

# Visualization at LLNL

April 22, 2014



**Richard Cook**

Information Management, Graphics and Security Group

**Integrated Computing and Communications Department**

Lawrence Livermore National Laboratory

# Visualization and Data Expertise at LLNL

## Data Group

- Provides data analysis and visualization support to customers
- Develops VisIt, a data analysis and visualization code



## Information Management, Graphics, and Security Group (HPC Center)

- Large scale data exploration
- Video production
- Maintains PowerWalls
- Develops software for searching and managing scientific data
- Provides easy-to-use interfaces to HPC resources

## Center for Applied Scientific Computing

- Develops visualization techniques for large scale data exploration
- Funded by the ASC program, LDRD, and others
- Collaborates with other lab researchers, academia, and industry



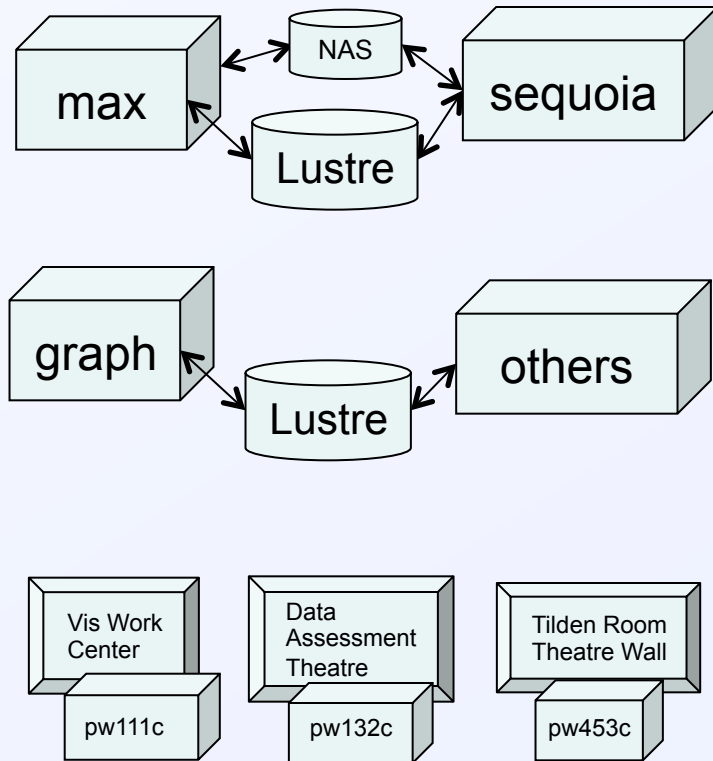
## Current LLNL Visualization Environment

- Three large visualization clusters
- Five small display drivers
- Linux with same admin support as compute clusters
- Four machine rooms
- Users access clusters over the network from workstations
- VNC is not used for visualization

