SEAGATE

Storage Consolidation Capabilities with End-to-end Multi-Level Security

Table of Contents

AUTHORS:

Neal A. Martin
Seagate Government Solutions
2300 Dulles Station Blvd., Suite 230
Herndon, VA 20170
nmartin@seagategov.com

Mark McKnight
Seagate Government Solutions
2300 Dulles Station Blvd., Suite 230
Herndon, VA 20170
Mark.McKnight@seagategov.com

REV: March 2017

SEAGATE

Storage Consolidation Capabilities with End-to-end Multi-Level Security

ABSTRACT— The use of multi-level security (MLS) software eliminates data silos, reduces hardware and personnel expenses, and speeds data access. In critical environments, this improved data access can save lives and becomes the difference between mission failure and success. The ClusterStor SL200 from Seagate is an MLS-enabled scale-out storage appliance featuring a high performance and linearly scalable parallel file system to provide faster solutions to the most demanding data intensive applications. MLS allows control of the data to which a user has access, based on the authentication information for the user; these capabilities are provided by the Red Hat SELinux operating system and the secure Lustre file system. The MLS ecosystem provides a number of platforms and tools that provide important capabilities for protecting various types of data at all security levels.

Keywords—multi-level security, HPC storage, parallel file system, Lustre, SELinux

I. Introduction

Government and commercial organizations have the need to segregate data either to protect customer information, to protect company confidential information, or to protect government confidential or classified information. One of the ways to eliminate the need for these separate security domains is through the use of multi-level security (MLS) software and processes with the goal of maintaining security classification levels throughout a customer's distributed systems. The rules are fairly straightforward and operate on sensitivity (for example, Secret or Top Secret) and a set of categories. System processes, acting on behalf of a user, can only access data at the same or lower sensitivity and only if that process is authorized for that particular set of categories. For distributed MLS systems, providing mission-critical information often involves making data at different security classification levels available on the network across security domains. Fully bidirectional transfer of information between security levels often involves the use of a Cross Domain Solution (CDS) with the CDS bridging several network levels and providing bi-directional traffic between any pair of levels as allowed by the security policy of the CDS.

II. Why MLS Is Needed

A. Data Silos

The requirement to segregate data often results in creating information silos, in which data of a particular classification is maintained separately to ensure that it is not inadvertently accessed by anyone not permitted to see data of that classification. Perhaps of even greater concern is loss of customer or company confidential data to hackers, such as happened with the breach at the Office of Personnel Management (OPM), in which it was discovered that the personnel data of 22 million current and former Federal government employees had been stolen.

B. Hardware Expenses

Separate information silos, holding data of a specific classification, require separate and distinct storage, networking, and in some instances, separate servers to process the data. In addition to tripling or quadrupling hardware components, there are power, cooling, data center space, and software needed to support the separate silos of data. Resource replication is particularly onerous when missions necessitate creating extremely expensive high performance computing, network backbone and storage networks for geospatial data collection and distribution. Another added cost is the time and personnel resources required to maintain these separate data silos.

C. Data Access Delays

With the separate data silos, there is often a need to move data from one classification level to another when the data stored in one of these silos is needed in another security domain. With the policies that can be set up with an MLS-enabled system, this data transfer is no longer needed and much of the human element in allowing or denying access to data of different classifications is removed. This results in a more secure environment and eliminates the costs of having to perform this data movement. Another important consideration is the ability to provide more timely access to the required data. This can be crucial when lives depend on having access to the information as quickly as possible. An example is surveillance data that must be matched against a database of terror suspects at a higher security classification; if the surveillance data and the database reside on the same storage, seconds or minutes can be saved resulting in the apprehension of the suspect before he can carry out his plans.

III. CS-SL200 Description

A. Hardware

Currently based on Seagate's proven ClusterStor 9000, the CS-SL200 is the high performance computing (HPC) industry's first scale-out storage appliance featuring a high performance and linearly scalable parallel file system with the ability to meet ICD 503 (DCID 6/3 PL4) requirements integrated within a single environment. In the ClusterStor system, the operating environment is integrated into distributed data storage elements called Scalable Storage Units (SSUs), which integrate block storage, network, RAID data protection, storage enclosure services and file system components into a single building block that delivers linear scaling in performance and data storage capacity. Please see figures 2 and 3 at the end of this paper for details on the architecture of the CS-SL200.

