Scalable Data Warehouse for the Operations Monitoring and Notification Infrastructure at NERSC

Meriam Gay Bautista, Thomas Davis, Elizabeth Bautista, Melissa Abdelbaky, Cary Whitney,
Operations and Technology Group, National Energy Research Scientific Computing Center
Lawrence Berkeley National Laboratory, Berkeley, California, USA

Abstract

The Operations Monitoring and Notification Infrastructure (OMNI) is a data collection system and warehouse that collects operational heterogeneous data at the National Energy Research Scientific Computer Center (NERSC) at Lawrence Berkeley National laboratory to provide computational resources to science user at high availability of the high-performance computing (HPC). OMNI assists NERSC in monitoring the health of most of the areas in the facilities that operate on a 24/7 basis. The data provides the team with a holistic view of the HPC data center that includes the building management system data, sensors data, and computer Syslog data. The data is used to plan, procure, build and remodel the next-generation systems.

With the delivery of the new HPC systems, and because of the scale of the new machine, the data rate available is expected to be 100 to 1000x faster. Thus, it is anticipated that the exascale size of data is to be sent to OMNI. To support the ability to collect more data, the team developed and instrumented a scalable and integrated network data collect automation strategy to scale OMNI's growth.

Using OMNI, the operational team are able to use real-time data to keep the HPC system highly available. The data has been used to lower costs, save hardware, assist with business decision and influence collaborations. OMNI system collected and stored years of data that can be used as training datasets, eventually enabling machine learning and automated optimization

Timeline

Problem statement:

- Back in 2012, the main issue the team encountered is that the monitoring and performance packages used would not scale to the desired size needed for the HPC System.

Solution :

- To build an an "instrument everything" monitoring infrastructure, that is reliable, scalable, maintainable and with high availability.

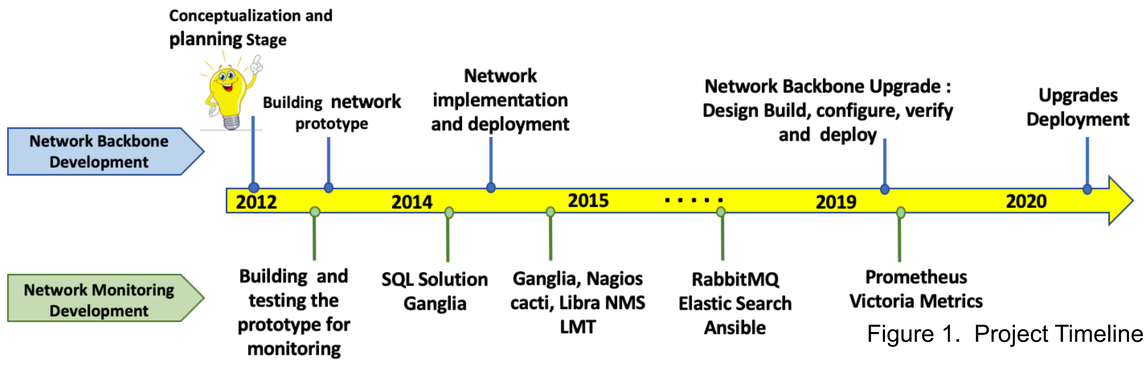

Figure 1. Project Timeline

Scalable and Automated Data Warehouse for OMNI

- The team adopted an "instrument everything" approach for collecting data from heterogenous sources and exposed it via the data warehouse. The key feature of OMNI is its ability to ingest and present data on a near real-time basis. The new backbone network adopted a scalable open-networking solution and integrated automation strategy to accelerate the workflow.

Network Infrastructure:

- 2-Spine - 4-Leaf topology
- Spine switch is implemented using 32-Port 40-GbE fiber SDN switch
- Leaf switch is implemented using 54T with 48-10GbE ports and 6-40GbE ports.

Key Technology feature of a Scalable network backbone:

- Centralized Configuration Management plane
- Ultra Power-over-Ethernet (PoE) for instrumentations
- Scalable Spine-Leaf network architecture
- Use of BGP-EVPN-VXLAN for the network configuration

