Upgrade Guide

FortiSIEM 7.5.0

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

03/05/2026
FortiSIEM 7.5.0 Upgrade Guide

TABLE OF CONTENTS

Change Log

Date	Change Description
12/23/2025	Initial version of the 7.5.0 Upgrade Guide.

Target Upgrade Version

Check here to choose the appropriate target FortiSIEM upgrade version based on your current running version.

Upgrades may fail if you decide to upgrade to an unlisted version.

If your system has been patched by Fortinet Support, then contact Fortinet Support before upgrading.

Upgrade Paths

Please follow the proceeding upgrade paths to upgrade existing FortiSIEM installs to the latest 7.5.0 release.

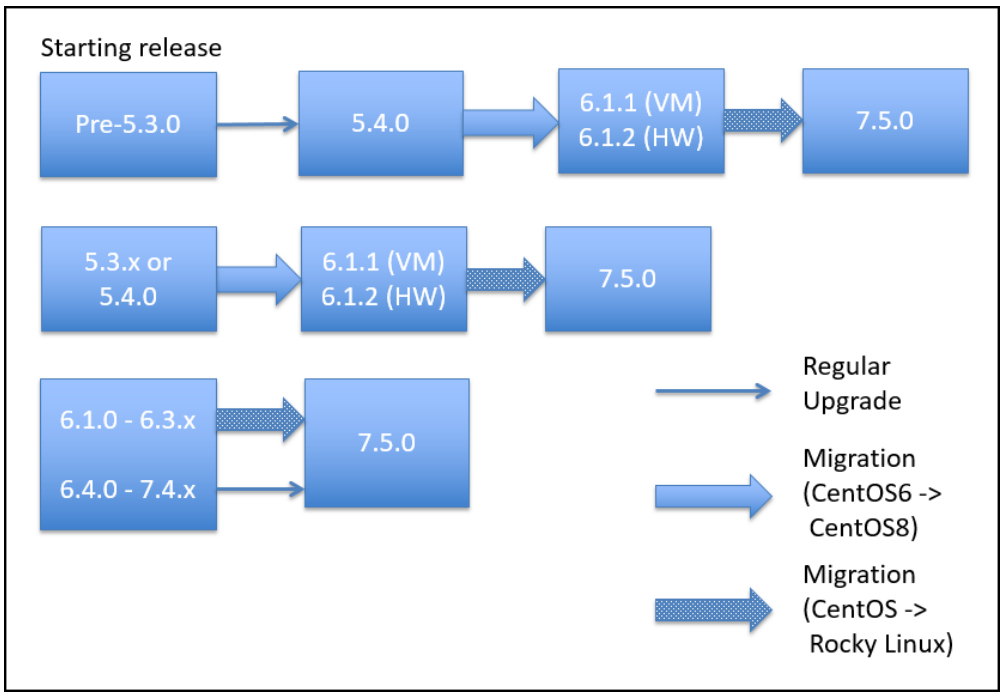

This applies for all environments.

Upgrade Supervisors and Workers

Upgrade method depends on

- FortiSIEM Version being upgraded from
- Whether High Availability (HA) is deployed

Note that FortiSIEM has released 3 versions of Supervisor HA:

- **HA V1**: released in FortiSIEM 6.7.0. In this version, *manual operations are needed* to failover to another Supervisor.
- **HA V2**: released in FortiSIEM 7.3.0. This version enables *automated* HA using Virtual IP (VIP) or DNS.
- **HA V3**: released in FortiSIEM 7.4.1. This version further improves HA V2 and does *not* need VIP or DNS Configuration.

From Version	HA Configured?	DR Configured?	Upgrade Method
7.3.1-7.4.x	Yes (HA V2 or HA V3)	Does not matter	**Highly Recommended: Use Cluster Upgrade script (see Case 1 below) to upgrade all nodes.** Manual upgrade (see Case 2 below) can be error prone, and nodes have to be upgraded in a specific order.
	No	Does not matter	**Recommended: Use Cluster Upgrade script (see Case 1 below)** or Manual Upgrade (see Case 2 below).
7.3.0	Yes (HA V2)	Does not matter	**Highly Recommended: Use Cluster Upgrade script (see Case 1 below) to upgrade all Supervisor nodes. Then upgrade Worker nodes manually (**See **Case 2 - Step 3. Upgrade Supervisor and Worker Nodes)** for steps. Fully Manual upgrade (see Case 2 below) can be error prone, and nodes have to be upgraded in a specific order.
	No	Does not matter	**Recommended: Use Cluster Upgrade script (see Case 1 below).** Manual Upgrade (see Case 2 below) can be followed.
6.7.0-7.2.x	Yes (HA V1)	Does not matter	Remove HA V1 by deleting the Followers from the Leader, upgrade manually (Case 2 below) and finally, re-do the HA V3 Configuration. When one of three appliances is in a different data center, DR functionality is provided by HA.
	No	Does not matter	Upgrade manually. See Case 2 below.
6.3.0-6.6.x	N/A	N/A	Upgrade manually. See Case 2 below.

Pre-Upgrade Checklist

To perform an upgrade, the following prerequisites must be met.

1. Make sure the FortiSIEM license is not expired.
2. Carefully consider the known issues, if any, in the Release Notes.
3. Make sure the Supervisor processes are all up, Cloud Health is Normal and that you can login to the FortiSIEM GUI and successfully discover your devices.
4. Before upgrading the ClickHouse Data Nodes, make sure that none of them are in read-only mode. Follow Advanced Operations - Recovering from Read-Only Mode to detect and then recover from read-only mode. Upgrade will fail if a node is in read-only mode.
5. Make sure the Supervisor can connect to the Internet on port 443 to the Rocky Linux 8 OS repositories (`os-pkgs-cdn.fortisiem.fortinet.com` and `os-pkgs-r8.fortisiem.fortinet.com`) hosted by Fortinet, to get the latest OS packages. Connectivity can be either directly or via a proxy. For proxy based upgrades, see Upgrade via Proxy. If Internet connectivity is not available, then follow the Offline Installation and Upgrade Guide.
6. If you have a VM based deployment, then take a snapshot of the running FortiSIEM instance.

Case 1: Cluster Upgrade

This section describes steps to upgrade a FortiSIEM deployment. The same procedures work whether it is an all-in-one deployment or a cluster with Supervisors and Workers.

Step 1. Check Licensed Supervisor Node's Public Keys

During cluster upgrade, Licensed Supervisor will need to SSH to Worker and other Supervisor nodes. To do this, Licensed Supervisor node's Public SSH Key needs to be defined in the GUI. Licensed Supervisor will push its Public SSH Key to other nodes, which will enable it to securely communicate to the other nodes. Follow these steps to make sure SSH Keys are defined.

1. Login to Licensed Supervisor via GUI.
2. Go to **Admin > License > Nodes**, select the Licensed Supervisor entry and click **Edit**.
3. Check if the Admin SSH Public Key and HA User SSH Public Key are defined.
4. If the two keys are defined, then proceed to Step 2.
5. Otherwise generate the two keys as follows:
   a. SSH to Licensed Supervisor.
   b. Get admin user's SSH Public Key by running:
      ```
      cat /opt/phoenix/bin/.ssh/id_rsa.pub
      ```
   c. If key not present, then generate it as follows
      ```
      su - admin
      ssh-keygen -t rsa -b 4096
      ```
   d. Get pghauser's SSH Public Key by running:
      ```
      cat /home/pghauser/.ssh/id_rsa.pub
      ```
   e. If key not present, then generate it as follows:
      ```
      su - pghauser
      ```

```
ssh-keygen -t rsa -b 4096
```

6. Go to **Admin > License > Nodes**, select the Licensed Supervisor entry and click **Edit**. Enter the two SSH Keys from Step 5.
7. Click **Save**.
8. Wait a few minutes for the keys to propagate to the other nodes.

Step 2. Remove DR if your Environment has DR Configured

1. Login to the Primary Licensed Supervisor node.
2. Go to **Admin > License > Nodes**.
3. Select Secondary DR node.
4. Click **Delete**.

Step 3. Stop App Servers on Non-Licensed Supervisor Nodes

If you are running HA, then you must stop the App Servers on all nodes other than the Licensed Supervisor node, to prevent accidental updates to PostgreSQL database while the database schema is being upgraded. Take the following steps.

1. Login to GUI and determine the Licensed Supervisor node.
2. Login as root to every non-Licensed Supervisor node and repeat these steps:
3. Run the following commands:

```
systemctl stop phxctl
killall -9 java
```

4. Make sure App Server is staying down.

Step 4. Run Cluster Upgrade

Follow these steps.

