LC Systems Update

LC User Meeting

December 8, 2020

David Smith, LC System Administration Group Lead

Lawrence Livermore National Laboratory

LLNL-PRES-817519

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

- OCF Systems
- SCF Systems
- HPC System Summary
- Storage
- Remote Visualization Tools

Corona Expansion Highlights (11.3 PF CPUs+GPUs)

- User Nodes
  - Login nodes: 3
  - Batch nodes: 285
  - Debug: 0
  - GPU nodes: 285
- CPUs (dual socket / node)
  - AMD EPYC 7401 & 7402
    - 48 core
    - 2.0 GHz (7401)
    - *2.8 GHz (7402)*
  - 96 cores / node
- GPUs (Radeon Instinct mi50 and mi60)
  - 82 - 4 x mi50, 82 - 4 x mi60
  - *123 - 8 x mi50 (AMD EPYC 7402)*
- Memory
  - 256 GB / node

- Interconnect
  - IB HDR
- Parallel job type:
  - multiple nodes per job

Ruby (5.9 PF)

User Nodes
  -- Login nodes: 8
  -- Batch nodes: 1480
  -- Debug: 24
CPUs (dual socket / node)
  -- Intel Xeon CLX-8276L
    - 2.2 GHz
  -- 56 cores / node
Memory
  -- 192 GB / node
Interconnect
  -- Omni-path
Parallel job type:
  -- multiple nodes per job

Mammoth (0.3 PF)

- User Nodes
  - Login nodes: 2
  - Batch nodes: 64
  - Debug: 0
- CPUs (dual socket / node)
  - AMD EPYC 7752 64 core
    - 2.250 GHz
  - 128 cores / node
- Memory
  - 2048 GB / node
- Interconnect
  - Omni-path
- Parallel job type:
  - multiple nodes per job

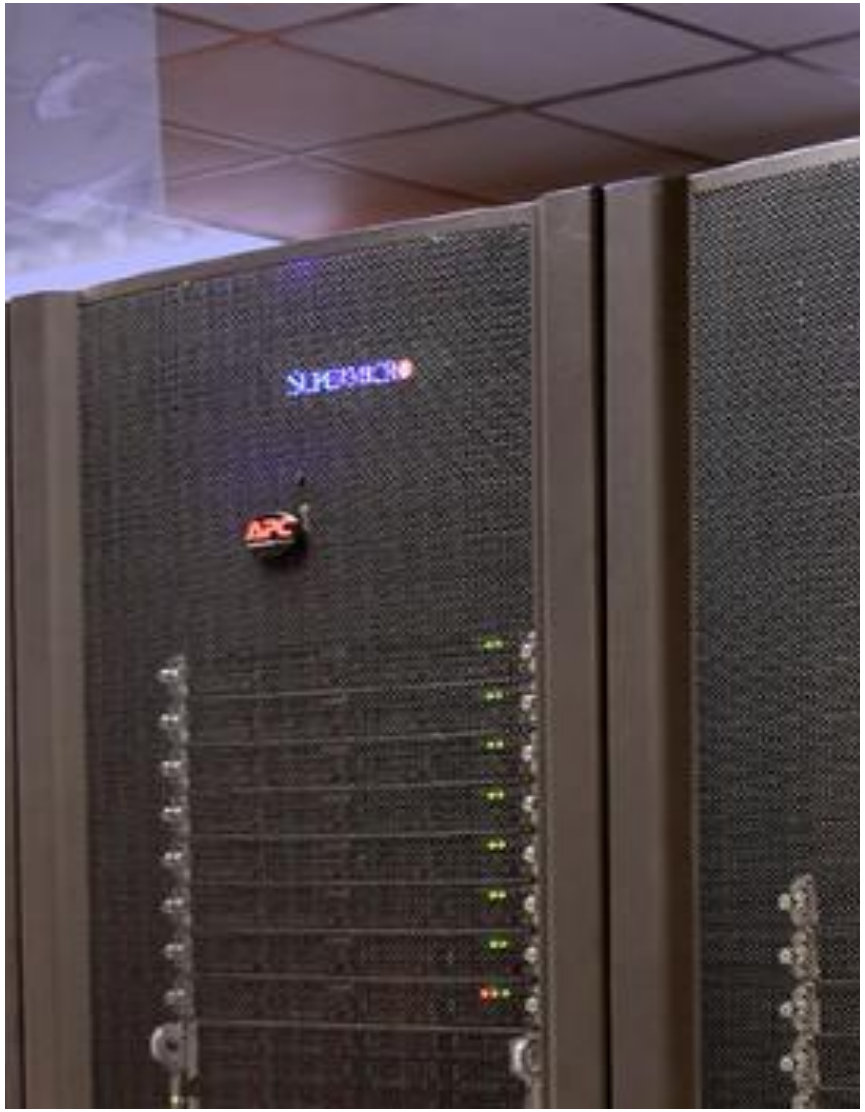

Boraxo (58.1 TF)