The CS-SL200 has the following capabilities for scalability:

- Raw file system capacity of 3.4 PB per rack using 6TB SAS HDDs
- Up to 7 SSUs per rack and up to 14 embedded server modules per rack
- Client access via InfiniBand (QDR or FDR) or Ethernet (10 GbE or 40 GbE)

Future SL200 releases are expected to have higher performance, high capacities, and other connectivity options to provide greater flexibility.

B. Software

The SL200 combines Seagate's Lustre [1] high performance file system solution with its ClusterStor Secured Operating Environment. The Secured Operating Environment is based on enhancements to SELinux to meet the security, performance and scalability needs of the most demanding customers. The ClusterStor storage is the result of adding optimizations to the Lustre file system and embedding it within its modular scale-out hardware architecture in order to address a wide spectrum of HPC and Big Data needs. The high performance of the Lustre file system allows users to more quickly solve the most demanding data intensive, large I/O applications.

1) SELinux

SELinux is the primary Mandatory Access Control (MAC) mechanism built into a number of GNU / Linux distributions; the specific distribution used on the CS-SL200 is Red Hat version 6.5. If SELinux is enabled, it functions to:

- Provide policy which defines what access to resources and operations on them (e.g. read, write) are allowed (i.e. SELinux stops all access unless allowed by policy). The policy design, implementation and testing against a defined security policy or requirements is needed to verify correct implementation.

- SELinux can confine an application within its own domain and allow it to have the minimum privileges required to do its job. Should the application require access to networks or other applications (or their data), then (as part of the security policy design), this access would need to be granted.
- If an application attempts to perform any actions that are not allowed by policy (intentional or otherwise), SELinux would stop these actions.
- If an application tries to perform an action that is allowed by policy, SELinux may contain any damage that can be done intentionally or otherwise.
- User login sessions can be confined to their domains; this allows the clients that they run to be given only the privileges that they need.
- It is very easy to add new rules to SELinux policy; however, this has the potential to start opening security holes, so it is best practice to be sure what rules are really required [2].

With Kerberos enhanced authentication and underlying data encryption capabilities included in the SL200, government and intelligence organizations gain the ability to better implement, deploy and enforce their end-to-end security access control policies and ensure strict need-to-know compliance. This enhanced level of protection guards against external and insider threats, as well as providing a way to deploy and enforce sophisticated mission policies such as CDS collaboration between cooperating agencies and multi-disciplinary mission teams.

2) Secure Lustre File System

The CS-SL200 utilizes a Lustre 2.5 base and provides a ClusterStor Secure Lustre file system. The functional overview of ClusterStor Secure Lustre is provided below:

- Enables per label authorization access permissions by checking at the Metadata Server (MDS).
- Enables the client label copy from the client to the corresponding MDS inode entry.
- Transfers labels from the MDS to the Object Storage Target (OST); the label on the Object Storage Server prevents unauthorized file access by bypassing client and MDS label checking.
- The Red Hat SELinux management interface creates and manages the label security policy; this allows the security administrator the capability to assign a security label per file and Linux user process label.
- The security policy provides file access authorization by comparing both per file and per user labels.
- ClusterStor has developed a unique Lustre client that creates a per file metadata extended attribute that is used to define security access permissions.
- The metadata extended attribute conforms to the SELinux label format and usage definition enabling the SELinux security administrator to create fine grain access permissions.

