TECHBYTE
Insights into Research Computing

12 March 2024

SECURING HIGH PERF0RMANCE COMPUTING INFRASTRUCTURE

Author:

Rakesh Sabharwal

Founder **ON DEMAND SYSTEMS PTE LTD**

ON DEMAND SYSTEMS PTE LTD

ON DEMAND systems

TECHBYTE
Insights into Research Computing
12 March 2024

SECURING HIGH PERFORMANCE COMPUTING (HPC) INFRASTRUCTURE

Historically, HPC and Big Data platforms (Hadoop) have been managed and operated in isolation.The researchers have to access the environment using a separate set of credentials.With HPC going mainstream and integrating with the enterprise environment, there is a need to delve into securing the HPC infrastructure to comply with audit and compliance requirements.

Almost all HPC/Big Data setups would have Linux as an operating environment. These environments have some inherent security issues.

Identity Silos = Increased Risk

Organizations typically have multiple HPC/Big Data Clusters for different applications or environments such Dev, Test, and Production as well as multiple Production environments for various lines of business.

High-Value Data = High-Risk Target

HPC/Big Data is typically used to analyze larger volumes of data that come from several different sources which often contain sensitive PII or PCI or legal or financial or Patient Data.

Risk of Failed Audits

IT staff require access and privileges to manage the cluster, and end users need to be able to submit jobs across the clusters.

ON DEMAND SYSTEMS PTE LTD

Author:
Rakesh Sabharwal
Founder ON DEMAND SYSTEMS PTE LTD

TECH BYTE
Insights into Research Computing
12 March 2024

SECURING HIGH PERFORMANCE COMPUTING (HPC) INFRASTRUCTURE

Enterprises and research organisations shall explore the following best practices to secure the HPC/Big Data Infrastructure.

Identity Consolidation

This should be seen as the first step, wherein all identities shall be consolidated into the corporate directory e.g. Active Directory.
"No Local User Accounts" shall be the established target.

Enforce Least Privilege Access

Implement Least Privileged Access to sensitive HPC infrastructure and data through access with just-in-time privilege for the just-enough privilege.

Secure Administrative Access

Implement privileged access management for administrative accounts. Securely store, rotate and strictly control access to secrets to limit exposure of credentials and reduce the attack surface. Leverage the existing PAM infrastructure to your HPC/Big Data infrastructure.

Secure Remote Access

Have your remote employees, outsourced IT and third parties login directly to HPC Infrastructure without a VPN, reducing the risk associated with access to the entire network.

ON DEMAND SYSTEMS PTE LTD

Author:
Rakesh Sabharwal
Founder ON DEMAND SYSTEMS PTE LTD

SECURING HIGH PERFORMANCE COMPUTING (HPC) INFRASTRUCTURE

Multi-Factor Authentication (MFA)

Multi-Factor Authentication for administrative access provides an extra layer of security that stops in-progress attacks on critical resources. Specific to HPC/Big Data infrastructure, MFA at system login is essential to ensure that only authorised users can submit jobs and access critical data. Once on the server, the user should be prompted for a second factor authentication when elevating privilege (Sudo) to run a highly privileged command. MFA can also be used when checking out a vaulted password; e.g., during a “break- glass” emergency where the root account password is required for console login.

Audit

Fully audit the user access, all privileged activities should be monitored and recorded so that any breach investigation or incident response has a visual forensic-level records to assist.

In conclusion, HPC/Big Data security needs to be reconsidered, which means consolidating identities, centralizing identity and privilege management, controlling access, and auditing all privileged activities as you do today for any other enterprise systems.

Author:
Rakesh Sabharwal
Founder ON DEMAND SYSTEMS PTE LTD