SEAGATE GOVERNMENT SOLUTIONS

MLS AND LUSTRE FOR SECURE ENVIRONMENTS

Henry Newman | July 25, 2017

Agenda

More than just changes to Lustre are required

- ➢ **ICD/NIST background**
- ➢ **Security components**
- ➢ **What happens in the Lustre data path**
- ➢ **Final Thoughts**

I/O is and was hard

It has never been easy

Though SGI CXFS supported MLS on both Trusted IRIX and SELinux scaling was always a problem even on the UV system

- Using a parallel distributed memory system with MLS is relatively new and difficult
  - Beside the hardware and software ecosystem Lockheed and Leidos are both providing integration services given:
    - **Complexity to system integration and maintenance**
    - **Paperwork needed to get ATO**

Even with this, MLS is still far more cost effective than multiple systems and multiple security levels and allows data fusion not possible without MLS given 1 way guard performance

What are the requirements

What is ICD-503 and what are the security controls

ICD 503 was issued by the DNI (Director of National Intelligence) in 2008 and replaced DCID 6/3 was effectively “rescinded and replaced” almost in its entirety

ICD-503 uses a RMF (Risk Management Framework) based on NIST 800-53 (Security and Privacy Controls for Federal Information Systems and Organizations) and also NIST 800-37 (Guide for Applying the Risk Management Framework to Federal Information Systems) and the Committee for National Security Systems (CNSS) Instruction 1253 (Security Categorization and Control for National Security Systems)

There are many NIST standards that must be followed and are in regular change based on new security threats

What does this translate to

It is all about requirements

The Bell–LaPadula Model (BLP) is framework used for enforcing access control in

- Security labels range from the most sensitive

SELinux conforms to this model

- SELinux was added to Linux in 2003

Other requirements include but not limited to:

- Logging each activity by all users
- Role based access control separating users, system administrators, security administrators and audit log administrators
- Correlation of logs to address insider threat

This translates to:

Functional Requirements

Red Hat SELinux running in enforcement mode

- You must run a Common Criteria Certified OS on clients, OSS, MDS and management framework
  - No one wants to redo all of the testing for Common Criteria
    - **CentOS, Debian etc. will not be considered**
  - You will not meet ICD-503 requirements nor will pass security inspection and get an ATI
- All logs must be collected and processed
  - Including management system logs

But it is more than just SELinux; you need databases, job schedulers…

- You need an ecosystem not just an operating system

What is an ecosystem

What is needed

Hardware

- Need a supply chain that allows the system to be flashed and loaded in the USA
  - Firmware
  - Likely FIPS-140v2 disk drives are required to support DAR NIST requirement

Software

- MLS aware software is important
  - Kerberos to prevent rogue clients
  - Databases that is MLS aware
  - Job Scheduler that is MLS aware
  - MPI communications that are MLS aware

Other considerations

Things on the short term horizon

Key management for NIST compliant FIPS-140v2 drives

- NIST mandates encryption at test
- Who manages the keys, rotates the keys and determines key entropy requirements

Common criteria for other system components

- Is this a requirement for ATO
- Other certifications like CSFC (Commercial Systems For Classified)
- Keeping up with requirement and changes is and will be ongoing