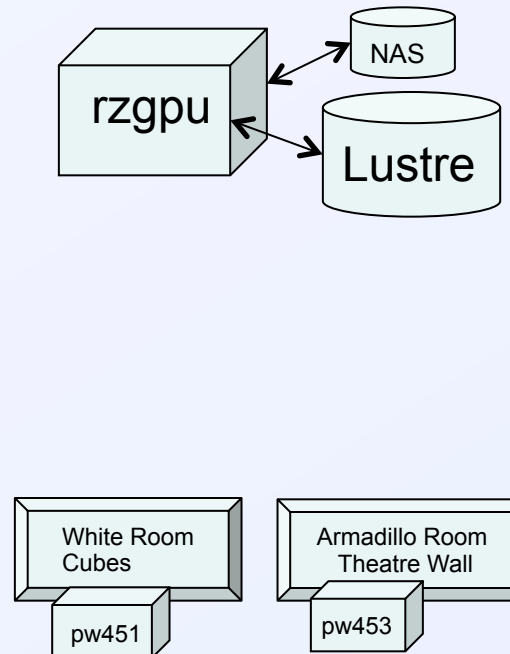



# Visualization Resources

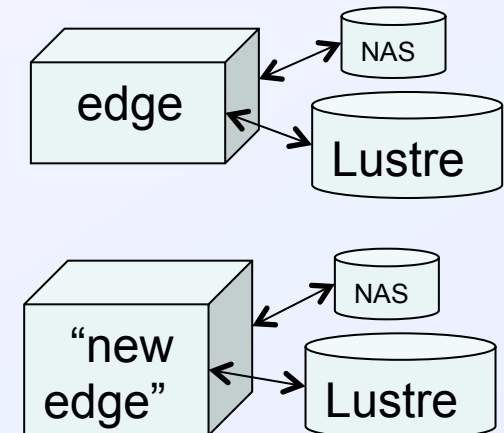
## Secure Computing Facility



## Restricted Zone “unclassified internal”



## Collaboration Zone “classified external”



# Advanced Technology Platforms: Sequoia and Vulcan

- 20 PF/s target
- 1.6 PB memory
- 1.6M cores
- 9.6MW Power, 4,000 ft<sup>2</sup>
- Hybrid cooled



## Vulcan (5 PF/s)

- Became available to users last May

## Sequoia

- Became available to users last summer



## Max: 300-node Sequoia Data Analysis Cluster

- 2 Scalable Units (TLCC2-like)
- 64GB/s bandwidth to Lustre
- 280 compute nodes
  - 16 Cores at 2.6GHz per node
  - 256 GB RAM per node
- 20 additional compute nodes
  - 2 Kepler K20X per node
  - 6GB RAM per Kepler
- Qlogic QDR Infiniband



## Specs For Vis Clusters and Wall Drivers

Cluster	Class	CPU	FLOPs	RAM	GPU
edge.llnl.gov	UC	206 x 2.8GHz Intel Xeon	29.0 TF	96G	2 Tesla M2050s per node
edge replacement	UC	160 nodes		256G	Kepler K20X
rzgpu.llnl.gov	UC	48 x 2.8GHz Intel Xeon	7.3 TF	96G	2 Tesla M2070s per node
graph.llnl.gov	C	576 x 2.0GHz AMD Operon	110.6 TF	128G	None
max.llnl.gov	C	302 x 2.6GHz Intel Xeon	107.0 TF	256G	40 Kepler K20X's

**Powerwalls each driven by single node with NVIDIA Quadroplex 7000**

## Vis hardware and software at LLNL

- Goodbye DMX, hello Quadroplex!
- GPGPU work might actually be gaining some steam this year as teams start to feel the memory pinch and exascale looms.
- GPU hardware is currently mostly doing non-vis stuff.
- New Kepler cluster is on its way. CUDA 6 looks nice.
- Blockbuster continues to be movie player of choice on the wall thanks mainly to sidecar.
- Streaming Movie format 3 now supports metadata.



