



SUSE Enterprise Storage 7

Security Hardening Guide

Security Hardening Guide

SUSE Enterprise Storage 7

by Tomáš Bažant, Alexandra Settle, and Liam Proven

Publication Date: 12 Dec 2024

<https://documentation.suse.com> 

Copyright © 2020–2024 SUSE LLC and contributors. All rights reserved.

Except where otherwise noted, this document is licensed under Creative Commons Attribution-ShareAlike 4.0 International (CC-BY-SA 4.0): <https://creativecommons.org/licenses/by-sa/4.0/legalcode> .

For SUSE trademarks, see <http://www.suse.com/company/legal/> . All third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors nor the translators shall be held liable for possible errors or the consequences thereof.

Contents

About this guide vi

- 1 Available documentation vi
- 2 Giving feedback vii
- 3 Documentation conventions viii
- 4 Support ix
 - Support statement for SUSE Enterprise Storage x • Technology previews xi
- 5 Ceph contributors xi
- 6 Commands and command prompts used in this guide xii
 - Salt-related commands xii • Ceph related commands xii • General Linux commands xiii • Additional information xiv

I INTRODUCTION 1

- 1 Understanding the threat 2
- 2 About this guide 3

II HARDENING MEASSURES 5

- 3 General 6
 - 3.1 Basic security hygiene 6
 - 3.2 General system hardening 6
 - 3.3 Monitoring 7

4	Network	8
5	Prevent Denial Of Service (DoS)	10
6	Authentication	12
6.1	Enabling strong authentication	12
6.2	Ensuring secure storage of keys	13
6.3	Account setup	13
7	Confidentiality	15
7.1	Data at rest	15
7.2	Data in flight	15
	Glossary	18

About this guide

This guide focuses on how to ensure that your Ceph cluster is secure.

SUSE Enterprise Storage 7 is an extension to SUSE Linux Enterprise Server 15 SP2. It combines the capabilities of the Ceph (<http://ceph.com/>) storage project with the enterprise engineering and support of SUSE. SUSE Enterprise Storage 7 provides IT organizations with the ability to deploy a distributed storage architecture that can support a number of use cases using commodity hardware platforms.

1 Available documentation



Note: Online documentation and latest updates

Documentation for our products is available at <https://documentation.suse.com>, where you can also find the latest updates, and browse or download the documentation in various formats. The latest documentation updates can be found in the English language version.

In addition, the product documentation is available in your installed system under `/usr/share/doc/manual`. It is included in an RPM package named `ses-manual_LANG_CODE`. Install it if it is not already on your system, for example:

```
# zypper install ses-manual_en
```

The following documentation is available for this product:

Deployment Guide (<https://documentation.suse.com/ses/html/ses-all/book-storage-deployment.html>)

This guide focuses on deploying a basic Ceph cluster, and how to deploy additional services. It also covers the steps for upgrading to SUSE Enterprise Storage 7 from the previous product version.

Administration and Operations Guide (<https://documentation.suse.com/ses/html/ses-all/book-storage-admin.html>)

This guide focuses on routine tasks that you as an administrator need to take care of after the basic Ceph cluster has been deployed (day 2 operations). It also describes all the supported ways to access data stored in a Ceph cluster.

Security Hardening Guide (<https://documentation.suse.com/ses/html/ses-all/book-storage-security.html>) ↗

This guide focuses on how to ensure your cluster is secure.

Troubleshooting Guide (<https://documentation.suse.com/ses/html/ses-all/book-storage-troubleshooting.html>) ↗

This guide takes you through various common problems when running SUSE Enterprise Storage 7 and other related issues to relevant components such as Ceph or Object Gateway.

SUSE Enterprise Storage for Windows Guide (<https://documentation.suse.com/ses/html/ses-all/book-storage-windows.html>) ↗

This guide describes the integration, installation, and configuration of Microsoft Windows environments and SUSE Enterprise Storage using the Windows Driver.

2 Giving feedback

We welcome feedback on, and contributions to, this documentation. There are several channels for this:

Service requests and support

For services and support options available for your product, see <http://www.suse.com/support/> ↗.

To open a service request, you need a SUSE subscription registered at SUSE Customer Center. Go to <https://scc.suse.com/support/requests> ↗, log in, and click *Create New*.

Bug reports

Report issues with the documentation at <https://bugzilla.suse.com/> ↗. Reporting issues requires a Bugzilla account.

