

Oracle® Exadata

Security Guide



26.1
F29252-36
July 2026

ORACLE®

Oracle Exadata Security Guide, 26.1

F29252-36

Copyright © 2008, 2026, Oracle and/or its affiliates.

Primary Authors: Peter Fusek, Glenn Maxey

Contributing Authors: Craig Brown, Dan Norris

Contributors: Yang Liu, Philip Newlan, Tina Rose, Kevin Simmons, Zheren Zhang

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Contents

Preface

Audience	i
Related Documents	i

1 Overview of Oracle Exadata Security

1.1	Survivability of Mission-Critical Workloads	2
1.2	Defense in Depth to Secure the Operating Environment	2
1.3	Least Privilege for Services and Users	3
1.4	Accountability of Events and Actions	3
1.5	Understanding Operating System Security of Oracle Exadata Storage Servers	3

2 Security Features of Oracle Exadata Database Machine

2.1	Using Security-Enhanced Linux	2
2.1.1	Enabling SELinux	3
2.1.2	Changing the SELinux Mode	3
2.1.3	Disabling SELinux	4
2.1.4	Displaying the SELinux Status	4
2.2	Restricting the Binaries Used to Boot the System	4
2.2.1	Enabling Secure Boot on Exadata System Hardware	5
2.2.2	Enabling KVM Guest Secure Boot	6
2.2.3	Managing Keys and Certificates Used with Secure Boot	7
2.2.3.1	Adding Keys for Secure Boot Using mokutil	9
2.2.3.2	Removing Keys for Secure Boot Using mokutil	11
2.2.4	Troubleshooting Secure Boot	13
2.3	Using Isolation Policies	13
2.3.1	Isolating Network Traffic	14
2.3.2	Isolating Databases	14
2.3.3	Isolating Storage	15
2.4	Network Access to Oracle Exadata Storage Servers	15
2.5	Controlling Access to Data	15
2.5.1	Controlling Network Access	16
2.5.2	Controlling Database Access	16

2.5.3	Controlling Storage Access	16
2.6	Using Cryptographic Services	17
2.7	Monitoring and Auditing of Oracle Exadata Database Machine	18
2.7.1	Monitoring and Auditing of Oracle Exadata Database Machine	18
2.7.2	Monitoring and Auditing Oracle Database Activity	19
2.7.3	Operating System Activity Monitoring on Oracle Exadata Servers	19
2.8	Maintaining Quality of Service	20
2.9	Using Oracle ILOM for Secure Management	20
2.10	Considerations for a Secure Environment	21
2.10.1	Identity and Access Management Considerations	21
2.10.2	Network Security Considerations	22
2.11	Understanding the Default Security Settings	24

3 User Security on Oracle Exadata

3.1	Default User Accounts for Oracle Exadata	1
3.2	Default Password Requirements	5
3.3	Default User Security Settings	5
3.4	Managing Password and Authentication Policies	6
3.5	Creating Oracle Exadata System Software Users and Roles	10
3.5.1	Overview of Creating Exadata System Software Users	11
3.5.2	Creating Roles and Getting Information about Roles	11
3.5.3	Granting and Revoking Privileges	12
3.5.4	Creating Users	12
3.5.5	Configuring Password Expiration for Users Accessing the Server Remotely	13
3.5.6	Granting and Revoking Roles	14
3.6	Security Policies for Oracle Exadata Storage Server Operating System Users	14
3.6.1	Changing a Password	14
3.6.2	Enabling the Security Policies for Operating System Users	15
3.6.3	Viewing Failed Operating System Password Attempts	15
3.6.4	Resetting a Locked Operating System User Account	15
3.7	Enhanced Default Security Settings for Storage Server Operating System Users	16

4 Keeping the Oracle Exadata Secure

4.1	Securing the Hardware	1
4.1.1	Getting the Rack Serial Number	2
4.1.2	Getting the Serial Numbers for Rack Components	2
4.1.3	Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch	5
4.1.4	Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36	5
4.1.5	Getting the Serial Number for a Cisco 4948 Ethernet Switch	6
4.2	Securing the Software	6

4.3	Managing SSH on Storage Servers	7
4.3.1	Locking a Cell	8
4.3.2	Unlocking a Cell Temporarily	9
4.3.3	Unlocking a Cell Permanently	10
4.3.4	Checking the Current Access Level for a Cell	11
4.3.5	Access Level Alerts from the Management Server	11
4.3.6	Managing SSH Equivalence for Storage Server OS Users	11
4.4	Configuring Data Security for Exadata Storage Servers	12
4.4.1	About Exadata Storage Server Security Modes	13
4.4.2	Best Practices for ASM-Scoped Security and DB-Scoped Security	14
4.4.3	About Security Keys	14
4.4.4	Setting Up ASM-Scoped Security on Oracle Exadata Storage Servers	17
4.4.5	Setting Up DB-Scoped Security on Oracle Exadata Database Machine	19
4.4.6	Changing Security Keys for ASM-Scoped Security or DB-Scoped Security	24
4.4.6.1	Upgrading ASM-Scoped Security Key for ASMCLUSTER	24
4.4.6.2	Changing the Assigned Key Value for ASM-Scoped Security	25
4.4.6.3	Changing the Assigned Key Value for DB-Scoped Security	26
4.4.7	Enabling Cell-to-Cell Operations	28
4.4.7.1	Configuring Simple Cell Access	28
4.4.7.2	Configuring LOCAL and REMOTE Cell Keys	29
4.4.7.3	Changing Between Simple Cell Keys and LOCAL and REMOTE Keys	30
4.4.8	Removing ASM-Scoped Security or DB-Scoped Security	31
4.4.8.1	Removing DB-Scoped Security	31
4.4.8.2	Removing ASM-Scoped Security	33
4.5	Maintaining a Secure Environment	35
4.5.1	Maintaining Network Security	35
4.5.2	Encrypting System Log Information	36
4.5.2.1	Overview of syslog File Encryption	37
4.5.2.2	Configure CA Server and Central rsyslogd Server	37
4.5.2.3	Configure a Client for SYSLOG Encryption	43
4.5.2.4	Confirming Syslog Encryption is Enabled	44
4.5.3	Guarding Against Unauthorized Operating System Access	44
4.5.3.1	About Advanced Intrusion Detection Environment (AIDE)	45
4.5.3.2	Managing AIDE Components	45
4.5.3.3	Adding Custom AIDE Rules	46
4.5.3.4	Managing AIDE Alerts when Updating Exadata Software	46
4.5.4	Updating Software and Firmware	47
4.5.4.1	Regenerate SSH Keys for ILOM Version 5	47
4.5.5	Ensuring Data Security Outside of Oracle Exadata Database Machine	48

5 Securely Erasing Oracle Exadata

5.1	Overview of Secure Eraser	1
5.2	Securely Erasing Database Servers and Storage Servers	4
5.3	Automatic Secure Eraser through PXE Boot	4
5.3.1	Automatic Secure Eraser through PXE Boot for X7 and Later Systems	5
5.3.2	Automatic Secure Eraser through PXE Boot for X6 and Earlier Systems	11
5.4	Interactive Secure Eraser through PXE Boot	16
5.5	Interactive Secure Eraser through Network Boot	21
5.6	Secure Eraser Syntax	27
5.7	Resetting Network Switches and Power Distribution Units to Factory Default	29
5.7.1	Resetting a Cisco Nexus 9336C-FX2 RoCE Network Fabric Switch to Factory Default Settings	30
5.7.2	Resetting InfiniBand Network Fabric Switches to Factory Default	32
5.7.3	Resetting the Cisco Management Network Switch to Factory Default Settings	32
5.7.4	Resetting Power Distribution Units to Factory Default	34
5.8	Actions After Using Secure Eraser	35

Preface

This guide describes security for an Oracle Exadata. It includes information about the components, the recommended password policies, and best practices for securing the Oracle Exadata environment.

Oracle Exadatas supported are Oracle Exadata and Oracle Zero Data Loss Recovery Appliance.

- [Audience](#)
- [Related Documents](#)

Audience

This document is intended for database administrators and network administrators responsible for securing an Oracle Exadata.

Related Documents

For more information, see the following documents:

- *Oracle Exadata System Overview*
- *Oracle Exadata Installation and Configuration Guide*
- *Oracle Exadata System Software User's Guide*
- *Oracle AI Database Security Guide*
- [Sun Datacenter InfiniBand Switch 36 Hardware Security Guide](#)
- [Oracle ILOM Security Guide For Firmware Releases 3.x and 4.x](#)
- [Oracle Server X8-2 Security Guide](#)
- [Oracle Server X7-2 Security Guide](#)
- [Oracle Server X6-2 Security Guide](#)
- [Oracle Server X5-2 Security Guide](#)
- [Sun Server X4-2 Security Guide](#)
- [Sun Server X3-2 Security Guide](#)

1

Overview of Oracle Exadata Security

Oracle Exadata is an engineered system that combines the optimized database performance of Oracle Database integrated with Oracle Exadata Storage Servers.

These core components are connected over a redundant RDMA Network Fabric that enables low latency, and high throughput network communication. There are 3 networks:

- Private Network - uses the RDMA Network Fabric for communications between database and storage servers in one or many physical racks.
- Client Network - used for communication from the client applications to services running on Oracle Exadata.
- Administration Network - used for managing the hardware of the Oracle Exadata components including database and storage servers, PDUs, and switches.

Within this framework, there are basic security principles that should be adhered to for all software and hardware. The following are the principles:

- Authentication: Authentication is how a user is identified, typically through confidential information such as user name and password, or shared keys. All components use authentication to ensure that users are who they say they are. By default, local user names and passwords are used for authentication. Shared key-based authentication is also available.
- Authorization: Authorization allows administrators to control what tasks or privileges a user may perform or use. Personnel can only access the tasks and privileges that have been given to them. Oracle Exadata system administrators can configure resources with permissions to control user access to commands, disk space, devices, and applications.
- Accounting and Auditing: Accounting and auditing maintain a record of a user's activity on the system. Oracle Exadata software and hardware features allow administrators to monitor login activity, and maintain hardware inventories.
 - User logons are monitored through system logs. System administrators and service accounts have access to commands that used incorrectly could cause harm and data loss. Access and commands should be carefully monitored through system logs.
 - Hardware assets are tracked through serial numbers. Oracle part numbers are electronically recorded on all cards, modules, and mother boards, and can be used for inventory purposes.

In addition to the basic security principles, Oracle Exadata addresses survivability, defense in depth, least privilege, and accountability. Oracle Exadata delivers a well-integrated set of security capabilities that help organizations address their most-pressing security requirements and concerns.

An example of how these and following security principles should be applied to the separate networks is:

- Management Network requires a boundary level security, where only trusted administrators can access this network
- Data Network requires securing data flow using encryption when multiple tenants or secure information is sent across this network

- As the "front end" to the database, the Client Network requires the strongest security requirements, restricting access to this network connection ensures outside threats can be mitigated.
- [Survivability of Mission-Critical Workloads](#)
Oracle Exadata Database Machine can prevent or minimize the damage caused from accidental and malicious actions taken by internal users or external parties.
- [Defense in Depth to Secure the Operating Environment](#)
Oracle Exadata Database Machine employs multiple, independent, and mutually-reinforcing security controls to help organizations create a secure operating environment for their workloads and data.
- [Least Privilege for Services and Users](#)
Oracle Exadata Database Machine promotes the principle of least-privilege.
- [Accountability of Events and Actions](#)
When an incident occurs, a system must be able to detect and report the incident.
- [Understanding Operating System Security of Oracle Exadata Storage Servers](#)

1.1 Survivability of Mission-Critical Workloads

Oracle Exadata Database Machine can prevent or minimize the damage caused from accidental and malicious actions taken by internal users or external parties.

As part of the Oracle Maximum Availability Architecture best practices, survivability is increased by the following:

- Ensuring that the components used have been designed, engineered, and tested to work well together in support of secure deployment architectures. Oracle Exadata Database Machine supports secure isolation, access control, cryptographic services, monitoring and auditing, quality of service, and secure management.
- Reducing the default attack surface of its constituent products to help minimize the overall exposure of the machine. Organizations can customize the security settings of Oracle Exadata Database Machine based upon the organization's policies and needs.
- Protecting the machine, including its operational and management interfaces, using a complement of open and vetted protocols, and APIs capable of supporting traditional security goals of strong authentication, access control, confidentiality, integrity, and availability.
- Verifying that software and hardware contain features that keep the service available even when failures occur. These capabilities help in cases where attackers attempt to disable one or more individual components in the system.

1.2 Defense in Depth to Secure the Operating Environment

Oracle Exadata Database Machine employs multiple, independent, and mutually-reinforcing security controls to help organizations create a secure operating environment for their workloads and data.

Oracle Exadata Database Machine supports the principle of defense in depth as follows:

- Offering a strong complement of protections to secure information in transit, in use, and at rest. Security controls are available at the server, storage, network, database, and application layers. Each layer's unique security controls can be integrated with the others to enable the creation of strong, layered security architectures.

- Supporting the use of well-defined and open standards, protocols, and interfaces. Oracle Exadata Database Machine can be integrated into an organization's existing security policies, architectures, practices and standards. Integration is critical as applications and devices do not exist in isolation. The security of IT architectures is only as strong as its weakest component.
- Conducting multiple security scans using industry-leading security analyzers to implement all high-priority security items prior to the release of each new Oracle Exadata System Software release.

1.3 Least Privilege for Services and Users

Oracle Exadata Database Machine promotes the principle of least-privilege.

Ensuring that applications, services and users have access to the capabilities that they need to perform their tasks is only one side of the least-privilege principle. It is equally important to ensure that access to unnecessary capabilities, services, and interfaces are limited. Oracle Exadata Database Machine promotes the principle of least-privilege as follows:

- Ensuring that access to individual servers, storage, operating system, databases, and other components can be granted based upon the role of each user and administrator. The use of role-based and multi-factor access control models with fine-grained privileges ensures that access can be limited to only what is needed.
- Constraining applications so that their access to information, underlying resources, network communications, and local or remote service access is restricted based upon need.

Whether caused by an accident or malicious attack, applications can misbehave, and without enforcement of least privilege, those applications may be able to cause harm beyond their intended use.

1.4 Accountability of Events and Actions

When an incident occurs, a system must be able to detect and report the incident.

Similarly, when an event cannot be prevented, it is imperative that an organization be able to detect that the event occurred so that proper responses can be taken. Oracle Exadata Database Machine supports the principle of accountability as follows:

- Ensuring each of the components used in Oracle Exadata Database Machine supports activity auditing and monitoring, including the ability to record login and logout events, administrative actions, and other events specific to each component.
- Leveraging features in Oracle Database to support fine-grained, auditing configurations. This allows organizations to tune audit configurations in response to their standards and goals. Administrators can ensure that critical information is captured, while minimizing the amount of unnecessary audit events.

1.5 Understanding Operating System Security of Oracle Exadata Storage Servers

The security of the operating system on Oracle Exadata Storage Servers consists of the following:

- Enforcing security policies

- Protecting network access paths to the cells
- Monitoring operating system-level activities

Oracle Exadata System Software includes security features to ensure the operating system and network access to the Oracle Exadata Storage Servers are secure.

2

Security Features of Oracle Exadata Database Machine

Oracle Exadata Database Machine hardware and software are hardened.

The following steps have been done to harden Oracle Exadata Database Machine:

- Trimmed the list of installed packages so that unnecessary packages are not installed on the servers.
- Turned on only essential services on the Oracle Exadata Storage Servers.
- Enabled firewalls (iptables) on the storage servers.
- Enabled auditing of the operating system user.
- Enforced hardened password policies.

Oracle also provides recommended secure configurations for services such as NTP and SSH. In addition, the Oracle Exadata Database Machine architecture provides the following security capabilities to the core components. These security capabilities are most often applied by organizations seeking to deploy a layered security strategy.

- [Using Security-Enhanced Linux](#)
- [Restricting the Binaries Used to Boot the System](#)
Secure Boot supports a chain of trust that goes down to the kernel module level.
- [Using Isolation Policies](#)
Oracle Exadata Database Machine supports multiple isolation levels.
- [Network Access to Oracle Exadata Storage Servers](#)
Oracle Exadata System Software includes the cellwall service, which implements an iptables firewall on each cell.
- [Controlling Access to Data](#)
To protect application data, workloads, and the underlying infrastructure on which it runs, Oracle Exadata Database Machine offers comprehensive yet flexible access control capabilities for both users and administrators.
- [Using Cryptographic Services](#)
- [Monitoring and Auditing of Oracle Exadata Database Machine](#)
Whether for compliance reporting or incident response, monitoring and auditing are critical functions that organizations must use to gain increased visibility into their IT environment.
- [Maintaining Quality of Service](#)
There are many ways that applications can be attacked besides breaching a boundary or subverting an access control policy.
- [Using Oracle ILOM for Secure Management](#)
Collections of security controls and capabilities are necessary to properly secure individual applications and services.
- [Considerations for a Secure Environment](#)
Oracle Exadata Database Machine includes many layered security controls that can be tailored to meet an organization's specific policies and requirements.

- [Understanding the Default Security Settings](#)
Oracle Exadata System Software is installed with many default security settings.

2.1 Using Security-Enhanced Linux

SELinux is a Linux kernel security module that provides a mechanism for supporting access control security policies, including Mandatory Access Control (MAC).

Commencing with Oracle Exadata System Software release 21.2.0, every Linux installation on Oracle Exadata is equipped with a predefined SELinux policy that implements granular permissions for Exadata users, programs, processes, files, and devices. Each server also contains the `host_access_control` utility (`/opt/oracle.cellos/host_access_control`), which provides simple management and monitoring functions for SELinux on Oracle Exadata.

Furthermore, starting with Oracle Exadata System Software release 25.2.0, all new Exadata implementations use SELinux in permissive mode by default.

SELinux can operate in permissive mode or enforcing mode. In permissive mode, SELinux actively monitors the system and records policy violations in the audit logs. However, no actions are blocked in permissive mode. When enforcing mode is enabled, SELinux actively blocks policy violations and records interventions in the audit logs.

Note the following regarding SELinux on Oracle Exadata:

- It is strongly recommended to initially use permissive mode for some time to ensure that the predefined Exadata security policy works in conjunction with all the applications and procedures at your site.

While the system is in permissive mode, you should use standard SELinux commands, such as `aureport`, to monitor policy violations. If thorough system operation in permissive mode yields no policy violations, you can easily transition to enforcing mode.
- The source files for the predefined Exadata security policy are located in `/opt/oracle.SupportTools/selinux`. The file context definitions are contained in `/opt/oracle.SupportTools/selinux/Exadata.fc` and the type enforcement definitions are contained in `/opt/oracle.SupportTools/selinux/Exadata.te`. You can view these files to understand the Exadata security policy.
- Modifying the predefined Exadata security policy is not supported, and any unsupported modifications are overwritten the next time you upgrade the Oracle Exadata System Software.
- You can implement additional security policies at your discretion. For example, you can hand-craft additional policies using standard SELinux operating system commands or use the `audit2allow` command to generate a policy definition that addresses policy violations recorded in the audit logs.

Use the following procedures to manage SELinux on Oracle Exadata:

- [Enabling SELinux](#)
- [Changing the SELinux Mode](#)
- [Disabling SELinux](#)
- [Displaying the SELinux Status](#)

Related Topics

- [SELinux Project](#)

2.1.1 Enabling SELinux

You can enable SELinux on Exadata by using the `host_access_control` utility.

1. Set the SELinux mode.

You can configure SELinux to operate in permissive mode or enforcing mode:

- In permissive mode, SELinux actively monitors the system and records policy violations in the audit logs. However, no actions are blocked in permissive mode.

To specify permissive mode, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --permissive
```

It is strongly recommended to initially enable SELinux in permissive mode for some time to ensure that the predefined Exadata security policy works in conjunction with all of the applications and procedures at your site.

- When enforcing mode is enabled, SELinux actively blocks policy violations and records interventions in the audit logs.

To specify enforcing mode, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --enforcing
```

2. Command the system to perform relabeling.

Under SELinux, each file is associated with a context, or security label, that governs its use. When SELinux is first enabled, you must relabel the system to ensure that the files are associated with the appropriate security context.

The following command instructs the system to perform relabeling.

```
# /opt/oracle.cellos/host_access_control selinux --relabel
```

The command instructs the system to perform relabeling by touching the file located at `/.autorelabel`. However, the relabeling process occurs as part of the next system reboot.

3. Reboot the system.

A system reboot is required to perform relabeling and enable SELinux.

2.1.2 Changing the SELinux Mode

While SELinux is enabled, you can immediately transition between permissive mode and enforcing mode by using the `host_access_control` utility. In this case, no relabeling or reboot is required.

- To change to permissive mode, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --permissive
```

- To change to enforcing mode, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --enforcing
```

2.1.3 Disabling SELinux

You can disable SELinux on Exadata by using the `host_access_control` utility.

1. Set the SELinux mode to disabled.

```
# /opt/oracle.cellos/host_access_control selinux --disabled
```

2. Reboot the system.

A system reboot is required to deactivate SELinux.

2.1.4 Displaying the SELinux Status

You can display the SELinux status by using the `host_access_control` utility. You can display the current status or the configured status:

- The current status indicates the SELinux mode that is currently active on the system.

To display the current status, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --status
```

- The configured status indicates the SELinux mode set by the most recent `host_access_control` command. The configured status differs from the current status after you change the SELinux mode setting but before you reboot the system to activate the new mode.

To display the configured status, use the following command:

```
# /opt/oracle.cellos/host_access_control selinux --config
```

2.2 Restricting the Binaries Used to Boot the System

Secure Boot supports a chain of trust that goes down to the kernel module level.

Secure Boot leverages the UEFI boot architecture to restrict which binaries can boot the system, only allowing boot loaders that carry the cryptographic signature of trusted entities. In other words, anything run in the UEFI boot sequence must be *signed* with a key that the system recognizes as trustworthy. During each system reboot, every component in the boot sequence is verified, preventing malware from hiding embedded code in the boot sequence.

Loadable kernel modules must be signed with a trusted key or they cannot be loaded into the kernel.

The following trusted keys are maintained under the UEFI boot framework:

- Database (DB) - Signature database that contains well-known keys. Only binaries that can be verified against the DB are allowed.
- Forbidden Database (DBX) - Keys that are blocked. Attempting to load an object with a key that matches an entry in the DBX will be denied. This is a list of keys that are bad.
- Machine Owner Key (MOK) - User added keys for kernel modules you want to install.
- Platform Key (PK) - The key installed by the system vendor. This key is installed by the vendor and is in the ILOM firmware. This key is not accessible from the host.

- Key Exchange Key (KEK) - The key required to update the signature database.

The user must have access to the system console to add keys, modify keys, or enable and disable Secure Boot through the UEFI configuration menu. The default boot loader on most UEFI-enabled servers running Linux is grub2. With Secure Boot enabled, an additional *shim* boot loader is needed. When booting in Secure Boot mode, the shimloader is called first because it contains a trusted signature. The shimloader then loads grub2, which then loads the OS kernel, which is also signed.

The minimum system requirement for Secure Boot on Exadata system hardware is Oracle Exadata X7-2 running Oracle Exadata System Software release 19.1.0 or later. During Exadata system deployment, Secure Boot is enabled by default on Exadata storage servers, bare metal database servers, and KVM hosts that meet this requirement.

Oracle Exadata System Software release 24.1.0 extends Secure Boot to Oracle Linux KVM guests.

Secure Boot is not supported on Xen-based Oracle VM Servers (Dom0) and Oracle VM guests (DomU).

- [Enabling Secure Boot on Exadata System Hardware](#)
Secure Boot is enabled by default on supported Exadata storage servers, bare metal database servers, and KVM hosts.
- [Enabling KVM Guest Secure Boot](#)
Oracle Exadata System Software release 24.1.0 extends Secure Boot to Oracle Linux KVM guests.
- [Managing Keys and Certificates Used with Secure Boot](#)
- [Troubleshooting Secure Boot](#)
You might encounter the following problems when Secure Boot is enabled.

2.2.1 Enabling Secure Boot on Exadata System Hardware

Secure Boot is enabled by default on supported Exadata storage servers, bare metal database servers, and KVM hosts.

Note

Secure Boot is not supported on Xen-based Oracle VM Servers (Dom0) and Oracle VM guests (DomU).

Oracle recommends that you leave the default Secure Boot option setting.

You can enable (or disable) Secure Boot by pressing F12 during the boot process, navigating to the EFI boot menu and changing the Secure Boot setting.

To verify the status of the Secure Boot option, use the following command:

```
# mokutil --sb-state
SecureBoot enabled
```


2.2.2 Enabling KVM Guest Secure Boot

Oracle Exadata System Software release 24.1.0 extends Secure Boot to Oracle Linux KVM guests.

KVM Guest Secure Boot leverages the UEFI boot framework in Oracle Linux KVM to restrict which binaries can boot the KVM guest. Exadata support for UEFI in KVM is introduced in Oracle Exadata System Software release 24.1.0. Consequently, KVM Guest Secure Boot is available only on new KVM guests starting with Oracle Exadata System Software release 24.1.0.

To use KVM Guest Secure Boot, you must enable UEFI booting in the system image underpinning the KVM guest. For example, if the KVM host is using a fresh system image (which has not been updated using `patchmgr`), you can run the following command on the KVM host to enable UEFI booting:

```
# vm_maker --make-base-image --uefi-boot
```

After command completion, all new KVM guests will automatically use the UEFI boot framework in Oracle Linux KVM and have KVM Guest Secure Boot enabled. It is no longer possible to create a new guest without UEFI.

If the KVM host has been updated using `patchmgr`, then you must download and install a UEFI-enabled system image.

If required, you can disable KVM Guest Secure Boot on UEFI boot-enabled KVM guests in the following ways:

- To disable KVM Guest Secure Boot on an existing UEFI boot-enabled KVM guest, run the following command on the KVM host:

```
# vm_maker --secure-boot disable --domain domain-name
```

In the command, *domain-name* specifies the name of the existing KVM guest.

After the command, KVM Guest Secure Boot will be disabled when the guest is rebooted.

- To create a new KVM guest with KVM Guest Secure Boot disabled, run the `vm_maker --start-domain` command on the KVM host and include the `--secure-boot disable` option. For example:

```
# vm_maker --start-domain XML-config-file --secure-boot disable
```

In the command, *XML-config-file* specifies the name of the XML configuration file for the new guest.

- To disable KVM Guest Secure Boot on all new UEFI boot-enabled KVM guests, run the following command on the KVM host:

```
# vm_maker --secure-boot disable --system
```

After running this command, all new UEFI boot-enabled KVM guests are created with KVM Guest Secure Boot disabled. This command does not affect existing guests.

To verify the boot configuration for a KVM guest, run the following command on the KVM host:

```
# parted -s /EXAVMIMAGES/GuestImages/guest-name/System.img p
```

In the command, *guest-name* specifies the name of the KVM guest.

Examine the command output for one of the following flags, which indicates the boot configuration for the KVM guest:

- `efi`: Indicates that the guest uses the UEFI boot framework in Oracle Linux KVM.
- `bios_grub`: Indicates that the guest uses the non-UEFI virtual bios.

To verify the status of KVM Guest Secure Boot, run the following command in the KVM guest:

```
# mokutil --sb-state
SecureBoot enabled
```

2.2.3 Managing Keys and Certificates Used with Secure Boot

You can use the Machine Owner Keys (MOK) utility (`mokutil`) to manage the keys and certificates used with Secure Boot.

The certificates are signed by DigiCert. By default, a certificate is valid for one year from the date of signing. Even though a certificate may expire, the validation is based on the date on which the grub and kernel were signed and if the certificate was valid at that time.

To renew the certificates, you update the kernel, grub, and ILOM on the secured servers with a new, signed version.

- To query the existing keys, run the `mokutil` command as the root user.

For example:

```
# mokutil --list-enrolled
[key 1]
SHA1 Fingerprint:
5f:f4:35:5a:49:ec:8d:f1:56:d1:ee:9b:ac:f6:19:54:08:77:d3:59
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number:
      21:b3:c1:01:19:dc:af:44:43:15:8b:0f:33:6b:18:be
    Signature Algorithm: sha256WithRSAEncryption
    Issuer: C=US, O=Symantec Corporation, OU=Symantec Trust Network,
    CN=Symantec Class 3 Extended Validation Code Signing CA - G2
  Validity
    Not Before: Jun 30 00:00:00 2020 GMT
    Not After : Jul  1 23:59:59 2021 GMT
  Subject: jurisdictionC=US/jurisdictionST=Delaware/
  businessCategory=Private Organization/serialNumber=2101822, C=US/
  postalCode=94065, ST=California, L=Redwood City/street=500 Oracle Parkway,
  O=Oracle America Inc., OU=Winqual, CN=Oracle America Inc.
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    Public-Key: (2048 bit)
    Modulus:
```

```

00:8d:3e:e0:3b:35:99:fb:11:c0:2a:12:ac:07:40:
f7:90:d4:d3:62:5e:85:2d:ea:94:af:5f:26:33:98:
c8:03:33:0e:30:5e:4d:44:ca:fa:1a:3a:49:88:64:
89:16:5c:39:f3:35:86:ed:25:eb:0f:ca:fa:2c:3d:
d6:23:2a:b3:1e:62:fb:45:88:1a:05:be:95:d6:6a:
d9:c5:f2:81:7a:cc:63:71:3c:37:a0:23:1c:eb:20:
1a:3d:13:89:6a:9e:47:a0:eb:ca:64:21:3f:7a:f4:
e6:09:bf:47:63:c8:b3:6b:a5:c6:1b:de:f6:06:12:
56:eb:ab:24:00:01:c9:80:db:be:66:49:64:ac:c8:
ce:1e:da:7a:c1:42:21:85:f9:67:81:a4:f0:6d:14:
01:9b:45:1e:9f:08:e5:18:b7:c5:34:e5:55:e2:11:
dc:fe:0c:36:32:f4:bb:cb:34:00:37:b2:41:05:5f:
0a:69:68:55:cb:4e:ec:ca:cc:1b:67:dc:05:f1:98:
95:c4:14:35:41:01:fe:f5:bd:63:1a:8d:cc:8a:1f:
b6:87:ac:02:ea:e2:2e:29:d6:11:b9:bc:aa:d6:44:
3e:32:3c:a9:12:a4:aa:09:ec:6e:ba:99:08:58:36:
6b:ef:40:c5:3e:47:36:93:53:f1:c9:f2:79:f2:53:
c9:9b
Exponent: 65537 (0x10001)
X509v3 extensions:
X509v3 Basic Constraints:
CA:FALSE
X509v3 Key Usage: critical
Digital Signature
X509v3 Extended Key Usage:
Code Signing
X509v3 Certificate Policies:
Policy: 2.23.140.1.3
CPS: https://d.symcb.com/cps
User Notice:
Explicit Text: https://d.symcb.com/rpa

X509v3 Subject Key Identifier:
BC:59:71:95:4C:74:9D:3D:30:98:52:EF:0F:3C:23:6F:A4:98:E8:F6
X509v3 Authority Key Identifier:

keyid:16:66:DE:4A:34:E3:50:A7:11:86:03:B1:6C:A9:C6:AC:CD:59:6E:9B

X509v3 CRL Distribution Points:

Full Name:
URI:http://sw.symcb.com/sw.crl

Authority Information Access:
OCSP - URI:http://sw.symcd.com
CA Issuers - URI:http://sw.symcb.com/sw.crt

Signature Algorithm: sha256WithRSAEncryption
38:4d:10:69:07:db:7c:ce:18:2b:1e:c5:89:1c:71:a9:b0:07:
19:43:2d:a0:88:c5:f5:bf:82:a9:4b:f9:45:fa:2c:7c:00:cb:
be:24:b0:a8:98:7d:f5:a3:c4:42:52:f4:75:fd:22:c5:0c:2e:
a2:13:7f:b9:24:79:04:d5:ea:0e:1a:e6:e8:4c:61:48:65:5b:
c7:30:81:90:fd:17:d5:39:d4:70:00:00:b8:c5:80:03:da:88:
e0:f1:39:aa:d9:1d:ef:2f:bf:c3:06:18:2a:1b:1f:ce:30:a2:
bb:dd:d0:46:0e:d5:e1:22:0c:a0:cc:df:00:fe:0a:99:d5:cc:
16:76:4b:ab:dc:bb:80:4b:0e:1b:f5:5e:04:22:3e:a9:d0:70:

```

```
56:87:9b:c1:2f:95:cf:36:34:e7:c7:2e:0c:56:f3:24:fa:7d:
f7:25:54:50:34:f6:e5:30:76:8b:fd:65:25:19:8a:54:f9:f1:
93:24:ad:22:25:4a:e0:a2:63:b6:d7:d1:82:4e:5a:fc:34:52:
b4:9e:7d:1a:e2:b7:a1:92:13:0f:9d:7b:ae:42:6f:64:a2:02:
47:c7:f9:11:12:e4:82:b9:f7:ed:ce:14:ac:c2:b4:e3:cc:c4:
ef:f8:9f:78:23:91:89:b0:37:24:f1:c6:61:0c:2e:cf:af:29:
e5:68:70:4d
```

- [Adding Keys for Secure Boot Using mokutil](#)
- [Removing Keys for Secure Boot Using mokutil](#)

2.2.3.1 Adding Keys for Secure Boot Using mokutil

On bare metal database servers and KVM hosts, you can add new keys for use with Secure Boot using the Machine Owner Keys (MOK) utility (mokutil).

You must run the mokutil command as the root user.

You can run mokutil --help to view additional details about the mokutil command.

1. Create a DER-formatted X509 certificate file for the key you want to add.

For example, you could use the following command sequence to create a new key and the associated certificate.

