Upgrade Guide

FortiSIEM 7.3.0

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

07/29/2025
FortiSIEM 7.3.0 Upgrade Guide

TABLE OF CONTENTS

Change Log

Date	Change Description
12/16/2024	Initial version of the 7.3.0 Upgrade Guide.
03/04/2025	Initial version of the 7.3.1 Upgrade Guide.
03/21/2025	Initial version of the 7.3.2 Upgrade Guide.
07/15/2025	Initial version of the 7.3.3 Upgrade Guide.
07/17/2025	Added Target Upgrade Version section.
07/24/2025	Initial version of the 7.3.4 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 7.3.0 release.

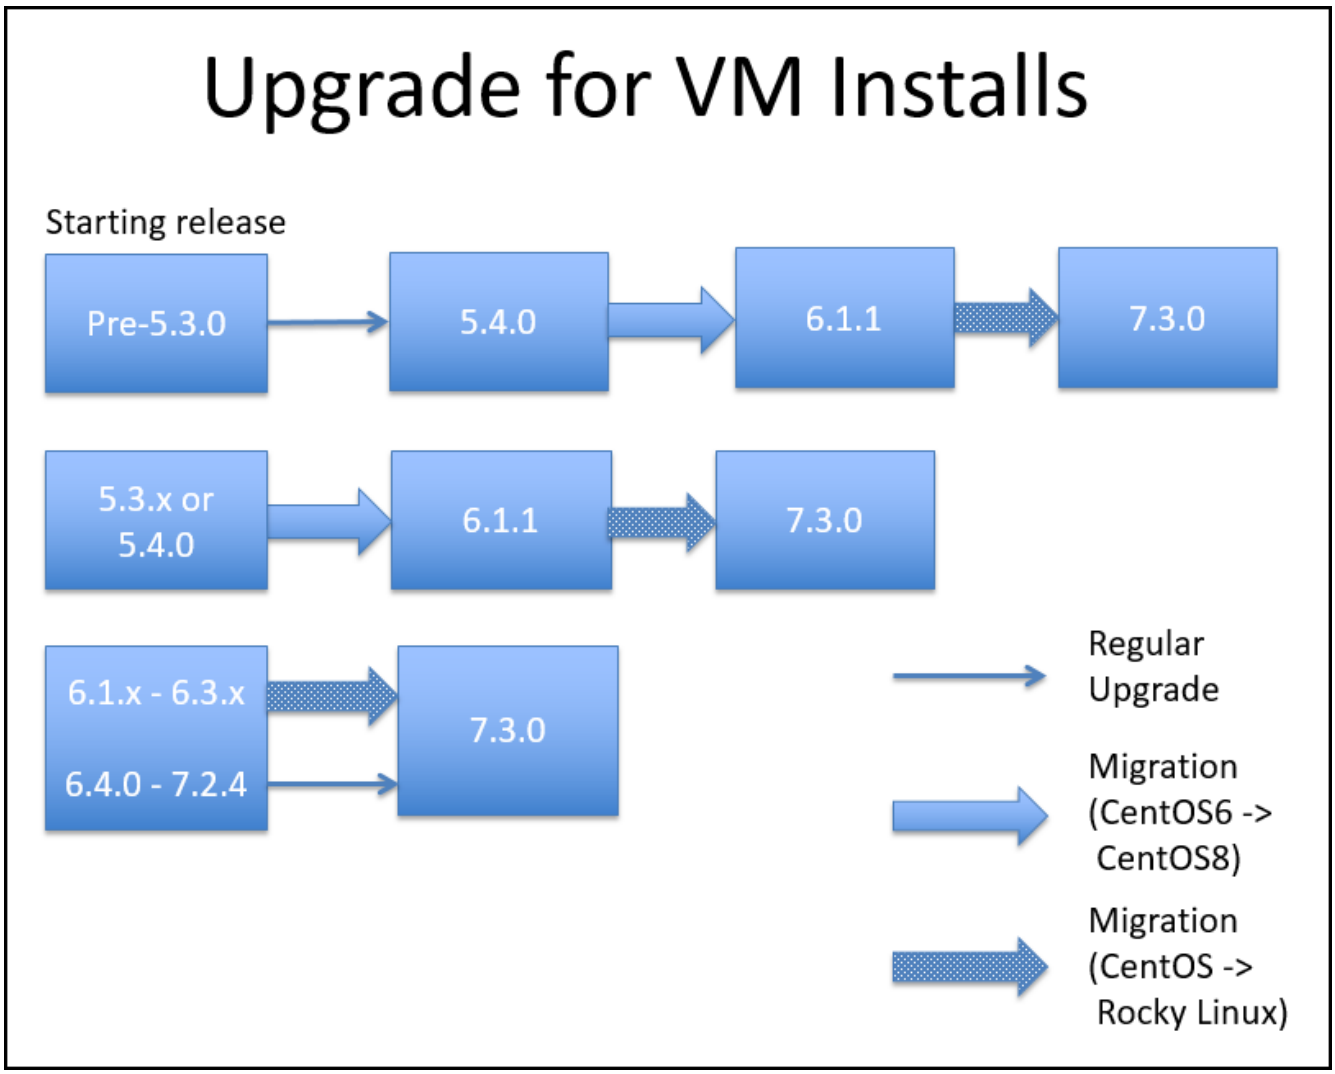

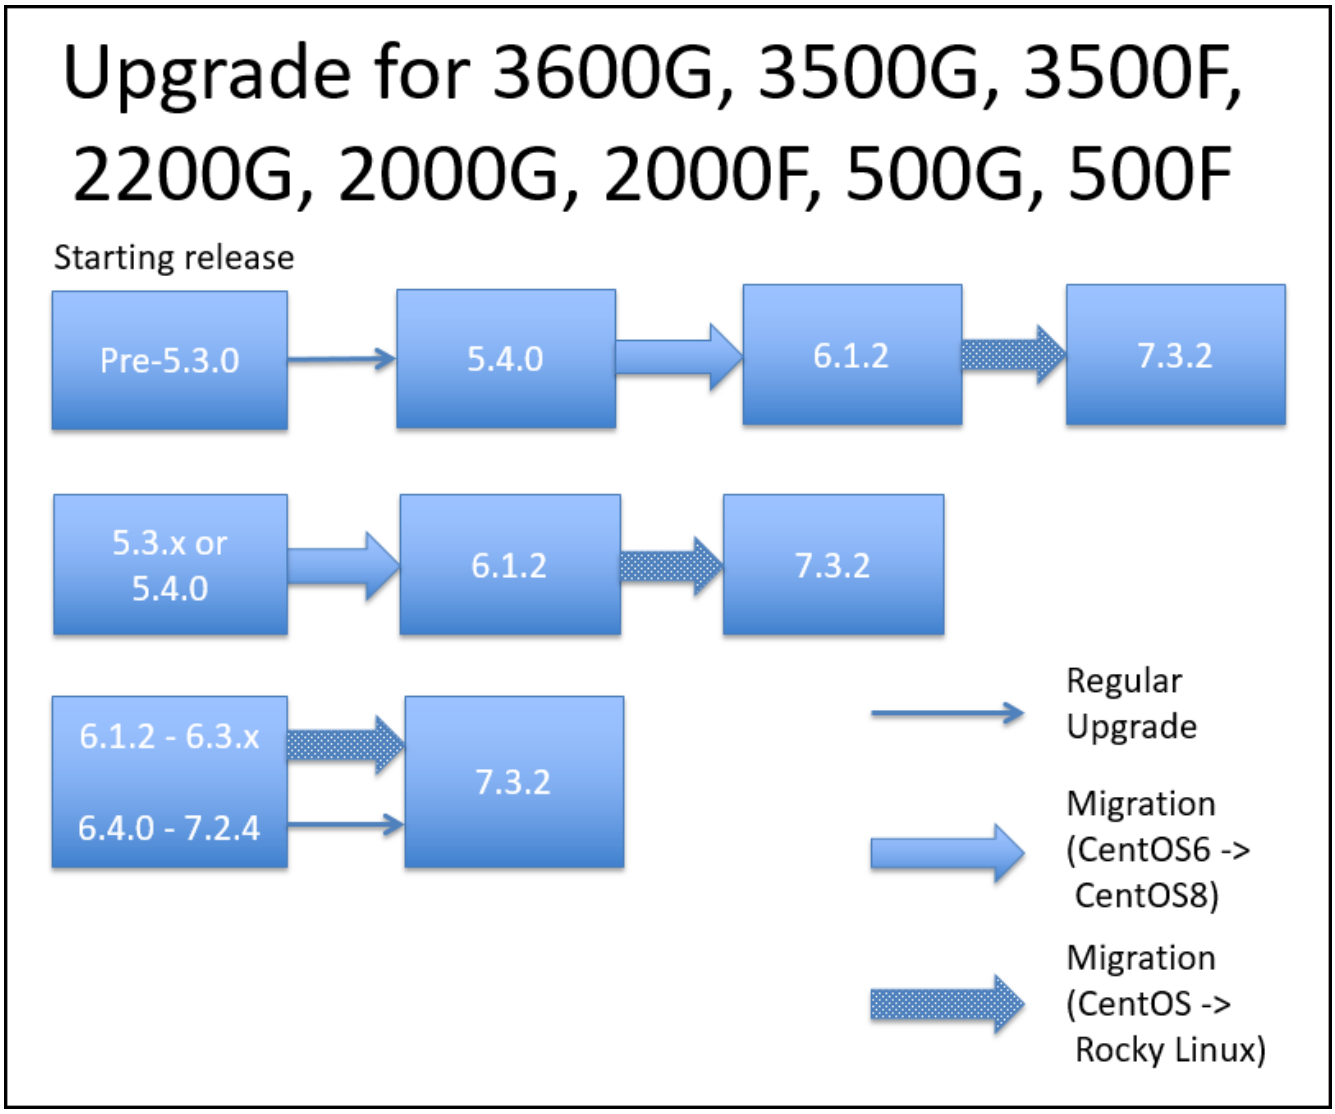
Upgrade for 3600G, 3500G, 3500F, 2200G, 2000G, 2000F, 500G, 500F
Starting release
Pre-5.3.0
5.4.0
6.1.2
7.3.2
5.3.x or 5.4.0
6.1.2
7.3.2
6.1.2 - 6.3.x
6.4.0 - 7.2.4
7.3.2
Regular Upgrade
Migration (CentOS6 -> CentOS8)
Migration (CentOS -> Rocky Linux)

Important Notes

Pre-Upgrade Checklist

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

1. Carefully consider the known issues, if any, in the Release Notes.
2. Make sure the Supervisor processes are all up.
3. Make sure you can login to the FortiSIEM GUI and successfully discover your devices.
4. Take a snapshot of the running FortiSIEM instance.
5. If you running FortiSIEM versions 6.2.0 or earlier and using Elasticsearch, then navigate to **ADMIN > Setup > Storage > Online >** and perform a **Test** and **Save** after the upgrade. This step is not required while upgrading from versions 6.2.1 or later.
6. From version 6.4.0 onwards, FortiSIEM runs on Rocky Linux. If upgrading from a release prior to 6.4.0, then FortiSIEM will automatically migrate the operating system from CentOS to Rocky Linux during the upgrade process. If upgrading from a FortiSIEM 6.4.0 release or later, then FortiSIEM will already be running Rocky Linux, so no additional migration is needed.
7. Make sure the FortiSIEM license is not expired.
8. Make sure the Supervisor, Workers and Collectors 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.

