



ALICE-USA T2 ORNL CADES

ALICE-USA
T2 Deployment at ORNL

February 11th, 2015

Pete Eby
ORNL





CADES

Compute & Data Environment for Science

- ORNL T2 site will leverage significant investments at ORNL:
 - 65,000 square feet of datacenter space
 - 40 Megawatts of power, 6K tons of cooling (after latest upgrade)
 - Connectivity to major research networks (ESNet 100G, Internet2 100G, XSEDEnet 10G)
 - Significant expertise across ORNL in delivering similar solutions

CADES Tier 2 Site

- Allows ease of expansion over time
- Provides potential for ALICE workloads to elastically scale to other CADES compute and storage resources



CADES home is shown in Magenta. ALICE system will reside in highlighted area.

CADES Alice Tier 2 Site



T2 Compute and Storage



- CADES will provide
 - High performance compute cores
 - Storage and compute network backplane at 10Gbe
 - High performance block storage
 - High performance WAN connectivity via ESNet
- Delivering to Alice T2
 - VO-Box images will be provisioned on **dedicated** ALICE compute and storage nodes
 - Block storage will be presented to the VO-Box images via SAS (high performance, low overhead protocols)
 - Integration in CADES will allow elastic compute and storage in the future while reducing administrative costs today

•Integration of OLCF resources within CADES provides

exciting future opportunities

T2 Overview

- T2 Compute overview
- T2 Storage overview
- Final config rack layout
- FY 2015 config
 - 56 compute nodes
 - 5 storage nodes
 - 5 - 60 bay storage chassis



FY 2015 T2 Compute Details

- Using the same compute platform we have deployed for many other CADES users providing economies of scale in managing the infrastructure
- Deploying Dell C6220 Quad Node Servers
 - 14 4U C6220s, 56 compute nodes
 - Provides 896 cores for compute (18 kHS06 min*)
 - Dual socket 8 core Intel Xeon E5-2640v2 2.0GHz
 - 64GB total, 4GB per core (8x8GB RDIMM, 1600MT/s)
 - Dual 1TB local drives
 - 10Gbit Ethernet connectivity for compute/data, 1Gbit management network

FY 2015 T2 Storage Details

- Deploying Sanmina Storage JBODS
 - 5 - 60 Bay storage JBODS
 - 4 Terabyte Enterprise SAS 7.2K drives
 - Provides 1.2 Petabytes raw, 1PB usable with 10+2 RAID6
 - Exploring 3 – 84 bay, 6TB, 1.5PB raw solution.
 - 5 - C6220 Storage nodes (same config as compute)
 - SAS connectivity – 1x SAS IO module with 4x x4 (24 Gbps) ports
 - 5 Storage server nodes total – each direct connected to the storage
 - 10 GByte/sec I/O throughput (limited by dual 10Gbe/server)
 - 555 KBytes/sec per HS06 (10 Gbyte/sec/19KHS06)
 - Drive zoning supported

OpenStack Tenancy Overview

The screenshot displays the OpenStack Dashboard interface. At the top, the browser address bar shows the URL `or-c13.ornl.gov/dashboard/project/`. The dashboard header includes the 'RED HAT OPENSTACK DASHBOARD' logo and navigation tabs for 'Project', 'Admin', and 'Settings'. The main navigation bar lists 'Compute', 'Network', and 'Orchestration', with 'Overview' selected. Below this, a sub-navigation bar includes 'Overview', 'Instances', 'Volumes', 'Images', and 'Access & Security'. The 'Overview' section is titled 'Limit Summary' and features seven circular gauges representing resource usage:

- Instances:** Used 2 of 10
- VCPUs:** Used 2 of 20
- RAM:** Used 4.0GB of 50.0GB
- Floating IPs:** Used 1 of 50
- Security Groups:** Used 2 of 10
- Volumes:** Used 2 of 10
- Volume Storage:** Used 40.0GB of 1000.0GB

Below the gauges is the 'Usage Summary' section, which includes a prompt to 'Select a period of time to query its usage:'. The form shows 'From: 2015-02-01' and 'To: 2015-02-11', with a 'Submit' button. A note states 'The date should be in YYYY-mm-dd format.' The usage summary text reads: 'Active Instances: 2 Active RAM: 4GB This Period's VCPU-Hours: 18.06 This Period's GB-Hours: 361.27'.

OpenStack Tenancy Instances

Instances - OpenStack Dashboard - Google Chrome

Instances - OpenStack | or-c13.ornl.gov/dashboard/project/instances/

RED HAT OPENSTACK DASHBOARD Project Admin Settings Current Project ALICE Red Hat Access Help admin

Compute Network Orchestration

Overview Instances Volumes Images Access & Security

Instances

Instances

Filter Filter [+ Launch Instance](#) [Soft Reboot Instances](#) [Terminate Instances](#)

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Uptime	Actions
<input type="checkbox"/>	alice-worker	-	CADES-public 128.219.150.157 CADES-external 172.23.64.38	m1.small 2GB RAM 1 VCPU 20.0GB Disk	or-c13	Active	nova	None	Running	38 minutes	Create Snapshot More
<input type="checkbox"/>	alice-scheduler	-	CADES-public 128.219.150.155 CADES-external 172.23.64.37	m1.small 2GB RAM 1 VCPU 20.0GB Disk	or-c13	Active	nova	None	Running	11 hours, 4 minutes	Create Snapshot More

Displaying 2 items

T2 Other Details

- DOE Review Completed in December
- T2 Funds Released 30 Dec
- Vobox instance and OpenStack VMs accessible now
- T2 Project SLA and WLCG MOU submitted and under approval process.
- CADES member Pete Eby will continue to support the T2 site. Michael Galloway has been added for additional design, implementation and support.

Current Summary

- ORNL VO Box (metal) accessible via 10G outside ORNL border firewall
- ALICE Tenancy created in CADES OpenStack
 - Hosting of management and workflow Vms
 - `alice-lbni.ornl.gov` – scheduler
 - `alice-worker.ornl.gov`
 - Accessible via ssh
- 56 compute cluster nodes ordered
- Storage JBODs final builds being designed
 - Will order before Feb. T1/T2 conference

Current Summary

- Deployment of storage infrastructure will be prioritized
- Storage and compute expected to be racked by mid-April to May 2015
- Rack space & power have been allocated