```
# openssl genpkey -algorithm EC -pkeyopt ec_paramgen_curve:P-256 >
my_private_key.pem
# openssl req -x509 -key my_private_key.pem -subj /CN=client.example.com >
my_certificate.pem
# openssl x509 -in my_certificate.pem -inform PEM -out my_certificate.der -
outform DER
```

2. Check to see if the key is already active.

```
# mokutil --test-key my_certificate.der
```

3. If the key is not currently active, then import the key certificate.

```
# mokutil --import my_certificate.der
```

When prompted, specify a MOK management password. Ensure that you remember the password because it is required to complete MOK enrollment after you reboot the system.

Note

- If required, you can use the `mokutil` command to import multiple certificates at once using the same password by providing a list of DER-formatted X509 certificate files. For example:

```
# mokutil --import my_certificate.der my_certificate2.der  
my_certificate3.der
```

If you import multiple certificates, the keys are later identified in the MOK management interfaces as Key 0, Key 1, Key 2, and so on.

- If you discover any problem with the certificate(s) you imported, you can revoke the imported certificates at any time before the keys are enrolled in the UEFI Secure Boot key database by running:

```
# mokutil --revoke-import
```

4. Record details about the key and associated certificate for identification purposes in the following steps.

Record the output from the following command:

```
# mokutil --list-new
```

5. Connect to the system console and reboot the system.

On a physical server, you can view the system console by connecting to the Integrated Lights Out Manager (ILOM) subsystem and then running the following command from the ILOM prompt:

```
-> start -script /SP/console
```

After connecting to the system console, you can reboot the system by logging in as the root user and running:

```
# shutdown -r now
```

6. Observe the system console and press any key to perform MOK management when prompted.
7. In the **Perform MOK management** screen, select the **Enroll MOK** menu option.
8. Select the **View key 0** menu option and verify that the key details match the new key that you want to enroll.

If you imported multiple keys, verify all of the key details using **View key 1**, **View key 2**, and so on.

9. Enroll the new key(s).
 - a. Return to the **Enroll MOK** screen.
 - b. Select the **Continue** option.
 - c. Select **Yes** in the **Enroll the key(s)?** dialog.

- d. When prompted, supply the MOK management password you specified earlier when importing the key certificate(s).

The key(s) are now enrolled in the UEFI Secure Boot key database.

10. In the **Perform MOK management** screen, select the **Reboot** menu option.
11. After the system reboots, confirm that the new key is enrolled by examining the output from the following command:

```
# mokutil --list-enrolled
```

2.2.3.2 Removing Keys for Secure Boot Using mokutil

On bare metal database servers and KVM hosts, you can remove keys associated with Secure Boot using the Machine Owner Keys (MOK) utility (mokutil).

You must run the mokutil command as the root user.

You can run mokutil --help to view additional details about the mokutil command.

Warning

Ensure that you do not alter or remove the Oracle-supplied keys and certificates included with the system. To identify an Oracle-supplied key, examine the corresponding certificate and look for Oracle (O=Oracle America Inc.) in the subject information. For example:

```
# mokutil --list-enrolled
[key 1]
SHA1 Fingerprint:
5f:f4:35:5a:49:ec:8d:f1:56:d1:ee:9b:ac:f6:19:54:08:77:d3:59
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number:
      21:b3:c1:01:19:dc:af:44:43:15:8b:0f:33:6b:18:be
    Signature Algorithm: sha256WithRSAEncryption
    Issuer: C=US, O=Symantec Corporation, OU=Symantec Trust
  Network,
    CN=Symantec Class 3 Extended Validation Code Signing CA -
  G2
  Validity
    Not Before: Jun 30 00:00:00 2020 GMT
    Not After : Jul  1 23:59:59 2021 GMT
  Subject: jurisdictionC=US/jurisdictionST=Delaware/
businessCategory=Private Organization/serialNumber=2101822,
C=US/postalCode=94065, ST=California, L=Redwood City/
street=500 Oracle Parkway,
O=Oracle America Inc., OU=Winqual, CN=Oracle America Inc.
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    Public-Key: (2048 bit)
...

```

1. Identify the key that you want to remove by examining the output from the following command:

```
# mokutil --list-enrolled
```

2. Export the key certificates using the following command:

```
# mokutil --export
```

The command outputs a DER-formatted X509 certificate file for every enrolled key. Examine the file(s) and locate the certificate(s) you want to remove.

3. Delete the certificate(s) associated with the key(s) that you want to remove.

For example:

```
# mokutil --delete my_certificate.der
```

When prompted, specify a MOK management password. Ensure that you remember the password because it is required to complete MOK deletion after you reboot the system.

Note

- If required, you can use the `mokutil` command to delete multiple certificates at once by providing a list of DER-formatted X509 certificate files. For example:

```
# mokutil --delete my_certificate.der my_certificate2.der  
my_certificate3.der
```

If you delete multiple certificates, the keys are later identified in the MOK management interfaces as Key 0, Key 1, Key 2, and so on.

- If you discover any issues with the certificate(s) you marked for deletion, you can revoke the deletion at any time before the keys are removed from the UEFI Secure Boot key database by running:

```
# mokutil --revoke-delete
```

4. Record details about the key and associated certificate for identification purposes in the following steps.

Record the output from the following command:

```
# mokutil --list-delete
```

5. Connect to the system console and reboot the system.

On a physical server, you can view the system console by connecting to the Integrated Lights Out Manager (ILOM) subsystem and then running the following command from the ILOM prompt:

```
-> start -script /SP/console
```

After connecting to the system console, you can reboot the system by logging in as the root user and running:

```
# shutdown -r now
```

6. Observe the system console and press any key to perform MOK management when prompted.
7. In the **Perform MOK management** screen, select the **Delete MOK** menu option.
8. Select the **View key 0** menu option and verify that the key details match the key that you want to delete.

If you are deleting multiple keys, verify all of the key details using **View key 1**, **View key 2**, and so on.

9. Delete the key(s).
 - a. Return to the **Delete MOK** screen.
 - b. Select the **Continue** option.
 - c. Select **Yes** in the **Delete the key(s)?** dialog.
 - d. When prompted, supply the MOK management password you specified earlier when deleting the key certificate(s).

The key(s) are now removed from the UEFI Secure Boot key database.

10. In the **Perform MOK management** screen, select the **Reboot** menu option.
11. After the system reboots, confirm that the key is deleted by examining the output from the following command:

```
# mokutil --list-enrolled
```

2.2.4 Troubleshooting Secure Boot

You might encounter the following problems when Secure Boot is enabled.

- **Secure boot violation: Invalid signature detected. Check Secure Boot Policy in Setup.** — The grub loader has an invalid signature.

Note

Even though a certificate may expire, the validation is based on the date on which the grub & kernel were signed and if the certificate was valid at that time.

- **error: file has invalid signature. error: You need to load the kernel first.** — The grub loader is signed, but the kernel is unsigned.
- **ERROR: Verification failed: (15) Access Denied. Failed to load image: Access Denied.start_image() returned Access Denied** — The ISO image being loaded to image the server is not signed.

2.3 Using Isolation Policies

Oracle Exadata Database Machine supports multiple isolation levels.

Organizations wanting to consolidate IT infrastructure, implement shared service architectures, and deliver secure multitenant services should isolate services, users, data, communications, and storage. Oracle Exadata Database Machine provides organizations the flexibility to implement the isolation policies and strategies based on their needs. The following are the secure isolation levels of Oracle Exadata Database Machine:

- [Isolating Network Traffic](#)
Oracle Exadata uses multiple networks to segregate network traffic.
- [Isolating Databases](#)
Use operating system controls and database features to enable database isolation.
- [Isolating Storage](#)
Oracle Exadata storage is isolated from the rest of the architecture through the use of a RDMA Network Fabric (InfiniBand or RDMA over Converged Ethernet (RoCE)).

2.3.1 Isolating Network Traffic

Oracle Exadata uses multiple networks to segregate network traffic.

At the physical network level, client access is isolated from device management and inter-device communication. Client and management network traffic are isolated on separate networks. Client access is provided over a bonded Ethernet network interface that ensures reliable, high-speed access to services running on the system. Management access is provided over a physically separate Ethernet network interface. This provides a separation between operational and management networks.

Organizations may choose to further segregate network traffic over the client access network by configuring virtual LANs (VLANs). VLANs segregate network traffic based on their requirements. Oracle recommends the use of encrypted protocols over VLANs to assure the confidentiality and integrity of communications.

Inter-device communication is provided by a RDMA Network Fabric (InfiniBand or RDMA over Converged Ethernet (RoCE)). The RDMA Network Fabric is a high-performance, low-latency backplane for communication between Oracle Exadata Storage Servers and database servers. By default, Oracle Exadata Storage Servers include a configured software firewall. The database servers can also be configured with a software firewall.

Note

Partitioning the InfiniBand private network does not protect an InfiniBand fabric. Partitioning only offers InfiniBand traffic isolation between machines.

2.3.2 Isolating Databases

Use operating system controls and database features to enable database isolation.

Physical separation by dedicating an entire environment to a single application or database is one of the best isolation methods. However, it is expensive. A more cost-effective isolation strategy uses multiple databases within the same operating system image. Multiple database isolation is achieved through a combination of database and operating system-level controls, such as dedicated credentials for users, groups, and resource controls.

All Oracle Database security options are available for Oracle Exadata Database Machine. Organizations wanting finer-grained database isolation can use software such as Oracle Database Vault, Oracle Virtual Private Database, and Oracle Label Security.

Oracle Database Vault includes a mandatory access control model to enforce isolation using logical realms within a single database. Logical realms form a protective boundary around existing application tables by blocking administrative accounts from having ad-hoc access to application data. Oracle Database Vault command rules enable policy-based controls that limit who, when, where, and how the database and application data is accessed. This creates a trusted path to application data. Oracle Database Vault can also be employed to restrict access based upon time, source IP address, and other criteria.

Oracle Virtual Private Database enables the creation of policies that enforce fine-grained access to database tables and views at the row and column levels. Oracle Virtual Private Database provides security portability because the policies are associated with database objects, and are automatically applied no matter how the data is accessed. Oracle Virtual Private Database can be used for fine-grained isolation within the database.

Oracle Label Security is used to classify data, and mediate access to that data based upon its classification. Organizations define classification strategies, such as hierarchical or disjoint, that best support their needs. This capability allows information stored at different classification levels to be isolated at the row level within a single tablespace.

2.3.3 Isolating Storage

Oracle Exadata storage is isolated from the rest of the architecture through the use of a RDMA Network Fabric (InfiniBand or RDMA over Converged Ethernet (RoCE)).

The storage managed by Oracle Exadata Storage Servers can be subdivided using Oracle Automatic Storage Management (Oracle ASM) to create individual disk groups. Each disk group can have its own security policies.

2.4 Network Access to Oracle Exadata Storage Servers

Oracle Exadata System Software includes the cellwall service, which implements an iptables firewall on each cell.

Additionally, the SSH server is configured to respond to connection requests only on the management network (NET0) and the RDMA Network Fabric.

To review the firewall rules, run the following command as the root user:

```
iptables --list
```

Note

There is no firewall automatically configured for the database servers. You can implement an iptables firewall on each database server to meet your network security requirements.

2.5 Controlling Access to Data

To protect application data, workloads, and the underlying infrastructure on which it runs, Oracle Exadata Database Machine offers comprehensive yet flexible access control capabilities for both users and administrators.

The control capabilities include network access, database access, and storage access.

- [Controlling Network Access](#)
Beyond simple network-level isolation, fine-grained access control policies can be instituted at the device level.
- [Controlling Database Access](#)
Separation of duties is critical at every layer of the architecture to reduce the risk of collusive behavior, and prevent inadvertent errors.
- [Controlling Storage Access](#)
Oracle Exadata System Software supports the access control modes of open security, Oracle ASM-scoped security, and database-scoped security.

2.5.1 Controlling Network Access

Beyond simple network-level isolation, fine-grained access control policies can be instituted at the device level.

All components in Oracle Exadata Database Machine include the ability to limit network access to services either using architectural methods, such as network isolation, or using packet filtering and access control lists to limit communication to, from, and between components and services.

2.5.2 Controlling Database Access

Separation of duties is critical at every layer of the architecture to reduce the risk of collusive behavior, and prevent inadvertent errors.

For example, use different operating system accounts to ensure role separation for database and storage administrators, including administrators supporting Oracle Automatic Storage Management (Oracle ASM). Within Oracle Database, users can be assigned specific privileges and roles to ensure that users have access to only those data objects that they are authorized to access. Data cannot be shared unless it is explicitly permitted.

In addition to password-based authentication, Oracle Database also supports public key credentials, RADIUS, and Kerberos. Using Oracle Enterprise User Security, the database can be integrated with existing LDAP repositories for authentication and authorization. These capabilities provide higher assurance of the identity of users connecting to the database.

Oracle Database Vault can be used to manage administrative and privileged user access, controlling how, when and where application data can be accessed. Oracle Database Vault protects against misuse of stolen login credentials, application bypass, and unauthorized changes to applications and data, including attempts to make copies of application data. Oracle Database Vault is transparent to most applications, and day-to-day tasks. It supports multi-factor authorization policies, allowing for secure enforcement of policy without disrupting business operations.

Oracle Database Vault can enforce separation of duties to ensure that account management, security administration, resource management, and other functions are granted only to those users authorized to have those privileges.

2.5.3 Controlling Storage Access

Oracle Exadata System Software supports the access control modes of open security, Oracle ASM-scoped security, and database-scoped security.

- Open security allows any database to access any of the grid disks.

- Oracle ASM-scoped security allows multiple databases assigned to one or more Oracle ASM clusters to share specific grid disks.

In addition to its overall access control mode, Oracle ASM supports access controls at the disk group and file level to ensure that access to content stored on disk is only available to authorized users.

Note

- The `/etc/oracle/cell/network-config/cellkey.ora` file needs to be readable only by the software installation owner of Oracle Grid Infrastructure with its specific unique group, such as `asmadmin`.
- Use the `kfod` utility in the Oracle Grid Infrastructure home to troubleshoot or verify which disks are accessible for your cluster.

- Database-scoped security, the most fine-grained level of access control, ensures that only specific databases are able to access specific grid disks.

Database-scoped security works on a container level. This means that grid disks must be made available to the `DB_UNIQUE_NAME` of the container database (CDB) or non-CDB. Because of this, it is not possible to have database-scoped security per pluggable database (PDB).

Note

You should only set up database-scoped security after configuring and testing Oracle ASM-scoped security.

By default, SSH is enabled on storage servers. If required, you can "lock" the storage servers to block SSH access. You can still perform operations on the storage server using `exaccli`, which runs on compute nodes and communicates using HTTPS and REST APIs to a web service running on the cell. At a high-level, this is accomplished by creating users and roles in `CellCLI` and then disabling `remoteLogin`.

Related Topics

- Disabling SSH on Storage Servers
- [Configuring Data Security for Exadata Storage Servers](#)
Data security for Oracle Exadata Storage Servers is implemented by controlling which Oracle Automatic Storage Management (Oracle ASM) clusters and Oracle Database clients can access specific grid disks on storage cells.

2.6 Using Cryptographic Services

The requirement to protect and validate information at rest, in transit, and in use often employs cryptographic services. From encryption and decryption to digital fingerprint and certificate validation, cryptography is one of the most-widely deployed security controls in IT organizations.

Whenever possible, Oracle Exadata Database Machine makes use of hardware-based cryptographic engines on processor chips provided by Intel AES-NI and Oracle SPARC. Using hardware for cryptographic operations provides significant performance improvement over performing the operations in software. Both engines provide the ability to perform

cryptographic operations in hardware, and both are leveraged by Oracle software on the database and storage servers.

Network cryptographic services protect the confidentiality and integrity of communications by using a cryptographically-secure protocol. For example, Secure Shell (SSH) access provides secure administrative access to systems and Integrated Lights Out Manager (ILOM). SSL/TLS can enable secure communications between applications and other services.

Database cryptographic services are available by using transparent data encryption (TDE), a component of Oracle Advanced Security. TDE supports the encryption of entire tablespaces and individual columns within a table. TDE is woven into the architecture of Oracle Database, with data encryption being maintained in temporary tablespaces and redo logs. Data encryption is also maintained in database backups, which protects the data regardless of the storage device in use. On Exadata, TDE-encrypted data also remains encrypted when it resides in Exadata Smart Flash Cache or in Exadata RDMA Memory Cache.

In addition, Oracle Advanced Security can encrypt Oracle Net Services and JDBC traffic using either native encryption or SSL to protect information while in transit over a network. Both administrative and application connections can be protected to ensure that data in transit is protected. The SSL implementation supports the standard set of authentication methods including anonymous (Diffie-Hellman), server-only authentication using X.509 certificates, and mutual (client-server) authentication with X.509.

2.7 Monitoring and Auditing of Oracle Exadata Database Machine

Whether for compliance reporting or incident response, monitoring and auditing are critical functions that organizations must use to gain increased visibility into their IT environment.

The degree to which monitoring and auditing is employed is often based upon the risk or criticality of the environment. Oracle Exadata has been designed to offer comprehensive monitoring and auditing functionality at the server, network, database, and storage layers ensuring that information can be made available to organizations in support of their audit and compliance requirements.

- [Monitoring and Auditing of Oracle Exadata Database Machine](#)
AIDE is a security feature that reports any malicious or unplanned change to the system.
- [Monitoring and Auditing Oracle Database Activity](#)
Oracle Database support of fine-grained auditing allows organizations to establish policies that selectively determine when audit records are generated.
- [Operating System Activity Monitoring on Oracle Exadata Servers](#)
Each Exadata server is configured with `auditd` to audit system-level activity.

2.7.1 Monitoring and Auditing of Oracle Exadata Database Machine

AIDE is a security feature that reports any malicious or unplanned change to the system.

Oracle Exadata System Software release 19.1.0 adds support for Advanced Intrusion Detection Environment (AIDE) to help guard against unauthorized access to the files on your Exadata system. AIDE creates a database of files on the system, and then uses that database to ensure file integrity and to detect system intrusions. To learn more about AIDE see https://en.wikipedia.org/wiki/Advanced_Intrusion_Detection_Environment

On Oracle Exadata, a Management Server (MS) alert is generated when AIDE identifies an unplanned change to the system (files or directories).

For non-production systems, or systems that are temporarily considered NON-PRODUCTION, where software installation or configuration is occurring, AIDE could generate a large number of alerts with false positives. While a system is in NON-PRODUCTION mode, the recommendation is to temporarily disable AIDE on each of the compute nodes by running the command `/opt/oracle.SupportTools/exadataAIDE -disable`.

For systems returned to PRODUCTION, after locking down software installation, one of the last steps should be the final update of the AIDE database, by executing following commands:

- `/opt/oracle.SupportTools/exadataAIDE -enable` — If AIDE was previously disabled
- `/opt/oracle.SupportTools/exadataAIDE -u` — To generate a new AIDE database baseline

If you must modify the configuration on any of the PRODUCTION database servers, then run an update of the AIDE database after the change, by executing the command `/opt/oracle.SupportTools/exadataAIDE -u`.

Note

Updating the AIDE database clears all open AIDE MS alerts.

2.7.2 Monitoring and Auditing Oracle Database Activity

Oracle Database support of fine-grained auditing allows organizations to establish policies that selectively determine when audit records are generated.

Establishing policies helps organizations focus on other database activities, and reduce the overhead that is often associated with audit activities.

Oracle Audit Vault centralizes the management of database audit settings and automates the consolidation of audit data into a secure repository. Oracle Audit Vault includes built-in reporting to monitor a wide range of activities including privileged user activity and changes to database structures. The reports generated by Oracle Audit Vault enable visibility into various application and administrative database activities, and provide detailed information to support accountability of actions.

Oracle Audit Vault enables the proactive detection and alerting of activities that may be indicative of unauthorized access attempts or abuse of system privileges. These alerts can include both system and user-defined events and conditions, such as the creation of privileged user accounts or the modification of tables containing sensitive information.

Oracle Database Firewall Remote Monitor can provide real-time database security monitoring. Oracle Database Firewall Remote Monitor queries database connections to detect malicious traffic, such as application bypass, unauthorized activity, SQL injection and other threats. Using an accurate SQL grammar-based approach, Oracle Database Firewall helps organizations quickly identify suspicious database activity.

2.7.3 Operating System Activity Monitoring on Oracle Exadata Servers

Each Exadata server is configured with `auditd` to audit system-level activity.

Starting with Oracle Exadata System Software release 19.1.0 and Oracle Linux 7, the audit rules for Oracle Exadata are stored in the `/etc/audit/rules.d/01-exadata_audit.rules` file.

Additional custom rules should be placed in separate audit rule files in the `/etc/audit/rules.d` directory, such as `/etc/audit/rules.d/20-customer_audit.rules`. These custom audit rule files are preserved across updates to Oracle Exadata System Software.

When the `auditd` service starts, it runs the `augenrules` utility. This utility merges all component audit rules files found in the audit rules directory, `/etc/audit/rules.d`, and places the merged results in the `/etc/audit/audit.rules` file. Component audit rule files, must end in `.rules` to be processed by `augenrules`. All other files in the `/etc/audit/rules.d` directory are ignored. The files are concatenated in order, based on their natural sort order and stripped of empty lines and comment (`#`) lines.

As in previous releases of Oracle Exadata System Software, the audit rules are immutable. A reboot is needed to effect changes to audit rules.

2.8 Maintaining Quality of Service

There are many ways that applications can be attacked besides breaching a boundary or subverting an access control policy.

Oracle Exadata provides a number of capabilities to help detect and prevent resource exhaustion attacks, denial of service attacks, and accidental or intentional faults that can impact the availability of services and data.

Oracle Exadata System Software includes I/O Resource Management (IORM) to manage interdatabase and intradatabase I/O resources. IORM allows different databases with different performance requirements to share a common Oracle Exadata Storage Server pool. Multiple workloads within the same database can have their own resource policies. This flexible architecture allows organizations to ensure that critical workloads and databases share I/O resources when operating on a consolidated architecture.

Oracle Database includes tools to enable multiple databases to operate under the same operating system. Oracle Database Resource Manager, and instance caging support the ability to dynamically control access to CPU resources using fine-grained methods. Oracle Database Resource Manager can control the degree of parallelism, the number of active sessions, and other shared resources to protect one database from monopolizing resources needed in shared database architectures.

Oracle Database Quality of Service Management (Oracle Database QoS Management) is an automated, policy-based solution that monitors the workload requests of an entire system. Oracle Database QoS Management correlates accurate run-time performance and resource metrics, analyzes the data to identify bottlenecks, and produces recommended resource adjustments to maintain performance objectives under dynamic load conditions.

2.9 Using Oracle ILOM for Secure Management

Collections of security controls and capabilities are necessary to properly secure individual applications and services.

It is equally important to have comprehensive management capabilities to sustain the security of the deployed services and systems. Oracle Exadata Database Machine uses the security management capabilities of ILOM.

ILOM is a service processor embedded in many Oracle Exadata Database Machine components. It is used to perform out-of-band management activities, such as the following:

- Provide secure access to perform secure lights-out management of the database and storage servers. Access includes web-based access protected by SSL, command-line access using Secure Shell (SSH), and protocol access using TLS and SNMPv3.
- Separate duty requirements using a role-based access control model. Individual users are assigned to specific roles that limit the functions that can be performed.
- Provide an audit record of all logins and configuration changes. Each audit log entry lists the user performing the action, and a timestamp. This allows organizations to detect unauthorized activity or changes, and attribute those actions back to specific users.

Related Topics

- [How to disable IPMI 2.0 on Exadata nodes \(My Oracle Support Doc ID KB515361\)](#)

2.10 Considerations for a Secure Environment

Oracle Exadata Database Machine includes many layered security controls that can be tailored to meet an organization's specific policies and requirements.

Organizations must evaluate how to best utilize these capabilities and integrate them into their existing IT security architecture. Effective IT security must consider the people, processes, and technology in order to provide solid risk management and governance practices. Practices and policies should be designed and reviewed during the planning, installation, and deployment stages of Oracle Exadata Database Machine.

While many of the features integrated into Oracle Exadata Database Machine are configured by default for secure deployment, organizations have their own security configuration standards. It is important to review Oracle security information before testing any security setting changes to Oracle Exadata Database Machine components. In particular, it is important to identify where existing standards can be improved, and where support issues may limit what changes can be made to a given component.

Note

To minimize the attack surface, Oracle Exadata Storage Servers do not support customization outside of their management interfaces. No custom users are permitted on the storage servers. The servers have been optimized and hardened for their specific purpose.

- [Identity and Access Management Considerations](#)
A unified approach should be used when integrating Oracle Exadata Database Machine components and deployed services with your organization's existing identity and access management architecture.
- [Network Security Considerations](#)
Before Oracle Exadata arrives at your location, network security considerations should be discussed.

2.10.1 Identity and Access Management Considerations

A unified approach should be used when integrating Oracle Exadata Database Machine components and deployed services with your organization's existing identity and access management architecture.

Oracle Database supports many open and standard protocols that allow it to be integrated with existing identity and access management deployments. To ensure application availability,

unified identity and access management systems must be available, or the availability of Oracle Exadata Database Machine may be compromised.

Before Oracle Exadata Database Machine arrives, the following security considerations should be discussed. These considerations are based on Oracle best practices for Oracle Exadata Database Machine.

- The ability to directly log in to common operating system accounts such as root, grid and oracle should be disabled. Individual user accounts should be created for each administrator. After logging in with their individual account, the administrator can use sudo to run privileged commands, when required.
- The use of host-based intrusion detection and prevention systems for increased visibility within Oracle Exadata Database Machine. By using the fine-grained auditing capabilities of Oracle Database, host-based systems have a greater likelihood of detecting inappropriate actions and unauthorized activity.
- The use of centralized audit and log repositories to aggregate the security-relevant information for improved correlation, analysis, and reporting. Oracle Exadata Storage Servers support this through the CELL attribute syslogConf. The database servers support centralized logging using the typical system configuration methods.
- The use of encryption features such as transparent data encryption (TDE), Oracle Recovery Manager (RMAN) encryption for backups.

The security of the data and system is diminished by user access and password security. Oracle recommends the following guidelines to maximize your user security:

- Create separate software owner accounts for Oracle Grid Infrastructure and Oracle Database software installations. These accounts should be used when deploying Oracle Exadata Database Machine. A separate software owner for Oracle Grid Infrastructure and Oracle Database software installations is required for implementing DB-scoped security.
- Implement a user password policy that enforces password complexity beyond the minimum requirements.
- Implement password aging and account locking. Starting with Oracle Exadata System Software release 19.1.0 you can use DBSERVER and CELL attributes to configure the following account security features:
 - A user's password expires after a specified number of days. The default user password expiration time is 0. 0 means passwords will not expire.
 - A user gets a warning message when logging in for a specified number of days before their password expires. The default user account password expiration warning time is 7 days.
 - The user is prompted to change their password when logging within a specified number of days after their password expires. If the remotePwdChangeAllowed attribute on the server indicates that a service request is not required to change the password, then the user can change the password immediately. Otherwise, the user must connect the server administrator to have their password changed.
 - A user account is locked a specified number of days after the password expires. The default user account lock time is 7 days. After the account is locked, the user must contact the server administrator to have the account unlocked.

2.10.2 Network Security Considerations

Before Oracle Exadata arrives at your location, network security considerations should be discussed.

The following considerations are based on Oracle best practices for Oracle Exadata.

- The use of intrusion prevention systems on database servers to monitor network traffic flowing to and from Oracle Exadata. Such systems enable the identification of suspicious communications, potential attack patterns, and unauthorized access attempts.
- The use of application and network-layer firewalls to protect information flowing to and from Oracle Exadata. Filtering network ports provides the first line of defense in preventing unauthorized access to systems and services.

Network-level segmentation using Ethernet virtual local area networks (VLANs) and host-based firewalls enforce inbound and outbound network policy at the host level. Using segmentation allows fine-grained control of communications between components of Oracle Exadata. Oracle Exadata Storage Servers include a configured software firewall by default. The database servers can be configured with a software firewall.

- The use of encryption features such as Oracle Advanced Security to encrypt traffic to Oracle Data Guard standby databases.

The security of the data and system is diminished by weak network security. Oracle recommends the following guidelines to maximize your Ethernet network security:

- Configure administrative and operational services to use encryption protocols and key lengths that align with current policies. Cryptographic services provided by Oracle Exadata benefit from hardware acceleration, which improves security without impacting performance.
- Manage and separate switches in Oracle Exadata from data traffic on the network. This separation is also referred to as "out-of-band."
- Separate sensitive clusters from the rest of the network by using virtual local area networks (VLANs). This decreases the likelihood that users can gain access to information on these clients and servers.
- Use a static VLAN configuration.
- Disable unused switch ports, and assign an unused VLAN number.
- Assign a unique native VLAN number to trunk ports.
- Limit the VLANs that can be transported over a trunk to only those that are strictly required.
- Disable VLAN Trunking Protocol (VTP), if possible. If it is not possible, then set the management domain, password and pruning for VTP. In addition, set VTP to transparent mode.
- Disable unnecessary network services, such as TCP small servers or HTTP. Enable only necessary network services, and configure these services securely.
- Network switches offer different levels of port security features. Use these port security features if they are available:
- Lock the Media Access Control (MAC) address of one or more connected devices to a physical port on a switch. If a switch port is locked to a particular MAC address, then super users cannot create back doors into the network with rogue access points.
- Disable a specified MAC address from connecting to a switch.
- Use each switch port's direct connections so the switch can set security based on its current connections.

Related Topics

- Understanding the Network Requirements for Oracle Exadata Database Machine
- Network Partitioning on Oracle Exadata Database Machine

2.11 Understanding the Default Security Settings

Oracle Exadata System Software is installed with many default security settings.

Whenever possible and practical, secure default settings should be chosen and configured. The following default settings are used in Oracle Exadata Database Machine:

- A minimal software installation to reduce attack surface.
- Oracle Database secure settings developed and implemented using Oracle best practices.
- A password policy that enforces a minimum password complexity.
- Failed log in attempts cause a lockout after a set number of failed attempts.
- All default system accounts in the operating system are locked and prohibited from logging in.
- Limited ability to use the su command.
- Password-protected boot loader installation.
- All unnecessary system services are disabled, including the Internet service daemon (inetd/xinetd).
- Software firewall configured on the storage cells.
- Restrictive file permissions on key security-related configuration files and executable files.
- SSH listen ports restricted to management and private networks.
- SSH limited to v2 protocol.
- Disabled insecure SSH authentication mechanisms.
- Configured specific cryptographic ciphers.
- Unnecessary protocols and modules are disabled from the operating system kernel.

3

User Security on Oracle Exadata

Increase the security of your data and system by limiting user access and developing strong password security policies.

- [Default User Accounts for Oracle Exadata](#)
Several user accounts regularly manage the components of Oracle Exadata.
- [Default Password Requirements](#)
Oracle Exadata Deployment Assistant (OEDA) implements a default password policy on Oracle Exadata Database Machine.
- [Default User Security Settings](#)
- [Managing Password and Authentication Policies](#)
- [Creating Oracle Exadata System Software Users and Roles](#)
You can control which Oracle Exadata System Software commands users can run by granting privileges to roles, and granting roles to users.
- [Security Policies for Oracle Exadata Storage Server Operating System Users](#)
User access to the operating system can be secured by the use of secure, hardened passwords.
- [Enhanced Default Security Settings for Storage Server Operating System Users](#)
Enhanced default security measures prevent storage server OS users from performing unnecessary and potentially damaging OS commands, either accidentally or maliciously.

3.1 Default User Accounts for Oracle Exadata

Several user accounts regularly manage the components of Oracle Exadata.

In addition to the root user, Oracle Exadata Storage Servers have two users, celladmin and cellmonitor. The celladmin user is used to run all services on the cell. The cellmonitor user is used for monitoring purposes. The cellmonitor user cannot run services on the cell. Other Oracle Exadata components have users for the management of the component.

Note

After Oracle Exadata has been deployed, the installation process disables all root SSH keys and expires all user passwords as a security measure for your system. If you do not want the SSH keys disabled or the passwords expired, advise the installation engineer before the deployment.

Starting with Oracle Exadata System Software release 19.1.0, two new users are created, to improve security of specific actions. The cellofl user runs query offload processes on the storage servers as a non-root user. The exawatch user is responsible for collecting and archiving system statistics on both the database servers and the storage servers.

The following table lists the default users and passwords for the Oracle Exadata components. All default passwords should be changed after installation of Oracle Exadata. Refer to My

Oracle Support document KB434220 for information about changing the default user accounts passwords.

Table 3-1 Default Oracle Exadata Users and Passwords

Account	Default Password	Account Type	Component(s)
root	welcome1	Operating system user	Oracle Exadata Database Servers Oracle Exadata Storage Servers RDMA Network Fabric switches Database server ILOMs Oracle Exadata Storage Server ILOMs RDMA Network Fabric ILOMs
oracle	Welcome\$	Operating system user	Oracle Exadata Database Servers
grid Note: This account exists only if role separation is chosen during deployment.	Welcome\$	Operating system user	Oracle Exadata Database Servers
celladmin	welcome Note: Commencing with the Oracle Exadata Deployment Assistant (OEDA) November 2019 release, the password of the celladmin user is set to a random string during deployment, which you must change on first use.	Operating system user	Oracle Exadata Storage Servers
CELLDIAG	Welcome12345 Note: The password of the CELLDIAG user is reset to a random password during the "Apply Security Fixes" step of OEDA.	Oracle Exadata System Software user	Oracle Exadata Storage Servers
cellmonitor	welcome Note: Commencing with the OEDA November 2019 release, the password of the cellmonitor user is set to a random string during deployment, which you must change on first use.	Operating system user	Oracle Exadata Storage Servers

Table 3-1 (Cont.) Default Oracle Exadata Users and Passwords

Account	Default Password	Account Type	Component(s)
cellofl Note: This account has no login privileges and exists only in release 19.1.0 and later.		Operating system user	Oracle Exadata Storage Servers
dbmadmin	welcome Note: Commencing with the OEDA November 2019 release, the password of the dbmadmin user is set to a random string during deployment, which you must change on first use.	Operating system user	Oracle Exadata Database Servers
dbmmonitor	welcome Note: Commencing with the OEDA November 2019 release, the password of the dbmmonitor user is set to a random string during deployment, which you must change on first use.	Operating system user	Oracle Exadata Database Servers
dbmsvc Note: This account has no login privileges and exists only in release 12.1.2.1.0 and later.		Operating system user	Oracle Exadata Database Servers
exawatch Note: This account has no login privileges and exists only in release 19.1.0 and later.		Operating system user	Oracle Exadata Database Servers Oracle Exadata Storage Servers
SYS	We1come\$	Oracle Database user	Oracle Exadata Database Servers
SYSTEM	We1come\$	Oracle Database user	Oracle Exadata Database Servers
GRUB boot loader administrator Note: Starting with Oracle Exadata System Software release 26.1.0, the default administrator username for the GRUB boot loader changes from root to grub-admin	sos1Exadata	Operating system user	Oracle Exadata Database Servers Oracle Exadata Storage Servers

Table 3-1 (Cont.) Default Oracle Exadata Users and Passwords

Account	Default Password	Account Type	Component(s)
nm2user	changeme	Firmware user	InfiniBand Network Fabric switches
ilom-admin	ilom-admin	ILOM user	InfiniBand Network Fabric switches
ilom-operator	ilom-operator	ILOM user	InfiniBand Network Fabric switches
admin	welcome1	Firmware/switch administrator	RoCE Network Fabric switches
admin	welcome1 Note: You should secure the enable mode <i>password</i> and <i>secret</i> values for the admin user.	Firmware user	Ethernet switches
admin	welcome1 Note: The password for the admin user is admin if you reset the PDU to factory default settings.	Firmware user	Power distribution units (PDUs) Keyboard, video, mouse (KVM)
MSUser Note: Management Server (MS) uses this account to reset the ILOM interface if it stops responding. Do not modify this account. This account is to be used by MS only.	The MSUser password is not persisted anywhere. Each time MS starts up, it deletes the previous MSUser account and re-creates the account with a randomly generated password.	ILOM user	Database server ILOMs Oracle Exadata Storage Server ILOMs
LocalMSV3user Note: Management Server (MS) uses this account for hardware monitoring and failure handling using an automatic ILOM SNMP notification rule. Do not modify this account or the associated ILOM SNMP notification rule. This account is to be used by MS only.	The LocalMSV3user password is not persisted anywhere. Each time MS starts up, it deletes the previous LocalMSV3user account and re-creates the account with a randomly generated password.	ILOM SNMP version 3 user	Database server ILOMs Oracle Exadata Storage Server ILOMs

Table 3-1 (Cont.) Default Oracle Exadata Users and Passwords

Account	Default Password	Account Type	Component(s)
rocedisc Note: By default, this account is disabled and cannot be used to log in to the RoCE Network Fabric switch. Do not delete this account. Otherwise, verification of the switch configuration will fail.		RoCE Network Fabric switch user	RoCE Network Fabric switches

Related Topics

- CREATE DIAGPACK
- [How to change OS user password for Cell Node, Database Node , ILOM, KVM , Infiniband Switch , GigaBit Ethernet Switch and PDU on Exadata Database Machine \(My Oracle Support Doc ID KB434220\)](#)

3.2 Default Password Requirements

Oracle Exadata Deployment Assistant (OEDA) implements a default password policy on Oracle Exadata Database Machine.

The last step of OEDA, "Secure Oracle Exadata Database Machine", implements the following password requirements:

- Dictionary words are not valid or accepted.
- Character classes for passwords are uppercase letters, lowercase letters, digits, and special characters.
- Passwords must contain characters from all four character classes. Passwords using only one, two, or three character classes are not allowed.
- The minimum length of a password is eight characters.
- Pass-phrases are allowed. A pass-phrase should contain at least three words, be 16 to 40 characters in length, and contain different character classes.
- A new password cannot be similar to old passwords. There must be at least eight characters in the new password that were not present in the old password.
- A maximum of three consecutive characters of the same value can be used in a password.
- A maximum of four consecutive characters of the same character class can be used in a password. For example, abcde1#6B cannot be used as a password because it uses five consecutive lower case letters.

3.3 Default User Security Settings

The following default user security settings are implemented on Oracle Exadata.

- The following password rules apply by default to all operating system operating system users on the database servers and storage servers:

- Non-root users must change their password during first login.
- The password complexity rules depend on the Oracle Linux version in use.

For systems with Oracle Linux 7 or later:

- * The minimum password length is 15 characters,
- * The password must contain at least one digit, one uppercase character, one lowercase character, and one other character.
- * The password must not contain the same character consecutively more than 3 times.
- * The password must not contain more than 4 consecutive characters from the same class (digits, lowercase letters, uppercase letters, or other characters).
- * For password changes, the new password must contain a minimum of 8 character changes.

For systems with Oracle Linux 6 or earlier, the minimum password length is 5 characters with no additional complexity requirements.

- The maximum password age is 60 days.
- The minimum amount of time between password changes is 1 day.
- Warning alerts are generated 7 days before password expiry.
- When changing a user password, the new password cannot match any of the 10 previous passwords.
- An operating system user account is locked for 15 minutes after three failed login attempts within a 15-minute period.
- Starting with Oracle Exadata System Software release 25.1.0, non-root operating system user accounts are automatically locked after 35 days of inactivity.
- Login sessions automatically terminate after 14400 seconds of no input.
- SSH sessions automatically terminate after 600 seconds of inactivity.
- For the root user, SSH equivalency is removed for all database servers and storage servers.

3.4 Managing Password and Authentication Policies

Each Oracle Exadata server contains the `host_access_control` utility (`/opt/oracle.cellos/host_access_control`), which provides simple interfaces to view and modify the password and authentication policies.

Oracle recommends using the `host_access_control` utility to view and modify the password and authentication policies. You may perform customizations outside the scope of the `host_access_control` utility at your own cost and risk.

- To manage the password aging policy settings for new account creation, use the `host_access_control` command with the `password-policy` option.
 - For a complete description of the available options and settings, use:

```
# /opt/oracle.cellos/host_access_control password-policy --help
```

- To view the current policy settings, use:

```
# /opt/oracle.cellos/host_access_control password-policy --status
```

- To reset the policy to the factory default settings, use:

```
# /opt/oracle.cellos/host_access_control password-policy --defaults
```

Under the factory default password aging policy:

- * The maximum password age is 60 days.
- * The minimum period allowed between password changes is 1 day.
- * The minimum password length is 8 characters.
- * The password expiry warning period is 7 days.

- To modify specific policy settings, specify one or more of the following:

- * --PASS_MAX_DAYS: Specifies the maximum password age (in days).
- * --PASS_MIN_DAYS: Specifies the minimum number of days allowed between password changes.
- * --PASS_MIN_LEN: Specifies the minimum password length.
- * --PASS_WARN_AGE: Specifies the password expiry warning period (in days).

For example, use the following command to set the maximum password age to 100 days and the minimum password length to 12 characters:

```
# /opt/oracle.cellos/host_access_control password-policy --  
PASS_MAX_DAYS 100 --PASS_MIN_LEN 12
```

- To manage the password aging policy for existing interactive user accounts, use the `host_access_control` command with the `password-aging` option.
 - For a complete description of the available options and settings, use:

```
# /opt/oracle.cellos/host_access_control password-aging --help
```

- To view the current policy settings, use:

```
# /opt/oracle.cellos/host_access_control password-aging --status
```

- To reset the password aging policy to the factory default settings, use:

```
# /opt/oracle.cellos/host_access_control password-aging --defaults
```

Under the factory default password aging policy:

- * The maximum password age is 60 days.
- * The minimum period allowed between password changes is 1 day.
- * The minimum password length is 8 characters.
- * The password expiry warning period is 7 days.

- To reset the password aging policy to the Exadata secure default settings, use:

```
# /opt/oracle.cellos/host_access_control password-aging --secdefaults
```

Under the Exadata secure default settings, the minimum password length is 15 characters. All other settings match the factory default policy.

- To modify existing users to use the policy settings for new account creation, which are the settings defined by using `host_access_control` with the `password-policy` option, use:

```
# /opt/oracle.cellos/host_access_control password-aging --policy
```

- To modify specific policy settings for a user, specify the user and one or more of the following attributes:

- * `--maxdays`: Specifies the maximum password age (in days).
- * `--mindays`: Specifies the minimum number of days allowed between password changes.
- * `--warndays`: Specifies the password expiry warning period (in days).

For example, use the following command to set the maximum password age to 80 days for the `oracle` OS user:

```
# /opt/oracle.cellos/host_access_control password-aging --maxdays 80 --user oracle
```

- To manage the system authentication policy settings, use the `host_access_control` command with the `pam-auth` option. The system authentication settings include the password complexity and password history rules that apply to all users.

Commencing with Oracle Exadata System Software release 23.1.0 and Oracle Linux 8, the security settings managed by the `pam-auth` option are encapsulated in a custom Exadata security profile using the Linux `authselect` utility.

- For a complete description of the available options and settings, use:

```
# /opt/oracle.cellos/host_access_control pam-auth --help
```

- To view the current authentication settings, use:

```
# /opt/oracle.cellos/host_access_control pam-auth --status
```

- To reset the authentication settings the factory default settings, use:

```
# /opt/oracle.cellos/host_access_control pam-auth --defaults
```

Under the factory default authentication settings:

- * A user account is locked for 15 minutes after three failed login attempts within a 15-minute period.
- * Starting with Oracle Exadata System Software release 25.1.0, non-root operating system user accounts are automatically locked after 35 days of inactivity.
- * When changing a user password, the new password cannot match any of the 10 previous passwords.

- * The password complexity rules depend on the Oracle Linux version in use.
For systems with Oracle Linux 7 or later:
 - * The minimum password length is 8 characters,
 - * The password must contain at least one digit, one uppercase character, one lowercase character, and one other character.
 - * The password must not contain the same character consecutively more than 3 times.
 - * The password must not contain more than 4 consecutive characters from the same class (digits, lowercase letters, uppercase letters, or other characters).
 - * For password changes, the new password must contain a minimum of 8 character changes.

For systems with Oracle Linux 6 or earlier, the minimum password length is 5 characters with no additional complexity requirements.

- To reset the authentication settings to the Exadata secure default settings, use:

```
# /opt/oracle.cellos/host_access_control pam-auth --secdefaults
```

Under the Exadata secure default settings:

- * A user account is locked for 15 minutes after three failed login attempts within a 15-minute period.
- * Starting with Oracle Exadata System Software release 25.2.0, non-root operating system user accounts are automatically locked after 35 days of inactivity.
- * When changing a user password, the new password cannot match any of the 10 previous passwords.
- * The password complexity rules depend on the Oracle Linux version in use.

For systems with Oracle Linux 7 or later:

- * The minimum password length is 15 characters,
- * The password must contain at least one digit, one uppercase character, one lowercase character, and one other character.
- * The password must not contain the same character consecutively more than 3 times.
- * The password must not contain more than 4 consecutive characters from the same class (digits, lowercase letters, uppercase letters, or other characters).
- * For password changes, the new password must contain a minimum of 8 character changes.

For systems with earlier Oracle Linux versions, the minimum password length is 8 characters and the password must contain at least one digit, one uppercase character, one lowercase character, and one other character. Alternatively, you can use a password with at least 12 characters that contains at least 3 out of the 4 character classes (digits, lowercase letters, uppercase letters, or other characters).

- To modify specific authentication settings, specify one or more of the following:
 - * --deny: Specifies the required number of consecutive failed login attempts within the interval (specified by --interval) to trigger an account lockout.

- * `--interval`: Specifies the number of seconds during which the consecutive failed login attempts must happen to trigger an account lockout.
- * `--lock`: Specifies the duration (in seconds) of an account lockout.
- * `--inactive`: Specifies the duration (in days) before an inactive non-root account is automatically locked.

This option is available beginning with Oracle Exadata System Software release 25.2.0.

- * `--passwdqc`: This setting applies only to systems with Oracle Linux 6 or earlier. The value is a comma-separated list defining the minimum allowed length for different types of passwords or passphrases. See the `pam_passwdqc` Linux man page for details about this setting.
- * `--pwquality`: This setting applies only to systems with Oracle Linux 7 or later. The value is either an integer that defines the minimum password length or a comma-separated list that defines the password complexity rules using the following attributes: `minlen`, `dcredit`, `ucredit`, `lcredit`, `ocredit`, `difok`, `maxrepeat`, `maxclassrepeat`, `minclass`, `maxsequence`, and `gecoscheck`. See the `pam_pwquality` Linux man page for details about the password complexity attributes.
- * `--remember`: Specifies the size of the password history list for each user. For password changes, the new password cannot match any of previous passwords in the password history list.

For example, use the following command to set the lockout period to 20 minutes after two failed login attempts within a 10-minute period:

```
# /opt/oracle.cellos/host_access_control pam-auth --lock 1200 --deny 2
--interval 600
```

Related Topics

- [Default User Security Settings](#)

3.5 Creating Oracle Exadata System Software Users and Roles

You can control which Oracle Exadata System Software commands users can run by granting privileges to roles, and granting roles to users.

For example, you can specify that a user can run the `LIST GRIDDISK` command but not `ALTER GRIDDISK`. This level of control is useful in Oracle Cloud environments, where you might want to allow full access to the system to only a few users.

- [Overview of Creating Exadata System Software Users](#)
- [Creating Roles and Getting Information about Roles](#)
Use the `CREATE ROLE` command to create roles for Oracle Exadata System Software users.
- [Granting and Revoking Privileges](#)
Use the `GRANT PRIVILEGE` command to grant privileges to roles for Oracle Exadata System Software users.
- [Creating Users](#)
Use the `CREATE USER` command to create Oracle Exadata System Software users.
- [Configuring Password Expiration for Users Accessing the Server Remotely](#)
You can configure `CELL` attributes to expire user passwords.

- [Granting and Revoking Roles](#)

Use the GRANT ROLE command to create roles to Oracle Exadata System Software users.

Related Topics

- Using the ExaCLI Utility

3.5.1 Overview of Creating Exadata System Software Users

Oracle Exadata System Software users are required when running ExaCLI in on-premise or Oracle Cloud environments. ExaCLI enables you to manage cells remotely from compute nodes. When you run ExaCLI on a compute node, you need to specify a user name to use to connect to the cell node. The Management Server (MS) authenticates the user credentials, then performs authorization checks on the commands issued by the user. If the user does not have the proper privileges to run a command, MS returns an error.

The password security key is encrypted using Password-Based Key Derivation Function 2 (PBKDF2) with HMAC-SHA1.

The high-level steps for creating users and roles for use with Oracle Exadata System Software are:

1. Create roles using the CREATE ROLE command.
2. Grant privileges to roles using the GRANT PRIVILEGE command.
3. Create users using the CREATE USER command.
4. Grant roles to users using the GRANT ROLE command.

You can also revoke privileges from roles using the REVOKE PRIVILEGE command. To revoke roles from users, use the REVOKE ROLE command.

3.5.2 Creating Roles and Getting Information about Roles

Use the CREATE ROLE command to create roles for Oracle Exadata System Software users.

For example, to create a role for administrators, you could use the following command:

```
CellCLI> CREATE ROLE admin
```

After you have created a role, you can then grant privileges to the role using the GRANT PRIVILEGE command. You can also grant the role to users, for example:

```
CellCLI> GRANT PRIVILEGE ALL ACTIONS ON ALL OBJECTS TO ROLE admin
```

```
CellCLI> GRANT ROLE admin TO USER username
```

To get detailed information about a role, use the LIST ROLE command. The following command returns all the attributes for the admin role.

```
CellCLI> LIST ROLE admin DETAIL
      name:                admin
      privileges:           object=all objects, verb=all actions,
attributes=all attributes, options=all options
```

Related Topics

- CREATE ROLE
- LIST ROLE

3.5.3 Granting and Revoking Privileges

Use the GRANT PRIVILEGE command to grant privileges to roles for Oracle Exadata System Software users.

- Grant privileges to roles using the GRANT PRIVILEGE command.
 - The following example grants all privileges to Oracle Exadata System Software users with the admin role.

```
CellCLI> GRANT PRIVILEGE ALL ACTIONS ON ALL OBJECTS TO ROLE admin
```

- You can also grant individual command privileges to a role.

```
CellCLI> GRANT PRIVILEGE list ON griddisk TO ROLE diskmonitor
```

- You can also grant all command privileges for specific objects to a role.

```
CellCLI> GRANT PRIVILEGE ALL ACTIONS ON griddisk TO ROLE diskadmin
```

- You can revoke privileges from roles using the REVOKE PRIVILEGE command.

```
CellCLI> REVOKE PRIVILEGE ALL ACTIONS ON griddisk FROM ROLE diskadmin
```

Related Topics

- GRANT PRIVILEGE
- REVOKE PRIVILEGE

3.5.4 Creating Users

Use the CREATE USER command to create Oracle Exadata System Software users.

A newly created user does not have any privileges. The Oracle Exadata System Software user is granted privileges through roles granted to the user.

1. Use the CREATE USER command to create a user and assign an initial password.

The following command creates a user called fred with password uq==A*2D\$_18.

```
CellCLI> CREATE USER fred PASSWORD = "uq==A*2D$_18"
```

2. To grant privileges to the new user fred, use the GRANT ROLE command for a role that has already been configured.

Related Topics

- CREATE USER
- [Granting and Revoking Privileges](#)

Use the GRANT PRIVILEGE command to grant privileges to roles for Oracle Exadata System Software users.

- [Default Password Requirements](#)
Oracle Exadata Deployment Assistant (OEDA) implements a default password policy on Oracle Exadata Database Machine.

3.5.5 Configuring Password Expiration for Users Accessing the Server Remotely

You can configure CELL attributes to expire user passwords.

In Oracle Exadata System Software release 19.1.0, there are new CELL attributes for configuring password security for users that access Oracle Exadata System Software servers remotely, such as with REST API or ExaCLI. These attributes determine if the user is able to change the password remotely, the amount of time before a user password expires, and the number of days prior to password expiration that the user receives warning messages. In the default configuration, user passwords do not expire.

Note

The CELL attributes for password expiration apply only to users created with Oracle Exadata System Software. Password expiration applies only to users that are displayed with the LIST USER command and does not apply to operating system users like celladmin or oracle.

- To allow the user to change the password remotely, use the ALTER CELL command to set the remotePwdChangeAllowed attribute to true.

If you set the value to false, then the user receives a message indicating that they must contact the server administrator to have their password changed.

```
CellCLI> ALTER CELL remotePwdChangeAllowed=true
```

- To change the length of time before a user password expires, use the ALTER CELL command to modify the pwdExpInDays attribute.

Set the value *n* to the number of days before the password expires. If pwdExpInDays is set to 0 (the default value), then the user password does not expire.

```
CellCLI> ALTER CELL pwdExpInDays=n
```

- To configure the length of the warning period before the password expires, use the ALTER CELL command to modify the pwdExpWarnInDays attribute.

Set the value *n* to the number of days to warn the user before the password expires. The default user account password expiration warning time is 7 days.

```
CellCLI> ALTER CELL pwdExpWarnInDays=n
```

- To specify the length of time before a user account is locked after the user password expires, use the ALTER CELL command to modify the accountLockInDays attribute.

Set the value *n* to the number of days before the user account is locked. The default user account lock time is 7 days.

```
CellCLI> ALTER CELL accountLockInDays=n
```


3.5.6 Granting and Revoking Roles

Use the GRANT ROLE command to create roles to Oracle Exadata System Software users.

Command privileges are granted to roles, and then the roles are granted to users. You do not grant command privileges directly to the Oracle Exadata System Software users.

- Use the GRANT ROLE command to grant roles to users.

The following example grants the admin role to the user fred.

```
CellCLI> GRANT ROLE admin TO USER fred
```

- You can revoke roles from users using the REVOKE ROLE command.

Related Topics

- GRANT ROLE
- REVOKE ROLE

3.6 Security Policies for Oracle Exadata Storage Server Operating System Users

User access to the operating system can be secured by the use of secure, hardened passwords.

The passwords for operating system users who administer Oracle Exadata System Software adhere to the security guidelines enacted by Oracle Exadata Deployment Assistant (OEDA). See [Default Security Setting Enacted by OEDA](#) for more information.

- [Changing a Password](#)
Use the operating system command passwd to change user passwords.
- [Enabling the Security Policies for Operating System Users](#)
The /opt/oracle.cell/os/RESECURED_NODE file enables the security policies.
- [Viewing Failed Operating System Password Attempts](#)
- [Resetting a Locked Operating System User Account](#)

3.6.1 Changing a Password

Use the operating system command passwd to change user passwords.

Operating system users are notified of the need to change their passwords 7 days before the expiration date.

- To change a password, use the passwd command, where *username* is the user name for which you want to change the password.

```
passwd username
```

3.6.2 Enabling the Security Policies for Operating System Users

The `/opt/oracle.celllos/RESECURED_NODE` file enables the security policies.

If the file does not exist, then you can reset the security policies for all operating system users by performing the following steps:

1. Shut down the Oracle Grid Infrastructure services on all database servers.
2. Shut down the cell services on the storage servers.

`cellcli -e alter cell shutdown services all`
3. Use the `harden_passwords_reset_root_ssh` script to reset the security policies.

Note

The `harden_passwords_reset_root_ssh` script restarts the cell.

`/opt/oracle.SupportTools/harden_passwords_reset_root_ssh`

4. All operating system users must set a new password the next time they log in.

3.6.3 Viewing Failed Operating System Password Attempts

Use the `faillock` operating system utility to view failed login attempts.

For example, to see the failed login attempts for the `celladmin` user:

```
# faillock --user celladmin
celladmin:
When                Type  Source
Valid
2022-11-06 18:23:18 RHOST
xxx.xxx.xxx.xxx      V
2022-11-06 18:23:23 RHOST
xxx.xxx.xxx.xxx      V
2022-11-06 18:23:27 RHOST
xxx.xxx.xxx.xxx      V
```

In the example output, `xxx.xxx.xxx.xxx` represents the IP address that is the source of the login failure.

3.6.4 Resetting a Locked Operating System User Account

The operating system user account is locked when the number of login failures exceeds the threshold in the security policy.

To reset a locked account, use the following command, where *username* is the name of the locked user:

```
# faillock --user username --reset
```

3.7 Enhanced Default Security Settings for Storage Server Operating System Users

Enhanced default security measures prevent storage server OS users from performing unnecessary and potentially damaging OS commands, either accidentally or maliciously.

Starting with Oracle Exadata System Software release 26.1.0, the Exadata storage server built-in operating system (OS) user accounts, *celladmin* and *cellmonitor*, are restricted to only run specific Cell Command Line Interface (CellCLI) and Exascale Command Line Interface (ESCLI) commands, along with a limited set of read-only Linux OS commands.

Additionally, for Exascale environments, each OS user now uses a separate Exascale keystore (wallet) with appropriate corresponding credentials. By default, the wallet associated with *celladmin* provides Exascale administrator privileges, while the *cellmonitor* wallet provides read-only access to Exascale objects and metadata.

These changes limit the ability to perform malicious operations using the *celladmin* and *cellmonitor* accounts if their credentials are compromised. They also reduce the possibility of accidental damage through mistakes by legitimate users.

All of these changes are configured and enabled by default on newly deployed systems using Oracle Exadata System Software release 26.1.0 or later. On systems updated from an earlier release, you must create the *cellmonitor* Exascale wallet. On each storage server, the wallet must be located at `/opt/oracle/cell/cellsrv/deploy/config/cellmonitor_wallet` and must be linked to an Exascale user account with the `cl_monitor` privilege.

Related Topics

- [Exascale Users](#)
- [Exascale User Credentials](#)
- [Exascale Cluster Level Privileges](#)

4

Keeping the Oracle Exadata Secure

This chapter describes policies and procedures to keep Oracle Exadata secure.

- [Securing the Hardware](#)
After installation of Oracle Exadata Database Machine, the hardware should be secured.
- [Securing the Software](#)
Frequently, hardware security is implemented through software measures.
- [Managing SSH on Storage Servers](#)
You can control remote SSH access to Exadata storage servers. By default, SSH is enabled on storage servers.
- [Configuring Data Security for Exadata Storage Servers](#)
Data security for Oracle Exadata Storage Servers is implemented by controlling which Oracle Automatic Storage Management (Oracle ASM) clusters and Oracle Database clients can access specific grid disks on storage cells.
- [Maintaining a Secure Environment](#)
After security measures are implemented, they must be maintained to keep the system secure.

4.1 Securing the Hardware

After installation of Oracle Exadata Database Machine, the hardware should be secured.

Hardware can be secured by restricting access to the hardware and recording the serial numbers. Oracle recommends the following practices to restrict access:

- Install Oracle Exadata Database Machine and related equipment in a locked, restricted-access room.
- Lock the rack door unless service is required on components within the rack.
- Restrict access to hot-pluggable or hot-swappable devices because the components can be easily removed by design. See
- Store spare field-replaceable units (FRUs) or customer-replaceable units (CRUs) in a locked cabinet. Restrict access to the locked cabinet to authorized personnel.
- Limit SSH listener ports to the management and private networks.
- Use SSH protocol 2 (SSH-2) and FIPS 140-2 approved ciphers.
- Limit SSH allowed authentication mechanisms. Inherently insecure methods are disabled.
- Mark all significant items of computer hardware, such as FRUs.
- Keep hardware activation keys and licenses in a secure location that is easily accessible to the system managers in the case of a system emergency.
- Record the serial numbers of the components in Oracle Exadata Database Machine, and keep a record in a secure place. All components in Oracle Exadata Database Machine have a serial number.
- [Getting the Rack Serial Number](#)
Use the `ipmitool` utility to get the serial number for the rack.

- [Getting the Serial Numbers for Rack Components](#)
The CheckHWnFWProfile command can be used to display the serial number of most of the system components.
- [Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch](#)
Use the show license host-id command on the switch to get the serial number.
- [Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36](#)
Use the showfruinfo command on the switch to get the serial number.
- [Getting the Serial Number for a Cisco 4948 Ethernet Switch](#)
Use the sh inventory command on the switch to get the serial number.

4.1.1 Getting the Rack Serial Number

Use the ipmitool utility to get the serial number for the rack.

When interacting with Oracle Support Services, the CSI number for a rack is based on the rack serial number.

1. Log in to one of the servers in the rack as the root user.
2. Use ipmitool to get the serial number for the rack.