To simplify this process, you can use the *Report Documentation Bug* links next to headlines in the HTML version of this document. These preselect the right product and category in Bugzilla and add a link to the current section. You can start typing your bug report right away.

Contributions

To contribute to this documentation, use the *Edit Source* links next to headlines in the HTML version of this document. They take you to the source code on GitHub, where you can open a pull request. Contributing requires a GitHub account.

For more information about the documentation environment used for this documentation, see the repository's README at <https://github.com/SUSE/doc-ses>.

Mail

You can also report errors and send feedback concerning the documentation to doc-team@suse.com. Include the document title, the product version, and the publication date of the document. Additionally, include the relevant section number and title (or provide the URL) and provide a concise description of the problem.

3 Documentation conventions

The following notices and typographic conventions are used in this document:

- `/etc/passwd`: Directory names and file names
- `PLACEHOLDER`: Replace `PLACEHOLDER` with the actual value
- `PATH`: An environment variable
- `ls`, `--help`: Commands, options, and parameters
- `user`: The name of user or group
- `package_name`: The name of a software package
- `Alt`, `Alt-F1`: A key to press or a key combination. Keys are shown in uppercase as on a keyboard.
- `File`, `File > Save As`: menu items, buttons
- `AMD/Intel` This paragraph is only relevant for the AMD64/Intel 64 architectures. The arrows mark the beginning and the end of the text block. `<`
- `IBM Z, POWER` This paragraph is only relevant for the architectures `IBM Z` and `POWER`. The arrows mark the beginning and the end of the text block. `<`
- *Chapter 1, “Example chapter”*: A cross-reference to another chapter in this guide.
- Commands that must be run with `root` privileges. Often you can also prefix these commands with the `sudo` command to run them as non-privileged user.

```
# command
> sudo command
```

- Commands that can be run by non-privileged users.

```
> command
```

- Notices



Warning: Warning notice

Vital information you must be aware of before proceeding. Warns you about security issues, potential loss of data, damage to hardware, or physical hazards.



Important: Important notice

Important information you should be aware of before proceeding.



Note: Note notice

Additional information, for example about differences in software versions.



Tip: Tip notice

Helpful information, like a guideline or a piece of practical advice.

- Compact Notices



Additional information, for example about differences in software versions.



Helpful information, like a guideline or a piece of practical advice.

4 Support

Find the support statement for SUSE Enterprise Storage and general information about technology previews below. For details about the product lifecycle, see <https://www.suse.com/lifecycle>. If you are entitled to support, find details on how to collect information for a support ticket at <https://documentation.suse.com/sles-15/html/SLES-all/cha-adm-support.html>.

4.1 Support statement for SUSE Enterprise Storage

To receive support, you need an appropriate subscription with SUSE. To view the specific support offerings available to you, go to <https://www.suse.com/support/> and select your product.

The support levels are defined as follows:

L1

Problem determination, which means technical support designed to provide compatibility information, usage support, ongoing maintenance, information gathering and basic troubleshooting using available documentation.

L2

Problem isolation, which means technical support designed to analyze data, reproduce customer problems, isolate problem area and provide a resolution for problems not resolved by Level 1 or prepare for Level 3.

L3

Problem resolution, which means technical support designed to resolve problems by engaging engineering to resolve product defects which have been identified by Level 2 Support.

For contracted customers and partners, SUSE Enterprise Storage is delivered with L3 support for all packages, except for the following:

- Technology previews.
- Sound, graphics, fonts, and artwork.
- Packages that require an additional customer contract.
- Some packages shipped as part of the module *Workstation Extension* are L2-supported only.
- Packages with names ending in `-devel` (containing header files and similar developer resources) will only be supported together with their main packages.

SUSE will only support the usage of original packages. That is, packages that are unchanged and not recompiled.

4.2 Technology previews

Technology previews are packages, stacks, or features delivered by SUSE to provide glimpses into upcoming innovations. Technology previews are included for your convenience to give you a chance to test new technologies within your environment. We would appreciate your feedback! If you test a technology preview, please contact your SUSE representative and let them know about your experience and use cases. Your input is helpful for future development.

Technology previews have the following limitations:

- Technology previews are still in development. Therefore, they may be functionally incomplete, unstable, or in other ways *not* suitable for production use.
- Technology previews are *not* supported.
- Technology previews may only be available for specific hardware architectures.
- Details and functionality of technology previews are subject to change. As a result, upgrading to subsequent releases of a technology preview may be impossible and require a fresh installation.
- SUSE may discover that a preview does not meet customer or market needs, or does not comply with enterprise standards. Technology previews can be removed from a product at any time. SUSE does not commit to providing a supported version of such technologies in the future.