1. Login to the **Licensed Supervisor** via SSH.
2. Create the path `/opt/upgrade`.
   `mkdir -p /opt/upgrade`
3. Download the upgrade zip package `FSM_Upgrade_All_7.5.0_build0590.zip`, then upload it to the Supervisor node under the `/opt/upgrade/` folder.
   Example (From Linux CLI):
   `scp FSM_Upgrade_All_7.5.0_build0590.zip root@10.10.10.15:/opt/upgrade/`
4. Go to `/opt/upgrade`.
   `cd /opt/upgrade`
5. Use 7za to extract the upgrade zip package.
   Note: 7za replaces unzip for FortiSIEM 7.1.0 and later to avert unzip security vulnerabilities.
   `7za x FSM_Upgrade_All_7.5.0_build0590.zip`
6. Go to the FSM_Upgrade_All_7.5.0_build0590 directory.
   `cd FSM_Upgrade_All_7.5.0_build0590`

7. Run a screen.
   `screen -S upgrade`
   Note: This is intended for situations where network connectivity is less than favorable. If there is any connection loss, log back into the SSH console and return to the virtual screen by using the following command.
   `screen -r`
8. If you are running 7.3.0 or higher, then run the following cluster upgrade script from Licensed Supervisor node:
   `python fsm_cluster_upgrade.py`
   and proceed to step 9.
   If you are upgrading from a version earlier than 7.3.0, then run the following command to perform the upgrade:
   `sh upgrade.sh`
   After the process is completed, perform a basic health check by running the following command.
   `phstatus`
   All processes should be up and running.
9. Upgrade Status of various nodes can be found in this file on the Licensed Supervisor node.
   `/var/tmp/upgrade_status_management.json`
10. Upgrade logs can be found in this file on the Licensed Supervisor node.
    `/var/log/fsm_cluster_upgrade.log`
    For additional troubleshooting, you can also review the following files.
    `/var/log/screen_*.log`
11. After upgrade is complete, check Cloud Health.
12. The upgrade logs are stored in compressed format in this location.
    `/var/log/*fsm_cluster_upgrade.log.xz`

Case 2: Manual Upgrade

Upgrade Notes

1. **Upgrade Order**: In manual mode, FortiSIEM nodes must be upgraded in the following order. This ensures that ClickHouse Cluster remains up and running during the upgrade.
   a. All Nodes running PostgreSQL Database
   b. All Supervisor Nodes without PostgreSQL Database
   c. Worker Nodes (ClickHouse Data Nodes)
   d. All Worker nodes in ClickHouse Keeper Cluster
2. Upgrade each ClickHouse Keeper node one by one to maintain quorum. This ensures that the ClickHouse Data Cluster will not become read only.
3. You can upgrade ClickHouse Data Nodes in parallel, but do not upgrade two Workers from the *same* shard in parallel, since that would render the shard ineffective for both read and write.

   As an example, suppose the ClickHouse Cluster consists of
   - Keeper Cluster – nodes K1, K2, K3
   - Data Cluster

- Shard1 – W1, W2
- Shard 2 – W3, W4
- Shard 3 – W5, W6

Then the Worker upgrade order is

(W1, W3, W5) -> (W2, W4, W6) -> K1-> K2-> K3

The nodes within the parentheses () can be upgraded in parallel.

Step 1. Remove DR if your Environment has DR Configured

1. Login to the Primary Licensed Supervisor node.
2. Go to **Admin > License > Nodes**.
3. Select Secondary DR node.
4. Click **Delete**.

Step 2. Stop App Servers on Non-Licensed Supervisor Nodes

If you are running HA, then you must stop the App Servers on all nodes other than the Licensed Supervisor node, to prevent accidental updates to PostgreSQL database while the database schema is being upgraded. Take the following steps.

1. Login to GUI and determine the Licensed Supervisor node.
2. Login as root to every non-Licensed Supervisor node and repeat these steps:
3. Run the following commands:

```
systemctl stop phxctl
killall -9 java
```

4. Make sure App Server is staying down.

Step 3. Upgrade Supervisor and Worker Nodes

Upgrade the Supervisor and Worker nodes in the correct order following the Upgrade Notes.

The steps to upgrade the Supervisor and Worker are identical. Follow these steps.

1. Login to the node via SSH.
2. Create the path `/opt/upgrade`.
   `mkdir -p /opt/upgrade`
3. Download the upgrade zip package `FSM_Upgrade_All_7.5.0_build0590.zip`, then upload it to the node under the `/opt/upgrade/` folder.
   Example (From Linux CLI):
   `scp FSM_Upgrade_All_7.5.0_build0590.zip root@10.10.10.15:/opt/upgrade/`
4. Go to `/opt/upgrade`.
   `cd /opt/upgrade`
5. Use 7za to extract the upgrade zip package.
   Note: 7za replaces unzip for FortiSIEM 7.1.0 and later to avert unzip security vulnerabilities.
   `7za x FSM_Upgrade_All_7.5.0_build0590.zip`
6. Go to the FSM_Upgrade_All_7.5.0_build0590 directory.
   `cd FSM_Upgrade_All_7.5.0_build0590`

   a. Run a screen.
      `screen -S upgrade`
      Note: This is intended for situations where network connectivity is less than favorable. If there is any connection loss, log back into the SSH console and return to the virtual screen by using the following command.
      `screen -r`
7. Start the upgrade process by entering the following.
   `sh upgrade.sh`
8. After the process is completed, perform a basic health check by running the following command.
   `phstatus`
   All processes should be up and running.

Upgrade Collectors

Main Upgrade Steps

Step 1: Download the Correct Collector Image

Download the Collector image from the Fortinet Support site. As an example, Collector 6.4.0 image file name is `FSM_Upgrade_All_6.4.0_build1412.zip` and matches the hash in the support site to the locally computed hash. This ensures that the file has not been corrupted in transit.

Step 2: Upload the Image to the Supervisor Node

Note: In this step, you will upload the image to the Supervisor, which will then internally create a URL for the Collectors to download the image. It is critical to set the host name in the URL correctly so that a Collector can resolve the host name. Otherwise, the image download in Step 3 will fail.

There are two solutions.

Solution 1	By default, the Supervisor's host name in **Admin > License > Nodes** is used to create the URL. If the host name is a Fully Qualified domain name and is resolvable by the Collectors, then no further action is necessary.
	For example, a host name like `c2-52-35-20-68.us-west-2.compute.amazonaws.com` is resolvable to an external IP address. A host name like `2-52-35-20-68.us-west-2.compute` is likely not resolvable.
	If the hostname is not resolvable, either create a DNS entry to allow the Collector to resolve the hostname, or add an entry to the Collector `/etc/hosts` file in the following format:
	`<ip> <host name>`
	For example:
	`10.0.1.21 2-52-35-20-68.us-west-2.compute`

Solution 2	If there is a load balancer in front of the Supervisors, or you want to override the Supervisor host name in the default image download URL, then you can enter the appropriate host name or IP after going to **Admin > Settings > System > Image Server > Custom Update** and then clicking **Save**. If you have entered a host name here, make sure that it is a Fully Qualified domain name and is resolvable by Collectors. Do this step first before proceeding to the remaining of Step 2. Note that if you create an entry in Custom Update, then it applies to ALL Collectors and Agents. This means that every Collector and Agent will get the URL with the Custom Update entry.

1. Go to **Admin > Settings > System > Image Server**.
2. Under **Collector**, in the **Version** field, enter the version you downloaded in Step 1. The format is #.#.#. Example: 6.4.0.
3. Under **Collector**, click **Select File** and select the Collector upgrade image you downloaded in Step 1.
4. Under **Collector**, click **Upload File** to upload the Collector upgrade image to the Supervisor. This may take a while depending on the network connection between your workstation and Supervisor node. FortiSIEM will validate the image hash and upload the image to Supervisor if the hash matches.

Note: If you do not want a FortiSIEM to perform a hash check, from the **Hash Check** drop-down, select **Disabled**.

5. Run the following SQL and make sure ImageSetup task is completed.

```
# psql phoenixdb phoenix -c "select type, progress from ph_task where type = 'ImageSetup'"
    type | progress
------------+----------
 ImageSetup | 100
 ImageSetup | 100
 ImageSetup | 100
(3 rows)
```

Step 3: Download the Image to the Collector

1. Go to **Admin > Health > Collector Health**.
2. From the **Columns** drop-down list, ensure **Download Status** is selected. If not, select it so the **Download Status** column is displayed.
3. Select the Collector(s) you wish to download the image to.
4. From the **Actions** drop-down list (⋮), select **Download Image**.
5. Check that the **Download Status** column shows **finished** to confirm that the download has been completed for the selected Collectors.

Step 4: Upgrade the Collectors

1. Go to **Admin > Health > Collector Health**.
2. From the **Columns** drop-down list, ensure **Version** is selected. If not, select it so the **Version** column is displayed.
3. Select the Collector(s) you wish to upgrade.
4. From the **Actions** drop-down list (⋮), select **Install Image**.
5. Check that the **Version** column shows the correct version number to confirm that the Collector(s) have upgraded successfully.

Post Upgrade Health Check

Note: If any of the checks fail, then the upgrade might have failed. In this case, contact Fortinet Support.

1. Check Cloud health and Collector health from the FortiSIEM GUI:
   - Versions display correctly.
   - All processes are up and running.
   - Resource usage is within limits.