```
# ipmitool sunoem cli "show /SP system_identifier"
Connected. Use ^D to exit.
-> show /SP system_identifier

/SP
Properties:
    system_identifier = Exadata Database Machine X2-8xxxxAKyyyy

-> Session closed
Disconnected
```

4.1.2 Getting the Serial Numbers for Rack Components

The CheckHWnFWProfile command can be used to display the serial number of most of the system components.

1. Log in to one of the servers in the rack as the root user.
2. On each server in the rack, use CheckHWnFWProfile with the -S option to display the serial number of the components for that server.

```
# /opt/oracle.SupportTools/CheckHWnFWProfile -S > /tmp/
CheckHWnFWProfile_hostname.txt
```

The result is specific to each server, so the command must be performed on every node. The following is a partial example of the output:

```
Server_Model=ORACLE_SERVER_X8-2L
====START SERIAL NUMBERS====
==Motherboard, from dmidecode==
--System serial--
```

```

1904XCA000
--Motherboard serial--
469996N+0000RD01RN
--Chassis serial--
1900XCA000
--Rack serial--
AK00400000
==Infiniband HCA==
ID:      CX354A - ConnectX-3 QSFP
PN:      7046442
EC:      XX
SN:      465000K-1800000000
V0:      PCIe Gen3 x8
==Motherboard, RAM etc from ipmitool==
FRU Device Description : Builtin FRU Device (LUN 0 ID 0)
...
Product Name           : ILOM
Product Version        : 4.0.4.38.a

FRU Device Description : BMC
...
Product Name           : ILOM
Product Version        : 4.0.4.38.a

FRU Device Description : /SYS (LUN 0 ID 3)
...
Product Part Number    : 8200669
Product Serial         : 1900XCA000

FRU Device Description : DBP (LUN 0 ID 210)

Board Part Number      : 7341141
Board Extra            : Rev 09

FRU Device Description : HDD0 (LUN 0 ID 47)
Device not present (Requested sensor, data, or record not found)

FRU Device Description : HDD1 (LUN 0 ID 48)
Device not present (Requested sensor, data, or record not found)

...

FRU Device Description : MB (LUN 0 ID 4)
Board Mfg Date         : Sun Jan 20 16:57:00 2019
Board Mfg              : Oracle Corporation
...

FRU Device Description : MB/BIOS (LUN 0 ID 5)
...

FRU Device Description : MB/CPLD (LUN 0 ID 8)
Product Manufacturer   : Oracle Corporation
Product Name           : Power Control FPGA
Product Version        : FW:3.9

FRU Device Description : M2R0/SSD0 (LUN 0 ID 211)

```

Device not present (Requested sensor, data, or record not found)

FRU Device Description : M2R1/SSD0 (LUN 0 ID 212)

Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/NET0 (LUN 0 ID 43)

Product Manufacturer : INTEL

Product Name : 1G Ethernet Controller

...

FRU Device Description : MB/P0 (LUN 0 ID 16)

Product Manufacturer : Intel

Product Name : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz

...

FRU Device Description : MB/P0/D0 (LUN 0 ID 24)

Product Manufacturer : Samsung

Product Name : 16384MB DDR4 SDRAM DIMM

...

FRU Device Description : MB/P0/D1 (LUN 0 ID 25)

Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/P0/D2 (LUN 0 ID 26)

Product Manufacturer : Samsung

Product Name : 16384MB DDR4 SDRAM DIMM

...

FRU Device Description : MB/P1 (LUN 0 ID 17)

Product Manufacturer : Intel

Product Name : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz

...

FRU Device Description : MB/P1/D0 (LUN 0 ID 36)

Product Manufacturer : Samsung

Product Name : 16384MB DDR4 SDRAM DIMM

...

FRU Device Description : PS0 (LUN 0 ID 63)

...

FRU Device Description : PS1 (LUN 0 ID 64)

...

FRU Device Description : SP/NET0 (LUN 0 ID 1)

...

FRU Device Description : SP/NET1 (LUN 0 ID 2)

...

FRU Device Description : /UUID (LUN 0 ID 6)

...

FRU Device Description : TOP_LEVEL_CH (LUN 0 ID 251)

Chassis Type : Rack Mount Chassis

Chassis Part Number : 8200669

Chassis Serial : 1900XCA0000

Chassis Extra : chassis_name:ORACLE SERVER X8-2L

FRU Device Description : TOP_LEVEL_PROD (LUN 0 ID 250)

Product Manufacturer : Oracle Corporation

```
Product Name       : Exadata X8-2
Product Part Number : Exadata X8-2
Product Serial     : AK00430000
```

```
====END SERIAL NUMBERS====
```

4.1.3 Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch

Use the `show license host-id` command on the switch to get the serial number.

1. Connect to the switch from a server with SSH equivalency configured, or log in as the admin user.
2. Obtain the serial number for the switch by entering the `show license host-id` command.

The host ID is also referred to as the device serial number.

```
# switch# show license host-id
License hostid: VDH=FLA12345678
```

Use the entire ID that appears after the equal sign (=). In this example, the host ID is FLA12345678.

4.1.4 Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36

Use the `showfruinfo` command on the switch to get the serial number.

1. Log in to the switch as root.

```
$ ssh root@switch_name
```

2. Use the `showfruinfo` command to view the serial number for the switch.

```
root@ib-switch-> showfruinfo
Sun_Man1R:
UNIX_Timestamp32 : Fri Mar 19 16:29:59 2010
Sun_Fru_Description : ASSY,NM2-GW
Vendor_ID_Code : 11 E1
Vendor_ID_Code_Source : 01
Vendor_Name_And_Site_Location : 4577 CELESTICA CORP. SAN JOSE CA US
Sun_Part_Number : 5111402
Sun_Serial_Number : 0110SJC-1010NG0040
Serial_Number_Format : 4V3F1-2Y2W2X4S
Initial_HW_Dash_Level : 03
Initial_HW_Rev_Level : 50
Sun_Fru_Shortname : NM2 gateway
Sun_Hazard_Class_Code : Y
Sun_SpecPartNo : 885-1655-01
Sun_FRU_LabelR:
Sun_Serial_Number : AK000XXXX2
FRU_Part_Dash_Number : 541-4188-01
```


4.1.5 Getting the Serial Number for a Cisco 4948 Ethernet Switch

Use the `sh inventory` command on the switch to get the serial number.

1. Log in to the Cisco Ethernet switch.
2. Obtain the serial number for the switch and its components by entering the `sh inventory` command.

```
# Switch# sh inventory
NAME: "Switch System", DESCR: "Cisco Systems, Inc. WS-C4948 1 slot switch "
PID: , VID: , SN: FOX0000G0B6
NAME: "Linecard(slot 1)", DESCR: "10/100/1000BaseT (RJ45), 1000BaseX
(SFP)
Supervisor with 48 10/100/1000BASE-T ports and 4 1000BASE-"
PID: WS-C4948 , VID: V09 , SN: FOX0000G0B6
NAME: "Power Supply 1", DESCR: "Power Supply ( AC 300W )"
PID: PWR-C49-300AC , VID: , SN: QCS0000B1XR
NAME: "Power Supply 2", DESCR: "Power Supply ( AC 300W )"
PID: PWR-C49-300AC , VID: , SN: QCS0000B1X5
```

4.2 Securing the Software

Frequently, hardware security is implemented through software measures.

Implement the following guidelines to protect hardware and software:

- Change all default passwords when the system is installed at the site. Oracle Exadata Database Machine uses default passwords for initial installation and deployment that are widely known. A default password could allow unauthorized access to the equipment. Devices such as the network switches have multiple user accounts. Be sure to change all account passwords on the components in the rack.
- Limit use of the root super user account. Create and use Integrated Lights Out Manager (ILOM) user accounts for individual users to ensure a positive identification in audit trails, and less maintenance when administrators leave the team or company.
- Ensure Oracle Exadata Database Machine is deployed with separate software owner accounts for Oracle Grid Infrastructure and Oracle Database software installations.

Note

Separate software owner accounts for Oracle Grid Infrastructure and Oracle Database software installations are required for enabling DB-scoped security.

- Disable unnecessary protocols and modules in the operating system.
- Restrict physical access to USB ports, network ports, and system consoles. Servers and network switches have ports and console connections, which provide direct access to the system.
- Restrict the capability to restart the system over the network.
- Refer to the documentation to enable available security features.

Oracle Exadata Database Machine can leverage all the security features available with Oracle Database releases installed on legacy platforms. Oracle Database security products and features include the following:

- Oracle Advanced Security
- Oracle Audit Vault
- Data Masking
- Oracle Database Firewall
- Oracle Database Vault
- Oracle Label Security
- Oracle Secure Backup
- Flashback Data Archive

Using the Oracle privileged user and multi-factor access control, data classification, transparent data encryption, auditing, monitoring, and data masking, customers can deploy reliable data security solutions that do not require any changes to existing applications.

Related Topics

- [Default User Accounts for Oracle Exadata](#)
Several user accounts regularly manage the components of Oracle Exadata.
- *Oracle AI Database Security Guide*

4.3 Managing SSH on Storage Servers

You can control remote SSH access to Exadata storage servers. By default, SSH is enabled on storage servers.

Two storage server CELL attributes control remote SSH access:

- `accessLevelPerm`: This attribute specifies the default cell access level. This setting applies at all times unless overridden by `accessLevelTemp`.
- `accessLevelTemp`: This attribute specifies a temporary access level for a specified duration. During the specified duration, `accessLevelTemp` overrides `accessLevelPerm`. After the specified duration expires, the cell access level reverts back to the `accessLevelPerm` value. This attribute is typically used to enable a scheduled maintenance operation, such as a software update.

The `accessLevelPerm` and `accessLevelTemp` attributes permit the following access level settings:

- `remoteLoginEnabled`: SSH service is enabled. You can access the cell using SSH or ExaCLI. This is the default value for `accessLevelPerm`.
- `remoteLoginDisabled`: SSH service is disabled.

If the SSH service is disabled, you can still perform operations on the storage server using ExaCLI, which runs on the database servers and communicates using HTTPS and REST APIs to a web service running on the storage server.

When you need to perform operations that require remote login to the storage server, you can temporarily unlock the storage server. After the operation is complete, you can relock the storage server.

- `remoteLoginEnabledRootLoginDisabled`: SSH service is enabled, but remote root login is not permitted. This setting enables remote SSH access to the storage server using the

celladmin and cellmonitor operating system (OS) user accounts while disabling remote root access.

This setting compliments the enhanced default security measures introduced in Oracle Exadata System Software release 26.1.0, which prevent the celladmin and cellmonitor storage server OS users from performing unnecessary and potentially damaging OS commands, either accidentally or maliciously. See [Enhanced Default Security Settings for Storage Server Operating System Users](#).

Note

remoteLoginEnabledRootLoginDisabled may only be specified starting with Oracle Exadata System Software release 26.1.0.

Regardless of the setting, the current access level persists across storage server reboots.

- [Locking a Cell](#)
You lock a cell by setting its accessLevelPerm attribute to remoteLoginDisabled.
- [Unlocking a Cell Temporarily](#)
You can unlock a locked storage server, or cell, for a specified time period to perform scheduled maintenance or upgrades that require SSH login to the storage server.
- [Unlocking a Cell Permanently](#)
You can unlock a cell by setting its accessLevelPerm attribute to remoteLoginEnabled.
- [Checking the Current Access Level for a Cell](#)
View the accessLevelPerm and accessLevelTemp attributes for a cell to determine the current access level.
- [Access Level Alerts from the Management Server](#)
A stateless alert is generated when the accessLevelPerm attribute is modified.
- [Managing SSH Equivalence for Storage Server OS Users](#)

4.3.1 Locking a Cell

You lock a cell by setting its accessLevelPerm attribute to remoteLoginDisabled.

You must use a user that has the privilege to alter the accessLevelPerm attribute.

1. Grant the necessary privileges to a user.

On the storage server, run these commands:

```
cellcli> create role administrator
cellcli> grant privilege all actions on all objects all attributes with
all options to role administrator
cellcli> create user celladministrator password=*
cellcli> grant role administrator to user celladministrator
```

2. Run ExaCLI as the celladministrator user and run the ALTER CELL command:

```
$ exacli -l celladministrator -c exam08cel01
Password=*****
```

```
exacli> alter cell accessLevelPerm = remoteLoginDisabled
```

4.3.2 Unlocking a Cell Temporarily

You can unlock a locked storage server, or cell, for a specified time period to perform scheduled maintenance or upgrades that require SSH login to the storage server.

You can specify the start time of a temporary access window and how long it should last by using the `ALTER CELL` command to set the `accessLevelTemp` attribute.

The `accessLevelTemp` attribute includes the following properties:

- `accessLevel`: (Mandatory) Specifies the access level using one of the following values:
 - `remoteLoginEnabled`: SSH service is enabled for all remote logins.
 - `remoteLoginDisabled`: SSH service is disabled.
 - `remoteLoginEnabledRootLoginDisabled`: SSH service is enabled, but remote root login is not permitted.

Note

`remoteLoginEnabledRootLoginDisabled` may only be specified starting with Oracle Exadata System Software release 26.1.0.

- `startTime`: Specifies when the access level window starts. The time is specified in the ISO 8601 format: "yyyy-MM-ddTHH:mm:ssZ". You can also specify the keyword `now` to indicate that the specified access level should start immediately. The default value for this attribute is `now`.
- `duration`: Specifies how long the access level window should last. The default value is 2h (2 hours).

You can specify the duration as a number of minutes, hours, or days using the following suffixes:

- `m`: for minutes. For example, to specify 90 minutes, use `90m`.
- `h`: for hours. For example, to specify 1 hour, use `1h`.
- `d`: for days. For example, to specify 3 days, use `3d`.

You can also use composite values. For example, to specify 1 day and 12 hours, use `1d12h`.

- `reason`: Specifies a reason for changing the access level. For example: "Performing an upgrade". The default value is `none`.

Note the following:

- Only one temporary access window is allowed at any time. An error is displayed if you create a new temporary access window when one is currently active.
- To modify a future temporary access window that is not yet active, run the `ALTER CELL` command with a new `accessLevelTemp` setting. In this case, the new setting replaces the previous definition.
- You can modify an existing temporary access window to change the duration or reason (or both). To do this, run the `ALTER CELL` command and specify the `accessLevelTemp` setting using the exact start time that was used to create the existing temporary access window, along with the updated duration or reason (or both).

Example 4-1 Creating a Temporary Access Window

The following example creates a two-hour temporary access window that starts immediately. The command uses the default values for start time and duration.

```
exacli> ALTER CELL accessLevelTemp=((accessLevel="remoteLoginEnabled", -  
                                     reason="Quarterly maintenance"))
```

Example 4-2 Creating a Temporary Access Window in the Future

The following example creates a 30 minute temporary access window beginning on June 20, 2026, at 1:01 AM.

```
exacli> ALTER CELL accessLevelTemp=((accessLevel="remoteLoginEnabled", -  
                                     startTime="2026-06-20T01:01:00-07:00", -  
                                     duration="30m", -  
                                     reason="Quarterly maintenance"))
```

Example 4-3 Extending a Temporary Access Window

The following example extends the temporary access window created in the previous example to 5 hours. Note that the start time must match the window that is being adjusted.

```
exacli> ALTER CELL accessLevelTemp=((accessLevel="remoteLoginEnabled", -  
                                     startTime="2026-06-20T01:01:00-07:00", -  
                                     duration="5h", -  
                                     reason="Quarterly maintenance window extended to 5 hrs"))
```

Example 4-4 Deleting a Temporary Access Window

The following example deletes the temporary access window. If the temporary access window is currently active, it is closed immediately and the access level will be set back to the permanent access level. If the temporary access window is in the future and not yet active, it is canceled.

```
exacli> ALTER CELL accessLevelTemp=''
```

4.3.3 Unlocking a Cell Permanently

You can unlock a cell by setting its `accessLevelPerm` attribute to `remoteLoginEnabled`.

You must use a user that has the privilege to alter the `accessLevelPerm` attribute.

1. Start ExaCLI as a user that has the privilege to alter the `accessLevelPerm` attribute.

For example:

```
$ exacli -l celladministrator -c exam08cel01  
Password=*****
```

```
exacli>
```

2. Run the `ALTER CELL` command to set the `accessLevelPerm` attribute to `remoteLoginEnabled`.

For example:

```
exaccli> alter cell accessLevelPerm=remoteLoginEnabled
Cell exam08cel01 successfully altered
```

4.3.4 Checking the Current Access Level for a Cell

View the `accessLevelPerm` and `accessLevelTemp` attributes for a cell to determine the current access level.

- To see what the current access level is, use the `LIST CELL` command.

```
exaccli> LIST CELL ATTRIBUTES name,accessLevelPerm,accessLevelTemp
```

4.3.5 Access Level Alerts from the Management Server

A stateless alert is generated when the `accessLevelPerm` attribute is modified.

A stateful alert is generated when the `accessLevelTemp` window is created. An alert email is sent out when the `accessLevelTemp` window is activated. The alert is cleared when the window expires.

4.3.6 Managing SSH Equivalence for Storage Server OS Users

Starting with Oracle Exadata System Software release 26.1.0, Exadata administrators can configure SSH equivalence for the `celladmin` and `cellmonitor` storage server operating system (OS) user accounts using CellCLI commands. This capability removes the requirement for root access to configure SSH equivalence and complements the ability to disable root SSH access.

To configure SSH equivalence for a storage server OS user account, use the CellCLI `CREATE SSHEQUIVALENCE` command. The command requires two arguments:

- `LOGINUSERNAME`: Specifies the storage server OS user account, either `celladmin` or `cellmonitor`.
- `PUBLICKEY`: Specifies the public key to associate with the specified storage server OS user account.

For example:

```
CellCLI> CREATE SSHEQUIVALENCE LOGINUSERNAME=celladmin,PUBLICKEY="ssh-rsa
AAAAB3NzaC1yc2EAAA...ADAQABAAQADgFU0hS"
```

At any time, Exadata administrators can view existing SSH equivalence configuration details using the `LIST SSHEQUIVALENCE` command.

For example:

```
CellCLI> LIST SSHEQUIVALENCE attributes loginusername,publickey
```

```
CellCLI> LIST SSHEQUIVALENCE where loginusername=celladmin
```

If required, configuration details for SSH equivalence can be removed from a storage server using the `DROP SSHEQUIVALENCE` command.

For example:

```
CellCLI> DROP SSHEQUIVALENCE where loginusername=cellmonitor
```

```
CellCLI> DROP SSHEQUIVALENCE ALL
```

Related Topics

- [CREATE SSHEQUIVALENCE](#)
- [LIST SSHEQUIVALENCE](#)
- [DROP SSHEQUIVALENCE](#)

4.4 Configuring Data Security for Exadata Storage Servers

Data security for Oracle Exadata Storage Servers is implemented by controlling which Oracle Automatic Storage Management (Oracle ASM) clusters and Oracle Database clients can access specific grid disks on storage cells.

By default, all Oracle Database and Oracle ASM instances have access to all grid disks on the storage servers. The storage for each database is provided by Oracle ASM. You can have multiple clustered or unclustered databases running in your Oracle Exadata Database Machine. You can also have more than one Oracle ASM cluster. An Oracle ASM cluster is a collection of interconnected nodes, each with an Oracle ASM instance, operating as a unified cluster. An Oracle ASM cluster presents a shared pool of storage to one or more Oracle Database instances that are also operating on the nodes.

- [About Exadata Storage Server Security Modes](#)
Oracle Exadata System Software security allows open security, ASM-scoped security, or DB-scoped security.
- [Best Practices for ASM-Scoped Security and DB-Scoped Security](#)
While setting up security, it is imperative that the configuration is the same across all the storage servers.
- [About Security Keys](#)
Oracle Exadata System Software uses keys to identify clients and determine access rights to the grid disks.
- [Setting Up ASM-Scoped Security on Oracle Exadata Storage Servers](#)
Configuring ASM-scoped security requires actions to be performed on both the database servers and storage servers.
- [Setting Up DB-Scoped Security on Oracle Exadata Database Machine](#)
When configuring DB-scoped security, you perform actions on both the database and storage servers.
- [Changing Security Keys for ASM-Scoped Security or DB-Scoped Security](#)
You can change the keys used with ASM-scoped security or DB-scoped security.
- [Enabling Cell-to-Cell Operations](#)
If you have configured ASM-scoped security or DB-scoped security for your Oracle Exadata Database Machine, then you must configure access control to ensure direct cell-to-cell operations are not restricted.

- [Removing ASM-Scoped Security or DB-Scoped Security](#)

If you want to revert to an open security model, you must remove DB-scoped security for grid disks before removing ASM-scoped security.

4.4.1 About Exadata Storage Server Security Modes

Oracle Exadata System Software security allows open security, ASM-scoped security, or DB-scoped security.

Open Security


The default security mode is open security, where any database client has access to all the grid disks. Open security mode is useful for test or development databases where there are no security requirements. This is the default security mode after creating a new storage cell.

ASM-Scoped Security

Using ASM-scoped security you can restrict access to only the grid disks used by the Oracle ASM cluster. All Oracle Database instances associated with that Oracle ASM cluster have access to the disk groups and underlying grid disks. Grid disks used in disk groups belonging to a different Oracle ASM cluster are not be accessible to these instances.

Use ASM-scoped security when you want all databases and Oracle ASM instances in a cluster to have access to the grid disks that comprise the Oracle ASM disk groups used by the cluster. This includes the case when there is only one database in an Oracle ASM cluster.


Figure 4-1 ASM-Scoped Security



DB-Scoped Security

Using DB-scoped security, you can restrict access for an Oracle Database instance to a specific set of grid disks. The database instance must have access to all the grid disks used by the Oracle ASM disk groups for the database. The grid disks used by these Oracle ASM disk groups cannot be used by any other Oracle ASM disk group.

DB-scoped security mode is appropriate when multiple databases are accessing the grid disks. DB-scoped security allows you to limit database access to only the grid disks that are used by the Oracle ASM disk groups associated with the database.



4.4.2 Best Practices for ASM-Scoped Security and DB-Scoped Security

While setting up security, it is imperative that the configuration is the same across all the storage servers.

To have consistent Oracle Exadata System Software security, ensure the following:

- Configure the same cell-side grid disk security for all grid disks that belong to the same Oracle ASM disk group to avoid confusion and errors.
- Ensure that all Oracle Real Application Clusters (Oracle RAC) servers in an Oracle ASM cluster have the same content, ownership, and security for the Oracle ASM cellkey.ora file.
- Ensure that all Oracle RAC servers in a database cluster have the same content, ownership, and security for the database cellkey.ora file.
- If DB-scoped security is implemented, then ensure it is implemented for all databases accessing the grid disks. Do not mix ASM-scoped security and DB-scoped security for any grid disks.
- Use the `dcli` utility to make configuration changes to ensure consistency and eliminate potential user errors.

Related Topics

- Using the `dcli` Utility

4.4.3 About Security Keys

Oracle Exadata System Software uses keys to identify clients and determine access rights to the grid disks.

To determine which clients have access to a grid disk, a key is generated using CellCLI and stored in a read-only file that is accessible only by the clients. The CellCLI `CREATE KEY` command generates a random hexadecimal string for use as a security key. This key is stored in a `cellkey.ora` file on the client side, and assigned to the targets on the storage servers using the `ASSIGN KEY` command.

Note

The client name or Oracle ASM cluster name not case-sensitive. For example, ASM1 and asm1 are treated as the same value.

The CREATE KEY command can be run on any cell. You should only run this command when you need to create a new unique key. If you are configuring ASM-scoped security, then only a single security key is needed for each Oracle ASM cluster. If you are configuring DB-scoped security, then one key is needed for each Oracle ASM cluster and an additional security key is needed for each database client.

For ASM-scoped security, the cellkey.ora file is only accessible by the Oracle ASM cluster and its database clients. For DB-scoped security multiple security keys are used. One key is assigned to the Oracle ASM cluster, and one key is assigned to each database client. These security keys are stored in separate cellkey.ora files, and each file is accessible only by the client.

The cellkey.ora file contains entries that configure security among Oracle ASM, database clients and cells. The key and asm values in the cellkey.ora files on the Oracle ASM and database host computers must match the values assigned to the clients on the cells.

The cellkey.ora file contains the following fields:

- key — This field is required.
 - For ASM-scoped security, this key must match the value of the key assigned to the Oracle ASM cluster with the CellCLI ASSIGN KEY command.
 - For DB-scoped security, this key must match the value of the key assigned to the database client with the CellCLI ASSIGN KEY command.
- asm — This field is required.

This field contains a unique identifier for each Oracle ASM cluster. This value must be unique among all the clusters that access the storage servers in the storage grid. This value is used in the ASSIGN KEY command run on each cell, and the ALTER GRIDDISK and CREATE GRIDDISK commands used to configure grid disks to allow access from only a specific Oracle ASM cluster.

Note

The asm field is used when configuring both ASM-scoped security and DB-scoped security.


Access to the cellkey.ora file is controlled by operating system privileges.

- For Oracle ASM clients, the file is stored in the /etc/oracle/cell/network-config directory, and is readable by the Oracle Grid Infrastructure software owner and the group to which the user belongs.
- For DB-scoped security, there is a cellkey.ora for Oracle ASM, readable only by the Oracle Grid Infrastructure software owner, and additional cellkey.ora files for each database client. For the database clients, the file is stored in the pfile directory of the Oracle home directory for each database. The cellkey.ora file for a database is readable only by the operating system user that owns the Oracle Database software installation.

Note

The Oracle Grid Infrastructure and Oracle Database software must be owned by different operating system users to implement DB-scoped security.

Figure 4-2 Security Keys and cellkey.ora Files



On the storage servers, you use the `ASSIGN KEY` command to store the security keys in an access control list (ACL) for the cells. You use the `ALTER GRIDDISK` command to set the access rights for individual grid disks with the `AvailableTo` attribute. In order to access a grid disk, the security key of the cell ACL must match the security key of the client and the unique name of the client must be included in the `AvailableTo` attribute of the grid disk.

The identifying names used for the Oracle ASM and database must be unique. However, in some environments, there is more than one database with the same value for the `DB_UNIQUE_ID`. Each database uses a different Oracle ASM cluster for storage. Starting with Oracle Exadata System Software release 12.2.1.1.0, you can define ASM-scoped security based on Oracle ASM clusters. You use the `ASMCLUSTER` keyword with the `ASSIGN KEY` command. When you use the `ASMCLUSTER` keyword, the database name is qualified by the Oracle ASM unique identifier, creating a unique ID for each database, even if they have the same `DB_UNIQUE_ID`. Within each Oracle ASM cluster, the database names still have to be unique.

If you configure ASM-scoped security or DB-scoped security, then you also need to configure security keys to enable cell-to-cell operations. You create a key for each storage server, or cell, and assign that key as a local key for the cell using the `ASSIGN KEY FOR LOCAL CELL` command. You then assign the keys for the other cells that perform cell-to-cell operations with the current cell using the `ASSIGN KEY FOR REMOTE CELL` command.

Example 4-5 Creating a Security Key for Exadata Storage Security

Use the following command on any storage server to generate a unique security key. This key will be used to configure ASM-scoped security or DB-scoped security.

```
CellCLI> CREATE KEY
66e12adb996805358bf82258587f5050
```

Example 4-6 cellkey.ora File Entries

This example shows the entries for the `cellkey.ora` file.

```
key=66e12adb996805358bf82258587f5050
asm=cluster1
```

Related Topics

- CREATE KEY
- ASSIGN KEY

4.4.4 Setting Up ASM-Scoped Security on Oracle Exadata Storage Servers

Configuring ASM-scoped security requires actions to be performed on both the database servers and storage servers.

The steps here are for configuring ASM-scoped security for a single Oracle ASM cluster and the databases that are clients of that Oracle ASM cluster. The examples in these steps assume that the Oracle ASM software owner is the user `grid`.

1. Shut down the Oracle Database and Oracle ASM instances.
2. On one of the storage servers, use the CellCLI `CREATE KEY` command to generate a random hexadecimal string. The command can be run on any storage server.

```
CellCLI> CREATE KEY

66e12adb996805358bf82258587f5050
```

3. Copy the key to the `cellkey.ora` file on one of the database servers.
 - a. Create a `cellkey.ora` in the home directory of the software owner, for example, `/home/grid/cellkey.ora`.
 - b. Using the format shown below, where `cluster1` is an alias for the Oracle ASM cluster, add the key to the `cellkey.ora` file.

The alias must be unique among all the clusters that access the storage servers. The Oracle Clusterware cluster name can be used if the clusters have unique names.

```
key=66e12adb996805358bf82258587f5050
asm=cluster1
```

4. Use the `ASSIGN KEY` command to assign the security key to the Oracle ASM cluster client on all the storage servers that you want the Oracle ASM cluster to access, for example:

```
CellCLI> ASSIGN KEY FOR ASMCLUSTER
'cluster1'='66e12adb996805358bf82258587f5050'
```

The above command must be repeated on all storage servers that you want the Oracle ASM cluster to access, or you can use a `dcli` command, as shown here:

```
dcli -g cell_group -l celladmin "cellcli -e \"ASSIGN KEY FOR ASMCLUSTER
'cluster1'=
'66e12adb996805358bf82258587f5050'\""
```

5. Configure security on the grid disks on all the storage servers that you want the Oracle ASM cluster to access.

You can configure the security key either when creating grid disks, or with the ALTER GRDDISK command.

- To create new grid disks with security configured, use a command similar to the following where `cluster1` is the unique name for the Oracle ASM client:

```
CellCLI> CREATE GRDDISK ALL PREFIX=sales, size=75G, -
          availableTo='cluster1'
```

- To change security on existing grid disks, use a command similar to the following **on every storage server**. In the following example, `cluster1` is the unique name for the Oracle ASM client. The grid disks specified in this command must include all the grid disks used by the Oracle ASM disk groups.

```
CellCLI> ALTER GRDDISK sales_CD_01_cell01,
sales_CD_02_cell01,          -
          sales_CD_03_cell01, sales_CD_04_cell01,
sales_CD_05_cell01,          -
          sales_CD_06_cell01          -
          availableTo='cluster1'
```

In the preceding commands, the alias used, for example, `cluster1`, is a unique and consistent value on all the storage servers and among all the clusters. When you specify a value for `availableTo`, then only the clients configured with the same alias in the `cellkey.ora` file have access to the grid disks.

6. Verify the key assignment on the storage servers.

```
CellCLI> LIST KEY
```

```
CellCLI> LIST GRDDISK ATTRIBUTES name,availableTo
```

7. Copy the `cellkey.ora` file created in Step 3 to the `/etc/oracle/cell/network-config` directory on each database server.

In the following command, `db_group` is a file that contains the names of the database servers that are clients of the Oracle ASM cluster.

```
dcli -g db_group -l grid -f /home/grid/cellkey.ora -d /etc/oracle/cell/
network-config/cellkey.ora
```

8. On each database server, configure the permissions for the `cellkey.ora`.

- If you use role-separated management for all Oracle software, then set the permissions on the file to be read-only by the owner and the group. Change the group ownership for the file to the `oinstall` group that contains the installation users for both Oracle Grid Infrastructure (`grid`) and Oracle Database (`oracle`).

You can use `dcli` to configure the file on multiple servers with a single command.

```
dcli -g db_group -l root chown grid:oinstall /etc/oracle/cell/
network-config/cellkey.ora
```

```
dcli -g dbs_group -l grid chmod 440 /etc/oracle/cell/network-config/  
cellkey.ora
```

- If you use a single installation user for all Oracle software, then set the permissions on the file to be read-only by the software owner (for example, oracle).

