INL/EXP-20-59238-Revision-0

High Performance Computing Systems Tools, Visualization, and Management

Paul W Spencer, Matthew R Sgambati

August 2020

The INL is a U.S. Department of Energy National Laboratory
operated by Battelle Energy Alliance

INL/EXP-20-59238-Revision-0

High Performance Computing Systems Tools, Visualization, and Management

Paul W Spencer, Matthew R Sgambati

August 2020

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy

Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517, DE-AC07-05ID14517

Final Edge of Poster

Paul Spencer - Brigham Young University
C520 - HPC Advanced Scientific Computing
Mentor - Matthew Sgambati

Idaho National Laboratory

High Performance Computing Systems Tools, Visualization, and Management

Software Module Usage

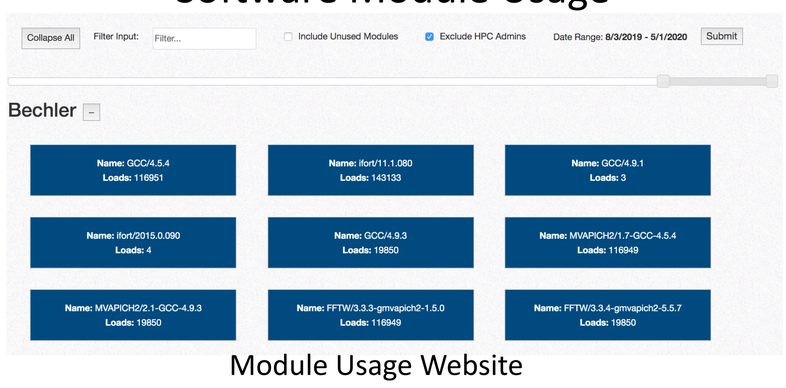

Module Usage Website

HPC software administrators manage hundreds of software packages, or modules. Some of these modules see more usage than others. Administrators need to know which ones can be removed and which ones need to be upgraded.

The Software Module Usage project is just that: a system to keep track of which modules are used and how frequently. It includes a website to visualize the information the system tracks.

Completed Summer 2019 with Nathan Johnson.

User Monitoring and History

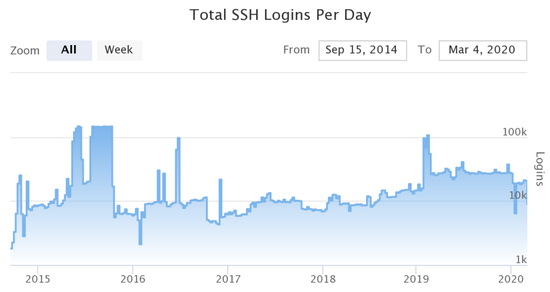

User Login Metrics

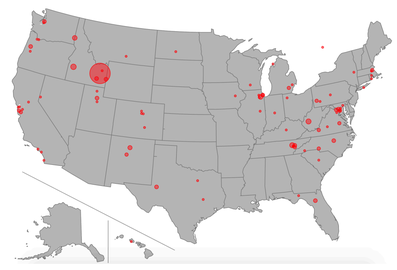

User Login Locations

HPC system administrators must be aware of each user using HPC resources as well as what the users are doing with the supercomputers. This is a difficult task without tools designed to track users.

The User Monitoring and History (UMH) project is a system to track user sessions in the HPC enclave. With it, system administrators can easily get a good idea of “who is doing what where”. Currently, it consists of a website to show current login data (including geographic user locations) as well as login data over time.

Still in development with Dylan Gardner.

Netmap

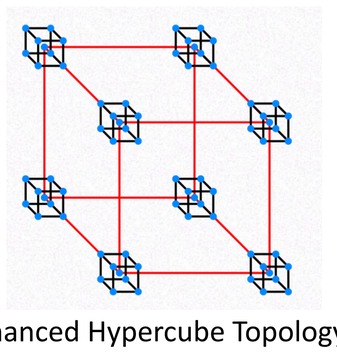

Enhanced Hypercube Topology

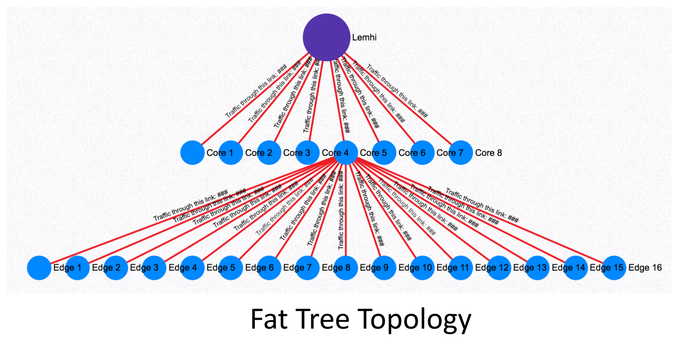

Fat Tree Topology

HPC centers require very complex network systems. Each supercomputer is composed of hundreds of servers that must be able to communicate with each other efficiently. This is cause for a complicated network topology, which can be difficult to manage.