For an overview of technology previews shipped with your product, see the release notes at https://www.suse.com/releasenotes/x86_64/SUSE-Enterprise-Storage/7.

5 Ceph contributors

The Ceph project and its documentation is a result of the work of hundreds of contributors and organizations. See <https://ceph.com/contributors/> for more details.

6 Commands and command prompts used in this guide

As a Ceph cluster administrator, you will be configuring and adjusting the cluster behavior by running specific commands. There are several types of commands you will need:

6.1 Salt-related commands

These commands help you to deploy Ceph cluster nodes, run commands on several (or all) cluster nodes at the same time, or assist you when adding or removing cluster nodes. The most frequently used commands are `ceph-salt` and `ceph-salt config`. You need to run Salt commands on the Salt Master node as `root`. These commands are introduced with the following prompt:

```
root@master #
```

For example:

```
root@master # ceph-salt config ls
```

6.2 Ceph related commands

These are lower-level commands to configure and fine tune all aspects of the cluster and its gateways on the command line, for example `ceph`, `cephadm`, `rbd`, or `radosgw-admin`.

To run Ceph related commands, you need to have read access to a Ceph key. The key's capabilities then define your privileges within the Ceph environment. One option is to run Ceph commands as `root` (or via `sudo`) and use the unrestricted default keyring 'ceph.client.admin.key'. The safer and recommended option is to create a more restrictive individual key for each administrator user and put it in a directory where the users can read it, for example:

```
~/ceph/ceph.client.USERNAME.keyring
```



Tip: Path to Ceph keys

To use a custom admin user and keyring, you need to specify the user name and path to the key each time you run the `ceph` command using the `-n client.USER_NAME` and `--keyring PATH/TO/KEYRING` options.

To avoid this, include these options in the `CEPH_ARGS` variable in the individual users' `~/.bashrc` files.

Although you can run Ceph-related commands on any cluster node, we recommend running them on the Admin Node. This documentation uses the `cephuser` user to run the commands, therefore they are introduced with the following prompt:

```
cephuser@adm >
```

For example:

```
cephuser@adm > ceph auth list
```



Tip: Commands for specific nodes

If the documentation instructs you to run a command on a cluster node with a specific role, it will be addressed by the prompt. For example:

```
cephuser@mon >
```

6.2.1 Running `ceph-volume`

Starting with SUSE Enterprise Storage 7, Ceph services are running containerized. If you need to run `ceph-volume` on an OSD node, you need to prepend it with the `cephadm` command, for example:

```
cephuser@adm > cephadm ceph-volume simple scan
```

6.3 General Linux commands

Linux commands not related to Ceph, such as `mount`, `cat`, or `openssl`, are introduced either with the `cephuser@adm >` or `#` prompts, depending on which privileges the related command requires.

6.4 Additional information

For more information on Ceph key management, refer to *Book "Administration and Operations Guide", Chapter 30 "Authentication with cephx", Section 30.2 "Key management"*.

I Introduction

- 1 Understanding the threat 2
- 2 About this guide 3

1 Understanding the threat

Before you start to harden your SUSE Enterprise Storage cluster you need to consider the threat landscape you try to control.

Depending on your exposure you will need to invest at different levels. You need to take differing measures when you use SUSE Enterprise Storage to provide storage on an internal network to a well-known group of employees, as opposed to deploying a cluster in a setting where arbitrary internet actors can access the cluster. For example, in a public cloud offering as storage solution. This is something that needs to happen as a first step as it provides a set of guidelines on how much effort you need to invest to get to the level of security you need.

If you have a mature IT security landscape, you should have already policies and standards that you can use to guide you here. You need to have a threat model for the planned system and implement measures that you find necessary for your situation. Look to you CISO or similar role for guidance on this. It is mandatory to understand the potential threats and security requirements before you continue.

Without a threat model you run the risk of not investing enough or you might spent too much on securing a resource than you should. A good approach to this is described in the [OWASP Threat Modeling Cheat Sheet \(https://cheatsheetseries.owasp.org/cheatsheets/Threat_Modeling_Cheat_Sheet.html\)](https://cheatsheetseries.owasp.org/cheatsheets/Threat_Modeling_Cheat_Sheet.html).