```
chown oracle:oinstall /etc/oracle/cell/network-config/cellkey.ora  
chmod 400 /etc/oracle/cell/network-config/cellkey.ora
```

9. On the database servers, restart Oracle Clusterware, the Oracle ASM instances, and the Oracle Database instances.

The commands to stop and start Oracle Clusterware are:

```
# crsctl stop crs  
# crsctl start crs
```

Then restart the Oracle ASM and Oracle Database instances if they are not started automatically.

10. Remove the temporary cellkey.ora file created in Step 3.

Related Topics

- CREATE KEY
- ASSIGN KEY
- LIST KEY
- ALTER GRDDISK
- CREATE GRDDISK
- Using the dcli Utility

4.4.5 Setting Up DB-Scoped Security on Oracle Exadata Database Machine

When configuring DB-scoped security, you perform actions on both the database and storage servers.

These steps are for configuring DB-scoped security for a single Oracle ASM cluster and the databases that are clients of that Oracle ASM cluster. The examples in these steps assume that the Oracle ASM software owner is the user `grid` and each Oracle Database home is owned by a different operating system user.

To configure DB-scoped security, the following conditions must exist:

- The combination of the Oracle Database client unique name and the Oracle ASM Cluster used by the database client must be unique in your environment.
- The Oracle Grid Infrastructure software installation and each Oracle Database software installation must be owned by different operating system users.
- Each database client must have distinct Oracle ASM disk groups. The grid disks used by Oracle ASM disk groups can only be assigned to one Oracle ASM disk group.

1. Get the `DB_UNIQUE_NAME` for each database. Names are case-sensitive.

```
SQL> SELECT db_unique_name FROM V$DATABASE;
```

2. Shut down the databases and Oracle ASM instances for the Oracle ASM cluster.
3. On one of the storage servers, use the CellCLI CREATE KEY command to generate a key for Oracle ASM and a key for each Oracle Database client.

The command can be run on any storage server.

```
CellCLI> CREATE KEY  
66e12adb996805358bf82258587f5050
```

```
CellCLI> CREATE KEY  
f3f21c625ff41ef479a1bb033e0839e5
```

```
CellCLI> CREATE KEY  
cf03b74de32a67a6c1ec87b9da72bd47
```

4. Copy the key for Oracle ASM to the cellkey.ora file on one of the database servers.
 - a. Create a cellkey.ora in the home directory of the Oracle Grid Infrastructure software owner, for example, /home/grid/cellkey.ora.
 - b. Using the format shown below, where asm1 is an alias for the Oracle ASM cluster, add the key to the cellkey.ora file.

The alias must be unique among all the clusters that access the storage servers. The Oracle Clusterware cluster name can be used if the clusters have unique names.

```
key=66e12adb996805358bf82258587f5050  
asm=asm1
```

5. Copy the key for each Oracle Database client to a cellkey.ora file in the user directory of the Oracle home software owner on one of the database servers.
 - a. Create a cellkey.ora in the home directory of each Oracle Database software owner, for example, /home/db1user/cellkey.ora and /home/db2user/cellkey.ora.
 - b. Add the security key for each database client to the cellkey.ora file.

Using the format shown here, where asm1 is an alias for the Oracle ASM cluster used by the database client. In this example, each database client uses the same Oracle ASM cluster.

```
$ cat /home/db1user/cellkey.ora  
key=f3f21c625ff41ef479a1bb033e0839e5  
asm=asm1  
$ cat /home/db2user/cellkey.ora  
key=cf03b74de32a67a6c1ec87b9da72bd47  
asm=asm1
```

Note

The asm field is used for both ASM-scoped security and DB-scoped security.

6. Use the ASSIGN KEY command to assign the security key for the Oracle ASM client on all the storage servers that you want the Oracle ASM cluster to access, for example:

```
CellCLI> ASSIGN KEY FOR 'asm1'='66e12adb996805358bf82258587f5050'
```

The above command must be repeated on all storage servers, or you can use a `dcli` command, as shown here:

```
dcli -g cell_group -l root "cellcli -e \"ASSIGN KEY FOR 'asm1'='66e12adb996805358bf82258587f5050'\""
```

7. Use the `ASSIGN KEY` command to assign the security key for each Oracle Database client on all the storage servers that contain grid disks used by the Oracle ASM disk groups of that database, for example:

```
CellCLI> ASSIGN KEY FOR 'db1'='f3f21c625ff41ef479a1bb033e0839e5'
```

```
CellCLI> ASSIGN KEY FOR 'db2'='cf03b74de32a67a6c1ec87b9da72bd47'
```

The above command must be repeated on all storage servers, or you can use a `dcli` command, as shown here, where `cell_group_db1` contains a list of the storage servers used by the first database and `cell_group_db2` contains the list of storage servers used by the second database:

```
dcli -g cell_group_db1 -l root "cellcli -e \"ASSIGN KEY FOR 'db1'='f3f21c625ff41ef479a1bb033e0839e5'\""
```

```
dcli -g cell_group_db2 -l root "cellcli -e \"ASSIGN KEY FOR 'db2'='cf03b74de32a67a6c1ec87b9da72bd47'\""
```

8. Configure security on the grid disks on all the storage servers that you want the database clients to access.

You can configure the security key either when creating grid disks, or with the `ALTER GRIDDISK` command.

- To create new grid disks with security configured, use a command similar to the following where `asm1` is the unique name for the Oracle ASM client, and `db1` and `db2` are the unique names for the database clients:

```
CellCLI> CREATE GRIDDISK data1_CD_00_cell01,data1_CD_01_cell01, -  
data1_CD_02_cell01 size=75G, availableTo='asm1,db1'
```

```
CellCLI> CREATE GRIDDISK reco1_CD_00_cell01,reco1_CD_01_cell01, -  
reco1_CD_02_cell01 size=75G, availableTo='asm1,db1'
```

```
CellCLI> CREATE GRIDDISK data2_CD_00_cell01,data2_CD_01_cell01, -  
data2_CD_02_cell01 size=75G, availableTo='asm1,db2'
```

```
CellCLI> CREATE GRIDDISK reco2_CD_00_cell01,reco2_CD_01_cell01, -  
reco2_CD_02_cell01 size=75G, availableTo='asm1,db2'
```

- To change security on existing grid disks, use a command similar to the following **on every storage server**. In the following example, `asm1` is the unique name for the Oracle ASM client, and `db1` and `db2` are the unique names for the database clients. The grid disks specified in this command must include all the grid disks used by the Oracle ASM disk groups for each database client.

```
CellCLI> ALTER GRIDDISK data1_CD_00_cell01, data1_CD_01_cell01, -  
data1_CD_02_cell01, reco1_CD_00_cell01, reco1_CD_01_cell01,
```



```
reco1_CD_02_cell01, -
availableTo='asm1,db1'
```

```
CellCLI> ALTER GRIDDISK data2_CD_00_cell01, data2_CD_01_cell02, -
data2_CD_02_cell01, reco2_CD_00_cell01, reco2_CD_01_cell01,
reco2_CD_02_cell01, -
availableTo='asm1,db2'
```

When you specify a value for `availableTo`, then only the clients configured with the same key assigned to that alias in their `cellkey.ora` file have access to the grid disks.

9. Verify the key assignment on the storage servers.

```
CellCLI> LIST KEY
asm1      66e12adb996805358bf82258587f5050
db1       f3f21c625ff41ef479a1bb033e0839e5
db2       cf03b74de32a67a6c1ec87b9da72bd47
```

```
CellCLI> LIST GRIDDISK ATTRIBUTES name,availableTo WHERE availableTo='''
DATA1_CD_00_cell01      asm1,db1
DATA1_CD_01_cell01      asm1,db1
DATA1_CD_02_cell01      asm1,db1
DATA2_CD_00_cell01      asm1,db2
DATA2_CD_01_cell01      asm1,db2
DATA2_CD_02_cell01      asm1,db2
...
```

10. Copy the `cellkey.ora` file created in step 4 to the `/etc/oracle/cell/network-config` directory on each database server.

In the following command, `dbs_group` is a file that contains the names of the database servers that are clients of the Oracle ASM cluster.

```
dcli -g dbs_group -l grid -f /home/grid/cellkey.ora -d /etc/oracle/cell/
network-config/cellkey.ora
```

11. On each database server, set the permissions for the `cellkey.ora` file for Oracle ASM to be read-only by the file owner and the group.

You can use `dcli` to configure the file permissions multiple servers with a single command. In the following example, `dbs_group` is a file that contains a list of all the database servers that are clients of the Oracle ASM cluster.

```
dcli -g dbs_group -l grid chmod 640 /etc/oracle/cell/network-config/
cellkey.ora
```

12. Copy each `cellkey.ora` file created in step 5 to its `$ORACLE_HOME/admin/db_unique_name/pfile` directory on each database server.

In the following command, `db1_group` is a file that contains the names of the database servers that host the database instances for the `db1` database, and `dbuser1` is the operating system user that owns the `$ORACLE_HOME/admin/db1` directory.

```
dcli -g db1_group -l dbuser1 -f /home/dbuser1/cellkey.ora
-d $ORACLE_HOME/admin/db1/pfile/cellkey.ora
```

In this example, db2_group is a file that contains the names of the database servers that host the database instances for the database with the unique database name of db2, and dbuser2 is the operating system user that owns the \$ORACLE_HOME/admin/db2 directory.

```
dcli -g db2_group -l dbuser2 -f /home/dbuser2/cellkey.ora  
-d $ORACLE_HOME/admin/db2/pfile/cellkey.ora
```

13. On each database server, set the permissions for the cellkey.ora file in each Oracle database home to be accessible by only the file owner.

You can use dcli to configure the file permissions multiple servers with a single command. In the following example, db1_group is a file that contains the names of the database servers that host the database instances for the database with the unique database name of db1, and dbuser1 is the operating system user that owns the \$ORACLE_HOME/admin/db1 directory.

```
dcli -g db1_group -l dbuser1 chmod 600 $ORACLE_HOME/admin/  
db1/pfile/cellkey.ora
```

In this example, db2_group is a file that contains the names of the database servers that host the database instances for the database with the unique database name of db2, and dbuser2 is the operating system user that owns the \$ORACLE_HOME/admin/db2 directory.

```
dcli -g db2_group -l dbuser2 chmod 600 $ORACLE_HOME/admin/db2/pfile/  
cellkey.ora
```

14. On the database servers, restart Oracle Clusterware, the Oracle ASM instances, and the Oracle Database instances.

The commands to stop and start Oracle Clusterware are:

```
# crsctl stop crs  
# crsctl start crs
```

Then restart the Oracle ASM and Oracle Database instances if they are not started automatically by Oracle Clusterware.

15. Remove the temporary cellkey.ora files created in step [4](#) and step [5](#).

Related Topics

- CREATE KEY
- ASSIGN KEY
- LIST KEY
- ALTER GRIDDISK
- CREATE GRIDDISK
- Using the dcli Utility

4.4.6 Changing Security Keys for ASM-Scoped Security or DB-Scoped Security

You can change the keys used with ASM-scoped security or DB-scoped security.

- [Upgrading ASM-Scoped Security Key for ASMCLUSTER](#)
Starting with Oracle Exadata System Software release 12.2.1.1.0, you can define ASM-scoped security based on Oracle ASM clusters.
- [Changing the Assigned Key Value for ASM-Scoped Security](#)
You can change the key value used for an Oracle ASM client configured to use ASM-scoped security.
- [Changing the Assigned Key Value for DB-Scoped Security](#)
You can change the key value used by an Oracle Database client configured to use DB-scoped security.

4.4.6.1 Upgrading ASM-Scoped Security Key for ASMCLUSTER

Starting with Oracle Exadata System Software release 12.2.1.1.0, you can define ASM-scoped security based on Oracle ASM clusters.

The identifying names used for the Oracle ASM and database instances must be unique. However, in some environments, there is more than one database with the same value for the DB_UNIQUE_ID. If each database uses a different Oracle ASM cluster for storage, then you can use the ASMCLUSTER keyword when assigning the security key to specify that access should be limited to the specified Oracle ASM cluster.

When you use the ASMCLUSTER keyword, the database name is qualified by the Oracle ASM unique identifier, creating a unique ID for each database, even if they have the same DB_UNIQUE_ID. Within each Oracle ASM cluster, the database names still have to be unique.

Note

Do not use the ASMCLUSTER keyword when assigning keys if you are using DB-scoped security. Only use the ASMCLUSTER keyword for ASM-scoped security.

If you have ASM-scoped security configured, but want to change the keys to be ASMCLUSTER keys, perform the following steps:

1. Obtain the key value that you want to upgrade to an ASMCLUSTER key.

```
$ dcli -c dm01cel01,dm01cel02,dm01cel03 cellcli -e list key
dm01cel01:      asm1      118af47c57ab8da650ab67d5587fe728
dm01cel02:      asm1      118af47c57ab8da650ab67d5587fe728
dm01cel03:      asm1      118af47c57ab8da650ab67d5587fe728
```

2. Re-issue the `ASSIGN KEY` command using the same key value but adding the `ASMCLUSTER` keyword.

```
$ dcli -c dm01cel01,dm01cel02,dm01cel03 "cellcli -e assign key for
ASMCLUSTER
\ 'asm1\ '=\'118af47c57ab8da650ab67d5587fe728\'"
```

The name and key are removed from the ASM-scoped security list, and added as an Oracle ASM cluster client. Grid disks with this Oracle ASM client in their ACL can remain online for this operation.

3. Verify the keys have been upgraded to `ASMCLUSTER` keys.

```
$ dcli -c dm01cel01,dm01cel02,dm01cel03 cellcli -e list key
dm01cel01:      asm1      118af47c57ab8da650ab67d5587fe728      ASMCLUSTER
dm01cel02:      asm1      118af47c57ab8da650ab67d5587fe728      ASMCLUSTER
dm01cel03:      asm1      118af47c57ab8da650ab67d5587fe728      ASMCLUSTER
```

4.4.6.2 Changing the Assigned Key Value for ASM-Scoped Security

You can change the key value used for an Oracle ASM client configured to use ASM-scoped security.

1. Shut down the Oracle Database and Oracle ASM instances for which you are changing the security configuration.
2. Use the `CellCLI CREATE KEY` command to generate a random hexadecimal string. The command can be run on any cell.

```
CellCLI> CREATE KEY

f3d15c0c5e854345bcb3c2b678b1de45
```

3. Update the `cellkey.ora` file.
 - a. On one of the database servers, copy the `/etc/oracle/cell/network-config/cellkey.ora` file to the grid owner home directory, for example, `/home/grid/cellkey.ora`,
 - b. Update the file to use the new key for the Oracle ASM client.

```
key=f3d15c0c5e854345bcb3c2b678b1de45
asm=asm1
```

4. Copy the `cellkey.ora` file to `/etc/oracle/cell/network-config` on each database server, overwriting the existing file.

In the following command, `dbs_group` is a file that contains the names of the database servers that are clients of the Oracle ASM cluster, and `grid` is the software owner for the Oracle ASM installation.

```
dcli -g dbs_group -l grid -f /home/grid/cellkey.ora -d /etc/oracle/cell/
network-config/cellkey.ora
```

5. Make sure the permissions for the `cellkey.ora` in `/etc/oracle/cell/network-config` are configured so that only the Oracle Grid Infrastructure software owner has access to the file.

Make sure the owner of the file is the Oracle Grid Infrastructure software owner. If the current file permissions are not `rw-----` (600), then modify the permissions, as shown in the following command:

```
dcli -g dbs_group -l grid chmod 600 /etc/oracle/cell/network-config/cellkey.ora
```

6. Use the `ASSIGN KEY` command to update the security key assigned to the Oracle ASM cluster client on all the cells that contains grid disks used by the Oracle ASM cluster.

Use the same identifier and key value for the Oracle ASM client that is used in the `cellkey.ora` file on the database servers.

```
dcli -g cell_group -l root "cellcli -e \"ASSIGN KEY FOR  
'asm1'='f3d15c0c5e8543  
45bcb3c2b678b1de45'"
```

Starting with Oracle Exadata System Software release 12.2.1.1.0, add the `ASMCLUSTER` keyword to the `ASSIGN KEY` command if the security is based only on Oracle ASM clusters. Specify the Oracle ASM cluster name for the unique name for the key. For example:

```
CellCLI> ASSIGN KEY FOR ASMCLUSTER '+asm1' -  
        ='f3d15c0c5e854345bcb3c2b678b1de45'
```

Note

If you are using DB-scoped security, do not add the `ASMCLUSTER` keyword to the `ASSIGN KEY` command.

7. On the database servers, restart Oracle Clusterware, the Oracle ASM instances, and the Oracle Database instances.

The commands to stop and start Oracle Clusterware are:

```
# crsctl stop crs  
# crsctl start crs
```

Then restart the Oracle ASM and Oracle Database instances if they are not started automatically by Oracle Clusterware.

8. Remove the temporary `cellkey.ora` files created in step [3](#).

Related Topics

- [Upgrading ASM-Scoped Security Key for ASMCLUSTER](#)

Starting with Oracle Exadata System Software release 12.2.1.1.0, you can define ASM-scoped security based on Oracle ASM clusters.

4.4.6.3 Changing the Assigned Key Value for DB-Scoped Security

You can change the key value used by an Oracle Database client configured to use DB-scoped security.

1. Shut down the Oracle Database instances for which you are changing the security configuration.

2. Use the CellCLI CREATE KEY command to generate a random hexadecimal string.
This command can be run on any cell.

```
CellCLI> CREATE KEY

fa292e11b31b210c4b7a24c5f1bb4d32
```

3. Update the cellkey.ora file for the database client.
 - a. On one of the database servers, copy the \$ORACLE_HOME/admin/db_unique_name/pfile/cellkey.ora file to the user directory for the Oracle software owner, for example, /home/oradba1/cellkey.ora.
 - b. Update the file to use the new key for the database client.

```
key=fa292e11b31b210c4b7a24c5f1bb4d32
asm=asm1
```

4. Copy the cellkey.ora file to \$ORACLE_HOME/admin/db_unique_name/pfile on each database server, overwriting the existing file.

In the following command, dba1_group is a file that contains the names of the database servers that host instances of the database client, and oradba1 is the software owner for the Oracle Database installation.

```
dcli -g dba1_group -l oradba1 -f /home/oradba1/cellkey.ora -d $ORACLE_HOME/
admin/
db_unique_name/pfile/cellkey.ora
```

5. Make sure the permissions for the cellkey.ora in \$ORACLE_HOME/admin/db_unique_name/pfile are configured so that only the Oracle Database software owner has access to the file.

Make sure the owner of the file is the Oracle Database software owner. If the current file permissions are not rw- - - - - (600), then modify the permissions, as shown in the following command:

```
dcli -g dba1_group -l oradba1 chmod 600 $ORACLE_HOME/admin/db_unique_name/
pfile/cellkey.ora
```

6. Use the ASSIGN KEY command to update the security key assigned to the Oracle Database client on all the cells that contains grid disks used by the database.

Use the DB_UNIQUE_NAME for the database and the new key value for the database client.

```
dcli -g cell_group -l celladmin "cellcli -e \"ASSIGN KEY FOR
'db_unique_name'=
'fa292e11b31b210c4b7a24c5f1bb4d32'
```

7. On the database servers, restart the Oracle Database instances.
8. Remove the temporary cellkey.ora file created in step 3.

Related Topics

- [Setting Up DB-Scoped Security on Oracle Exadata Database Machine](#)
When configuring DB-scoped security, you perform actions on both the database and storage servers.

- [About Security Keys](#)
Oracle Exadata System Software uses keys to identify clients and determine access rights to the grid disks.

4.4.7 Enabling Cell-to-Cell Operations

If you have configured ASM-scoped security or DB-scoped security for your Oracle Exadata Database Machine, then you must configure access control to ensure direct cell-to-cell operations are not restricted.

To ensure cells have access to other cells when they need to communicate directly with one another, for example for offload operations, you need to set up **cell keys** for each cell.

- [Configuring Simple Cell Access](#)
You can use a single key for both inbound and outbound cell-to-cell traffic.
- [Configuring LOCAL and REMOTE Cell Keys](#)
You can configure each cell to have a unique key and to accept multiple remote cell keys for granting access.
- [Changing Between Simple Cell Keys and LOCAL and REMOTE Keys](#)
You must remove the existing cell keys before assigning any new keys with a different format. This protects you from inadvertently breaking your configuration by having different remote and local keys on the storage servers.

4.4.7.1 Configuring Simple Cell Access

You can use a single key for both inbound and outbound cell-to-cell traffic.

To ensure cells have access to other cells when they need to communicate directly with one another, for offload operations for example, you can create a single key. Assign that key to all cells using the simple version of the ASSIGN KEY command.

It is not necessary to use the ALTER GRIDDISK command to update the availableTo attribute for the cell key. The cells use the existing access control policy by making sure the originating client is identified at the remote target cell.

Perform these steps only if you have already configured ASM-Scoped Security or DB-Scoped Security.

1. Generate a random hexadecimal string.

This command can be run on any cell.

```
CellCLI> CREATE KEY  
fa292e11b31b210c4b7a24c5f1bb4d32
```

2. Assign the key to each cell.

```
CellCLI> ASSIGN KEY FOR CELL 'fa292e11b31b210c4b7a24c5f1bb4d32'
```

To update all cells with a single command, use dcli.

```
dcli -g mycells -l celladmin "cellcli -e assign key for cell  
'fa292e11b31b210c4b7a24c5f1bb4d32\''"
```

4.4.7.2 Configuring LOCAL and REMOTE Cell Keys

You can configure each cell to have a unique key and to accept multiple remote cell keys for granting access.

To ensure cells have access to other cells when they need to communicate directly with one another, for offload operations for example, you create cell keys. You can create a single key used by all cells, or you can assign a unique key to individual cells using the LOCAL and REMOTE keywords.

You might want to use multiple cell keys temporarily during rekeying, or if you want to limit access to specific cells. In this case, you need to specify LOCAL or REMOTE in the ASSIGN KEY command.

Perform these steps only if you have already configured ASM-Scoped Security or DB-Scoped Security.

1. Generate a random hexadecimal string for each cell.

These commands can be run on any cell.

```
Cellcli> CREATE KEY
          fa292e11b31b210c4b7a24c5f1bb4d3

Cellcli> CREATE KEY
          b67d5587fe728118af47c57ab8da650a

Cellcli> CREATE KEY
          118af47c57ab8da650ab67d5587fe728
```

2. On each cell, set the key for the local cell.

Specify a unique identifier for each cell, for example, cell01, cell02, and so on.

```
[celladmin@dm01cel01 ~]$ cellcli -e assign key for local cell -
'cell01'='fa292e11b31b210c4b7a24c5f1bb4d3'
[celladmin@dm01cel02 ~]$ cellcli -e assign key for local cell -
'cell02'='b67d5587fe728118af47c57ab8da650a'
[celladmin@dm01cel03 ~]$ cellcli -e assign key for local cell -
'cell03'='118af47c57ab8da650ab67d5587fe728'
```

3. Set the cell keys that the local cell will accept.

For each cell that you want to grant access to the local cell, use the ASSIGN KEY FOR REMOTE CELL command and specify the local key of that cell. You can specify any name for the remote keys.

```
[celladmin@dm01cel01 ~]$ cellcli -e assign key for remote cell -
'rcell02'='b67d5587fe728118af47c57ab8da650a', -
'rcell03'='118af47c57ab8da650ab67d5587fe728'

[celladmin@dm01cel02 ~]$ cellcli -e assign key for remote cell -
'rcell01'='fa292e11b31b210c4b7a24c5f1bb4d3', -
'rcell03'='118af47c57ab8da650ab67d5587fe728'
```



```
[celladmin@dm01cel03 ~]$ cellcli -e assign key for remote cell -
'rcell01'='fa292e11b31b210c4b7a24c5f1bb4d3', -
'rcell02'='b67d5587fe728118af47c57ab8da650a'
```

4. Verify that the keys are set on each storage server.

```
CellCLI> LIST KEY
db1          c25a62472a160e28bf15a29c162f1d74
asm1         118af47c57ab8da650ab67d5587fe728    ASMCLUSTER
cell01       fa292e11b31b210c4b7a24c5f1bb4d32    LOCAL-CELL
rcell02      b67d5587fe728118af47c57ab8da650a    REMOTE-CELL
rcell03      118af47c57ab8da650ab67d5587fe728    REMOTE-CELL
```

4.4.7.3 Changing Between Simple Cell Keys and LOCAL and REMOTE Keys

You must remove the existing cell keys before assigning any new keys with a different format. This protects you from inadvertently breaking your configuration by having different remote and local keys on the storage servers.

If you attempt to assign a LOCAL or REMOTE key to an existing cell key (that is, you have run the simple ASSIGN KEY FOR CELL command, and you want to switch to the explicit LOCAL or REMOTE keys), you will get the following error:

CELL-02911: Remove existing CELL key before assigning LOCAL CELL key

Similarly, if you attempt run ASSIGN KEY FOR CELL when local or remote cell keys already exist, you will get the following error:

CELL-02912: Remove all existing LOCAL and REMOTE CELL keys before assigning a CELL key.

Use LIST KEY FOR CELL to show all LOCAL and REMOTE CELL keys, then use ASSIGN KEY to assign an empty value to each.

You must remove all existing LOCAL and REMOTE cell keys before assigning a simple cell key.

1. View the configured keys on each storage server.
 - a. For simple cell keys, you would see results similar to the following:

```
CellCLI> LIST KEY
db1          c25a62472a160e28bf15a29c162f1d74
asm1         b4095e91d67bbf68d2e2fbe1fc50530f    ASMCLUSTER
cellkey      fa292e11b31b210c4b7a24c5f1bb4d32    CELL
```

- b. For local and remote cell keys, you would see results similar to the following:

```
CellCLI> LIST KEY
db1          c25a62472a160e28bf15a29c162f1d74
asm1         b4095e91d67bbf68d2e2fbe1fc50530f    ASMCLUSTER
cell01       fa292e11b31b210c4b7a24c5f1bb4d32    LOCAL-CELL
rcell02      b67d5587fe728118af47c57ab8da650a    REMOTE-CELL
rcell03      118af47c57ab8da650ab67d5587fe728    REMOTE-CELL
```

2. Remove the existing cell keys.

- a. For simple cell keys, use a command similar to the following:

```
dcli -g mycells -l celladmin "cellcli -e assign key for cell
'cell01'='' "
```

- b. For local and remote cell keys:

Remove the local key on each cell.

```
[celladmin@dm01cel01 ~]$ cellcli -e assign key for local cell
'cell01'=''
[celladmin@dm01cel02 ~]$ cellcli -e assign key for local cell
'cell02'=''
[celladmin@dm01cel03 ~]$ cellcli -e assign key for local cell
'cell03'=''
```

Run the following commands from any server to remove the remote keys on each cell:

```
dcli -g mycells -l celladmin "cellcli -e assign key for remote cell
'rcell01'='' "
dcli -g mycells -l celladmin "cellcli -e assign key for remote cell
'rcell02'='' "
dcli -g mycells -l celladmin "cellcli -e assign key for remote cell
'rcell03'='' "
```

3. Recreate the keys in the desired format.
 - a. To change simple cell keys to local and remote keys, follow the steps in [Configuring LOCAL and REMOTE Cell Keys](#)
 - b. To change local and remote cell keys to a simple cell key, follow the steps in [Configuring Simple Cell Access](#).

4.4.8 Removing ASM-Scoped Security or DB-Scoped Security

If you want to revert to an open security model, you must remove DB-scoped security for grid disks before removing ASM-scoped security.

Before making updates to the security on cells, you must shut down the Oracle Database and Oracle ASM instances. After all of the changes to security configuration are complete, start the Oracle Database and Oracle ASM instances.

- [Removing DB-Scoped Security](#)
- [Removing ASM-Scoped Security](#)

After you have removed DB-scoped security, you can remove ASM-scoped security if you want open security for the grid disks on the storage servers.

4.4.8.1 Removing DB-Scoped Security

To remove DB-scoped security on grid disks, perform the following procedure:

1. Shut down the Oracle Database and Oracle ASM instances.
2. Remove any database clients named in the availableTo attribute of the grid disks for which you want to remove DB-scoped security. [Example 4-7](#) provides examples of how to do this.

Note

If you removing DB-scoped security for a database client for only some of the grid disks, but not all, do not complete any further steps.

3. If a database client is not configured for security with any other grid disks, then you can remove the key assigned to the database client on the storage servers. Use the CellCLI ASSIGN KEY command.

In the following command, *db_client* is the name of the database client. To the right of the equal sign are two single quote characters with no space between them.

```
CellCLI> ASSIGN KEY FOR 'db_client'=''
```

Repeat this step on each storage server (cell) for which the database client no longer needs DB-scoped security.

4. On each database server, remove the cellkey.ora file located in the \$ORACLE_HOME/admin/db_unique_name/pfile/ directory for the database client.
5. Verify the key assignment has been updated on the storage servers.

```
CellCLI> LIST KEY
asm1      66e12adb996805358bf82258587f5050
db2       cf03b74de32a67a6c1ec87b9da72bd47
```

```
CellCLI> LIST GRIDDISK ATTRIBUTES name,availableTo WHERE availableTo=''
DATA1_CD_00_cell01      asm1
DATA1_CD_01_cell01      asm1
DATA1_CD_02_cell01      asm1
DATA2_CD_00_cell01      asm1,db2
DATA2_CD_01_cell01      asm1,db2
DATA2_CD_02_cell01      asm1,db2
...
```

6. Restart the Oracle Database and Oracle ASM instances.

Example 4-7 Removing Database Clients from Grid Disks

The following command removes all database clients from a group of grid disks:

```
CellCLI> ALTER GRIDDISK DATA1_CD_00_cell01,
DATA1_CD_01_cell01,          -
DATA1_CD_02_cell01, REC01_CD_00_cell01,
REC01_CD_01_cell01,          -
REC01_CD_02_cell01          -
availableTo='asm1'
```

If there are multiple database clients configured for DB-scoped security, for example db1, db2, and db3, and you only want to remove security for one client (db1), you can use a command similar to the following for a group of grid disks:

```
CellCLI> ALTER GRIDDISK sales_CD_04_cell01,
sales_CD_05_cell01, -
sales_CD_06_cell01, sales_CD_07_cell01,
sales_CD_08_cell01, -
sales_CD_09_cell01, -
availableTo='+asm,db2,db3'
```

The following example removes all database clients from all grid disks:

```
ALTER GRIDDISK ALL availableTo='+asm'
```

Note

If you want open security for the grid disks and storage servers, then you must remove ASM-scoped security after removing DB-scoped security.

Related Topics

- [Removing ASM-Scoped Security](#)
After you have removed DB-scoped security, you can remove ASM-scoped security if you want open security for the grid disks on the storage servers.

4.4.8.2 Removing ASM-Scoped Security

After you have removed DB-scoped security, you can remove ASM-scoped security if you want open security for the grid disks on the storage servers.

1. Shut down the Oracle Database and Oracle ASM instances.
2. Use the LIST KEY command to view the unique alias used for the Oracle ASM client.

Run this command on any storage server where ASM-scoped security is configured for the Oracle ASM client.