# HPC Hardware at Lawrence Livermore

System	Top500 Rank	Program	Manufacture/ Model	Processor Architecture	OS	Inter-connect	Avg Power Demand (KW)	Nodes	Cores	Memory (GB)	Peak TFLOP/s
<b>Unclassified Network (OCF)</b>											
Vulcan	9	ASC+M&IC+HPCIC	IBM BGQ	IBM PowerPC A2	RHEL/CNK	5D Torus	TBD	24,576	393,216	393,216	5,033.2
Sierra	263	M&IC	Dell	Intel Xeon EP X5660	TOSS	IB QDR	TBD	1,944	23,328	46,656	243.7
Cab (TLCC2)	94	ASC+M&IC+HPCIC	Appro	Intel Xeon E5-2670	TOSS	IB QDR	564	1,296	20,736	41,472	426.0
Ansel		M&IC	Dell	Intel Xeon EP X5660	TOSS	IB QDR	TBD	324	3,888	7,776	43.5
RZMerl (TLCC2)		ASC+ICF	Appro	Intel Xeon E5-2670	TOSS	IB QDR	TBD	162	2,592	5,184	53.9
RZzeus		M&IC	Appro	Intel Xeon E5530	TOSS	IB DDR	143	267	2,144	6,408	20.6
Edge		M&IC	Appro	Intel Xeon EP X5660	TOSS	IB QDR	165	216	2,592	20,736	239.9
Aztec		M&IC	Dell	Intel Xeon EP X5660	TOSS	N/A	TBD	96	1,152	4,608	12.9
Herd		M&IC	Appro	AMD Opteron 8356, 6128 Intel Xeon E5-2670	TOSS	IB DDR	7	9	256	1,088	1.6
OCF Totals	Systems	9									<b>6,075.3</b>
<b>Classified Network (SCF)</b>											
Pinot(TLCC2, SNSI)		M&IC	Appro	Intel Xeon E5-2670	TOSS	IB QDR	TBD	162	2,592	5,184	53.9
Sequoia	3	ASC	IBM BGQ	IBM PowerPC A2	RHEL/CNK	5D Torus	TBD	98,304	1,572,864	1,572,864	20132.7
Zin (TLCC2)	41	ASC	Appro	Intel Xeon E5-2670	TOSS	IB QDR	TBD	2,916	46,656	93,312	961.1
Juno (TLCC)	399	ASC	Appro	AMD Opteron 8354	TOSS	IB DDR	600	1,152	18,432	36,864	162.2
Muir		ICF	Dell	Intel Xeon EP X5660	TOSS	IB QDR	TBD	1,296	15,552	31,104	168.0
Graph		ASC	Appro	AMD Opteron 8423	TOSS	IB DDR	429	576	13,824	72,960	107.5
Max		ASC	Appro	Intel Xeon E5-2670	TOSS	IB QDR	TBD	324	5,184	82,944	107.8
Inca		ASC	Dell	Intel Xeon EP X5660	TOSS	N/A	TBD	100	1,216	5,120	13.5
SCF Totals	Systems	8									<b>21,706.7</b>
Combined Totals		17									<b>27,782.0</b>

Unclassified	6,075.3	21.9%
Capability	5,276.9	86.9%
Capacity	544.0	9.0%
Visualization	239.9	3.9%
Serial	14.5	0.2%

Classified	21,706.7	78.1%
Capability	20,132.7	92.7%
Capacity	1,345.2	6.2%
Visualization	215.3	1.0%
Serial	13.5	0.1%

