LA-UR-12-22037
Approved for public release; distribution is unlimited.

Title:	LANL Site Update
Author(s):	Armstrong, Harold G
Intended for:	NNSA/CEA Computing Sciences Collaboration, 2012-06-12/2012-06-14 (Bordeaux, ---, France)

LANL Site Update

Hal Armstrong

June 13, 2012

Overview

- Capability cluster - Cielo
  - ACES agreement
  - History and cluster details
  - File system update
- Capacity clusters
  - Luna - classified system cluster
  - Moonlight – unclassified TLCC2 with GPGPU cluster

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

ACES and Cielo

- ACES: The New Mexico Alliance for Computing at Extreme Scale
  - 3/2008: LANL & SNL Memorandum of Understanding
  - Joint design, architecture, development, deployment and operation of production capability systems for NNSA
  - Driven by mission needs
  - Commitment to the development and use of world class computing
  - Continued leadership in high performance computing

- ACES developed Cielo RFP requirements, including input received from presentations at LANL, SNL, and LLNL
  - Technical Evaluation Team from all three laboratories

- Cielo provides petascale production capability computing for the ASC tri-lab community (LLNL, SNL, LANL)

Cielo Platform Schedule Highlights 2009-2011

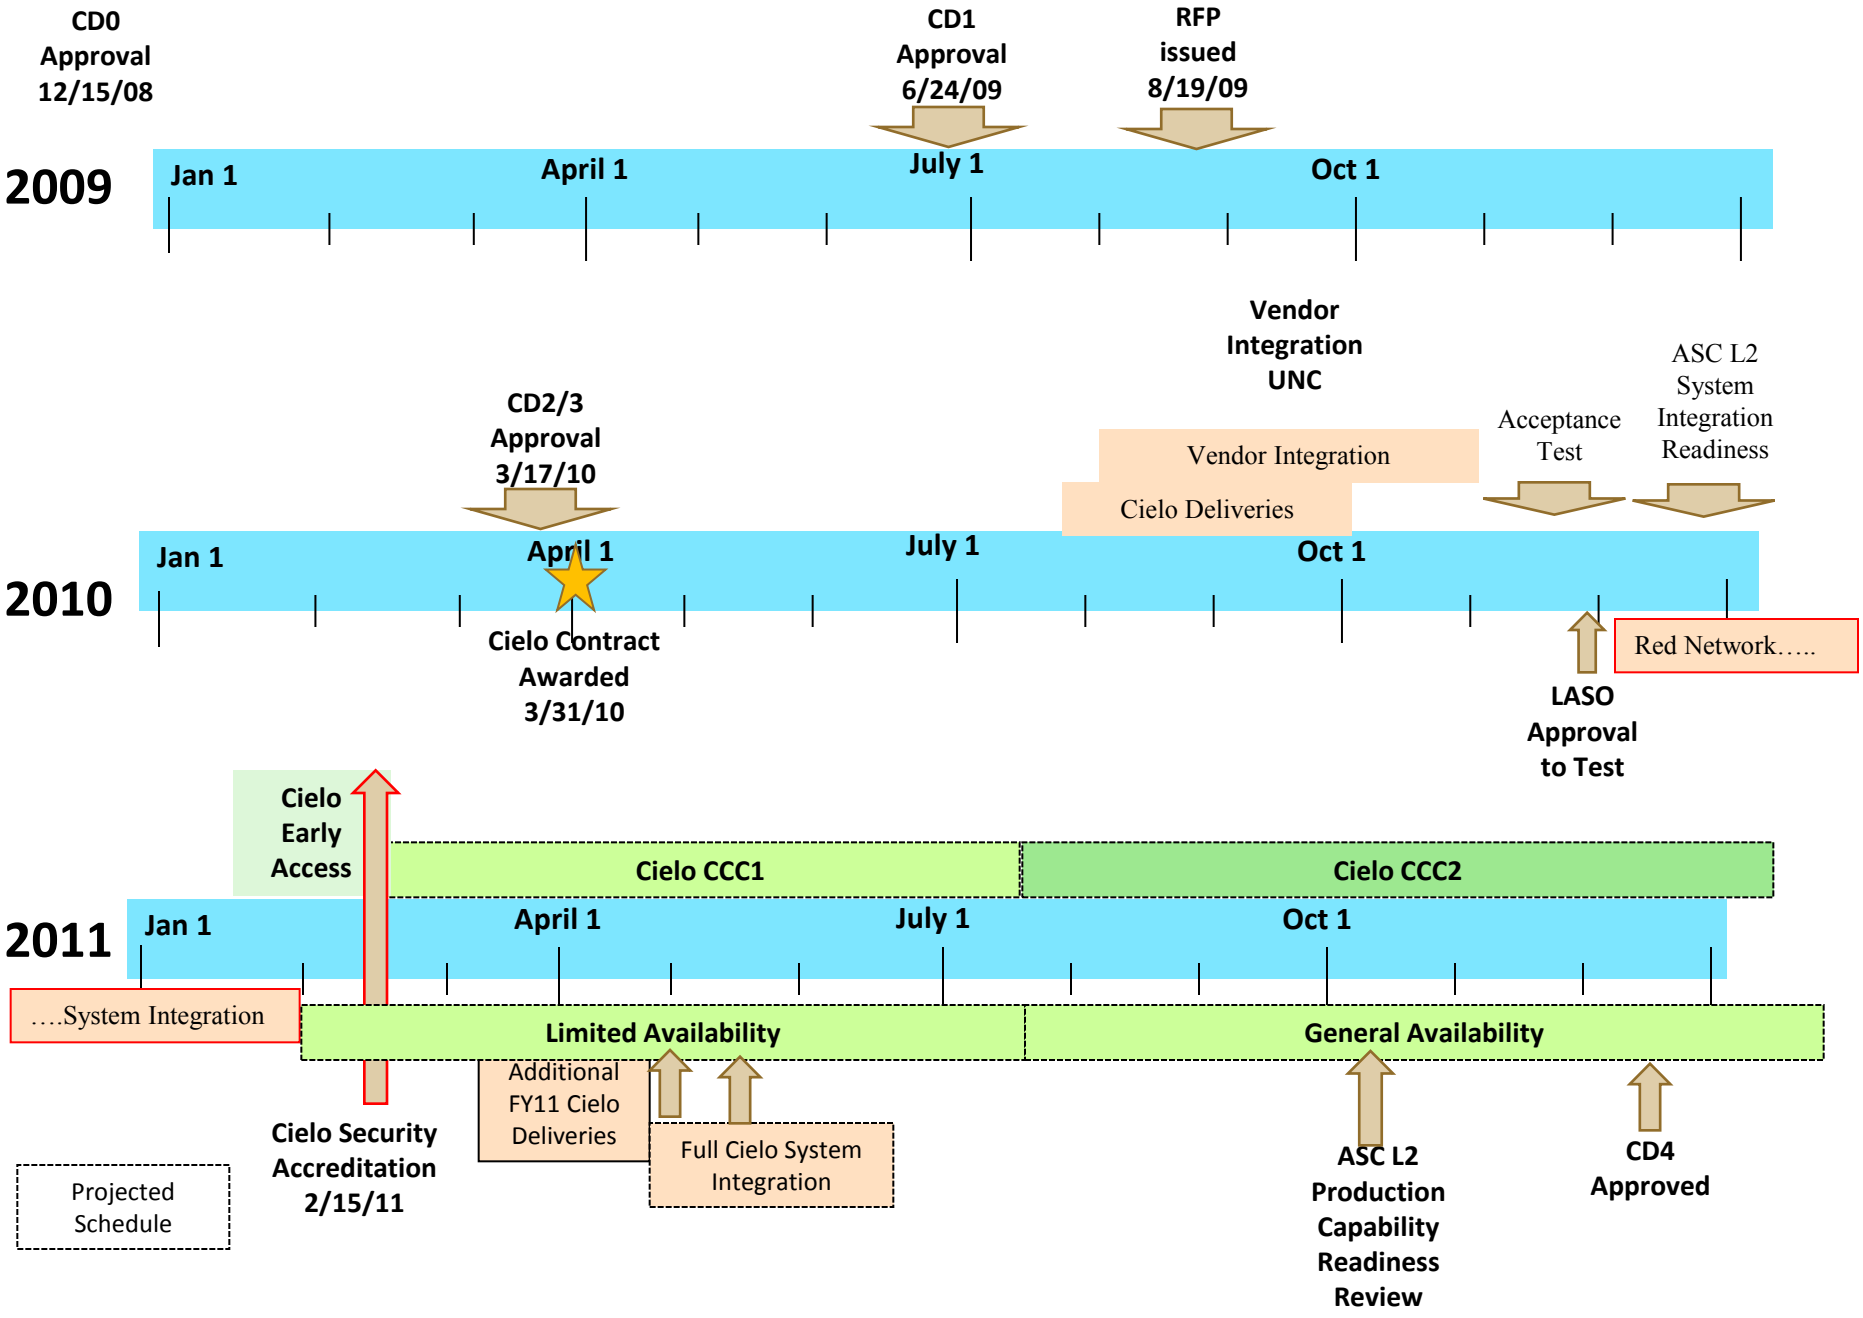

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

The Cielo Platform

- Cray, Inc. selected to deliver Cielo Platform by ACES
  - Cielo Platform (1.03 Peak PetaFlops in FY10)
    - 5% Cielo Nodes with Additional Memory (for Viz partition)
  - Additional Delivery Option (.33 additional PFs in FY11)
  - Total of 1.37 PF system
- Cielo - provide petascale production capability computing for the ASC tri-lab community (LLNL, SNL, LANL)
  - Capable of running a single application across the entire machine
  - Usage Model will follow the ASC's Capability Computing Campaign (CCC) process

Cielo Hardware Architecture

- AMD Magny-Cours Node
  - Dual-socket AMD 6136 Processors
  - 2 x 8 = 16 total cores
  - 2.4 GHz core frequency
  - 32 GB of 1333 DDR3 memory
    - 64 GB for Visualization Nodes
  - 153.6 peak DP GFLOPs
  - 85.3 peak GB/s memory BW
- Gemini High-Speed Interconnect
  - 3D Torus topology
  - Phase 1: 18x8x24
    - X bisection: > 4.38 TB/s
    - Y bisection: > 4.92 TB/s
    - Z bisection: > 3.92 TB/s
  - Phase 2: 16x12x24
    - X bisection: > 6.57 TB/s
    - Y bisection: > 4.38 TB/s
    - Z bisection: > 4.38 TB/s
  - Node Injection
    - > 6 GB/s/dir sustained BW
    - > 8 MMsgs/sec sustained