```
CellCLI> LIST KEY
asm1 66e12adb996805358bf82258587f5050
db2 cf03b74de32a67a6c1ec87b9da72bd47
```

3. Remove the Oracle ASM client named in the availableTo attribute of the grid disks for which you want to remove ASM-scoped security.

Use the alias from step 2. [Example 4-8](#) provides examples of how to do this.

4. If the Oracle ASM client is not configured for security with any other grid disks, then you can remove the key assigned to the Oracle ASM client on the storage servers. Use the CellCLI ASSIGN KEY command.
 - a. Determine if the Oracle ASM client in step 2 shows the ASMCLUSTER designation to the right of the key value.

- b. Use the ASMCLUSTER keyword in the ASSIGN KEY command only if the Oracle ASM client is listed as an ASMCLUSTER.

In the following command, *asm_cluster* is the alias from step 2. To the right of the equal sign are two single quote characters with no space between them.

```
CellCLI> ASSIGN KEY FOR ASMCLUSTER 'asm_cluster'=''
```

Run this command on all storage servers on which the key was assigned to the Oracle ASM client. You can alternatively use a command similar to the following:

```
dcli -g cell_group -l celladmin "cellcli -e \"ASSIGN KEY FOR ASMCLUSTER  
'asm_cluster'=''\\""
```

- 5. Update or delete the cellkey.ora file.

View the cellkey.ora file located in the /etc/oracle/cell/network-config/ directory on each database server.

If the Oracle ASM client for which you are removing ASM-scoped security is the only client listed in the file, then remove the cellkey.ora file on all servers with the same file contents.

If there is more than one Oracle ASM client listed in the cellkey.ora file, then perform the following steps:

- a. Remove the entry for the Oracle ASM client for which you are removing ASM-scoped security.
 - b. For all servers that list the Oracle ASM client in the cellkey.ora file, copy this file to those servers, or update the file on those servers.
- 6. Verify the key assignment has been updated on the storage servers.

```
CellCLI> LIST KEY
```

```
CellCLI> LIST GRIDDISK ATTRIBUTES name,availableTo WHERE availableTo != ''
```

- 7. Restart the Oracle Database and Oracle ASM instances.

Example 4-8 Removing the Oracle ASM Client from Grid Disks

The following command removes the Oracle ASM client from all grid disks on a cell:

```
CellCLI> ALTER GRIDDISK ALL availableTo=''
```

The following command removes the Oracle ASM client from all grid disks on a group of cells:

```
dcli -g cell_group -l celladmin "cellcli -e \"ALTER GRIDDISK ALL  
availableTo=''\\""
```

The following command removes the Oracle ASM client from a group of grid disks on a cell:

```
CellCLI> ALTER GRIDDISK sales_CD_01_cell01, sales_CD_02_cell01, -  
sales_CD_03_cell01, sales_CD_04_cell01, sales_CD_05_cell01, -
```

```
sales_CD_06_cell01  
availableTo= ' '
```

Related Topics

- ALTER GRIDDISK
- ASSIGN KEY

4.5 Maintaining a Secure Environment

After security measures are implemented, they must be maintained to keep the system secure.

Software, hardware and user access need to be updated and reviewed periodically. For example, organizations should review the users and administrators with access to Oracle Exadata Database Machine, and its deployed services to verify if the levels of access and privilege are appropriate. Without review, the level of access granted to individuals may increase unintentionally due to role changes or changes to default settings. It is recommended that access rights for operational and administrative tasks be reviewed to ensure that each user's level of access is aligned to their roles and responsibilities.

Organizations are encouraged to utilize tools to detect unauthorized changes, configuration drift, and prepare for security updates. Oracle Enterprise Manager provides an integrated solution for managing operational issues for hardware, deployed applications, and services.

- [Maintaining Network Security](#)
- [Encrypting System Log Information](#)
Management Server (MS) on database and storage servers supports the syslogconf attribute. Starting with Oracle Exadata System Software release 19.3.0, you can encrypt the log transfer.
- [Guarding Against Unauthorized Operating System Access](#)
AIDE is a utility that creates a database of files on the system, and then uses that database to ensure file integrity and to detect system intrusions.
- [Updating Software and Firmware](#)
Effective and proactive software management is a critical part of system security.
- [Ensuring Data Security Outside of Oracle Exadata Database Machine](#)
It is important to protect data stored outside of Oracle Exadata Database Machine, on backups or removed hard drives.

Related Topics

- [Responses to common Exadata security scan findings \(My Oracle Support Doc ID FAQ3926\)](#)
- *Oracle Exadata Maintenance Guide*

4.5.1 Maintaining Network Security

Follow these guidelines to ensure the security of local and remote access to the system:

- Network switch configuration files should be managed offline, and access to the configuration file should be limited to authorized administrators. The configuration file should contain descriptive comments for each setting. Consider keeping a static copy of the configuration file in a source code control system.

For more information on network switch configuration, refer to the vendor documentation for the network switch.

- Review the client access network to ensure that secure host and Integrated Lights Out Manager (ILOM) settings are in effect. Review the settings periodically to ensure that they remain intact.
- Create a login banner to state that unauthorized access is prohibited.
- Use access control lists to apply restrictions where appropriate.
- Set time-outs for extended sessions and set privilege levels.
- Use authentication, authorization, and accounting (AAA) features for local and remote access to a network switch.
- Use the port mirroring/switch port analyzer (SPAN) capability of the switch for intrusion detection system (IDS) access.
- Implement port security to limit access based upon a MAC address (MAC ACL).
- Limit remote configuration to specific IP addresses using SSH (IP ACL).
- Disable auto-trunking on all ports for any switch connected to Oracle Exadata Database Machine.
- Require users to use strong passwords by setting minimum password complexity rules and password expiration policies.
- Enable logging and send logs to a dedicated secure log host.
- Configure logging to include accurate time information, using NTP and timestamps.
- Review logs for possible incidents and archive them in accordance with the organization's security policy.
- Standard 140 of FIPS (Federal Information Processing Standards) relates to security and cryptography. FIPS 140 is a collection of standards published by NIST (National Institute of Standards and Technology), an agency of the United States federal government. FIPS 140 protects data during transit as well as at rest. It specifies security standards for cryptographic components within a computing environment. FIPS 140 is useful for organizations that need to document that their computing environment meets a published level of security. Many government agencies and financial institutions use FIPS 140 qualified systems.

Configuring FIPS 140 at the Oracle Database level enables the use of FIPS 140 cryptographic modules in the Secure Sockets Layer (SSL), transparent data encryption (TDE), DBMS_CRYPT PL/SQL package, and Exadata Smart Scan, which protects data while processing Smart Scan offload operations.

Related Topics

- *Oracle AI Database Security Guide*
- *Oracle AI Database Advanced Security Guide*
- *Oracle Exadata System Overview*

4.5.2 Encrypting System Log Information

Management Server (MS) on database and storage servers supports the `syslogconf` attribute. Starting with Oracle Exadata System Software release 19.3.0, you can encrypt the log transfer.

In the following topics, syslog and rsyslog are used interchangeably. They both refer to the message logger.

- [Overview of syslog File Encryption](#)
You can use certificates and the `syslogconf` attribute to configure encryption of the syslog information.
- [Configure CA Server and Central rsyslogd Server](#)
Before you can encrypt the syslog transfer, you must generate certificates and sign them by a host that acts as Certificate Authority (CA). This procedure only needs to be completed once.
- [Configure a Client for SYSLOG Encryption](#)
Configure the client so that it checks the server identity and sends messages only if the server identity is known.
- [Confirming Syslog Encryption is Enabled](#)
After configuring rsyslog encryption, you can perform basic checks to verify the encryption is working.

4.5.2.1 Overview of syslog File Encryption

You can use certificates and the `syslogconf` attribute to configure encryption of the syslog information.

The `syslogconf` attribute extends syslog rules for a database server. The attribute can be used to designate that syslog messages be forwarded to a specific remote rsyslogd service. On the MS, the forwarded messages are directed to a file, console, or management application, depending on the syslog configuration on the MS. This enables system logs from different servers to be aggregated and mined in a centralized logging server for security auditing, data mining, and so on.

The high-level steps required to enable rsyslog encryption are:

1. Setup a Certificate Authority (CA). This could be any node which has the `certtool` command. It is recommended to use a non-Exadata server. The CA creates a self-signed certificate. The certificate encryption key must be stored in a secure place. This certificate is used to sign other certificates.
2. Generate certificates on each participating node. If you do not have a central CA, then the Exadata administrator can generate both the private and public key on the CA and distribute copies to each trusted server. If you have a central CA, then the Exadata administrator generates the private key for each server.
3. If using a central CA, the Exadata administrator creates a certificate request. This request is then sent to the CA admin, which in turn generates the certificate (containing the public key). The CA admin then sends back the signed certificate to the Exadata administrator.
4. Install the signed certificates on each participating node. If using a central CA, the Exadata administrator installs the certificate signed by the CA. If you are not using a central CA, then the Exadata administrator installs a copy of private and public keys that were generated on the CA.
5. Setup a syslog central server. The central server needs `syslog.conf` setup. It also needs signed certificates.
6. Enable or disable the encryption of syslog on each client by using `CellCLI` or `DBMCLI`.

After completing these steps, the syslog chosen to be transported will be encrypted.

4.5.2.2 Configure CA Server and Central rsyslogd Server

Before you can encrypt the syslog transfer, you must generate certificates and sign them by a host that acts as Certificate Authority (CA). This procedure only needs to be completed once.

1. On the CA server, create a private key and certificate.

Follow steps in the rsyslog documentation for "Setting up the CA"

- a. Generate a private key.

```
# certtool--generate-privkey--outfile ca-key.pem --sec-param high
Generating a 3072 bit RSA private key...
```

- b. Create a self-signed certificate.

When prompted, supply the requested information about your organization for the certificate. Each certificate is valid for a specified period of time, after which you need to recreate all certificates. So you might want to use a long period, for example 3650 days (10 years).

To use this certificate to sign other certificates, when asked if this certificate belongs to an authority, you must specify Y. Also reply with a Y when asked:

- Is this certificate is a TLS web client (or server) certificate?
- Will the certificate be used for signing (DHE and RSA-EXPORT ciphersuites)?
- Will the certificate be used for encryption (RSA ciphersuites)?
- Will the certificate be used to sign other certificates?

```
# certtool--generate-self-signed --load-privkey ca-key.pem --outfile
ca.pem
Generating a self signed certificate...
Please enter the details of the certificate's distinguished name. Just
press enter to ignore a field.
Common name: common_name_for_CA
UID:
Organizational unit name: Org_unit_name
Organization name: Org_name
Locality name: CountyName_or_Locale
State or province name: state_prov
Country name (2 chars): Country_code
Enter the subject's domain component (DC):
This field should not be used in new certificates.
E-mail: CA_user_email_address
Enter the certificate's serial number in decimal (default:
6722248921586908930):
```

Activation/Expiration time.

The certificate will expire in (days): 3650

Extensions.Does the certificate belong to an authority? (y/N): Y

Path length constraint (decimal, -1 for no constraint):

Is this a TLS web client certificate? (y/N): Y

Will the certificate be used for IPsec IKE operations? (y/N):

Is this a TLS web server certificate? (y/N): Y

Enter a dnsName of the subject of the certificate: CA_hostname

Enter a dnsName of the subject of the certificate:

Enter a URI of the subject of the certificate:

Enter the IP address of the subject of the certificate: CA_IP_Address

Will the certificate be used for signing (DHE and RSA-EXPORT ciphersuites)? (Y/n): Y

Will the certificate be used for encryption (RSA ciphersuites)? (Y/n): Y

Will the certificate be used to sign OCSP requests? (y/N):

```

Will the certificate be used to sign code? (y/N):
Will the certificate be used for time stamping? (y/N):
Will the certificate be used to sign other certificates? (y/N): Y
Will the certificate be used to sign CRLs? (y/N):
Enter the URI of the CRL distribution point:
X.509 Certificate Information:
  Version: 3
  Serial Number (hex): 5d4a39c736f0bf02
  Validity:
    Not Before: Wed Aug 07 02:39:03 UTC 2019
    Not After: Sat Aug 04 02:39:03 UTC 2029
  Subject:
    CN=common_name_for_CA,OU=Org_unit_name,O=Org_name,L=CountyName_or_Locale
    ,ST=state_prov,C=Country_code,CA_user_email_address
  Subject Public Key Algorithm: RSA
  Algorithm Security Level: High (3072 bits)
  Modulus (bits 3072):

00:a5:b2:d6:5d:33:2c:79:2d:9c:79:f4:7b:0b:27:ef
20:29:ff:21:0c:19:11:22:c1:17:26:fc:46:5c:bb:c0
f6:9d:d0:ff:0d:4d:9e:25:18:62:53:8b:c6:4e:8b:05
...
03:21:7d:87:af:2b:a2:0b:42:ee:45:36:d7:14:aa:e8
6e:c1:25:4d:5d:66:db:fc:82:0c:92:69:66:04:70:a7
      5b
    Exponent (bits 24):
      01:00:01
  Extensions:
    Basic Constraints (critical):
      Certificate Authority (CA): TRUE
    Key Purpose (not critical):
      TLS WWW Client.
      TLS WWW Server.
    Subject Alternative Name (not critical):
      DNSname: CA_host_name
      IPAddress: CA_IP_Address
    Key Usage (critical):
      Digital signature.
      Key encipherment.
      Certificate signing.
    Subject Key Identifier (not critical):
      2b3c1e34e5a0347b6e62fd893430fa0b20d2d0a3
Other Information:
  Public Key ID:
    2b3c1e34e5a0347b6e62fd893430fa0b20d2d0a3
  Public key's random art:
    +--[ RSA 3072 ]-----+
    |          |
    | .        |
    | . 0 0 . . |
    | + 0 + +   |
    | E . = + S |

```

```

|0. . 0 . . |
| 0 0 @ . |
| + = B . |
| 0.0 0 |
+-----+

```

Is the above information ok? (y/N): y

Signing certificate...

#

c. Secure the ca-key.pem file.

Place the file in a secure place. Ensure that no user except root can access the certificates (not even read permissions).

```
# chmod 600 ca-key.pem
```

2. Generate the machine certificate.

Follow the step in the rsyslog document for "Generating the machine certificate".

```
# certtool --generate-privkey --outfile machine-key.pem --sec-param high
Generating a 3072 bit RSA private key...
```

This command can be run by the Exadata administrator. The output file is sent to the CA.

```
# certtool --generate-request --load-privkey machine-key.pem --outfile
request.pem
Generating a PKCS #10 certificate request...
Common name: Trusted_server
Organizational unit name: Org_unit_name
Organization name: Org_name
Locality name: CountryName_or_Locale
State or province name: state_prov
Country name (2 chars): Country_code
Enter the subject's domain component (DC):
UID:
Enter a dnsName of the subject of the certificate: Trusted_server_hostname
Enter a dnsName of the subject of the certificate:
Enter a URI of the subject of the certificate:
Enter the IP address of the subject of the certificate: Trusted_server_IP
Enter the e-mail of the subject of the certificate:
Enter a challenge password:
Does the certificate belong to an authority? (y/N):
Will the certificate be used for signing (DHE and RSA-EXPORT
ciphersuites)? (Y/n): Y
Will the certificate be used for encryption (RSA ciphersuites)? (Y/n): Y
Will the certificate be used to sign code? (y/N):
Will the certificate be used for time stamping? (y/N):
Will the certificate be used for IPsec IKE operations? (y/N):
Will the certificate be used to sign OCSP requests? (y/N):
Is this a TLS web client certificate? (y/N): Y
Is this a TLS web server certificate? (y/N): Y
```

This command is run by the CA administrator, using the request generated by the previous command. Review this example to see how to answer each prompt.

```
#certtool --generate-certificate --load-request request.pem --outfile
machine-cert.pem
```

```
--load-ca-certificate ca.pem --load-ca-privkey ca-key.pem
```

```
Generating a signed certificate...
```

```
Enter the certificate's serial number in decimal (default:
6722252284267403216):
```

```
Activation/Expiration time.
```

```
The certificate will expire in (days): 3650
```

```
Extensions.
```

```
Do you want to honour the extensions from the request? (y/N):
```

```
Does the certificate belong to an authority? (y/N):
```

```
Is this a TLS web client certificate? (y/N): y
```

```
Will the certificate be used for IPsec IKE operations? (y/N):
```

```
Is this a TLS web server certificate? (y/N): y
```

```
Enter a dnsName of the subject of the certificate: Trusted_server
```

```
Enter a dnsName of the subject of the certificate:
```

```
Enter a URI of the subject of the certificate:
```

```
Enter the IP address of the subject of the certificate:
```

```
Trusted_server_IP_addr
```

```
Will the certificate be used for signing (DHE and RSA-EXPORT
ciphersuites)? (Y/n): y
```

```
Will the certificate be used for encryption (RSA ciphersuites)? (Y/n): y
```

```
Will the certificate be used to sign OCSP requests? (y/N):
```

```
Will the certificate be used to sign code? (y/N):
```

```
Will the certificate be used for time stamping? (y/N):
```

```
X.509 Certificate Information:
```

```
Version: 3
```

```
Serial Number (hex): 5d4a3cd6265117d0
```

```
Validity:
```

```
Not Before: Wed Aug 07 02:52:06 UTC 2019
```

```
Not After: Sat Aug 04 02:52:06 UTC 2029
```

```
Subject:
```

```
OU=Org_unit_name,O=Org_name,L=CountryName_or_Locale,ST=state_prov,C=Country
_code
```

```
Subject Public Key Algorithm: RSA
```

```
Algorithm Security Level: High (3072 bits)
```

```
Modulus (bits 3072):
```

```
00:cf:f6:44:d4:e0:a8:b5:e6:48:8b:26:cb:59:c4:47
```

```
c5:f7:03:5f:99:88:ac:ed:94:d4:90:92:e4:61:75:4c
```

```
67:c4:16:c2:bf:31:40:f4:92:1e:94:73:08:d1:d5:3a
```

```
...
```

```
14:2f:08:02:74:f2:43:40:37:29:bd:6e:92:a6:07:6e
```

```
99:1e:e5:67:b8:0c:eb:a7:3d:9b:a5:35:46:8c:d3:e4
```

```
f7
```

```
Exponent (bits 24):
```

```
01:00:01
```

```

Extensions:
    Basic Constraints (critical):
        Certificate Authority (CA): FALSE
    Key Purpose (not critical):
        TLS WWW Client.
        TLS WWW Server.
    Subject Alternative Name (not critical):
        DNsname: Trusted_server
        IPAddress: Trusted_server_IP_addr
    Key Usage (critical):
        Digital signature.
        Key encipherment.
    Subject Key Identifier (not critical):
        7c343773a33cdbc6113fd05b3418ad129e9c4a64
    Authority Key Identifier (not critical):
        2b3c1e34e5a0347b6e62fd893430fa0b20d2d0a3

Other Information:
    Public Key ID:
        7c343773a33cdbc6113fd05b3418ad129e9c4a64
    Public key's random art:
    +--[ RSA 3072 ]-----+
    |                .+..|
    |                E . .0.|
    |                o = B.=..|
    |                . o X *.+o|
    |                S o = .o.|
    |                o  = ..|
    |                . + |
    |                . |
    +-----+

Is the above information ok? (y/N): y
Signing certificate...

```

3. Configure the rsyslogd server.

Install the certificates on the designated rsyslogd server. The server needs machine-cert.pem, machine-key.pem, and a copy of ca.pem. Add these certificates to the /etc/pki/rsyslog/rsyslog.conf file.

Ensure that no user except root can access the certificates (not even read permissions).

```
# chmod 600 cert_name.pem
```

Configure the server so that it accepts messages from all machines in your domain that have certificates from your CA. In this setup, you can use wildcards to ease adding new systems. Using wildcards permits the server to accept messages from systems whose names match *.domain. For example, if your domain is example.net, to allow permitted peers from different domain trees, you could use the following configuration:

```
$InputTCPServerStreamDriverPermittedPeer "*.example.net","*.example.com"
```

The following example shows a sample `/etc/pki/rsyslog/rsyslog.conf` file for the rsyslogd central server. This example configures the rsyslogd server to accept messages from any server on port 10514.

```
$ModLoad imtcp
# make gtls driver the default and set certificate files
$DefaultNetstreamDriver="gtls"
$DefaultNetstreamDriverCAFile="/etc/pki/rsyslog/ca.pem"
$DefaultNetstreamDriverCertFile="/etc/pki/rsyslog/machine-cert.pem"
$DefaultNetstreamDriverKeyFile="/etc/pki/rsyslog/machine-key.pem")

$InputTCPServerStreamDriverAuthMode x509/name
$InputTCPServerStreamDriverPermittedPeer *
$InputTCPServerStreamDriverMode 1 # run driver in TLS-only mode
$InputTCPServerRun 10514 # start up listener at port 10514
```

4. Restart the rsyslog process.

```
#service rsyslog stop

#service rsyslog start
```

4.5.2.3 Configure a Client for SYSLOG Encryption

Configure the client so that it checks the server identity and sends messages only if the server identity is known.

This configuration prevents malicious actors gaining access to the syslog data. These steps need to be performed on each server running the syslog client.

Before starting this task, you must have completed the steps in [Configure CA Server and Central rsyslogd Server](#). You will need the IP address and port number for the rsyslogd central server in this procedure.

1. Copy the `ca.pem`, `machine-key.pem` and `machine-cert.pem` certificates from the central rsyslogd server to the `/etc/pki/rsyslog/` directory on the client server.
2. Use CellCLI or DBMCLI to modify the `syslogconf` attribute on the client server.

The CellCLI or DBMCLI command appends the value you specify for `syslogconf` to the `rsyslog.conf` file and restarts the syslog process.

For a storage server client, you would use a command similar to the following:

```
ALTER CELL syslogconf=(' $DefaultNetstreamDrivergtls', \
'$DefaultNetstreamDriverCAFile /etc/pki/rsyslog/ca.pem', \
'$DefaultNetstreamDriverCertFile /etc/pki/rsyslog/machine-cert.pem', \
'$DefaultNetstreamDriverKeyFile /etc/pki/rsyslog/machine-key.pem', \
'$ActionSendStreamDriverAuthMode x509/name', \
'$ActionSendStreamDriverPermittedPeer *', \
'$ActionSendStreamDriverMode 1', 'user.* @@rsyslogd_server_IP_address:port')
```

If you are configuring syslog encryption for a database server, then use DBMCLI and replace `ALTER CELL` with `ALTER DBSERVER` in the above command.

3. Verify the syslogconf attribute has been updated correctly.

```
CellCLI> LIST CELL ATTRIBUTES syslogconf
$DefaultNetstreamDriver          gtls
$DefaultNetstreamDriverCAFile    /etc/pki/rsyslog/ca.pem
$DefaultNetstreamDriverCertFile  /etc/pki/rsyslog/machine-
cert.pem
$DefaultNetstreamDriverKeyFile    /etc/pki/rsyslog/machine-
key.pem
$ActionSendStreamDriverAuthMode  x509/name
$ActionSendStreamDriverPermittedPeer *
$ActionSendStreamDriverMode      1
user.*
@@rsyslogd_server_IP_address:port
```

If you are configuring syslog encryption for a database server, then use DBMCLI and replace LIST CELL with LIST DBSERVER in the above command.

4. Repeat these steps for each client server that needs to encrypt the syslog information.

4.5.2.4 Confirming Syslog Encryption is Enabled

After configuring rsyslog encryption, you can perform basic checks to verify the encryption is working.

1. Validate the Syslog configuration.

For a database server, use the following command:

```
DBMCLI> ALTER DBSERVER VALIDATE SYSLOGCONF 'kern.info'
```

For a storage server, use the following command:

```
CellCLI> ALTER CELL VALIDATE SYSLOGCONF 'kern.info'
```

2. Check the message in /var/log/messages.
3. Check for rsyslog error messages in /var/log/messages.
4. To verify the messages are transmitted in encrypted form, use the tcpdump utility.

From the target server, use the following command:

```
% tcpdump -A src source-server-IP-address
```

The output from the tcpdump command should not be readable text.

4.5.3 Guarding Against Unauthorized Operating System Access

AIDE is a utility that creates a database of files on the system, and then uses that database to ensure file integrity and to detect system intrusions.

- [About Advanced Intrusion Detection Environment \(AIDE\)](#)
AIDE helps to track down which file has been affected in case the system was compromised.

- [Managing AIDE Components](#)
You can use the `exadataAIDE` utility to manage AIDE.
- [Adding Custom AIDE Rules](#)
You can instruct AIDE to not check for changes in specific directories during AIDE the metadata initialize step and also during the daily cron check.
- [Managing AIDE Alerts when Updating Exadata Software](#)
Software and hardware updates and installs tend to change the size and nature of operating system files. Therefore, you should re-generate the AIDE database after making changes.

4.5.3.1 About Advanced Intrusion Detection Environment (AIDE)

AIDE helps to track down which file has been affected in case the system was compromised.

AIDE runs a daily cron job that monitors the system for changes to files in specific directories. It takes a snapshot of all files in the system defined by rules specified in its configuration file. AIDE compares the current file with the snapshot of files taken previously. If any content changes in the snapshot file, AIDE automatically raises CRITICAL software alerts. AIDE uses the default alerting destination email and sends the alert email to the configured SMTP email address. The results of the daily AIDE scan are written to `/var/log/aide/aide.log`.

The file snapshot database created by AIDE is stored at `/var/lib/aide/aide.db.gz`. You can backup this file daily if you want to audit what happened on a given system on daily basis.

4.5.3.2 Managing AIDE Components

You can use the `exadataAIDE` utility to manage AIDE.

AIDE comes pre-configured with Exadata System Software release 19.1; you do not have to perform any setup tasks to use this feature.

exadataAIDE Syntax

The utility is located at `/opt/oracle.SupportTools/exadataAIDE`.

```
exadataAIDE [-s|-status] [-e|enable] [-d|disable] [-u|-update] [-h|help]
```

Description of syntax options:

- `-s[tatus]` : Print the current status of the AIDE daily cron job
- `-e[nable]` : Enable the AIDE daily cron job
- `-d[isable]` : Disable the AIDE daily cron job
- `-u[pdate]` : Update the AIDE database metadata and run the daily scan
- `-h[elp]` : Print the command syntax and help information
- Get the current status of the aide cron job.

```
exadataAIDE -status
```

- Disable the daily AIDE scan.

```
exadataAIDE -disable
```


- Enable the daily AIDE scan.

```
exadataAIDE -enable
```

- Update the AIDE database after making changes to the system.

```
exadataAIDE -update
```

4.5.3.3 Adding Custom AIDE Rules

You can instruct AIDE to not check for changes in specific directories during AIDE the metadata initialize step and also during the daily cron check.

1. Log in to the server or virtual machine as the root user.
2. Edit the file `/etc/aide.conf`.

Add the directories you want AIDE to skip during its scan. Prefix the directory path with an exclamation point.

```
# Ignore /opt/myapp directory content  
!/opt/myapp
```

3. Update the AIDE database metadata.

```
# /opt/oracle.SupportTools/exadataAIDE -u
```

AIDE will not raise any alerts for `/opt/myapp` directory content changes going forward.

4.5.3.4 Managing AIDE Alerts when Updating Exadata Software

Software and hardware updates and installs tend to change the size and nature of operating system files. Therefore, you should re-generate the AIDE database after making changes.

Use the following steps to reduce false alarms when updating software:

Note

Oracle Exadata Deployment Assistant (OEDA) has intelligence built-in to avoid false alerts when installing or updating software.

1. Disable AIDE monitoring.
2. Update the software on your system.
3. Re-enable AIDE monitoring.

```
exadataAIDE -disable
```

```
exadataAIDE -enable
```

4. Update the AIDE database with the recent file changes.

```
exadataAIDE -update
```

4.5.4 Updating Software and Firmware

Effective and proactive software management is a critical part of system security.

Security enhancements are introduced through new releases and software updates. Oracle recommends installing the latest release of the software, and all necessary security updates on the equipment. The application of Oracle recommended and security updates is a best practice for the establishment of baseline security.

Operating system and kernel updates for Oracle Exadata database servers and storage servers are delivered with Oracle Exadata System Software updates. Power distribution unit (PDU) firmware updates are handled separately from the software and other firmware updates. Ensure that the PDU is running the latest approved firmware for Oracle Exadata. As PDU firmware updates are not issued frequently, it is usually sufficient to check the PDU firmware release when upgrading Oracle Exadata System Software.

Note

Devices such as network switches that contain firmware may require patches and firmware updates.

- [Regenerate SSH Keys for ILOM Version 5](#)
On systems with ILOM version 5, you can create an SSH key with a key size of 3072 bits.

Related Topics

- About Updating Exadata Software

4.5.4.1 Regenerate SSH Keys for ILOM Version 5

On systems with ILOM version 5, you can create an SSH key with a key size of 3072 bits.

Previous versions of Integrated Lights Out Manager (ILOM) support a 1024 bit SSH key, while ILOM version 5 supports a 3072 bit SSH key.

If you upgrade an existing Exadata system to ILOM version 5, then the upgraded system preserves the original 1024 bit SSH key. To use a 3072 bit SSH key, you must manually regenerate the SSH key in the ILOM service processor by using the following command:

```
-> set /SP/services/ssh generate_new_key_type=rsa generate_new_key_action=true
```

After you regenerate the SSH key, you can query ILOM to report the public key value.

