# CIS Docker Community Edition Benchmark

v1.1.0 - 07-06-2017

Table of Contents

# Overview

This document is intended for system and application administrators, security specialists, auditors, help desk, and platform deployment personnel who plan to develop, deploy, assess, or secure solutions that incorporate Docker CE 17.06 or later technology.

## Intended Audience

This document, CIS Docker CE 17.06 Benchmark, provides prescriptive guidance for establishing a secure configuration posture for Docker CE container version 17.06. This guide was tested against Docker CE 17.06 on RHEL 7 and Debian 8. To obtain the latest version of this guide, please visit http://benchmarks.cisecurity.org. If you have questions, comments, or have identified ways to improve this guide, please write us at feedback@cisecurity.org.

## Consensus Guidance

This benchmark was created using a consensus review process comprised of subject matter experts. Consensus participants provide perspective from a diverse set of backgrounds including consulting, software development, audit and compliance, security research, operations, government, and legal.

Each CIS benchmark undergoes two phases of consensus review. The first phase occurs during initial benchmark development. During this phase, subject matter experts convene to discuss, create, and test working drafts of the benchmark. This discussion occurs until consensus has been reached on benchmark recommendations. The second phase begins after the benchmark has been published. During this phase, all feedback provided by the Internet community is reviewed by the consensus team for incorporation in the benchmark. If you are interested in participating in the consensus process, please visit https://community.cisecurity.org.

# Typographical Conventions

The following typographical conventions are used throughout this guide:

| Convention | Meaning |
|---|---|
| `Stylized Monospace font` | Used for blocks of code, command, and script examples. Text should be interpreted exactly as presented. |
| Monospace font | Used for inline code, commands, or examples. Text should be interpreted exactly as presented. |
| *<italic font in brackets>* | Italic texts set in angle brackets denote a variable requiring substitution for a real value. |
| *Italic font* | Used to denote the title of a book, article, or other publication. |
| **Note** | Additional information or caveats |

# Scoring Information

A scoring status indicates whether compliance with the given recommendation impacts the assessed target's benchmark score. The following scoring statuses are used in this benchmark:

**Scored**

Failure to comply with "Scored" recommendations will decrease the final benchmark score. Compliance with "Scored" recommendations will increase the final benchmark score.

**Not Scored**

Failure to comply with "Not Scored" recommendations will not decrease the final benchmark score. Compliance with "Not Scored" recommendations will not increase the final benchmark score.

# Profile Definitions

The following configuration profiles are defined by this Benchmark:

- **Level 1 - Docker**

  Items in this profile intend to:

  - Be practical and prudent;
  - Provide a clear security benefit; and
  - Not inhibit the utility of the technology beyond acceptable means.

- **Level 1 - Linux Host OS**

  Items in this profile pertain to the Linux Host OS and intend to:

  - Be practical and prudent;
  - Provide a clear security benefit; and
  - Not inhibit the utility of the technology beyond acceptable means.

- **Level 2 - Docker**

  Items in this profile exhibit one or more of the following characteristics:

  - Are intended for environments or use cases where security is paramount
  - Acts as defense in depth measure
  - May negatively inhibit the utility or performance of the technology

# Acknowledgements

# Recommendations

## *1 Host Configuration*

This section covers security recommendations that you should follow to prepare the host machine that you plan to use for executing containerized workloads. Securing the Docker host and following your infrastructure security best practices would build a solid and secure foundation for executing containerized workloads.

## *1.1 Ensure a separate partition for containers has been created (Scored)*

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

All Docker containers and their data and metadata is stored under `/var/lib/docker` directory. By default, `/var/lib/docker` would be mounted under `/` or `/var` partitions based on availability.

**Rationale:**

Docker depends on `/var/lib/docker` as the default directory where all Docker related files, including the images, are stored. This directory might fill up fast and soon Docker and the host could become unusable. So, it is advisable to create a separate partition (logical volume) for storing Docker files.

**Audit:**

At the Docker host execute the below command:

```
grep /var/lib/docker /etc/fstab
```

This should return the partition details for `/var/lib/docker` mount point.

**Remediation:**

For new installations, create a separate partition for `/var/lib/docker` mount point. For systems that were previously installed, use the Logical Volume Manager (LVM) to create partitions.

**Impact:**

None.

**Default Value:**

By default, `/var/lib/docker` would be mounted under `/` or `/var` partitions based on availability.

**References:**

1. https://www.projectatomic.io/docs/docker-storage-recommendation/

**CIS Controls:**

14 Controlled Access Based on the Need to Know
Controlled Access Based on the Need to Know

## 1.2 Ensure the container host has been Hardened (Not Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Containers run on a Linux host. A container host can run one or more containers. It is of utmost importance to harden the host to mitigate host security misconfiguration.

**Rationale:**

You should follow infrastructure security best practices and harden your host OS. Keeping the host system hardened would ensure that the host vulnerabilities are mitigated. Not hardening the host system could lead to security exposures and breaches.

**Audit:**

Ensure that the host specific security guidelines are followed. Ask the system administrators which security benchmark does current host system comply with. Ensure that the host systems actually comply with that host specific security benchmark.

**Remediation:**

You may consider various CIS Security Benchmarks for your container host. If you have other security guidelines or regulatory requirements to adhere to, please follow them as suitable in your environment.

Additionally, you can run a kernel with `grsecurity` and `PaX`. This would add many safety checks, both at compile-time and run-time. It is also designed to defeat many exploits and has powerful security features. These features do not require Docker-specific configuration, since those security features apply system-wide, independent of containers.

**Impact:**

None.

**Default Value:**

By default, host has factory settings. It is not hardened.

**References:**

1. https://docs.docker.com/engine/security/security/

2. https://learn.cisecurity.org/benchmarks
3. https://docs.docker.com/engine/security/security/#other-kernel-security-features
4. https://grsecurity.net/
5. https://en.wikibooks.org/wiki/Grsecurity
6. https://pax.grsecurity.net/
7. http://en.wikipedia.org/wiki/PaX

**CIS Controls:**

3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

## 1.3 Ensure Docker is up to date (Not Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

There are frequent releases for Docker software that address security vulnerabilities, product bugs and bring in new functionality. Keep a tab on these product updates and upgrade as frequently as when new security vulnerabilities are fixed or deemed correct for your organization.

**Rationale:**

By staying up to date on Docker updates, vulnerabilities in the Docker software can be mitigated. An educated attacker may exploit known vulnerabilities when attempting to attain access or elevate privileges. Not installing regular Docker updates may leave you with running vulnerable Docker software. It might lead to elevation privileges, unauthorized access or other security breaches. Keep a track of new releases and update as necessary.

**Audit:**

Execute the below command and verify that the Docker version is up to date as deemed necessary. It is not a mandate to be on the latest one, though.

```
docker version
```

**Remediation:**

Keep a track of Docker releases and update as necessary.

**Impact:**

Perform a risk assessment for docker version updates in how they may impact your Docker operations. Be aware that several third-party products that use Docker may require older major version of Docker to be supported.

**Default Value:**

Not Applicable

**References:**

1. https://docs.docker.com/engine/installation/
2. https://github.com/moby/moby/releases/latest
3. https://github.com/docker/docker-ce/releases/latest

**CIS Controls:**

4 Continuous Vulnerability Assessment and Remediation
Continuous Vulnerability Assessment and Remediation

## 1.4 Ensure only trusted users are allowed to control Docker daemon (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

The Docker daemon currently requires `root` privileges. A user added to the `docker` group gives him full `root` access rights.

**Rationale:**

Docker allows you to share a directory between the Docker host and a guest container without limiting the access rights of the container. This means that you can start a container and map the `/` directory on your host to the container. The container will then be able to alter your host file system without any restrictions. In simple terms, it means that you can attain elevated privileges with just being a member of the `docker` group and then starting a container with mapped `/` directory on the host.

**Audit:**

Execute the below command on the docker host and ensure that only trusted users are members of the `docker` group.

```
getent group docker
```

**Remediation:**

Remove any users from the `docker` group that are not trusted. Additionally, do not create a mapping of sensitive directories on host to container volumes.

**Impact:**

Rights to build and execute containers as normal user would be restricted.

**Default Value:**

Not Applicable

**References:**

1. https://docs.docker.com/engine/security/security/#docker-daemon-attack-surface
2. https://www.andreas-jung.com/contents/on-docker-security-docker-group-considered-harmful
3. http://www.projectatomic.io/blog/2015/08/why-we-dont-let-non-root-users-run-docker-in-centos-fedora-or-rhel/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 1.5 Ensure auditing is configured for the docker daemon (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit all Docker daemon activities.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit Docker daemon as well. Docker daemon runs with `root` privileges. It is thus necessary to audit its activities and usage.

**Audit:**

Verify that there is an audit rule for Docker daemon. For example, execute below command:

```
auditctl -l | grep /usr/bin/docker
```

This should list a rule for Docker daemon.

**Remediation:**

Add a rule for Docker daemon.

For example,

Add the line as below line in `/etc/audit/audit.rules` file:

```
-w /usr/bin/docker -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker daemon is not audited.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

6.2 Ensure Audit Log Settings Support Appropriate Log Entry Formatting
Validate audit log settings for each hardware device and the software installed on it, ensuring that logs include a date, timestamp, source addresses, destination addresses, and various other useful elements of each packet and/or transaction. Systems should record logs in a standardized format such as syslog entries or those outlined by the Common Event Expression initiative. If systems cannot generate logs in a standardized format, log normalization tools can be deployed to convert logs into such a format.

## 1.6 Ensure auditing is configured for Docker files and directories - /var/lib/docker (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/var/lib/docker`.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/var/lib/docker` is one such directory. It holds all the information about containers. It must be audited.

**Audit:**

Verify that there is an audit rule corresponding to `/var/lib/docker` directory.

For example, execute below command:

```
auditctl -l | grep /var/lib/docker
```

This should list a rule for `/var/lib/docker` directory.

**Remediation:**

Add a rule for `/var/lib/docker` directory.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /var/lib/docker -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited.

**References:**

1. https://access.redhat.com/documentation/en-
US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for
sensitive data.

## 1.7 Ensure auditing is configured for Docker files and directories - /etc/docker (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/etc/docker`.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/etc/docker` is one such directory. It holds various certificates and keys used for TLS communication between Docker daemon and Docker client. It must be audited.

**Audit:**

Verify that there is an audit rule corresponding to `/etc/docker` directory.

For example, execute below command:

```
auditctl -l | grep /etc/docker
```

This should list a rule for `/etc/docker` directory.

**Remediation:**

Add a rule for `/etc/docker` directory.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /etc/docker -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 1.8 Ensure auditing is configured for Docker files and directories - docker.service (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `docker.service`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `docker.service` is one such file. The `docker.service` file might be present if the daemon parameters have been changed by an administrator. It holds various parameters for Docker daemon. It must be audited, if applicable.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.service
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, verify that there is an audit rule corresponding to the file:

For example, execute the below command:

```
auditctl -l | grep docker.service
```

This should list a rule for `docker.service` as per its location.

**Remediation:**

If the file exists, add a rule for it.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /usr/lib/systemd/system/docker.service -k docker
```
Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file `docker.service` may not be available on the system.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 1.9 Ensure auditing is configured for Docker files and directories - docker.socket (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `docker.socket`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `docker.socket` is one such file. It holds various parameters for Docker daemon socket. It must be audited, if applicable.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.socket
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, verify that there is an audit rule corresponding to the file:

For example, execute the below command:

```
auditctl -l | grep docker.socket
```

This should list a rule for `docker.socket` as per its location.

**Remediation:**

If the file exists, add a rule for it.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /usr/lib/systemd/system/docker.socket -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file `docker.socket` may not be available on the system.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 1.10 Ensure auditing is configured for Docker files and directories - /etc/default/docker (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/etc/default/docker`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/etc/default/docker` is one such file. It holds various parameters for Docker daemon. It must be audited, if applicable.

**Audit:**

Verify that there is an audit rule corresponding to `/etc/default/docker` file.

For example, execute below command:

```
auditctl -l | grep /etc/default/docker
```

This should list a rule for `/etc/default/docker` file.

**Remediation:**

Add a rule for `/etc/default/docker` file.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /etc/default/docker -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file
`/etc/default/docker` may not be available on the system. In that case, this
recommendation is not applicable.

**References:**

1. https://access.redhat.com/documentation/en-
   US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for
sensitive data.

## 1.11 Ensure auditing is configured for Docker files and directories - /etc/docker/daemon.json (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/etc/docker/daemon.json`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/etc/docker/daemon.json` is one such file. It holds various parameters for Docker daemon. It must be audited, if applicable.

**Audit:**

Verify that there is an audit rule corresponding to `/etc/docker/daemon.json` file.

For example, execute below command:

```
auditctl -l | grep /etc/docker/daemon.json
```

This should list a rule for `/etc/docker/daemon.json` file.

**Remediation:**

Add a rule for `/etc/docker/daemon.json` file.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /etc/docker/daemon.json -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file `/etc/docker/daemon.json` may not be available on the system. In that case, this recommendation is not applicable.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html
2. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-configuration-file

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 1.12 Ensure auditing is configured for Docker files and directories - /usr/bin/docker-containerd (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/usr/bin/docker-containerd`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/usr/bin/docker-containerd` is one such file. Docker now relies on `containerd`and `runC` to spawn containers. It must be audited, if applicable.

**Audit:**

Verify that there is an audit rule corresponding to `/usr/bin/docker-containerd` file.

For example, execute below command:

```
auditctl -l | grep /usr/bin/docker-containerd
```

This should list a rule for `/usr/bin/docker-containerd` file.

**Remediation:**

Add a rule for `/usr/bin/docker-containerd` file.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /usr/bin/docker-containerd -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file `/usr/bin/docker-containerd` may not be available on the system. In that case, this recommendation is not applicable.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html
2. https://github.com/docker/docker/pull/20662
3. https://containerd.tools/

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 1.13 Ensure auditing is configured for Docker files and directories - /usr/bin/docker-runc (Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Audit `/usr/bin/docker-runc`, if applicable.

**Rationale:**

Apart from auditing your regular Linux file system and system calls, audit all Docker related files and directories. Docker daemon runs with `root` privileges. Its behavior depends on some key files and directories. `/usr/bin/docker-runc` is one such file. Docker now relies on `containerd` and `runC` to spawn containers. It must be audited, if applicable.

**Audit:**

Verify that there is an audit rule corresponding to `/usr/bin/docker-runc` file.

For example, execute below command:

```
auditctl -l | grep /usr/bin/docker-runc
```

This should list a rule for `/usr/bin/docker-runc` file.

**Remediation:**

Add a rule for `/usr/bin/docker-runc` file.

For example,

Add the line as below in `/etc/audit/audit.rules` file:

```
-w /usr/bin/docker-runc -k docker
```

Then, restart the audit daemon. For example,

```
service auditd restart
```

**Impact:**

Auditing generates quite big log files. Ensure to rotate and archive them periodically. Also, create a separate partition of audit to avoid filling root file system.

**Default Value:**

By default, Docker related files and directories are not audited. The file `/usr/bin/docker-runc` may not be available on the system. In that case, this recommendation is not applicable.

**References:**

1. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Security_Guide/chap-system_auditing.html
2. https://github.com/docker/docker/pull/20662
3. https://containerd.tools/
4. https://github.com/opencontainers/runc

**CIS Controls:**

14.6 Enforce Detailed Audit Logging For Sensitive Information
Enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.

## 2 Docker daemon configuration

This section lists the recommendations that alter and secure the behavior of the Docker daemon. The settings that are under this section affect ALL container instances.

**Note:** Docker daemon options can also be controlled using files such as `/etc/sysconfig/docker`, `/etc/default/docker`, the `systemd` unit file or `/etc/docker/daemon.json`. Also, note that Docker in daemon mode can be identified as `/usr/bin/dockerd`, or having `-d` or `daemon` as the argument to `docker` service.

## 2.1 Ensure network traffic is restricted between containers on the default bridge (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, all network traffic is allowed between containers on the same host on the default network bridge. If not desired, restrict all the inter-container communication. Link specific containers together that require communication. Alternatively, you can create custom network and only join containers that need to communicate to that custom network.

**Rationale:**

By default, unrestricted network traffic is enabled between all containers on the same host on the default network bridge. Thus, each container has the potential of reading all packets across the container network on the same host. This might lead to an unintended and unwanted disclosure of information to other containers. Hence, restrict the inter-container communication on the default network bridge.

**Audit:**

Run the below command and verify that the default network bridge has been configured to restrict inter-container communication.

```
docker network ls --quiet | xargs docker network inspect --format '{{ .Name }}: {{ .Options }}'
```

It should return `com.docker.network.bridge.enable_icc:false` for the default network bridge.

**Remediation:**

Run the docker in daemon mode and pass `--icc=false` as an argument.

For Example,

```
dockerd --icc=false
```

Alternatively, you can follow the Docker documentation and create a custom network and only join containers that need to communicate to that custom network. The `--icc` parameter only applies to the default docker bridge, if custom networks are used then the approach of segmenting networks should be adopted instead.

**Impact:**

The inter-container communication would be disabled on the default network bridge. If any communication between containers on the same host is desired, then it needs to be explicitly defined using container linking or alternatively custom networks have to be defined.

**Default Value:**

By default, all inter-container communication is allowed on the default network bridge.

**References:**

1. https://docs.docker.com/engine/userguide/networking/
2. https://docs.docker.com/engine/userguide/networking/default_network/container-communication/#communication-between-containers

## 2.2 Ensure the logging level is set to 'info' (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Set Docker daemon log level to `info`.

**Rationale:**

Setting up an appropriate log level, configures the Docker daemon to log events that you would want to review later. A base log level of `info` and above would capture all logs except debug logs. Until and unless required, you should not run Docker daemon at `debug` log level.

**Audit:**

```
ps -ef | grep docker
```

Ensure that either the `--log-level` parameter is not present or if present, then it is set to `info`.

**Remediation:**

Run the Docker daemon as below:

```
dockerd --log-level="info"
```

**Impact:**

None.

**Default Value:**

By default, Docker daemon is set to log level of `info`.

**References:**

1. https://docs.docker.com/edge/engine/reference/commandline/dockerd/

**CIS Controls:**

6.2 Ensure Audit Log Settings Support Appropriate Log Entry Formatting
Validate audit log settings for each hardware device and the software installed on it,

ensuring that logs include a date, timestamp, source addresses, destination addresses, and various other useful elements of each packet and/or transaction. Systems should record logs in a standardized format such as syslog entries or those outlined by the Common Event Expression initiative. If systems cannot generate logs in a standardized format, log normalization tools can be deployed to convert logs into such a format.