DOE Computer Graphics Forum

Blue == Visualization

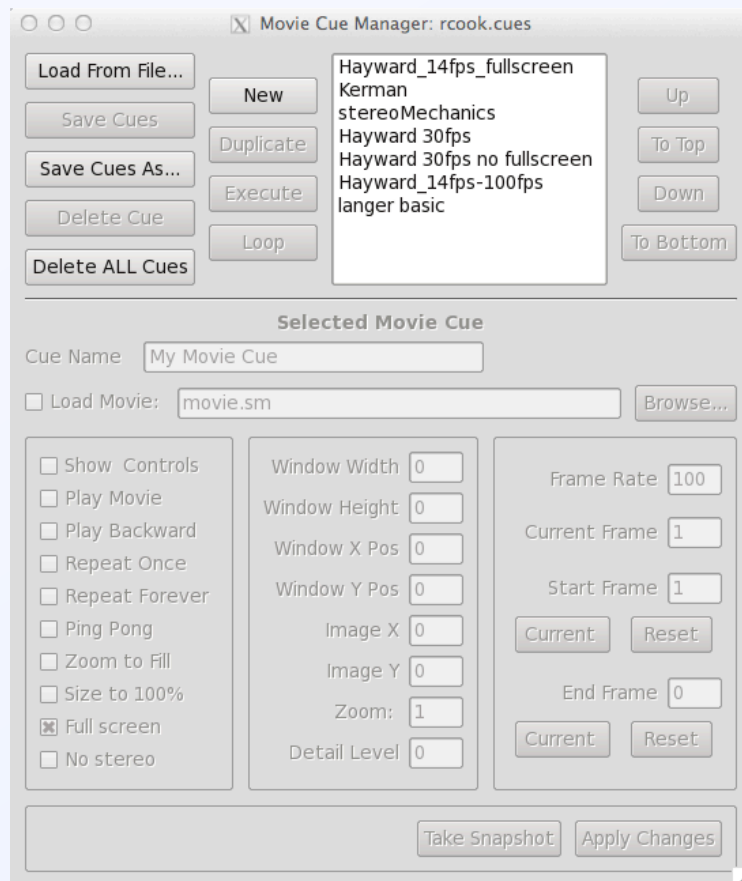


# Installed Software

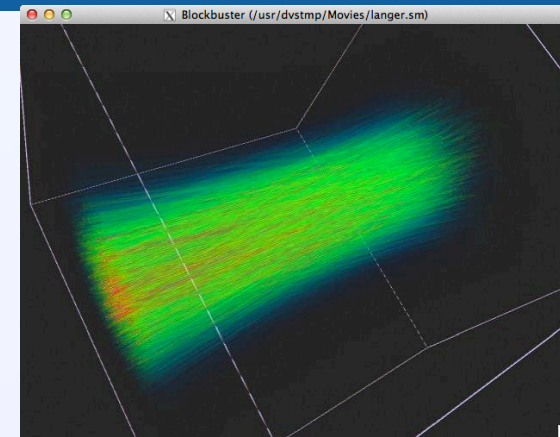
Software Name	Version	# of licenses	Yearly cost	Usage rate
<b>AVS/Express</b>	8.1	1 OCF only	0	occasional
<b>blockbuster and smtools</b>	2.8.5	N/A	0	unmeasured
<b>OpenDX</b>	4.4.4	N/A	0	3-5 per day
<b>EnSight</b>	10.0.3	7 OCF, 2 SCF	27,120	0-dozens per day
<b>GMT</b>	4.2.0	N/A	0	7-43 per day
<b>gnuplot</b>	4.4.0	N/A	0	100s per day
<b>CUDA</b>	4.0	N/A	0	unmeasured
<b>xmgrace</b>	5.1.20	N/A	0	occasional
<b>IDL</b>	8.2	9+2 OCF , 6+2 SCF	13,575	50%
<b>ImageMagick</b>	6.8.0	N/A	0	unknown
<b>Maya</b>	2014	1 OCF, 1 SCF	1,790	occasional
<b>NCAR/NCL</b>	6.1.0	N/A	0	7 per day
<b>ParaView</b>	4.1.0	N/A	0	unmeasured
<b>POV-Ray</b>	3.6.1	N/A	0	80 per day!
<b>Qt</b>	5.1	1	1,755	unmeasured
<b>RasMol</b>	2.7.5.2	N/A	0	unmeasured
<b>Tecplot 360</b>	2013	7 OCF, 2 SCF	15,840	frequent
<b>VisIt</b>	2.7.2	N/A	0	unmeasured
<b>vmd</b>	1.9.1	N/A	0	12 per day
<b>xv</b>	3.10a	N/A	0	occasional
<b>TOTAL</b>			<b>60,080</b>	



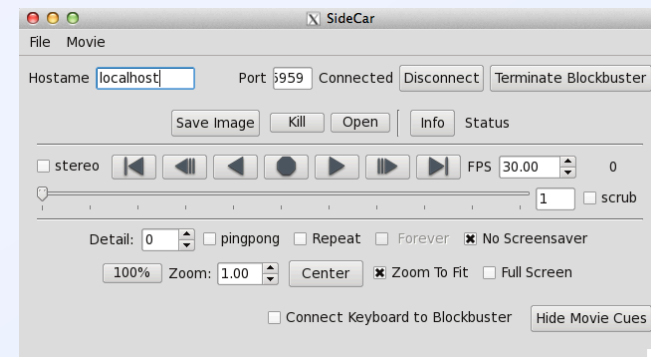
# Playing movies on our Powerwalls: Blockbuster/Sidecar