```
-> show /SP/services/ssh/keys/rsa
```

```
/SP/services/ssh/keys/rsa
```

```
Targets:
```

```
Properties:
```

```
  fingerprint = hex-id  
  fingerprint_algorithm = SHA1  
  length = 3072  
  privatekey = (Cannot show property)  
  publickey = public-key-value
```

Note

Systems originally deployed with ILOM version 5 use 3072 bit keys by default.

4.5.5 Ensuring Data Security Outside of Oracle Exadata Database Machine

It is important to protect data stored outside of Oracle Exadata Database Machine, on backups or removed hard drives.

Data located outside of Oracle Exadata Database Machine can be secured by backing up important data. The data should then be stored in an off-site, secure location. Retain the backups according to organizational policies and requirements.

When disposing of an old hard drive, physically destroy the drive or completely erase all the data on the drive. Deleting the files or reformatting the drive removes only the address tables on the drive. The information can still be recovered from a drive after deleting files or reformatting the drive. The Oracle Exadata Database Machine disk retention support option allows the retention of all replaced hard drives and flash drives, instead of returning them to Oracle.

The CellCLI command `DROP CELLDISK` includes an option to securely erase data by overwriting the data. If Oracle Exadata Storage Server drives contain sensitive data that needs to be erased for redeployment or another purpose, then the secure erase feature should be used on the storage cell. The `ERASE` option ensures that all data is overwritten with random data, and erased up to seven times. This ensures that the data cannot be recovered, and that the data is permanently erased.

Starting with Oracle Exadata System Software release 19.1.0, if you use `DROP CELLDISK` and select to erase disks using 1pass, 3pass, or 7pass method, Oracle Exadata System Software uses the better and faster Secure Eraser if supported by the underlying hardware.

5

Securely Erasing Oracle Exadata

If you are repurposing or removing an Oracle Exadata from your environment, it is critical to securely erase all the information on the servers.

Starting with Oracle Exadata System Software release 19.1.0, Secure Eraser is automatically started during re-imaging if the hardware supports Secure Eraser. This significantly simplifies the re-imaging procedure while maintaining performance. Now, when re-purposing a rack, you only have to image the rack and the secure data erasure is taken care of transparently as part of the process.

- [Overview of Secure Eraser](#)
Oracle Exadata System Software provides a secure erasure solution, called Secure Eraser, for every component within Oracle Exadata.
- [Securely Erasing Database Servers and Storage Servers](#)
- [Automatic Secure Eraser through PXE Boot](#)
In this procedure, you configure Secure Eraser to run automatically when you reboot the nodes.
- [Interactive Secure Eraser through PXE Boot](#)
On Exadata systems prior to Oracle Exadata X7-2, you can use Preboot Execution Environment (PXE) Boot when performing a Secure Eraser.
- [Interactive Secure Eraser through Network Boot](#)
Starting with Oracle Exadata X7-2, you can use EFI Network Boot when using Secure Eraser.
- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.
- [Resetting Network Switches and Power Distribution Units to Factory Default](#)
- [Actions After Using Secure Eraser](#)
After performing a secure erase, the system is ready for return or re-imaging.

5.1 Overview of Secure Eraser

Oracle Exadata System Software provides a secure erasure solution, called Secure Eraser, for every component within Oracle Exadata.

Secure Eraser is a comprehensive solution that covers all Oracle Exadatas (V2 or higher), including both 2-socket and 8-socket servers. The solution securely erases all data on both database servers and storage servers, and resets the internal network switches, the Ethernet switches, and the power distribution units back to factory defaults.

To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).

To achieve the best possible performance, secure erasure is performed in parallel at every layer on an Oracle Exadata. All Oracle Exadata Database Servers and Oracle Exadata Storage Servers are securely erased in parallel. Within a server, all device types (such as hard drives, flash devices, persistent memory and internal USBs) are securely erased in parallel. For each device type, all devices are further securely erased in parallel. This means that the total time to securely erase an entire rack is the same regardless of whether it's a quarter, half, or full rack, and that the total time should be approximately the time it takes to erase whichever component takes the longest time.

Secure Eraser automatically detects the hardware capability of a storage device and picks the best erasure method supported by the device. Cryptographic erasure is used whenever possible to provide better security and faster speed. The cryptographic erasure method used by Secure Eraser is fully compliant with the NIST SP-800-88r1 standard.

Secure Eraser comes with flexible options. The entire process can be completely automated without any user intervention. Or, you can choose to do it interactively and choose to erase specific types of storage devices.

Secure Eraser periodically generates a progress report every 10 seconds so that you can easily monitor the progress.

When the secure erasure is completed, a certificate is generated for each server with a list of devices that have been securely erased. By default, the certificate for each secure erasure is generated on the associated server under `/var/log/cellos/`.

The following figure shows a sample certificate from Secure Eraser.

Figure 5-1 Sample Certificate from Secure Eraser

Data Erasure Certificate

Start Time: 2025-08-27 21:51:58 UTC

End Time: 2025-08-27 21:54:19 UTC

Chassis Number: 2440XLN4HL

This is to certify that the components identified below have been securely erased in accordance with the applicable guidelines of NIST SP-800-88r1 standard. Secure erasure on storage devices has been further validated by random data sampling.

ID	Type	Model	Serial Number	Size	Erasure Level	Status
1	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR430200306P8AGN-2	3.09TB	Crypto Erase	Succeeded
2	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR430200306P8AGN-1	3.09TB	Crypto Erase	Succeeded
3	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303005W6P8AGN-2	3.09TB	Crypto Erase	Succeeded
4	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303005W6P8AGN-1	3.09TB	Crypto Erase	Succeeded
5	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303007Q6P8AGN-2	3.09TB	Crypto Erase	Succeeded
6	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303007Q6P8AGN-1	3.09TB	Crypto Erase	Succeeded
7	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303003N6P8AGN-2	3.09TB	Crypto Erase	Succeeded
8	Flash	Flash Accelerator F680 v2 PCIe Card	PHCR4303003N6P8AGN-1	3.09TB	Crypto Erase	Succeeded
9	Flash	MZVL2480HBJD-00A07	S78UNN0X901420	447.13GB	Crypto Erase	Succeeded
10	Flash	MZVL2480HBJD-00A07	S78UNN0X901422	447.13GB	Crypto Erase	Succeeded
11	Disk	W7222A520ORA022T	0023147ZRXUS	20.00TB	Crypto Erase	Succeeded
12	Disk	W7222A520ORA022T	0023147E6GYS	20.00TB	Crypto Erase	Succeeded
13	Disk	W7222A520ORA022T	0023147E3UXS	20.00TB	Crypto Erase	Succeeded
14	Disk	W7222A520ORA022T	0023147DKGUS	20.00TB	Crypto Erase	Succeeded
15	Disk	W7222A520ORA022T	0023147DT4LS	20.00TB	Crypto Erase	Succeeded
16	Disk	W7222A520ORA022T	0023147B3BKZ	20.00TB	Crypto Erase	Succeeded
17	Disk	W7222A520ORA022T	0023147BPKXZ	20.00TB	Crypto Erase	Succeeded
18	Disk	W7222A520ORA022T	0023147E46ZS	20.00TB	Crypto Erase	Succeeded
19	Disk	W7222A520ORA022T	0023147E4U9S	20.00TB	Crypto Erase	Succeeded
20	Disk	W7222A520ORA022T	00231476KESZ	20.00TB	Crypto Erase	Succeeded
21	Disk	W7222A520ORA022T	0023147BNZYZ	20.00TB	Crypto Erase	Succeeded
22	Disk	W7222A520ORA022T	0023147BNTHZ	20.00TB	Crypto Erase	Succeeded

ERASURE PERFORMED BY:

Name _____

Date _____

Related Topics

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#)

5.2 Securely Erasing Database Servers and Storage Servers

Oracle Exadata System Software includes the Secure Eraser utility, which securely erases data on hard drives, flash devices, persistent memory, and internal USBs. It also resets ILOM to factory settings.

Secure Eraser sanitizes all content, not only user data but also operating system, Oracle Exadata System Software, and user configurations.

⚠ Caution

The server will become unbootable after the system devices are securely erased, and ILOM will no longer be remotely accessible after being reset to factory default. ILOM will remain accessible through serial console.

The Secure Eraser utility works on both database servers and storage servers and covers all supported Oracle Exadata systems.

Based on hardware capabilities, different secure erasure methods are applied. In general, Secure Eraser has two types of erasure methods: 3-pass erase and crypto erase. The 3-pass erase method overwrites all addressable locations with a character, its complement, then a random character, and finally verifies the results. The crypto erase method erases all user data present on instant secure erase (ISE) devices by deleting the encryption keys with which the user data was previously encrypted. The crypto erase method is used by default on all hardware devices that support it.

Refer to the table "Estimated Erasure Times for Disks by Erasure Method" in the topic DROP CELL for a summary of the secure erasure methods used and their approximate time. Note that the time for 3-pass erase varies from drives to drives based on their size and speed. It is approximately equal to the time required to overwrite the entire device three times and read it one more time. Hard drives, flash devices, persistent memory, and internal USBs are securely erased in parallel: the time required to erase one device is the same as that required for erasing multiple devices of the same kind.

5.3 Automatic Secure Eraser through PXE Boot

In this procedure, you configure Secure Eraser to run automatically when you reboot the nodes.

📘 Note

Starting with Oracle Exadata System Software release 19.1.0, the Secure Eraser package (`secureeraser_label.zip`) contains ISO images instead of NFS images.

Use one of the following procedures, depending on your system:

- [Automatic Secure Eraser through PXE Boot for X7 and Later Systems](#)
In this procedure, you configure Secure Eraser to run automatically when you reboot Oracle Exadata X7-2 and later nodes.
- [Automatic Secure Eraser through PXE Boot for X6 and Earlier Systems](#)
In this procedure, you configure Secure Eraser to run automatically when you reboot the nodes.

Related Topics

- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

5.3.1 Automatic Secure Eraser through PXE Boot for X7 and Later Systems

In this procedure, you configure Secure Eraser to run automatically when you reboot Oracle Exadata X7-2 and later nodes.

Note

Starting with Oracle Exadata System Software release 19.1.0, the Secure Eraser package (`secureeraser_label.zip`) contains ISO images instead of NFS images.

Before you begin:

- Download and install the latest Secure Eraser package available for your system.
To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).
 - Make sure you have access to a Preboot Execution Environment (PXE) server where the nodes to be erased can boot from.
 - Make sure you have access to a NFS server that is accessible from all the nodes to be erased.
 - Make sure you have access to one of the nodes to be erased.
1. Copy the PXE image files `initrd (initrd-<version>)` and `kernel (vmlinux-<version>)` from the Secure Eraser package to the `/tftpboot` directory on the PXE server.
 2. Create a file containing the names of the database servers and storage servers you want to erase.

To generate this file, you can run the following command from one of the nodes to be erased, and verify the nodes in the files are the ones to be erased.

```
# ibhosts | awk '/S [0-9\\.\\,]*/ || /C [0-9\\.\\,]*/ {print $6}' |  
sed "s/\\/\\/g" > nodes_to_be_erased
```

If you only want to erase one server, enter the name of the server into the `nodes_to_be_erased` file, for example `Exa01celadm04`.

3. Copy the `dcli` utility from the Secure Eraser package and the `nodes_to_be_erased` file generated in step 2 to the PXE server.
4. Create a PXE configuration template called `pxe_cfg.template` to contain the following lines:

Note

In the following example, the following parameters must be updated to match your environment:

- `kernel` (the `vmlinux` file)
- `initrd` (the `initrd*.img` file)
- `logpath`

- For Oracle Exadata System Software 18c (18.1.0):

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-nfs-18.1.0.0.0-170915.1 dhcp pxe boot-from=uefi
    quiet loglevel=0 secureeraser bootarea=diagnostics
    console=ttyS0,115200n8
    logpath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-nfs-18.1.0.0.0-170915.1.img
    echo "Booting installation kernel"
}
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-iso-19.1.2.0.0-190111 dhcp pxe boot-from=uefi
    quiet loglevel=0 secureeraser bootarea=diagnostics
    console=ttyS0,115200n8
    logpath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-iso-19.1.2.0.0-190111.img
    echo "Booting installation kernel"
}
```

- The first line (default) identifies a menu entry that should be selected by default, after the timeout value specified by the second line.
- The third line (menuentry) represents the Linux kernel that will be used in the Secure Eraser environment.
- The fifth line (linuxefi) indicates the kernel is on an UEFI-based system. The `linuxefi` statement must be on a single line in the configuration file.
 - The `dhcp` option specifies to use DHCP to discover the `eth0` interface.

- The pxe option suppresses search for the image on virtual CD and USB devices.
 - The boot-from=uefi option indicates the system is booting from UEFI.
 - The quiet option disables excessive kernel log messages.
 - The loglevel=0 option suppresses non-critical kernel messages.
 - The secureeraser option indicates PXE boot will automatically trigger the Secure Eraser utility to sanitize all media installed on the node, including hard drives, flash devices, persistent memory, internal USBs, and ILOM.
 - The bootarea option indicates that the boot mode is diagnostic and not imaging install or rescue.
 - The console options indicate standard output and standard error messages are printed to both ILOM web console and serial console.
 - The logpath option specifies the NFS share directory where Secure Eraser will save the certificate.
- The seventh line (initrdefi) specifies the initrd file to load. In this case it is the initrd file copied over in step 1.

By default, the examples shown above cause Secure Eraser to erase all components. You can use `secureeraser-options` to specify command-line options for Secure Eraser to change the default behavior and securely erase certain components only. For example, to erase hard drives and USBs only during the PXE boot, the template would look like this for grub2 / Secure Boot on Oracle Exadata Database Machine X7 and later systems:

Note

In the following example, the following parameters must be updated to match your environment:

- kernel (the vmlinux file)
- initrd (the initrd*img file)
- logpath

- For Oracle Exadata System Software 18c (18.1.0):

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinuz-nfs-18.1.0.0-170915.1 stit dhcp pxe boot-from=
uefi quiet loglevel=0 secureeraser secureeraser-options="--hdd --usb"
bootarea=diagnostics console=ttyS0,115200n8 logpath=10.133.42
.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-nfs-18.1.0.0-170915.1.img
    echo "Booting installation kernel"
}
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinuz-iso-19.1.2.0.0-190111 stit dhcp pxe boot-from=
uefi quiet loglevel=0 secureeraser secureeraser-options="--hdd --usb"
bootarea=diagnostics console=ttyS0,115200n8 logpath=10.133.42
.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-iso-19.1.2.0.0-190111.img
    echo "Booting installation kernel"
}
```

- On the PXE server, use the template file to generate a PXE configuration file in the /tftpboot/pxelinux.cfg/ directory for each of the nodes to be erased.

The PXE configuration file name is the dash-separated MAC address of the node with the prefix 01-.

If the nodes to be erased are accessible, use the following steps to automatically generate a PXE configuration file for each node based on the template.

- Set up SSH equivalence with the nodes to be erased from the PXE server. The command will prompt for the root password of each node.

```
pxe_server# dcli -g nodes_to_be_erased -k -l root
```

- Create PXE configuration files, one for each node to be erased based on the configuration template.

```
pxe_server# dcli -g nodes_to_be_erased -l root "ip addr show eth0" |
awk '/link\/ether/ {print "01:"$3}' | sed "s:/-:/g" |
xargs -I {} cp pxe_cfg.template {}
```

If the nodes are not accessible, use the following step to generate a PXE configuration file for each node to be erased:

- Manually collect the MAC address of the eth0 interface from each node and write them into a text file called mac_addresses. Write one MAC address per line. For example:

```
00:10:e0:62:c4:fa
00:10:e0:62:c2:8a
00:10:e0:62:b8:7c
00:10:e0:62:b8:3a
00:10:e0:62:c6:bc
```

- Use the following command to create a list of PXE configuration files, one for each node to be erased based on the configuration template.

```
pxe_server# cat mac_addresses | sed "s:/-:/g;s/^/01-/g" |
xargs -I {} cp pxe_cfg.template {}
```

In both cases, you should get a list of PXE configuration files, one for each node to be erased. For example, if the MAC addresses of the nodes in a quarter rack are

00:10:e0:62:c4:fa, 00:10:e0:62:c2:8a, 00:10:e0:62:b8:7c, 00:10:e0:62:b8:3a, and 00:10:e0:62:c6:bc, you should get the following files:

```
01-00-10-e0-62-c4-fa
01-00-10-e0-62-c2-8a
01-00-10-e0-62-b8-7c
01-00-10-e0-62-b8-3a
01-00-10-e0-62-c6-bc
```

The files have the same content as the configuration template.

Check your specific PXE server requirements. Your PXE server may need slightly different names or settings.

6. Configure the nodes to boot from PXE and reboot the nodes.

If the nodes to be erased are accessible, run the following commands:

```
pxe_server# dcli -g nodes_to_be_erased -l root "ipmitool chassis bootdev
pxe"
```

```
pxe_server# dcli -g nodes_to_be_erased -l root "reboot"
```

If the nodes to be erased are not remotely accessible but the ILOMs are, use the following steps

a. Create a file called `iloms_to_be_reset` containing the names of ILOMs. For example:

```
db1-ilom
db2-ilom
cell1-ilom
cell2-ilom
cell3-ilom
```

b. Configure the nodes to boot from PXE through ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{} -U root chassis bootdev pxe
```

c. Reboot the nodes from ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{} -U root chassis power cycle
```

If neither host nor ILOM is remotely accessible, log into ILOM using a serial console and run the following commands

```
ILOM> set /HOST/boot_device=pxe
```

```
ILOM> reset /SYS
```

7. The Secure Eraser utility will be automatically called to sanitize all installed storage media, including hard drives, flash devices, persistent memory, and internal USBs, and to reset ILOM to factory default for all nodes in parallel.

Secure Eraser creates a file called `secureeraser_node_chassis_number_date_time.certificate` in the specified logpath location. `node_chassis_number` is the ID attribute of the storage server or database server in CellCLI or DBMCLI.

The file contains a progress report that is updated every 10 seconds. The progress report is also output to the console on each node. The following is an example of the progress report:

ID	Type	Model	Serial Number	Size	Status
1	Flash	Flash Accel F640 PCIe Card v2	PHLN8BQ6P4EGN-1	2.91TB	To Be Erased (0%)
2	Flash	Flash Accel F640 PCIe Card v2	PHLN8BQ6P4EGN-2	2.91TB	To Be Erased (0%)
3	Flash	Flash Accel F640 PCIe Card v2	PHLN8BL6P4EGN-2	2.91TB	To Be Erased (0%)
4	Flash	Flash Accel F640 PCIe Card v2	PHLN8AX6P4EGN-1	2.91TB	To Be Erased (0%)
5	Flash	Flash Accel F640 PCIe Card v2	PHLN8AX6P4EGN-2	2.91TB	To Be Erased (0%)
6	Flash	Flash Accel F640 PCIe Card v2	PHLN88S6P4EGN-1	2.91TB	To Be Erased (0%)
7	Flash	Flash Accel F640 PCIe Card v2	PHLN88S6P4EGN-2	2.91TB	To Be Erased (0%)
8	Flash	Flash Accel F640 PCIe Card v2	PHLN88S6P4EGN-1	2.91TB	To Be Erased (0%)
9	Flash	Flash Accel F640 PCIe Card v2	PHLN88S6P4EGN-2	2.91TB	To Be Erased (0%)
10	Flash	Flash Accel F640 PCIe Card v2	PHLN8BL6P4EGN-1	2.91TB	To Be Erased (0%)
11	Flash	Flash Accel F640 PCIe Card v2	PHLN88W6P4EGN-2	2.91TB	To Be Erased (0%)
12	Flash	Flash Accel F640 PCIe Card v2	PHLN88W6P4EGN-1	2.91TB	To Be Erased (0%)
13	Flash	Flash Accel F640 PCIe Card v2	PHLN89F6P4EGN-2	2.91TB	To Be Erased (0%)
14	Flash	Flash Accel F640 PCIe Card v2	PHLN8DQ6P4EGN-2	2.91TB	To Be Erased (0%)
15	Flash	Flash Accel F640 PCIe Card v2	PHLN89F6P4EGN-1	2.91TB	To Be Erased (0%)
16	Flash	Flash Accel F640 PCIe Card v2	PHLN88G6P4EGN-2	2.91TB	To Be Erased (0%)
17	M.2	INTEL SSDSCKKB24	PHYH88H240J	139.69GB	To Be Erased (0%)
18	M.2	INTEL SSDSCKKB24	PHYH84060035240J	139.69GB	To Be Erased (0%)
19	PM	NMA1XBD128GQS	8089-a2-0000028a	126.37GB	To Be Erased (0%)
20	PM	NMA1XBD128GQS	8089-a2-000002f4	126.37GB	To Be Erased (0%)
21	PM	NMA1XBD128GQS	8089-a2-000009d9	126.37GB	To Be Erased (0%)
22	PM	NMA1XBD128GQS	8089-a2-00000a27	126.37GB	To Be Erased (0%)
23	PM	NMA1XBD128GQS	8089-a2-00000231	126.37GB	To Be Erased (0%)
24	PM	NMA1XBD128GQS	8089-a2-0000039e	126.37GB	To Be Erased (0%)
25	PM	NMA1XBD128GQS	8089-a2-000006be	126.37GB	To Be Erased (0%)
26	PM	NMA1XBD128GQS	8089-a2-00000916	126.37GB	To Be Erased (0%)
27	PM	NMA1XBD128GQS	8089-a2-00000105	126.37GB	To Be Erased (0%)
28	PM	NMA1XBD128GQS	8089-a2-00000216	126.37GB	Being Erased (0%)
29	PM	NMA1XBD128GQS	8089-a2-00000151	126.37GB	Being Erased (0%)
30	PM	NMA1XBD128GQS	8089-a2-000002f5	126.37GB	To Be Erased (0%)
31	ILOM		1824XCA004		To Be Reset

As the sample progress report shows, Secure Eraser erases all storage devices in parallel. After the storage devices are securely erased, Secure Eraser will reset the ILOM to the factory default. This is to ensure that in the case that secure erasure fails on some storage device, the web console is still accessible for remote debugging, and ILOM is still accessible to control the host.

Once secure erasure is complete, a certificate called `secureeraser_node_chassis_number_date_time.certificate.pdf` is generated at the NFS share location specified by the `logpath` option in step 4. If secure erasure is successful, the nodes will be shut down automatically. If Secure Eraser does not succeed on some components, then the node will be left in diagnostic shell for further debugging. Assuming all previous steps are successful, and you have resolved the issue, you can go back to step 6 and rerun Secure Eraser.

Related Topics

- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

5.3.2 Automatic Secure Eraser through PXE Boot for X6 and Earlier Systems

In this procedure, you configure Secure Eraser to run automatically when you reboot the nodes.

Note

Starting with Oracle Exadata System Software release 19.1.0, the Secure Eraser package (`secureeraser_label.zip`) contains ISO images instead of NFS images.

Before you begin:

- Download and install the latest Secure Eraser package available for your system.
To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).
 - Make sure you have access to a Preboot Execution Environment (PXE) server where the nodes to be erased can boot from.
 - Make sure you have access to a NFS server that is accessible from all the nodes to be erased.
 - Make sure you have access to one of the nodes to be erased.
1. Copy the PXE image files `initrd (initrd-<version>)` and `kernel (vmlinix-<version>)` from the Secure Eraser package to the `/tftpboot` directory on the PXE server.
 2. Create a file containing the names of the database servers and storage servers you want to erase.

To generate this file, you can run the following command from one of the nodes to be erased, and verify the nodes in the files are the ones to be erased.

```
# ibhosts | awk '/S [0-9\\.\\,]*/ || /C [0-9\\.\\,]*/ {print $6}' |  
sed "s/\\/\\/g" > nodes_to_be_erased
```

If you only want to erase one server, enter the name of the server into the `nodes_to_be_erased` file, for example `Exa01celadm04`.

3. Copy the `dcli` utility from the Secure Eraser package and the `nodes_to_be_erased` file generated in step 2 to the PXE server.
4. Create a PXE configuration template called `pxe_cfg.template` to contain the following lines:

- For releases earlier than Oracle Exadata System Software release 19.1.0:

```
default linux
label linux
kernel vmlinux-nfs-12.2.1.1.0-161015-cell
append initrd=initrd-nfs-12.2.1.1.0-161015-cell.img dhcp pxe quiet
loglevel=
0 secureeraser bootarea=diagnostics console=tty1 console=ttyS0,115200n8
logp
ath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
default linux
label linux
kernel vmlinux-iso-19.1.2.0.0-190111-cell
append initrd=initrd-iso-19.1.2.0.0-190111-cell.img dhcp pxe quiet
loglevel=
0 secureeraser bootarea=diagnostics console=tty1 console=ttyS0,115200n8
logp
ath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
```

- The first line (default) indicates that the default label to use is called linux.
- The second line (label) defines the linux label.
- The third line (kernel) identifies the kernel file to load. In this case it is the file copied over in step [1](#).
- The fourth line (append) adds more options to the kernel command line. The append statement must be on a single line in the configuration file.
 - The `initrd` option specifies the `initrd` file to load. In this case it is the `initrd` file copied over in step [1](#).
 - The `dhcp` option specifies to use DHCP to discover the `eth0` interface.
 - The `pxe` option suppresses search for the image on virtual CD and USB devices.
 - The `quiet` option disables excessive kernel log messages.
 - The `loglevel=0` option suppresses non-critical kernel messages.
 - The `secureeraser` option indicates PXE boot will automatically trigger the Secure Eraser utility to sanitize all media including hard drives, flash devices, internal USBs, and ILOM on the node.
 - The `bootarea` option indicates that the boot mode is diagnostic and not imaging install or rescue.
 - The `console` options indicate standard output and standard error messages are printed to both ILOM web console and serial console.
 - The `logpath` option specifies the NFS share directory where Secure Eraser will save the certificate.

By default, the examples shown above cause Secure Eraser to erase all components. You can use `secureeraser-options` to specify command-line options for Secure Eraser to change the default behavior and securely erase certain components only. For example, to erase hard drives and USBs only during the PXE boot, the template would look like this:

- For releases earlier than Oracle Exadata System Software release 19.1.0:

```
default linux
label linux
kernel vmlinux-nfs-12.2.1.1.0-161015-cell
append initrd=initrd-nfs-12.2.1.1.0-161015-cell.img dhcp pxe quiet
loglevel=0
secureeraser secureeraser-options="--hdd --usb" bootarea=diagnostics
console=tty1 console=ttyS0,115200n8 logpath=10.133.42.221:/export/
exadata_
secure_eraser_certificate_dir
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
default linux
label linux
kernel vmlinux-iso-19.1.2.0.0-190111-cell
append initrd=initrd-iso-19.1.2.0.0-190111-cell.img dhcp pxe quiet
loglevel=0
secureeraser secureeraser-options="--hdd --usb" bootarea=diagnostics
console=tty1 console=ttyS0,115200n8 logpath=10.133.42.221:/export/
exadata_
secure_eraser_certificate_dir
```

- On the PXE server, use the template file to generate a PXE configuration file in the /tftpboot/pxelinux.cfg/ directory for each of the nodes to be erased.

The PXE configuration file name is the dash-separated MAC address of the node with the prefix 01-.

If the nodes to be erased are accessible, use the following steps to automatically generate a PXE configuration file for each node based on the template.

- Set up SSH equivalence with the nodes to be erased from the PXE server. The command will prompt for the root password of each node.

```
pxe_server# dcli -g nodes_to_be_erased -k -l root
```

- Create PXE configuration files, one for each node to be erased based on the configuration template.

```
pxe_server# dcli -g nodes_to_be_erased -l root "ip addr show eth0" |
awk '/link\|ether/ {print "01:"$3}' | sed "s:/-:/g" |
xargs -I {} cp pxe_cfg.template {}
```

If the nodes are not accessible, use the following step to generate a PXE configuration file for each node to be erased:

- Manually collect the MAC address of the eth0 interface from each node and write them into a text file called mac_addresses. Write one MAC address per line. For example:

```
00:10:e0:62:c4:fa
00:10:e0:62:c2:8a
00:10:e0:62:b8:7c
00:10:e0:62:b8:3a
00:10:e0:62:c6:bc
```


- b. Use the following command to create a list of PXE configuration files, one for each node to be erased based on the configuration template.

```
pxe_server# cat mac_addresses | sed "s/:/-/g;s/^/01-/g" |  
xargs -I {} cp pxe_cfg.template {}
```

In both cases, you should get a list of PXE configuration files, one for each node to be erased. For example, if the MAC addresses of the nodes in a quarter rack are 00:10:e0:62:c4:fa, 00:10:e0:62:c2:8a, 00:10:e0:62:b8:7c, 00:10:e0:62:b8:3a, and 00:10:e0:62:c6:bc, you should get the following files:

```
01-00-10-e0-62-c4-fa  
01-00-10-e0-62-c2-8a  
01-00-10-e0-62-b8-7c  
01-00-10-e0-62-b8-3a  
01-00-10-e0-62-c6-bc
```

The files have the same content as the configuration template.

Check your specific PXE server requirements. Your PXE server may need slightly different names or settings.

6. Configure the nodes to boot from PXE and reboot the nodes.

If the nodes to be erased are accessible, run the following commands:

```
pxe_server# dcli -g nodes_to_be_erased -l root "ipmitool chassis bootdev  
pxe"
```

```
pxe_server# dcli -g nodes_to_be_erased -l root "reboot"
```

If the nodes to be erased are not remotely accessible but the ILOMs are, use the following steps

- a. Create a file called `iloms_to_be_reset` containing the names of ILOMs. For example:

```
db1-ilom  
db2-ilom  
cell1-ilom  
cell2-ilom  
cell3-ilom
```

- b. Configure the nodes to boot from PXE through ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H  
{ } -U root chassis bootdev pxe
```

- c. Reboot the nodes from ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H  
{ } -U root chassis power cycle
```

If neither host nor ILOM is remotely accessible, log into ILOM using a serial console and run the following commands

```
ILOM> set /HOST/boot_device=pxe
```

```
ILOM> reset /SYS
```

7. The Secure Eraser utility will be automatically called to sanitize all storage media including hard drives, flash devices, and internal USBs, and to reset ILOM to factory default for all nodes in parallel.

Secure Eraser creates a file called `secureeraser_node_chassis_number_date_time.certificate` in the specified logpath location. `node_chassis_number` is the ID attribute of the storage server or database server in CellCLI or DBMCLI.

The file contains a progress report that is updated every 10 seconds. The progress report is also output to the console on each node. The following is an example of the progress report:

ID	Type	Model	Serial Number	Size
1	Flash	Flash Accelerator F80 PCIe Card	FL00A96H	200.00GB
Being Erased (6%)				
2	Flash	Flash Accelerator F80 PCIe Card	FL00A84Y	200.00GB
Being Erased (5%)				
3	Flash	Flash Accelerator F80 PCIe Card	FL00A7D4	200.00GB
Being Erased (5%)				
4	Flash	Flash Accelerator F80 PCIe Card	FL00A6WG	200.00GB
Being Erased (6%)				
5	Flash	Flash Accelerator F80 PCIe Card	FL008KSE	200.00GB
Being Erased (5%)				
6	Flash	Flash Accelerator F80 PCIe Card	FL008KS3	200.00GB
Being Erased (5%)				
7	Flash	Flash Accelerator F80 PCIe Card	FL008KL7	200.00GB
Being Erased (5%)				
8	Flash	Flash Accelerator F80 PCIe Card	FL008KQR	200.00GB
Being Erased (6%)				
9	Flash	Flash Accelerator F80 PCIe Card	FL00A812	200.00GB
Being Erased (5%)				
10	Flash	Flash Accelerator F80 PCIe Card	FL00A79G	200.00GB
Being Erased (5%)				
11	Flash	Flash Accelerator F80 PCIe Card	FL00A80C	200.00GB
Being Erased (6%)				
12	Flash	Flash Accelerator F80 PCIe Card	FL00A79F	200.00GB
Being Erased (6%)				
13	Flash	Flash Accelerator F80 PCIe Card	FL00A5WD	200.00GB
Being Erased (5%)				
14	Flash	Flash Accelerator F80 PCIe Card	FL00A5XS	200.00GB
Being Erased (5%)				
15	Flash	Flash Accelerator F80 PCIe Card	FL00A7N1	200.00GB
Being Erased (5%)				
16	Flash	Flash Accelerator F80 PCIe Card	FL00A62G	200.00GB
Being Erased (5%)				
17	Disk	H7240AS60SUN4.0T	1352E5XHWX	4.00TB
Being Erased (1%)				
18	Disk	H7240AS60SUN4.0T	1352E60SYX	4.00TB

```

Being Erased (1%)
19 Disk H7240AS60SUN4.0T 1352E60U4X 4.00TB
Being Erased (1%)
20 Disk H7240AS60SUN4.0T 1352E5UPAX 4.00TB
Being Erased (1%)
21 Disk H7240AS60SUN4.0T 1352E5XK3X 4.00TB
Being Erased (1%)
22 Disk H7240AS60SUN4.0T 1352E62M7X 4.00TB
Being Erased (1%)
23 Disk H7240AS60SUN4.0T 1352E5PSPX 4.00TB
Being Erased (1%)
24 Disk H7240AS60SUN4.0T 1352E60TJX 4.00TB
Being Erased (1%)
25 Disk H7240AS60SUN4.0T 1352E5LYDX 4.00TB
Being Erased (1%)
26 Disk H7240AS60SUN4.0T 1352E602WX 4.00TB
Being Erased (1%)
27 Disk H7240AS60SUN4.0T 1352E5LY9X 4.00TB
Being Erased (1%)
28 Disk H7240AS60SUN4.0T 1352E5VX4X 4.00TB
Being Erased (1%)
29 USB SSM 1900638EA8BFB749 8.00GB
Being Erased (5%)
30 ILOM
1403NM50CA To Be Reset

```

As the sample progress report shows, Secure Eraser erases all storage devices in parallel. After the storage devices are securely erased, Secure Eraser will reset the ILOM to the factory default. This is to ensure that in the case that secure erasure fails on some storage device, the web console is still accessible for remote debugging, and ILOM is still accessible to control the host.

Once secure erasure is complete, a certificate called `secureeraser_node_chassis_number_date_time.certificate.pdf` is generated at the NFS share location specified by the `logpath` option in step 4. If secure erasure is successful, the nodes will be shut down automatically. If Secure Eraser does not succeed on some components, then the node will be left in diagnostic shell for further debugging. Assuming all previous steps are successful, and you have resolved the issue, you can go back to step 6 and rerun Secure Eraser.

Related Topics

- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

5.4 Interactive Secure Eraser through PXE Boot

On Exadata systems prior to Oracle Exadata X7-2, you can use Preboot Execution Environment (PXE) Boot when performing a Secure Eraser.

Note

Starting with Oracle Exadata System Software release 19.1.0, the Secure Eraser package (`secureeraser_label.zip`) contains ISO images instead of NFS images.

Before you begin:

- Download and install the latest Secure Eraser package available for your system.
To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).
- Make sure you have access to a PXE server where the nodes to be erased can boot from.
- Make sure you have access to a NFS server that is accessible from all the nodes to be erased.
- Make sure you have access to one of the nodes to be erased.

1. Copy the cell PXE image files `initrd` (`initrd-version`) and `kernel` (`vmlinux-version`) from the Secure Eraser package to the `/tftpboot` directory on the PXE server. For Oracle Exadata Database Machine X7 and later systems, the directory is `/tftpboot/efi` for `grub2` and Secure Boot.
2. Create a file containing the names of the database servers and storage servers to be erased.

To generate this file, you can run the following command from one of the nodes to be erased, and verify the nodes in the files are the ones to be erased.