Secure Data Appliance

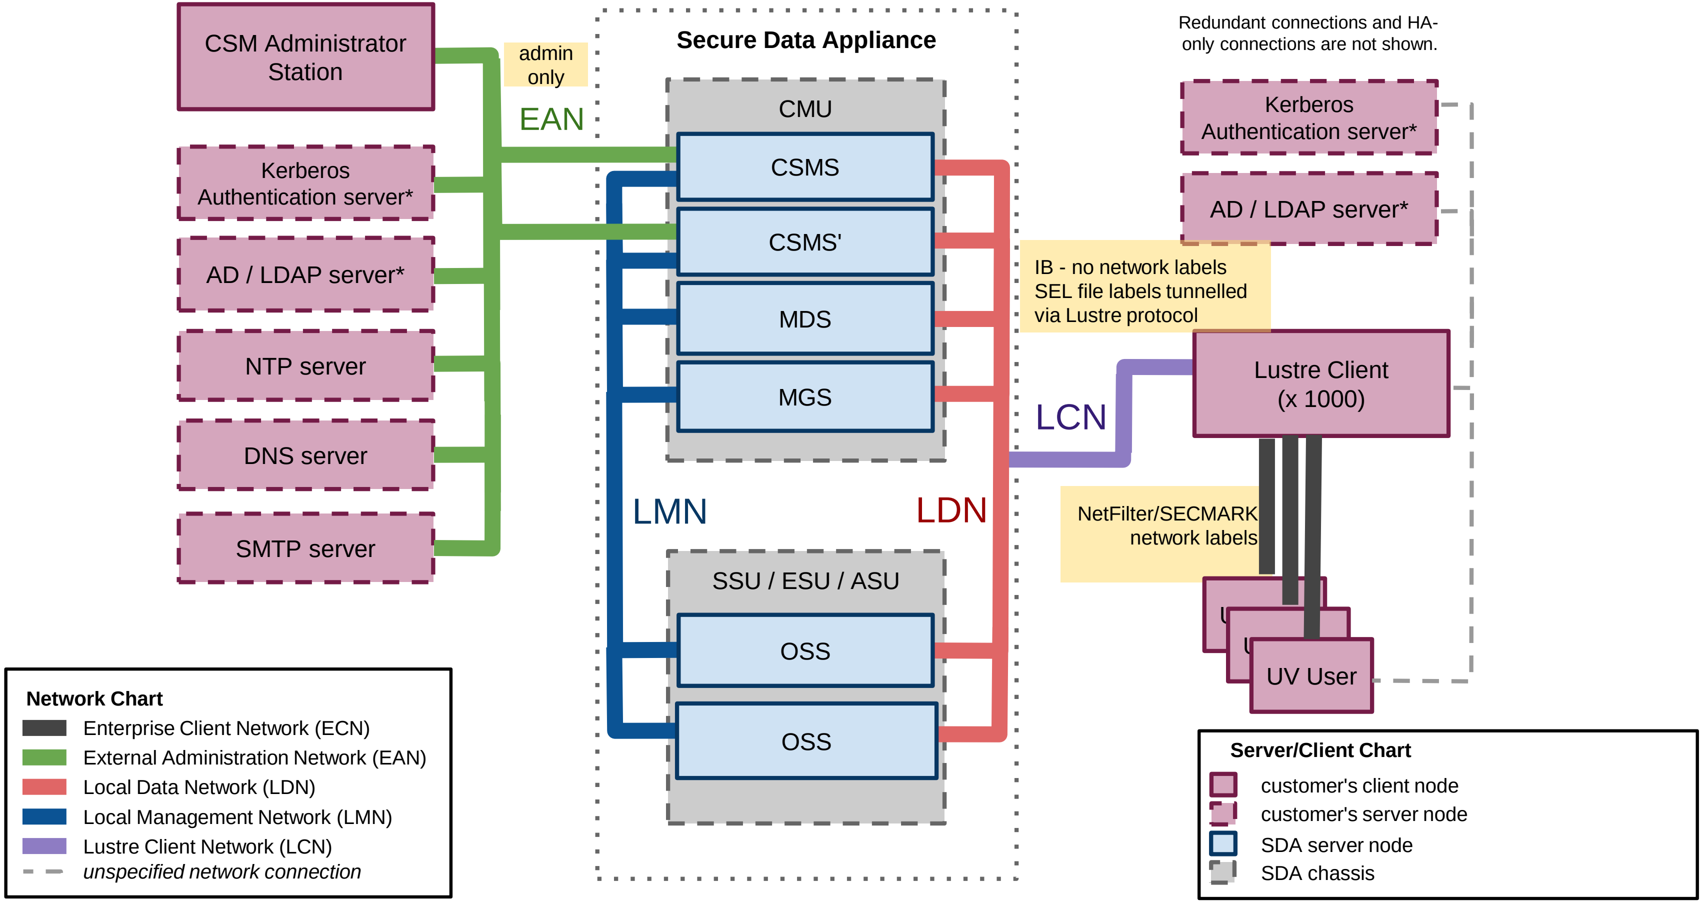

Redundant connections and HA-only connections are not shown.

