

CADES at ORNL



Michael Galloway & Pete Eby & CADES Team March 2017

CADES Short Overview

Compute And Data Environment for Science

CADES Brief History

What are we talking about when we talk about CADES?

Means different things to different {people, organizations, labs}

CADES organizationally was ITSD now is NCCS

Who is a CADES 'customer': who's getting science done?

The Early Years (2013–2015)

Project Initiation

Initial Investments

Initial Deployments

Chaos

Middle Period (2015 - 2016)

Initial Deployments

Focus on Operational Readiness

Initial Offerings to “Customers”

The Harsh Hand of Reality sets in

Current State

Moving Towards Production

Increasing Collaborative/Integrative Workflows

More Nonconforming Platforms

Move from ORNL/ITSD to NCCS

Team Focused on CADES

CADES Today

Network

Cloud


HPC

Storage

Platforms

ALICE











CADES Network

Layered/Tiered Topology

Firewalled PZ's Inside the Border (RAN/OR/ETC)

Border Connections - ALICE, LHCOne (in progress), KBASE

Science DMZ/Data Fabrics - Performance tuned DTN pools, Globus

CADES Storage

Lustre Open and Mod - 2PB Open, 0.5PB Moderate 56Gb/s

NFS Open and Mod - 700TB Shared Open/Moderate 10Gb/s

Object Store - 2PB Shared Open/Moderate - S3/NFS

Providing Unified Fabrics

Evaluating additional solutions and protocols

CADES Cloud

Private Cloud running OpenStack Mitaka (Red Hat OSP9)

Host Aggregates Provide Mixture of Node Types - General/Highmem/etc

Data Fabric Connectivity (NFS/Luster/Globus/etc.)

“RUC” Allocations and “Birthright”

Self Provisioning - User have Horizon and API access, ORNL images provide standard environments, Ansible roles, etc.

AWS Available to Projects - Billed to project accounts

CADES HPC

7,500 Core Open

4,500 Cores Moderate

20 GPU (K40x80) Nodes Open

10 GPU (K80x20) Nodes Moderate

Migration of Some OLC Components

CADES Platforms

CRAY GX - Analytics Engine (Spark/Hadoop/Graph)

CRAY XK7 - Two Rack 'Baby Titan'

SGI UV300 - Large Shared Memory Systems (3TB/24TB)

NVIDIA DGX-1 - Deep Learning Platform (1 Open/2 Moderate)

CADES Alice/USA

ALICE Tier-2 Service

1,500 Core HPC

1.4PB Data Store

Connection to Network Edge 10Gb/s uplinks

LHCONE 'Real Soon Now'

CADES Customers

Heterogeneous Mix of ORNL R&D Divisions

Getting to Scale in small bites

- ALICE/USA
- ACME
- Nuc Phys/GIST
- SNS/BSD/CNMS/EDDE/ACME/ARM/OLCF

CADES What We Do

High Level Engineering/Consultative Service

Application Integration

Workflow/Dataflow Assistance

CADES Future

Where is this CADES thing heading

Gluing together heterogeneous hardware platforms

Development of innovative workflows for diverse environment

Integration of analytics for these complex workflows

Increased collaborations with internal and external entities

Data life cycle management

CADES Team

Ajun Shankar - Director

Brian Zachary - Deputy Director/Team Lead

Susan Hicks, Ryan Prout -- Networking


Chris Layton, Pete Eby -- Cloud

Nathan Grodowitz, Michael Galloway -- SHPC/Storage

Steve Moulton -- Monitoring/Performance Reporting

Rich Mohr, Hong Liu, Kwai Wong, Jason Charcalla -- NICS Subcontractors

CADES Team



Questions?

We're happy to answer anything we can.