



Configuring SELinux for File Systems based on Intel* EE for Lustre* Software

Partner Guide

High Performance Data Division

Software Version: 3.0.0.0 or later

World Wide Web: <http://www.intel.com>

Disclaimer and legal information

Copyright 2016 Intel® Corporation. All Rights Reserved.

The source code contained or described herein and all documents related to the source code ("Material") are owned by Intel® Corporation or its suppliers or licensors. Title to the Material remains with Intel® Corporation or its suppliers and licensors. The Material contains trade secrets and proprietary and confidential information of Intel® or its suppliers and licensors. The Material is protected by worldwide copyright and trade secret laws and treaty provisions. No part of the Material may be used, copied, reproduced, modified, published, uploaded, posted, transmitted, distributed, or disclosed in any way without Intel's prior express written permission.

No license under any patent, copyright, trade secret or other intellectual property right is granted to or conferred upon you by disclosure or delivery of the Materials, either expressly, by implication, inducement, estoppel or otherwise. Any license under such intellectual property rights must be express and approved by Intel® in writing.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL® ASSUMES NO LIABILITY WHATSOEVER AND INTEL® DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel® Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL® AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL® OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL® PRODUCT OR ANY OF ITS PARTS.

Intel® may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel® reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Before using any third party software referenced herein, please refer to the third party software provider's website for more information, including without limitation, information regarding the mitigation of potential security vulnerabilities in the third party software.

Contact your local Intel® sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel® literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>.

Intel® and the Intel® logo are trademarks of Intel® Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit.
(<http://www.openssl.org/>)

Contents

| | |
|-------------------------------|----|
| About this Document | ii |
| Document Purpose | ii |
| Intended Audience..... | ii |
| Conventions Used | ii |
| Related Documentation | ii |
| Introduction..... | 1 |
| An Overview of SELinux | 1 |
| Requirements..... | 2 |
| Configuring SELinux..... | 2 |
| Status of SELinux | 2 |
| SELinux Security Context..... | 3 |

About this Document

Document Purpose

This document describes how to use SELinux with a Lustre* file system created with Intel® Enterprise Edition for Lustre* software and Intel® Manager for Lustre* dashboard. This document introduces SELinux, its functionality, and discusses how to perform simple configuration. For more advanced information on SELinux, see:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/SELinux_Users_and_Administrators_Guide/

Intended Audience

This guide is intended for systems integrators with a strong technical background in Linux system administration as well as Lustre file system deployment and management.

It is expected that readers have:

- Experience administering file systems and storage infrastructure, and familiarity with storage concepts such as RAID, SAN, and LVM.
- System management experience sufficient to install and configure a storage platform compatible with the requirements as defined in this guide.
- Proficiency in setting up, administering, and maintaining computer networks, including Ethernet, TCP/IP and InfiniBand where necessary. Some knowledge of Lustre networking (LNET) is required.
- Some experience with installing and managing Lustre file systems.

This document is not intended for end-users or application developers.

Conventions Used

Conventions used in this document include:

- # preceding a command indicates the command is to be entered as root
- \$ indicates a command is to be entered as a user
- <variable_name> indicates the placeholder text that appears between the angle brackets is to be replaced with an appropriate value

Related Documentation

- *Intel® Enterprise Edition for Lustre* Software, Version 3.0.0.0 Release Notes*
- *Intel® Manager for Lustre* Software User Guide*

- *Hierarchical Storage Management Configuration Guide*
- *Configuring LNet Routers for File Systems based on Intel® EE for Lustre® Software*
- *Installing Hadoop, the Hadoop Adapter for Intel® EE for Lustre®, and the Job Scheduler Integration*
- *Creating an HBase Cluster and Integrating Hive on an Intel® EE for Lustre® File System*
- *Creating a High-Availability Lustre® Storage Solution over a ZFS File System*
- *Upgrading a Lustre file system to Intel® Enterprise Edition for Lustre® software (Lustre only)*
- *Creating a Scalable File Service for Windows Networks using Intel® EE for Lustre® Software*
- *Intel® EE for Lustre® Hierarchical Storage Management Framework White Paper*
- *Architecting a High-Performance Storage System White Paper*

Introduction

A Lustre® file system is a network-based, distributed parallel storage platform consisting of metadata servers (MDS) and object storage servers (OSS). Metadata servers manage the file system namespace (directory structure, file information – inodes), while object storage servers contain the file contents. A management server (MGS) provides directory services, storing the configuration information for Lustre file systems and presenting that information to the population of Lustre servers and clients.

