LC Systems Update

LC User Meeting

March 28, 2019

David Smith, LC System Administration Group Lead

Lawrence Livermore National Laboratory

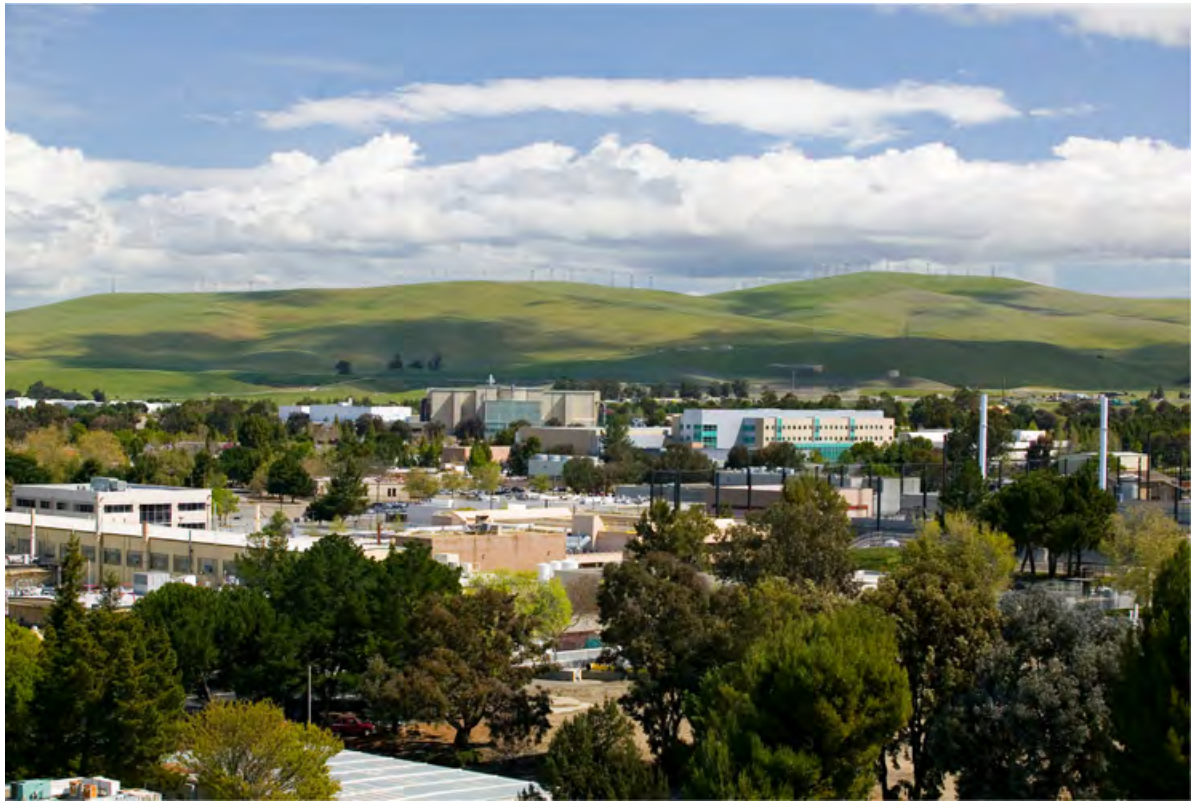

LLNL-PRES-770939

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Agenda

- New Systems
  - Quartz expansion
  - Corona
  - SCF CTS-1 system (future – under consideration)
- Systems Retiring
- HPC System Summary

Quartz Expansion (CTS-1)