Network Chart

- ⬛ Enterprise Client Network (ECN)
- 🟩 External Administration Network (EAN)
- 🟥 Local Data Network (LDN)
- 🟦 Local Management Network (LMN)
- 🟪 Lustre Client Network (LCN)
- - - *unspecified network connection*

Server/Client Chart

- □ customer's client node
- ▨ customer's server node
- □ SDA server node
- ▨ SDA chassis

* Kerberos and AD/LDAP servers, if present, must be common to the SL220 and Lustre clients

Lustre flow diagram with SELinux: open(2)

- ❑ Open is open RPC with embedded client SELinux context label
- ❑ ✓ is SELinux module verification of context with EA
- ❑ EA is the security selinux extended attribute that holds a file's security context

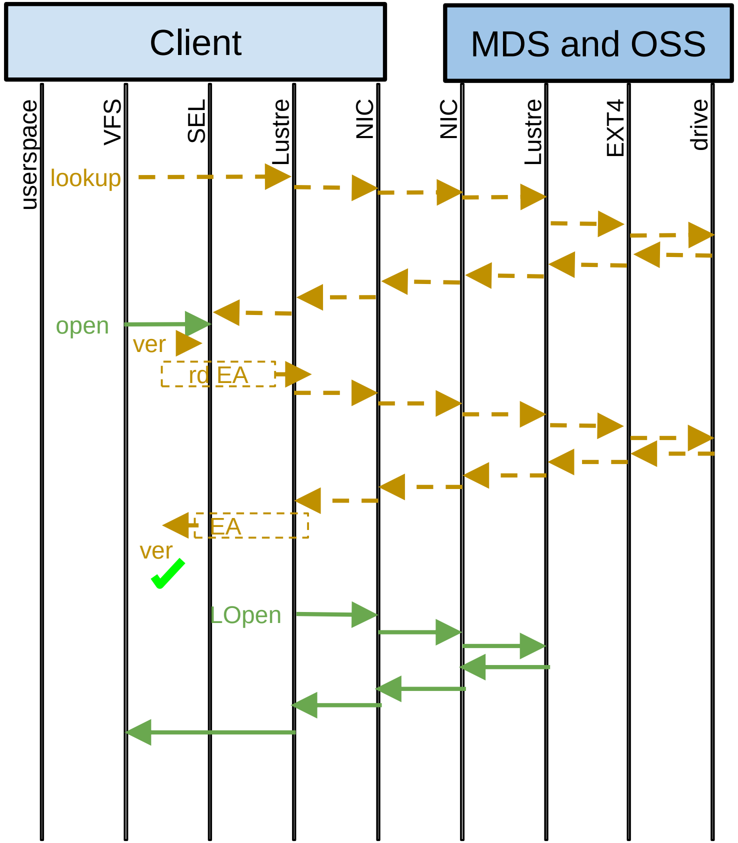

Dashed lines show potentially cached operations.
- lookup and getxattr

SELinux label check on client

SDA flow diagram: write(2)