The Netmap project provides a model for each supercomputer’s topology including metrics through each connection, as well as a view of the physical layout of the racks of servers and their network connections.

Still in development with Bradlee Rothwell.

IPAM (IP Address Management)

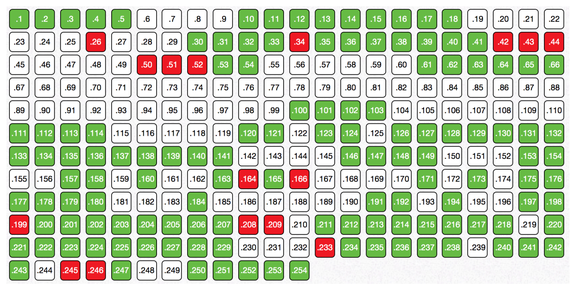

IP Address Status View

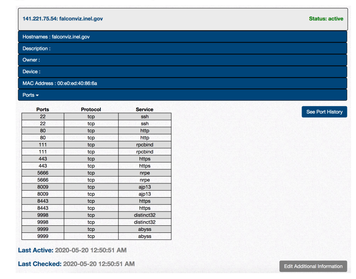

Machine Information View

The extensive network systems in INL’s HPC center involve hundreds of IP addresses for various servers. System administrators are tasked with keeping track of which ones are in use, which ones are free, and managing every machine on the network.

The IPAM project provides an easy way for system administrators to view and manage any IP address inside HPC’s network. They can easily associate IP addresses with hostnames, check the history of that IP address, view the open ports on that server, and see other useful information.

Completed Fall 2019 through Winter 2020 with Bradlee Rothwell.

GPU Metrics

GPU Utilization over Time

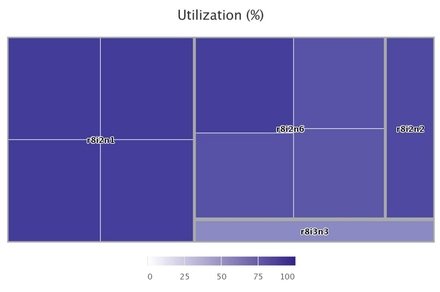

Current GPU Utilization

Because of the power of AI and Machine Learning in nuclear research as well as the efficiency of GPUs in these areas, INL HPC has purchased systems with large numbers of GPUs.

The GPU Metrics project provides a way to monitor important metrics associated with HPC’s GPU systems. Utilization, power consumption, temperature, and memory usage are examples of metrics that are collected by the GPU monitoring system. The project also includes a website that can be used to visualize this data and a history of the metrics.

Completed Winter 2020 with Bradlee Rothwell.

Sensitive Text Obfuscator

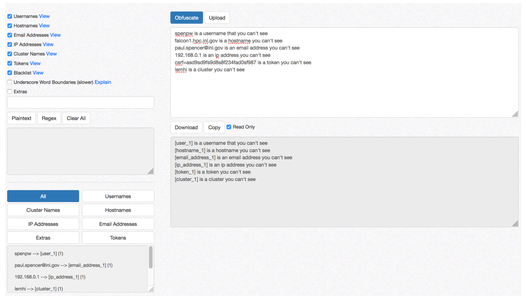

Web Obfuscation

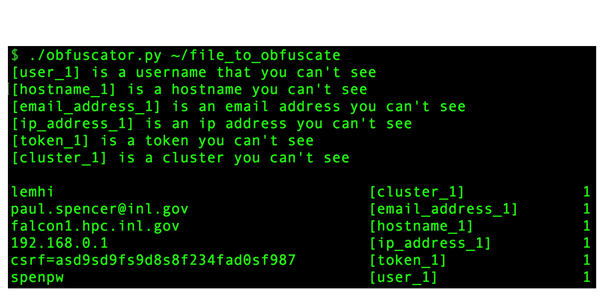

Command Line Obfuscation

HPC centers have high quantities of sensitive data about users and datacenter specifications. When sharing logs or other information with vendors or outside entities, it can sometimes be difficult for HPC administrators to properly filter out the sensitive information and avoid security risks.

The Sensitive Text Obfuscator is a tool that obfuscates information such as usernames, email addresses, IP addresses, hostnames, security tokens, etc. It is available through a website or a command line utility.

Completed Spring/Summer 2020 with Dylan Gardner.

Tools used:

Final Edge of Poster