## 2.3 Ensure Docker is allowed to make changes to iptables (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Iptables are used to set up, maintain, and inspect the tables of IP packet filter rules in the Linux kernel. Allow the Docker daemon to make changes to the `iptables`.

**Rationale:**

Docker will never make changes to your system `iptables` rules if you choose to do so. Docker server would automatically make the needed changes to iptables based on how you choose your networking options for the containers if it is allowed to do so. It is recommended to let Docker server make changes to `iptables` automatically to avoid networking misconfiguration that might hamper the communication between containers and to the outside world. Additionally, it would save you hassles of updating `iptables` every time you choose to run the containers or modify networking options.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--iptables` parameter is either not present or not set to `false`.

**Remediation:**

Do not run the Docker daemon with `--iptables=false` parameter. For example, do not start the Docker daemon as below:

```
dockerd --iptables=false
```

**Impact:**

Docker daemon service start requires iptables rules to be enabled before it starts. Any restarts of iptables during docker daemon operation may result in losing docker-created rules. Adding iptables-persistent to your iptables install can assist with mitigating this impact.

**Default Value:**

By default, `iptables` is set to `true`.

**References:**

1. https://docs.docker.com/engine/userguide/networking/default_network/container-communication/
2. https://fralef.me/docker-and-iptables.html

**CIS Controls:**

5 <u>Controlled Use of Administration Privileges</u>
Controlled Use of Administration Privileges

## 2.4 Ensure insecure registries are not used (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Docker considers a private registry either secure or insecure. By default, registries are considered secure.

**Rationale:**

A secure registry uses TLS. A copy of registry's CA certificate is placed on the Docker host at `/etc/docker/certs.d/<registry-name>/` directory. An insecure registry is the one not having either valid registry certificate or is not using TLS. You should not be using any insecure registries in the production environment. Insecure registries can be tampered with leading to possible compromise to your production system.

Additionally, If a registry is marked as insecure then `docker pull`, `docker push`, and `docker search` commands will not result in an error message and the user might be indefinitely working with insecure registries without ever being notified of potential danger.

**Audit:**

Run `docker info` or execute the below command to find out if any insecure registries are used:

```
ps -ef | grep dockerd
```

Ensure that the `--insecure-registry` parameter is not present.

**Remediation:**

Do not use any insecure registries.

For example, do not start the Docker daemon as below:

```
dockerd --insecure-registry 10.1.0.0/16
```

**Impact:**

None.

**Default Value:**

By default, Docker assumes all, but local, registries are secure.

**References:**

1. https://docs.docker.com/registry/insecure/

**CIS Controls:**

14.2 Encrypt All Sensitive Information Over Less-trusted Networks
All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## 2.5 Ensure aufs storage driver is not used (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Do not use `aufs` as storage driver for your Docker instance.

**Rationale:**

The `aufs` storage driver is the oldest storage driver. It is based on a Linux kernel patch-set that is unlikely to be merged into the main Linux kernel. `aufs` driver is also known to cause some serious kernel crashes. `aufs` just has legacy support from Docker. Most importantly, `aufs` is not a supported driver in many Linux distributions using latest Linux kernels.

**Audit:**

Execute the below command and verify that `aufs` is not used as storage driver:

```
docker info | grep -e "^Storage Driver:\s*aufs\s*$"
```

The above command should not return anything.

**Remediation:**

Do not explicitly use `aufs` as storage driver.

For example, do not start Docker daemon as below:

```
dockerd --storage-driver aufs
```

**Impact:**

`aufs` is the only storage driver that allows containers to share executable and shared library memory. It might be useful if you are running thousands of containers with the same program or libraries.

**Default Value:**

By default, Docker uses `devicemapper` as the storage driver on most of the platforms. Default storage driver can vary based on your OS vendor. You should use the storage driver that is best supported by your preferred vendor.

**References:**

1. https://docs.docker.com/engine/userguide/storagedriver/selectadriver/#support ed-backing-filesystems
2. http://muehe.org/posts/switching-docker-from-aufs-to-devicemapper/
3. http://jpetazzo.github.io/assets/2015-03-05-deep-dive-into-docker-storage-drivers.html#1
4. https://docs.docker.com/engine/userguide/storagedriver/

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.6 Ensure TLS authentication for Docker daemon is configured (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

It is possible to make the Docker daemon to listen on a specific IP and port and any other Unix socket other than default Unix socket. Configure TLS authentication to restrict access to Docker daemon via IP and port.

**Rationale:**

By default, Docker daemon binds to a non-networked Unix socket and runs with `root` privileges. If you change the default docker daemon binding to a TCP port or any other Unix socket, anyone with access to that port or socket can have full access to Docker daemon and in turn to the host system. Hence, you should not bind the Docker daemon to another IP/port or a Unix socket.

If you must expose the Docker daemon via a network socket, configure TLS authentication for the daemon and Docker Swarm APIs (if using). This would restrict the connections to your Docker daemon over the network to a limited number of clients who could successfully authenticate over TLS.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the below parameters are present:

```
--tlsverify
--tlscacert
--tlscert
--tlskey
```

**Remediation:**

Follow the steps mentioned in the Docker documentation or other references.

**Impact:**

You would need to manage and guard certificates and keys for Docker daemon and Docker clients.

**Default Value:**

By default, TLS authentication is not configured.

**References:**

1. https://docs.docker.com/engine/security/https/

**CIS Controls:**

9.1 <u>Limit Open Ports, Protocols, and Services</u>
Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 2.7 Ensure the default ulimit is configured appropriately (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Set the default ulimit options as appropriate in your environment.

**Rationale:**

`ulimit` provides control over the resources available to the shell and to processes started by it. Setting system resource limits judiciously saves you from many disasters such as a fork bomb. Sometimes, even friendly users and legitimate processes can overuse system resources and in-turn can make the system unusable.

Setting default ulimit for the Docker daemon would enforce the ulimit for all container instances. You would not need to setup ulimit for each container instance. However, the default ulimit can be overridden during container runtime, if needed. Hence, to control the system resources, define a default ulimit as needed in your environment.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--default-ulimit` parameter is set as appropriate.

**Remediation:**

Run the docker in daemon mode and pass `--default-ulimit` as argument with respective ulimits as appropriate in your environment.

For Example,

```
dockerd --default-ulimit nproc=1024:2048 --default-ulimit nofile=100:200
```

**Impact:**

If the ulimits are not set properly, the desired resource control might not be achieved and might even make the system unusable.

**Default Value:**

By default, no ulimit is set.

**References:**

1. https://docs.docker.com/edge/engine/reference/commandline/dockerd/#default-ulimits

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.8 Enable user namespace support (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Enable user namespace support in Docker daemon to utilize container user to host user re-mapping. This recommendation is beneficial where containers you are using do not have an explicit container user defined in the container image. If container images that you are using have a pre-defined non-root user, this recommendation may be skipped since this feature is still in its infancy and might give you unpredictable issues and complexities.

**Rationale:**

The Linux kernel user namespace support in Docker daemon provides additional security for the Docker host system. It allows a container to have a unique range of user and group IDs which are outside the traditional user and group range utilized by the host system.

For example, the root user will have expected administrative privilege inside the container but can effectively be mapped to an unprivileged UID on the host system.

**Audit:**

```
ps -p $(docker inspect --format='{{ .State.Pid }}' <CONTAINER ID>) -o
pid,user
```

The above command would find the PID of the container and then would list the host user associated with the container process. If the container process is running as `root`, then this recommendation is non-compliant.

Alternatively, you can run `docker info` to ensure that the `userns` is listed under `Security Options`:

```
docker info --format '{{ .SecurityOptions }}'
```

**Remediation:**

Please consult Docker documentation for various ways in which this can be configured depending upon your requirements. Your steps might also vary based on platform - For example, on Red Hat, sub-UIDs and sub-GIDs mapping creation does not work automatically. You might have to create your own mapping.

However, the high-level steps are as below:

**Step 1:** Ensure that the files `/etc/subuid` and `/etc/subgid` exist.

```
touch /etc/subuid /etc/subgid
```

**Step 2:** Start the docker daemon with `--userns-remap` flag

```
dockerd --userns-remap=default
```

**Impact:**

User namespace remapping makes quite a few Docker features incompatible and also currently breaks a few functionalities. Check out the Docker documentation and referenced links for details.

**Default Value:**

By default, user namespace is not remapped.

**References:**

1. http://man7.org/linux/man-pages/man7/user_namespaces.7.html
2. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-user-namespace-options
3. http://events.linuxfoundation.org/sites/events/files/slides/User%20Namespaces%20-%20ContainerCon%202015%20-%2016-9-final_0.pdf

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.9 Ensure the default cgroup usage has been confirmed (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

The `--cgroup-parent` option allows you to set the default cgroup parent to use for all the containers. If there is no specific use case, this setting should be left at its default.

**Rationale:**

System administrators typically define cgroups under which containers are supposed to run. Even if cgroups are not explicitly defined by the system administrators, containers run under `docker` cgroup by default.

It is possible to attach to a different cgroup other than that is the default. This usage should be monitored and confirmed. By attaching to a different cgroup than the one that is a default, it is possible to share resources unevenly and thus might starve the host for resources.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--cgroup-parent` parameter is either not set or is set as appropriate non-default cgroup.

**Remediation:**

The default setting is good enough and can be left as-is. If you want to specifically set a non-default cgroup, pass `--cgroup-parent` parameter to the docker daemon when starting it.

For Example,

```
dockerd --cgroup-parent=/foobar
```

**Impact:**

None.

**Default Value:**

By default, docker daemon uses `/docker` for fs cgroup driver and `system.slice` for systemd cgroup driver.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#default-cgroup-parent

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.10 Ensure base device size is not changed until needed (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

In certain circumstances, you might need containers bigger than 10G in size. In these cases, carefully choose the base device size.

**Rationale:**

The base device size can be increased at daemon restart. Increasing the base device size allows all future images and containers to be of the new base device size. A user can use this option to expand the base device size however shrinking is not permitted. This value affects the system-wide "base" empty filesystem that may already be initialized and inherited by pulled images.

Though the file system does not allot the increased size if it is empty, it will use more space for the empty case depending upon the device size. This may cause a denial of service by ending up in file system being over-allocated or full.

**Audit:**

```
ps -ef | grep dockerd
```

Execute the above command and it should not show any `--storage-opt dm.basesize` parameters.

**Remediation:**

Do not set `--storage-opt dm.basesize` until needed.

**Impact:**

None.

**Default Value:**

The default base device size is 10G.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#storage-driver-options

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.11 Ensure that authorization for Docker client commands is enabled (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Use native Docker authorization plugins or a third party authorization mechanism with Docker daemon to manage access to Docker client commands.

**Rationale:**

Docker's out-of-the-box authorization model is all or nothing. Any user with permission to access the Docker daemon can run any Docker client command. The same is true for callers using Docker's remote API to contact the daemon. If you require greater access control, you can create authorization plugins and add them to your Docker daemon configuration. Using an authorization plugin, a Docker administrator can configure granular access policies for managing access to Docker daemon.

Third party integrations of Docker may implement their own authorization models to require authorization with the Docker daemon outside of docker's native authorization plugin (i.e. Kubernetes, Cloud Foundry, Openshift).

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--authorization-plugin` parameter is set as appropriate if using docker native authorization.

```
docker search hello-world
```

Ensure that docker daemon requires authorization to perform the above command.

**Remediation:**

**Step 1**: Install/Create an authorization plugin.

**Step 2**: Configure the authorization policy as desired.

**Step 3**: Start the docker daemon as below:

```
dockerd --authorization-plugin=<PLUGIN_ID>
```

**Impact:**

Each docker command specifically passes through authorization plugin mechanism. This might introduce a slight performance drop.

Third party use of alternative container engines that utilize the docker daemon may provide alternative mechanisms to provide this security control.

**Default Value:**

By default, authorization plugins are not set up.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#access-authorization
2. https://docs.docker.com/engine/extend/plugins_authorization/
3. https://github.com/twistlock/authz

**Notes:**

As a scored control, focus should be on a PASS/FAIL if the authentication occurs when a docker client command is executed against docker daemon to enforce authentication. The native docker authentication plugin is just *one* method to enforce this control.

**CIS Controls:**

16 Account Monitoring and Control
Account Monitoring and Control

## *2.12 Ensure centralized and remote logging is configured (Scored)*

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Docker now supports various log drivers. A preferable way to store logs is the one that supports centralized and remote logging.

**Rationale:**

Centralized and remote logging ensures that all important log records are safe despite catastrophic events. Docker now supports various such logging drivers. Use the one that suits your environment the best.

**Audit:**

Run `docker info` and ensure that the `Logging Driver` property set as appropriate.

```
docker info --format '{{ .LoggingDriver }}'
```

Alternatively, the below command would give you the `--log-driver` setting, if configured. Ensure that it is set as appropriate.

```
ps -ef | grep dockerd
```

**Remediation:**

**Step 1**: Setup the desired log driver by following its documentation.

**Step 2**: Start the docker daemon with that logging driver.

For example,

```
dockerd --log-driver=syslog --log-opt syslog-address=tcp://192.xxx.xxx.xxx
```

**Impact:**

None.

**Default Value:**

By default, container logs are maintained as json files

**References:**

1. https://docs.docker.com/engine/admin/logging/overview/

**CIS Controls:**

6.6 <u>Deploy A SIEM OR Log Analysis Tools For Aggregation And Correlation/Analysis</u>
Deploy a SIEM (Security Information and Event Management) or log analytic tools for log aggregation and consolidation from multiple machines and for log correlation and analysis. Using the SIEM tool, system administrators and security personnel should devise profiles of common events from given systems so that they can tune detection to focus on unusual activity, avoid false positives, more rapidly identify anomalies, and prevent overwhelming analysts with insignificant alerts.

## 2.13 Ensure operations on legacy registry (v1) are Disabled (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The latest Docker registry is v2. All operations on the legacy registry version (v1) should be restricted.

**Rationale:**

Docker registry v2 brings in many performance and security improvements over v1. It supports container image provenance and other security features such as image signing and verification. Hence, operations on Docker legacy registry should be restricted.

**Audit:**

```
ps -ef | grep dockerd
```

The above command should list `--disable-legacy-registry` as an option passed to the docker daemon.

**Remediation:**

Start the docker daemon as below:

```
dockerd --disable-legacy-registry
```

**Impact:**

Legacy registry operations would be restricted.

**Default Value:**

By default, legacy registry operations are allowed.

**References:**

1. https://docs.docker.com/edge/engine/reference/commandline/dockerd/#legacy-registries
2. https://docs.docker.com/registry/spec/api/
3. https://the.binbashtheory.com/creating-private-docker-registry-2-0-with-token-authentication-service/

4. https://blog.docker.com/2015/07/new-tool-v1-registry-docker-trusted-registry-v2-open-source/
5. http://www.slideshare.net/Docker/docker-registry-v2

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.14 Ensure live restore is Enabled (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The `--live-restore` enables full support of daemon-less containers in docker. It ensures that docker does not stop containers on shutdown or restore and properly reconnects to the container when restarted.

**Rationale:**

One of the important security triads is availability. Setting `--live-restore` flag in the docker daemon ensures that container execution is not interrupted when the docker daemon is not available. This also means that it is now easier to update and patch the docker daemon without execution downtime.

**Audit:**

Run `docker info` and ensure that the `Live Restore Enabled` property is set to `true`.

```
docker info --format '{{ .LiveRestoreEnabled }}'
```

Alternatively run the below command and ensure that --live-restore is used.

```
ps -ef | grep dockerd
```

**Remediation:**

Run the docker in daemon mode and pass `--live-restore` as an argument.

For Example,

```
dockerd --live-restore
```

**Impact:**

None.

**Default Value:**

By default, `--live-restore` is not enabled.

**References:**

1. https://docs.docker.com/engine/admin/live-restore/

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.15 Ensure Userland Proxy is Disabled (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The docker daemon starts a userland proxy service for port forwarding whenever a port is exposed. Where hairpin NAT is available, this service is generally superfluous to requirements and can be disabled.

**Rationale:**

Docker engine provides two mechanisms for forwarding ports from the host to containers, hairpin NAT, and a userland proxy. In most circumstances, the hairpin NAT mode is preferred as it improves performance and makes use of native Linux iptables functionality instead of an additional component.

Where hairpin NAT is available, the userland proxy should be disabled on startup to reduce the attack surface of the installation.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--userland-proxy` parameter is set to `false`.

**Remediation:**

Run the Docker daemon as below:

```
dockerd --userland-proxy=false
```

**Impact:**

Some systems with older Linux kernels may not be able to support hairpin NAT and therefore require the userland proxy service. Also, some networking setups can be impacted by the removal of the userland proxy.

**Default Value:**

By default, the userland proxy is enabled.

**References:**

1. http://windsock.io/the-docker-proxy/
2. https://github.com/docker/docker/issues/14856
3. https://github.com/docker/docker/issues/22741
4. https://docs.docker.com/engine/userguide/networking/default_network/binding/

**CIS Controls:**

9.1 Limit Open Ports, Protocols, and Services

Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 2.16 Ensure daemon-wide custom seccomp profile is applied, if needed (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

You can choose to apply your custom seccomp profile at the daemon-wide level if needed and override Docker's default seccomp profile.

**Rationale:**

A large number of system calls are exposed to every userland process with many of them going unused for the entire lifetime of the process. Most of the applications do not need all the system calls and thus benefit by having a reduced set of available system calls. The reduced set of system calls reduces the total kernel surface exposed to the application and thus improvises application security.

You could apply your own custom seccomp profile instead of Docker's default seccomp profile. Alternatively, if Docker's default profile is good for your environment, you can choose to ignore this recommendation.

**Audit:**

Run the below command and review the seccomp profile listed in the `Security Options` section. If it is `default`, that means, Docker's default seccomp profile is applied.

```
docker info --format '{{ .SecurityOptions }}'
```

**Remediation:**

By default, Docker's default seccomp profile is applied. If this is good for your environment, no action is necessary. Alternatively, if you choose to apply your own seccomp profile, use the `--seccomp-profile` flag at daemon start or put it in the daemon runtime parameters file.

```
dockerd --seccomp-profile </path/to/seccomp/profile>
```

**Impact:**

A misconfigured seccomp profile could possibly interrupt your container environment. Docker-default blocked calls have been carefully scrutinized. These address some critical

vulnerabilities/issues within container environments (for example, kernel key ring calls). So, you should be very careful while overriding the defaults.

**Default Value:**

By default, Docker applies a seccomp profile.