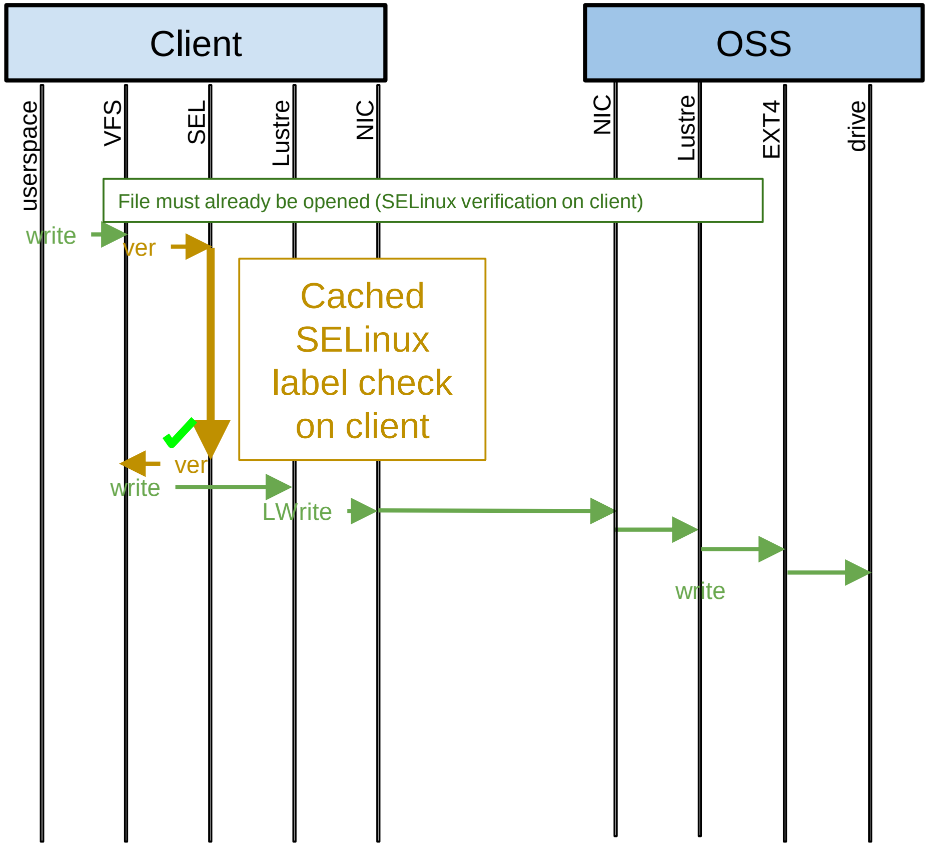
Client
OSS
userspace
VFS
SEL
Lustre
NIC
NIC
Lustre
EXT4
drive
File must already be opened (SELinux verification on client)
write
ver
Cached SELinux label check on client
ver
write
LWrite
write

- ❑ LWrite is RPC with embedded client SELinux context label
- ❑ ✓is SELinux module verification of context with EA
- ❑ EA is the security selinux extended attribute that holds a file’s security context

Maintenance is hard and complex

Lots to think about

Need to deal with patches both critical and non-critical

- Too often requires lots of approvals from security and work to re-certify
- Too little and security gets nervous about the need for patches
  - Think the 3 little bears (Just right)

Supply chain issues are becoming more and more important

- As are FIPS-140 DAR and key management

In some ways just having the Lustre client working is the easy part

- You need to think big picture

But the rewards in terms of cost of running and MLS environment is huge

- Otherwise we would not have done it

Final Thoughts

If it was easy then it would have been done long ago

Just having a CC OS is not enough

Just running SELinux in enforcement mode is not enough

Just having the Lustre client check SELinux context is not enough

Nor adding the MDS

Nor adding the OSS

Nor adding the management framework

You need all of these plus the logging and security policy to meet the requirements and the ecosystem to support and manage the system

Final Thoughts

MLS is hard with Lustre or any other file system

SELinux has been around for almost 15 year but was only used in guards until recently

- The reason in my and others opinion is the lack of an eco-system

MLS has had a long history in HPC and a long history in general

- From DOD Orange Book in the 1980s and Secure UNICOS in the 1990s along with Trusted IRIX, Trusted Solaris and Trusted AIX, to DCID-6/3 and now ICD-503 and SELinux
- Which each of these implementations there were things missing or market pressures that prevented wide spread adoption

SEAGATE | GOVERNMENT SOLUTIONS

THANK YOU FOR LISTENING