2 About this guide

There are three main pillars of security:

- **Confidentiality:** Protect information from unauthorized access.
- **Integrity:** Ensure that information is accurate, consistent and only changed via authorized operations.
- **Availability:** The resource is available when needed.

This guide is mainly concerned with confidentiality and integrity. Availability is something you can configure in SUSE Enterprise Storage easily depending on your requirements. It is important that you consider the availability requirements you have and create the cluster accordingly. For example, high availability requirements ensure that you do not have a single point of failure. Especially Ceph Monitor nodes are critical and need to be available at all times. SUSE Enterprise Storage makes it easy for you to have more nodes for a given service than what you need during normal operations. The more important a service is to you the more redundancy you need to build into the system.

When you plan for availability, make sure the SUSE Enterprise Storage environment is not isolated. The requirements also affect other systems that interact with SUSE Enterprise Storage transitively. For example, if you use LDAP for authentication, then a highly available SUSE Enterprise Storage cluster does not help you if the LDAP server is not reachable.

Confidentiality needs to consider the life cycle of the data in question and the hardware that is used to hold the data. You not only need to take measures to ensure the confidentiality of data while it's kept in SUSE Enterprise Storage, you also need to consider how to safely discard data once you remove hardware.

Integrity requires that the cluster is in a trusted state and that data can only be modified by authorized subjects. Keeping the cluster up to date is the most important step in ensuring the integrity of the cluster. Ensuring that only authorized subjects can change data requires that permissions are handed out in a controlled and granular way. To ensure that this does not deteriorate over time, ensure that you regularly review existing permissions and have processes in place that revoke them if necessary.

This guide will not give you a set of commands you can run to ensure security. The new system needs to be integrated into your organizational security framework and the concrete steps often depend on your local configurations. For example, the recommendations on how to structure the network and which ports need to be made available then need to be translated into changes of your existing networking hardware, such as firewalls.

Some suggested changes have performance implications. Changing a plain text to an encrypted protocol causes the cluster to have more work. This may not be noticeable (such as full disk encryption for OSDs, where the CPU is not the limiting factor), but you need to check for your setup if this causes issues for you with workloads that are realistic for your environment.

II Hardening measures

- 3 General 6
- 4 Network 8
- 5 Prevent Denial Of Service (DoS) 10
- 6 Authentication 12
- 7 Confidentiality 15

3 General

Hardening your SUSE Enterprise Storage installation involves reducing the attack surface presented to potential attackers. But this is only the tip of the iceberg. All the basic tasks of securing a system applied to SUSE Enterprise Storage as well.

3.1 Basic security hygiene

As with any other system it is important that you practice proper security hygiene for you SUSE Enterprise Storage installation. This includes monitoring a suitable channel for security notices (<https://www.suse.com/security/cve/>) and incorporate this in your security tracking.

It is mandatory that you install updates in a timely manner. If available, you can use threat intelligence to guide you in your update strategy, but the sooner you install updates the better. Most organizations do not get hacked via 0-day exploits but through long known security issues. If you keep your cluster current you improve the security posture dramatically.

Installing updates in a SUSE Enterprise Storage context means that you keep the base operating system and the SUSE Enterprise Storage images up to date. For the base operating system you can either use basic command line tools like **zypper** or use SUSE Manager to conveniently manage a large fleet of machines. Refer to *Book "Administration and Operations Guide", Chapter 13 "Operational tasks", Section 13.7 "Updating Ceph"* on how to keep the SUSE Enterprise Storage images up to date.

3.2 General system hardening

Ensuring that the base system is hardened is helping to provide a proper base for further hardening measures more specific to SUSE Enterprise Storage. SUSE published a hardening guide for SUSE Linux Enterprise Server at <https://documentation.suse.com/de-de/sles/15-SP1/html/SLES-all/book-hardening.html>. As SUSE Enterprise Storage is based on SUSE Linux Enterprise Server this contains tips that you can incorporate in your security strategy. For example, we recommend that you ensure that the systems that host SUSE Enterprise Storage are physically secure and that the boot process is protected is important to have a solid base for further hardenings.

We also recommend that you do not add any other workloads on the machines that you use for your SUSE Enterprise Storage cluster. Not only can this negatively impact the performance of your SUSE Enterprise Storage cluster, but you also introduce additional risk to your data. If an attacker is able to exploit a vulnerability in the unrelated workload, they may be able to use this access to compromise your SUSE Enterprise Storage cluster.