Intel® Manager for Lustre® software is an additional service supplied with Intel® Enterprise Edition for Lustre® software that enables administrators to create, manage, and monitor Lustre file systems. With this software, operators can configure and manage servers, as well as monitor file system health and performance.

Multi-level security (MLS) and role-based access control (RBAC) security and data protection policies are enforced by the Red Hat Enterprise Linux (RHEL) operating system (or CentOS) using software called Security Enhanced Linux (SELinux). SELinux was originally developed in the USA by the National Security Agency (NSA) and is one of the most common platforms used for providing mandatory access control (MAC) in Linux operating environments. MAC means that the operating system enforces the security contexts for applications and data running on a host – the user cannot override the context (e.g., by changing a file's permissions). SELinux support is provided to allow enforcement of policies on Lustre clients. This will allow an operator to implement MLS and RBAC policies on RHEL or CentOS clients running Lustre.

An Overview of SELinux

Security is an essential feature for IT infrastructure, services and applications in Life Sciences and Health Care, Government departments such as defense, and any organizations where sensitive information must be secured. In many industry sectors, regulatory compliance mandates the implementation of security and encryption protocols.

SELinux is a Linux security module that is built into the Linux kernel and provides fine-grained administrative controls for accessing files, based on the application, user, group or other role attributes and policy.

Lustre only supports SELinux enforcement on the client and not the servers. This means that the Lustre servers totally trust the client for access enforcement purposes and does not perform any additional verification. Lustre with SELinux enabled properly, initiates the security context of a file created on a SELinux-enabled client, and stores this security context on the metadata server side (MDS) via the `security.selinux` extended attribute.

Requirements

There are no requirements to using Lustre with SELinux.

Configuring SELinux

Configuring SELinux for client requires root or super-user access. This procedure is performed at the client.

SELinux is either enabled or, by default, disabled. In order to enable SELinux at boot time, you must edit the `SELINUX` variable in `/etc/selinux/config` to `enforcing` or `permissive`. If SELinux is set to `enforcing`, all SELinux security policies will be enforced. If SELinux is set to `permissive`, SELinux policy is not enforced, but all actions that would be denied if SELinux were being enforced are logged. Once SELinux is enabled at boot time, you can change between `enforcing` and `permissive` using the `setenforce` utility. The options for `SELINUXTYPE` are the default value of `targeted`, where only targeted network daemons are protected, and `strict`, which provides full SELinux protection. Note that Lustre only supports the `SELINUXTYPE` policy `targeted`.

SELinux policy rules are checked after the traditional UNIX permissions of user, group, and world permissions for read, write and execute. If the traditional UNIX permissions deny access, SELinux policy rules are not checked. When SELinux is enabled, the default action is deny; if no policy rule exists to allow access, access is denied.

To configure a Lustre client with SELinux, change `SELINUX=disabled` to `SELINUX=enforcing` in `/etc/selinux/config`. Save the file and reboot following any reboot procedures recommended by Lustre.

Status of SELinux

The following commands can be performed by any user; root permission is not required.

To check the status of SELinux, use the `sestatus` command at the client. When SELinux is enabled and set to `enforcing` mode, you should see:

```
# sestatus
SELinux status:                enabled
SELinuxfs mount:              /selinux
Current mode:                  enforcing
Mode from config file:        enforcing
Policy version:                24
Policy from config file:      targeted
```

You can get the current mode of SELinux on the command line using the utility `getenforce` and it will return `enforcing`, `permissive`, or `disabled`. For example, if SELinux is set to `permissive` mode, then `getenforce` returns:

```
# getenforce
Permissive
```

When SELinux is not enabled, you will see:

```
# getenforce
Disabled

# sestatus
SELinux status:   disabled
```

SELinux Security Context

All files created under SELinux are assigned labels that convey security-relevant information called the SELinux context. To view this information for an example file we call `dd_file_2`, use the following command:

```
# ls -lZ /lustre/scratch/dd_file_2
-rw-r--r--. root root unconfined_u:object_r:file_t:s0
/lustre/scratch/dd_file_2
```

In this case, SELinux provides a security context comprised of a user (`unconfined_u`), a role (`object_r`), a type (`file_t`), and a category/level (`s0`).

In Lustre, the security context is saved in the file `security.selinux xattr` on the MDT. To change or assign more categories to a file and to assign categories to users, follow the directions provided by the OS distribution.

For more advanced information about configuring and using SELinux, see:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/SELinux_Users_and_Administrators_Guide/