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.

   To upgrade 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.

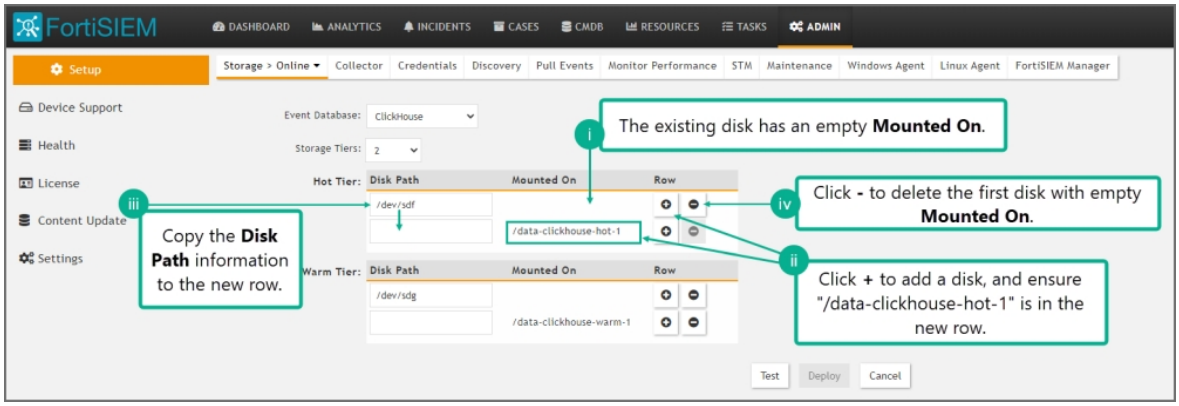

Do this for all disks you have configured in 6.5.0. After your changes, the disks should be ordered /data-clickhouse-hot-1, /data-clickhouse-hot-2, /data-clickhouse-hot-3 from top to bottom.

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**.

6.2.0 to 7.3.0 Upgrade Notes

This note applies only if you are upgrading from 6.2.0.

Before upgrading Collectors to 7.3.0, you will need to copy the `phcollectorimageinstaller.py` file from the Supervisor to the Collectors. See steps 1-3 in Upgrade Collectors.

6.1.x to 7.3.0 Upgrade Notes

These notes apply only if you are upgrading from 6.1.x to 7.3.0.

1. The 7.3.0 upgrade will attempt to migrate existing SVN files (stored in `/svn`) from the old svn format to the new svn-lite format. During this process, it will first export `/svn` to `/opt` and then import them back to `/svn` in the new svn-lite format. If your `/svn` uses a large amount of disk space, and `/opt` does not have enough disk space left, then migration will fail. Fortinet recommends doing the following steps before upgrading:
   - Check `/svn` usage
   - Check if there is enough disk space left in `/opt` to accommodate `/svn`
   - Expand `/opt` by the size of `/svn`
   - Begin upgrade

   See Steps for Expanding /opt Disk for more information.

2. If you are using AWS Elasticsearch, then after upgrading to 7.3.0, take the following steps:
   a. Go to **ADMIN > Setup > Storage> Online**.
   b. Select "ES-type" and re-enter the credential.

General Upgrade Notes

These notes apply to all upgrades in general.

1. For the Supervisor and Worker, do not use the upgrade menu item in configFSM.sh to upgrade from 6.2.0 to 7.3.0. This is deprecated, so it will not work. Use the new method as instructed in this guide (See **Upgrade Supervisor** for the appropriate deployment under Upgrade Single Node Deployment or Upgrade Cluster Deployment).
2. In 6.1.x releases, new 5.x collectors could not register to the Supervisor. This restriction has been removed in 6.2.x so long as the Supervisor is running in non-FIPS mode. However, 5.x collectors are not recommended since CentOS 6 has been declared End of Life.
3. Remember to remove the browser cache after logging on to the 7.3.0 GUI and before doing any operations.
4. 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 7.3.0 after you have upgraded all Workers.
5. If you are running FortiSIEM versions 6.2.0 or earlier and using Elasticsearch, then you must redo your Elasticsearch configuration after your upgrade by taking the following steps:
   a. Navigate to **ADMIN > Setup > Storage > Online**.
   b. Redo your configuration.
   c. Click **Test** to verify.
   d. Click **Save**.

   Note: These steps (5a-d) are not required while upgrading from versions 6.2.1 or later.
6. 5.x Collector will not work with FortiSIEM 6.7.2 or later. This step is taken for improved security. Follow these steps to make the 5.x Collectors operational after upgrade.
   a. Upgrade the Supervisor to the latest version: 7.0.0 or higher.
   b. Copy `phProvisionCollector.collector` from the Supervisor to all 5.x Collectors.
      i. Login to Supervisor.
      ii. Run the following command.

      ```
      scp /opt/phoenix/phscripts/bin/phProvisionCollector.collector
      root@<Collector_IP>:/opt/phoenix/bin/phProvisionCollector
      ```
   c. Update 5.x Collector password.
      i. SSH to the Collector.
      ii. Run the following command.

      ```
      phProvisionCollector --update <Organization-user-name> <Organization-user-
      password> <Supervisor-IP> <Organization-name> <Collector-name>
      ```
      iii. Make sure the Collector ID and password are present in the file `/etc/httpd/accounts/passwds` on

Supervisors and Workers.

d. Reboot the Collector.