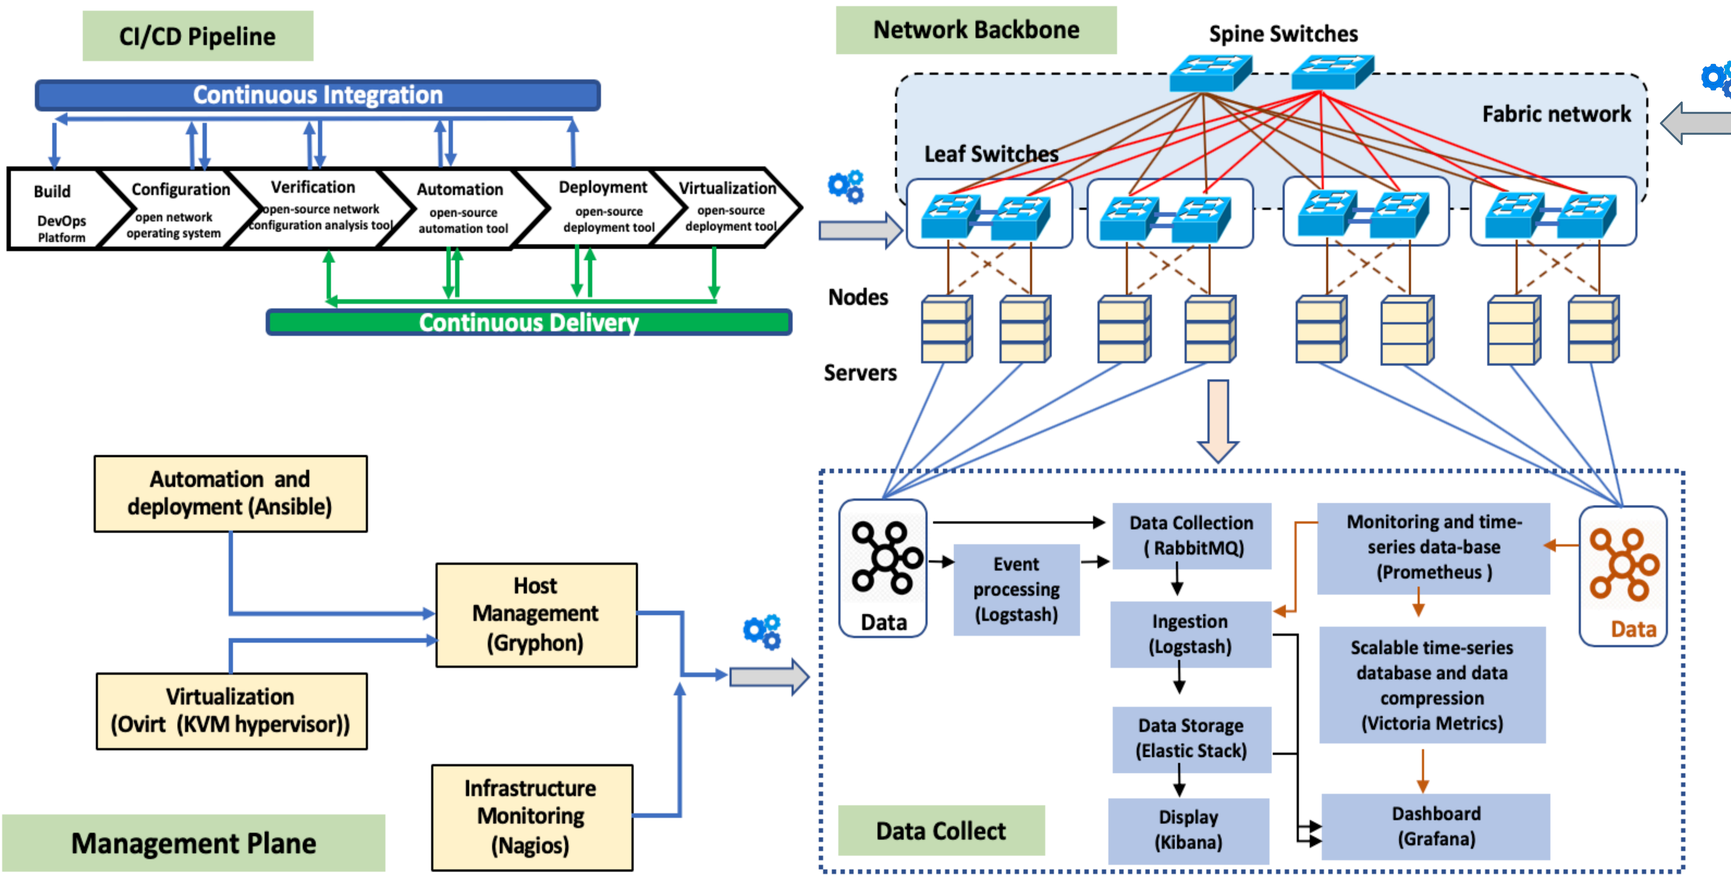

Figure 2: OMNI Network backbone and data collect infrastructure

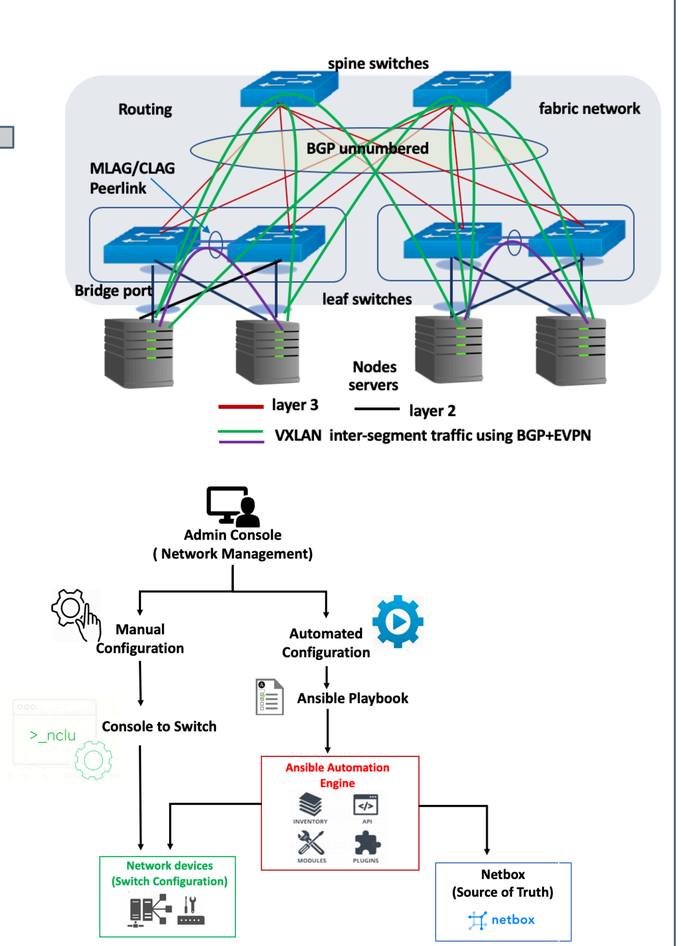

Figure 3: Automation Workflow

Current Stats:

- Elastic Search: 550 Billion records (125 TB data)
- Victoria Metrics: 4.6 Trillion datapoints

Data Monitoring and Visualization

Temperature Monitoring

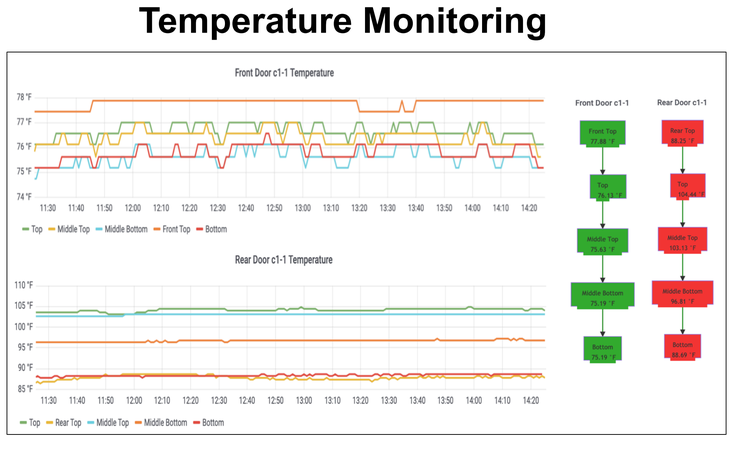

Figure 4. Rack temperature Dashboard.

- The temperature data is collected using T-String temperature sensors mounted at the front and rear cabinet doors. This dashboard shows the time-series average temperature of the sensor nodes in each rocks of NERSC supercomputers.