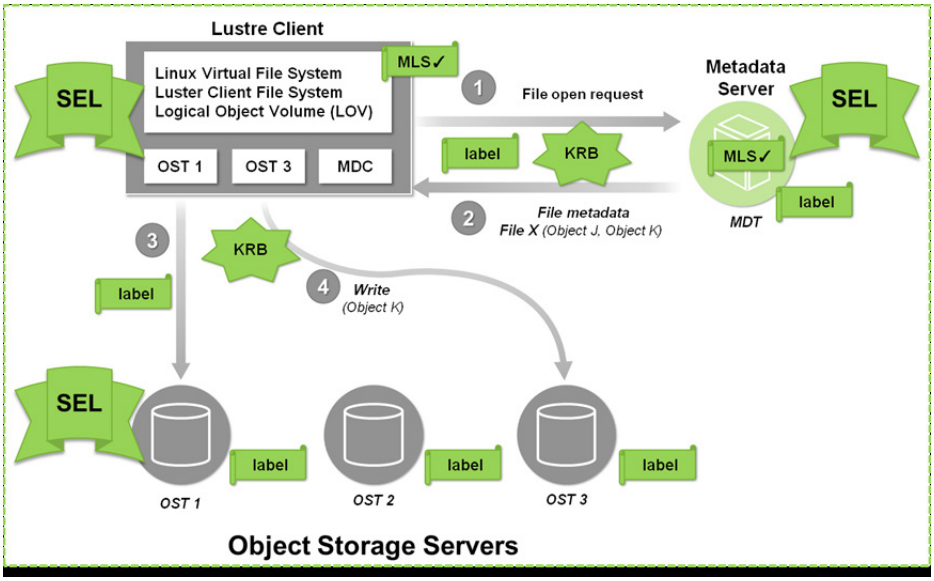

Fig. 1. Secure Lustre Components

IV. How MLS Works

The following sections describe the MLS capabilities offered with the ClusterStor SL200.

A. Compartmented Systems

The CS-SL200 has the distinctive advantage of providing a multilevel security environment for storing large amounts of data on a single appliance and file system. This security allows for a multi-tenant system while also ensuring that data is compartmented (separated) based on sensitivity or other factors, such as category or need-to-know. The ClusterStor Secured Operating Environment is based on Red Hat Enterprise Linux with SELinux enabled, enforcing the MLS policy. SELinux provides Mandatory Access Controls (MAC) that can give subjects roles (including programs and tenants) and secure data against unauthorized access. In addition, type enforcement provides the advantage of limiting access to, and use of, different domains within the system. This means that programs intended for use in specific domains cannot access other domains without a special exception written into the security policy, preventing accidental or intentional damage to that domain or access to unauthorized data. Data stored on the system also has the advantage of being labeled using security and category labels. This means that the data can be compartmented based on these labels, and if a subject's access does not match up with the data's label, that subject cannot access the data. These labels provide separation of data without separate networks or systems.

Kerberos authenticates users so that external and internal threats cannot circumvent MAC, type enforcement, or labels. Kerberos security servers, after authenticating a user or client, grant them tickets that allow access into, and authenticate the user to, the larger system. This means that external threats will not be able to authenticate themselves to the system and access is denied, while internal threats granted a ticket will only have limited access to the system. This provides a highly secure environment where data of different types and sensitivities can be stored separately and secured against unauthorized access in a single file system with multiple tenants.

B. MLS Policy

MLS policies in SELinux follow a model of least-privilege, in which everything is denied and policies are written to provide exceptions for each function in the system such that it is only given the minimal access required to function. If a user tries to modify or use a service, program or function that the user does not have permission to use, the access is denied and the activity is logged. A security context, or security label, is the mechanism used by SELinux to classify resources, such as processes and files, on a SELinux-enabled system. SELinux contexts allow enforcement of rules for how and by whom a given resource should be accessed. SELinux has three basic modes of operation, of which "enforcing" is set as the installation default mode in which the SELinux security policy is enforced on the system, denying unauthorized access and logging actions.

With MLS, there is no need for separate security domains and data guards to control data flow. Data at different classification levels and compartments can co-exist on the same storage and the MLS software controls the data to which a user has access, based on the authentication information for the user. SELinux requires that policies be defined correctly and complemented with other file permission controls to ensure the best protection. For proper functioning of MLS, there must be a configuration set up by the system administrator to define the policy. The security policy is extremely flexible to allow the definition of the policy that is needed, but it does require good familiarity with how SELinux functions to ensure appropriately defined security controls. In addition, the system administrator is responsible for tasks such as enabling kernel auditing as well as dumping and viewing audit logs.

