SAND2019-11239C

Uniform Monitoring of Heterogeneous Architectures

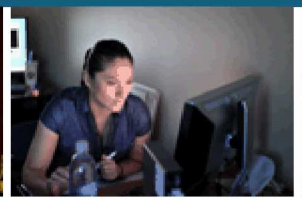

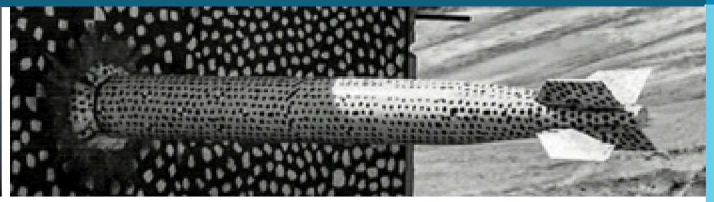

Connor Brown, Ben Schwaller, Nathan Gauntt, Benjamin Allan, Kevin Davis

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Outline

- Introduction
- Design
- Implementation
- Future Work
- Questions

Introduction

- HPC centers becoming increasingly heterogeneous with the need for monitoring ever present
- Heterogeneous Advanced Architecture Platforms (HAAPs) team administers next-generation testbeds
  - Unique mix of many standalone clusters residing on multiple networks
- Needed an image management solution that allowed for rapid deployment of our system configurations

Design

Summary of our clusters

- Eleven individual testbed clusters of a variety of emerging technologies
  - Each platform cluster with its own admin and login node
  - Diskless boot
- Full monitoring stack on five clusters

System Architecture	*Network Fabric*	*Node Count*
Cray XC-40 with Intel Haswell x86_64	Aries	26
Cray XC-40 with Intel Knight's Landing x86_64	Aries	22
Intel Haswell x86_64	Infiniband	4
Intel Skylake x86_64	OmniPath	30
AMD x86_64	Infiniband	4
IBM Power9	Infiniband	10

LDMS

- Monitoring with Lightweight Distributed Metric Service (LDMS)
  - Portable across multiple architectures
  - Wide range of metric collection options suitable for many networks and architectures
- 1Hz sample rate collecting 105GB per day across all of our testbeds
- Monitoring box with local RAID stores for all LDMS data