- The temperature monitoring is critical to NERSC building management system. It is important that the temperature are in nominal values and within threshold level.

Air Handling Unit Monitoring

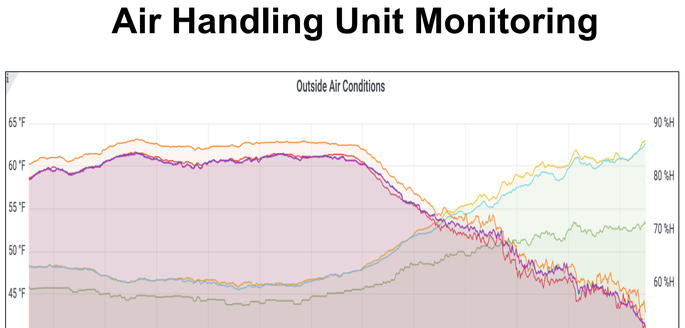

Figure 5. Outside Air Conditions Dashboard

- This dashboard shows the time-series monitoring data of the Air handler Unit (AHU) and the wet and dry bulb. The AHU controls the humidity in the facilities and the wet and dry bulb indicate how much moisture is in the air.
- The moisture in the air reduces the air temperature and lowers the temperature of the building, thus less energy is used by the AHU.
- They are inversely related, higher relative humidity will decrease the temperature of the AHU.

This work used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.