Secure4 Research Computing Environment Security Plan

Revision 2018-02-14

Audience

With the increasing applicability bioinformatics and genomics to clinical medicine, there is an increasing need for specialized and secure compute. Computational analysis, even on a small scale, often requires significant storage, compute capacity, and memory and even minimal configurations are well beyond the capability of the most advanced workstations. Information security requirements make the acquisition, configuration and operation of this equipment even more complicated. Moving beyond "small scale" requires large investments in infrastructure. By developing and sharing centralized capacity, researchers can use equipment on an occasional basis without investing in dedicated infrastructure. Centralized resources can also be sanctioned as meeting information security requirements of the various funding agencies and will facilitate IRB and other approvals.

This security plan is a resource for Principal Investigators (PIs) when developing and submitting grant proposals. This plan outlines how the Secure4 research computing environment and the data stored within will be secured and also explains important processes such as how access is granted to research teams and how research data is segregated from other research projects co-located within the same environment.

Environment

The Secure4 environment is a high performance research (HPC) cluster suited for researchers who need to perform computations with Data Classification Level 4 (DCL4) data and who also need sizeable disk space, considerable memory, large compute capabilities or access to the more than 200 applications available in the environment. Please contact rcss-support@missouri.edu to schedule a consultation with Research Computing Support Services.

Project Data

Access to the Secure4 environment is restricted to eligible projects that contain data that is classified as DCL4. Project data must fall into one or more of the following data source categories:

1. Data that is a part of a project approved by the University of Missouri Institutional Review Board (IRB).
2. Data that is a part of sponsored research project that is managed by the Office of Research.
3. Data that is solely owned by the University of Missouri.

The Secure4 environment explicitly does not allow project data that is classified as Controlled Unclassified Information (CUI) or is controlled by International Traffic in Arms Regulations (ITAR).

PI's are responsible for ensuring that this security plan meets the requirements for all data use agreements and data security plans for all project data. PI's are also responsible for compliance with any University, College/Unit, and data owner policies, data management plan, data use agreement, and all other agreements and policies associated with the use of the data. If researchers have questions they should contact Research Computing Support Services.

Researchers must indicate the source of the data. Changes (addition, removal, and modification) to controlled data sources (DCL3 and DCL4) must be reported. The changes in DCL1 (Public) and DLC2 (Sensitive) data sources does not constitute a change in the project.

Security Plan

Description of the Environment

The Secure4 environment is modeled after a traditional HPC computing environment (cluster) with security enhancements and is comprised of four components, the "login" node for user interactions, the "head node" to control the cluster, a storage node for network storage, and "compute nodes" to perform calculations.

Users login to a "login" node to spawn computational tasks on "compute nodes". The compute nodes are schedule by the "head" node that also provides networking, file, and configuration services internal to the cluster. Users access a dedicated account for each user and project (user-project) via a secure shell connection that provides remote console, display, and file access to the cluster. The security measures put into place include a remote audit log for all security events, a highly restrictive file permissions policy to prevent accidental exposure of data to other accounts in the system, encrypted storage system (drive encryption), and no user access to the head and storage node that runs file, scheduler, and configuration services for the cluster.

All the nodes in the system run CentOS7 (https://www.centos.org/) Linux operating system. The head node runs the SLURM job scheduler (http://slurm.schedmd.com/), the Puppet configuration management system http://puppet.com), the Razor provisioning system (https://docs.puppet.com/pe/latest/razor_intro.html) and the Network File System (NFS). The storage node runs ZFS to export large project storage to the cluster and encryption is handled on a per-project basis.

Data Center Physical Security and environment

The Secure4 environment resides in the University of Missouri Data Center. The MU Data Center, operated by the Division of Information Technology (DoIT), provides a physically secure and environmentally controlled facility that houses computing systems for MU and other University of Missouri organizations. The center offers the following standards services and features:

- Environmentally controlled cooling and humidity
- Redundant power to all rack locations
- Uninterrupted power supply and generator backup
- Fire Suppression systems
- Physical security with video surveillance
- 24x7x365 system monitoring, alerts and staff response

Network Security

All systems in the data center are located behind a firewall with a default configuration of DENY ALL. Exceptions to the default firewall configuration are managed on an IP by IP basis. The Secure4 environment will be reviewed annually by the University's Information Security & Access Management team (ISAM), consisting of certified security professionals who use industry standard best practices to evaluate the security posture of the environment.

In addition to the data center firewall, a separate firewall will be used to segment the Secure4 environment from other systems within the data center. The use of a privileged Virtual Private Network (VPN) group will be required to gain access through this firewall irrespective of whether researchers are on or off the campus network. Principal Investigators (PIs) are the only authorized approver of accounts accessing their

projects. Day to day account management of the VPN group will be MU's information security team. Annual review of the VPN group members will be handled by the MU Research Computing Support Services team.

Data Segmentation

Users and data will be isolated. Users login to the system with a user-project account and can only access a single project with the user-project account, thus, preventing the spread of data from one project to another. The account will belong primarily to the user-project Unix group and a "project" group for the project.