**References:**

1. https://docs.docker.com/engine/security/seccomp/
2. https://github.com/docker/docker/pull/26276

**CIS Controls:**

18 Application Software Security
Application Software Security

## 2.17 Ensure experimental features are avoided in production (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Avoid experimental features in production.

**Rationale:**

Experimental is now a runtime docker daemon flag instead of a separate build. Passing `--experimental` as a runtime flag to the docker daemon, activates experimental features. Experimental is now considered a stable release, but with a couple of features which might not have tested and guaranteed API stability.

**Audit:**

Run the below command and ensure that the `Experimental` property is set to `false` in the Server section.

```
docker version --format '{{ .Server.Experimental }}'
```

**Remediation:**

Do not pass `--experimental` as a runtime parameter to the docker daemon.

**Impact:**

None

**Default Value:**

By default, experimental features are not activated on the docker daemon.

**References:**

1. https://docs.docker.com/edge/engine/reference/commandline/dockerd/#options

**CIS Controls:**

18 <u>Application Software Security</u>
Application Software Security

## 2.18 Ensure containers are restricted from acquiring new privileges (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Restrict containers from acquiring additional privileges via suid or sgid bits, by default.

**Rationale:**

A process can set the `no_new_priv` bit in the kernel. It persists across fork, clone and execve. The `no_new_priv` bit ensures that the process or its children processes do not gain any additional privileges via suid or sgid bits. This way a lot of dangerous operations become a lot less dangerous because there is no possibility of subverting privileged binaries.

Setting this at the daemon level ensures that by default all new containers are restricted from acquiring new privileges.

**Audit:**

```
ps -ef | grep dockerd
```

Ensure that the `--no-new-privileges` parameter is present and is not set to `false`.

**Remediation:**

Run the Docker daemon as below:

```
dockerd --no-new-privileges
```

**Impact:**

`no_new_priv` prevents LSMs like SELinux from transitioning to process labels that have access not allowed to the current process.

**Default Value:**

By default, containers are not restricted from acquiring new privileges.

**References:**

1. https://github.com/moby/moby/pull/29984
2. https://github.com/moby/moby/pull/20727

**CIS Controls:**

5 Controlled Use of Administration Privileges
Controlled Use of Administration Privileges

# 3 Docker daemon configuration files

This section covers Docker related files and directory permissions and ownership. Keeping the files and directories, that may contain sensitive parameters, secure is important for correct and secure functioning of Docker daemon.

## 3.1 Ensure that docker.service file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `docker.service` file ownership and group-ownership are correctly set to `root`.

**Rationale:**

`docker.service` file contains sensitive parameters that may alter the behavior of Docker daemon. Hence, it should be owned and group-owned by `root` to maintain the integrity of the file.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.service
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to verify that the file is owned and group-owned by `root`.

For example,

```
stat -c %U:%G /usr/lib/systemd/system/docker.service | grep -v root:root
```

The above command should not return anything.

**Remediation:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.service
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to set the ownership and group ownership for the file to `root`.

For example,

```
chown root:root /usr/lib/systemd/system/docker.service
```

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable. By default, if the file is present, the ownership and group-ownership for this file is correctly set to `root`.

**References:**

1. https://docs.docker.com/engine/admin/systemd/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.2 Ensure that docker.service file permissions are set to 644 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `docker.service` file permissions are correctly set to `644` or more restrictive.

**Rationale:**

`docker.service` file contains sensitive parameters that may alter the behavior of Docker daemon. Hence, it should not be writable by any other user other than `root` to maintain the integrity of the file.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.service
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to verify that the file permissions are set to `644` or more restrictive.

For example,

```
stat -c %a /usr/lib/systemd/system/docker.service
```

**Remediation:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.service
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to set the file permissions to `644`.

For example,

```
chmod 644 /usr/lib/systemd/system/docker.service
```

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable. By default, if the file is present, the file permissions are correctly set to `644`.

**References:**

1. https://docs.docker.com/articles/systemd/

**CIS Controls:**

14.4 Protect Information With Access Control Lists
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.3 Ensure that docker.socket file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `docker.socket` file ownership and group ownership is correctly set to `root`.

**Rationale:**

`docker.socket` file contains sensitive parameters that may alter the behavior of Docker remote API. Hence, it should be owned and group-owned by `root` to maintain the integrity of the file.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.socket
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to verify that the file is owned and group-owned by `root`.

For example,

```
stat -c %U:%G /usr/lib/systemd/system/docker.socket | grep -v root:root
```

The above command should not return anything.

**Remediation:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.socket
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to set the ownership and group ownership for the file to `root`.

For example,

```
chown root:root /usr/lib/systemd/system/docker.socket
```

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable. By default, if the file is present, the ownership and group-ownership for this file is correctly set to `root`.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-socket-option
2. https://github.com/docker/docker-ce/blob/master/components/packaging/deb/systemd/docker.socket

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.4 Ensure that docker.socket file permissions are set to 644 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `docker.socket` file permissions are correctly set to `644` or more restrictive.

**Rationale:**

`docker.socket` file contains sensitive parameters that may alter the behavior of Docker remote API. Hence, it should be writable only by `root` to maintain the integrity of the file.

**Audit:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.socket
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to verify that the file permissions are set to `644` or more restrictive.

For example,

```
stat -c %a /usr/lib/systemd/system/docker.socket
```

**Remediation:**

**Step 1**: Find out the file location:

```
systemctl show -p FragmentPath docker.socket
```

**Step 2**: If the file does not exist, this recommendation is not applicable. If the file exists, execute the below command with the correct file path to set the file permissions to `644`.

For example,

```
chmod 644 /usr/lib/systemd/system/docker.socket
```

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable. By default, if the file is present, the file permissions for this file are correctly set to `644`.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#bind-docker-to-another-hostport-or-a-unix-socket
2. https://github.com/YungSang/fedora-atomic-packer/blob/master/oem/docker.socket
3. http://daviddaeschler.com/2014/12/14/centos-7rhel-7-and-docker-containers-on-boot/

**CIS Controls:**

14.4 Protect Information With Access Control Lists
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.5 Ensure that /etc/docker directory ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `/etc/docker` directory ownership and group-ownership is correctly set to `root`.

**Rationale:**

`/etc/docker` directory contains certificates and keys in addition to various sensitive files. Hence, it should be owned and group-owned by `root` to maintain the integrity of the directory.

**Audit:**

Execute the below command to verify that the directory is owned and group-owned by `root`:

```
stat -c %U:%G /etc/docker | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root /etc/docker
```

This would set the ownership and group-ownership for the directory to `root`.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for this directory is correctly set to `root`.

**References:**

1. https://docs.docker.com/engine/security/https/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.6 Ensure that /etc/docker directory permissions are set to 755 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `/etc/docker` directory permissions are correctly set to `755` or more restrictive.

**Rationale:**

`/etc/docker` directory contains certificates and keys in addition to various sensitive files. Hence, it should only be writable by `root` to maintain the integrity of the directory.

**Audit:**

Execute the below command to verify that the directory has permissions of `755` or more restrictive:

```
stat -c %a /etc/docker
```

**Remediation:**

```
chmod 755 /etc/docker
```

This would set the permissions for the directory to `755`.

**Impact:**

None.

**Default Value:**

By default, the permissions for this directory are correctly set to `755`.

**References:**

1. https://docs.docker.com/engine/security/https/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.7 Ensure that registry certificate file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that all the registry certificate files (usually found under `/etc/docker/certs.d/<registry-name>` directory) are owned and group-owned by `root`.

**Rationale:**

`/etc/docker/certs.d/<registry-name>` directory contains Docker registry certificates. These certificate files must be owned and group-owned by `root` to maintain the integrity of the certificates.

**Audit:**

Execute the below command to verify that the registry certificate files are owned and group-owned by `root`:

```
stat -c %U:%G /etc/docker/certs.d/* | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root /etc/docker/certs.d/<registry-name>/*
```

This would set the ownership and group-ownership for the registry certificate files to `root`.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for registry certificate files is correctly set to `root`.

**References:**

1. https://docs.docker.com/registry/insecure/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.8 Ensure that registry certificate file permissions are set to 444 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that all the registry certificate files (usually found under `/etc/docker/certs.d/<registry-name>` directory) have permissions of `444` or more restrictive.

**Rationale:**

`/etc/docker/certs.d/<registry-name>` directory contains Docker registry certificates. These certificate files must have permissions of `444` to maintain the integrity of the certificates.

**Audit:**

Execute the below command to verify that the registry certificate files have permissions of `444` or more restrictive:

```
stat -c %a /etc/docker/certs.d/<registry-name>/*
```

**Remediation:**

```
chmod 444 /etc/docker/certs.d/<registry-name>/*
```

This would set the permissions for registry certificate files to `444`.

**Impact:**

None.

**Default Value:**

By default, the permissions for registry certificate files might not be `444`. The default file permissions are governed by the system or user specific `umask` values.

**References:**

1. https://docs.docker.com/registry/insecure/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.9 Ensure that TLS CA certificate file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the TLS CA certificate file (the file that is passed alongwith `--tlscacert` parameter) is owned and group-owned by `root`.

**Rationale:**

The TLS CA certificate file should be protected from any tampering. It is used to authenticate Docker server based on given CA certificate. Hence, it must be owned and group-owned by `root` to maintain the integrity of the CA certificate.

**Audit:**

Execute the below command to verify that the TLS CA certificate file is owned and group-owned by `root`:

```
stat -c %U:%G <path to TLS CA certificate file> | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root <path to TLS CA certificate file>
```

This would set the ownership and group-ownership for the TLS CA certificate file to `root`.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for TLS CA certificate file is correctly set to `root`.

**References:**

1. https://docs.docker.com/registry/insecure/

2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.10 Ensure that TLS CA certificate file permissions are set to 444 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the TLS CA certificate file (the file that is passed alongwith `--tlscacert` parameter) has permissions of `444` or more restrictive.

**Rationale:**

The TLS CA certificate file should be protected from any tampering. It is used to authenticate Docker server based on given CA certificate. Hence, it must have permissions of `444` to maintain the integrity of the CA certificate.

**Audit:**

Execute the below command to verify that the TLS CA certificate file has permissions of `444` or more restrictive:

```
stat -c %a <path to TLS CA certificate file>
```

**Remediation:**

```
chmod 444 <path to TLS CA certificate file>
```

This would set the file permissions of the TLS CA file to `444`.

**Impact:**

None.

**Default Value:**

By default, the permissions for TLS CA certificate file might not be `444`. The default file permissions are governed by the system or user specific `umask` values.

**References:**

1. https://docs.docker.com/registry/insecure/
2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.11 Ensure that Docker server certificate file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker server certificate file (the file that is passed alongwith `--tlscert` parameter) is owned and group-owned by `root`.

**Rationale:**

The Docker server certificate file should be protected from any tampering. It is used to authenticate Docker server based on the given server certificate. Hence, it must be owned and group-owned by `root` to maintain the integrity of the certificate.

**Audit:**

Execute the below command to verify that the Docker server certificate file is owned and group-owned by `root`:

```
stat -c %U:%G <path to Docker server certificate file> | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root <path to Docker server certificate file>
```

This would set the ownership and group-ownership for the Docker server certificate file to `root`.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for Docker server certificate file is correctly set to `root`.

**References:**

1. https://docs.docker.com/registry/insecure/
2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.12 Ensure that Docker server certificate file permissions are set to 444 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker server certificate file (the file that is passed alongwith `--tlscert` parameter) has permissions of `444` or more restrictive.

**Rationale:**

The Docker server certificate file should be protected from any tampering. It is used to authenticate Docker server based on the given server certificate. Hence, it must have permissions of `444` to maintain the integrity of the certificate.

**Audit:**

Execute the below command to verify that the Docker server certificate file has permissions of `444` or more restrictive:

```
stat -c %a <path to Docker server certificate file>
```

**Remediation:**

```
chmod 444 <path to Docker server certificate file>
```

This would set the file permissions of the Docker server file to `444`.

**Impact:**

None.

**Default Value:**

By default, the permissions for Docker server certificate file might not be `444`. The default file permissions are governed by the system or user specific `umask` values.

**References:**

1. https://docs.docker.com/registry/insecure/
2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.13 Ensure that Docker server certificate key file ownership is set to root:root (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker server certificate key file (the file that is passed alongwith `--tlskey` parameter) is owned and group-owned by `root`.

**Rationale:**

The Docker server certificate key file should be protected from any tampering or unneeded reads. It holds the private key for the Docker server certificate. Hence, it must be owned and group-owned by `root` to maintain the integrity of the Docker server certificate.

**Audit:**

Execute the below command to verify that the Docker server certificate key file is owned and group-owned by `root`:

```
stat -c %U:%G <path to Docker server certificate key file> | grep -v
root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root <path to Docker server certificate key file>
```

This would set the ownership and group-ownership for the Docker server certificate key file to `root`.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for Docker server certificate key file is correctly set to `root`.

**References:**

1. https://docs.docker.com/registry/insecure/
2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.14 Ensure that Docker server certificate key file permissions are set to 400 (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker server certificate key file (the file that is passed alongwith `--tlskey` parameter) has permissions of `400`.

**Rationale:**

The Docker server certificate key file should be protected from any tampering or unneeded reads. It holds the private key for the Docker server certificate. Hence, it must have permissions of `400` to maintain the integrity of the Docker server certificate.

**Audit:**

Execute the below command to verify that the Docker server certificate key file has permissions of `400`:

```
stat -c %a <path to Docker server certificate key file>
```

**Remediation:**

```
chmod 400 <path to Docker server certificate key file>
```

This would set the Docker server certificate key file permissions to `400`.

**Impact:**

None.

**Default Value:**

By default, the permissions for Docker server certificate key file might not be `400`. The default file permissions are governed by the system or user specific `umask` values.

**References:**

1. https://docs.docker.com/registry/insecure/
2. https://docs.docker.com/engine/security/https/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 3.15 Ensure that Docker socket file ownership is set to root:docker (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker socket file is owned by `root` and group-owned by `docker`.

**Rationale:**

Docker daemon runs as `root`. The default Unix socket hence must be owned by `root`. If any other user or process owns this socket, then it might be possible for that non-privileged user or process to interact with Docker daemon. Also, such a non-privileged user or process might interact with containers. This is neither secure nor desired behavior.

Additionally, the Docker installer creates a Unix group called `docker`. You can add users to this group, and then those users would be able to read and write to default Docker Unix socket. The membership to the `docker` group is tightly controlled by the system administrator. If any other group owns this socket, then it might be possible for members of that group to interact with Docker daemon. Also, such a group might not be as tightly controlled as the `docker` group. This is neither secure nor desired behavior.

Hence, the default Docker Unix socket file must be owned by `root` and group-owned by `docker` to maintain the integrity of the socket file.

**Audit:**

Execute the below command to verify that the Docker socket file is owned by `root` and group-owned by `docker`:

```
stat -c %U:%G /var/run/docker.sock | grep -v root:docker
```

The above command should not return anything.

**Remediation:**

```
chown root:docker /var/run/docker.sock
```

This would set the ownership to `root` and group-ownership to `docker` for default Docker socket file.

**Impact:**

None.

**Default Value:**

By default, the ownership and group-ownership for Docker socket file is correctly set to `root:docker`.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-socket-option
2. https://docs.docker.com/engine/reference/commandline/dockerd/#bind-docker-to-another-hostport-or-a-unix-socket

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.16 Ensure that Docker socket file permissions are set to 660 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the Docker socket file has permissions of `660` or more restrictive.

**Rationale:**

Only `root` and members of `docker` group should be allowed to read and write to default Docker Unix socket. Hence, the Docket socket file must have permissions of `660` or more restrictive.

**Audit:**

Execute the below command to verify that the Docker socket file has permissions of `660` or more restrictive:

```
stat -c %a /var/run/docker.sock
```

**Remediation:**

```
chmod 660 /var/run/docker.sock
```

This would set the file permissions of the Docker socket file to `660`.

**Impact:**

None.

**Default Value:**

By default, the permissions for Docker socket file is correctly set to `660`.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-socket-option
2. https://docs.docker.com/engine/reference/commandline/dockerd/#bind-docker-to-another-hostport-or-a-unix-socket

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## *3.17 Ensure that daemon.json file ownership is set to root:root (Scored)*

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `daemon.json` file ownership and group-ownership is correctly set to `root`.

**Rationale:**

`daemon.json` file contains sensitive parameters that may alter the behavior of docker daemon. Hence, it should be owned and group-owned by `root` to maintain the integrity of the file.

**Audit:**

Execute the below command to verify that the file is owned and group-owned by `root`:

```
stat -c %U:%G /etc/docker/daemon.json | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root /etc/docker/daemon.json
```

This would set the ownership and group-ownership for the file to `root`.

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-configuration-file

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 3.18 Ensure that daemon.json file permissions are set to 644 or more restrictive (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `daemon.json` file permissions are correctly set to `644` or more restrictive.

**Rationale:**

`daemon.json` file contains sensitive parameters that may alter the behavior of docker daemon. Hence, it should be writable only by `root` to maintain the integrity of the file.

**Audit:**

Execute the below command to verify that the file permissions are correctly set to `644` or more restrictive:

```
stat -c %a /etc/docker/daemon.json
```

**Remediation:**

```
chmod 644 /etc/docker/daemon.json
```

This would set the file permissions for this file to `644`.

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable.

**References:**

1. https://docs.docker.com/engine/reference/commandline/dockerd/#daemon-configuration-file

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## *3.19 Ensure that /etc/default/docker file ownership is set to root:root (Scored)*

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `/etc/default/docker` file ownership and group-ownership is correctly set to `root`.

**Rationale:**

`/etc/default/docker` file contains sensitive parameters that may alter the behavior of docker daemon. Hence, it should be owned and group-owned by `root` to maintain the integrity of the file.

**Audit:**

Execute the below command to verify that the file is owned and group-owned by `root`:

```
stat -c %U:%G /etc/default/docker | grep -v root:root
```

The above command should not return anything.

**Remediation:**

```
chown root:root /etc/default/docker
```

This would set the ownership and group-ownership for the file to `root`.

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable.

**References:**

1. https://docs.docker.com/engine/admin/configuring/

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## *3.20 Ensure that /etc/default/docker file permissions are set to 644 or more restrictive (Scored)*

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Verify that the `/etc/default/docker` file permissions are correctly set to `644` or more restrictive.

**Rationale:**

`/etc/default/docker` file contains sensitive parameters that may alter the behavior of docker daemon. Hence, it should be writable only by `root` to maintain the integrity of the file.