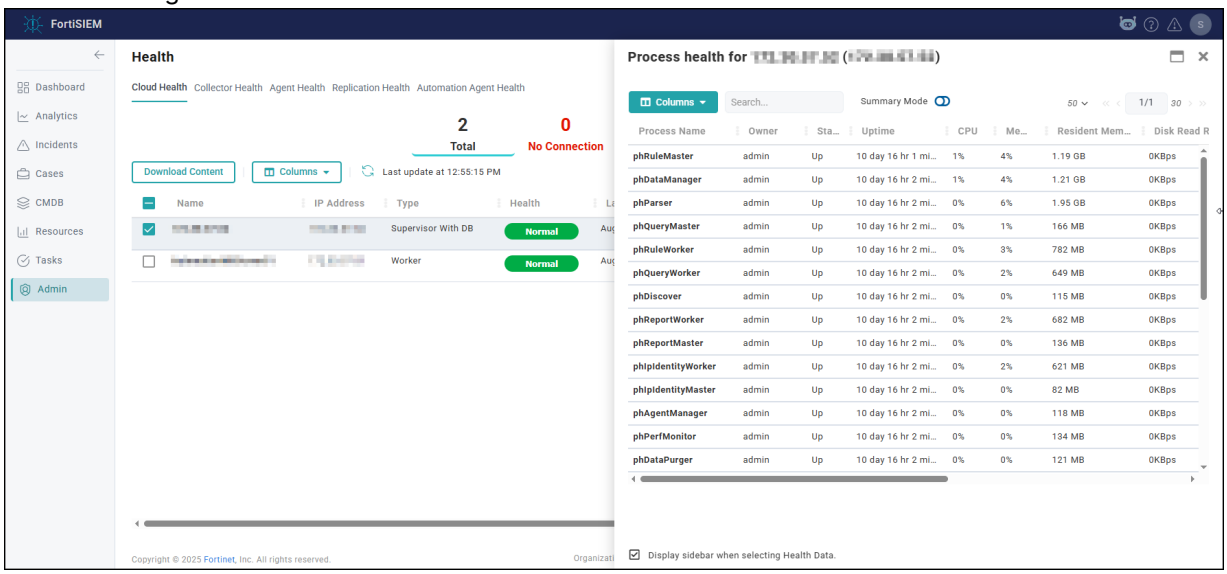

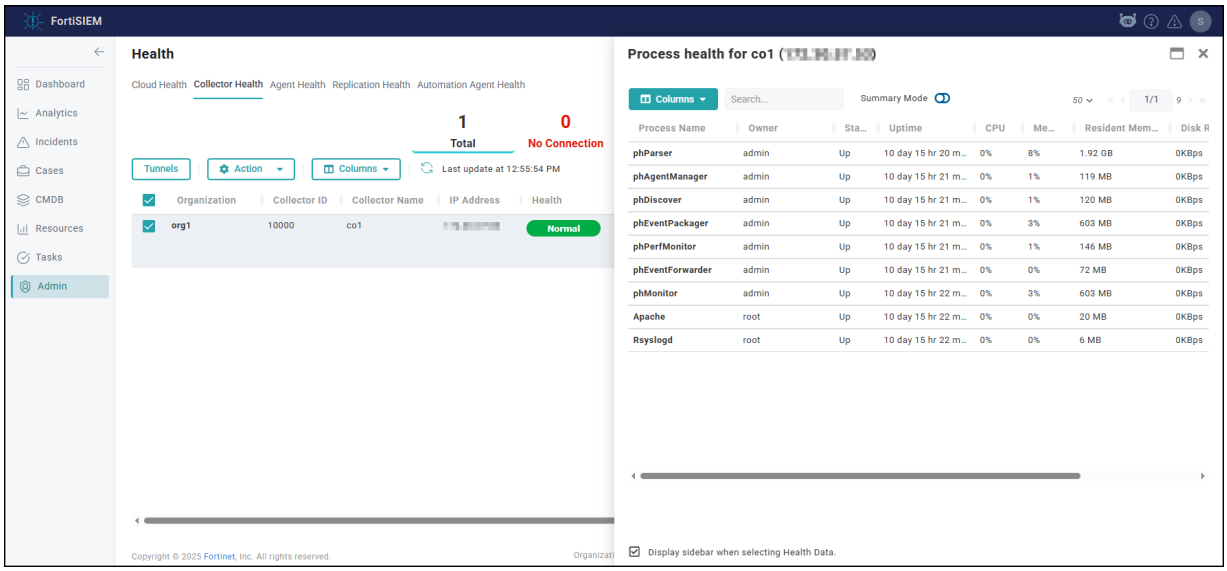

2. Check that the Redis passwords match on the Supervisor and Workers:
   - Supervisor: run the command `phLicenseTool --showRedisPassword`
   - Worker: run the command `grep -i auth /opt/node-rest-service/ecosystem.config.js`

```
[root@offlinesuper ~]# grep -i auth /opt/node-rest-service/ecosystem.config.js
      REDIS_AUTH: '4CiVtA9n1Fh2KPlkDWCjsLTzJCwiwg7F3Yok@SWhVYAnGjSB66pR1v743v5zGNJYXyB%ZB5ScQfk6ihx8L^DzhjAYO%tWQff554ERhE%U1jBtBZ%chxCLYqcvqvzswQ9',
      REDIS_AUTH: '4CiVtA9n1Fh2KPlkDWCjsLTzJCwiwg7F3Yok@SWhVYAnGjSB66pR1v743v5zGNJYXyB%ZB5ScQfk6ihx8L^DzhjAYO%tWQff554ERhE%U1jBtBZ%chxCLYqcvqvzswQ9',
[root@offlinesuper ~]# ssh root@172.30.57.231
root@172.30.57.231's password:
Last login: Thu Jul  1 13:17:46 2021 from 172.30.57.230
[root@offlineworker ~]# grep -i auth /opt/node-rest-service/ecosystem.config.js
      REDIS_AUTH: '4CiVtA9n1Fh2KPlkDWCjsLTzJCwiwg7F3Yok@SWhVYAnGjSB66pR1v743v5zGNJYXyB%ZB5ScQfk6ihx8L^DzhjAYO%tWQff554ERhE%U1jBtBZ%chxCLYqcvqvzswQ9',
      REDIS_AUTH: '4CiVtA9n1Fh2KPlkDWCjsLTzJCwiwg7F3Yok@SWhVYAnGjSB66pR1v743v5zGNJYXyB%ZB5ScQfk6ihx8L^DzhjAYO%tWQff554ERhE%U1jBtBZ%chxCLYqcvqvzswQ9',
```

3. Check that the database passwords match on the Supervisor and Workers:
   - Supervisor: run the command `phLicenseTool --showDatabasePassword`
   - Worker: run the command `phLicenseTool --showDatabasePassword`
4. Elasticsearch case: check the Elasticsearch health
   Note: Status should be Normal, not Warning as illustrated here.