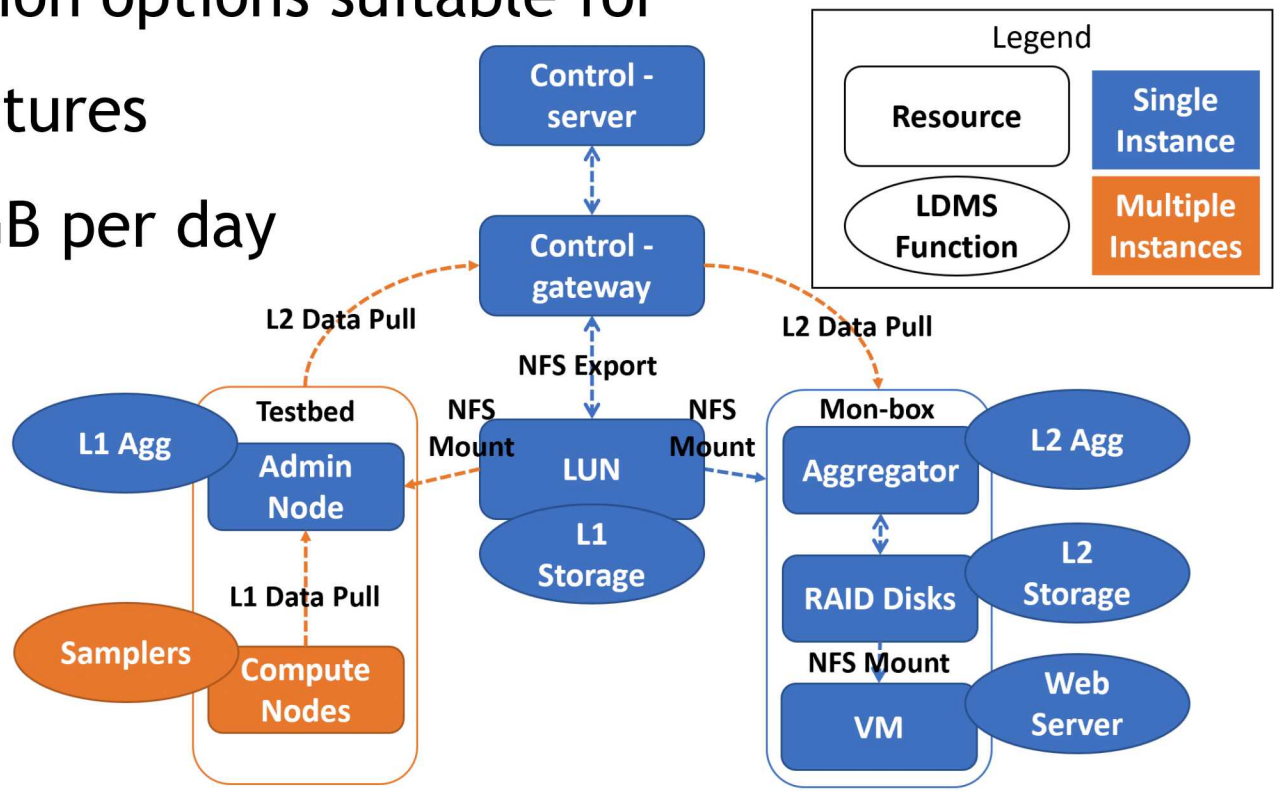

Golden Image

- Common location for maintaining custom environment and package repositories
- Rapid deployment of new testbed systems with minimal configuration
- Out of band configuration testing with real hardware
- Relatively simple imaging process without many additional software packages needed

Implementation

Networking

- Custom VLAN routing allowing network isolation
- Scalable ruleset for many cluster monitoring
- Simple configuration allowing for easy integration of new testbeds