Upgrade Pre-5.3.0 Deployment

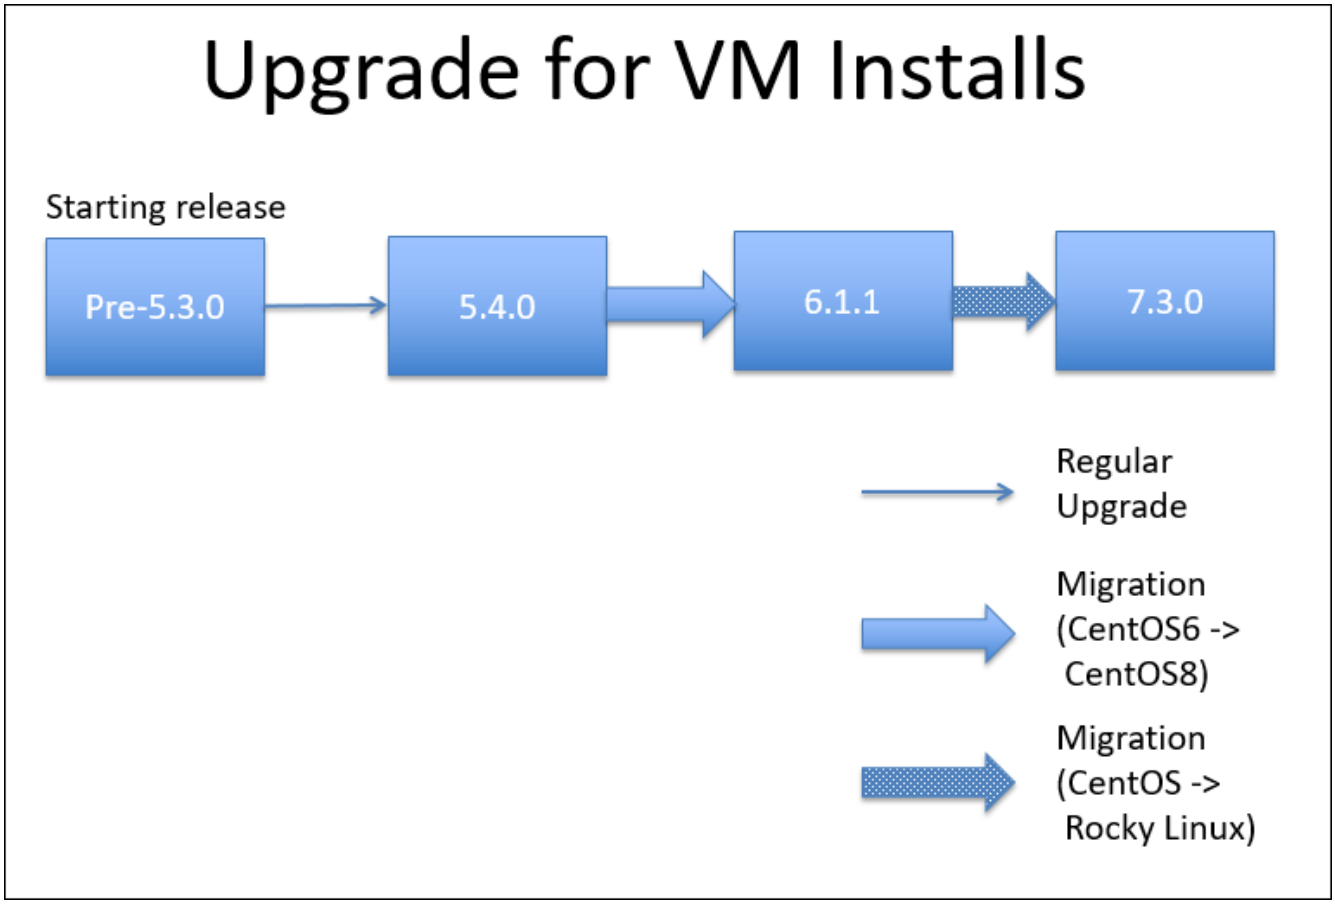

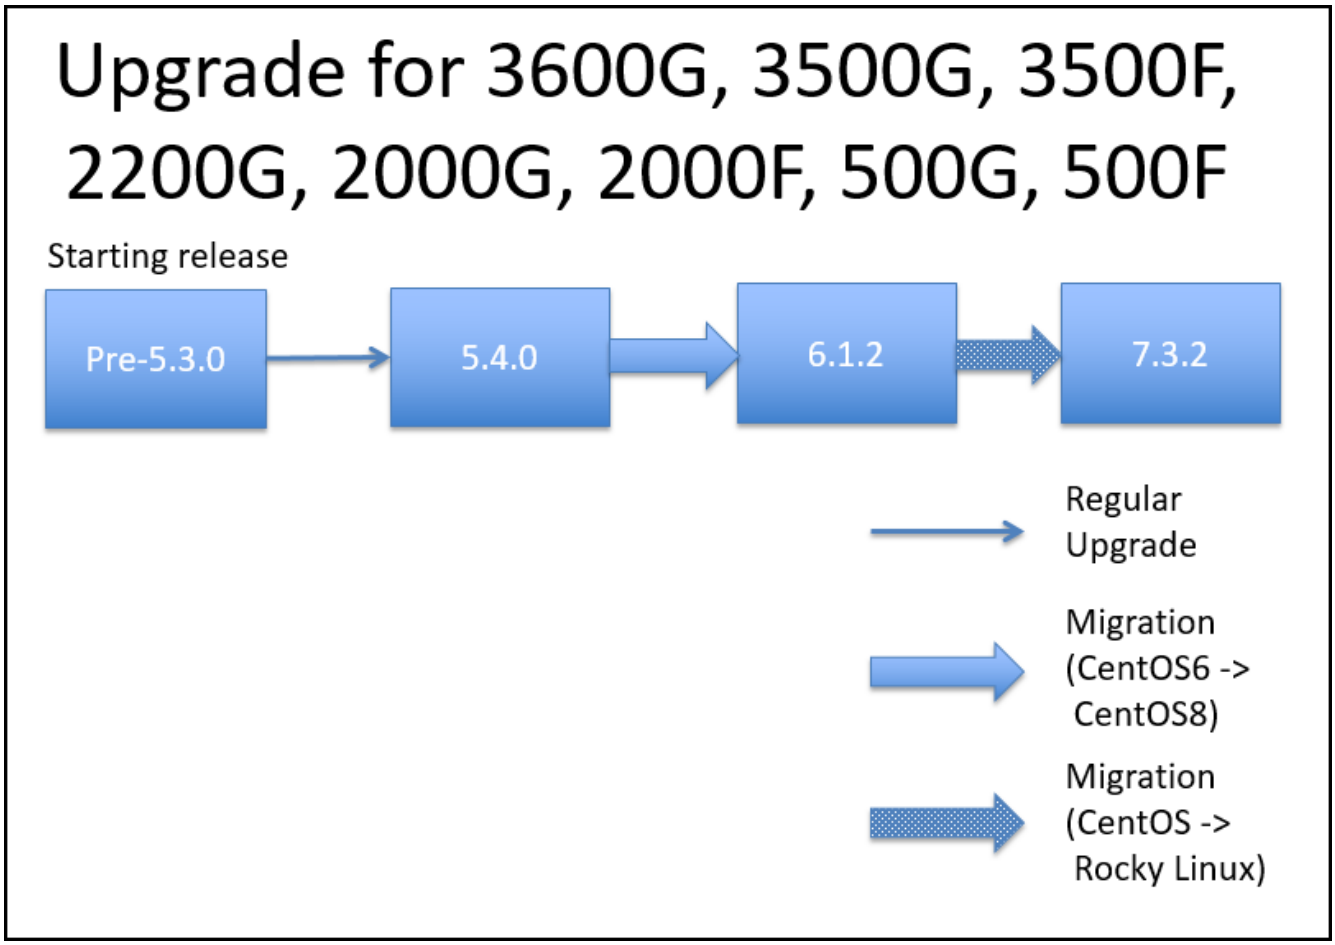

1. Upgrade to 5.4.0 by using the 5.4.0 Upgrade Guide: Single Node Deployment / Cluster Deployment.
2. Perform a health check to make sure the system has upgraded to 5.4.0 successfully.
3. If you are running a Software Virtual Appliance, you must migrate to 6.1.1. Since the base OS changed from CentOS 6 to CentOS 8, the steps are platform specific. Use the appropriate 6.1.1 guide and follow the migration instructions.
   - AWS Installation and Migration Guide
   - ESX Installation and Migration Guide
   - KVM Installation and Migration Guide
   - HyperV Installation and Migration Guide
   - Azure Installation and Migration Guide

   If you are running a hardware appliance (3500G, 3500F, 2000F, 500F), you must migrate to 6.1.2. Since the base OS changed from CentOS 6 to CentOS 8, the steps are platform specific. Follow the "Migrating from 5.3.x or 5.4.x to 6.1.2" instructions from the appropriate appliance specific documents listed here.

   Note: If you are upgrading from a 2000F, 3500F, or 3500G appliance, make sure to follow the instructions at Fix After Upgrading 2000F, 3500F, or 3500G From 5.3.x or 5.4.0 to 6.1.2 after migrating to 6.1.2.
   - 3500G Hardware Configuration Guide
   - 3500F Hardware Configuration Guide
   - 2000F Hardware Configuration Guide
   - 500F Hardware Configuration Guide

4. Perform a health check to make sure the system is upgraded to 6.1.1 or 6.1.2 successfully.
5. Upgrade to 6.x/7.x by following the steps in Upgrade 6.x/7.x Single Node Deployment or Upgrade 6.x/7.x Cluster Deployment.

