Upgrade Guide

FortiSIEM 6.2.1

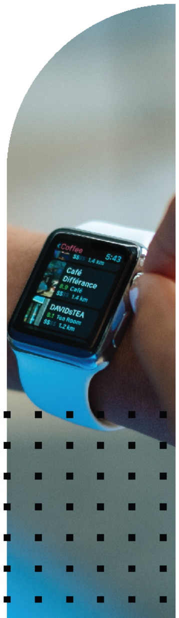

01/31/2022
FortiSIEM 6.2.1 Upgrade Guide

TABLE OF CONTENTS

Change Log

Date	Change Description
03/22/2021	Initial version of the 6.2.0 Upgrade Guide.
03/29/2021	Added Upgrade via Proxy and Post Upgrade Health Check.
03/31/2021	Added Reference section with additional DNS information.
04/05/2021	Updated Pre-Upgrade Checklist.
04/22/2021	Added Upgrade and Migrate Log sections.
05/06/2021	Initial version of the 6.2.1 Upgrade Guide.
05/12/2021	Updated Upgrade via Proxy section.
05/17/2021	Updated existing heading, added Sizing Guide link, removed DNS check for 6.2.1 Upgrade Guide.
05/19/2021	Added "Fix After Upgrading 2000F or 3500F From 5.3.x or 5.4.0 to 6.1.2" section for 6.2.x Upgrade Guides.
05/21/2021	Update to "After Upgrading 2000F or 3500F From 5.3.x or 5.4.0 to 6.1.2" section for 6.2.x Upgrade Guides.
05/24/2021	Update to "Upgrade Collectors" sections for 6.2.x Upgrade Guides.
06/03/2021	Known Issue after 6.2.1 Upgrade added to 6.2.1 Upgrade Guide.
06/07/2021	Update to "Upgrade Collectors" sections for 6.2.1 Upgrade Guide.
06/23/2021	Added 3500G Hardware for 6.2.x Upgrade Guides.
06/24/2021	Upgrading From 6.1.x or 6.2.0 to 6.2.1 section updated for 6.2.x Upgrade Guides.
07/21/2021	Updated Pre-Upgrade Checklist section.
07/22/2021	Updated Upgrade via Proxy section.
12/01/2021	Updated Pre-Upgrade Checklist section.

Important Notes for 6.2.0 to 6.2.1 Upgrade

1. For your Supervisor and Worker, do not use the upgrade menu item in configFSM.sh to upgrade from 6.2.0 to 6.2.1. This is deprecated, so it will not work. Use the new method as instructed in this guide (See **Upgrade Supervisor** for your appropriate deployment in Upgrading From 6.1.x or 6.2.0 to 6.2.1).
2. Before upgrading Collectors to 6.2.1, you will need to copy the `phcollectorimageinstaller.py` file from your Supervisor to your Collectors. See steps 1-3 in Upgrade Collectors.
3. 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.
4. Remember to remove the browser cache after logging on to the 6.2.1 GUI and before doing any operations.
5. 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 6.2.1 after you have upgraded all Workers.

Important Notes for 6.1.x to 6.2.1 Upgrade

1. For your Supervisor and Worker, do not use the upgrade menu item in configFSM.sh to upgrade from 6.1.x to 6.2.1. This is deprecated, so it will not work. Use the new method as instructed in this guide (See **Upgrade Supervisor** for your appropriate deployment in Upgrading From 6.1.x or 6.2.0 to 6.2.1).
2. Before upgrading Collectors to 6.2.1, you will need to copy the `phcollectorimageinstaller.py` file from your Supervisor to your Collectors. See steps 1-3 in Upgrade Collectors.
3. 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.
4. The 6.2.1 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.
5. If you are using AWS Elasticsearch, then after upgrading to 6.2.1, take the following steps:
   a. Go to **ADMIN > Setup > Storage> Online**.
   b. Select "ES-type" and re-enter the credential.
6. If you have more than 5 Workers, Fortinet recommends using at least 16 vCPU for the Supervisor and to increase the number of notification threads for RuleMaster (See the sizing guide for more information). To do this, SSH to the Supervisor and take the following steps:
   a. Modify the `phoenix_config.txt` file, located at `/opt/phoenix/config/` with
   ```
   #notification will open threads to accept connections
   #FSM upgrade preserves customer changes to the parameter value
   notification_server_thread_num=50
   ```
   Note: The default notification_server_thread_num is 20.
   b. Restart phRuleMaster.
7. Upgrading Elasticsearch Transport Client usage - The Transport Client option has been removed as Elasticsearch no longer supports that client. If you are using Transport Client in pre-6.2.1, you will need to modify the existing URL by adding "http://" or "https://" in front of the **URL** field after upgrading, as displayed in **ADMIN > Setup > Storage > Online** > with **Elasticsearch** selected.
   a. When upgrading to 6.2.1, the Elasticsearch **Cluster IP/Host** field changes to the **URL** field. In the **URL** field, add "http://" or "https://" to your IP address.