The umask will be set to 007 making data available by default to only the user-project account or the "project" group (default not world readable). This policy is enforced and logged using SELinux on every system. NFS will carry these attributes across the cluster. All policy events will be logged and transmitted to and stored in a secure remote logging service.

Encryption

Data in transit is encrypted using Secure Shell (SSH) through the creation of a secure SSH key and passphrase on the researcher's workstations. This will be generated uniquely for the Secure4 environment prior to authorization being granted.

The Secure4 environment will utilize hard disk encryption to protect the data at rest using Linux Unified Key Setup (LUKS) encryption. LUKS encryption conforms to the Transmission Sleeve Kit 1 (TSK1) secure key setup scheme and operates based on an enhanced version of cryptsetup, using dm-crypt. A portable USB drive containing the decryption key is stored in a locked safe with limited access privileges, managed by the University's information security team.

Access to the Secure4 Environment

Only University of Missouri projects that meet the eligibility requirements in the Project Data section are eligible to use the Secure4 environment. Principal Investigators (PI) are the only individuals authorized to request access for members of their research team. Only requests originating from a PI are trusted. PI's and their research team members must work with their information technology professionals to create the secure SSH key pair and passphrase necessary to access the Secure4 environment. Accounts that have been granted access to the Secure4 environment are renewed on an annual basis. PIs are responsible for quickly communicating changes in research team/account status to the MU Research Support Computing team.

All members of the research team having access to the Secure4 environment must minimally have received IT Security training, prior to using the Secure4 environment. It is the responsibility of the PI to ensure this training has been completed.

Researcher Workstation Security

Only University issued devices are allowed to access the Secure4 environment. It is the responsibility of the researcher's IT Professional to ensure that workstations with access to the Secure4 environment comply with University polices for DCL4 workstation management as well as with the access/encryption requirements in this document. Requirements include but are not limited to the following:

- Local Firewall must be turned on
- Remote desktop access to workstations accessing Secure4 must be disabled.
- External services such as ftp, http, https, ssh, RDP, VLC, and other remote shell or file capabilities are not allowed
- Use of MU provided or authorized encryption
- Workstations and user accounts must be managed by UM/MU Active Directory and users must use their SSO (Pawprint) credentials to login to the workstation.

- Up-to-date anti-virus
- Up-to-date system patches
- Compliance with all other DCL level 4 standards

IT-Professionals will review and report that workstations comply with requirements quarterly to Research Computing Support Services.

While not prohibited, researchers are encouraged to use dedicated workstations free from unnecessary applications and programs. When dedicated computers are not feasible, researchers should refrain from personal or unnecessary web browsing and should also be careful to not open untrustworthy emails, download executables or open suspicious files as doing so could expose workstations to many forms of malware including viruses, worms and key loggers.

Data Auditing and Logging

Security event logs for the head node, login node, and compute nodes will be immediately shipped to the University of Missouri's secure logging system and the MU Research Computing Support Services logging system, including but not limited to the following events:

Security events (/var/log/secure)

- Login and Logout events
- Login failures
- Elevated privilege commands (sudo)

Audit events (/var/log/audit/audit.log)

- Secure shell events (login/logout/failures)
- SELinux policy violations.

Data Destruction

Data will be deleted from all user-project accounts and shared project storage associated with the project when the project is done. Decryption keys for large project storage will be destroyed when the project is done.

Drives are encrypted with Luks with the decryption key on a USB drive that is only present during boot. Data is not accessible without both the USB key and the drive. When a node is decommissioned the USB drive will first be securely erased and destroyed and after destruction the hard disk will removed from the machine and the secure enclosure and will be shredded. At no time will the USB drive and the hard disk be in the same location outside the secure rack. All movements will be logged.

Procedures

To Request the following activities, please send an email to rcss-support@missouri.edu

- Request a project
- Request, modify, delete secure4 accounts
- Request VPN set up information
- Request Data/Project deletion

UM and MU Policies and Procedures

Institutional Research

- http://ir.missouri.edu/
- https://www.umsystem.edu/ums/rules/collected_rules/research/ch410/410.010_research_involving_humans_in_experiments

Security policies

- https://www.umsystem.edu/ums/is/infosec
- https://www.umsystem.edu/ums/is/infosec/sections-portable-storage
- https://www.umsystem.edu/ums/is/infosec/standards-password
- https://www.umsystem.edu/ums/is/infosec/standards-logging

Data classification

- https://www.umsystem.edu/ums/is/infosec/classification
- https://www.umsystem.edu/ums/is/infosec/classification-definitions

Workstation standards

- https://www.umsystem.edu/ums/is/infosec/sections-workstation
- https://www.umsystem.edu/ums/is/infosec/standards-workstation-management

System and application standards

- https://www.umsystem.edu/ums/is/infosec/sections-sysapp

References

- Controlled Unclassified Information (CUI): https://www.federalregister.gov/d/2016-21665
- NIST SP 800.171: https://csrc.nist.gov/publications/detail/sp/800-171a/draft

Revision History

- 2016-06-14: https://hdl.handle.net/10355/52851
- 2018-02-14: Expand scope to sponsored research, technology transfer, and internal data. Increase security requirements on workstations.