Upgrade 5.3.x or 5.4.0 Deployment

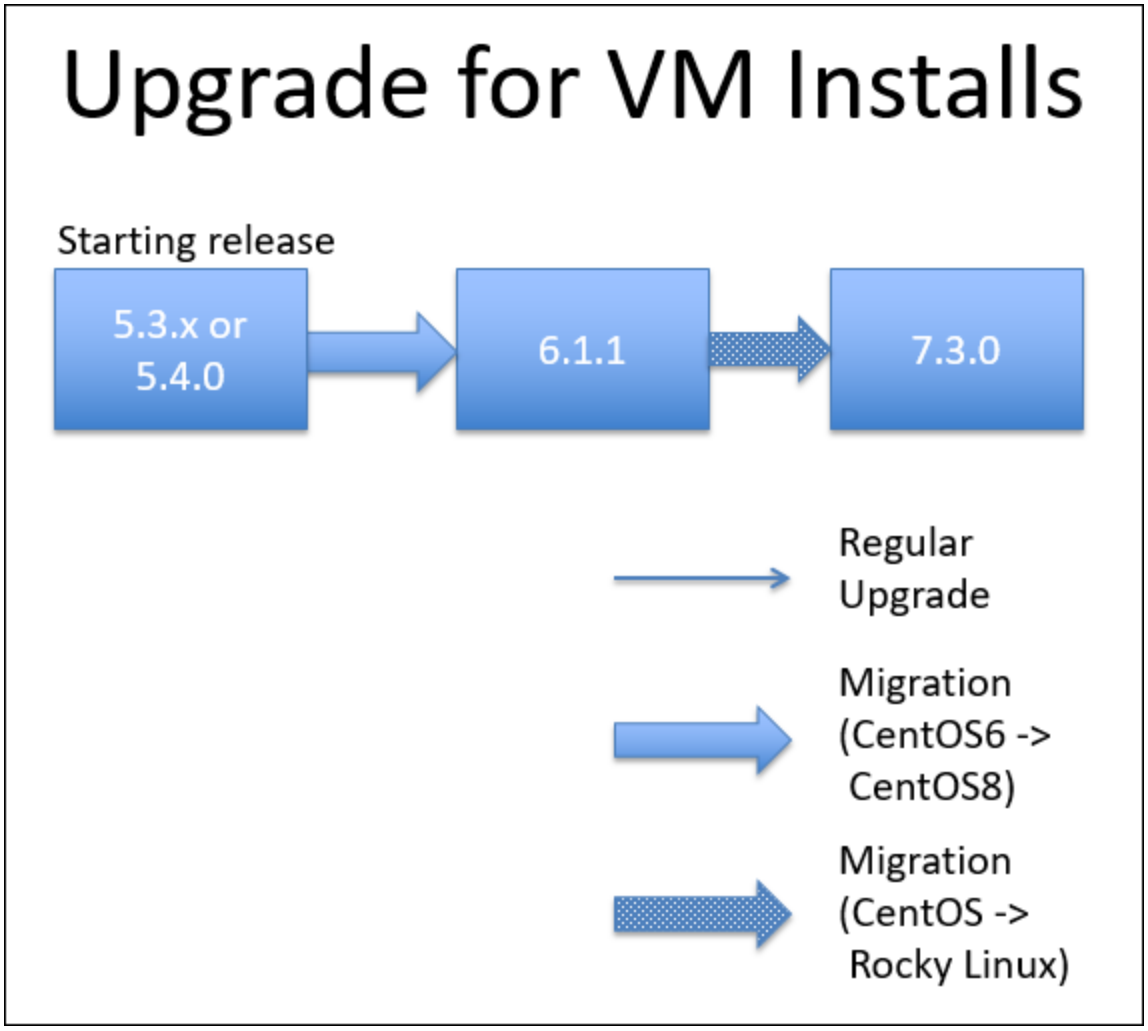

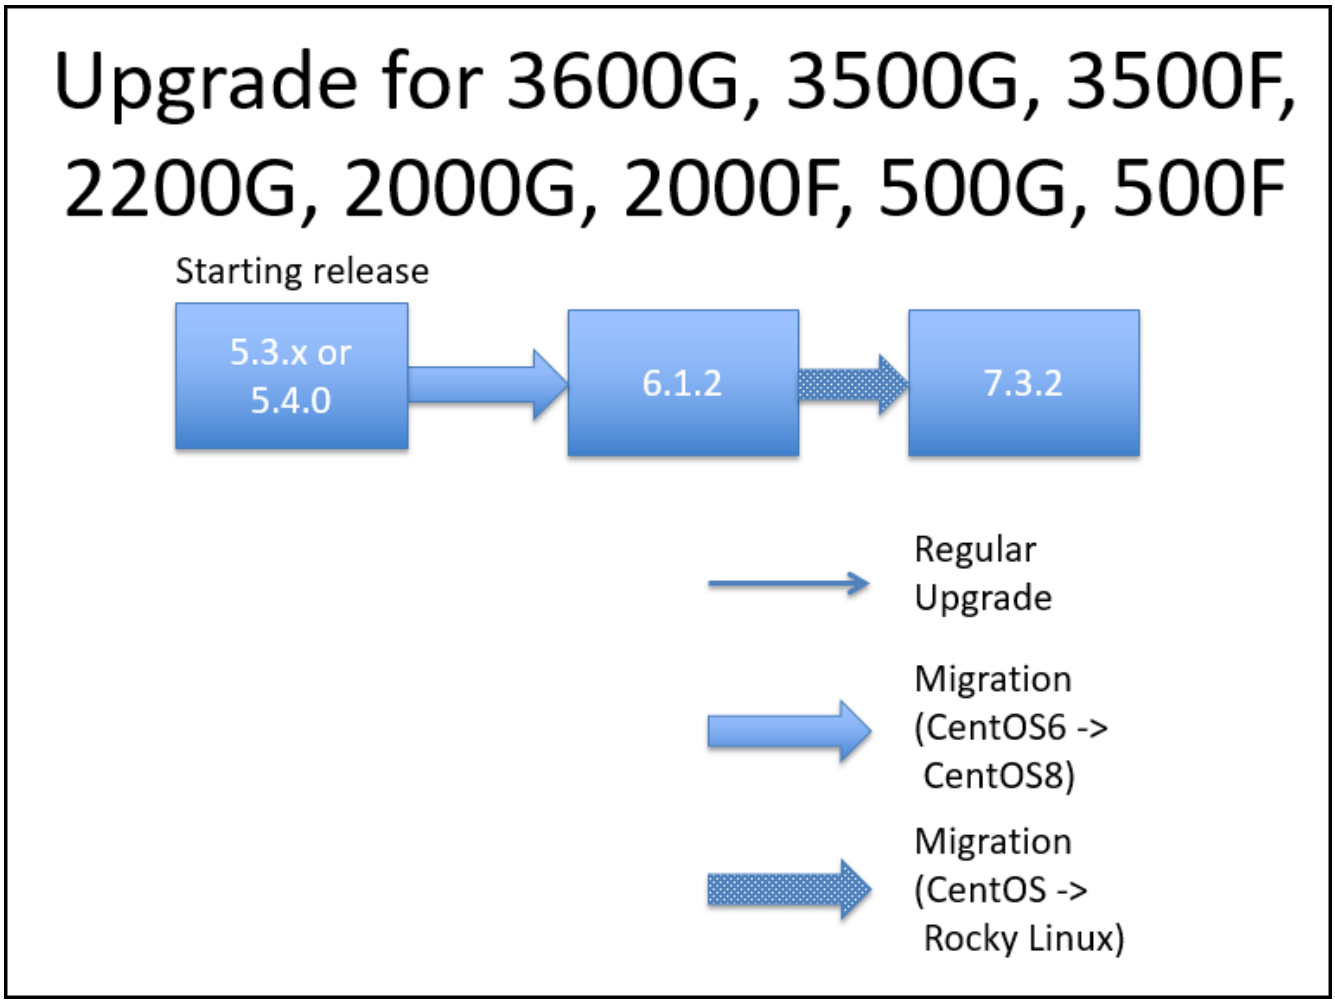

Start at step 3 from Upgrade Pre-5.3.0 Deployment, and follow the progressive steps.

Note: If you are upgrading from a 2000F, 3500F, 3500G appliance, make sure to follow the instructions at Fix After Upgrading 2000F, 3500F, or 3500G From 5.3.x or 5.4.0 to 6.1.2 after migrating to 6.1.2.

Upgrade 6.x/7.x Single Node Deployment

Prior to the 6.x/7.x Deployment 7.3.0 upgrade, ensure that the Supervisor, and all Workers are running on at least 6.x versions.

If a proxy is needed for the FortiSIEM Supervisor, Worker or Hardware appliances (FSM-2000F, 2000G, 3500F, 3500G and 3600G) to access the Internet, please refer to Upgrade via Proxy before starting.

Upgrading a single node deployment requires upgrading the Supervisor. If you have any Collectors, the Supervisor is a required upgrade before the Collectors. After completion of the upgrade, follow the appropriate steps in Post Upgrade Health Check.

- Upgrade Supervisor
- Upgrade Collectors

Upgrade Supervisor

To upgrade the Supervisor, take the following steps.

1. Login to the Supervisor via SSH.
2. Create the path `/opt/upgrade`.
   `mkdir -p /opt/upgrade`
3. Download the upgrade zip package `FSM_Upgrade_All_7.3.0_build0338.zip`, then upload it to the Supervisor node under the `/opt/upgrade/` folder.
   Example (From Linux CLI):
   `scp FSM_Upgrade_All_7.3.0_build0338.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_7.3.0_build0338.zip`
6. Go to the FSM_Upgrade_All_7.3.0_build0338 directory.
   `cd FSM_Upgrade_All_7.3.0_build0338`
   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. All processes should be up and running.
   `phstatus`
   Example output:

```
System uptime:  13:31:19 up 1 day,  2:44,  1 user,  load average: 0.95, 1.00, 1.20
Tasks: 29 total, 0 running, 29 sleeping, 0 stopped, 0 zombie
```

```
Cpu(s): 8 cores, 15.4%us, 0.5%sy, 0.0%ni, 83.6%id, 0.0%wa, 0.4%hi, 0.1%si, 0.0%st
Mem: 24468880k total, 12074704k used, 10214416k free, 5248k buffers
Swap: 26058744k total, 0k used, 26058744k free, 2931812k cached


PROCESS                 UPTIME          CPU%            VIRT_MEM        RES_MEM

phParser                23:57:06        0               2276m           695m
phQueryMaster           1-02:40:44      0               986m            99m
phRuleMaster            1-02:40:44      0               1315m           650m
phRuleWorker            1-02:40:44      0               1420m           252m
phQueryWorker           1-02:40:44      0               1450m           113m
phDataManager           1-02:40:44      0               1195m           101m
phDiscover              1-02:40:44      0               542m            59m
phReportWorker          1-02:40:44      0               1482m           193m
phReportMaster          1-02:40:44      0               694m            84m
phIpIdentityWorker      1-02:40:44      0               1044m           85m
phIpIdentityMaster      1-02:40:44      0               505m            43m
phAgentManager          1-02:40:44      0               1526m           71m
phCheckpoint            1-02:40:44      0               305m            49m
phPerfMonitor           1-02:40:44      0               820m            82m
phReportLoader          1-02:40:44      0               826m            327m
phDataPurger            1-02:40:44      0               613m            88m
phEventForwarder        1-02:40:44      0               534m            37m
phMonitor               1-02:40:49      0               1322m           629m
Apache                  1-02:43:50      0               305m            15m
Rsyslogd                1-02:43:49      0               192m            4224k
Node.js-charting        1-02:43:43      0               614m            80m
Node.js-pm2             1-02:43:41      0               681m            61m
phFortiInsightAI        1-02:43:50      0               13996m          374m
AppSvr                  1-02:43:38      14              11149m          4459m
DBSvr                   1-02:43:50      0               425m            37m
JavaQueryServer         1-02:40:49      0               10881m          1579m
phAnomaly               1-02:40:29      0               982m            61m
SVNLite                 1-02:43:50      0               9870m           450m
Redis                   1-02:43:43      0               107m            70m
```

Upgrade Collectors

To upgrade Collectors, take the following steps.

Extra Upgrade Steps from 6.2.0 to 7.3.0

From version 6.2.0 to 7.3.0, take the following steps before initiating the upgrade. Otherwise, go to Main Upgrade Steps.

1. Login to the Collector via SSH as root.
2. Copy `/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py` from the Supervisor by running the following command. (**Note**: This is copied from the 6.2.1 or 7.3.0 Supervisor.)
   `scp root@<SupervisorIP>:/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py /opt/phoenix/phscripts/bin/`
3. Change permission by running the following command.
   `chmod 755 /opt/phoenix/phscripts/bin/phcollectorimageinstaller.py`

Main Upgrade Steps

To upgrade your FortiSIEM Collectors, follow the steps in Installation Through Image Server Page.

Note: Installation through Image Server requires FortiSIEM 6.4.0 or higher.

Installation Through Image Server Page

To install through the Image Server GUI, take the following steps:

Note: Installation through Image Server requires FortiSIEM 6.4.0 or higher.

1. Navigate to Click **ADMIN > Settings > System > Image Server**.
2. Follow the instructions here.

Upgrade 6.x/7.x Cluster Deployment

Prior to the 6.x/7.x Deployment 7.3.0 upgrade, ensure that the Supervisor, and all Workers are running on at least 6.x versions.

If a proxy is needed for the FortiSIEM Supervisor, Worker or Hardware appliances (FSM-2000F, 2000G, 3500F, 3500G and 3600G) to access the Internet, please refer to Upgrade via Proxy before starting.

It is critical to review Overview prior to taking the detailed steps to upgrade your FortiSIEM cluster.

Overview

1. On the **worker nodes ONLY**, stop the backend processes on each worker node (See Detailed Steps for the actual steps to do this). Make sure not to shutdown or reboot the workers at this stage, as this prevents the workers from communicating with upgraded supervisor until they have been upgraded as well. Collectors can be up and running and buffering events.
2. Upgrade Primary Leader Supervisor.
3. After the Primary Leader Supervisor upgrade is complete, verify the Supervisor's health is good. If you have multiple Supervisors, the key is to upgrade the Primary Leader Supervisor first.
4. If you have multiple Supervisors, then upgrade all Primary Follower Supervisors. You can upgrade them one by one or in parallel.
5. After the Primary Leader Supervisor upgrade is complete, verify the health of all Supervisors is good.
6. Upgrade each Worker individually, then verify the Worker's health.
7. If your online storage is Elasticsearch, take the following steps:
   a. Navigate to **ADMIN > Setup > Storage > Online**.
   b. Click **Test** to verify the space.
   c. Click **Save** to save.
8. Upgrade each Collector individually.

Notes:

- Step 1 prevents the accumulation of Report files when the Supervisor is not available during its upgrade. If these steps are not followed, the Supervisor may not come up after the upgrade because of excessive unprocessed report file accumulation.
- Both the Supervisor and Workers must be on the same FortiSIEM version, otherwise various software modules may not work properly. However, Collectors can be in an older version, one version older to be exact. These Collectors will work, however they may not have the latest discovery and performance monitoring features offered in the latest Supervisor/Worker versions. FortiSIEM recommends that you upgrade the Collectors as soon as possible. If you have Collectors in your deployment, make sure you have configured an image server to use as a repository for them.

Detailed Steps

Take the following steps to upgrade your FortiSIEM cluster.