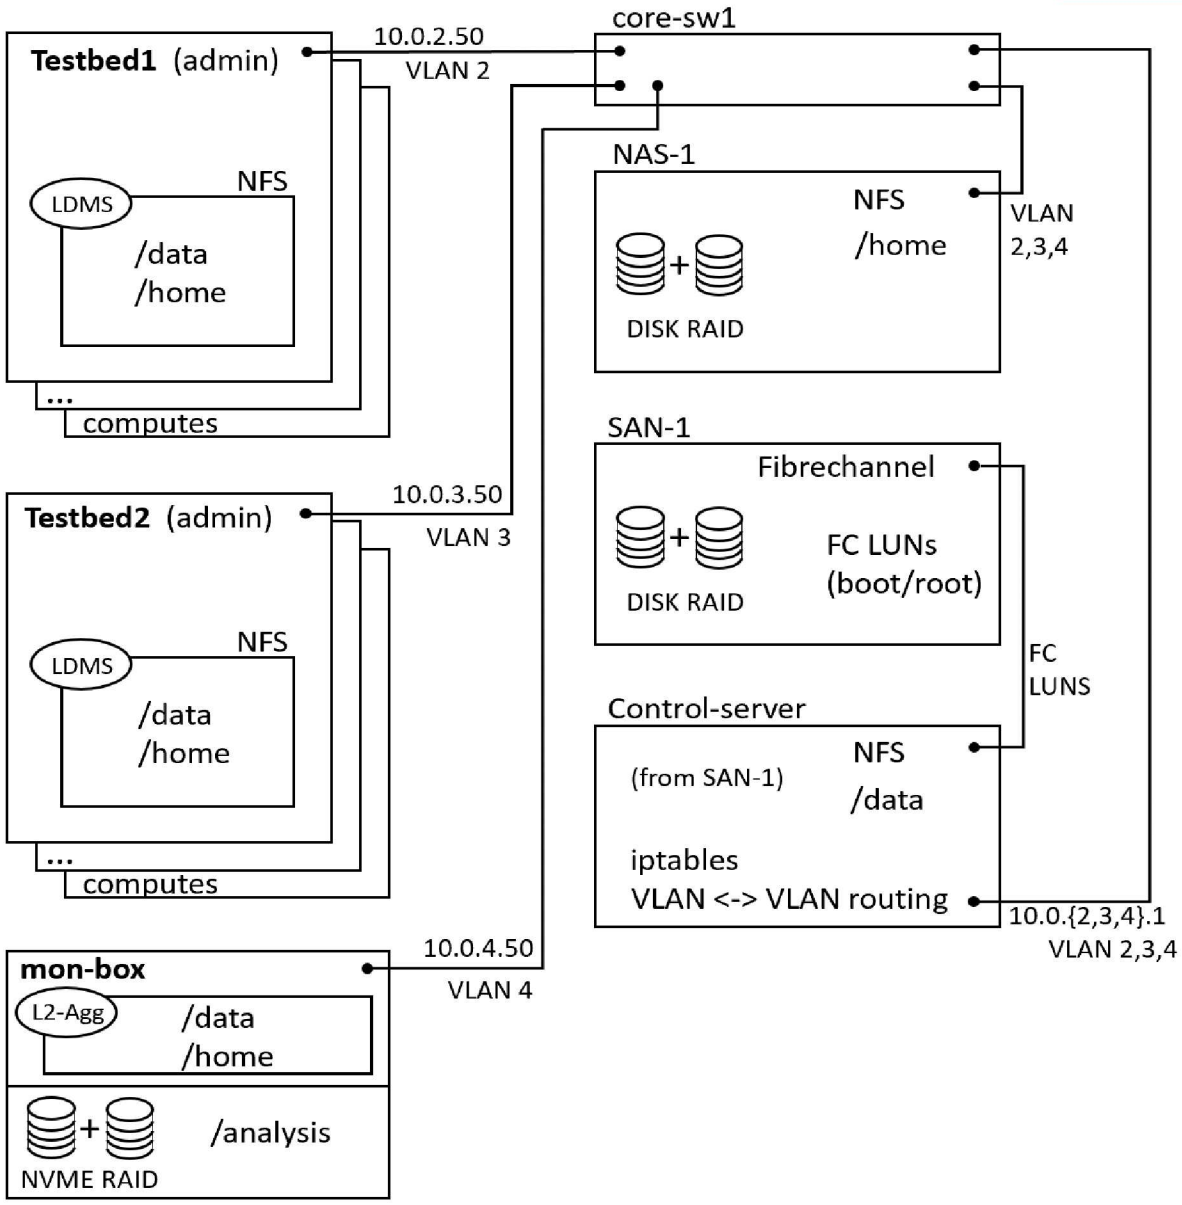

Webserver

- Django framework that uses Scalable Object Store (SOS) queries
- Provides performance and resource utilization about applications at a per job level
- Located within a VM on our monitoring box

From: 06/26/2019

To:

Show 25 entries

Search:

Job ID	App ID	Node ID	Runtime (s)	Back Pressure	Mem Score	Anomalies	PAPI Perf	App Perf
42093	miniAMR	nid000[52-55]	439	0.0	2	None	Back	1.45
42092	miniGhost	nid000[21,29-31]	1043	49.07	2	Cache	Back	-1.93
42091	miniMD	nid000[57-60]	617	5.24	3	Cache	Back	No data
42090	kripke	nid000[21,29-31]	66	0.0	1	None	Back	No data
42089	CoMD	nid000[52-55]	742	91.68	1	Cache	Back	1.52
42088	miniAMR	nid000[21,29-31]	447	0.0	2	Cache	Back	1.45
42087	miniGhost	nid000[57-60]	1043	73.88	2	Cache	Back	-0.27
42086	miniMD	nid000[21,29-31]	619	13.33	3	Cache	Back	No data

Future Work

Future Work

Baler/Nagios

- Better interpretation of syslog data
- High level view into system status

Automation Scripting

- Additional scripting to create custom cluster configuration
- Decreased system standup time
- Package installation

Future Work

Webserver/Data Analytics

- Openly accessible VM to allow users to access the webserver
- Add cluster-wide metrics
- Add the data from the remainder of our clusters to the webserver
- Grafana testing and integration to allow independent user analysis
- Additional figures of merit that the user and administrator community will find helpful

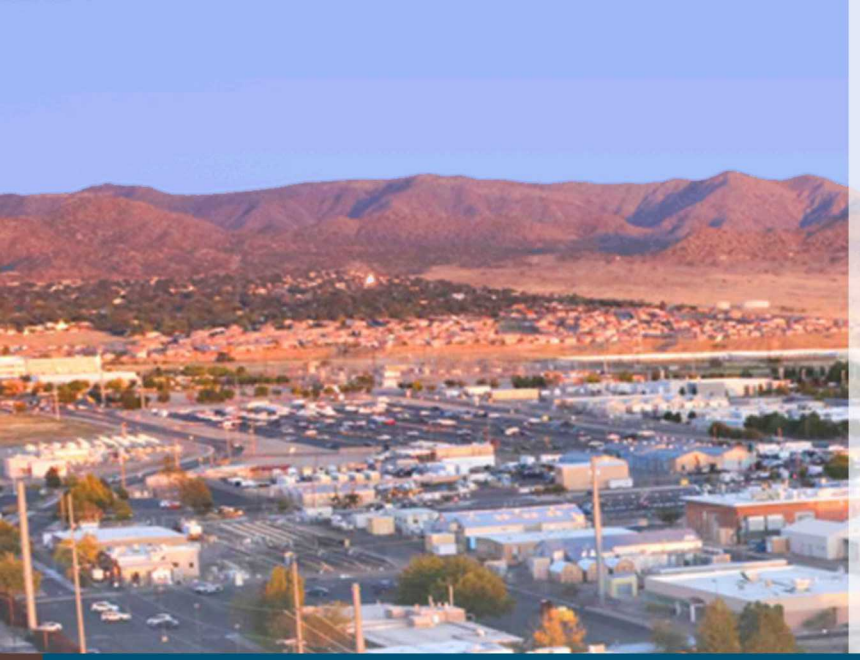

Questions?