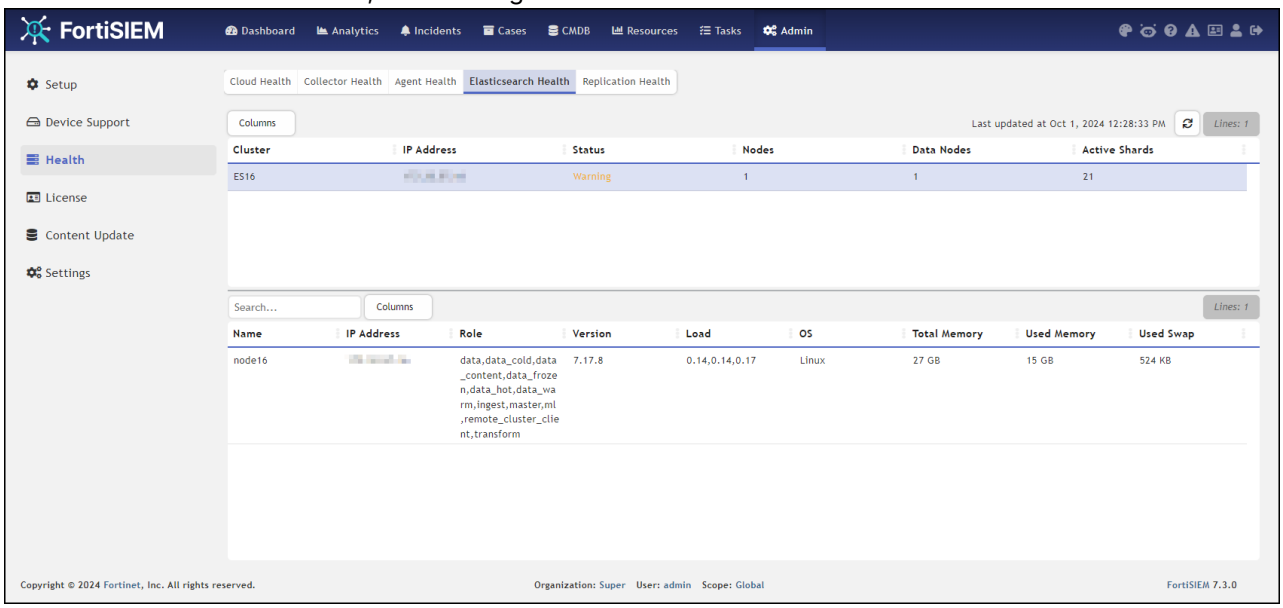

5. Check that events are received correctly:
   a. Search All Events in last 10 minutes and make sure there is data.

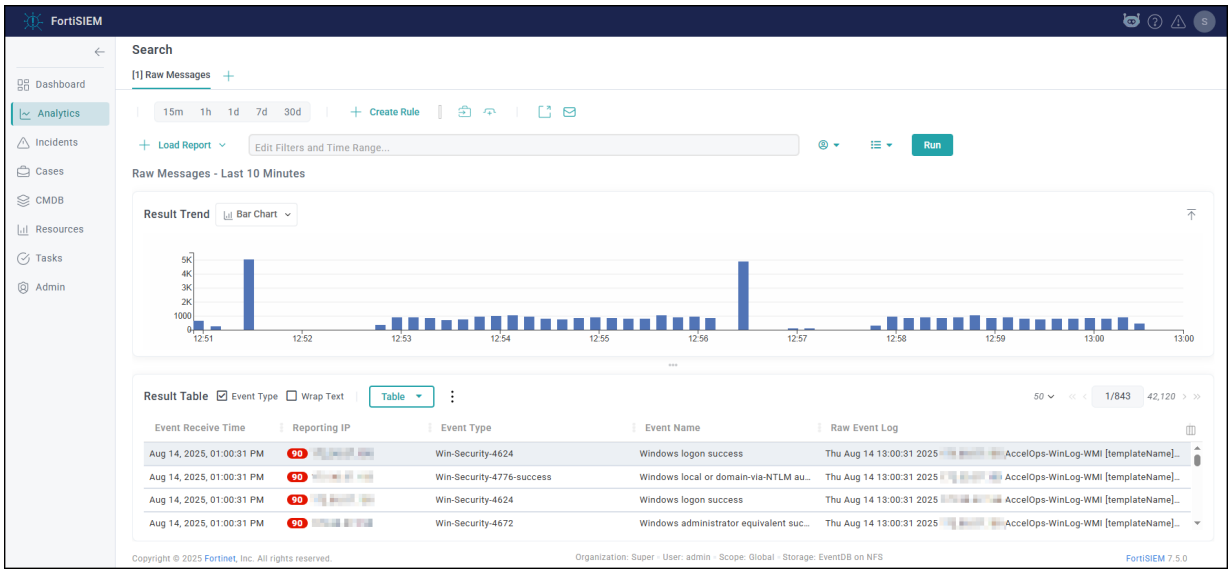

   b. Search for events from Collector and Agents and make sure there is data. Both old and new collectors and agents must work.

c. Search for events using CMDB Groups (Windows, Linux, Firewalls, etc.) and make sure there is data.