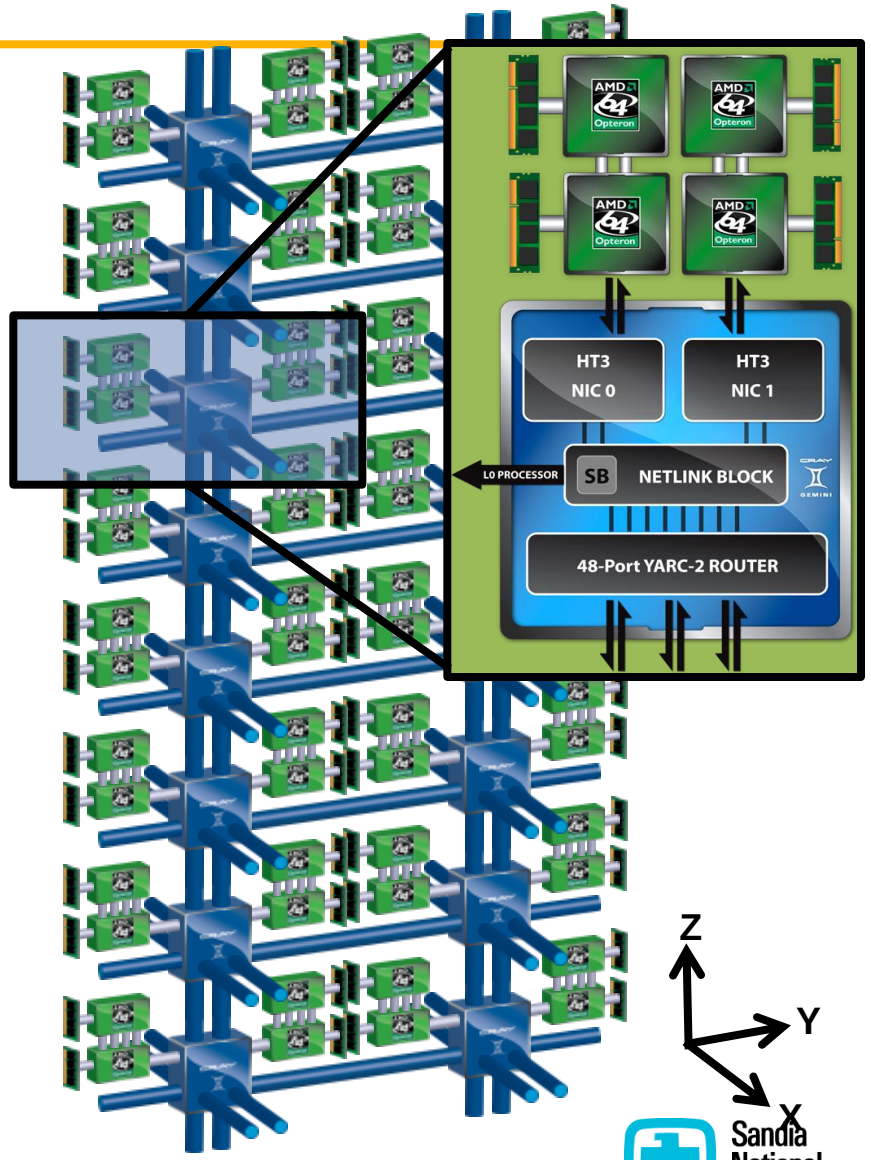

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

Cielo By Numbers

	Phase 1	Phase 2	Cielito
# of Cabinets	72	96	1
# of Service Nodes	208	272	14
# of Compute Nodes	6,704*	8,944*	68
# of Visualization Nodes	(376)	(376)	(4)
# of Compute Cores	107,264	143,104	1,088
Peak Memory BW	572 TB/s	763 TB/s	5.8 TB/s
Memory Capacity per Core	2 GB (4 GB)	2 GB (4 GB)	2 GB (4 GB)
Compute Memory Capacity	226.6 TB	298.2 TB	2.3 TB
Peak Compute FLOPS	1.03 PF	1.37 PF	10.4 TF
Sustained PFS BW	> 160 GB/s		TBD
System Power	~3 MW w/HPL	~4 MW w/HPL	
Full System Job MTBI	> 25 hours		
System MTBI	> 200 hours		

* Total compute nodes including Viz nodes and nodes allocated for other services

File system background

Cielo operational and running production workload via CCC-1 and CCC-2. However, file system reliability and performance was less than expected. Through 2011, a file system risk mitigation process was executed.

- April 13 Cielo Quarterly
  - Cielo File System
  - **Project Team and Cielo and ASC Execs initiated Risk Planning and high level interactions with Panasas**
- July 13 Cielo Quarterly
  - Cielo File System reliability and performance improvements
  - Panasas software release in June had some problems
  - Start of CCC2 in July was impacted by File system issues
  - **Risk Planning accelerated by the Project Team and Cielo and ASC Execs**
- October 25 Cielo Quarterly
  - Users have been frustrated with lower availability on Cielo due to file system problems.
  - **Risk Planning continued with development of decision criteria, etc.**

Cielo File System Risk Mitigation Process Timeline

Plan of Record: Panasas

Continued project and technical interactions

Track production usage on Cielo

Set performance and reliability targets

Test and Evaluate new Panasas software

Develop test environments

Risk Mitigation

Select Option(s)

Decision

Identify Risk Mitigation Options

Evaluate Technical Feasibility, cost, schedule, etc.

Collect data, analyze

Develop criteria for triggering (or not) risk mitigation

Complete criteria for risk mitigation

Change in Cielo File System

The ACES team, in conjunction with ASC Program Management, has decided to make a change in the file system for Cielo. Improving the reliability and performance of the Cielo file system has been the number one priority for the ACES team, which included risk mitigation planning for a possible transition. We believe that transitioning to the Lustre file system on Cielo will provide the stability and performance necessary to support the CCC projects over the next few years.

Transition to a vendor (Cray) supported Lustre-based solution with an esFS file systems infrastructure similar to that used at the National Energy Research Scientific Computing (NERSC) Center, while at the same time, preserving investment in the current Cielo file system by migrating hardware infrastructure from Panasas to Lustre.