If you have a virtualized environment and can easily provision machines, we recommend using one machine for each role. Especially the Ceph Monitor should be stand-alone as they have access to all the key material and running other services on them increases their risk profile.

3.3 Monitoring

Without visibility into your systems it is tough to ensure that they run in a secure state. You have to either monitor the SUSE Enterprise Storage cluster itself or hook it into your existing monitoring framework to ensure that you are aware of changes in the cluster. This mainly helps with availability, but is also important for other security goals. For example, you need to notice if someone is trying to brute force credentials on the machines by collecting and analyzing the logs showing this behavior.

You should at least include `/var/log/ceph/cephadm.log` into your log analysis setup to make sure you notice changes on your SUSE Enterprise Storage cluster.

4 Network

SUSE Enterprise Storage is a complex system that communicates internally and externally via networks. As with all other systems careful design and control of this network access is vital for ensuring the security of your cluster.

While SUSE Enterprise Storage can be run with a single network connected to all nodes, it is important for security to use the setup recommended in *Book "Deployment Guide", Chapter 2 "Hardware requirements and recommendations", Section 2.1 "Network overview"* and have two separate networks connected to your cluster. It is to be preferred to have physically separate networks. If this is not possible, you can use VLANs to logically separate them.

The internal network is used for replication and recovery and should not be available to clients. Unless special measures are taken (described in [Chapter 7, Confidentiality](#)) data stored on SUSE Enterprise Storage is transferred in cleartext on this network. Even when you ensure that data is transferred only in encrypted form, we highly recommend to use a dedicated network.

The public network is used as interface for clients and can be restricted to the minimal necessary access and also be monitored for anomalies.

These are the TCP ports that are necessary for various services:

- 3300, 6789: Monitor nodes
- 6800-7300: OSD nodes
- 6800-7300: MGR nodes
- 6800: MDS nodes
- 80,443,7480: Radosgw
- 8080,8443: Dashboard
- 4505,4506: Administration via salt

You should ensure on a network and on a host level that these ports are only accessible to the strictest possible set of clients. All other ports should be blocked by a default-deny rule. Remember to block the ports you want to deny access to for IPv4 and IPv6 if that is enabled in your environment.

Consider your use case and then analyze what network access is necessary on each network. For example, the Ceph Dashboard usually does not need to be accessible to users and access to it can be restricted via firewalls. The RADOS Block Device, CephFS, and the Object Gateway must

be available to the clients that use them. If certain services are not used, or only a limited set of users use it, you can prevent access to these ports in general or for groups that do not need access. This limits the damage that can be done if a vulnerability in this component is found.

5 Prevent Denial Of Service (DoS)

The most important piece in preventing Denial Of Service (DoS) is to put proper quotas on users and groups to ensure that clients can not exhaust resources easily. While this is not the only way a client can impact your cluster, it's the easiest one and also can happen by accident. For details on how to setup quotas please refer to *Book "Administration and Operations Guide", Chapter 23 "Clustered file system", Section 23.6 "Setting CephFS quotas"* and *Book "Administration and Operations Guide", Chapter 21 "Ceph Object Gateway", Section 21.5.2.4 "Enabling user quota management"*.



Important

Be aware that CephFS quotas are enforced client side, so a malicious client can ignore them and exceed the limitations. If this is a concern in your environment, do not use CephFS.

To set the quotas conveniently you can use the Ceph Dashboard.