1. On the **Worker nodes ONLY**, stop the backend processes by running the following commands for each worker node:
   a. `# systemctl stop phxctl`
      `# phtools --stop all`
   b. Do not shutdown or reboot the workers at this stage, as this prevents the workers from communicating with the upgraded supervisor until they have been upgraded as well.
      Collectors can be up and running and buffering events. After the backend processes for all worker nodes have been stopped, proceed to the next step.
2. Upgrade the Primary Leader Supervisor using the steps in Upgrade Supervisor. Make sure the Supervisor is running the version you have upgraded to and that all processes are up and running.
   `# phshowVersion.sh`
   `# phstatus`
3. If you have Primary Follower Supervisors, then upgrade them now. The steps are the same as Upgrade Supervisor. You can upgrade them one by one or in parallel.
4. If you are running Elasticsearch, and upgrading from 6.1.x to 7.3.0, then take the following steps, else skip this step and proceed to Step 5.
   a. Navigate to **ADMIN > Storage > Online > Elasticsearch**.
   b. Verify that the Elasticsearch cluster has enough nodes (each type node >= replica + 1).
   c. Go to **ADMIN > Setup > Storage > Online**.
   d. Select "ES-type" and re-enter the credential of the Elasticsearch cluster.
   e. Click **Test and Save**. This important step pushes the latest event attribute definitions to Elasticsearch.
5. Upgrade each Worker one by one, using the procedure in Upgrade Workers.
6. Login to the Supervisor and go to **ADMIN > Health > Cloud Health** to ensure that all Workers and Supervisor have been upgraded to the intended version.
   Note: The Supervisor and Workers must be on the same version.
7. Upgrade Collectors using the steps in Upgrade Collectors.
8. After completion of the upgrade, follow the appropriate steps in Post Upgrade Health Check.

Upgrade Supervisor

To upgrade the Supervisor, take the following steps.

1. Make sure Worker processes are stopped on all workers by running the following commands on each worker. Collectors can remain up and running.**Note**: Make sure the commands from step 1 here are only run on Worker nodes, **NOT** on the Supervisor.
   `# systemctl stop phxctl`
   `# phtools --stop all`
2. Login to the Supervisor via SSH.
3. Create the path `/opt/upgrade`.
   `mkdir -p /opt/upgrade`
4. Download the upgrade zip package `FSM_Upgrade_All_7.3.0_build0338.zip`, then upload it to the Supervisor node under the `/opt/upgrade/` folder.
   Example (From Linux CLI):
   `scp FSM_Upgrade_All_7.3.0_build0338.zip root@10.10.10.15:/opt/upgrade/`

5. Go to `/opt/upgrade`.
   `cd /opt/upgrade`
6. 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_7.3.0_build0338.zip`
7. Go to the FSM_Upgrade_All_7.3.0_build0338 directory.
   `cd FSM_Upgrade_All_7.3.0_build0338`
   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`
8. Start the upgrade process by entering the following.
   `sh upgrade.sh`
9. After the process is completed, perform a basic health check. All processes should be up and running.
   `phstatus`
   Example output:

```
System uptime:  13:31:19 up 1 day,  2:44,  1 user,  load average: 0.95, 1.00, 1.20
Tasks: 29 total, 0 running, 29 sleeping, 0 stopped, 0 zombie
Cpu(s): 8 cores, 15.4%us, 0.5%sy, 0.0%ni, 83.6%id, 0.0%wa, 0.4%hi, 0.1%si, 0.0%st
Mem: 24468880k total, 12074704k used, 10214416k free, 5248k buffers
Swap: 26058744k total, 0k used, 26058744k free, 2931812k cached


PROCESS                 UPTIME          CPU%            VIRT_MEM        RES_MEM

phParser                23:57:06        0               2276m           695m
phQueryMaster           1-02:40:44      0               986m            99m
phRuleMaster            1-02:40:44      0               1315m           650m
phRuleWorker            1-02:40:44      0               1420m           252m
phQueryWorker           1-02:40:44      0               1450m           113m
phDataManager           1-02:40:44      0               1195m           101m
phDiscover              1-02:40:44      0               542m            59m
phReportWorker          1-02:40:44      0               1482m           193m
phReportMaster          1-02:40:44      0               694m            84m
phIpIdentityWorker      1-02:40:44      0               1044m           85m
phIpIdentityMaster      1-02:40:44      0               505m            43m
phAgentManager          1-02:40:44      0               1526m           71m
phCheckpoint            1-02:40:44      0               305m            49m
phPerfMonitor           1-02:40:44      0               820m            82m
phReportLoader          1-02:40:44      0               826m            327m
phDataPurger            1-02:40:44      0               613m            88m
phEventForwarder        1-02:40:44      0               534m            37m
phMonitor               1-02:40:49      0               1322m           629m
Apache                  1-02:43:50      0               305m            15m
Rsyslogd                1-02:43:49      0               192m            4224k
Node.js-charting        1-02:43:43      0               614m            80m
Node.js-pm2             1-02:43:41      0               681m            61m
phFortiInsightAI        1-02:43:50      0               13996m          374m
AppSvr                  1-02:43:38      14              11149m          4459m
DBSvr                   1-02:43:50      0               425m            37m
JavaQueryServer         1-02:40:49      0               10881m          1579m
```

```
phAnomaly                1-02:40:29     0              982m           61m
SVNLite                  1-02:43:50     0              9870m          450m
Redis                    1-02:43:43     0              107m           70m
```

Upgrade Workers

To upgrade Workers, take the following steps for each Worker.

1. Login to a worker via SSH.
2. Create the path `/opt/upgrade`.
   `mkdir -p /opt/upgrade`
3. Download the upgrade zip package `FSM_Upgrade_All_7.3.0_build0338.zip` to `/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_7.3.0_build0338.zip`
6. Go to the FSM_Upgrade_All_7.3.0_build0338 directory.
   `cd FSM_Upgrade_All_7.3.0_build0338`
   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. All processes should be up and running.
9. After all Workers are upgraded, perform this extra set of steps if you were running FortiSIEM versions 6.2.0 or earlier and using Elasticsearch after the upgrade.
   a. Navigate to **ADMIN > Setup > Storage > Online**.
   b. Redo your configuration.
   c. Perform a **Test** to verify it is working.
   d. Click **Save**.

Note: These steps (9a-d) are not required while upgrading from versions 6.2.1 or later.

Upgrade Collectors

Extra Upgrade Steps from 6.2.0 to 7.3.0

From version 6.2.0 to 7.3.0, take the following steps before initiating the upgrade. Otherwise, go to Main Upgrade Steps.

1. Login to the Collector via SSH as root.
2. Copy `/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py` from the Supervisor by running the following command. (**Note**: This is copied from the 6.2.1 or 7.3.0 Supervisor.)
   ```
   scp root@<SupervisorIP>:/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py
   /opt/phoenix/phscripts/bin/
   ```
3. Change permission by running the following command.
   ```
   chmod 755 /opt/phoenix/phscripts/bin/phcollectorimageinstaller.py
   ```

Main Upgrade Steps

- Installation Through Image Server Page
  Note: Installation through Image Server requires FortiSIEM 6.4.0 or higher.

Installation Through Image Server Page

To install through the Image Server GUI, take the following steps:

Note: Installation through Image Server requires FortiSIEM 6.4.0 or higher.

1. Navigate to Click **ADMIN > Settings > System > Image Server**.
2. Follow the instructions here.

Upgrading with Disaster Recovery Enabled

Health Check before Upgrade for Disaster Recovery

Prior to upgrading a Disaster Recovery (DR) environment, take the following steps to verify the health of your DR environment.

1. On the Primary, go to `/opt/phoenix/cache/replication`, and run the following commands to confirm the Primary is functioning correctly.

   a.
   ```
   # ls -al
   total 12
   drwxrwxr-x  2 admin admin   46 Dec 12 17:02 .
   drwxr-xr-x 24 admin admin 4096 Dec 12 17:04 ..
   -rw-------  1 admin admin 1317 Dec 12 17:02 complete_status.xml
   -rw-rw-r--  1 admin admin    8 Dec 11 16:27 .role
   ```

   b.
   ```
   # xmllint --format complete_status.xml
   <status><phoenixServer ip="192.0.2.0" role="Secondary"><cmdb><replication_delay_
   kb>0</replication_delay_kb><replication_delay_seconds>0</replication_delay_
   seconds><last_replication_time>1702429352080</last_replication_time><replication_
   paused>false</replication_paused></cmdb><svnlite><primary_size_kb>0</primary_size_
   kb><secondary_size_kb>0</secondary_size_kb><replication_delay_kb>0</replication_
   delay_kb><replication_delay_seconds>0</replication_delay_seconds><replication_
   progress>100</replication_progress><last_replication_time>1702427552</last_
   replication_time></svnlite><profiledb><primary_size_kb>472</primary_size_
   kb><secondary_size_kb>472</secondary_size_kb><replication_delay_kb>0</replication_
   delay_kb><replication_delay_seconds>0</replication_delay_seconds><replication_
   progress>100</replication_progress><last_replication_time>1702427564</last_
   replication_time></profiledb><eventdb><primary_size_kb>-1</primary_size_
   kb><secondary_size_kb>-1</secondary_size_kb><replication_delay_kb>0</replication_
   delay_kb><replication_progress>0</replication_progress><last_replication_
   time>0</last_replication_time></eventdb><elastic><replication_delay_
   ops>0</replication_delay_ops><replication_progress>0</replication_progress><last_
   replication_time>0</last_replication_time></elastic></phoenixServer></status>
   ```

   c.
   ```
   cat .role
   Primary
   ```

2. On the Secondary, go to `/opt/phoenix/cache/replication`, and run the following commands to confirm the Secondary is functioning correctly.

   a.
   ```
   # ls -la
   total 20
   drwxrwxr-x  2 admin admin   91 Dec 12 17:07 .
   drwxr-xr-x 24 admin admin 4096 Dec 12 16:32 ..
   -rw-rw-r--  1 admin admin  111 Dec 12 17:07 cmdbstatus
   -rw-------  1 admin admin 1317 Dec 12 17:07 complete_status.xml
   ```

```
-rw-rw-r--  1 admin admin   11 Dec 12 17:02 last_finish_svnlite
-rw-rw-r--  1 admin admin   10 Dec 11 16:27 .role
```

b.
```
[root@SECONDARY replication]# cat *
# cat cmdbstatus
replication_delay_bytes=0
replication_delay_seconds=0
last_replication_time=1702429617078
replication_paused=f
```

c.
```
# xmllint --format complete_status.xml
<status><phoenixServer ip="192.0.2.0" role="Secondary"><cmdb><replication_delay_
kb>0</replication_delay_kb><replication_delay_seconds>0</replication_delay_
seconds><last_replication_time>1702429617078</last_replication_time><replication_
paused>false</replication_paused></cmdb><svnlite><primary_size_kb>0</primary_size_
kb><secondary_size_kb>0</secondary_size_kb><replication_delay_kb>0</replication_
delay_kb><replication_delay_seconds>0</replication_delay_seconds><replication_
progress>100</replication_progress><last_replication_time>1702429354</last_
replication_time></svnlite><profiledb><primary_size_kb>472</primary_size_
kb><secondary_size_kb>472</secondary_size_kb><replication_delay_kb>0</replication_
delay_kb><replication_delay_seconds>0</replication_delay_seconds><replication_
progress>100</replication_progress><last_replication_time>1702427564</last_
replication_time></profiledb><eventdb><primary_size_kb>-1</primary_size_
kb><secondary_size_kb>-1</secondary_size_kb><replication_delay_kb>0</replication_
delay_kb><replication_progress>0</replication_progress><last_replication_
time>0</last_replication_time></eventdb><elastic><replication_delay_
ops>0</replication_delay_ops><replication_progress>0</replication_progress><last_
replication_time>0</last_replication_time></elastic></phoenixServer></status>
```

d.
```
# cat last_finished_svnlite
1702429354
```

e.
```
# cat .role
Secondary
```

f. Run the following command to see in a converted time if replication status is current.

```
# date -d @<time from d>
```