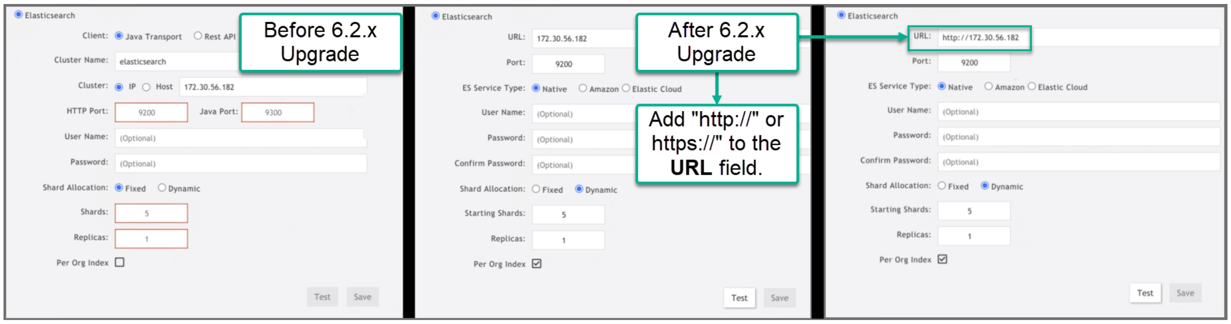

Before 6.2.x Upgrade

● Elasticsearch

Client: ● Java Transport ○ Rest API

Cluster Name: elasticsearch

Cluster: ● IP ○ Host 172.30.56.182

HTTP Port: 9200 Java Port: 9300

User Name: (Optional)

Password: (Optional)

Shard Allocation: ● Fixed ○ Dynamic

Shards: 5

Replicas: 1

Per Org Index □

After 6.2.x Upgrade

● Elasticsearch

URL: 172.30.56.182

Port: 9200

ES Service Type: ● Native ○ Amazon ○ Elastic Cloud

User Name: (Optional)

Password: (Optional)

Confirm Password: (Optional)

Shard Allocation: ○ Fixed ● Dynamic

Starting Shards: 5

Replicas: 1

Per Org Index ☑

Add "http://" or https://" to the **URL** field.

● Elasticsearch

URL: http://172.30.56.182

Port: 9200

ES Service Type: ● Native ○ Amazon ○ Elastic Cloud

User Name: (Optional)

Password: (Optional)

Confirm Password: (Optional)

Shard Allocation: ○ Fixed ● Dynamic

Starting Shards: 5

Replicas: 1

Per Org Index ☑

b. Next, select **Test** to confirm functionality, and select **Save** to save the updated settings.

8. Prior to upgrading, ensure that hot node and warm node counts are both greater than the number of replicas. Failure to do so will result in Test and Save operation failure after an upgrade. This basic requirement check has been added for version 6.2.0 and later.

9. Remember to remove the browser cache after logging on to the 6.2.1 GUI and before doing any operations.

Known Issue after 6.2.1 Upgrade

As part of a fix for excessive SSL communication errors between phReportWorker and phReportMaster (Bug 710109) in 6.2.1, the value for `count_distinct_precision` in the `/opt/phoenix/config/phoenix_config.txt` file is reduced from 14 to 9 for the Supervisor and Worker nodes. While 6.2.1 fresh install sets the value correctly, 6.2.1 upgrade may keep the old value (14) and fail to set the new value (9). Because of this, you can still have excessive SSL communication errors between phReportWorker and phReportMaster and it may appear that bug 710109 is not fixed.

To fix this issue, follow the instructions in Modification on the Supervisor and Modification on each Worker for your Supervisor and Workers.

Modification on the Supervisor

1. SSH into the super as root and edit the `/opt/phoenix/config/phoenix_config.txt` file.

   ```
   ssh root@<supervisor FQDN/IP>
   su admin
   vi /opt/phoenix/config/phoenix_config.txt
   ```

2. Find "count_distinct_precision=".
3. Modify the value to that shown here.

   ```
   count_distinct_precision=9 # in range 4-18
   ```

4. Save the configuration.
5. Stop phRerportMaster/Worker by running the following commands.

   ```
   phtools --stop phReportWorker
   phtools --stop phReportMaster
   ```

6. Start phReportMaster/Worker by running the following commands.

   ```
   phtools --start phReportMaster
   phtools --start phReportWorker
   ```

7. Monitor Stability by running the following command.

   ```
   phstatus
   ```

Modification on each Worker

1. SSH into each Worker as root and edit the `/opt/phoenix/config/phoenix_config.txt` file by running the following commands.