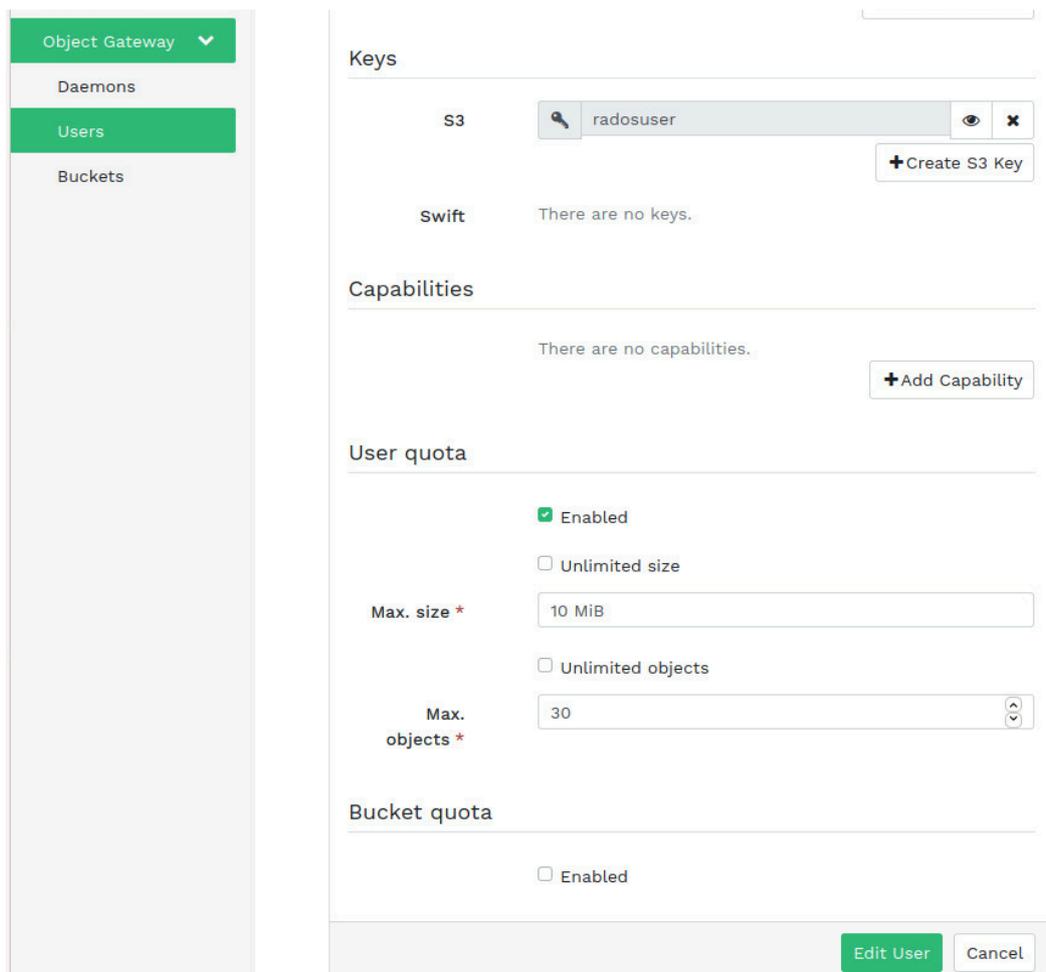


FIGURE 5.1: QUOTAS IN THE DASHBOARD

Current Ceph versions do not offer advanced ways of preventing malicious clients from attacking the availability of the cluster (for example, with many open connections). To ensure you notice an attack or a misconfiguration, you need to setup proper monitoring that will alert you if the cluster gets into a problematic state so you can investigate and if necessary act.

6 Authentication

Ensuring that clients need to authenticate before accessing resources on SUSE Enterprise Storage is important to the security of the system. Allowing anonymous usage or weak authentication schemes should be avoided.

6.1 Enabling strong authentication

Enforce strong authentication whenever possible: `cephx` works by having a secret key shared between the client and the service. This way both sides can prove to each other that they are who they claim to be. `cephx` is the default authentication scheme and should not be replaced by weaker methods.

`cephx` is enabled by default for communication in the cluster itself (`auth_cluster_required`) and for the communication between the client and the cluster (`auth_service_required`). You can check this by running the following:

```
cephuser@adm > for option_name in auth_cluster_required auth_service_required
auth_client_required ; do
    echo -n "$option_name: "
    ceph config get mon $option_name
done
auth_cluster_required: cephx
auth_service_required: cephx
auth_client_required: cephx, none
```

You can also enable `auth_client_required` to force the SUSE Enterprise Storage cluster to authenticate towards clients to prevent malicious actors from impersonating services. This is done by setting `auth_client_required` to `cephx` only with the `ceph config set global auth_client_required cephx` command.

`cephx` is only concerned with authentication. It does not ensure that the data is encrypted when sent to the cluster (in transport) or when stored in the cluster (at rest). When choosing a way for clients to access the cluster, select a access method that ensures the confidentiality in transport (for example, use HTTPS when using RADOS). For more details about `cephx`, see *Book "Administration and Operations Guide", Chapter 30 "Authentication with cephx", Section 30.1 "Authentication architecture"*.

You can enforce message signing to harden the authentication protocol. With the `ceph config set global cephx_require_signatures true` command, you can force that signatures are used on all messages between the client and the cluster and in between the daemons in

the cluster. If you run into issue with clients not being able to properly sign their messages you can enable it only for use within the cluster with the `ceph config set global cephx_cluster_require_signatures` command.