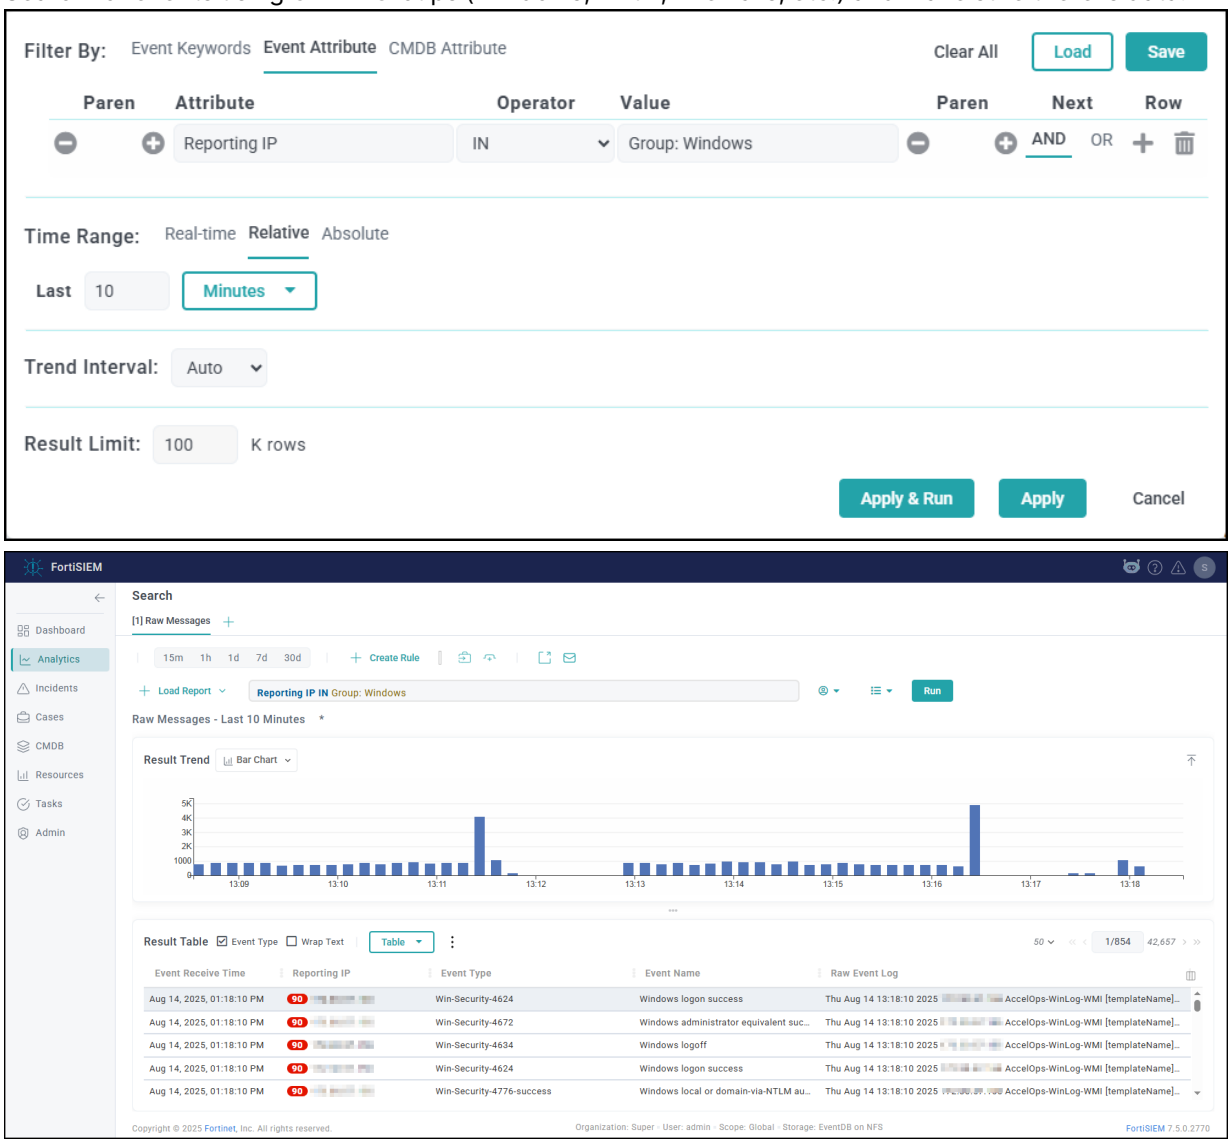

6. Make sure there are no SVN authentication errors in CMDB when you click any device name.

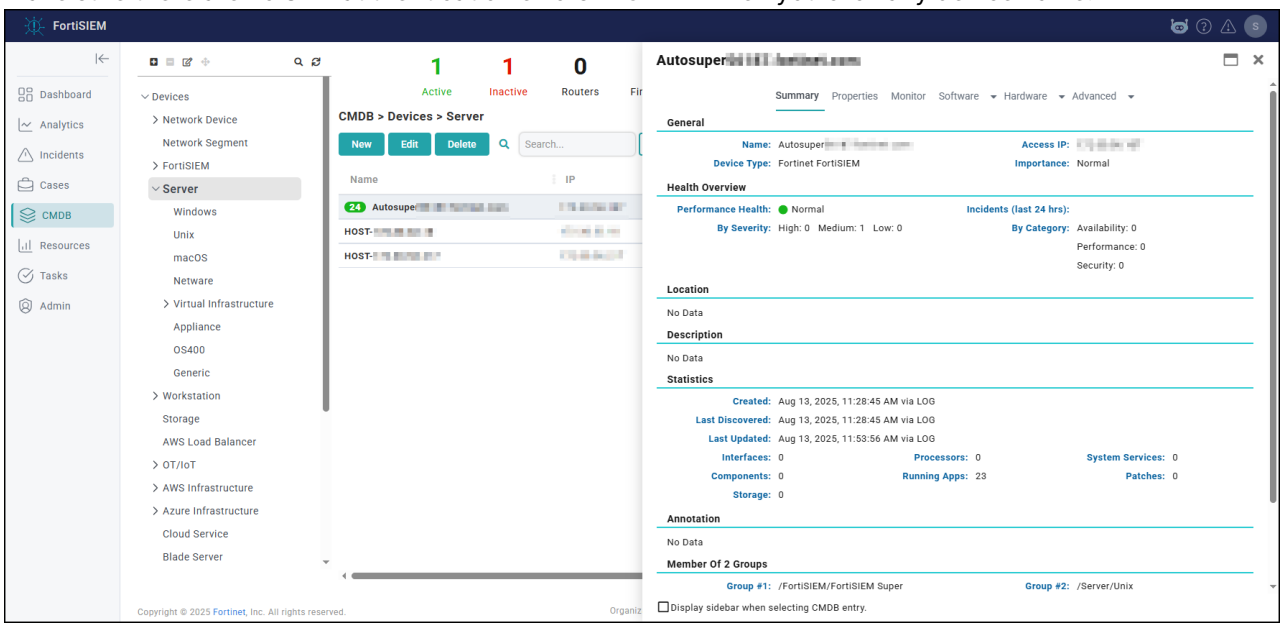

7. Make sure recent Incidents and their triggering events are displayed.

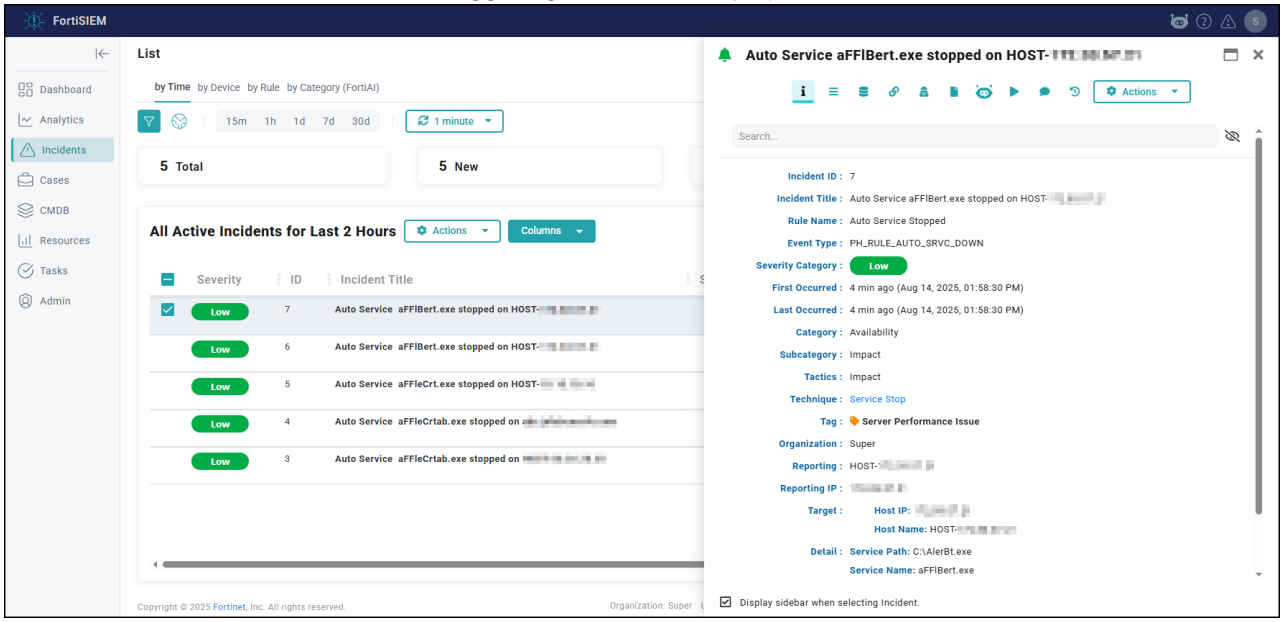

8. Check Worker for Collector Credentials by running the following command:
   `cat /etc/httpd/accounts/passwds`
   This validates that all workers contain collector credentials to log in and upload logs.
9. Run the following script on the Supervisor.
   `get-fsm-health.py --local`
   Your output should appear similar to the example output in Post Upgrade Health Check get-fsm-health.py --local Example Output.

Upgrading with FortiSIEM Manager

If you have FortiSIEM and FortiSIEM Manager deployed in your environment, then take the following steps.

1. Upgrade the FortiSIEM Manager.
2. After the FortiSIEM Manager is fully upgraded, then upgrade each FortiSIEM Cluster.

Upgrade via Proxy

During upgrade, the FortiSIEM Supervisor, Worker, or Hardware appliances (FSM-2000F, 2000G, 2200G, 3500F, 3500G, or 3600G) must be able to communicate with the Rocky Linux 8 OS repositories (`os-pkgs-cdn.fortisiem.fortinet.com` and `os-pkgs-r8.fortisiem.fortinet.com`) hosted by Fortinet, to get the latest OS packages. Follow these steps to set up this communication via proxy, before initiating the upgrade.

1. SSH to the node.
2. Create this file `etc/profile.d/proxy.sh` with the following content and then save the file.

```
PROXY_URL="<proxy-ip-or-hostname>:<proxy-port>"
export http_proxy="$PROXY_URL"
export https_proxy="$PROXY_URL"
export ftp_proxy="$PROXY_URL"
export no_proxy="127.0.0.1,localhost"
```

3. Run `source /etc/profile.d/proxy.sh`.
4. Test that you can use the proxy to successfully communicate with the two sites here: `os-pkgs-cdn.fortisiem.fortinet.com` `os-pkgs.fortisiem.fortinet.com`.
5. Begin the upgrade.

Restoring Hardware from Backup After a Failed Upgrade

Background Information

Introduced in 6.3.1, when upgrading a FortiSIEM system on hardware (500F, 500G, 2000F, 2000G, 2200G, 3500F, 3500G, 3600G) to a later version, the upgrade automatically creates a system backup of root disk, boot disk, opt disk, and for the Supervisor, also CMDB disk, and SVN disks.

This backup is stored in `/opt/hwbackup` if the `/opt` partition has 300GB or more free space. Once the backup pre-upgrade task is complete, the logs are stored at `/opt/phoenix/log/backup-upg.stdout.log` and `/opt/phoenix/log/backup-upg.stderr.log`.

The actual backup may be much smaller depending on the size of your CMDB and SVN partitions. Backups are also compressed using XZ compression. The partition itself is 500GB in size, so in most installations, you will have this much available space.

In case you do not have 300GB free space in `/opt`, the upgrade will abort quickly. In this case, you can also externally store the backup. For this, you will need to mount an external disk and create a symlink like this:

```
ln -s <external-disk-mount-point> /opt/hwbackup
```

Here is a sample listing of `/opt/hwbackup`:

```
[root@sp5747 hwbackup]# pwd
/opt/hwbackup
[root@sp5747 hwbackup]# ls -lh
total 19G
-rw-r--r-- 1 root root  824 Aug 24 17:08 fsm_backup_sha256sum_6.3.0.0331_2021-08-24-17-01.txt
-rw-r--r-- 1 root root 803M Aug 24 17:05 fsm_boot_disk_6.3.0.0331_2021-08-24-17-01.img.xz
-rw-r--r-- 1 root root  61M Aug 24 17:07 fsm_cmdb_6.3.0.0331_2021-08-24-17-01.xfsdump.xz
-rwxr-xr-x 1 root root 6.0K Aug 19 16:12 fsm_hw_restore_from_backup.sh
-rw-r--r-- 1 root root  14G Aug 24 17:05 fsm_opt_6.3.0.0331_2021-08-24-17-01.tar.xz
-rw-r--r-- 1 root root 5.0G Aug 24 17:07 fsm_root_disk_6.3.0.0331_2021-08-24-17-01.xfsdump.xz
-rw-r--r-- 1 root root  192 Aug 24 17:07 fsm_root_disk_partition_table_6.3.0.0331_2021-08-24-17-01.txt
-rw------- 1 root root 1.8K Aug 24 17:07 fsm_root_disk_vg_cfg_backup_6.3.0.0331_2021-08-24-17-01.txt
-rw-r--r-- 1 root root  13K Aug 24 17:07 fsm_svn_6.3.0.0331_2021-08-24-17-01.xfsdump.xz
-rw-r--r-- 1 root root  30K Aug 24 17:08 MegaSAS.log
[root@sp5747 hwbackup]# ./fsm_hw_restore_from_backup.sh
```

If there was a previous attempt at an upgrade, then there will already be a `/opt/hwbackup` directory. A new attempt will rename `/opt/hwbackup` to `/opt/hwbackup.1` and continue the new backup and upgrade. This means that the system will keep at most 2 backups. For instance, if you upgrade from 6.3.0 to 6.3.1 and in the future to 6.3.2, then you will have a backup of both the 6.3.0 system as well as 6.3.1 system. If you perform a hardware restore after upgrading to 7.5.0, the appliance will reboot twice during the restoration process.

Restoring from Backup

Restoring from backup will restore the root disk, boot disk, opt disk, and for the Supervisor, also CMDB disk, and SVN disks. The event data is not modified as part of an upgrade and therefore requires no restoration.

To restore from a backup, take the following steps:

1. Switch the running system to rescue mode. You will need to do the following on the VGA or serial console of the hardware.
2. Switch to rescue mode as follows after logging into the system as the 'root' user.

   `systemctl isolate rescue.target`
3. You will be prompted to type the root administrator password as shown here.