C. Ease of Access

As noted previously, data movement from one classification level to another is no longer needed. This area is sometimes overlooked but that has an impact on both network traffic and personnel resources for movement of data. With the flexible policies that can be set up with an MLS-enabled system, this data transfer is eliminated and much of the human element in allowing or denying access to data of different classifications is removed. This results in a more secure environment and eliminates the costs of having to perform this data movement. This improved data access can result in more rapid availability of critical data that can save lives and money.

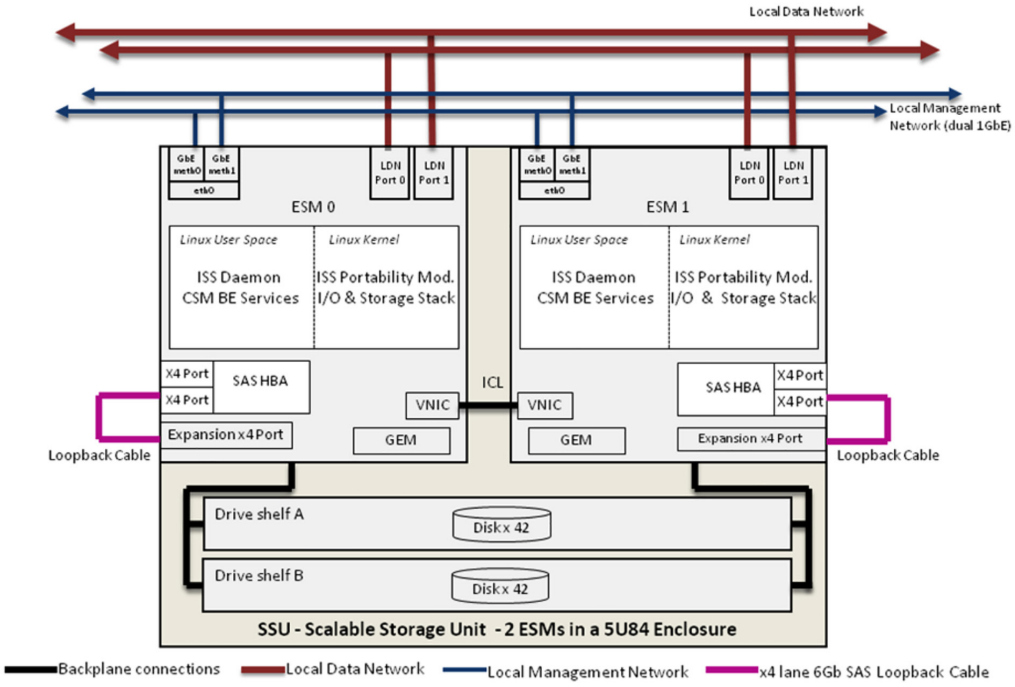

Fig. 2. SSU Block Diagram

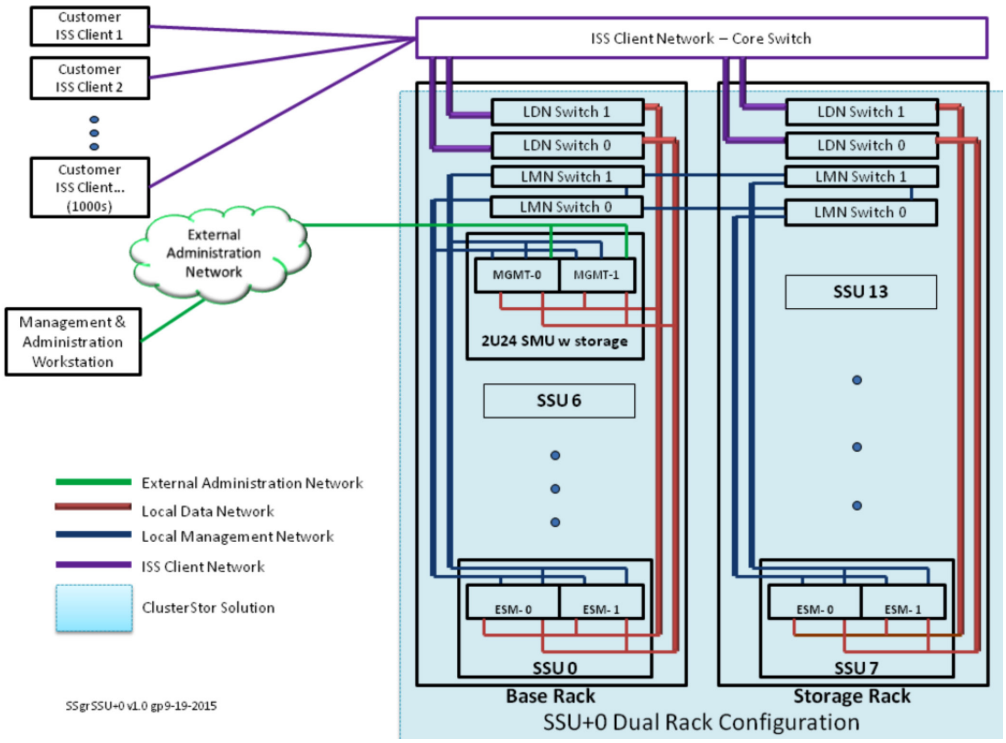

Fig. 3. CS-SL200 Rack Configuration

V. MLS Ecosystem

A consortium of companies has been assembled by Lockheed Martin, under contract and guidance from the NRO Central Super Computing Facility (CSCF), to create a commercially available MLS cloud solution with virtualization for a US Government customer. This consortium provides a full MLS solution ecosystem with all software and hardware components that have ICD 503 certification to handle multiple levels of secure data. This solution is authorized to be connected simultaneously to different networks operating at different security levels and separates the functions of system administration, security administration, and users so that there is no “superuser” access. This solution provides for fully encrypted data paths and full data forensics for automatic intrusion detection and long-term correlation of events. In addition to the CSCF and its leadership role in ICD 503 certification, MLS ecosystem participants include Seagate, Cray Computer [3], SGI [4], Lockheed Martin [5], Red Hat [6], Splunk [7], Bay Microsystems [8], Crunchy Data Solutions [9], Semper Fortis [10], and Fornetix [11].

VI. Conclusion