Strong authentication also requires proper key handling from the creation to the destruction of keys. Each client should receive a separate key which shouldn't be reused, at least not for clients with different security properties. With that you can then give the least amount of privileges to each account that is necessary, therefor limiting the risk. Creating users is covered in *Book "Administration and Operations Guide", Chapter 20 "RADOS Block Device", Section 20.2.1 "Creating a Ceph user account"*.

6.2 Ensuring secure storage of keys

Keys must be stored with safe permissions on the client machine to prevent credential leakage. In most cases this means that only the user and root is able to read the key material. If you have a setup where you need to provide broader access you need to think through the security implications that accidental or malicious leaks of the key material has in your environment.

By default, key material is stored in a safe way on SUSE Enterprise Storage and you need to make sure that you do the same when transferring key material to clients. This means that you use secure transport mechanisms (HTTPS, SSH) to transfer the key and set strict permissions of files storing key material. Depending on your security requirements the use of vault services or hardware security modules might be appropriate.

Also consider your backup scheme to ensure that keys are secure during the whole life cycle. The backup of keys must meet the same security standards as the other systems that store a key during its lifetime.

6.3 Account setup

In the initial configuration, you have administrative accounts that hold all the power. Depending on your regulatory and other requirements, you need to split up these accounts into several accounts that hold different privileges. This allows to assign the least amount of privilege needed to fulfill a task.

You should create a process that regularly reviews the privileges users have and adjust them if necessary. Especially for highly privileged accounts, we recommend this happens on a regular basis and every time a user that could have modified these settings is removed from the settings. For example, if someone changes role and is not the administrator for SUSE Enterprise Storage anymore.

7 Confidentiality

Confidentiality is a common requirement. There are different ways of ensuring data stays confidential at different times of its lifecycle.

7.1 Data at rest

The data stored in OSDs is not encrypted by default. If you have high confidentiality requirements, we recommend that you encrypt the storage that you provide SUSE Enterprise Storage for use. This protects the data if a disk is stolen or when you decommission disk drives.

The easiest way to do this is to enable disk encryption directly when you install the system. The SUSE installer allows you to create encrypted partitions.

Alternatively you can use `cryptsetup` to encrypt individual partitions: `cryptsetup luksFormat --key-size 256 /dev/sda1` allows you to create an encrypted partition that you can open with `cryptsetup luksOpen /dev/sda1 osd_data`. The resulting `osd_data` device can then be given to SUSE Enterprise Storage to store data that will be transparently encrypted.

This only protects data if the disk is offline. To protect data at rest when the system is running you can use client-side encryption.

- For RADOS Block Device, you can encrypt the block device at the client (such as, `cryptsetup`).
- For CephFS you can use file level encryption (such as, EncFS).
- For RADOS, you need to encrypt the data in the application before storing it via the REST interface.

With this the data is encrypted on the client and is never available in cleartext on SUSE Enterprise Storage. This also protects in case the transport protocol is either not encrypted or an attacker manages to circumvent the transport encryption.

7.2 Data in flight

When creating services on your SUSE Enterprise Storage cluster you should enable encryption if possible. For example, RADOS can communicate via HTTP or HTTPS. Configure it to use HTTPS and use certificates that can be checked by the client.

If possible, do not use self signed certificates. Either use certificates signed by trusted authorities or create your own PKI to issue trusted certificates. Using self signed certificates is better than using plaintext protocols, but it can still allow attackers to get between the communicating nodes.

To secure the communication within the SUSE Enterprise Storage cluster you have several options:

- You can use a dedicated network that is not reachable externally. Depending on your security needs and regulatory requirements that can be acceptable.
- You encrypt the links connecting the SUSE Enterprise Storage machines with an external mechanism, for example using IPsec to setup secure tunnels between the machines that takes care of encryption.
- You use the encryption capabilities in msgr2.

On fresh SUSE Enterprise Storage installs msgr2 is available, which also allows for transport encryption for data. Unfortunately, many clients still do not support this, but Ceph clients using librbd starting with Nautilus can already benefit from this.

The previous message protocol had no guarantee of data authenticity or data encryption. msgr2 uses port 3300, port 6789 is used for the old version. When you want to make sure that only msgr2 is used you can block 6789 to guarantee that the old protocol will not be used.