**Audit:**

Execute the below command to verify that the file permissions are correctly set to `644` or more restrictive:

```
stat -c %a /etc/default/docker
```

**Remediation:**

```
chmod 644 /etc/default/docker
```

This would set the file permissions for this file to `644`.

**Impact:**

None.

**Default Value:**

This file may not be present on the system. In that case, this recommendation is not applicable.

**References:**

1. https://docs.docker.com/engine/admin/configuring/

**CIS Controls:**

14.4 <u>Protect Information With Access Control Lists</u>
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## *4 Container Images and Build File*

Container base images and build files govern the fundamentals of how a container instance from a particular image would behave. Ensuring that you are using proper base images and appropriate build files can be very important for building your containerized infrastructure. Below are some of the recommendations that you should follow for container base images and build files to ensure that your containerized infrastructure is secure.

## *4.1 Ensure a user for the container has been created (Scored)*

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Create a non-root user for the container in the Dockerfile for the container image.

**Rationale:**

It is a good practice to run the container as a non-root user, if possible. Though user namespace mapping is now available, if a user is already defined in the container image, the container is run as that user by default and specific user namespace remapping is not required.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}: User={{ .Config.User }}'
```

The above command should return container username or user ID. If it is blank it means, the container is running as `root`.

**Remediation:**

Ensure that the Dockerfile for the container image contains below instruction:

```
USER <username or ID>
```

where username or ID refers to the user that could be found in the container base image. If there is no specific user created in the container base image, then add a `useradd` command to add the specific user before `USER` instruction.

For example, add the below lines in the Dockerfile to create a user in the container:

```
RUN useradd -d /home/username -m -s /bin/bash username
USER username
```

**Note:** If there are users in the image that the containers do not need, consider deleting them. After deleting those users, commit the image and then generate new instances of containers for use.

**Impact:**

None.

**Default Value:**

By default, the containers are run with `root` privileges and as user `root` inside the container.

**References:**

1. https://github.com/docker/docker/issues/2918
2. https://github.com/docker/docker/pull/4572
3. https://github.com/docker/docker/issues/7906

**CIS Controls:**

5.1 Minimize And Sparingly Use Administrative Privileges
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 4.2 Ensure that containers use trusted base images (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Ensure that the container image is written either from scratch or is based on another established and trusted base image downloaded over a secure channel.

**Rationale:**

Official repositories are Docker images curated and optimized by the Docker community or the vendor. There could be other potentially unsafe public repositories. Caution should be exercised when obtaining container images from Docker and third parties to how they will be used for your organization's data.

**Audit:**

**Step 1** - Inspect the Docker host for Docker images used by executing the below command:

```
docker images
```

This would list all the container images that are currently available for use on the Docker host. Interview the system administrator and obtain a proof of evidence that the list of images was obtained from trusted source over a secure channel or from a trusted, secure private Docker registry.

**Step 2** - For each Docker image found on the Docker host, inspect the image for how it was built to verify if from trusted sources and hardened configuration:

```
docker history <imageName>
```

**Remediation:**

- Configure and use Docker Content trust.
- Inspect Docker image history to evaluate their risk to operate on your network.
- Scan Docker images for vulnerabilities in their dependencies and configurations they will impose upon your network.

**Impact:**

None.

**Default Value:**

Not Applicable.

**References:**

1. https://titanous.com/posts/docker-insecurity
2. https://registry.hub.docker.com/
3. http://blog.docker.com/2014/10/docker-1-3-signed-images-process-injection-security-options-mac-shared-directories/
4. https://github.com/docker/docker/issues/8093
5. https://docs.docker.com/engine/reference/commandline/pull/
6. https://github.com/docker/docker/pull/11109
7. https://blog.docker.com/2015/11/docker-trusted-registry-1-4/

**CIS Controls:**

3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

## 4.3 Ensure unnecessary packages are not installed in the container (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Containers tend to be minimal and slim down versions of the Operating System. Do not install anything that does not justify the purpose of container.

**Rationale:**

Bloating containers with unnecessary software could possibly increase the attack surface of the container. This also voids the concept of minimal and slim down versions of container images. Hence, do not install anything else apart from what is truly needed for the purpose of the container.

**Audit:**

**Step 1**: List all the running instances of containers by executing below command:

```
docker ps --quiet
```

**Step 2**: For each container instance, execute the below or equivalent command:

```
docker exec $INSTANCE_ID rpm -qa
```

The above command would list the packages installed on the container. Review the list and ensure that it is legitimate.

**Remediation:**

At the outset, do not install anything on the container that does not justify the purpose. If the image had some packages that your container does not use, uninstall them.

Consider using a minimal base image rather than the standard Redhat/Centos/Debian images if you can. Some of the options include BusyBox and Alpine.

Not only does this trim your image size from >150Mb to ~20 Mb, there are also fewer tools and paths to escalate privileges. You can even remove the package installer as a final hardening measure for leaf/production containers.

**Impact:**

None.

**Default Value:**

Not Applicable.

**References:**

1. https://docs.docker.com/userguide/dockerimages/
2. http://www.livewyer.com/blog/2015/02/24/slimming-down-your-docker-containers-alpine-linux
3. https://github.com/progrium/busybox

**CIS Controls:**

18 Application Software Security
Application Software Security

## 4.4 Ensure images are scanned and rebuilt to include security patches (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Images should be scanned "frequently" for any vulnerabilities. Rebuild the images to include patches and then instantiate new containers from it.

**Rationale:**

Vulnerabilities are loopholes/bugs that can be exploited and security patches are updates to resolve these vulnerabilities. We can use image vulnerability scanning tools to find any kind of vulnerabilities within the images and then check for available patches to mitigate these vulnerabilities. Patches update the system to the most recent code base. Being on the current code base is important because that's where vendors focus on fixing problems. Evaluate the security patches before applying and follow the patching best practices.

Also, it would be better if, image vulnerability scanning tools could perform binary level analysis or hash based verification instead of just version string matching.

**Audit:**

**Step 1**: List all the running instances of containers by executing below command:

```
docker ps --quiet
```

**Step 2**: For each container instance, execute the below or equivalent command to find the list of packages installed within the container. Ensure that the security updates for various affected packages are installed.

```
docker exec $INSTANCE_ID rpm -qa
```

Alternatively, you could run image vulnerability scanning tools which can scan all the images in your ecosystem and then apply patches for the detected vulnerabilities based on your patch management procedures.

**Remediation:**

Follow the below steps to rebuild the images with security patches:

**Step 1**: Pull all the base images (i.e., given your set of Dockerfiles, extract all images declared in FROM instructions, and re-pull them to check for an updated/patched versions). Patch the packages within the images too.

```
docker pull
```

**Step 2**: Force a rebuild of each image:

```
docker build --no-cache
```

**Step 3**: Restart all containers with the updated images.

You could also use ONBUILD directive in the Dockerfile to trigger particular update instructions for images that you know are used as base images frequently.

**Impact:**

None

**Default Value:**

By default, containers and images are not updated of their own.

**References:**

1. https://docs.docker.com/userguide/dockerimages/
2. https://docs.docker.com/docker-cloud/builds/image-scan/
3. https://blog.docker.com/2016/05/docker-security-scanning/
4. https://docs.docker.com/engine/reference/builder/#/onbuild

**CIS Controls:**

18.1 Use Only Vendor-supported Software
For all acquired application software, check that the version you are using is still supported by the vendor. If not, update to the most current version and install all relevant patches and vendor security recommendations.

## 4.5 Ensure Content trust for Docker is Enabled (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Content trust is disabled by default. You should enable it.

**Rationale:**

Content trust provides the ability to use digital signatures for data sent to and received from remote Docker registries. These signatures allow client-side verification of the integrity and publisher of specific image tags. This ensures provenance of container images.

**Audit:**

```
echo $DOCKER_CONTENT_TRUST
```

This should return 1.

**Remediation:**

To enable content trust in a bash shell, enter the following command:

```
export DOCKER_CONTENT_TRUST=1
```

Alternatively, set this environment variable in your profile file so that content trust in enabled on every login.

**Impact:**

In an environment where `DOCKER_CONTENT_TRUST` is set, you are required to follow trust procedures while working with images - `build`, `create`, `pull`, `push` and `run`. You can use the `--disable-content-trust` flag to run individual operations on tagged images without content trust on an as-needed basis but that defeats the purpose of enabling content trust and hence, should be avoided wherever possible.

Note: Content trust is currently only available for users of the public Docker Hub. It is currently not available for the Docker Trusted Registry or for private registries.

**Default Value:**

By default, content trust is disabled.

**References:**

1. https://docs.docker.com/engine/security/trust/content_trust/
2. https://docs.docker.com/engine/reference/commandline/cli/#notary
3. https://docs.docker.com/engine/reference/commandline/cli/#environment-variables

**CIS Controls:**

18 Application Software Security
Application Software Security

## 4.6 Ensure HEALTHCHECK instructions have been added to the container image (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Add `HEALTHCHECK` instruction in your docker container images to perform the health check on running containers.

**Rationale:**

One of the important security triads is availability. Adding `HEALTHCHECK` instruction to your container image ensures that the docker engine periodically checks the running container instances against that instruction to ensure that the instances are still working.

Based on the reported health status, the docker engine could then exit non-working containers and instantiate new ones.

**Audit:**

Run the below command and ensure that the docker image has appropriate `HEALTHCHECK` instruction set up.

```
docker inspect --format='{{ .Config.Healthcheck }}' <IMAGE>
```

**Remediation:**

Follow Docker documentation and rebuild your container image with `HEALTHCHECK` instruction.

**Impact:**

None.

**Default Value:**

By default, `HEALTHCHECK` is not set.

**References:**

1. https://docs.docker.com/engine/reference/builder/#healthcheck

**CIS Controls:**

18 <u>Application Software Security</u>
Application Software Security

## 4.7 Ensure update instructions are not use alone in the Dockerfile (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Do not use update instructions such as `apt-get update` alone or in a single line in the Dockerfile.

**Rationale:**

Adding the update instructions in a single line on the Dockerfile will cache the update layer. Thus, when you build any image later using the same instruction, previously cached update layer will be used. This could potentially deny any fresh updates to go in the later builds.

**Audit:**

**Step 1:** Run the below command to get the list of images:

```
docker images
```

**Step 2:** Run the below command for each image in the list above, and look for any update instructions being in a single line:

```
docker history <Image_ID>
```

Alternatively, if you have access to Dockerfile for the image, verify that there are no update instructions as described above.

**Remediation:**

Use update instructions along with install instructions (or any other) and version pinning for packages while installing them. This would bust the cache and force to extract the required versions.

Alternatively, you could use `--no-cache` flag during `docker build` process to avoid using cached layers.

**Impact:**

None

**Default Value:**

By default, docker does not enforce any restrictions on using update instructions.

**References:**

1. https://docs.docker.com/engine/userguide/eng-image/dockerfile_best-practices/#run
2. https://github.com/docker/docker/issues/3313

**CIS Controls:**

18 Application Software Security
Application Software Security

## 4.8 Ensure setuid and setgid permissions are removed in the images (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Removing setuid and setgid permissions in the images would prevent privilege escalation attacks in the containers.

**Rationale:**

setuid and setgid permissions could be used for elevating privileges. While these permissions are at times legitimately needed, these could potentially be used in privilege escalation attacks. Thus, you should consider dropping these permissions for the packages which do not need them within the images.

**Audit:**

Run the below command on the image to list the executables having setuid and setgid permissions:

```
docker run <Image_ID> find / -perm +6000 -type f -exec ls -ld {} \; 2>
/dev/null
```

Carefully, review the list and ensure that it is legitimate.

**Remediation:**

Allow setuid and setgid permissions only on executables which need them. You could remove these permissions during build time by adding the following command in your Dockerfile, preferably towards the end of the Dockerfile:

```
RUN find / -perm +6000 -type f -exec chmod a-s {} \; || true
```

**Impact:**

Above command breaks all the executables that depend on setuid or setgid permissions including the legitimate ones. Hence, be careful to modify the command to suit your requirements so that it does not drop the permissions of legitimate programs. This requires a careful examination of each executable and fine tuning the permissions.

**Default Value:**

Not Applicable

**References:**

1. http://www.oreilly.com/webops-perf/free/files/docker-security.pdf
2. http://container-solutions.com/content/uploads/2015/06/15.06.15_DockerCheatSheet_A2.pdf
3. http://man7.org/linux/man-pages/man2/setuid.2.html
4. http://man7.org/linux/man-pages/man2/setgid.2.html

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 4.9 Ensure COPY is used instead of ADD in Dockerfile (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Use `COPY` instruction instead of `ADD` instruction in the Dockerfile.

**Rationale:**

`COPY` instruction just copies the files from the local host machine to the container file system. `ADD` instruction potentially could retrieve files from remote URLs and perform operations such as unpacking. Thus, `ADD` instruction introduces risks such as adding malicious files from URLs without scanning and unpacking procedure vulnerabilities.

**Audit:**

**Step 1:** Run the below command to get the list of images:

```
docker images
```

**Step 2:** Run the below command for each image in the list above and look for any `ADD` instructions:

```
docker history <Image_ID>
```

Alternatively, if you have access to Dockerfile for the image, verify that there are no `ADD` instructions.

**Remediation:**

Use `COPY` instructions in Dockerfiles.

**Impact:**

You would need to take care of the functionalities provided by `ADD` instructions such as fetching files from remote URLs.

**Default Value:**

Not Applicable

**References:**

1. https://docs.docker.com/engine/userguide/eng-image/dockerfile_best-practices/#add-or-copy

**CIS Controls:**

18 Application Software Security
Application Software Security

## 4.10 Ensure secrets are not stored in Dockerfiles (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Do not store any secrets in Dockerfiles.

**Rationale:**

Dockerfiles could be backtracked easily by using native Docker commands such as `docker history` and various tools and utilities. Also, as a general practice, image publishers provide Dockerfiles to build the credibility for their images. Hence, the secrets within these Dockerfiles could be easily exposed and potentially be exploited.

**Audit:**

**Step 1:** Run the below command to get the list of images:

```
docker images
```

**Step 2:** Run the below command for each image in the list above, and look for any secrets:

```
docker history <Image_ID>
```

Alternatively, if you have access to Dockerfile for the image, verify that there are no secrets as described above.

**Remediation:**

Do not store any kind of secrets within Dockerfiles.

**Impact:**

You would need to identify a way to handle secrets for your Docker images.

**Default Value:**

By default, there are no restrictions on storing config secrets in the Dockerfiles.

**References:**

1. https://github.com/docker/docker/issues/13490
2. http://12factor.net/config

3. https://avicoder.me/2016/07/22/Twitter-Vine-Source-code-dump/

**CIS Controls:**

14 Controlled Access Based on the Need to Know
Controlled Access Based on the Need to Know

## 4.11 Ensure verified packages are only Installed (Not Scored)

**Profile Applicability:**

• Level 2 - Docker

**Description:**

Verify authenticity of the packages before installing them in the image.

**Rationale:**

Verifying authenticity of the packages is essential for building a secure container image. Tampered packages could potentially be malicious or have some known vulnerabilities that could be exploited.

**Audit:**

**Step 1:** Run the below command to get the list of images:

```
docker images
```

**Step 2:** Run the below command for each image in the list above, and look for how the authenticity of the packages is determined. This could be via the use of GPG keys or other secure package distribution mechanisms

```
docker history <Image_ID>
```

Alternatively, if you have access to Dockerfile for the image, verify that the authenticity of the packages is checked.

**Remediation:**

Use GPG keys for downloading and verifying packages or any other secure package distribution mechanism of your choice.

**Impact:**

None

**Default Value:**

Not Applicable

**References:**

1. http://www.oreilly.com/webops-perf/free/files/docker-security.pdf
2. https://github.com/docker-library/httpd/blob/12bf8c8883340c98b3988a7bade8ef2d0d6dcf8a/2.4/Dockerfile
3. https://github.com/docker-library/php/blob/d8a4ccf4d620ec866d5b42335b699742df08c5f0/7.0/alpine/Dockerfile
4. https://access.redhat.com/security/team/key

**CIS Controls:**

18.1 <u>Use Only Vendor-supported Software</u>

For all acquired application software, check that the version you are using is still supported by the vendor. If not, update to the most current version and install all relevant patches and vendor security recommendations.

# 5 Container Runtime

The ways in which a container is started governs a lot of security implications. It is possible to provide potentially dangerous runtime parameters that might compromise the host and other containers on the host. Verifying container runtime is thus very important. Various recommendations to assess the container runtime are as below:

## 5.1 Ensure AppArmor Profile is Enabled (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

AppArmor is an effective and easy-to-use Linux application security system. It is available on quite a few Linux distributions by default such as Debian and Ubuntu.

**Rationale:**

AppArmor protects the Linux OS and applications from various threats by enforcing security policy which is also known as AppArmor profile. You can create your own AppArmor profile for containers or use the Docker's default AppArmor profile. This would enforce security policies on the containers as defined in the profile.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
AppArmorProfile={{ .AppArmorProfile }}'
```

The above command should return a valid AppArmor Profile for each container instance.

**Remediation:**

If AppArmor is applicable for your Linux OS, use it. You may have to follow below set of steps:

1. Verify if AppArmor is installed. If not, install it.
2. Create or import a AppArmor profile for Docker containers.
3. Put this profile in enforcing mode.
4. Start your Docker container using the customized AppArmor profile. For example,

```
docker run --interactive --tty --security-opt="apparmor:PROFILENAME" centos
/bin/bash
```

Alternatively, you can keep the docker's default apparmor profile.

**Impact:**

The container (process) would have set of restrictions as defined in AppArmor profile. If your AppArmor profile is mis-configured, then the container may not entirely work as expected.

**Default Value:**

By default, `docker-default` AppArmor profile is applied for running containers and this profile can be found at `/etc/apparmor.d/docker`.

**References:**

1. https://docs.docker.com/engine/security/apparmor/
2. https://docs.docker.com/engine/reference/run/#security-configuration
3. https://docs.docker.com/engine/security/security/#other-kernel-security-features

**CIS Controls:**

14.4 Protect Information With Access Control Lists
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 5.2 Ensure SELinux security options are set, if applicable (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

SELinux is an effective and easy-to-use Linux application security system. It is available on quite a few Linux distributions by default such as Red Hat and Fedora.

