High-Performance Computing Security at The University of Memphis
Implementation of NIST SP 800-223

This document is written in response to the NIST Special Publication, NIST SP 800-223, which aims to standardize and facilitate the high-performance computing (HPC) security, architecture, threat analysis, and security posture. Included are recommendations for the HPC environment at the University of Memphis (UofM).

The document references four security zones which are: Compute, Data Storage, Access, and Management.

Compute

NIST SP 800-223 recommends, where possible, "Compute Node Sanitation" – which is rebuilding the nodes after jobs complete. In general, the UofM's HPC is utilized in a way that nodes cannot be easily sanitized after each job. The configuration allows multiple jobs to share resources across the cluster as available instead of a "one node, one job" policy. That is, multiple jobs from different users will be running on the same node at the same time and finishing at different times. Scheduled jobs use various amounts of CPU and RAM. Sanitizing the node after a job completes would remove resources in use by other jobs running on the node.

There are two solutions available to meet this requirement.

The first solution would be to set aside, on an as-needed basis, nodes that can be used for single jobs. The nodes will be rebuilt upon job completion. Users working on regulated research would request this configuration. Using existing HPC resources, a job scheduler will be configured to separate out nodes into their own job queue.

A second solution would be to use the few whole nodes already set aside for jobs that require an entire node by themselves (or multiple nodes using MPI). Users working on regulated research would request the job scheduler for designated nodes would be set to re-image after use.

Costs to meet the Compute security requirements for NIST SP 800-223 include system administrator time and wait-time for nodes to finish currently running jobs so the nodes can be configured to securely support regulated research. When the need has been met, the nodes can be returned to the original configuration.

Data Storage

Data Protection is considered in two ways: how data is physically protected and considerations for encryption by the end-users.

Storage attached to the HPC has several layers of physical protections to guard against loss. Controls include limiting access to the data center to authorized personnel that are required to be logged in and escorted or use logged key-fob access. Data center access logs are reviewed monthly. The data center is alarmed and is equipped with cameras at entrances and on each aisle. In addition to the data center's physical controls, HPC data is stored across hard drives in a regular RAID fashion. This renders individual hard drives' data unusable in the event of individual drive loss.

Software is available to meet the data encryption requirements in NIST SP 800-223. Researchers can select encryption software that will best meet their needs and budget. Research Computing can assist researchers in identifying and installing encryption software and acquiring keys/tokens needed to encrypt and decrypt data. The researcher will be responsible for deciding which encryption software will be used, procuring the software, and for the management of keys/tokens.

Access

NIST SP 800-223 recommends using network segmentation to limit access to HPC nodes by IP address or IP range. Access control for the HPC clusters nodes is implemented as follows.

Outward-facing nodes for the HPC cluster are the single head node and two login nodes which are round-robined in DNS. Each node is assigned an IP address.

All campus IP addresses (ranges include 141.225.0.0/16 and 10.0.0.0/8) can access the HPC login nodes using the SSH protocol. To increase access control, plans include constraining access to HPC login nodes to the campus VPN for members of the hpc-users LDAP group.

Access to the head node is limited to the members of the Enterprise Infrastructure team and the hpc-admins LDAP group by authentication to the campus VPN. Head node access to the two groups is limited to the SSH protocol. Authentication is accomplished by the sssd service and access is controlled by group memberships. To increase access control, plans include implementing local firewall rules to block “east-west” traffic from the network segment on which the head node resides and to limit access to members of hpc-sysadmins group.

The compute nodes are not routable from outside the HPC cluster making compute nodes only available from within the cluster itself. Access to the compute nodes is further restricted to only those users who have an active job running.

Management

Management of the HPC is secured by limiting access to the head node by IP addresses and limiting login access and configuration privileges on the head node to select accounts. Limitations for access are described above. There are two HPC administrators, with root-level privilege, and a Graduate Assistant account, with some sudo privileges, that are authorized to make configuration changes to the HPC cluster.

Networking to the remote access control of the nodes (iDRAC, iLO, IPMI) is not routable and only available from the head node.