The use of MLS eliminates data silos, reduces hardware and personnel expenses, and speeds data access. In critical environments, this improved data access can save lives and becomes the difference between mission failure and success. The ClusterStor SL200 from Seagate is an MLS-enabled scale-out storage appliance featuring a high performance and linearly scalable parallel file system to provide faster solutions to the most demanding data intensive applications. With its compartmented data and MLS policy capabilities, it allows control of the data to which a user has access, based on the authentication information for the user. These MLS capabilities are provided by the Red Hat SELinux operating system and the secure Lustre file system. The MLS ecosystem provides a number of platforms and tools for extending the MLS capabilities provided by the CS-SL200. This shared storage platform, along with other computing platforms and applications to be included in the MLS ecosystem, provide a powerful set of tools for protecting various types of data at all security levels.

Acknowledgment

The authors wish to acknowledge the help and support on this paper of the following individuals:

Mike Moritzkat, Bill Downer, and Henry Newman

References

[1] Inside The Lustre File System; http://www.seagate.com/files/www-content/solutions-content/cloud-systems-and-solutions/high-performance-computing/_shared/docs/clusterstor-inside-the-lustre-file-system-ti.pdf.

[2] The SELinux Notebook - 4th Edition; http://freecomputerbooks.com/books/The_SELinux_Notebook-4th_Edition.pdf.

[3] Cray Computer website; http://www.cray.com/.

[4] SGI website: http://www.sgi.com/.

[5] Lockheed Martin website; http://www.lockheedmartin.com/ .

[6] Red Hat website; https://www.redhat.com/en.

[7] Splunk website; http://www.splunk.com/.

[8] Bay Microsystms website; http://www.baymicrosystems.com/.

[9] Crunchy Data website; http://crunchydata.com/.

[10] Semper Fortis website; http://www.semper-fortis.com/

[11] Fornetix website; https://www.fornetix.com/