**Rationale:**

SELinux provides a Mandatory Access Control (MAC) system that greatly augments the default Discretionary Access Control (DAC) model. You can thus add an extra layer of safety by enabling SELinux on your Linux host, if applicable.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
SecurityOpt={{ .HostConfig.SecurityOpt }}'
```

The above command should return all the security options currently configured for the containers.

**Remediation:**

If SELinux is applicable for your Linux OS, use it. You may have to follow below set of steps:

1. Set the SELinux State.
2. Set the SELinux Policy.
3. Create or import a SELinux policy template for Docker containers.
4. Start Docker in daemon mode with SELinux enabled. For example,

```
docker daemon --selinux-enabled
```

5. Start your Docker container using the security options. For example,

```
docker run --interactive --tty --security-opt label=level:TopSecret centos
/bin/bash
```

**Impact:**

The container (process) would have set of restrictions as defined in SELinux policy. If your SELinux policy is mis-configured, then the container may not entirely work as expected.

**Default Value:**

By default, no SELinux security options are applied on containers.

**References:**

1. https://docs.docker.com/engine/security/security/#other-kernel-security-features
2. https://docs.docker.com/engine/reference/run/#security-configuration
3. http://docs.fedoraproject.org/en-US/Fedora/13/html/Security-Enhanced_Linux/
4. https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux_atomic_host/7/html/container_security_guide/docker_selinux_security_policy

**CIS Controls:**

14.4 Protect Information With Access Control Lists
All information stored on systems shall be protected with file system, network share, claims, application, or database specific access control lists. These controls will enforce the principle that only authorized individuals should have access to the information based on their need to access the information as a part of their responsibilities.

## 5.3 Ensure Linux Kernel Capabilities are restricted within containers (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, Docker starts containers with a restricted set of Linux Kernel Capabilities. It means that any process may be granted the required capabilities instead of root access. Using Linux Kernel Capabilities, the processes do not have to run as root for almost all the specific areas where root privileges are usually needed.

**Rationale:**

Docker supports the addition and removal of capabilities, allowing the use of a non-default profile. This may make Docker more secure through capability removal, or less secure through the addition of capabilities. It is thus recommended to remove all capabilities except those explicitly required for your container process.

For example, capabilities such as below are usually not needed for container process:

```
NET_ADMIN
SYS_ADMIN
SYS_MODULE
```

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}: CapAdd={{
.HostConfig.CapAdd }} CapDrop={{ .HostConfig.CapDrop }}'
```

Verify that the added and dropped Linux Kernel Capabilities are in line with the ones needed for container process for each container instance.

**Remediation:**

Execute the below command to add needed capabilities:

$> docker run --cap-add={"Capability 1","Capability 2"}

For example,

```
docker run --interactive --tty --cap-add={"NET_ADMIN","SYS_ADMIN"}
centos:latest /bin/bash
```

Execute the below command to drop unneeded capabilities:

$> docker run --cap-drop={"Capability 1","Capability 2"}

For example,

```
docker run --interactive --tty --cap-drop={"SETUID","SETGID"} centos:latest
/bin/bash
```

Alternatively,

You may choose to drop all capabilities and add only add the needed ones:

$> docker run --cap-drop=all --cap-add={"Capability 1","Capability 2"}

For example,

```
docker run --interactive --tty --cap-drop=all --cap-
add={"NET_ADMIN","SYS_ADMIN"} centos:latest /bin/bash
```

**Impact:**

Based on what Linux Kernel Capabilities were added or dropped, restrictions within the container would apply.

**Default Value:**

By default, below capabilities are available for containers:

```
AUDIT_WRITE
CHOWN
DAC_OVERRIDE
FOWNER
FSETID
KILL
MKNOD
NET_BIND_SERVICE
NET_RAW
SETFCAP
SETGID
SETPCAP
SETUID
SYS_CHROOT
```

**References:**

1. https://docs.docker.com/engine/security/security/#linux-kernel-capabilities
2. http://man7.org/linux/man-pages/man7/capabilities.7.html
3. http://www.oreilly.com/webops-perf/free/files/docker-security.pdf

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 5.4 Ensure privileged containers are not used (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Using the `--privileged` flag gives all Linux Kernel Capabilities to the container thus overwriting the `--cap-add` and `--cap-drop` flags. Ensure that it is not used.

**Rationale:**

The `--privileged` flag gives all capabilities to the container, and it also lifts all the limitations enforced by the device cgroup controller. In other words, the container can then do almost everything that the host can do. This flag exists to allow special use-cases, like running Docker within Docker.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Privileged={{ .HostConfig.Privileged }}'
```

The above command should return `Privileged=false` for each container instance.

**Remediation:**

Do not run container with the `--privileged` flag.

For example, do not start a container as below:

```
docker run --interactive --tty --privileged centos /bin/bash
```

**Impact:**

Linux Kernel Capabilities other than defaults would not be available for use within container.

**Default Value:**

False.

**References:**

1. https://docs.docker.com/engine/reference/run/#runtime-privilege-and-linux-capabilities

**CIS Controls:**

5.1 <u>Minimize And Sparingly Use Administrative Privileges</u>
Minimize administrative privileges and only use administrative accounts when they are required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 5.5 Ensure sensitive host system directories are not mounted on containers (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Sensitive host system directories such as below should not be allowed to be mounted as container volumes especially in read-write mode.

```
/
/boot
/dev
/etc
/lib
/proc
/sys
/usr
```

**Rationale:**

If sensitive directories are mounted in read-write mode, it would be possible to make changes to files within those sensitive directories. The changes might bring down security implications or unwarranted changes that could put the Docker host in compromised state.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Volumes={{ .Mounts }}'
```

The above commands would return the list of current mapped directories and whether they are mounted in read-write mode for each container instance.

**Remediation:**

Do not mount host sensitive directories on containers especially in read-write mode.

**Impact:**

None.

**Default Value:**

Docker defaults to a read-write volume but you can also mount a directory read-only. By default, no sensitive host directories are mounted on containers.

**References:**

1. https://docs.docker.com/engine/tutorials/dockervolumes/

**CIS Controls:**

14 <u>Controlled Access Based on the Need to Know</u>
Controlled Access Based on the Need to Know

## 5.6 Ensure ssh is not run within containers (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

SSH server should not be running within the container. You should SSH into the Docker host, and use nsenter tool to enter a container from a remote host.

**Rationale:**

Running SSH within the container increases the complexity of security management by making it

- Difficult to manage access policies and security compliance for SSH server
- Difficult to manage keys and passwords across various containers
- Difficult to manage security upgrades for SSH server

It is possible to have shell access to a container without using SSH, the needlessly increasing the complexity of security management should be avoided.

**Audit:**

**Step 1**: List all the running instances of containers by executing below command:

```
docker ps --quiet
```

**Step 2**: For each container instance, execute the below command:

```
docker exec $INSTANCE_ID ps -el
```

Ensure that there is no process for SSH server.

**Remediation:**

Uninstall SSH server from the container and use `nsenter` or any other commands such as `docker exec` or `docker attach` to interact with the container instance.

```
docker exec --interactive --tty $INSTANCE_ID sh
```

OR

```
docker attach $INSTANCE_ID
```

**Impact:**

None.

**Default Value:**

By default, SSH server is not running inside the container. Only one process per container is allowed.

**References:**

1. http://blog.docker.com/2014/06/why-you-dont-need-to-run-sshd-in-docker/

**CIS Controls:**

9.1 Limit Open Ports, Protocols, and Services
Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 5.7 Ensure privileged ports are not mapped within containers (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The TCP/IP port numbers below `1024` are considered privileged ports. Normal users and processes are not allowed to use them for various security reasons. Docker allows a container port to be mapped to a privileged port.

**Rationale:**

By default, if the user does not specifically declare the container port to host port mapping, Docker automatically and correctly maps the container port to one available in `49153-65535` block on the host. But, Docker allows a container port to be mapped to a privileged port on the host if the user explicitly declared it. This is so because containers are executed with `NET_BIND_SERVICE` Linux kernel capability that does not restrict the privileged port mapping. The privileged ports receive and transmit various sensitive and privileged data. Allowing containers to use them can bring serious implications.

**Audit:**

List all running containers instances and their port mapping by executing the below command:

```
docker ps --quiet | xargs docker inspect --format '{{ .Id }}: Ports={{
.NetworkSettings.Ports }}'
```

Review the list and ensure that container ports are not mapped to host port numbers below `1024`.

**Remediation:**

Do not map the container ports to privileged host ports when starting a container. Also, ensure that there is no such container to host privileged port mapping declarations in the Dockerfile.

**Impact:**

None.

**Default Value:**

By default, mapping a container port to a privileged port on the host is allowed.

**Note:** There might be certain cases where you want to map privileged ports, because if you forbid it, then the corresponding application has to run outside of a container.

For example: HTTP and HTTPS load balancers have to bind 80/tcp and 443/tcp respectively. Forbidding to map privileged ports effectively forbids from running those in a container, and mandates using an external load balancer. In such cases, those containers instances should be marked as exceptions for this recommendation.

**References:**

1. https://docs.docker.com/engine/userguide/networking/

**CIS Controls:**

9.1 <u>Limit Open Ports, Protocols, and Services</u>
Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 5.8 Ensure only needed ports are open on the container (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Dockerfile for a container image defines the ports to be opened by default on a container instance. The list of ports may or may not be relevant to the application you are running within the container.

**Rationale:**

A container can be run just with the ports defined in the Dockerfile for its image or can be arbitrarily passed run time parameters to open a list of ports. Additionally, Overtime, Dockerfile may undergo various changes and the list of exposed ports may or may not be relevant to the application you are running within the container. Opening unneeded ports increase the attack surface of the container and the containerized application. As a recommended practice, do not open unneeded ports.

**Audit:**

List all the running instances of containers and their port mapping by executing the below command:

```
docker ps --quiet | xargs docker inspect --format '{{ .Id }}: Ports={{
.NetworkSettings.Ports }}'
```

Review the list and ensure that the ports mapped are the ones that are really needed for the container.

**Remediation:**

Fix the Dockerfile of the container image to expose only needed ports by your containerized application. You can also completely ignore the list of ports defined in the Dockerfile by **NOT** using `-P` (UPPERCASE) or `--publish-all` flag when starting the container. Use the `-p` (lowercase) or `--publish` flag to explicitly define the ports that you need for a particular container instance.

For example,

```
docker run --interactive --tty --publish 5000 --publish 5001 --publish 5002
centos /bin/bash
```

**Impact:**

None.

**Default Value:**

By default, all the ports that are listed in the Dockerfile under `EXPOSE` instruction for an image are opened when a container is run with `-P` or `--publish-all` flag.

**References:**

1. https://docs.docker.com/engine/userguide/networking/

**CIS Controls:**

9.1 <u>Limit Open Ports, Protocols, and Services</u>
Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 5.9 Ensure the host's network namespace is not shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The networking mode on a container when set to `--net=host`, skips placing the container inside separate network stack. In essence, this choice tells Docker to not containerize the container's networking. This would network-wise mean that the container lives "outside" in the main Docker host and has full access to its network interfaces.

**Rationale:**

This is potentially dangerous. It allows the container process to open low-numbered ports like any other `root` process. It also allows the container to access network services like D-bus on the Docker host. Thus, a container process can potentially do unexpected things such as shutting down the Docker host. You should not use this option.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
NetworkMode={{ .HostConfig.NetworkMode }}'
```

If the above command returns `NetworkMode=host`, it means that `--net=host` option was passed when container was started. This would be non-compliant.

**Remediation:**

Do not pass `--net=host` option when starting the container.

**Impact:**

None.

**Default Value:**

By default, container connects to Docker `bridge`.

**References:**

1. https://docs.docker.com/engine/userguide/networking/
2. https://docs.docker.com/engine/reference/run/#network-settings

**CIS Controls:**

12 Boundary Defense
Boundary Defense

## 5.10 Ensure memory usage for container is limited (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, all containers on a Docker host share the resources equally. By using the resource management capabilities of Docker host, such as memory limit, you can control the amount of memory that a container may consume.

**Rationale:**

By default, container can use all of the memory on the host. You can use memory limit mechanism to prevent a denial of service arising from one container consuming all of the host's resources such that other containers on the same host cannot perform their intended functions. Having no limit on memory can lead to issues where one container can easily make the whole system unstable and as a result unusable.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}: Memory={{ .HostConfig.Memory }}'
```

If the above command returns `0`, it means the memory limits are not in place. If the above command returns a non-zero value, it means memory limits are in place.

**Remediation:**

Run the container with only as much memory as required. Always run the container using the `--memory` argument.

For example, you could run a container as below:

```
docker run --interactive --tty --memory 256m centos /bin/bash
```

In the above example, the container is started with a memory limit of 256 MB.

Note: Please note that the output of the below command would return values in scientific notation if memory limits are in place.

```
docker inspect --format='{{.Config.Memory}}' 7c5a2d4c7fe0
```

For example, if the memory limit is set to `256 MB` for the above container instance, the output of the above command would be `2.68435456e+08` and NOT 256m. You should convert this value using a scientific calculator or programmatic methods.

**Impact:**

If you do not set proper limits, the container process may have to starve.

**Default Value:**

By default, all containers on a Docker host share the resources equally. No memory limits are enforced.

**References:**

1. https://goldmann.pl/blog/2014/09/11/resource-management-in-docker/
2. https://docs.docker.com/engine/reference/commandline/run/#options
3. https://docs.docker.com/engine/admin/runmetrics/

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.11 Ensure CPU priority is set appropriately on the container (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, all containers on a Docker host share the resources equally. By using the resource management capabilities of Docker host, such as CPU shares, you can control the host CPU resources that a container may consume.

**Rationale:**

By default, CPU time is divided between containers equally. If it is desired, to control the CPU time amongst the container instances, you can use CPU sharing feature. CPU sharing allows to prioritize one container over the other and forbids the lower priority container to claim CPU resources more often. This ensures that the high priority containers are served better.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
CpuShares={{ .HostConfig.CpuShares }}'
```

If the above command returns `0` or `1024`, it means the CPU shares are not in place. If the above command returns a non-zero value other than `1024`, it means CPU shares are in place.

**Remediation:**

Manage the CPU shares between your containers. To do so start the container using the `--cpu-shares` argument.

For example, you could run a container as below:

```
docker run --interactive --tty --cpu-shares 512 centos /bin/bash
```

In the above example, the container is started with CPU shares of 50% of what the other containers use. So, if the other container has CPU shares of 80%, this container will have CPU shares of 40%.

Note: Every new container will have `1024` shares of CPU by default. However, this value is shown as `0` if you run the command mentioned in the audit section.

Alternatively,

1. Navigate to `/sys/fs/cgroup/cpu/system.slice/` directory.
2. Check your container instance ID using `docker ps`.
3. Now, inside the above directory (in step 1), you would have a directory by name `docker-<Instance ID>.scope`. For example, `docker-4acae729e8659c6be696ee35b2237cc1fe4edd2672e9186434c5116e1a6fbed6.scope`. Navigate to this directory.
4. You will find a file named `cpu.shares`. Execute `cat cpu.shares`. This will always give you the CPU share value based on the system. So, even if there is no CPU shares configured using `-c` or `--cpu-shares` argument in the `docker run` command, this file will have a value of `1024`.

If we set one container's CPU shares to `512` it will receive half of the CPU time compared to the other container. So, take `1024` as 100% and then do quick math to derive the number that you should set for respective CPU shares. For example, use `512` if you want to set 50% and `256` if you want to set 25%.

**Impact:**

If you do not set proper CPU shares, the container process may have to starve if the resources on the host are not available. If the CPU resources on the host are free, CPU shares do not place any restrictions on the CPU that the container may use.

**Default Value:**

By default, all containers on a Docker host share the resources equally. No CPU shares are enforced.

**References:**

1. https://goldmann.pl/blog/2014/09/11/resource-management-in-docker/
2. https://docs.docker.com/engine/reference/commandline/run/#options
3. https://docs.docker.com/engine/admin/runmetrics/

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.12 Ensure the container's root filesystem is mounted as read only (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The container's root filesystem should be treated as a 'golden image' by using Docker run's `--read-only` option. This prevents any writes to the container's root filesystem at container runtime and enforces the principle of immutable infrastructure.

**Rationale:**

Enabling this option forces containers at runtime to explicitly define their data writing strategy to persist or not persist their data.

This also reduces security attack vectors since the container instance's filesystem cannot be tampered with or written to unless it has explicit read-write permissions on its filesystem folder and directories.

**Audit:**

Run the following command on the docker host:

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
ReadonlyRootfs={{ .HostConfig.ReadonlyRootfs }}'
```

If the above command returns `true`, it means the container's root filesystem is mounted read-only.

If the above command returns `false`, it means the container's root filesystem is writable.

**Remediation:**

Add a `--read-only` flag at a container's runtime to enforce the container's root filesystem to be mounted as read only.

```
docker run <Run arguments> --read-only <Container Image Name or ID> <Command>
```

Enabling the `--read-only` option at a container's runtime should be used by administrators to force a container's executable processes to only write container data to explicit storage locations during the container's runtime.

Examples of explicit storage locations during a container's runtime include, but not limited to:

1. Use the `--tmpfs` option to mount a temporary file system for non-persistent data writes.

```
docker run --interactive --tty --read-only --tmpfs "/run" --tmpfs "/tmp"
centos /bin/bash
```

2. Enabling Docker `rw` mounts at a container's runtime to persist container data directly on the Docker host filesystem.

```
docker run --interactive --tty --read-only -v /opt/app/data:/run/app/data:rw
centos /bin/bash
```

3. Utilizing Docker shared-storage volume plugins for Docker data volume to persist container data.

```
docker volume create -d convoy --opt o=size=20GB my-named-volume
docker run --interactive --tty --read-only -v my-named-volume:/run/app/data
centos /bin/bash
```

3. Transmitting container data outside of the docker during the container's runtime for container data to persist container data. Examples include hosted databases, network file shares, and APIs.

**Impact:**

Enabling `--read-only` at container runtime may break some container OS packages if a data writing strategy is not defined.

Define what the container's data should and should not persist at runtime to determine which recommendation procedure to utilize.

Example:

- Enable use `--tmpfs` for temporary file writes to `/tmp`
- Use Docker shared data volumes for persistent data writes

**Default Value:**

By default, a container will have its root filesystem writable allowing all container processes to write files owned by the container's runtime user.

**References:**

1. http://docs.docker.com/reference/commandline/cli/#run
2. https://docs.docker.com/engine/tutorials/dockervolumes/
3. http://www.projectatomic.io/blog/2015/12/making-docker-images-write-only-in-production/
4. https://docs.docker.com/engine/reference/commandline/run/#mount-tmpfs-tmpfs
5. https://docs.docker.com/engine/tutorials/dockervolumes/#creating-and-mounting-a-data-volume-container

**CIS Controls:**

14 Controlled Access Based on the Need to Know
Controlled Access Based on the Need to Know

## 5.13 Ensure incoming container traffic is binded to a specific host interface (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, Docker containers can make connections to the outside world, but the outside world cannot connect to containers. Each outgoing connection will appear to originate from one of the host machine's own IP addresses. Only allow container services to be contacted through a specific external interface on the host machine.