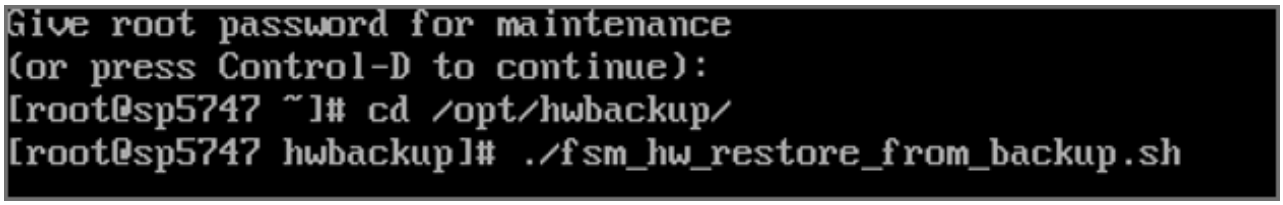

4. If the backup is stored on `/opt/hwbackup`, you can `chdir` to this. If the backup is stored on an external disk, mount the disk and symlink it again to `/opt/hwbackup`.
5. Run the restore command:

   `cd /opt/hwbackup`

   `./fsm_hw_restore_from_backup.sh`

   Note: If you run the restore program in normal multi-user mode, the script exits with an error like this:

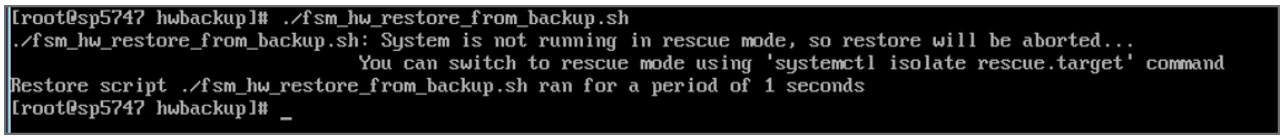

   The whole restore may take anywhere from 15 minutes to more than an hour depending on how large the CMDB/SVN partitions are. The restore script will make sure that the SHA 256 checksums for the backup files match and only then, will it proceed. If this fails, then it will stop the restore process immediately. Here are screenshots for a sample Supervisor restore from 6.3.1 to 6.3.0.0331:

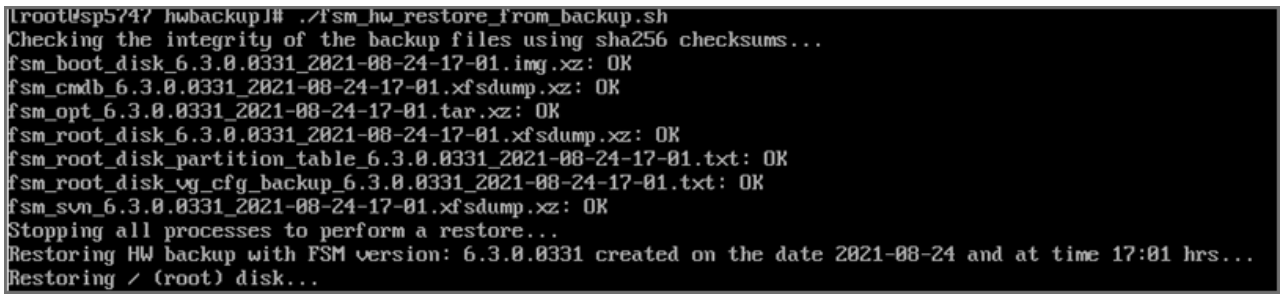

```
Restoring HW backup with FSM version: 6.3.0.0331 created on the date 2021-08-24 and at time 17:01 hrs...
Restoring / (root) disk...
xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.8 (dump format 3.0)
xfsrestore: searching media for dump
xfsrestore: examining media file 0
xfsrestore: dump description:
xfsrestore: hostname: sp5747.fortinet.com
xfsrestore: mount point: /
xfsrestore: volume: /dev/mapper/cl-root
xfsrestore: session time: Tue Aug 24 17:05:16 2021
xfsrestore: level: 0
xfsrestore: session label: "cl-root"
xfsrestore: media label: "cl-root"
xfsrestore: file system id: 511c435d-0ada-4b94-8125-6b80a63574ad
xfsrestore: session id: a9b57771-ac25-40c2-b453-a4b79e5b5ed3
xfsrestore: media id: 07670986-ce72-4f66-a4c0-2c1f74a52e0d
xfsrestore: searching media for directory dump
xfsrestore: reading directories
xfsrestore: 19595 directories and 175075 entries processed
xfsrestore: directory post-processing
xfsrestore: WARNING: unable to set secure extended attribute for proc: Operation not supported (95)
xfsrestore: restoring non-directory files
xfsrestore: status at 20:46:28: 21442/146457 files restored, 14.0% complete, 30 seconds elapsed
xfsrestore: status at 20:46:58: 38507/146457 files restored, 57.5% complete, 60 seconds elapsed
xfsrestore: status at 20:47:28: 38546/146457 files restored, 57.5% complete, 90 seconds elapsed
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of run/blkid/blkid.tab failed
 Inappropriate ioctl for device
xfsrestore: status at 20:47:58: 53052/146457 files restored, 65.0% complete, 120 seconds elapsed
xfsrestore: status at 20:48:28: 68088/146457 files restored, 68.7% complete, 150 seconds elapsed
xfsrestore: status at 20:48:58: 72511/146457 files restored, 70.2% complete, 180 seconds elapsed
xfsrestore: status at 20:49:28: 73913/146457 files restored, 73.6% complete, 210 seconds elapsed
xfsrestore: status at 20:49:58: 87298/146457 files restored, 85.1% complete, 240 seconds elapsed
xfsrestore: status at 20:50:28: 105103/146457 files restored, 88.2% complete, 270 seconds elapsed
xfsrestore: status at 20:50:58: 127998/146457 files restored, 97.4% complete, 300 seconds elapsed
```

```
xfsrestore: status at 20:50:58: 127998/146457 files restored, 97.4% complete, 300 seconds elapsed
xfsrestore: WARNING: open_by_handle of data failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of data failed: Bad file descr
iptor
xfsrestore: WARNING: open_by_handle of querydata failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of querydata failed: Bad file
descriptor
xfsrestore: WARNING: open_by_handle of cmdb failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of cmdb failed: Bad file descr
iptor
xfsrestore: WARNING: open_by_handle of svn failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of svn failed: Bad file descri
ptor
xfsrestore: WARNING: open_by_handle of opt failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of opt failed: Bad file descri
ptor
xfsrestore: WARNING: path_to_handle of var/lib/nfs/rpc_pipefs failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of var/lib/nfs/rpc_pipefs fail
ed: Bad file descriptor
xfsrestore: WARNING: path_to_handle of sys failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of sys failed: Bad file
 descriptor
xfsrestore: WARNING: path_to_handle of run/blkid failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of run/blkid failed: Bad file
descriptor
```

Note: These WARNING messages can be ignored. These are likely to be temporary system files at the Linux level when the backup was taken. At the time of backup, all FSM services are stopped.