The default configuration allows SUSE Enterprise Storage to use CRC mode for msgr2. You can check this with:

```
cephuser@adm > for option_name in ms_cluster_mode ms_service_mode ms_client_mode; do
    echo -n "$option_name: "
    ceph config get mon $option_name
done
ms_cluster_mode: crc secure
ms_service_mode: crc secure
ms_client_mode: crc secure
```

Currently only clients build on librbd support secure mode, but the kernel client does not. You can set secure mode only for the cluster internal communication by setting the `ms_cluster_mode` option to `secure`. If you have a client landscape that allows you to enforce secure mode you can also set the `ms_service_mode` and `ms_client_mode` options to `secure`.

This might cause performance issues for your setup, so you need to test this first. If you run into performance issues you can enable secure mode only for select daemons, for example if the ms_cluster_mode option allows you to force secure mode for Ceph Monitor while keeping a different setting for other services.

Glossary

General

Admin node

The host from which you run the Ceph-related commands to administer cluster hosts.

Alertmanager

A single binary which handles alerts sent by the Prometheus server and notifies the end user.

archive sync module

Module that enables creating an Object Gateway zone for keeping the history of S3 object versions.

Bucket

A point that aggregates other nodes into a hierarchy of physical locations.

Ceph Client

The collection of Ceph components which can access a Ceph Storage Cluster. These include the Object Gateway, the Ceph Block Device, the CephFS, and their corresponding libraries, kernel modules, and FUSE clients.

Ceph Dashboard

A built-in Web-based Ceph management and monitoring application to administer various aspects and objects of the cluster. The dashboard is implemented as a Ceph Manager module.

Ceph Manager

Ceph Manager or MGR is the Ceph manager software, which collects all the state from the whole cluster in one place.

Ceph Monitor

Ceph Monitor or MON is the Ceph monitor software.

Ceph Object Storage

The object storage "product", service or capabilities, which consists of a Ceph Storage Cluster and a Ceph Object Gateway.

Ceph OSD Daemon

The **ceph-osd** daemon is the component of Ceph that is responsible for storing objects on a local file system and providing access to them over the network.

Ceph Storage Cluster

The core set of storage software which stores the user's data. Such a set consists of Ceph monitors and OSDs.

ceph-salt

Provides tooling for deploying Ceph clusters managed by cephadm using Salt.

cephadm

cephadm deploys and manages a Ceph cluster by connecting to hosts from the manager daemon via SSH to add, remove, or update Ceph daemon containers.

CephFS

The Ceph file system.

CephX

The Ceph authentication protocol. Cephx operates like Kerberos, but it has no single point of failure.

CRUSH rule

The CRUSH data placement rule that applies to a particular pool or pools.

CRUSH, CRUSH Map

Controlled Replication Under Scalable Hashing: An algorithm that determines how to store and retrieve data by computing data storage locations. CRUSH requires a map of the cluster to pseudo-randomly store and retrieve data in OSDs with a uniform distribution of data across the cluster.

DriveGroups

DriveGroups are a declaration of one or more OSD layouts that can be mapped to physical drives. An OSD layout defines how Ceph physically allocates OSD storage on the media matching the specified criteria.

Grafana

Database analytics and monitoring solution.

Metadata Server

Metadata Server or MDS is the Ceph metadata software.

Multi-zone

Node

Any single machine or server in a Ceph cluster.

Object Gateway

The S3/Swift gateway component for Ceph Object Store. Also known as the RADOS Gateway (RGW).

OSD

Object Storage Device: A physical or logical storage unit.

OSD node

A cluster node that stores data, handles data replication, recovery, backfilling, rebalancing, and provides some monitoring information to Ceph monitors by checking other Ceph OSD daemons.

PG

Placement Group: a sub-division of a *pool*, used for performance tuning.

Point Release

Any ad-hoc release that includes only bug or security fixes.

Pool

Logical partitions for storing objects such as disk images.

Prometheus

Systems monitoring and alerting toolkit.

RADOS Block Device (RBD)

The block storage component of Ceph. Also known as the Ceph block device.

Reliable Autonomic Distributed Object Store (RADOS)

The core set of storage software which stores the user's data (MON + OSD).

Routing tree

A term given to any diagram that shows the various routes a receiver can run.

Rule Set

Rules to determine data placement for a pool.

Samba

Windows integration software.

Samba Gateway

The Samba Gateway joins the Active Directory in the Windows domain to authenticate and authorize users.

zonegroup