Operated by the Los Alamos National Security, LLC for the DOE/NNSA

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

File System Transition Objectives

- Support CCC2/CCC3 production computing throughout migration
- Minimize interruptions experienced by users
  - Couple work to scheduled downtime
  - Most bring-up is strictly offline
  - Actual downtimes are to support test and validation
- Reduce risk
  - Adopt a conservative Lustre configuration
  - ACES partner has significant experience supporting Lustre
  - Operate Lustre and Panasas in parallel during stabilization period
- Effectively deploy Lustre into the Cielo Environment
  - File systems, FTA access, external node access
  - Purge clusters
  - Monitoring and management

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

File System Migration Stages

- Scratch file systems on Cielo in January
  - 2 Primary scratch file systems using Panasas
  - 1 for production (/scratch5, 5 PB)
  - 1 for development and test (/scratch4, 5PB)
- Phase 1: Upgrade test system
  - Small test file system
  - Test validation
- Phase 2: Cielo, Rebuild /scratch4 as 2 Lustre: /lscratch2, /lscratch4
  - 2 smaller file systems, each approx 2.4 PB
  - 12 racks each
- Shift CCC3 campaigns to Lustre
- Phase 3: Cielo, Rebuild /scratch5 as Lustre: /lscratch3
  - Single large file system, approximately 4.8 PB
  - 24 racks