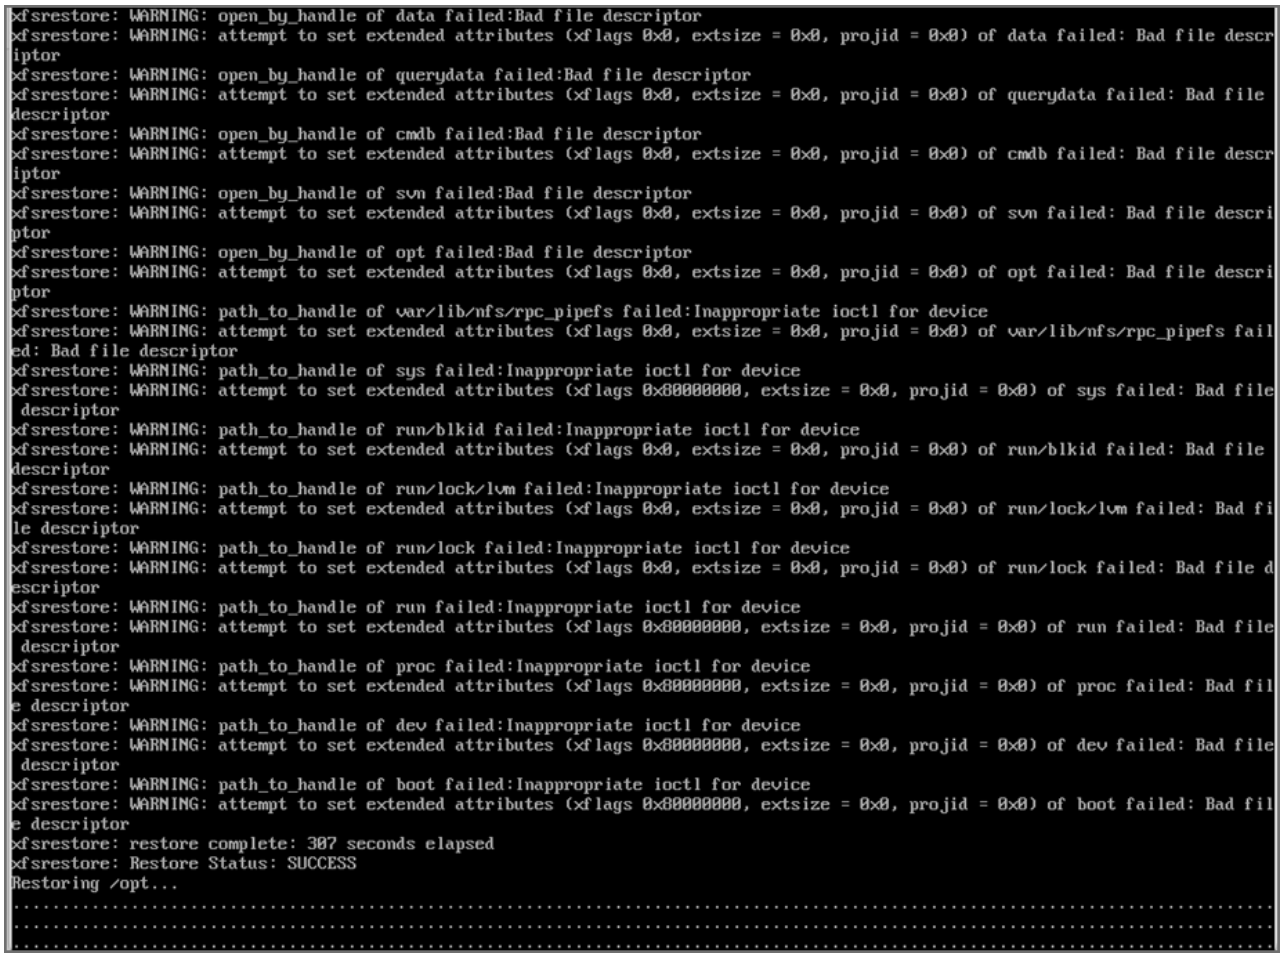
xfsrestore: WARNING: open_by_handle of data failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of data failed: Bad file descr
iptor
xfsrestore: WARNING: open_by_handle of querydata failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of querydata failed: Bad file
descriptor
xfsrestore: WARNING: open_by_handle of cmdb failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of cmdb failed: Bad file descr
iptor
xfsrestore: WARNING: open_by_handle of svn failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of svn failed: Bad file descri
ptor
xfsrestore: WARNING: open_by_handle of opt failed:Bad file descriptor
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of opt failed: Bad file descri
ptor
xfsrestore: WARNING: path_to_handle of var/lib/nfs/rpc_pipefs failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of var/lib/nfs/rpc_pipefs fail
ed: Bad file descriptor
xfsrestore: WARNING: path_to_handle of sys failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of sys failed: Bad file
descriptor
xfsrestore: WARNING: path_to_handle of run/blkid failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of run/blkid failed: Bad file
descriptor
xfsrestore: WARNING: path_to_handle of run/lock/lvm failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of run/lock/lvm failed: Bad fi
le descriptor
xfsrestore: WARNING: path_to_handle of run/lock failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x0, extsize = 0x0, projid = 0x0) of run/lock failed: Bad file d
escriptor
xfsrestore: WARNING: path_to_handle of run failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of run failed: Bad file
descriptor
xfsrestore: WARNING: path_to_handle of proc failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of proc failed: Bad fil
e descriptor
xfsrestore: WARNING: path_to_handle of dev failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of dev failed: Bad file
descriptor
xfsrestore: WARNING: path_to_handle of boot failed:Inappropriate ioctl for device
xfsrestore: WARNING: attempt to set extended attributes (xflags 0x80000000, extsize = 0x0, projid = 0x0) of boot failed: Bad fil
e descriptor
xfsrestore: restore complete: 307 seconds elapsed
xfsrestore: Restore Status: SUCCESS
Restoring /opt...

```
Restoring /opt...
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
...........................................................................................................................
.....................................................................................
Removing existing directories in cmdb
Restoring /cmdb disk...
xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.8 (dump format 3.0)
xfsrestore: searching media for dump
xfsrestore: examining media file 0
xfsrestore: dump description:
xfsrestore: hostname: sp5747.fortinet.com
xfsrestore: mount point: /cmdb
xfsrestore: volume: /dev/mapper/FSIEM3500F-phx_cmdb
xfsrestore: session time: Tue Aug 24 17:07:23 2021
xfsrestore: level: 0
xfsrestore: session label: "cmdb"
xfsrestore: media label: "cmdb"
xfsrestore: file system id: c9c48a15-b50e-4ff4-ad98-0b2b9fba0840
xfsrestore: session id: 210648e6-9f98-4cbe-96ca-5efd3c209206
xfsrestore: media id: 555b895a-1b3e-4bf8-b49f-432021e70e07
xfsrestore: searching media for directory dump
xfsrestore: reading directories
xfsrestore: 30 directories and 2643 entries processed
xfsrestore: directory post-processing
xfsrestore: restoring non-directory files
xfsrestore: restore complete: 6 seconds elapsed
xfsrestore: Restore Status: SUCCESS
Removing existing directories in svn
Restoring /svn disk...
xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.8 (dump format 3.0)
xfsrestore: searching media for dump
xfsrestore: examining media file 0
xfsrestore: dump description:
xfsrestore: hostname: sp5747.fortinet.com
xfsrestore: mount point: /svn
xfsrestore: volume: /dev/mapper/FSIEM3500F-phx_svn
```

```
xfsrestore: volume: /dev/mapper/FSIEM3500F-phx_svn
xfsrestore: session time: Tue Aug 24 17:07:39 2021
xfsrestore: level: 0
xfsrestore: session label: "svn"
xfsrestore: media label: "svn"
xfsrestore: file system id: 5dbb7da5-e17b-40cb-a3f4-5b0e59496b3d
xfsrestore: session id: 429b02ec-e8ea-4084-9590-cde65b467384
xfsrestore: media id: 26273d79-9025-4108-89f6-4616e6ad4e1b
xfsrestore: searching media for directory dump
xfsrestore: reading directories
xfsrestore: 12 directories and 43 entries processed
xfsrestore: directory post-processing
xfsrestore: restoring non-directory files
xfsrestore: restore complete: 0 seconds elapsed
xfsrestore: Restore Status: SUCCESS
Restoring /boot disk after umount...
1033060352 bytes (1.0 GB, 985 MiB) copied, 10 s, 103 MB/s
0+130005 records in
0+130005 records out
[root@sp5747 hwbackup]# 1073741824 bytes (1.1 GB, 1.0 GiB) copied, 29.1323 s, 36.9 MB/s
Restore 6.3.0.0331 complete.
Please reboot the system...
Restore script ./fsm_hw_restore_from_backup.sh ran for a period of 9 minutes and 27 seconds

[root@sp5747 hwbackup]#
```

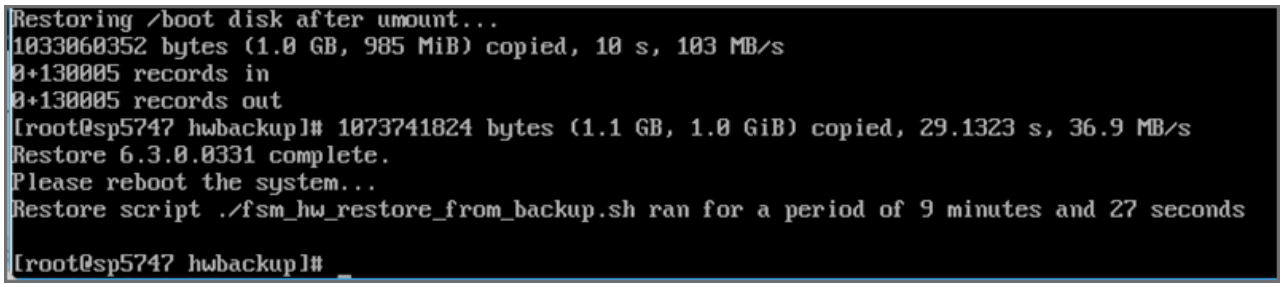

6. Once the restore is complete, it will print how long the restore took and will ask you to reboot the system. Run the command to reboot your system:

   `reboot`

   The system should now come up with your pre-upgrade version. Wait at least 15 minutes for all processes to come up.

   If you are using 3500F, 2000F, 2000G, 2200G, 3500G or 3600G as a worker node, or 500F, 500G as a collector node, then the restore of CMDB and SVN is skipped.

   The restore logs are stored in this location
   `/opt/hwbackup/fsm-hw-restore-<date>-<hour-minute>.log`

   If the restore fails for any reason or if processes do not come up after reboot, then please contact technical support.

Upgrade Log

The 7.5.0.0590 Upgrade ansible log file is located here: `/usr/local/upgrade/logs/ansible.log`.

Errors can be found at the end of the file.

Important Notes for Upgrading from Older Versions

Upgrading from 6.5.0 Running ClickHouse Event Database

1. **This applies only if you are upgrading from 6.5.0 and using ClickHouse.** FortiSIEM 6.5.0 ran ClickHouse on a single node and used the Merge Tree engine. FortiSIEM 6.6.0 onwards runs Replicated Merge Tree engine, even if Replication is not turned on. So after upgrading to FortiSIEM 6.6.0, you will need to do the following steps to migrate the event data previously stored in Merge Tree to Replicated Merge Tree. Without these steps, old events in 6.5.0 will not be searchable in 6.6.0. Once you are on post 6.5.0 release, you will not need to do this procedure again.

   After upgrading your FortiSIEM from 6.5.0 to 6.6.0 or later, take the following steps.

   a. Navigate to **Admin >Settings > Database > ClickHouse Config**.
   b. Click **Test**, then click **Deploy** to enable the ClickHouse Keeper service which is new in 6.6.0.
   c. Migrate the event data in 6.5.0 to 6.6.0 by running the script `/opt/phoenix/phscripts/clickhouse/clickhouse-migrate-650.sh`.
2. **This applies only if you are upgrading from 6.5.0 and using ClickHouse.** Go to Storage > Online Settings and click **Test**, it will fail. Fortinet introduced a new disk attribute called "Mounted On" to facilitate disk addition/deletion that was not present in 6.5.0. Follow these steps to fix the problem.

   a. Go to **Admin > Setup > Storage > Online**. ClickHouse should be the selected database.
   b. For Hot tier and for every configured disk within the tier, do the following:
      i. The existing disk should have empty Mounted On.
      ii. Click + to add a disk. For the new disk, Disk Path should be empty and Mounted On set to /data-clickhouse-hot-1.
      iii. Copy the Disk Path from the existing disk into this newly disk. The new disk should have the proper Disk Path and Mounted On fields.
      iv. Delete the first disk with empty Mounted On.