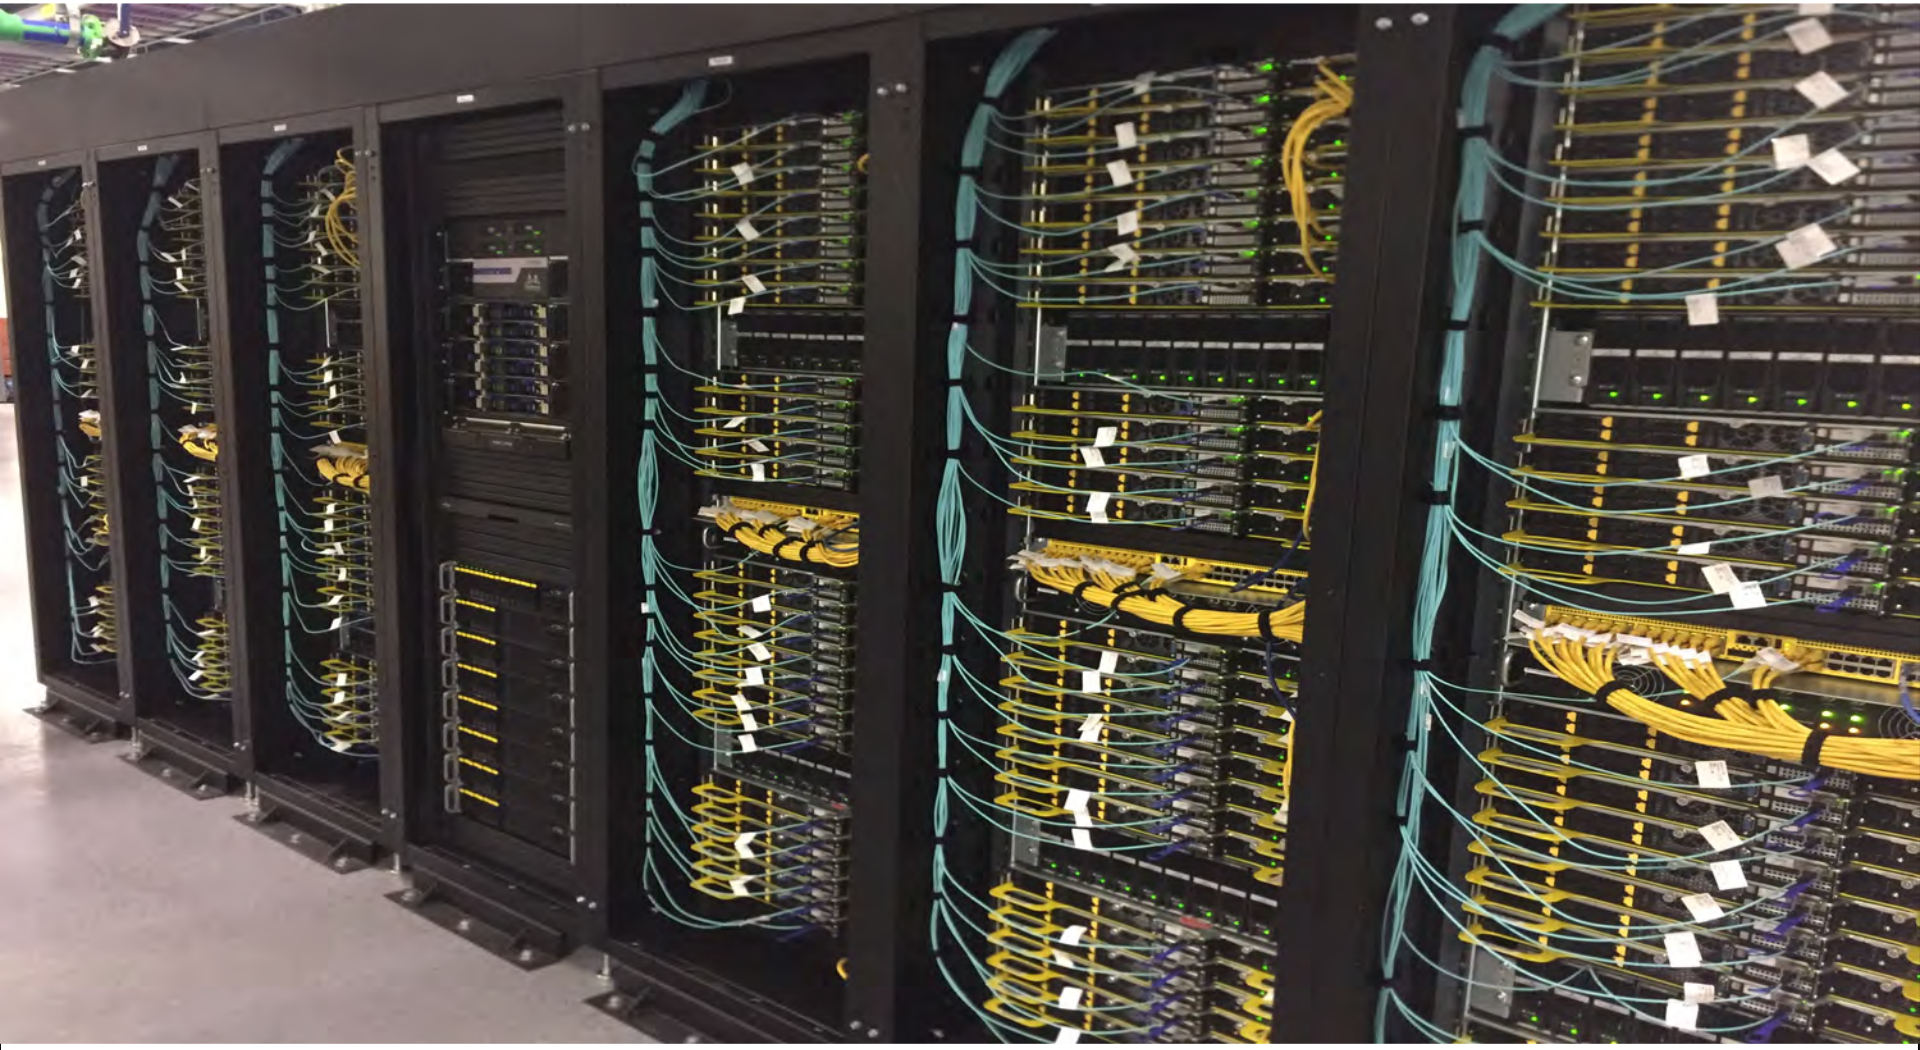

Quartz Expansion Systems Highlights

- **Quartz 2SU Expansion (~508 TF)**
  - 372 compute nodes. 2 mgmt, 2 login, 4 GW, 4 router

  Compute Node contains:
  - Dual socket Intel Skylake Gold 6140 CPUs, 18 cores @ 2.3 GHz per socket (CTS1 CPUs)
  - 192 GB memory (6 memory channels vs 4 on older CTS nodes = more memory bandwidth)
  - Intel Omni-Path 100GB/s interconnect