Presentation “cues” stored in “cue files”



Movie displayed remotely on powerwall



Movie controlled on local console

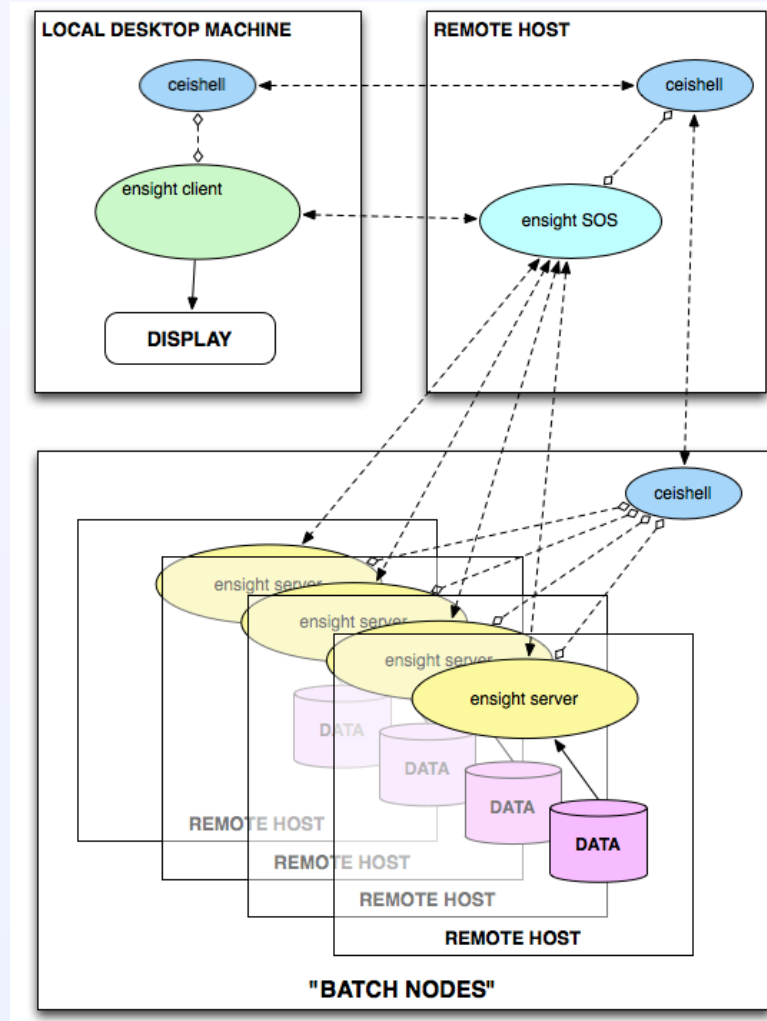
# EnSight in parallel at LLNL

This is way too complicated for us to ask users to manage due to firewalls, portals, and two-factor authentication.

Solutions:

- `ensight_desktop_cz.py`
- `ensight_cluster_rz.py`

“EnSight babysitters”



[https://computing.llnl.gov/vis/screensteps/EnSight\\_parallel\\_computing\\_at\\_LLNL.html](https://computing.llnl.gov/vis/screensteps/EnSight_parallel_computing_at_LLNL.html)



## Data Group Activities

- Released several new versions of VisIt
  - Covered in detail in the “VisIt Update” talk
- Provided data analysis and visualization services to ASC customers
- Movie making
- Hired a new VisIt developer to replace Brad Whitlock