3. Ensure that `walsender` is running on the Primary by running the following command:

```
# ps -ef | grep walsender
```

```
postgres 1547502    1151  0 Dec11 ?        00:00:36 postgres: walsender phoenix
192.0.2.0(33686) streaming 0/DD012E40
```

4. Ensure that `walreceiver` is running on the Secondary by running the following command:

```
# ps -ef | grep walreceiver
```

```
postgres 1550135 1550129  0 Dec11 ?        00:01:30 postgres: walreceiver streaming
0/DD079E30
```

5. Check that the Secondary contains the files `backup_label.old` and `backup_manifest` following under `/cmdb/data/`:

```
-rwx------  1 postgres postgres    226 Dec 11 16:27 backup_label.old
-rwx------  1 postgres postgres 409026 Dec 11 16:27 backup_manifest
```

6. Ensure the Secondary always sees the Primary as available by checking the log located in `/opt/phoenix/log/phoenix.log`:

```
2023-12-12T10:21:54.735502-08:00 SECONDARY phMonitorSupervisor[2205175]: [PH_GENERIC_
INFO]:[eventSeverity]=PHL_INFO,[procName]=phMonitorSupervisor,
[fileName]=phMonitorProcess.cpp,[lineNumber]=12587,[phLogDetail]=Periodic ReplHealth:
Storage type: clickhouse, All Supers: 192.0.2.0,192.0.2.1
```

7. Ensure that the Primary always sees the Secondary as available by checking the log located in `/opt/phoenix/log/phoenix.log`:

   ```
   2023-12-12T17:17:32.335756-08:00 PRIMARY phMonitorSupervisor[7430]: [PH_GENERIC_INFO]:
   [eventSeverity]=PHL_INFO,[procName]=phMonitorSupervisor,[fileName]=phMonitorProcess.cpp,
   [lineNumber]=12587,[phLogDetail]=Periodic ReplHealth: Storage type: clickhouse, All
   Supers: 192.0.2.1,192.0.2.0
   ```

8. In psql, run the following query to validate the last few update periods to ensure system consistency:

   ```
   select id,to_timestamp(creation_time/1000) as creation_time, to_timestamp(last_modified_
   time/1000) as last_modified_time, owner_id,health_status_id, delay_kb, delay_seconds,
   progress, last_replication_time from ph_health_replication order by last_modified_time
   desc limit 20;
   ```

   Example Output:

   ```
     id    |     creation_time      |   last_modified_time   | owner_id | health_
   status_id | delay_kb | delay_seconds | progress | last_replication_time
   ---------+------------------------+------------------------+----------+--------
   ----------+----------+---------------+----------+----------------------
    1262661 | 2023-12-13 01:23:23+00 | 2023-12-13 01:23:23+00 |        0 |
     972651 |         0 |             0 |        0 |                       0
    1262660 | 2023-12-13 01:22:03+00 | 2023-12-13 01:22:03+00 |        0 |
     972651 |         0 |             0 |        0 |                       0
    1262659 | 2023-12-13 01:20:42+00 | 2023-12-13 01:20:42+00 |        0 |
     972651 |         0 |             0 |        0 |                       0
   ```

9. Ensure that the connection between Primary and Secondary Supervisor can connect without issues by running the following command:

   ```
   # netstat -anp | grep <secondary ip>
   ```

   Example Output:

   ```
   tcp        0      0 primary:5432      secondary:33686     ESTABLISHED 1547502/postgres:
   w
   tcp6       0      0 primary:7900      secondary:39560     ESTABLISHED
   1550991/phDataPurge
   ```

Note: Fortinet recommends running these checks as needed in addition to using **Admin > Health > Replication Health** to ensure a healthy Disaster Recovery environment.

Disaster Recovery Upgrade Steps

Refer to the appropriate High Availability and Disaster Recovery Procedures Guide - Upgrading with High Availability and Disaster Recovery topic located in the FortiSIEM Document Library before proceeding with Disaster Recovery upgrade.

To upgrade your FortiSIEMs in a Disaster Recovery environment, take the following steps.

1. Upgrade the Primary Supervisor and Workers
2. After the Primary is fully upgraded, upgrade the Secondary Supervisor and Workers.
   See Upgrade 6.x/7.x Single Node Deployment or Upgrade 6.x/7.x Cluster Deployment for more information.

After Step 1, the Secondary Supervisor database schema is already upgraded. Step 2 simply upgrades the executables in Site 2.

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.

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.

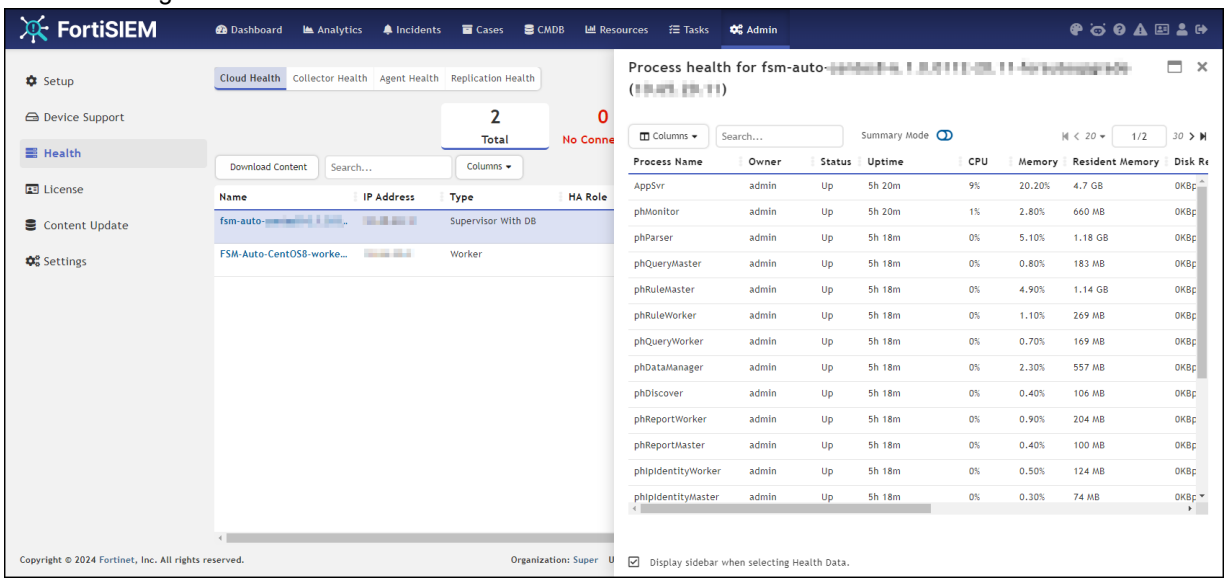

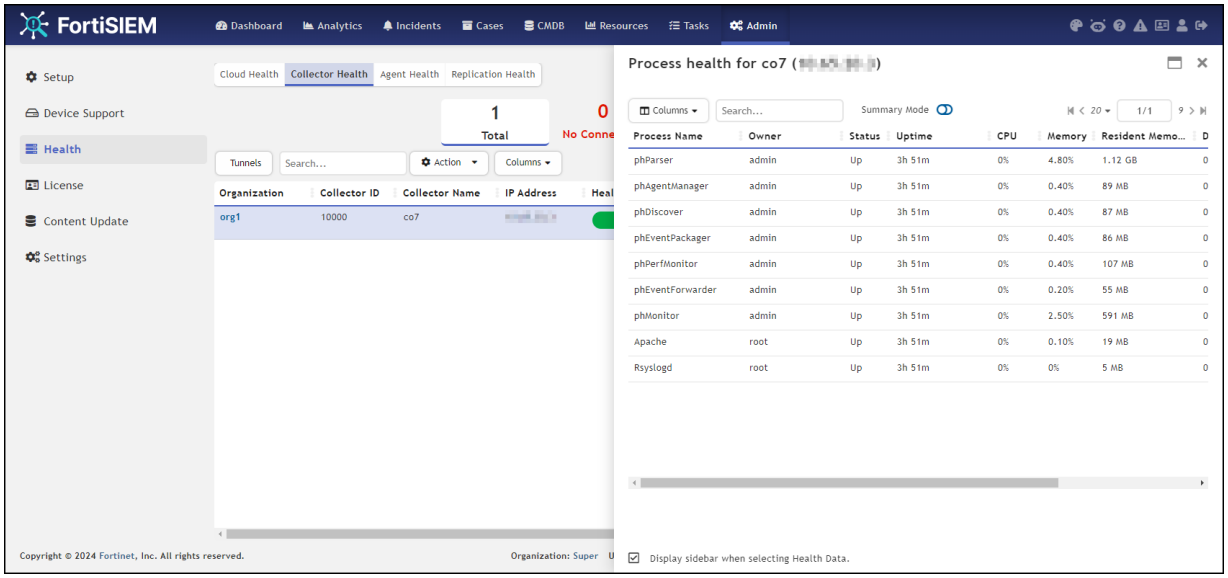

2. Check that the Redis passwords match on the Supervisor and Workers:
   - Supervisor: run the command `phLicenseTool --showRedisPassword`
   - Worker: run the command `grep -i auth /opt/node-rest-service/ecosystem.config.js`