cielo
NNSA New Mexico Alliance for Computing at Extreme Scale

Notational Transition Timeline

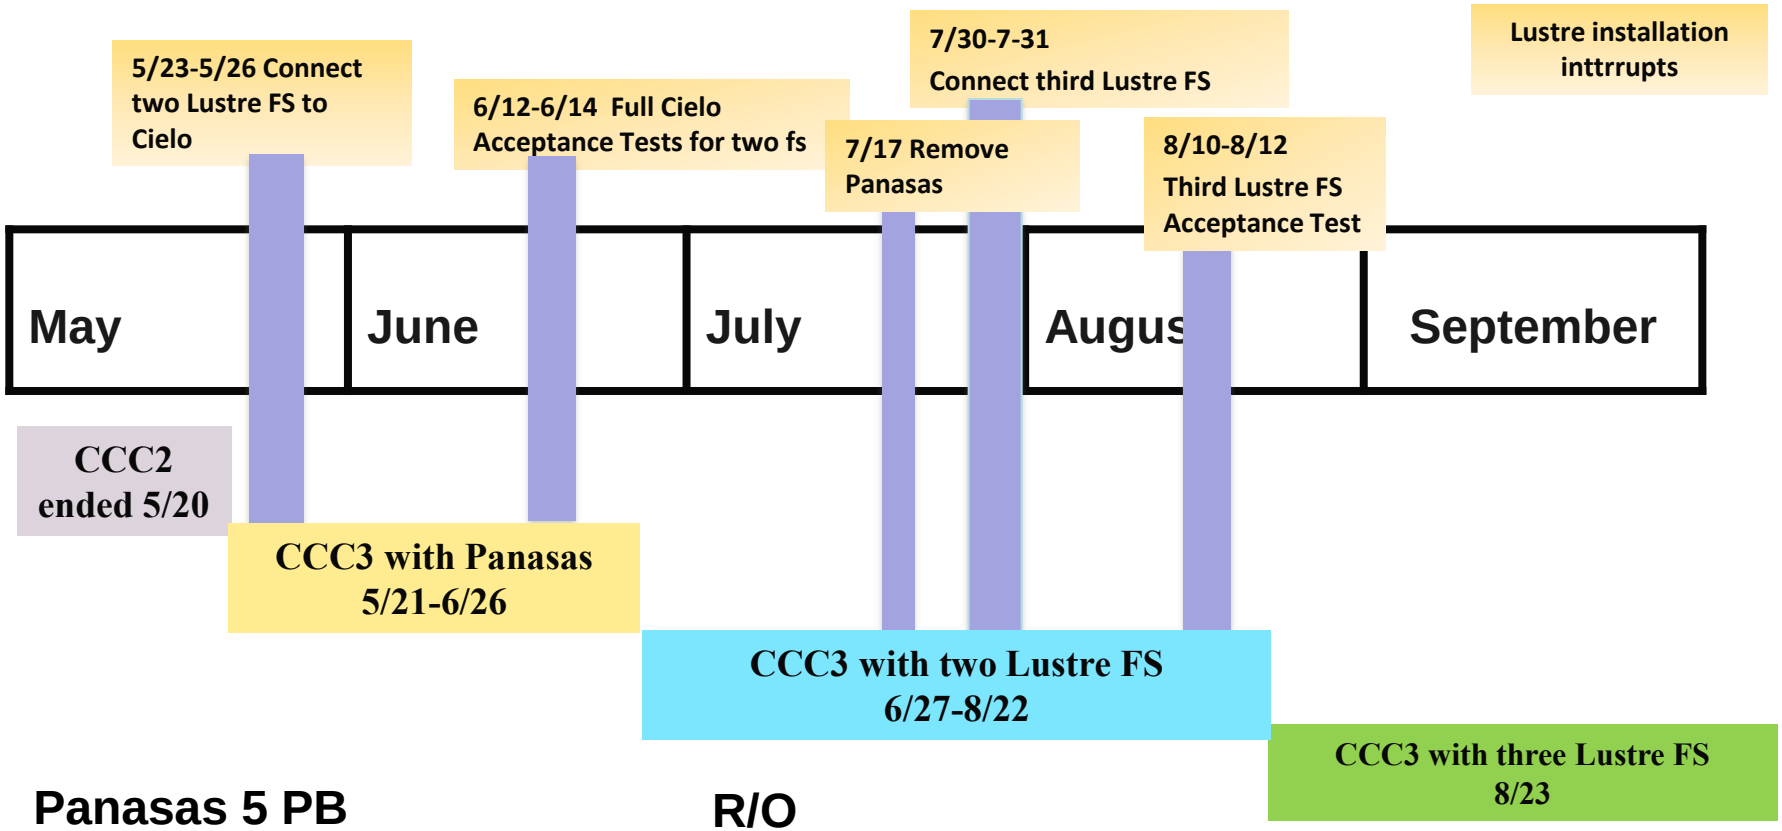

Los Alamos NATIONAL LABORATORY EST.1943

Sandia National Laboratories

NNSA

Luna & Moonlight Summary

Luna (Classified)

- Intel Sandy Bridge processor
- Peak Performance: 473 TF/s
- Peak Memory: 49,280 GB
- Compute Nodes: 1,540
- Cores: 24,640

Moonlight (Unclassified)

- Intel Sandy Bridge processor + Nvidia M2090 GPU
- Peak Performance: 488 TF/s
- Peak Memory: 9,856 GB + 3,696 GB GPU memory
- Compute nodes: 308
- Cores: 4,928 CPU + 315,392 GPU

Luna architecture

- **10 SU's (Scalable Units) (1540 Nodes, 539.1 TFlop/s peak)**
  - 154 Intel Xeon E5 Nodes (16 PEs)
    - Intel Sandy Bridge 2.6GHz
  - 32GB memory DDR3 (2GB per PE)
  - ~50TF / SU
- **Single InfiniBand 4x QDR Interconnect**
  - Qlogic
  - 324 Port IB QDR Switch
- **Delivered in two phases**
  - Phase 1 – 6 SU's – Arrived 10/27/2011
  - Phase 2 – 4 SU's – Arrived 12/12/2011
- **Running TOSS 2.0 O/S**

Moonlight architecture (w/GPGPUs)

- **2 SU's (Scalable Units) (308 nodes, 616 GPGPU nodes, ~ 500 TFlop/s peak)**
  - 154 Intel Xeon E5 Nodes (16 PEs)
    - Intel Sandy Bridge 2.6GHz
  - 32GB memory DDR3 (2GB per PE)
  - 308 NVidia M2090 Tesla GPUs (6GB / GPU)
  - ~250TF / SU
- **Single InfiniBand 4x QDR Interconnect**
  - Qlogic
- **Delivered March 12, 2012**
- **Running TOSS 2.0 O/S**

User experience with Luna

“Luna is an exceptionally well balanced machine. The ASC1 and ASC2 user communities have seen large, and largely unexpected, improvements in our throughput of DSW problems. Problems that used to take three weeks can now be completed n a few days. Thus Luna is changing the way that we do work.”

“Highly resolved simulations take only a few days.”

“Quick and dirty calculations (i.e., state-of-the-art only five years ago) finish in a day.”

Questions & Discussion