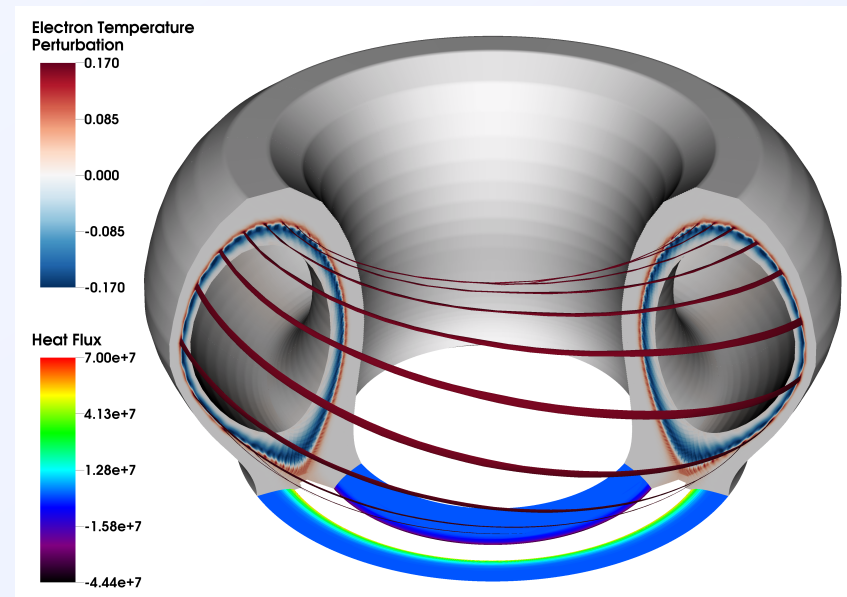


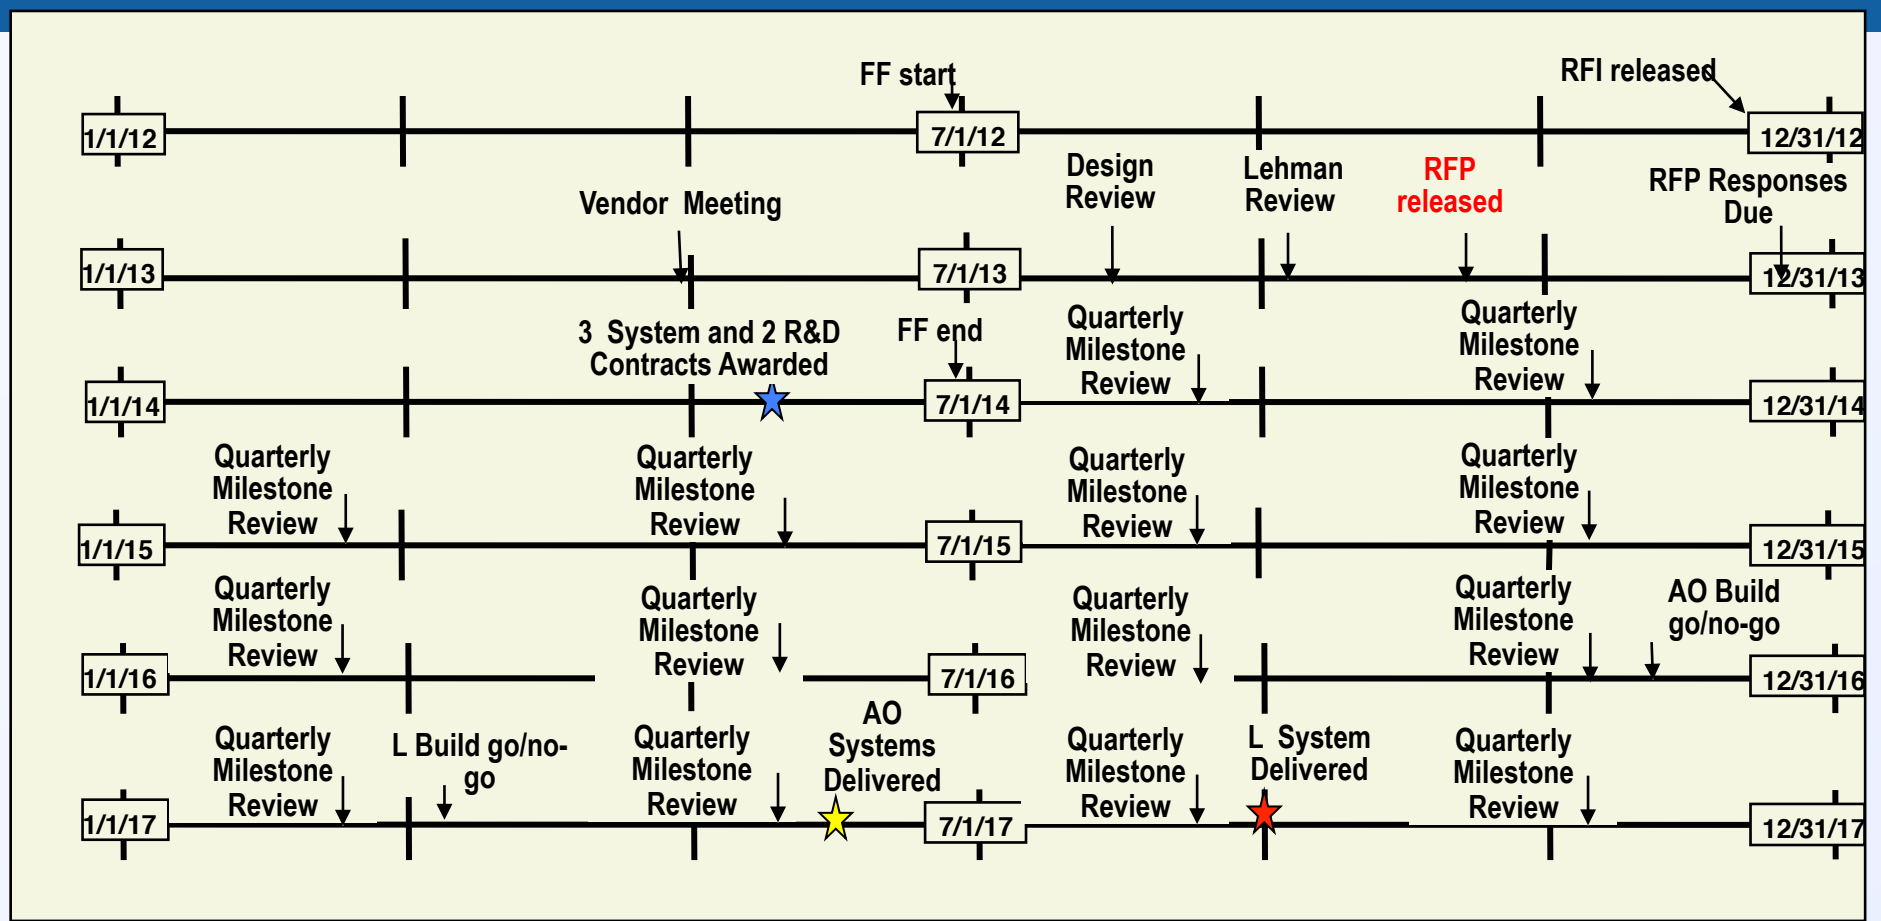
Image from the BOUT++ fusion code

## What is CORAL

- CORAL is a **C**ollaboration of **O**ak **R**idge, **A**rgonne, and **L**awrence Livermore Labs to acquire three systems for delivery in 2017.
- DOE's Office of Science and National Nuclear Security Administration signed an MOU agreeing to collaborate on HPC research and acquisitions
- Collaboration grouping was done based on common acquisition timings. It reduces the number of RFPs vendors have to respond to and number of reviews by Labs and allows pooling of R&D funds
- Los Alamos, Sandia, and Lawrence Berkeley National Labs are collaborating on the first of these joint acquisitions for the Trinity and NERSC-8 systems.



# CORAL Draft Timeline



AO=ANL-ORNL

L=LLNL

★ CORAL contracts awarded

★ AO systems delivered

★ L system delivered



## 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.