  Quartz 2U Expansion: End of May

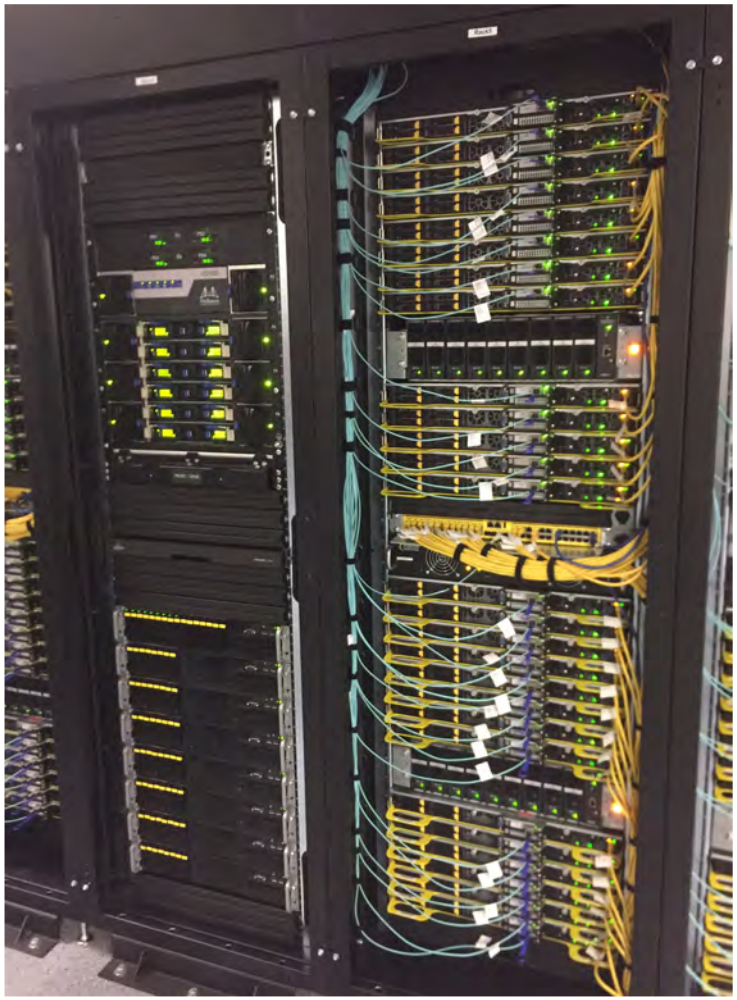

CTS Highlights

- Compute Nodes
  - Dual Socket Intel E5-2695 v4 (Broadwell)
  - 18 cores @ 2.1GHz per socket (36 cores/node)
  - 128GB DDR4 @ 2.4GHz DRAM
- Intel Omni-Path 100GB/s interconnect
- 1SU = 192 total nodes, 232 TF/s, 23.5TB DRAM
  - 1 Management node
  - 1 Login node
  - 4 Network router/gateway nodes
  - 186 Compute nodes

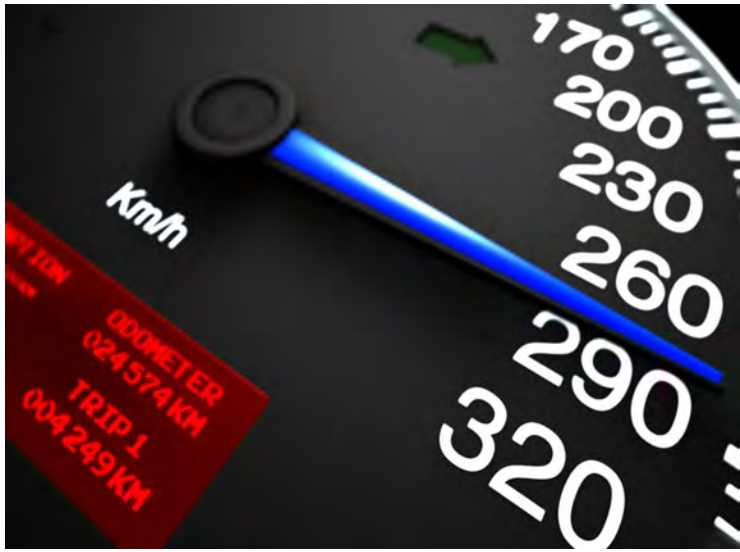

CTS Summary

Name	Size	Network
Borax	48 nodes (serial)	CZ
Quartz	14 SU → 16SU	CZ
Pascal	~1 SU	CZ
RZTopaz	4 SU	RZ
RZTrona	20 nodes (serial)	RZ
Jade	14 SU	SCF
Mica	2 SU	SCF
Agate	48 nodes (serial)	SCF
???	2 SU (??)	SCF • Future (in review) • Power/Cooling

Systems Retiring

Retiring Systems	Retirement Date	Replacement System
Vulcan	02/20/2019 11/30/2018 (original)	Lassen RZAnsel
Zin*	TBD	Jade
Sequoia*	TBD	Sierra