**Rationale:**

If you have multiple network interfaces on your host machine, the container can accept connections on the exposed ports on any network interface. This might not be desired and may not be secured. Many a times a particular interface is exposed externally and services such as intrusion detection, intrusion prevention, firewall, load balancing, etc. are run on those interfaces to screen incoming public traffic. Hence, you should not accept incoming connections on any interface. You should only allow incoming connections from a particular external interface.

**Audit:**

List all the running instances of containers and their port mapping by executing the below command:

```
docker ps --quiet | xargs docker inspect --format '{{ .Id }}: Ports={{
.NetworkSettings.Ports }}'
```

Review the list and ensure that the exposed container ports are tied to a particular interface and not to the wildcard IP address - `0.0.0.0`.

For example, if the above command returns as below, then this is non-compliant and the container can accept connections on any host interface on the specified port `49153`.

```
Ports=map[443/tcp:<nil> 80/tcp:[map[HostPort:49153 HostIp:0.0.0.0]]]
```

However, if the exposed port is tied to a particular interface on the host as below, then this recommendation is configured as desired and is compliant.

```
Ports=map[443/tcp:<nil> 80/tcp:[map[HostIp:10.2.3.4 HostPort:49153]]]
```

**Remediation:**

Bind the container port to a specific host interface on the desired host port.

For example,

```
docker run --detach --publish 10.2.3.4:49153:80 nginx
```

In the example above, the container port `80` is bound to the host port on `49153` and would accept incoming connection only from `10.2.3.4` external interface.

**Impact:**

None.

**Default Value:**

By default, Docker exposes the container ports on `0.0.0.0`, the wildcard IP address that will match any possible incoming network interface on the host machine.

**References:**

1. https://docs.docker.com/engine/userguide/networking/

**CIS Controls:**

9 Limitation and Control of Network Ports, Protocols, and Services
Limitation and Control of Network Ports, Protocols, and Services

## 5.14 Ensure 'on-failure' container restart policy is set to '5' (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Using the `--restart` flag in `docker run` command you can specify a restart policy for how a container should or should not be restarted on exit. You should choose the `on-failure` restart policy and limit the restart attempts to `5`.

**Rationale:**

If you indefinitely keep trying to start the container, it could possibly lead to a denial of service on the host. It could be an easy way to do a distributed denial of service attack especially if you have many containers on the same host. Additionally, ignoring the exit status of the container and `always` attempting to restart the container leads to non-investigation of the root cause behind containers getting terminated. If a container gets terminated, you should investigate on the reason behind it instead of just attempting to restart it indefinitely. Thus, it is recommended to use `on-failure` restart policy and limit it to maximum of `5` restart attempts.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
RestartPolicyName={{ .HostConfig.RestartPolicy.Name }} MaximumRetryCount={{
.HostConfig.RestartPolicy.MaximumRetryCount }}'
```

- If the above command returns `RestartPolicyName=always`, then the system is not configured as desired and hence this recommendation is non-compliant.
- If the above command returns `RestartPolicyName=no` or just `RestartPolicyName=`, then the restart policies are not being used and the container would never be restarted of its own. This recommendation is then Not Applicable and can be assumed to be compliant.
- If the above command returns `RestartPolicyName=on-failure`, then verify that the number of restart attempts is set to `5` or less by looking at `MaximumRetryCount`.

**Remediation:**

If a container is desired to be restarted of its own, then, for example, you could start the container as below:

```
docker run --detach --restart=on-failure:5 nginx
```

**Impact:**

The container would attempt to restart only for 5 times.

**Default Value:**

By default, containers are not configured with restart policies. Hence, containers do not attempt to restart of their own.

**References:**

1. https://docs.docker.com/engine/reference/commandline/run/#restart-policies-restart

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.15 Ensure the host's process namespace is not shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Process ID (PID) namespaces isolate the process ID number space, meaning that processes in different PID namespaces can have the same PID. This is process level isolation between containers and the host.

**Rationale:**

PID namespace provides separation of processes. The PID Namespace removes the view of the system processes, and allows process ids to be reused including PID 1. If the host's PID namespace is shared with the container, it would basically allow processes within the container to see all of the processes on the host system. This breaks the benefit of process level isolation between the host and the containers. Someone having access to the container can eventually know all the processes running on the host system and can even kill the host system processes from within the container. This can be catastrophic. Hence, do not share the host's process namespace with the containers.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
PidMode={{ .HostConfig.PidMode }}'
```

If the above command returns `host`, it means the host PID namespace is shared with the container else this recommendation is compliant.

**Remediation:**

Do not start a container with `--pid=host` argument.

For example, do not start a container as below:

```
docker run --interactive --tty --pid=host centos /bin/bash
```

**Impact:**

Container processes cannot see the processes on the host system. In certain cases, you want your container to share the host's process namespace. For example, you could build a container with debugging tools like `strace` or `gdb`, but want to use these tools when

debugging processes within the container. If this is desired, then share only one (or needed) host process by using the `-p` switch.

For example,

```
docker run --pid=host rhel7 strace -p 1234
```

**Default Value:**

By default, all containers have the PID namespace enabled and the host's process namespace is not shared with the containers.

**References:**

1. https://docs.docker.com/engine/reference/run/#pid-settings-pid
2. http://man7.org/linux/man-pages/man7/pid_namespaces.7.html

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.16 Ensure the host's IPC namespace is not shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

IPC (POSIX/SysV IPC) namespace provides separation of named shared memory segments, semaphores and message queues. IPC namespace on the host thus should not be shared with the containers and should remain isolated.

**Rationale:**

IPC namespace provides separation of IPC between the host and containers. If the host's IPC namespace is shared with the container, it would basically allow processes within the container to see all of the IPC on the host system. This breaks the benefit of IPC level isolation between the host and the containers. Someone having access to the container can eventually manipulate the host IPC. This can be catastrophic. Hence, do not share the host's IPC namespace with the containers.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
IpcMode={{ .HostConfig.IpcMode }}'
```

If the above command returns `host`, it means the host IPC namespace is shared with the container. If the above command returns nothing, then the host's IPC namespace is not shared. This recommendation is then compliant.

**Remediation:**

Do not start a container with `--ipc=host` argument. For example, do not start a container as below:

```
docker run --interactive --tty --ipc=host centos /bin/bash
```

**Impact:**

Shared memory segments are used to accelerate inter-process communication. It is commonly used by high-performance applications. If such applications are containerized into multiple containers, you might need to share the IPC namespace of the containers to achieve high performance. In such cases, you should still be sharing container specific IPC

namespaces only and not the host IPC namespace. You may share the container's IPC namespace with another container as below:

For example,

```
docker run --interactive --tty --ipc=container:e3a7a1a97c58 centos /bin/bash
```

**Default Value:**

By default, all containers have the IPC namespace enabled and host IPC namespace is not shared with any container.

**References:**

1. https://docs.docker.com/engine/reference/run/#ipc-settings-ipc
2. http://man7.org/linux/man-pages/man7/namespaces.7.html

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.17 Ensure host devices are not directly exposed to containers (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Host devices can be directly exposed to containers at runtime. Do not directly expose host devices to containers especially for containers that are not trusted.

**Rationale:**

The `--device` option exposes the host devices to the containers and consequently, the containers can directly access such host devices. You would not require the container to run in `privileged` mode to access and manipulate the host devices. By default, the container will be able to read, write and mknod these devices. Additionally, it is possible for containers to remove block devices from the host. Hence, do not expose host devices to containers directly.

If at all, you would want to expose the host device to a container, use the sharing permissions appropriately:

- r - read only
- w - writable
- m - mknod allowed

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Devices={{ .HostConfig.Devices }}'
```

The above command would list out each device with below information:

- `CgroupPermissions` - For example, `rwm`
- `PathInContainer` - Device path within the container
- `PathOnHost` - Device path on the host

Verify that the host device is needed to be accessed from within the container and the permissions required are correctly set. If the above command returns [], then the container does not have access to host devices. This recommendation can be assumed to be compliant.

**Remediation:**

Do not directly expose the host devices to containers. If at all, you need to expose the host devices to containers, use the correct set of permissions:

For example, do not start a container as below:

```
docker run --interactive --tty --device=/dev/tty0:/dev/tty0:rwm --
device=/dev/temp_sda:/dev/temp_sda:rwm centos bash
```

For example, share the host device with correct permissions:

```
docker run --interactive --tty --device=/dev/tty0:/dev/tty0:rw --
device=/dev/temp_sda:/dev/temp_sda:r centos bash
```

**Impact:**

You would not be able to use the host devices directly within the containers.

**Default Value:**

By default, no host devices are exposed to containers. If you do not provide sharing permissions and choose to expose a host device to a container, the host device would be exposed with `read`, `write` and `mknod` permissions.

**References:**

1. https://docs.docker.com/engine/reference/commandline/run/#options

**CIS Controls:**

14 Controlled Access Based on the Need to Know
Controlled Access Based on the Need to Know

## 5.18 Ensure the default ulimit is overwritten at runtime, only if needed (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The default ulimit is set at the Docker daemon level. However, you may override the default ulimit setting, if needed, during container runtime.

**Rationale:**

`ulimit` provides control over the resources available to the shell and to processes started by it. Setting system resource limits judiciously saves you from many disasters such as a fork bomb. Sometimes, even friendly users and legitimate processes can overuse system resources and in-turn can make the system unusable.

The default ulimit set at the Docker daemon level should be honored. If the default ulimit settings are not appropriate for a particular container instance, you may override them as an exception. But, do not make this a practice. If most of the container instances are overriding default ulimit settings, consider changing the default ulimit settings to something that is appropriate for your needs.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Ulimits={{ .HostConfig.Ulimits }}'
```

The above command should return `Ulimits=<no value>` for each container instance until and unless there is an exception and a need to override the default ulimit settings.

**Remediation:**

Only override the default ulimit settings if needed.

For example, to override default ulimit settings start a container as below:

```
docker run --ulimit nofile=1024:1024 --interactive --tty centos /bin/bash
```

**Impact:**

If the ulimits are not set properly, the desired resource control might not be achieved and might even make the system unusable.

**Default Value:**

Container instances inherit the default ulimit settings set at the Docker daemon level.

**References:**

1. https://docs.docker.com/engine/reference/commandline/run/#set-ulimits-in-container-ulimit
2. http://www.oreilly.com/webops-perf/free/files/docker-security.pdf

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.19 Ensure mount propagation mode is not set to shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Mount propagation mode allows mounting volumes in shared, slave or private mode on a container. Do not use shared mount propagation mode until needed.

**Rationale:**

A shared mount is replicated at all mounts and the changes made at any mount point are propagated to all mounts. Mounting a volume in shared mode does not restrict any other container to mount and make changes to that volume. This might be catastrophic if the mounted volume is sensitive to changes. Do not set mount propagation mode to shared until needed.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Propagation={{range $mnt := .Mounts}} {{json $mnt.Propagation}} {{end}}'
```

The above command would return the propagation mode for mounted volumes. Propagation mode should not be set to `shared` unless needed. The above command might throw errors if there are no mounts. In that case, this recommendation is not applicable.

**Remediation:**

Do not mount volumes in shared mode propagation.

For example, do not start container as below:

```
docker run <Run arguments> --volume=/hostPath:/containerPath:shared
<Container Image Name or ID> <Command>
```

**Impact:**

None.

**Default Value:**

By default, the container mounts are private.

**References:**

1. https://github.com/docker/docker/pull/17034
2. https://docs.docker.com/engine/reference/run/#volume-shared-filesystems
3. https://www.kernel.org/doc/Documentation/filesystems/sharedsubtree.txt

**CIS Controls:**

14 Controlled Access Based on the Need to Know
Controlled Access Based on the Need to Know

## 5.20 Ensure the host's UTS namespace is not shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

UTS namespaces provide isolation of two system identifiers: the hostname and the NIS domain name. It is used for setting the hostname and the domain that is visible to running processes in that namespace. Processes running within containers do not typically require to know hostname and domain name. Hence, the namespace should not be shared with the host.

**Rationale:**

Sharing the UTS namespace with the host provides full permission to the container to change the hostname of the host. This is insecure and should not be allowed.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
UTSMode={{ .HostConfig.UTSMode }}'
```

If the above command returns `host`, it means the host UTS namespace is shared with the container and this recommendation is non-compliant. If the above command returns nothing, then the host's UTS namespace is not shared. This recommendation is then compliant.

**Remediation:**

Do not start a container with `--uts=host` argument.

For example, do not start a container as below:

```
docker run --rm --interactive --tty --uts=host rhel7.2
```

**Impact:**

None.

**Default Value:**

By default, all containers have the UTS namespace enabled and host UTS namespace is not shared with any container.

**References:**

1. https://docs.docker.com/engine/reference/run/#uts-settings-uts
2. http://man7.org/linux/man-pages/man7/namespaces.7.html

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.21 Ensure the default seccomp profile is not Disabled (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Seccomp filtering provides a means for a process to specify a filter for incoming system calls. The default Docker seccomp profile works on whitelist basis and allows 311 system calls blocking all others. It should not be disabled unless it hinders your container application usage.

**Rationale:**

A large number of system calls are exposed to every userland process with many of them going unused for the entire lifetime of the process. Most of the applications do not need all the system calls and thus benefit by having a reduced set of available system calls. The reduced set of system calls reduces the total kernel surface exposed to the application and thus improvises application security.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
SecurityOpt={{ .HostConfig.SecurityOpt }}'
```

The above command should return `<no value>` or your modified seccomp profile. If it returns `[seccomp:unconfined]`, that means this recommendation is non-compliant and the container is running without any seccomp profiles.

**Remediation:**

By default, seccomp profiles are enabled. You do not need to do anything unless you want to modify and use the modified seccomp profile.

**Impact:**

With Docker 1.10 and greater, the default seccomp profile blocks syscalls, regardless of `--cap-add` passed to the container. You should create your own custom seccomp profile in such cases. You may also disable the default seccomp profile by passing `--security-opt=seccomp:unconfined` on `docker run`.

**Default Value:**

When you run a container, it uses the default profile unless you override it with the `--security-opt` option.

**References:**

1. http://blog.scalock.com/new-docker-security-features-and-what-they-mean-seccomp-profiles
2. https://docs.docker.com/engine/reference/run/#security-configuration
3. https://github.com/docker/docker/blob/master/profiles/seccomp/default.json
4. https://docs.docker.com/engine/security/seccomp/
5. https://www.kernel.org/doc/Documentation/prctl/seccomp_filter.txt
6. https://github.com/docker/docker/issues/22870

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.22 Ensure docker exec commands are not used with privileged option (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Do not `docker exec` with `--privileged` option.

**Rationale:**

Using `--privileged` option in `docker exec` gives extended Linux capabilities to the command. This could potentially be insecure and unsafe to do especially when you are running containers with dropped capabilities or with enhanced restrictions.

**Audit:**

If you have auditing enabled as prescribed in Section 1, you can use the below command to filter out `docker exec` commands that used `--privileged` option.

```
ausearch -k docker | grep exec | grep privileged
```

**Remediation:**

Do not use `--privileged` option in `docker exec` command.

**Impact:**

None. If you need enhanced capabilities within the container, then run the container with the needed capabilities.

**Default Value:**

By default, `docker exec` command runs without `--privileged` option.

**References:**

1. https://docs.docker.com/engine/reference/commandline/exec/

**CIS Controls:**

5.1 Minimize And Sparingly Use Administrative Privileges
Minimize administrative privileges and only use administrative accounts when they are

required. Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior.

## 5.23 Ensure docker exec commands are not used with user option (Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Do not `docker exec` with `--user` option.

**Rationale:**

Using `--user` option in `docker exec` executes the command within the container as that user. This could potentially be insecure and unsafe to do especially when you are running containers with dropped capabilities or with enhanced restrictions.

For example, suppose your container is running as tomcat user (or any other non-root user), it would be possible to run a command through `docker exec` as `root` with `--user=root` option. This could potentially be dangerous.

**Audit:**

If you have auditing enabled as prescribed in Section 1, you can use the below command to filter out `docker exec` commands that used `--user` option.

```
ausearch -k docker | grep exec | grep user
```

**Remediation:**

Do not use `--user` option in `docker exec` command.

**Impact:**

None.

**Default Value:**

By default, `docker exec` command runs without `--user` option.

**References:**

1. https://docs.docker.com/engine/reference/commandline/exec/

**CIS Controls:**

5 <u>Controlled Use of Administration Privileges</u>
Controlled Use of Administration Privileges

## 5.24 Ensure cgroup usage is confirmed (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

It is possible to attach to a particular cgroup on container run. Confirming cgroup usage would ensure that containers are running under defined cgroups.

**Rationale:**

System administrators typically define cgroups under which containers are supposed to run. Even if cgroups are not explicitly defined by the system administrators, containers run under `docker` cgroup by default.

At run-time, it is possible to attach to a different cgroup other than the one that was expected to be used. This usage should be monitored and confirmed. By attaching to a different cgroup than the one that is expected, excess permissions and resources might be granted to the container and thus, can prove to be unsafe.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
CgroupParent={{ .HostConfig.CgroupParent }}'
```

The above command would return the cgroup under which the containers are running. If it is blank, it means containers are running under default docker cgroup. In that case, this recommendation is compliant. If the containers are found to be running under cgroup other than the one that was expected, this recommendation is non-compliant.

**Remediation:**

Do not use `--cgroup-parent` option in `docker run` command unless needed.

**Impact:**

None.

**Default Value:**

By default, containers run under `docker` cgroup.

**References:**

1. https://docs.docker.com/engine/reference/run/#specify-custom-cgroups
2. https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Resource_Management_Guide/ch01.html

**CIS Controls:**

18 Application Software Security
Application Software Security

## 5.25 Ensure the container is restricted from acquiring additional privileges (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Restrict the container from acquiring additional privileges via suid or sgid bits.

**Rationale:**

A process can set the `no_new_priv` bit in the kernel. It persists across fork, clone and execve. The `no_new_priv` bit ensures that the process or its children processes do not gain any additional privileges via suid or sgid bits. This way a lot of dangerous operations become a lot less dangerous because there is no possibility of subverting privileged binaries.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
SecurityOpt={{ .HostConfig.SecurityOpt }}'
```

The above command should return all the security options currently configured for the containers. `no-new-privileges` should also be one of them.