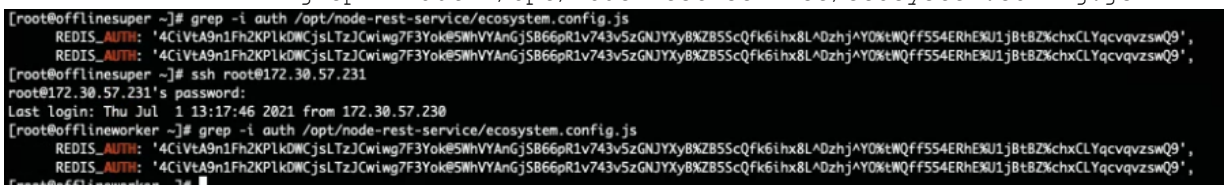

3. Check that the database passwords match on the Supervisor and Workers:
   - Supervisor: run the command `phLicenseTool --showDatabasePassword`
   - Worker: run the command `phLicenseTool --showDatabasePassword`
4. Elasticsearch case: check the Elasticsearch health
   Note: Status should be Normal,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

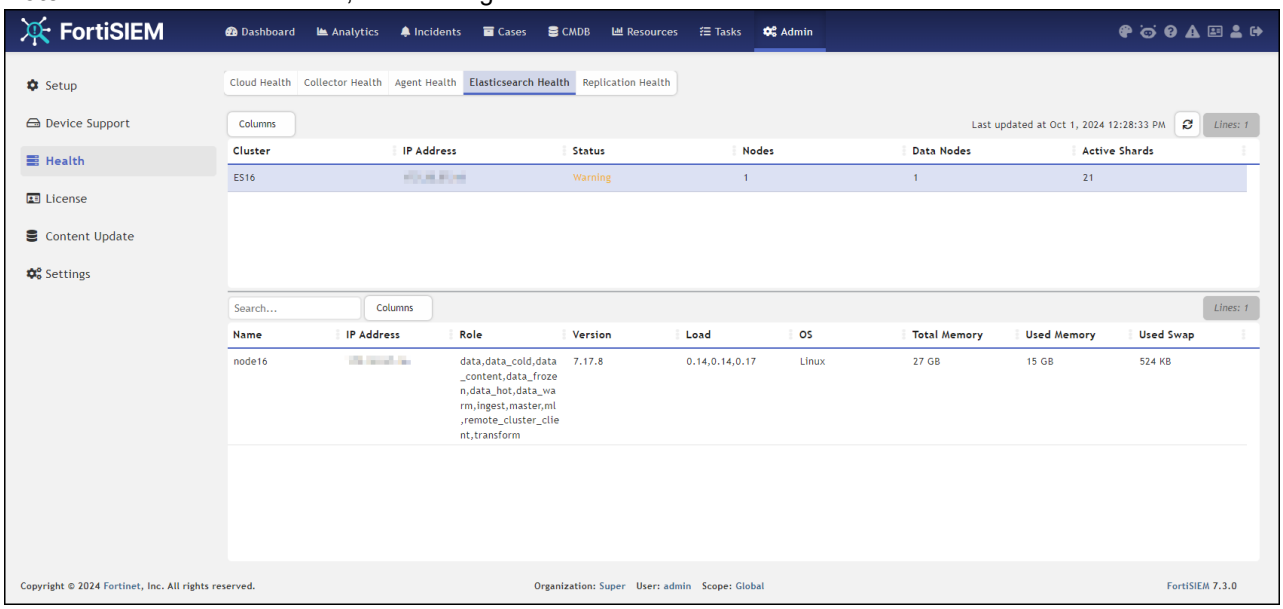

5. Check that events are received correctly:
   a. Search All Events in last 10 minutes and make sure there is data.

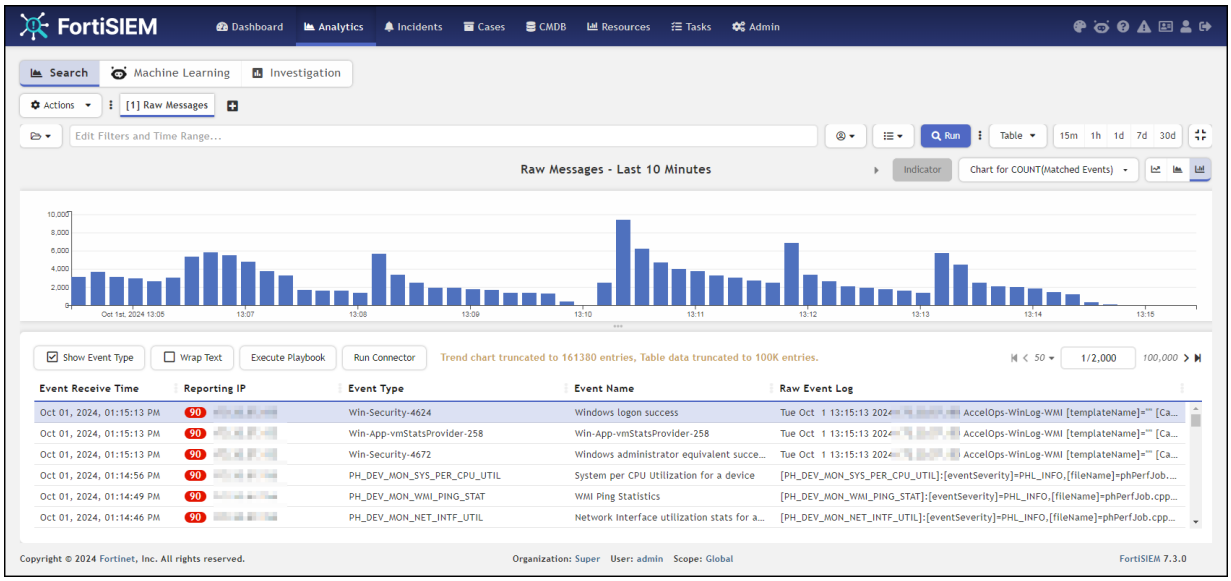

   b. Search for events from Collector and Agents and make sure there is data. Both old and new collectors and agents must work.