The Life of Zin

- **Extending service of Zin**
  - **Zin's environment has been very stable**
  - **Service life is typically: 5 years (Zin is 7.5 years)**
- **Zin timeline of operation**
  - **TOSS 3 (running TOSS 3 since November 2018)**
    - **Continue to use system (system node count may be reduced for parts)**
- **Retirement – subject to:**
  - **Hardware issues**
  - **Floor space needed for new system deployments**
- **Zin Facts**
  - **First received October 2011 (Tri-Lab Linux Capacity Cluster (TLCC2))**
  - **18 Scalable Units**
  - **970 teraFLOPs**

LC HPC System Summary – March 2019
(https://hpc.llnl.gov/hardware/platforms)

System	Rank	Program	/ Model	OS	connect	Nodes	Cores	(GB)	TFLOP/s
Unclassified Network (OCF)									
Lassen	11	ASC+M&IC	IBM P9	RHEL	2x IB EDR	684	30,096	218,880	19,886.0
Quartz	71	ASC+M&IC	Penguin	TOSS	Omni-Path	2,688	96,768	344,064	3251.4
Pascal		ASC+M&IC	Penguin	TOSS	IB EDR	163	5,868	41,728	1,700
RZTopaz		ASC	Penguin	TOSS	Omni-Path	768	27,648	98,304	929.0
RZManta		ASC	IBM P8	RHEL	IB EDR	36	720	11,520	597.6
Ray		ASC+M&IC	IBM P8	RHEL	IB EDR	54	1,080	17,280	896.4
RZAnsel	247	ASC	IBM P9	RHEL	2x IB EDR	54	2,376	17,280	1570.0
Catalyst		ASC+M&IC	Cray	TOSS	IB QDR	324	7,776	41,472	149.3
Syrah		ASC+M&IC	Cray	TOSS	IB QDR	324	5,184	20,736	107.8
Surface		ASC+M&IC	Cray	TOSS	IB FDR	162	2,592	41,500	451.9
Borax		ASC+M&IC	Penguin	TOSS	N/A	48	1,728	6,144	58.1
RZTrona		ASC	Penguin	TOSS	N/A	48	1,728	6,144	58.1
OCF Totals	Systems	13							29,655.6
Classified Network (SCF)									
Pinot(SNSI)		M&IC	Penguin	TOSS	Omni-Path	187	6,732	23,936	232.2
Sequoia	10	ASC	IBM BGQ	RHEL/CNK	5D Torus	98,304	1,572,864	1,572,864	20132.7
Sierra	2	ASC	IBM P9	RHEL	2x IB EDR	4,320	190,080	1,382,400	125626.0
Zin (TLCC2)		ASC	Appro	TOSS	IB QDR	2,916	46,656	93,312	961.1
Jade+Jadeita	70	ASC	Penguin	TOSS	Omni-Path	2,688	96,768	344,064	3251.4
Mica		ASC	Penguin	TOSS	Omni-Path	384	13,824	49,152	464.5
Shark		ASC	IBM P8	RHEL	IB EDR	36	720	11,520	597.6
Max		ASC	Appro	TOSS	IB FDR	324	5,184	82,944	107.8
Agate		ASC	Penguin	TOSS	N/A	48	1,728	6,144	58.1
SCF Totals	Systems	9							151,431.4
Combined Totals		22							181,087.0

System Category	TFLOP/s	% of Total
Unclassified	29,655.6	16.4%
Capability	22,950.0	77.4%
Capacity	4,437.5	15.0%
Visualization	2,151.9	7.3%
Serial	116.2	0.4%
Classified	151,431.4	83.6%
Capability	146,356.3	96.6%
Capacity	4,909.2	3.2%
Visualization	107.8	0.1%
Serial	58.1	0.0%

Other

- **Vendor discussions have started on next generation of systems (CTS-2).**
- **VNC – deployment as a system**
  - **OCF – Completed**
  - **SCF – May 2019**
- **SNSI Environment – Pinot**
  - **Lustre file system will be replaced (July 2019)**
  - **CTS-1 based system installed Oct 2018**
  - **NFS server updated Oct 2018**

Questions?

David Smith
smith107@llnl.gov
925-422-9256

Lawrence Livermore National Laboratory

Disclaimer
This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.