**Remediation:**

For example, you should start your container as below:

```
docker run --rm -it --security-opt=no-new-privileges ubuntu bash
```

**Impact:**

`no_new_priv` prevents LSMs like SELinux from transitioning to process labels that have access not allowed to the current process.

**Default Value:**

By default, new privileges are not restricted.

**References:**

1. https://github.com/projectatomic/atomic-site/issues/269

2. https://github.com/docker/docker/pull/20727
3. https://www.kernel.org/doc/Documentation/prctl/no_new_privs.txt
4. https://lwn.net/Articles/475678/
5. https://lwn.net/Articles/475362/

**CIS Controls:**

5 Controlled Use of Administration Privileges
Controlled Use of Administration Privileges

## 5.26 Ensure container health is checked at runtime (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

If the container image does not have an `HEALTHCHECK` instruction defined, use `--health-cmd` parameter at container runtime for checking container health.

**Rationale:**

One of the important security triads is availability. If the container image you are using does not have a pre-defined `HEALTHCHECK` instruction, use the `--health-cmd` parameter to check container health at runtime.

Based on the reported health status, you could take necessary actions.

**Audit:**

Run the below command and ensure that all the containers are reporting health status:

```
docker ps --quiet | xargs docker inspect --format '{{ .Id }}: Health={{
.State.Health.Status }}'
```

**Remediation:**

Run the container using `--health-cmd` and the other parameters.

For example,

```
docker run -d --health-cmd='stat /etc/passwd || exit 1' nginx
```

**Impact:**

None.

**Default Value:**

By default, health checks are not done at container runtime.

**References:**

1. https://docs.docker.com/engine/reference/run/#healthcheck

**CIS Controls:**

18 <u>Application Software Security</u>
Application Software Security

## 5.27 Ensure docker commands always get the latest version of the image (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Always ensure that you are using the latest version of the image within your repository and not the cached older versions.

**Rationale:**

Multiple docker commands such as `docker pull`, `docker run`, etc. are known to have an issue that by default, they extract the local copy of the image, if present, even though there is an updated version of the image with the "same tag" in the upstream repository. This could lead to using older and vulnerable images.

**Audit:**

**Step 1**: Open your image repository and list the image version history for the image you are inspecting.

**Step 2**: Observe the status when the `docker pull` command is triggered.

If the status is shown as `Image is up to date`, it means that you are getting the cached version of the image.

**Step 3**: Match the version of the image you are running with the latest version reported in your repository which tells if you are running the cached version or the latest copy.

**Remediation:**

Use proper version pinning mechanisms (the latest tag which is assigned by default is still vulnerable to caching attacks) to avoid extracting the cached older versions. Version pinning mechanisms should be used for base images, packages, and entire images too. You can customize version pinning rules as per your requirements.

**Impact:**

None

**Default Value:**

By default, docker commands extract the local copy unless version pinning mechanisms are used or the local cache is cleared.

**References:**

1. https://github.com/docker/docker/pull/16609

**CIS Controls:**

18.1 Use Only Vendor-supported Software
For all acquired application software, check that the version you are using is still supported by the vendor. If not, update to the most current version and install all relevant patches and vendor security recommendations.

## 5.28 Ensure PIDs cgroup limit is used (Scored)

**Profile Applicability:**

• Level 1 - Docker

**Description:**

Use `--pids-limit` flag at container runtime.

**Rationale:**

Attackers could launch a fork bomb with a single command inside the container. This fork bomb can crash the entire system and requires a restart of the host to make the system functional again. PIDs cgroup `--pids-limit` will prevent this kind of attacks by restricting the number of forks that can happen inside a container at a given time.

**Audit:**

Run the below command and ensure that `PidsLimit` is not set to 0 or -1. A `PidsLimit` of 0 or -1 means that any number of processes can be forked inside the container concurrently.

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
PidsLimit={{ .HostConfig.PidsLimit }}'
```

**Remediation:**

Use `--pids-limit` flag while launching the container with an appropriate value.

For example,

```
docker run -it --pids-limit 100 <Image_ID>
```

In the above example, the number of processes allowed to run at any given time is set to 100. After a limit of 100 concurrently running processes is reached, docker would restrict any new process creation.

**Impact:**

Set the PIDs limit value as appropriate. Incorrect values might leave the containers unusable.

**Default Value:**

The Default value for `--pids-limit` is `0` which means there is no restriction on the number of forks. Also, note that PIDs cgroup limit works only for the kernel versions 4.3+.

**References:**

1. https://github.com/docker/docker/pull/18697
2. https://docs.docker.com/engine/reference/commandline/run/#options

**CIS Controls:**

18 <u>Application Software Security</u>
Application Software Security

## 5.29 Ensure Docker's default bridge docker0 is not used (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Do not use Docker's default bridge `docker0`. Use docker's user-defined networks for container networking.

**Rationale:**

Docker connects virtual interfaces created in the bridge mode to a common bridge called `docker0`. This default networking model is vulnerable to ARP spoofing and MAC flooding attacks since there is no filtering applied.

**Audit:**

Run the below command, and verify that containers are on a user-defined network and not the default `docker0` bridge.

```
docker network ls --quiet | xargs xargs docker network inspect --format '{{
.Name }}: {{ .Options }}'
```

**Remediation:**

Follow Docker documentation and setup a user-defined network. Run all the containers in the defined network.

**Impact:**

You have to manage the user-defined networks.

**Default Value:**

By default, docker runs containers on its `docker0` bridge.

**References:**

1. https://github.com/nyantec/narwhal
2. https://arxiv.org/pdf/1501.02967
3. https://docs.docker.com/engine/userguide/networking/

**CIS Controls:**

9 <u>Limitation and Control of Network Ports, Protocols, and Services</u>
Limitation and Control of Network Ports, Protocols, and Services

## 5.30 Ensure the host's user namespaces is not shared (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Do not share the host's user namespaces with the containers.

**Rationale:**

User namespaces ensure that a root process inside the container will be mapped to a non-root process outside the container. Sharing the user namespaces of the host with the container thus does not isolate users on the host with users on the containers.

**Audit:**

Run the below command and ensure that it does not return any value for `UsernsMode`. If it returns a value of `host`, it means the host user namespace is shared with the containers.

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
UsernsMode={{ .HostConfig.UsernsMode }}'
```

**Remediation:**

Do not share user namespaces between host and containers.

For example, do not run a container as below:

```
docker run --rm -it --userns=host ubuntu bash
```

**Impact:**

None

**Default Value:**

By default, the host user namespace is shared with the containers until user namespace support is enabled.

**References:**

1. https://docs.docker.com/engine/security/userns-remap/
2. https://docs.docker.com/engine/reference/commandline/run/#options
3. https://github.com/docker/docker/pull/12648

4. https://events.linuxfoundation.org/sites/events/files/slides/User%20Namespaces%20-%20ContainerCon%202015%20-%2016-9-final_0.pdf

**CIS Controls:**

12 <u>Boundary Defense</u>
Boundary Defense

## 5.31 Ensure the Docker socket is not mounted inside any containers (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

The docker socket `docker.sock` should not be mounted inside a container.

**Rationale:**

If the docker socket is mounted inside a container it would allow processes running within the container to execute docker commands which effectively allows for full control of the host.

**Audit:**

```
docker ps --quiet --all | xargs docker inspect --format '{{ .Id }}:
Volumes={{ .Mounts }}' | grep docker.sock
```

The above command would return any instances where `docker.sock` had been mapped to a container as a volume.

**Remediation:**

Ensure that no containers mount `docker.sock` as a volume.

**Impact:**

None

**Default Value:**

By default, `docker.sock` is not mounted inside containers.

**References:**

1. https://raesene.github.io/blog/2016/03/06/The-Dangers-Of-Docker.sock/
2. https://forums.docker.com/t/docker-in-docker-vs-mounting-var-run-docker-sock/9450/2
3. https://github.com/docker/docker/issues/21109

**CIS Controls:**

9 <u>Limitation and Control of Network Ports, Protocols, and Services</u>
Limitation and Control of Network Ports, Protocols, and Services

# 6 Docker Security Operations

This sections covers some of the operational security aspects for Docker deployments. These are best practices that should be followed. Most of the recommendations here are just reminders that organizations should extend their current security best practices and policies to include containers.

## 6.1 Ensure image sprawl is avoided (Not Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Do not keep a large number of container images on the same host. Use only tagged images as appropriate.

**Rationale:**

Tagged images are useful to fall back from "latest" to a specific version of an image in production. Images with unused or old tags may contain vulnerabilities that might be exploited, if instantiated. Additionally, if you fail to remove unused images from the system and there are various such redundant and unused images, the host filesystem may become full and could lead to denial of service.

**Audit:**

**Step 1** Make a list of all image IDs that are currently instantiated by executing below command:

```
docker images --quiet | xargs docker inspect --format '{{ .Id }}: Image={{ .Config.Image }}'
```

**Step 2**: List all the images present on the system by executing below command:

```
docker images
```

**Step 3**: Compare the list of image IDs populated from Step 1 and Step 2 and find out images that are currently not being instantiated. If any such unused or old images are found, discuss with the system administrator the need to keep such images on the system. If such a need is not justified enough, then this recommendation is non-compliant.

**Remediation:**

Keep the set of the images that you actually need and establish a workflow to remove old or stale images from the host. Additionally, use features such as pull-by-digest to get specific images from the registry.

Additionally, you can follow below set of steps to find out unused images on the system and delete them.

**Step 1** Make a list of all image IDs that are currently instantiated by executing below command:

```
docker images --quiet | xargs docker inspect --format '{{ .Id }}: Image={{
.Config.Image }}'
```

**Step 2**: List all the images present on the system by executing below command:

```
docker images
```

**Step 3**: Compare the list of image IDs populated from Step 1 and Step 2 and find out images that are currently not being instantiated.

**Step 4**: Decide if you want to keep the images that are not currently in use. If not delete them by executing below command:

```
docker rmi $IMAGE_ID
```

**Impact:**

None

**Default Value:**

Images and layered filesystems remain accessible on the host until the administrator removes all tags that refer to those images or layers.

**References:**

1. http://craiccomputing.blogspot.in/2014/09/clean-up-unused-docker-containers-and.html
2. https://forums.docker.com/t/command-to-remove-all-unused-images/20/8
3. https://github.com/docker/docker/issues/9054
4. https://docs.docker.com/engine/reference/commandline/rmi/
5. https://docs.docker.com/engine/reference/commandline/pull/
6. https://github.com/docker/docker/pull/11109

**CIS Controls:**

18 <u>Application Software Security</u>
Application Software Security

## 6.2 Ensure container sprawl is avoided (Not Scored)

**Profile Applicability:**

- Level 1 - Linux Host OS

**Description:**

Do not keep a large number of containers on the same host.

**Rationale:**

The flexibility of containers makes it easy to run multiple instances of applications and indirectly leads to Docker images that exist at varying security patch levels. It also means that you are consuming host resources that otherwise could have been used for running 'useful' containers. Having more than just the manageable number of containers on a particular host makes the situation vulnerable to mishandling, misconfiguration and fragmentation. Thus, avoid container sprawl and keep the number of containers on a host to a manageable total.

**Audit:**

**Step 1** - Find the total number of containers you have on the host:

```
docker info --format '{{ .Containers }}'
```

**Step 2** - Execute the below commands to find the total number of containers that are actually running or in the stopped state on the host.

```
docker info --format '{{ .ContainersStopped }}'
docker info --format '{{ .ContainersRunning }}'
```

If the difference between the number of containers that are stopped on the host and the number of containers that are actually running on the host is large (say 25 or more), then perhaps, the containers are sprawled on the host.

**Remediation:**

Periodically check your container inventory per host and clean up the stopped containers using the below command:

```
docker container prune
```

**Impact:**

If you keep way too few number of containers per host, then perhaps you are not utilizing your host resources very adequately.

**Default Value:**

By default, Docker does not restrict the number of containers you may have on a host.

**References:**

1. https://zeltser.com/security-risks-and-benefits-of-docker-application/
2. http://searchsdn.techtarget.com/feature/Docker-networking-How-Linux-containers-will-change-your-network

**CIS Controls:**

18 Application Software Security
Application Software Security

# 7 Docker Swarm Configuration

This section lists the recommendations that alter and secure the behavior of the Docker Swarm. If you are not using Docker Swarm then the recommendations in this section do not apply.

## 7.1 Ensure swarm mode is not Enabled, if not needed (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Do not enable swarm mode on a docker engine instance unless needed.

**Rationale:**

By default, a Docker engine instance will not listen on any network ports, with all communications with the client coming over the Unix socket. When Docker swarm mode is enabled on a docker engine instance, multiple network ports are opened on the system and made available to other systems on the network for the purposes of cluster management and node communications.

Opening network ports on a system increase its attack surface and this should be avoided unless required.

**Audit:**

Review the output of the `docker info` command. If the output includes `Swarm: active` it indicates that swarm mode has been activated on the Docker engine. Confirm if swarm mode on the docker engine instance is actually needed.

**Remediation:**

If swarm mode has been enabled on a system in error, run

```
docker swarm leave
```

**Impact:**

None.

**Default Value:**

By default, docker swarm mode is not enabled.

**References:**

1. https://docs.docker.com/engine/reference/commandline/swarm_init/

**CIS Controls:**

9.1 Limit Open Ports, Protocols, and Services
Ensure that only ports, protocols, and services with validated business needs are running on each system.

## 7.2 Ensure the minimum number of manager nodes have been created in a swarm (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Ensure that the minimum number of required manager nodes is created in a swarm.

**Rationale:**

Manager nodes within a swarm have control over the swarm and change its configuration modifying security parameters. Having excessive manager nodes could render the swarm more susceptible to compromise.

If fault tolerance is not required in the manager nodes, a single node should be elected as a manger. If fault tolerance is required then the smallest practical odd number to achieve the appropriate level of tolerance should be configured.

**Audit:**

Run `docker info` and verify the number of managers.

```
docker info --format '{{ .Swarm.Managers }}'
```

Alternatively run the below command.

```
docker node ls | grep 'Leader'
```

**Remediation:**

If an excessive number of managers is configured, the excess can be demoted as worker using the following command:

```
docker node demote <ID>
```

Where is the node ID value of the manager to be demoted.

**Impact:**

None

**Default Value:**

A single manager is all that is required to start a given cluster.

**References:**

1. https://docs.docker.com/engine/swarm/manage-nodes/
2. https://docs.docker.com/engine/swarm/admin_guide/#/add-manager-nodes-for-fault-tolerance

**CIS Controls:**

5 Controlled Use of Administration Privileges
Controlled Use of Administration Privileges

## 7.3 Ensure swarm services are binded to a specific host interface (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

By default, the docker swarm services will listen to all interfaces on the host, which may not be necessary for the operation of the swarm where the host has multiple network interfaces.

**Rationale:**

When a swarm is initialized the default value for the `--listen-addr` flag is `0.0.0.0:2377` which means that the swarm services will listen on all interfaces on the host. If a host has multiple network interfaces this may be undesirable as it may expose the docker swarm services to networks which are not involved in the operation of the swarm.

By passing a specific IP address to the `--listen-addr`, a specific network interface can be specified limiting this exposure.

**Audit:**

List the network listener on port 2377/TCP (the default for docker swarm) and confirm that it is only listening on specific interfaces. For example, using ubuntu this could be done with the following command:

```
netstat -lt | grep -i 2377
```

**Remediation:**

Remediation of this requires re-initialization of the swarm specifying a specific interface for the `--listen-addr` parameter.

**Impact:**

None

**Default Value:**

By default, docker swarm services listen on all available host interfaces.

**References:**

1. https://docs.docker.com/engine/reference/commandline/swarm_init/#--listen-addr
2. https://docs.docker.com/engine/swarm/admin_guide/#recover-from-disaster

**Notes:**

A couple of points I noted looking at this one. there doesn't seem to be a parameter for docker swarm update to change the listen-addr. For the remediation I did wonder if --force-new-swarm could be used to change this, but I'm not sure what other effects that would have on the swarm so just left with a general requirement to re-initialize the swarm.

Also interestingly the node communication service running on 7946/TCP doesn't respect the --listen-addr parameter. this seems like a bug to me, I'll likely file an issue on github for it after a bit more exploration.

**CIS Controls:**

9 Limitation and Control of Network Ports, Protocols, and Services
Limitation and Control of Network Ports, Protocols, and Services

## 7.4 Ensure data exchanged between containers are encrypted on different nodes on the overlay network (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Encrypt data exchanged between containers on different nodes on the overlay network.

**Rationale:**

By default, data exchanged between containers on different nodes on the overlay network is not encrypted. This could potentially expose traffic between the container nodes.

**Audit:**

Run the below command and ensure that each overlay network has been `encrypted`.

```
docker network ls --filter driver=overlay --quiet | xargs docker network
inspect --format '{{.Name}} {{ .Options }}'
```

**Remediation:**

Create overlay network with `--opt encrypted` flag.

**Impact:**

None

**Default Value:**

By default, data exchanged between containers on different nodes on the overlay network are not encrypted in the Docker swarm mode.

**References:**

1. https://docs.docker.com/engine/userguide/networking/overlay-security-model/
2. https://github.com/docker/docker/issues/24253

**CIS Controls:**

14.2 Encrypt All Sensitive Information Over Less-trusted Networks
All communication of sensitive information over less-trusted networks should be

encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## 7.5 Ensure Docker's secret management commands are used for managing secrets in a Swarm cluster (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Use Docker's in-built secret management command.

**Rationale:**

Docker has various commands for managing secrets in a Swarm cluster. This is the foundation for future secret support in Docker with potential improvements such as Windows support, different backing stores, etc.

**Audit:**

On a swarm manager node, run the below command and ensure `docker secret` management is used in your environment, if applicable.

```
docker secret ls
```

**Remediation:**

Follow `docker secret` documentation and use it to manage secrets effectively.

**Impact:**

None

**Default Value:**

Not Applicable

**References:**

1. https://docs.docker.com/engine/reference/commandline/secret/

**CIS Controls:**

18 Application Software Security
Application Software Security

## 7.6 Ensure swarm manager is run in auto-lock mode (Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Run Docker swarm manager in auto-lock mode.

**Rationale:**

When Docker restarts, both the TLS key used to encrypt communication among swarm nodes, and the key used to encrypt and decrypt Raft logs on disk, are loaded into each manager node's memory. You should protect the mutual TLS encryption key and the key used to encrypt and decrypt Raft logs at rest. This protection could be enabled by initializing swarm with `--autolock` flag.

