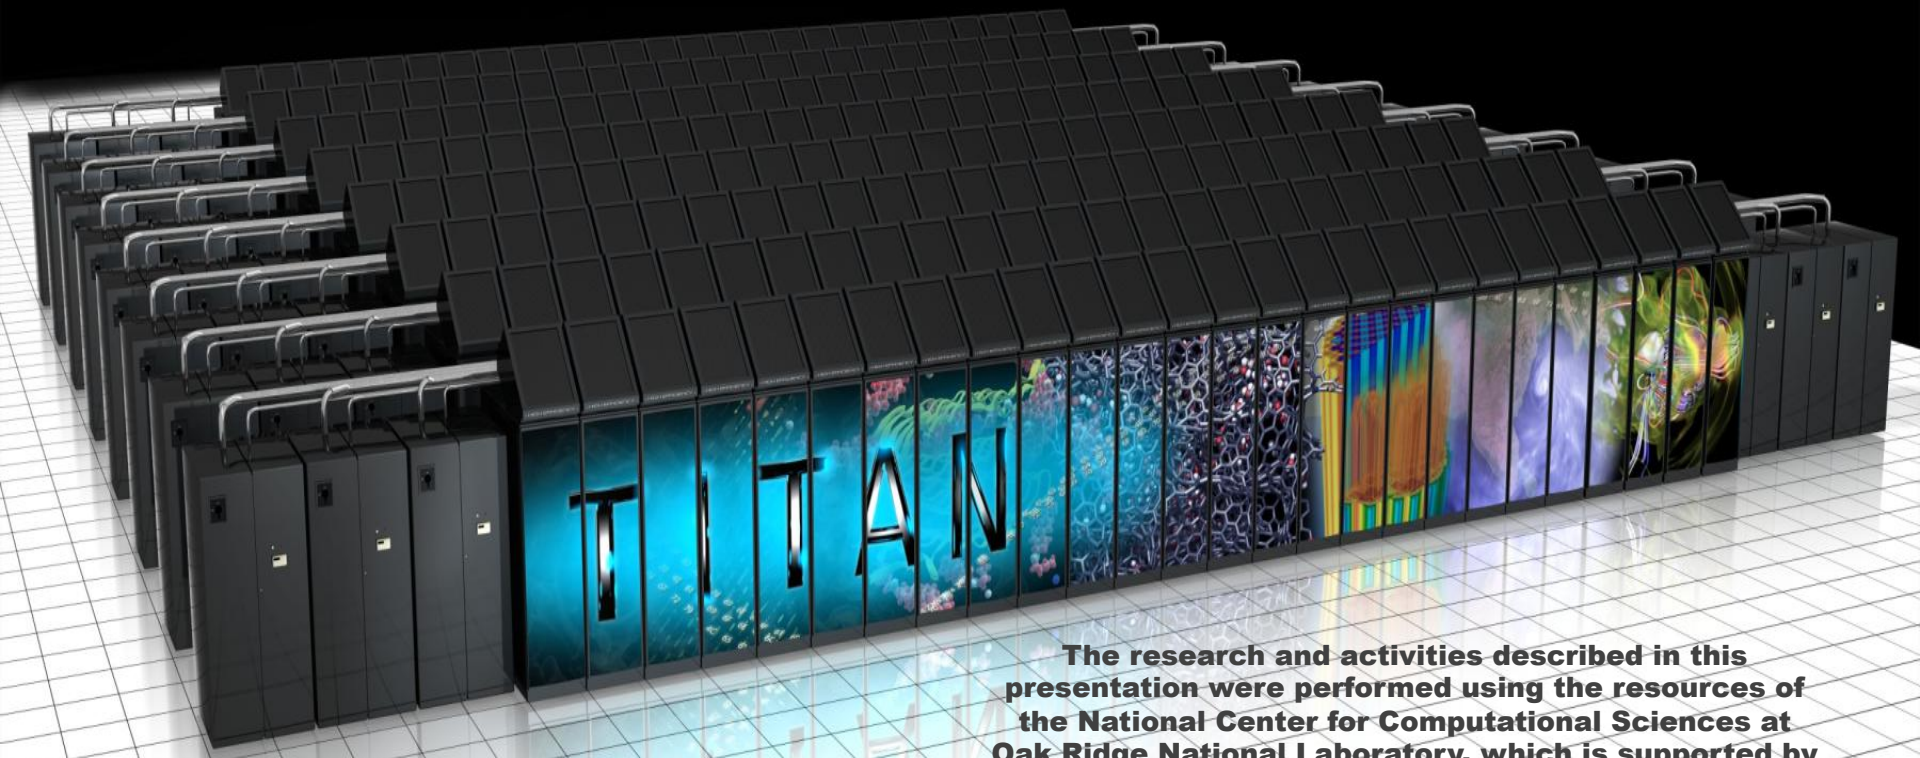

The research and activities described in this presentation were performed using the resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC0500OR22725.

OAK RIDGE National Laboratory