- User Nodes
  - Login nodes: 2
  - Batch nodes: 53
  - Debug: 4
- CPUs (dual socket / node)
  - Intel Gold 6140
    - 2.3 GHz
  - 36 cores / node
- Memory
  - 192 GB / node
- Interconnect
  - N/A
- Parallel job type:
  - multiple nodes per job

SCF -- Magma (5.3 PF)

- User Nodes
  - Login nodes: 8
  - Batch nodes: 732
  - Debug: 20
- CPUs (dual socket / node)
  - Intel Cascade Lake AP 9242
    - 2.3 GHz
  - 96 cores / node
- Memory
  - 384 GB / node
- Interconnect
  - Omni-path
- Parallel job type:
  - multiple nodes per job

SCF -- Tron (434 TF)

- User Nodes
  - Login nodes: 3
  - Batch nodes: 144
- CPUs (dual socket / node)
  - Intel Cascade Lake
    - 2.3 GHz
  - 32 cores / node
- Memory
  - 384 GB / node
- Interconnect
  - Mellanox EDR

LC HPC System Summary – December 2020 (https://hpc.llnl.gov/hardware/platforms)

	Rank	Program	/ Model	OS	connect	Nodes	Cores	(GB)	TFLOP/s
Unclassified Network (OCF)									
Lassen	17	ASC+M&IC	IBM P9	RHEL	2x IB EDR	792	34,848	253,440	**23,047.1**
Quartz	135	ASC+M&IC	Penguin	TOSS	Omni-Path	3,136	112,896	401,408	**3793.3**
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		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**
Mammoth		ASC+M&IC	Supermicro	TOSS	Omni-Path	69	8,832	131,072	**294.0**
Ruby	79	ASC+M&IC+CARES	Supermicro	TOSS	Omni-Path	1,512	84,672	290,304	**5959.2**
Corona		ASC+M&IC+CARES	Penguin, Super	TOSS	IB HDR	291	13,968	127,488	**11335.0**
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	15							**50,946.8**
Classified Network (SCF)									
Pinot(SNSI)		M&IC	Penguin	TOSS	Omni-Path	187	6,732	23,936	**232.2**
Sierra	3	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	134	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**
Magma	94	ASC	Penguin	TOSS	Omni-Path	722	69,312	296,448	**5454.0**
Shark		ASC	IBM P8	RHEL	IB EDR	36	720	11,520	**597.6**
Tron		ASC	Supermicro	TOSS	IB EDR	146	4,672	56,064	**433.6**
Agate		ASC	Penguin	TOSS	N/A	48	1,728	6,144	**58.1**
SCF Totals	Systems	9							**137,078.5**
Combined Totals		24							**188,025.3**

Storage

OCF and SCF home directory and project space refresh

-- Increases I/O operational capabilities

-- Allows for transparent data migration reducing downtime

-- OCF completed, SCF estimated completion is January

VAST (/p/vast1)

-- Provides ~5 PB of persistent storage

-- NFS mounted but can be used like Lustre

-- Currently directory-based quotas of 20 TB

-- User-level quotes in CY21

Lustre – Adaptive Storage Platform (ASP)

-- New deployments start in CY21

-- Larger capacity, increased quota limits and improved performance

-- Progressive file layouts – automatically stripe files across multiple disk volumes depending on file size

Remote Visualization Tools

- RealVNC
  - Deployed on CZ, RZ, SCF (czvnc/rzvnc/scfvnc)
  - VNC servers strained at current load
    - Seen up to 204 simultaneous running desktops per server
    - New RealVNC Session Manager in CY21 will load balance across multiple servers
    - LC working with RealVNC on development contract to adapt product for RSA OTP
- NICE DCV
  - Support GL/3D graphics in virtual desktop
  - Updated version supports standalone clients on Windows, MacOS, and Linux
    - See https://hpc.llnl.gov/software/visualization-software/vnc-nice-dcv for instructions
    - Default method remains via web browser

Questions?

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

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