With `--autolock`enabled, when Docker restarts, you must unlock the swarm first, using a key encryption key generated by Docker when the swarm was initialized.

**Audit:**

Run the below command. If it outputs the key, it means swarm was initialized with the `--autolock` flag. If the output is `no unlock key is set`, it means that swarm was NOT initialized with the `--autolock` flag and is non-compliant with respect to this recommendation.

```
docker swarm unlock-key
```

**Remediation:**

If you are initializing swarm, use the below command.

```
docker swarm init --autolock
```

If you want to set `--autolock` on an existing swarm manager node, use the below command.

```
docker swarm update --autolock
```

**Impact:**

A swarm in auto-lock mode won't recover from a re-start without manual intervention from a user to enter the unlock key. In some deployments, this might not be good for availability.

**Default Value:**

By default, swarm manager does not run in auto-lock mode.

**References:**

1. https://docs.docker.com/engine/swarm/swarm_manager_locking/

**CIS Controls:**

14.2 Encrypt All Sensitive Information Over Less-trusted Networks
All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## 7.7 Ensure swarm manager auto-lock key is rotated periodically (Not Scored)

**Profile Applicability:**

- Level 1 - Docker

**Description:**

Rotate swarm manager auto-lock key periodically.

**Rationale:**

Swarm manager auto-lock key is not automatically rotated. You should rotate them periodically as a best practice.

**Audit:**

Currently, there is no mechanism to find out when the key was last rotated on a swarm manager node. You should check with the system administrator if there is a key rotation record and the keys were rotated at a pre-defined frequency.

**Remediation:**

Run the below command to rotate the keys.

```
docker swarm unlock-key --rotate
```

Additionally, to facilitate audit for this recommendation, maintain key rotation records and ensure that you establish a pre-defined frequency for key rotation.

**Impact:**

None

**Default Value:**

By default, keys are not rotated automatically.

**References:**

1. https://docs.docker.com/engine/reference/commandline/swarm_unlock-key/

**CIS Controls:**

14.2 <u>Encrypt All Sensitive Information Over Less-trusted Networks</u>
All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## *7.8 Ensure node certificates are rotated as appropriate (Not Scored)*

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Rotate swarm node certificates as appropriate.

**Rationale:**

Docker Swarm uses mutual TLS for clustering operations amongst its nodes. Certificate rotation ensures that in an event such as compromised node or key, it is difficult to impersonate a node. By default, node certificates are rotated every 90 days. You should rotate it more often or as appropriate in your environment.

**Audit:**

Run the below command and ensure that the node certificate `Expiry Duration` is set as appropriate.

```
docker info | grep "Expiry Duration"
```

**Remediation:**

Run the below command to set the desired expiry time.

For example,

```
docker swarm update --cert-expiry 48h
```

**Impact:**

None

**Default Value:**

By default, node certificates are rotated automatically every 90 days.

**References:**

1. https://docs.docker.com/engine/reference/commandline/swarm_update/#examples

**CIS Controls:**

14.2 <u>Encrypt All Sensitive Information Over Less-trusted Networks</u>
All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## 7.9 Ensure CA certificates are rotated as appropriate (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Rotate root CA certificates as appropriate.

**Rationale:**

Docker Swarm uses mutual TLS for clustering operations amongst its nodes. Certificate rotation ensures that in an event such as compromised node or key, it is difficult to impersonate a node. Node certificates depend upon root CA certificates. For operational security, it is important to rotate these frequently. Currently, root CA certificates are not rotated automatically. You should thus establish a process to rotate it at the desired frequency.

**Audit:**

Based on your installation path, check the time stamp on the root CA certificate file.

For example,

```
ls -l /var/lib/docker/swarm/certificates/swarm-root-ca.crt
```

The certificate should have been rotated at the established frequency.

**Remediation:**

Run the below command to rotate the certificate.

```
docker swarm ca --rotate
```

**Impact:**

None

**Default Value:**

By default, root CA certificates are not rotated.

**References:**

1. https://docs.docker.com/engine/swarm/how-swarm-mode-works/pki/#rotating-the-ca-certificate

**CIS Controls:**

14.2 Encrypt All Sensitive Information Over Less-trusted Networks
All communication of sensitive information over less-trusted networks should be encrypted. Whenever information flows over a network with a lower trust level, the information should be encrypted.

## 7.10 Ensure management plane traffic has been separated from data plane traffic (Not Scored)

**Profile Applicability:**

- Level 2 - Docker

**Description:**

Separate management plane traffic from data plane traffic.

**Rationale:**

Separating the management plane traffic from data plane traffic ensures that these traffics are on their respective paths. These paths could then be individually monitored and could be tied to different traffic control policies and monitoring. It also ensures that management plane is always reachable despite the huge volume of data flow.

**Audit:**

Run the below command on each swarm node and ensure that the management plane address is different from data plane address.

```
docker node inspect  --format '{{ .Status.Addr }}' self
```

Note: At the time of writing of this benchmark, there is no way to inspect data plane address. An issue has been raised and is in the reference link.

**Remediation:**

Initialize Swarm with dedicated interfaces for management and data planes respectively.

For example,

```
docker swarm init --advertise-addr=192.168.0.1 --data-path-addr=17.1.0.3
```

**Impact:**

You would require 2 network interface cards per node.

**Default Value:**

By default, the data plane traffic is not separated from management plane traffic.

**References:**

1. https://docs.docker.com/engine/reference/commandline/swarm_init/#--data-path-addr
2. https://github.com/moby/moby/issues/33938
3. https://github.com/moby/moby/pull/32717

**CIS Controls:**

18 Application Software Security
Application Software Security

# Appendix: Summary Table

| Control | | Set Correctly | |
|---|---|---|---|
| | | Yes | No |
| **1** | **Host Configuration** | | |
| 1.1 | Ensure a separate partition for containers has been created (Scored) | ☐ | ☐ |
| 1.2 | Ensure the container host has been Hardened (Not Scored) | ☐ | ☐ |
| 1.3 | Ensure Docker is up to date (Not Scored) | ☐ | ☐ |
| 1.4 | Ensure only trusted users are allowed to control Docker daemon (Scored) | ☐ | ☐ |
| 1.5 | Ensure auditing is configured for the docker daemon (Scored) | ☐ | ☐ |
| 1.6 | Ensure auditing is configured for Docker files and directories - /var/lib/docker (Scored) | ☐ | ☐ |
| 1.7 | Ensure auditing is configured for Docker files and directories - /etc/docker (Scored) | ☐ | ☐ |
| 1.8 | Ensure auditing is configured for Docker files and directories - docker.service (Scored) | ☐ | ☐ |
| 1.9 | Ensure auditing is configured for Docker files and directories - docker.socket (Scored) | ☐ | ☐ |
| 1.10 | Ensure auditing is configured for Docker files and directories - /etc/default/docker (Scored) | ☐ | ☐ |
| 1.11 | Ensure auditing is configured for Docker files and directories - /etc/docker/daemon.json (Scored) | ☐ | ☐ |
| 1.12 | Ensure auditing is configured for Docker files and directories - /usr/bin/docker-containerd (Scored) | ☐ | ☐ |
| 1.13 | Ensure auditing is configured for Docker files and directories - /usr/bin/docker-runc (Scored) | ☐ | ☐ |
| **2** | **Docker daemon configuration** | | |
| 2.1 | Ensure network traffic is restricted between containers on the default bridge (Scored) | ☐ | ☐ |
| 2.2 | Ensure the logging level is set to 'info' (Scored) | ☐ | ☐ |
| 2.3 | Ensure Docker is allowed to make changes to iptables (Scored) | ☐ | ☐ |
| 2.4 | Ensure insecure registries are not used (Scored) | ☐ | ☐ |
| 2.5 | Ensure aufs storage driver is not used (Scored) | ☐ | ☐ |
| 2.6 | Ensure TLS authentication for Docker daemon is configured (Scored) | ☐ | ☐ |
| 2.7 | Ensure the default ulimit is configured appropriately (Not Scored) | ☐ | ☐ |
| 2.8 | Enable user namespace support (Scored) | ☐ | ☐ |

| | | | |
|---|---|---|---|
| 2.9 | Ensure the default cgroup usage has been confirmed (Scored) | ☐ | ☐ |
| 2.10 | Ensure base device size is not changed until needed (Scored) | ☐ | ☐ |
| 2.11 | Ensure that authorization for Docker client commands is enabled (Scored) | ☐ | ☐ |
| 2.12 | Ensure centralized and remote logging is configured (Scored) | ☐ | ☐ |
| 2.13 | Ensure operations on legacy registry (v1) are Disabled (Scored) | ☐ | ☐ |
| 2.14 | Ensure live restore is Enabled (Scored) | ☐ | ☐ |
| 2.15 | Ensure Userland Proxy is Disabled (Scored) | ☐ | ☐ |
| 2.16 | Ensure daemon-wide custom seccomp profile is applied, if needed (Not Scored) | ☐ | ☐ |
| 2.17 | Ensure experimental features are avoided in production (Scored) | ☐ | ☐ |
| 2.18 | Ensure containers are restricted from acquiring new privileges (Scored) | ☐ | ☐ |
| **3** | **Docker daemon configuration files** | | |
| 3.1 | Ensure that docker.service file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.2 | Ensure that docker.service file permissions are set to 644 or more restrictive (Scored) | ☐ | ☐ |
| 3.3 | Ensure that docker.socket file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.4 | Ensure that docker.socket file permissions are set to 644 or more restrictive (Scored) | ☐ | ☐ |
| 3.5 | Ensure that /etc/docker directory ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.6 | Ensure that /etc/docker directory permissions are set to 755 or more restrictive (Scored) | ☐ | ☐ |
| 3.7 | Ensure that registry certificate file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.8 | Ensure that registry certificate file permissions are set to 444 or more restrictive (Scored) | ☐ | ☐ |
| 3.9 | Ensure that TLS CA certificate file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.10 | Ensure that TLS CA certificate file permissions are set to 444 or more restrictive (Scored) | ☐ | ☐ |
| 3.11 | Ensure that Docker server certificate file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.12 | Ensure that Docker server certificate file permissions are set to 444 or more restrictive (Scored) | ☐ | ☐ |

| 3.13 | Ensure that Docker server certificate key file ownership is set to root:root (Scored) | ☐ | ☐ |
|---|---|---|---|
| 3.14 | Ensure that Docker server certificate key file permissions are set to 400 (Scored) | ☐ | ☐ |
| 3.15 | Ensure that Docker socket file ownership is set to root:docker (Scored) | ☐ | ☐ |
| 3.16 | Ensure that Docker socket file permissions are set to 660 or more restrictive (Scored) | ☐ | ☐ |
| 3.17 | Ensure that daemon.json file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.18 | Ensure that daemon.json file permissions are set to 644 or more restrictive (Scored) | ☐ | ☐ |
| 3.19 | Ensure that /etc/default/docker file ownership is set to root:root (Scored) | ☐ | ☐ |
| 3.20 | Ensure that /etc/default/docker file permissions are set to 644 or more restrictive (Scored) | ☐ | ☐ |
| **4** | **Container Images and Build File** | | |
| 4.1 | Ensure a user for the container has been created (Scored) | ☐ | ☐ |
| 4.2 | Ensure that containers use trusted base images (Not Scored) | ☐ | ☐ |
| 4.3 | Ensure unnecessary packages are not installed in the container (Not Scored) | ☐ | ☐ |
| 4.4 | Ensure images are scanned and rebuilt to include security patches (Not Scored) | ☐ | ☐ |
| 4.5 | Ensure Content trust for Docker is Enabled (Scored) | ☐ | ☐ |
| 4.6 | Ensure HEALTHCHECK instructions have been added to the container image (Scored) | ☐ | ☐ |
| 4.7 | Ensure update instructions are not use alone in the Dockerfile  (Not Scored) | ☐ | ☐ |
| 4.8 | Ensure setuid and setgid permissions are removed in the images (Not Scored) | ☐ | ☐ |
| 4.9 | Ensure COPY is used instead of ADD in Dockerfile  (Not Scored) | ☐ | ☐ |
| 4.10 | Ensure secrets are not stored in Dockerfiles (Not Scored) | ☐ | ☐ |
| 4.11 | Ensure verified packages are only Installed (Not Scored) | ☐ | ☐ |
| **5** | **Container Runtime** | | |
| 5.1 | Ensure AppArmor Profile is Enabled (Scored) | ☐ | ☐ |
| 5.2 | Ensure SELinux security options are set, if applicable (Scored) | ☐ | ☐ |
| 5.3 | Ensure Linux Kernel Capabilities are restricted within containers (Scored) | ☐ | ☐ |
| 5.4 | Ensure privileged containers are not used (Scored) | ☐ | ☐ |
| 5.5 | Ensure sensitive host system directories are not mounted on containers (Scored) | ☐ | ☐ |
| 5.6 | Ensure ssh is not run within containers (Scored) | ☐ | ☐ |

| 5.7 | Ensure privileged ports are not mapped within containers (Scored) | ☐ | ☐ |
|------|------|------|------|
| 5.8 | Ensure only needed ports are open on the container (Scored) | ☐ | ☐ |
| 5.9 | Ensure the host's network namespace is not shared (Scored) | ☐ | ☐ |
| 5.10 | Ensure memory usage for container is limited (Scored) | ☐ | ☐ |
| 5.11 | Ensure CPU priority is set appropriately on the container (Scored) | ☐ | ☐ |
| 5.12 | Ensure the container's root filesystem is mounted as read only (Scored) | ☐ | ☐ |
| 5.13 | Ensure incoming container traffic is binded to a specific host interface (Scored) | ☐ | ☐ |
| 5.14 | Ensure 'on-failure' container restart policy is set to '5' (Scored) | ☐ | ☐ |
| 5.15 | Ensure the host's process namespace is not shared (Scored) | ☐ | ☐ |
| 5.16 | Ensure the host's IPC namespace is not shared (Scored) | ☐ | ☐ |
| 5.17 | Ensure host devices are not directly exposed to containers (Not Scored) | ☐ | ☐ |
| 5.18 | Ensure the default ulimit is overwritten at runtime, only if needed (Not Scored) | ☐ | ☐ |
| 5.19 | Ensure mount propagation mode is not set to shared (Scored) | ☐ | ☐ |
| 5.20 | Ensure the host's UTS namespace is not shared (Scored) | ☐ | ☐ |
| 5.21 | Ensure the default seccomp profile is not Disabled (Scored) | ☐ | ☐ |
| 5.22 | Ensure docker exec commands are not used with privileged option (Scored) | ☐ | ☐ |
| 5.23 | Ensure docker exec commands are not used with user option (Scored) | ☐ | ☐ |
| 5.24 | Ensure cgroup usage is confirmed (Scored) | ☐ | ☐ |
| 5.25 | Ensure the container is restricted from acquiring additional privileges (Scored) | ☐ | ☐ |
| 5.26 | Ensure container health is checked at runtime (Scored) | ☐ | ☐ |
| 5.27 | Ensure docker commands always get the latest version of the image (Not Scored) | ☐ | ☐ |
| 5.28 | Ensure PIDs cgroup limit is used (Scored) | ☐ | ☐ |
| 5.29 | Ensure Docker's default bridge docker0 is not used (Not Scored) | ☐ | ☐ |
| 5.30 | Ensure the host's user namespaces is not shared (Scored) | ☐ | ☐ |
| 5.31 | Ensure the Docker socket is not mounted inside any containers (Scored) | ☐ | ☐ |
| **6** | **Docker Security Operations** | | |
| 6.1 | Ensure image sprawl is avoided (Not Scored) | ☐ | ☐ |

| 6.2 | Ensure container sprawl is avoided (Not Scored) | ☐ | ☐ |
|------|--------------------------------------------------|-----|-----|
| **7** | **Docker Swarm Configuration** | | |
| 7.1 | Ensure swarm mode is not Enabled, if not needed (Scored) | ☐ | ☐ |
| 7.2 | Ensure the minimum number of manager nodes have been created in a swarm (Scored) | ☐ | ☐ |
| 7.3 | Ensure swarm services are binded to a specific host interface (Scored) | ☐ | ☐ |
| 7.4 | Ensure data exchanged between containers are encrypted on different nodes on the overlay network (Scored) | ☐ | ☐ |
| 7.5 | Ensure Docker's secret management commands are used for managing secrets in a Swarm cluster (Not Scored) | ☐ | ☐ |
| 7.6 | Ensure swarm manager is run in auto-lock mode (Scored) | ☐ | ☐ |
| 7.7 | Ensure swarm manager auto-lock key is rotated periodically (Not Scored) | ☐ | ☐ |
| 7.8 | Ensure node certificates are rotated as appropriate (Not Scored) | ☐ | ☐ |
| 7.9 | Ensure CA certificates are rotated as appropriate (Not Scored) | ☐ | ☐ |
| 7.10 | Ensure management plane traffic has been separated from data plane traffic (Not Scored) | ☐ | ☐ |

# Appendix: Change History

| Date | Version | Changes for this version |
| --- | --- | --- |
| 01-19-17 | 1.0.0 | Initial Release |
| 07-06-17 | 1.1.0 | NEW - Ensure containers are restricted from acquiring new privileges. Ticket #5222 |
| 07-06-17 | 1.1.0 | NEW - Ensure node certificates are rotated as appropriate. Ticket #5227 |
| 07-06-17 | 1.1.0 | NEW - Ensure CA certificates are rotated as appropriate. Ticket #5228 |
| 07-06-17 | 1.1.0 | NEW - Ensure management plane traffic has been separated from data plane traffic. Ticket # 5229 |
| 07-06-17 | 1.1.0 | MODIFY -  clarified the intent of the recommendation - Ensure network traffic is restricted between containers on the default bridge. Ticket # 5210 |
| 07-06-17 | 1.1.0 | DELETED – 6.1 - Perform regular security audits of your host system and containers. Ticket # 5223 |
| 07-06-17 | 1.1.0 | DELETED – 6.2 - Monitor Docker containers usage, performance and metering. Ticket # 5224 |
| 07-06-17 | 1.1.0 | DELETED – 6.3 -Backup container data. Ticket # 5225 |
| 07-06-17 | 1.1.0 | UPDATES - Updated several reference URLs |

| 07-06-17 | 1.1.0 | NEW - New Section - "7 Docker Swarm Configuration" |
|----------|-------|-----------------------------------------------------|
| 07-06-17 | 1.1.0 | NEW – CIS CONTROLS Mappings |
| 07-06-17 | 1.1.0 | UPDATES – Updated all recommendation titles to conform to CIS standard. |
| 07-06-17 | 1.1.0 | UPDATES – Cleaned up formatting of benchmark |