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

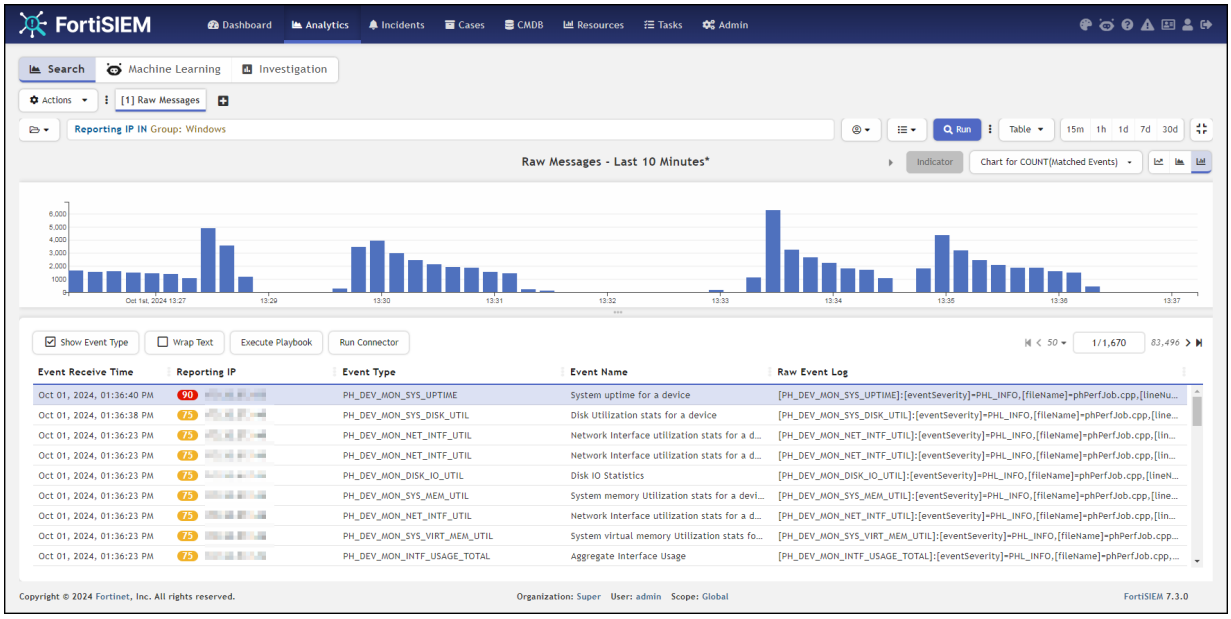

6. Make sure there are no SVN authentication errors in CMDB when you click any device name.

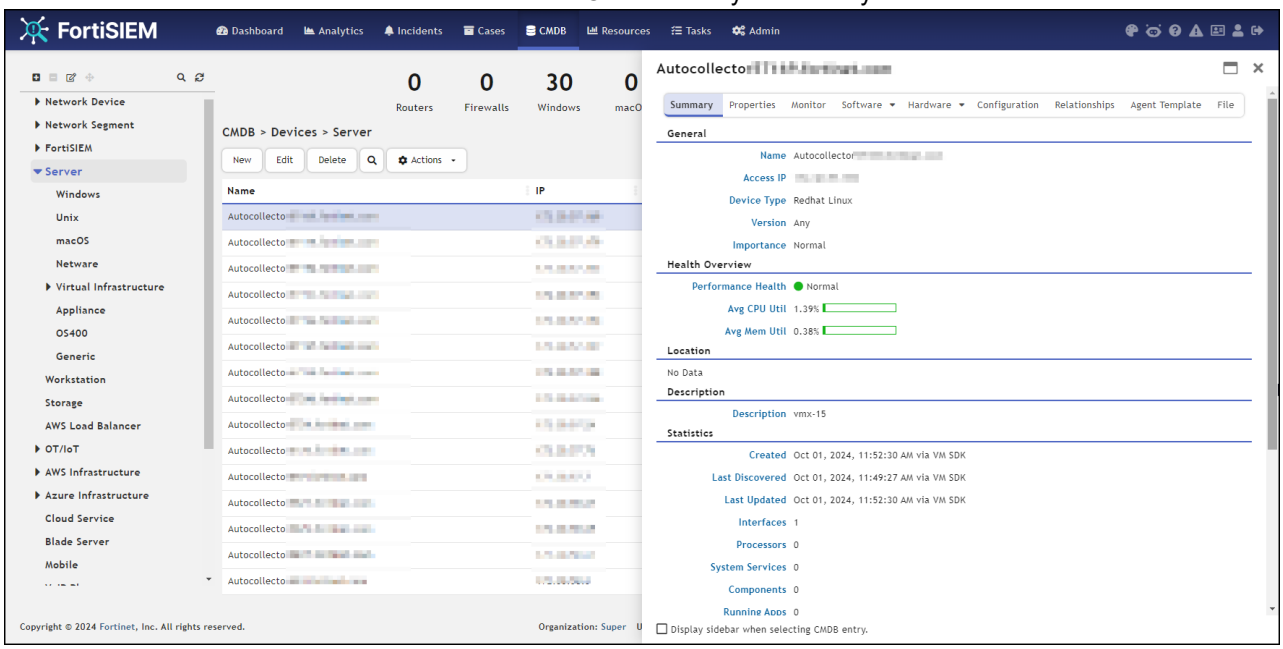

7. Make sure recent Incidents and their triggering events are displayed.

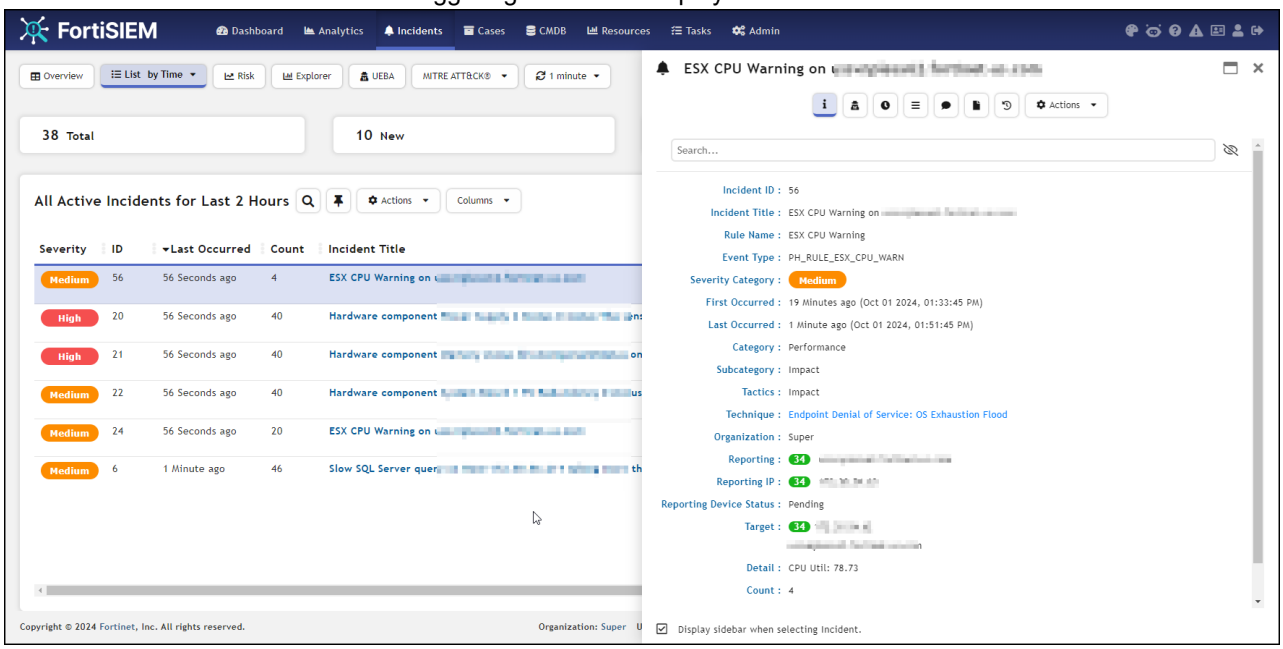

8. Check Worker for Collector Credentials by running the following command:
   `cat /etc/httpd/accounts/passwds`
   This validates that all workers contain collector credentials to log in and upload logs.
9. Run the following script on the Supervisor.
   `get-fsm-health.py --local`
   Your output should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

Upgrade via Proxy

During upgrade, the FortiSIEM Supervisor, Worker, or Hardware appliances (FSM-2000F, 2000G, 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

When you upgrade a FortiSIEM system running on hardware (2000F, 3500F, 3500G, 3600G, 500F) to 6.3.1 and later, the upgrade automatically makes a system backup of root disk, boot disk, opt disk, and in case of 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.

Restoring from Backup

Restoring from backup will restore the root disk, boot disk, opt disk, and in case of 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 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.

   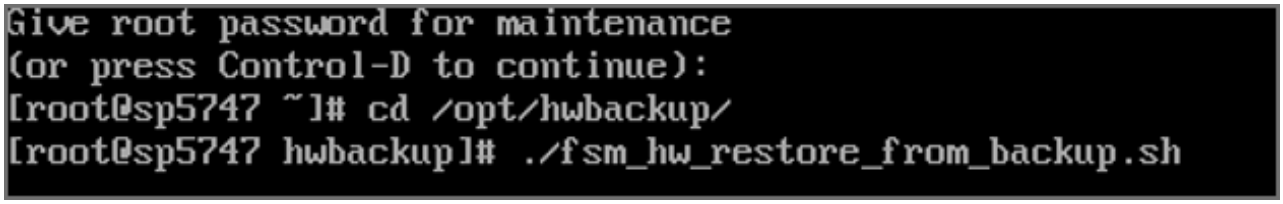

4. If the backup is stored on `/opt/hwbackup`, you can `chdir` to this. If the backup is stored on an external disk, mount the disk and symlink it again to `/opt/hwbackup`.
5. Run the restore command:

   ```
   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:

   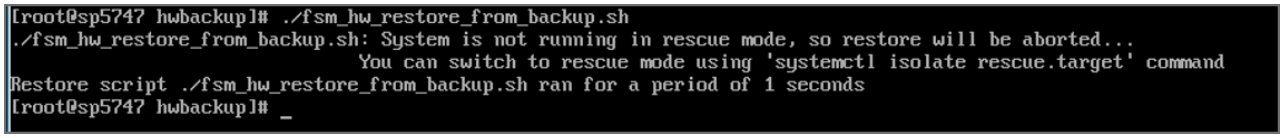

   The whole restore may take anywhere from 15 minutes to more than an hour depending on how large the CMDB/SVN partitions are. The restore script will make sure that the SHA 256 checksums for the backup files match and only then, will it proceed. If this fails, then it will stop the restore process immediately.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

   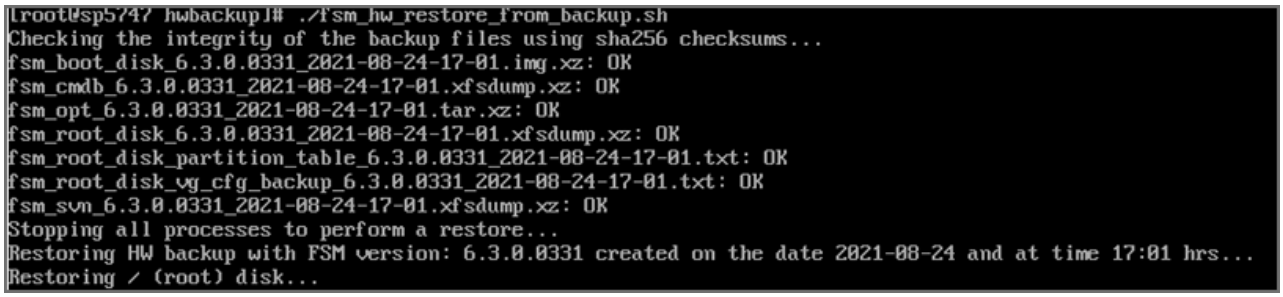


```
Restoring HW backup with FSM version: 6.3.0.0331 created on the date 2021-08-24 and at time 17:01 hrs...
Restoring / (root) disk...
xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.8 (dump format 3.0)
xfsrestore: searching media for dump
xfsrestore: examining media file 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]#
```

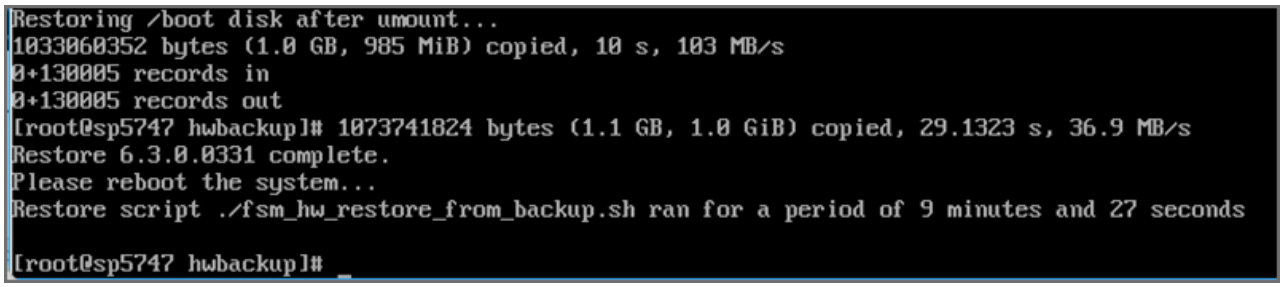

6. Once the restore is complete, it will print how long the restore took and will ask you to reboot the system. Run the command to reboot your system:

   `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, 3500G or 3600G as a worker node, or 500F 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 7.3.0.0338 Upgrade ansible log file is located here: `/usr/local/upgrade/logs/ansible.log`.

Errors can be found at the end of the file.

Migrate Log

The 5.3.x/5.4.x to 6.1.x Migrate ansible log file is located here: `/usr/local/migrate/logs/ansible.log`.

Errors can be found at the end of the file.

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
```

Fix After Upgrading 2000F, 3500F, 3500G from 5.3.x or 5.4.0 to 6.1.2

After upgrading hardware appliances 2000F, 3500F, or 3500G from 5.3.x or 5.4.0 to 6.1.2, the swap is reduced from 24GB to 2GB. Note that the upgrade from 6.1.2 to 6.2.x does not have this problem. This will impact performance. To fix this issue, take the following steps.

1. First, run the following command based on your hardware appliance model.
   For 2000F
   ```
   swapon -s /dev/mapper/FSIEM2000F-phx_swap
   ```
   For 3500F
   ```
   swapon -s /dev/mapper/FSIEM3500F-phx_swap
   ```
   For 3500G
   ```
   swapon -s /dev/mapper/FSIEM3500G-phx_swap
   ```
2. Add the following line to `/etc/fstab` for the above swap partition based on your hardware appliance model.
   For 2000F
   ```
   /dev/FSIEM2000F/phx_swap /swapfile swap defaults 0 0
   ```
   For 3500F
   ```
   /dev/FSIEM3500F/phx_swap /swapfile swap defaults 0 0
   ```
   For 3500G
   ```
   /dev/FSIEM3500G/phx_swap /swapfile swap defaults 0 0
   ```
3. Reboot the hardware appliance.
4. Run the following command
   ```
   swapon --show
   ```
   and make sure there are 2 swap partitions mounted instead of just 1, as shown here.

```
[root@sp5753 ~]# swapon --show
NAME      TYPE      SIZE USED PRIO
/dev/dm-5 partition  30G   0B   -3
/dev/dm-0 partition 2.5G   0B   -2
```

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