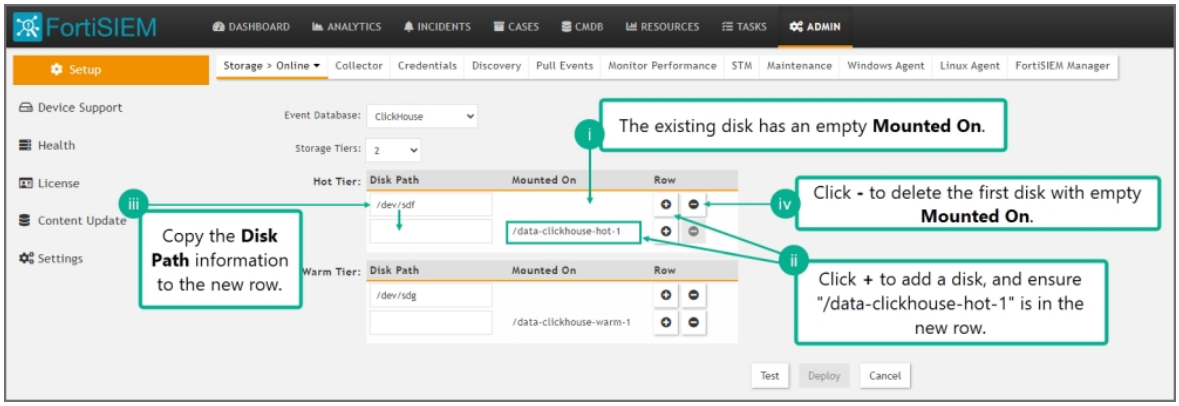

Do this for all disks you have configured in 6.5.0. After your changes, the disks should be ordered /data-clickhouse-hot-1, /data-clickhouse-hot-2, /data-clickhouse-hot-3 from top to bottom.

c. Repeat the same steps for the Warm tier (if one was configured in 6.5.0), except that the Mounted On fields should be /data-clickhouse-warm-1, /data-clickhouse-warm-2, /data-clickhouse-warm-3 from top to bottom.

d. When done, click **Test**, then click **Deploy**.

General Upgrade Notes

These notes apply to all upgrades in general.

1. Remember to remove the browser cache after logging on to the 7.5.0 GUI and before doing any operations.
2. Make sure to follow the listed upgrade order.
   a. Upgrade the Supervisor first. It must be upgraded prior to upgrading any Workers or Collectors.
   b. Upgrade all existing Workers next, after upgrading the Supervisor. The Supervisor and Workers must be on the same version.
   c. Older Collectors will work with the upgraded Supervisor and Workers. You can decide to upgrade Collectors to get the full feature set in 7.5.0 after you have upgraded all Workers.

Reference

Steps for Expanding /opt Disk

1. Go to the Hypervisor and increase the size of `/opt` disk or the size of `/svn` disk
2. # ssh into the supervisor as `root`
3. `# lsblk`

```
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
...
sdb           8:16   0  100G  0 disk            << old size
├─sdb1        8:17   0 22.4G  0 part [SWAP]
└─sdb2        8:18   0 68.9G  0 part /opt
      ...
```

4. `# yum -y install cloud-utils-growpart gdisk`
5. `# growpart /dev/sdb 2`
   `CHANGED: partition=2 start=50782208 old: size=144529408 end=195311616 new: size=473505759 end=524287967`
6. `# lsblk`

```
Changed the size to 250GB for example:
#lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
...
sdb           8:16   0  250G  0 disk            <<< NOTE the new size for the disk in /opt
├─sdb1        8:17   0 22.4G  0 part [SWAP]
└─sdb2        8:18   0 68.9G  0 part /opt
...
```

7. `# xfs_growfs /dev/sdb2`

```
meta-data=/dev/sdb2              isize=512    agcount=4, agsize=4516544 blks
         =                       sectsz=512   attr=2, projid32bit=1
         =                       crc=1        finobt=1, sparse=1, rmapbt=0
         =                       reflink=1
data     =                       bsize=4096   blocks=18066176, imaxpct=25
         =                       sunit=0      swidth=0 blks
naming   =version 2              bsize=4096   ascii-ci=0, ftype=1
```

```
log      =internal log           bsize=4096   blocks=8821, version=2
         =                       sectsz=512   sunit=0 blks, lazy-count=1
realtime =none                   extsz=4096   blocks=0, rtextents=0
data blocks changed from 18066176 to 59188219
```

8. `# df -hz`

```
Filesystem           Size  Used Avail Use% Mounted on
...
/dev/sdb2            226G  6.1G  220G   3% /  << NOTE the new disk size
```

Post Upgrade Health Check get-fsm-health.py --local Example Output

Here is an example of a successful output when running `get-fsm-health.py --local`.

```
                        Health Check
============================================================
Wed Jul 07 17:35:26 PDT 2021
--------------------
Fetching Information from Local.
- Host Info ..................................... succeeded.
- FortiSIEM Version ............................. succeeded.
- FortiSIEM License Info ........................ succeeded.
- Configuration ................................. succeeded.
- CMDB Info ..................................... succeeded.
- Largest CMDB Tables ........................... succeeded.
- EPS Info ...................................... succeeded.
- Worker Upload Event Queue Info ................ succeeded.
- Inline Report Queue ........................... succeeded.
- Active Queries ................................ succeeded.
- Load Average .................................. succeeded.
- CPU Usage Details ............................. succeeded.
- Top 5 Processes by CPU ........................ succeeded.
- Memory Usage .................................. succeeded.
- Swap Usage .................................... succeeded.
- Top 5 Processes by Resident Memory ............ succeeded.
- Disk Usage .................................... succeeded.
- IOStat ........................................ succeeded.
- Top 5 Processes by IO ......................... succeeded.
- NFSIOStat ..................................... succeeded.
- NFS Disk Operations Time (second) ............. succeeded.
- Top 10 Slow EventDB Queries ( > 1 min) Today ..... succeeded.
- Top 5 Rule with Large Memory Today ............ succeeded.
- FortiSIEM Process Uptime Less Than 1 day ..... succeeded.
- Top 5 log files in /var/log  ................. succeeded.
- FortiSIEM Shared Store Status ................. succeeded.
- App Server Exceptions Today ................... succeeded.
```

```
- Backend Errors Today ............................ succeeded.
- Backend Segfaults Today ......................... succeeded.
- Patched files ................................... succeeded.
- Outstanding Discovery Jobs ...................... succeeded.
- FortiSIEM Log File Size ......................... succeeded.
- FortiSIEM Fall Behind Jobs ...................... succeeded.
- FortiSIEM Jobs Distribution ..................... succeeded.


------------------------------------------------------------
                      Data Collection
============================================================
All data was collected.


------------------------------------------------------------
                     Health Assessment
============================================================
Overall health: **Critical**

CPU Utilization: Normal
  - 15 min Load average: 1.05
  - System CPU: 4.5%
Memory Utilization: Normal
  - Memory utilization: 48%
  - Swap space utilization: 0.0%
  - Swap in rate: 0B/s
  - Swap out rate: 0B/s
I/O Utilization: Normal
  - CPU Idle Wait: 0.0%
  - Local disk IO util: 0.2%
  - NFS latency (/data): 2.2ms
Disk Utilization: Normal
  - Disk Utilization: 33%
Event Ingestion: Normal
  - Worker event upload queue: 1
  - Shared store status: Nobody is falling behind
Event Analysis: Normal
  - Inline report queue: 4
  - Active query queue: 0
System Errors: Normal
  - Process down. See details.
  - App server errors: 0
  - Backend error: 2
Performance Monitoring: **Critical**
  - 1250 jobs are falling behind. (Super)  *****


------------------------------------------------------------
                           Details
============================================================
```

```
#######################  Host Info  ########################
NodeType   Host Name                     IP Address

Super      sp156                         172.30.56.156



###################  FortiSIEM Version  ####################
NodeType   Version         Commit Hash          Built On

Super      6.3.0.0331      6e29f46b382          Thu Jul 01 15:58:02 PDT 2021



#################  FortiSIEM License Info  #################
License Information:
Attribute                               Value                                    Expiration Date

Serial Number                           FSMTEST8888888888
Hardware ID                             88888888-8888-8888-8888-888888888888
License Type                            Service Provider
Devices                                 1000                                     Dec 31, 2021
Endpoint Devices                        1000                                     Dec 31, 2021
Additional EPS                          10000                                    Dec 31, 2021
Total EPS                               22000                                    Dec 31, 2021
Agents                                  2000                                     Dec 31, 2021
UEBA Telemetry License                  1000                                     Dec 31, 2021
IOC Service                             Valid                                    Dec 31, 2021
Maintenance and Support                 Valid                                    Dec 31, 2021

……
```

www.fortinet.com

Copyright© 2026 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's Chief Legal Officer, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.