```
# ibhosts | awk '/S [0-9\\.\\,]*/ || /C [0-9\\.\\,]*/ {print $6}' | sed
"s/\\/g"
> nodes_to_be_erased
```

If you only want to erase one server, enter the name of the server into the `nodes_to_be_erased` file, for example `Exa01celadm04`.

3. Copy the `dcli` utility from the Secure Eraser package along with the file generated in step 2 to the PXE server.
4. Create a PXE configuration template called `pxe_cfg.template` to contain the following lines:
 - For all models prior to Oracle Exadata Database Machine X7-2 — `grub`:
 - For releases earlier than Oracle Exadata System Software release 19.1.0:

```
default linux
label linux
kernel vmlinux-nfs-12.2.1.1.0-161015-cell
append initrd=initrd-nfs-12.2.1.1.0-161015-cell.img dhcp pxe quiet
loglevel=
0 secureeraser bootarea=diagnostics console=tty1
console=ttyS0,115200n8 logp
ath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
default linux
label linux
kernel vmlinux-iso-19.1.2.0.0-190111-cell
append initrd=initrd-iso-19.1.2.0.0-190111-cell.img dhcp pxe quiet
loglevel=
0 secureeraser bootarea=diagnostics console=tty1
console=ttyS0,115200n8 logp
ath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
```

For a description of each component of the above task, refer to the appropriate topic in [Automatic Secure Eraser through PXE Boot](#) for your system.

- For Oracle Exadata Database Machine X7-2 and newer models —grub2 / Secure Boot:

Note

In the following example, the following parameters must be updated to match your environment:

- kernel (the vmlinux file)
- initrd (the initrd*.img file)
- logpath

- For Oracle Exadata System Software 18c (18.1.0):

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-nfs-18.1.0.0.0-170915.1 stit dhcp pxe boot-
from=uefi
quiet loglevel=0 bootarea=diagnostics console=ttyS0,115200n8
logpath=10.133.42.221:/export/
exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-nfs-18.1.0.0.0-170915.1.img
    echo "Booting installation kernel"
}
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-iso-19.1.2.0.0-190111 stit dhcp pxe boot-
from=uefi
quiet loglevel=0 bootarea=diagnostics console=ttyS0,115200n8
logpath=10.133.42.221:/export/
exadata_secure_eraser_certificate_dir
```

```

    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-iso-19.1.2.0.0-190111.img
    echo "Booting installation kernel"
}

```

For a description of each component of the task in this step, see [Automatic Secure Eraser through PXE Boot](#).

This configuration file differs from the one in [Automatic Secure Eraser through PXE Boot](#) in that the `secureeraser` option is left out to indicate that the Secure Eraser process should not be triggered automatically. The rest of the file is the same.

5. On the PXE server, use the template file to generate a PXE configuration file. For all systems up to Oracle Exadata Database Machine X6-2, save the file in the `/tftpboot/pxelinux.cfg/` directory for each of the nodes to be erased. For Oracle Exadata Database Machine X7-2 and newer systems, save the file in the `/tftpboot/efi/pxelinux.cfg/` directory for each of the nodes to be erased.

The PXE configuration file name is the dash-separated MAC address of the node with the prefix `01-`.

If the nodes to be erased are accessible, perform the following steps to automatically generate a PXE configuration file for each node based on the template:

- a. Set up SSH equivalence with the nodes to be erased from the PXE server. The command will prompt for the root password of each node.

```
pxe_server# dcli -g nodes_to_be_erased -k -l root
```

- b. Create a list of PXE configuration files, one for each node to be erased based on the configuration template.

```

pxe_server# dcli -g nodes_to_be_erased -l root "ip addr show eth0" |
awk '/link\/ether/ {print "01:"$3}' | sed "s/:/-/g" | xargs -I {}
cp pxe_cfg.template {}

```

If the nodes are not accessible, perform the following steps to generate a PXE configuration file for each node.

- a. Manually collect the MAC address of the `eth0` interface from each node and write them into a text file called `mac_addresses`. Write one MAC address per line. For example:

```

00:10:e0:62:c4:fa
00:10:e0:62:c2:8a
00:10:e0:62:b8:7c
00:10:e0:62:b8:3a
00:10:e0:62:c6:bc

```

- b. Use the following command to create a list of PXE configuration file, one for each node to be erased based on the configuration template.

```

pxe_server# cat mac_addresses | sed "s/:/-/g;s/^/01-/g" | xargs -I {}
cp
pxe_cfg.template {}

```

In both cases, you should have a list of PXE configuration files, one for each node to be erased. For example, if the MAC addresses of the nodes in a quarter rack are

00:10:e0:62:c4:fa, 00:10:e0:62:c2:8a, 00:10:e0:62:b8:7c, 00:10:e0:62:b8:3a, and 00:10:e0:62:c6:bc, then you should get the following files:

```
01-00-10-e0-62-c4-fa
01-00-10-e0-62-c2-8a
01-00-10-e0-62-b8-7c
01-00-10-e0-62-b8-3a
01-00-10-e0-62-c6-bc
```

The files have the same content as the configuration template.

Check your specific PXE server requirements. Your PXE server may need slightly different names or settings.

6. Configure the nodes to boot from PXE and reboot the nodes.

If the nodes to be erased are accessible, run the following commands:

```
pxe_server# dcli -g nodes_to_be_erased -l root "ipmitool chassis bootdev
pxe"
```

```
pxe_server# dcli -g nodes_to_be_erased -l root "reboot"
```

If the nodes are not accessible, then perform the following steps:

a. Create a file called `iloms_to_be_reset` containing the names of ILOMs. For example:

```
db1-ilom
db2-ilom
cell1-ilom
cell2-ilom
cell3-ilom
```

b. Configure the nodes to boot from PXE through ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{}
-U root chassis bootdev pxe
```

c. Reboot the nodes from ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{}
-U root chassis power cycle
```

7. If you get the following prompt on the remote or serial console, enter `e` at the prompt to enter the diagnostic shell:

Choose from following by typing letter in '()':
(e)nter interactive diagnostics shell. Must use credentials from Oracle support to login (reboot or power cycle to exit the shell),
Select:`e`

8. If prompted, log in to the system as the root user.

If you require the password for the root user and do not have it, then contact Oracle Support Services.

```
localhost login: root
Password: *****
-sh-3.1#
```

9. Run the Secure Eraser utility to sanitize all devices or one type of device.

```
-sh-3.1# /usr/sbin/secureeraser --erase --all --output=REMOTE_NFS_LOCATION
```

REMOTE_NFS_LOCATION is the remote NFS location in the format of *IP:FILE_PATH*. The Secure Eraser utility will automatically mount the remote NFS location and save the certificate there.

For example, to erase all installed devices, including hard drives, flash devices, persistent memory, internal USBs, and ILOM, and save the certificate at this NFS location:
10.133.42.221:/export/exadata_secure_eraser_certificate_dir:

```
-sh-3.1# /usr/sbin/secureeraser --erase --all --output=10.133.42.221:/
export
/exadata_secure_eraser_certificate_dir
```

To erase just the hard drives:

```
-sh-3.1# /usr/sbin/secureeraser --erase --hdd --output=10.133.42.221:/
export
/exadata_secure_eraser_certificate_dir
```

Note that it is important to point the output option to an NFS location so that the certificate can be saved properly.

You will be prompted with a list of devices to be erased and to confirm that you want to proceed with Secure Eraser.

A progress report, as shown in step 7 of [Automatic Secure Eraser through PXE Boot for X7 and Later Systems](#), will be printed to the console every 10 seconds.

In interactive mode, the server will be left on after the specified devices are securely erased. You can power off the node from the diagnostic shell.

The web console will no longer be accessible if ILOM is reset. You can power off the server from the serial console or with the power button.

Related Topics

- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

5.5 Interactive Secure Eraser through Network Boot

Starting with Oracle Exadata X7-2, you can use EFI Network Boot when using Secure Eraser.

Note

Starting with Oracle Exadata System Software release 19.1.0, the Secure Eraser package (`secureeraser_label.zip`) contains ISO images instead of NFS images.

Before you begin:

- Download and install the latest Secure Eraser package available for your system.
To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).
- Make sure you have access to a NFS server that is accessible from all the nodes to be erased.
- Where necessary, configure appropriate firewall access rules to allow:
 - NFS traffic between the NFS server and all nodes being erased.
 - TFTP traffic between the EFI Network Boot server and all nodes being erased.
- Make sure you have access to one of the nodes to be erased.

1. Copy the cell `initrd` file (`initrd-version`) and kernel file (`vmlinux-version`) from the Secure Eraser package to the `/tftpboot/efi` on the network boot server.

The location does not have to be `/tftpboot/efi/`. The location is determined by the TFTP Server configuration.

2. Create a file containing the names of the database servers and storage servers to be erased.

To generate this file, you can run the following command from one of the nodes to be erased, and verify the nodes in the files are the ones to be erased.

```
# ibhosts | awk '/S [0-9\\.\\,]* / || /C [0-9\\.\\,]* / {print $6}' | sed
"s/\\/g"
> nodes_to_be_erased
```

If you only want to erase one server, enter the name of the server into the `nodes_to_be_erased` file, for example `Exa01celadm04`.

3. Copy the `dcli` utility from the Secure Eraser package along with the file generated in step 2 to the network boot server.
4. Create a configuration template called `pxe_cfg.template` to contain the following lines for grub2 and Secure Boot on Oracle Exadata Database Machine X7-2 and newer systems:

Note

In the following example, the following parameters must be updated to match your environment:

- kernel (the `vmlinux` file)
- `initrd` (the `initrd*.img` file)
- `logpath`

- For releases earlier than Oracle Exadata System Software release 19.1.0:

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-nfs-18.1.0.0-170915.1 stit dhcp pxe boot-
from=uefi quiet
    loglevel=0 bootarea=diagnostics console=ttyS0,115200n8
    logpath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-nfs-18.1.0.0-170915.1.img
    echo "Booting installation kernel"
}
```

- For Oracle Exadata System Software release 19.1.0 or later:

```
set default 0
set timeout=10
menuentry 'ExadataLinux' {
    echo "Loading efi/vmlinuz"
    linuxefi efi/vmlinux-nfs-19.1.2.0-190111 stit dhcp pxe boot-
from=uefi quiet
    loglevel=0 bootarea=diagnostics console=ttyS0,115200n8
    logpath=10.133.42.221:/export/exadata_secure_eraser_certificate_dir
    echo "Loading efi/initrd.img"
    initrdefi efi/initrd-nfs-19.1.2.0-190111.img
    echo "Booting installation kernel"
}
```

- The first line (`default`) identifies a menu entry that should be selected by default, after the timeout value specified by the second line.
- The third line (`menuentry`) represents the Linux kernel that will be used in the Secure Eraser environment.
- The fifth line (`linuxefi`) indicates the kernel is on an UEFI-based system. The `linuxefi` statement must be on a single line in the configuration file.
 - The `stit` option specifies INSTALL imaging mode, plus BARE METAL conditions, plus ERASING ADD DATA PARTITIONS
 - The `dhcp` option specifies to use DHCP to discover the `eth0` interface.
 - The `pxe` option suppresses search for the image on virtual CD and USB devices.
 - The `boot-from=uefi` option indicates the system is booting from UEFI.

- The `quiet` option disables excessive kernel log messages.
 - The `loglevel=0` option suppresses non-critical kernel messages.
 - The `secureeraser` option indicates the network boot will automatically trigger the Secure Eraser utility to sanitize all installed media, including hard drives, flash devices, persistent memory, internal USBs, and ILOM on the node.
 - The `bootarea` option indicates that the boot mode is diagnostic and not imaging install or rescue.
 - The console options indicate standard output and standard error messages are printed to both the ILOM web console and the serial console.
 - The `logpath` option specifies the NFS share directory where Secure Eraser will save the certificate.
 - The seventh line (`initrdefi`) specifies the `initrd` file to load. In this case it is the `initrd` file copied over in step 1.
5. On the network boot server, use the template file to generate a network boot configuration file in the `/tftpboot/efi/pxelinux.cfg/` directory (Oracle Exadata Database Machine X7-2 and newer) for each of the nodes to be erased.

The network boot configuration file name is the dash-separated MAC address of the node with the prefix `01-`.

If the nodes to be erased are accessible, perform the following steps to automatically generate a network boot configuration file for each node based on the template:

- a. Set up SSH equivalence with the nodes to be erased from the network boot server. The command will prompt for the root password of each node.

```
pxe_server# dcli -g nodes_to_be_erased -k -l root
```

- b. Create a list of network boot configuration files, one for each node to be erased based on the configuration template.

```
pxe_server# dcli -g nodes_to_be_erased -l root "ip addr show eth0" |
awk '/link\|ether/ {print "01:"$3}' | sed "s:/-/g" | xargs -I {}
cp pxe_cfg.template {}
```

If the nodes are not accessible, perform the following steps to generate a network boot configuration file for each node.

- a. Manually collect the MAC address of the `eth0` interface from each node and write them into a text file called `mac_addresses`. Write one MAC address per line. For example:

```
00:10:e0:62:c4:fa
00:10:e0:62:c2:8a
00:10:e0:62:b8:7c
00:10:e0:62:b8:3a
00:10:e0:62:c6:bc
```

- b. Use the following command to create a list of network boot configuration file, one for each node to be erased based on the configuration template.

```
pxe_server# cat mac_addresses | sed "s:/-/g;s/^/01-/g" | xargs -I {}
cp
pxe_cfg.template {}
```

In both cases, you should have a list of network boot configuration files, one for each node to be erased. For example, if the MAC addresses of the nodes in a quarter rack are 00:10:e0:62:c4:fa, 00:10:e0:62:c2:8a, 00:10:e0:62:b8:7c, 00:10:e0:62:b8:3a, and 00:10:e0:62:c6:bc, then you should get the following files:

```
01-00-10-e0-62-c4-fa
01-00-10-e0-62-c2-8a
01-00-10-e0-62-b8-7c
01-00-10-e0-62-b8-3a
01-00-10-e0-62-c6-bc
```

The files have the same content as the configuration template.

Check your specific network boot server requirements. Your network boot server may need slightly different names or settings.

6. Configure the nodes to boot from the network boot server and reboot the nodes.

If the nodes to be erased are accessible, run the following commands:

```
pxe_server# dcli -g nodes_to_be_erased -l root "ipmitool chassis bootdev
pxe"

pxe_server# dcli -g nodes_to_be_erased -l root "reboot"
```

If the nodes are not accessible, then perform the following steps:

a. Create a file called `iloms_to_be_reset` containing the names of ILOMs. For example:

```
db1-ilom
db2-ilom
cell1-ilom
cell2-ilom
cell3-ilom
```

b. Configure the nodes to boot from the network boot server through ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{}
-U root chassis bootdev pxe
```

c. Reboot the nodes from ILOMs. The command will prompt for ILOM root password.

```
pxe_server# cat iloms_to_be_reset | xargs -I {} ipmitool -I lanplus -H
{}
-U root chassis power cycle
```

7. If you get the following prompt on the remote or serial console, enter `e` at the prompt to enter the diagnostic shell:

Choose from following by typing letter in '()':
(e)nter interactive diagnostics shell. Must use credentials from Oracle support to login (reboot or power cycle to exit the shell),
Select:**e**

8. If prompted, log in to the system as the root user.

If you require the password for the root user and do not have it, then contact Oracle Support Services.

```
localhost login: root
Password: *****
-sh-3.1#
```

9. Run the Secure Eraser utility to sanitize all devices or one type of device.

```
-sh-3.1# /usr/sbin/secureeraser --erase --all --output=REMOTE_NFS_LOCATION
```

REMOTE_NFS_LOCATION is the remote NFS location in the format of *IP:FILE_PATH*. The Secure Eraser utility will automatically mount the remote NFS location and save the certificate there.

For example, to erase all installed devices, including hard drives, flash devices, persistent memory, internal USBs, and ILOM, and save the certificate at this NFS location:
10.133.42.221:/export/exadata_secure_eraser_certificate_dir:

```
-sh-3.1# /usr/sbin/secureeraser --erase --all --output=10.133.42.221:/
export
/exadata_secure_eraser_certificate_dir
```

To erase just the hard drives:

```
-sh-3.1# /usr/sbin/secureeraser --erase --hdd --output=10.133.42.221:/
export
/exadata_secure_eraser_certificate_dir
```

Note that it is important to point the output option to an NFS location so that the certificate can be saved properly.

You will be prompted with a list of devices to be erased and to confirm that you want to proceed with the Secure Eraser.

A progress report, as shown in step 7 of [Automatic Secure Eraser through PXE Boot for X7 and Later Systems](#), will be printed to the console every 10 seconds.

In interactive mode, the server will be left on after the specified devices are securely erased. You can power off the node from the diagnostic shell.

The web console will no longer be accessible if ILOM is reset. You can power off the server from the serial console or with the power button.

Related Topics

- [Secure Eraser Syntax](#)
Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

5.6 Secure Eraser Syntax

Secure Eraser securely erases all data on both database servers and storage servers, and resets InfiniBand Network Fabric or RDMA over Converged Ethernet (RoCE) switches, Ethernet switches, and power distribution units back to the factory default.

Syntax

`secureeraser options`

Command-Line Options for Secure Eraser

- `--all`
Perform the action (`--list` or `--erase`) on all devices on the system. Devices include hard drives, flash devices, persistent memory, USB devices, and ILOM.
- `--devices_to_erase`
This option allows you specify individual disks to be erased by providing their serial numbers. Multiple serial numbers can be provided separated with commas. Introduced with Oracle Exadata System Software release 19.1.0.
- `--erase`
Perform a secure erase of data.
- `--erasure_method_optional`
If a device is not eligible to be erased with the provided erasure method, the erasure method will silently fall back to the default method. Otherwise erasure will fail. This option can be used with all types of disks. Introduced with Oracle Exadata System Software release 19.1.0.
- `--flash`
Erase all flash devices.
- `--flash_erasure_method=FLASH_ERASURE_METHOD`
Force all flash devices to be erased using the specified method. The following values are supported:
 - `3pass` (Only on Oracle Exadata X4-8 or earlier systems)
 - `7pass` (Only on Oracle Exadata X4-8 or earlier systems)
 - `crypto` (Only on Oracle Exadata X5-2 or later systems, and requires Oracle Exadata System Software release 19.1.0 or later.)
- `--help, -h`
Show this help message and exit.
- `--hdd`
Erase all hard drives.
- `--hdd_erasure_method=HDD_ERASURE_METHOD`
Force all hard drives to be erased using the specified method. The following values are supported:
 - `3pass`

- 7pass
- crypto (Only on Oracle Exadata X6-2 or later systems, and requires Oracle Exadata System Software release 19.1.0 or later.)
- --ilom
Reset Integrated Lights Out Manager (ILOM) to factory default.
- --list
List devices (hard drives, flash devices, persistent memory, USB devices, and ILOM) on the system.
- --m2
Erase all M.2 devices.
- --m2_erasure_method=M2_ERASURE_METHOD
Force all M.2 devices to be erased using the specified method. The following values are supported:
 - 3pass
 - 7pass
 - crypto (Requires Oracle Exadata System Software release 19.1.0 or later)
- --output=CERTIFICATE_DIRECTORY, -o
Specify a full path to the directory for the certificate output location. The default is /var/log/cellos.
- --pmem
Erase all persistent memory (PMEM) devices using cryptographic erasure.
- --quiet, -q
Quietly skip prompts.
- --technician=TECHNICIAN_NAME, -t TECHNICIAN_NAME
Specify the name of the technician performing the erasure. This name will be recorded in the certificate.
- --usb
Erase all internal USB devices.
- --usb_erasure_method=USB_ERASURE_METHOD
Force all internal USB devices to be erased using the specified method. The following values are supported:
 - 3pass
 - 7pass
- --witness=WITNESS_NAME, -w WITNESS_NAME
Specify the name of the person witnessing the erasure. This name will be recorded in the certificate.

Usage Notes

Before you begin, download and install the latest Secure Eraser package available for your system.

To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).

Examples of Secure Eraser Syntax

List all devices (hard drives, flash devices, persistent memory, USB devices, and ILOM) on the system.

```
secureeraser --list --all
```

List all hard drives.

```
secureeraser --list --hdd
```

Securely erase all devices, and enter the names of the technician and witness in the certificate.

```
secureeraser --erase --all --technician="jdoe" --witness="jsmith"
```

Reset ILOM to factory default.

```
secureeraser --erase --ilom
```

Securely erase all hard drives.

```
secureeraser --erase --hdd
```

Securely erase all hard drives, all flash devices, and all internal USB devices. Force "crypto" method on flash devices.

```
secureeraser --erase --hdd --flash --usb --flash_erasure_method crypto
```

5.7 Resetting Network Switches and Power Distribution Units to Factory Default

Before you begin:

- Download and install the latest Secure Eraser package available for your system.
To download the Secure Eraser package, examine the Supplemental Readme associated with your current Oracle Exadata System Software version. To find the Supplemental Readme associated with each Oracle Exadata System Software version, see [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID KB153930\)](#).
- Print out the Exadata Factory Reset Certificate template in the Secure Eraser package.

The following figure shows the Factory Reset certificate:

Figure 5-2 Factory Reset Certificate

ORACLE
EXADATA

Factory Reset Certificate

This is to certify the following components have been reset to factory default.

Component	Serial Number	Technician's Signature	Date
Sun Datacenter InfiniBand Switch 36	_____	_____	_____
Sun Datacenter InfiniBand Switch 36	_____	_____	_____
Sun Datacenter InfiniBand Switch 36	_____	_____	_____
Cisco Catalyst 4948 Ethernet Switch	_____	_____	_____
Power Distribution Unit	_____	_____	_____
Power Distribution Unit	_____	_____	_____

Use the following procedures:

- [Resetting a Cisco Nexus 9336C-FX2 RoCE Network Fabric Switch to Factory Default Settings](#)
- [Resetting InfiniBand Network Fabric Switches to Factory Default](#)
- [Resetting the Cisco Management Network Switch to Factory Default Settings](#)
You can reset the Cisco Management Network Switch configuration to the original default factory settings.
- [Resetting Power Distribution Units to Factory Default](#)
You can reset the power distribution units (PDUs) configuration to the original default factory settings.

5.7.1 Resetting a Cisco Nexus 9336C-FX2 RoCE Network Fabric Switch to Factory Default Settings

The Cisco Nexus 9336C-FX2 RoCE Network Fabric switch comes preconfigured with specific configurations for RDMA over Converged Ethernet (RoCE). If you plan to reuse this switch in Oracle Exadata you need to save this configuration to:

- Local bootflash
- A remote server

To reset a Cisco Nexus 9336C-FX2 RoCE Network Fabric switch:

1. Make up backup of the current switch configuration.

You must save the current configuration if you plan to use the switch after resetting the configuration back to the factory settings. If you are returning the switch, then you do not need to save the RoCE-specific configuration.

Follow the steps documented in *Backing Up Settings on the ROCE Switch*

2. Delete the files in all directories on the switch.
 - a. List the directories available on the switch.

```
switch# dir ?
```

- b. For each directory listed in the above output (represented as *dir_name*), view the directory contents.

```
switch# dir dir_name:
```

- c. If any files are found in a directory, then delete the files.

```
switch# del dir_name:* no-prompt
```

3. Use the `write erase` command on the switch to erase the current configuration.

The `write erase` command erases the entire startup configuration, except for the following:

- Boot variable definitions.
- The IPv4 and IPv6 configuration on the `mgmt0` interface, including the subnet mask and default gateway/route in the management VR.

To also remove the boot variable definitions and the IPv4/IPv6 configuration on the `mgmt0` interface, use the `write erase boot` command.

Note

- The `write erase` command deletes the access authentication settings. After running this command, you must use the console to access the switch.
- Also, after you run the `write erase` command, you must reload the ASCII configuration twice to apply the breakout configuration.

4. Record the serial numbers of the switches that have been reset to factory default in the Factory Reset certificate template. Sign and date the entries.

You can identify the serial number of an RoCE Network Fabric switch by running the following command on the switch:

```
switch# show license host-id  
License hostid:VDH=FOX064317SQ
```

The host ID is also referred to as the device serial number. In the above example, you use all the text that appears after the equal sign (=), so the switch serial number is FOX064317SQ.

① See Also

Cisco NX-OS Licensing Guide at https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/licensing/guide/b_Cisco_NX-OS_Licensing_Guide.pdf

5.7.2 Resetting InfiniBand Network Fabric Switches to Factory Default

To reset the InfiniBand Network Fabric switches to factory default, refer to My Oracle Support document KB393468.

Record the serial numbers of the switches that have been reset to factory default in the Factory Reset certificate template. Sign and date the entries.

You can identify the serial number of an InfiniBand Network Fabric switch by running the following command on the switch:

```
[root@switch1 ~]# version | grep "Serial Number"
```

Related Topics

- [Sun Datacenter InfiniBand Switch Reset to Factory Default Setting \(My Oracle Support Doc ID KB393468\)](#)

5.7.3 Resetting the Cisco Management Network Switch to Factory Default Settings

You can reset the Cisco Management Network Switch configuration to the original default factory settings.

For Oracle Exadata X7-2 or later

To reset the Cisco Management Network Switch to factory default:

1. Display the start up configuration.

```
switch# show startup-config
```

2. Display the boot configuration.

```
switch# show boot
```

3. Display the debug configuration.

```
switch# show debug
```

4. Delete the files in all directories on the switch.

- a. List the directories available on the switch.

```
switch# dir ?
```

- b. For each directory listed in the above output (represented as *dir_name*), view the directory contents.

```
switch# dir dir_name:
```

- c. If any files are found in a directory, then delete the files.

```
switch# del dir_name:* no-prompt
```

5. Erase the startup-configuration file.

```
switch# write erase
```

Warning: This command will erase the startup-configuration.

Do you wish to proceed anyway? (y/n) [n] y

6. Erase the boot variable definitions.

```
switch# write erase boot
```

Warning: This command will erase the boot variables and the ip configuration of interface mgmt 0

Do you wish to proceed anyway? (y/n) [n] y

7. Erase the debugging configuration.

```
switch# write erase debug
```

8. Reload the Cisco Nexus 93108-1G or Cisco Nexus 9348 Ethernet switch.

```
switch# reload
```

This command will reboot the system. (y/n) [n] y

① See Also

["Erasing a Configuration"](#) in *Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 6.x*

For Oracle Exadata X6-2 or earlier

To reset the Cisco Management Network Switch to factory default, refer to "Reset Catalyst Switches Running Cisco IOS Software" in the Cisco Troubleshooting TechNotes Document ID: 24328.

Record the serial number of the Ethernet switch that has been reset to factory default in the Exadata Factory Reset certificate template. Sign and date the entry.

The serial number of an Ethernet switch can be identified by the "Processor board ID" field in the "show version" command output.

```
switch# show version
```

① See Also

"Reset Catalyst Switches Running Cisco IOS Software" at http://www.cisco.com/c/en/us/support/docs/switches/catalyst-2900-xl-series-switches/24328-156.html#reset_ios

5.7.4 Resetting Power Distribution Units to Factory Default

You can reset the power distribution units (PDUs) configuration to the original default factory settings.

There are two types of power distribution units (PDUs): original PDUs and enhanced PDUs. Enhanced PDUs have SER MGT port that can be connected to a host using an RS-232 cable, whereas the original PDUs do not have SER MGT port. Typically, Oracle Exadata V2 to Oracle Exadata X3 racks have the original PDUs, and Oracle Exadata X4-2 and later have the enhanced PDUs.

You can reset both the original power distribution units and the enhanced power distribution units, as described in the *Sun Rack II Power Distribution Units User's Guide* or the *Oracle Rack Cabinet 1242 Power Distribution Units User's Guide* (for Oracle Exadata X7-2 and later systems).

Record the serial numbers of the power distribution units that have been reset to factory default in the Exadata Factory Reset certificate template. Sign and date the entries.

The serial number can be found on the "View Module Information" topic of the appropriate *Power Distribution Units User's Guide*.

For enhanced PDUs, the serial number can also be retrieved through the following CLI command:

```
pducli -> get pdu_serial_number
```

① See Also

- "[View Module Information](#)" (Original or Enhanced PDU) in *Sun Rack II Power Distribution Units User's Guide*
- "[View Module Information](#)" (Original or Enhanced PDU) in *Oracle Rack Cabinet 1242 Power Distribution Units User's Guide*
- "[Restore the PDU to Factory Default Settings](#)" (Original or Enhanced PDU) in *Sun Rack II Power Distribution Units User's Guide*
- "[Restore the PDU to Factory Default Settings](#)" (Original or Enhanced PDU) in *Oracle Rack Cabinet 1242 Power Distribution Units User's Guide*

5.8 Actions After Using Secure Eraser

After performing a secure erase, the system is ready for return or re-imaging.

If you plan to re-image the machine, you must perform the following steps:

1. Connect to and configure ILOM. See "[Oracle ILOM – Quick Start](#)" in the *Oracle ILOM Getting Started Guide Firmware Release 4.0.x*.

Perform the following tasks:

- a. Connect to Oracle ILOM
 - b. Log In to Oracle ILOM
 - c. Modify Default Network Connectivity Settings
2. Re-image the system. Refer to Imaging a New System in *Oracle Exadata Installation and Configuration Guide*.

If you are preparing to return the machine, refer to <http://www.oracle.com/us/products/servers-storage/take-back-and-recycling/index.html>