   ```
   ssh root@<Worker FQDN/IP>
   ```

```
su admin
```

```
vi /opt/phoenix/config/phoenix_config.txt
```

2. Find "count_distinct_precision=".
3. Modify the value to that shown here.

   ```
   count_distinct_precision=9 # in range 4-18
   ```

4. Save the configuration.
5. Stop phReportWorker by running the following command.

   ```
   phtools --stop phReportWorker
   ```

6. Start phReportWorker by running the following command.

   ```
   phtools --start phReportWorker
   ```

7. Monitor Stability by running the following command.

   ```
   phstatus
   ```

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. Make sure the FortiSIEM license is not expired.
6. Make sure the Supervisor, Workers and Collectors can connect to the Internet on port 443 to the CentOS OS repositories (os-pkgs-cdn.fortisiem.fortinet.com and os-pkgs.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 Upgrade Guide.

Upgrade Overview

The general upgrade paths are:

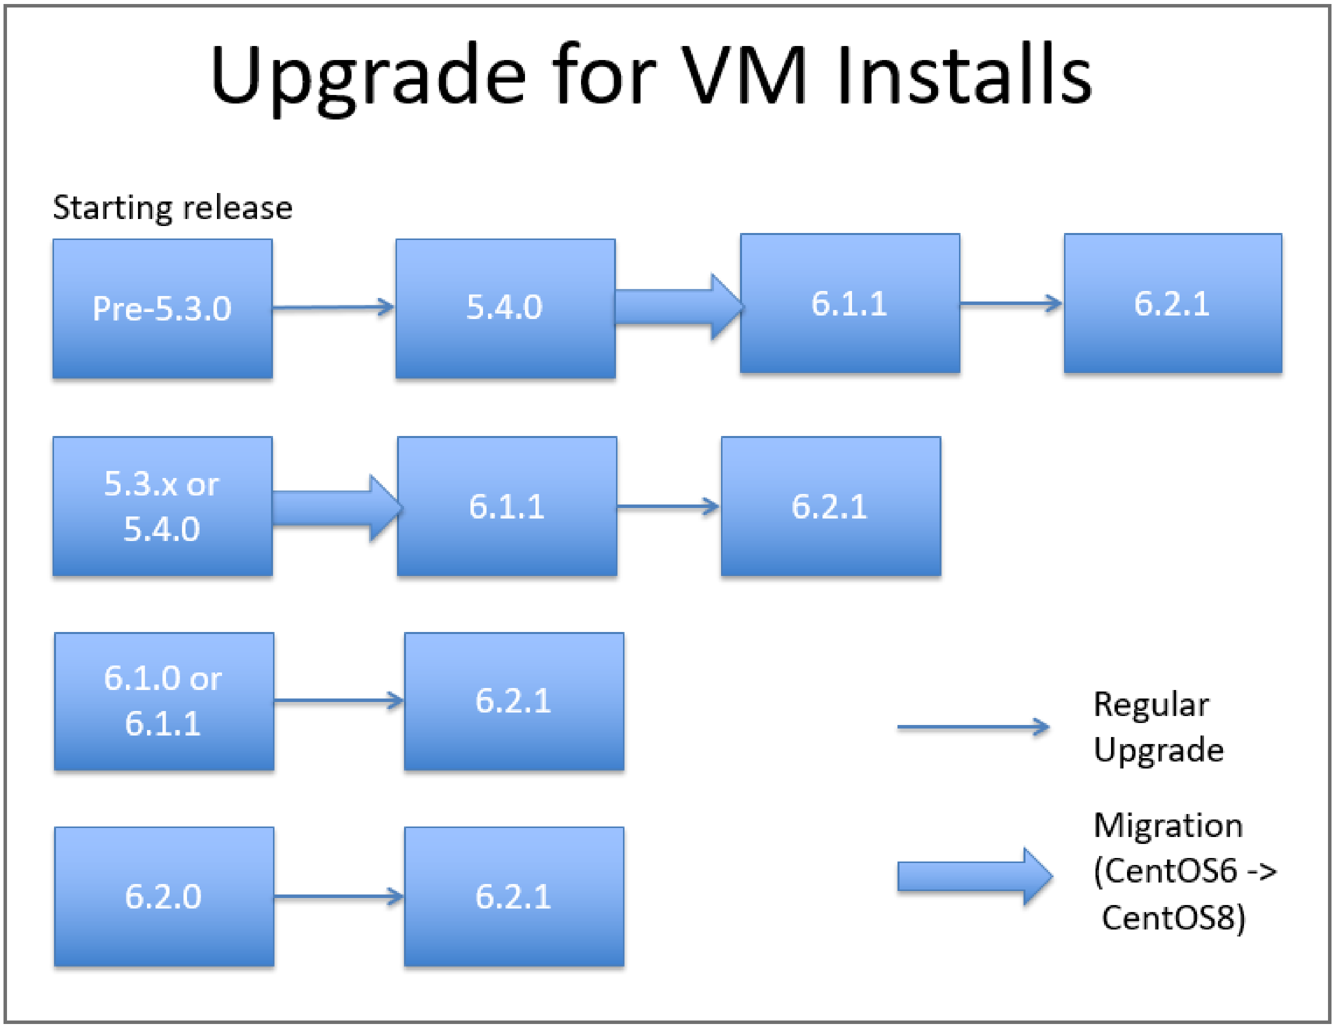

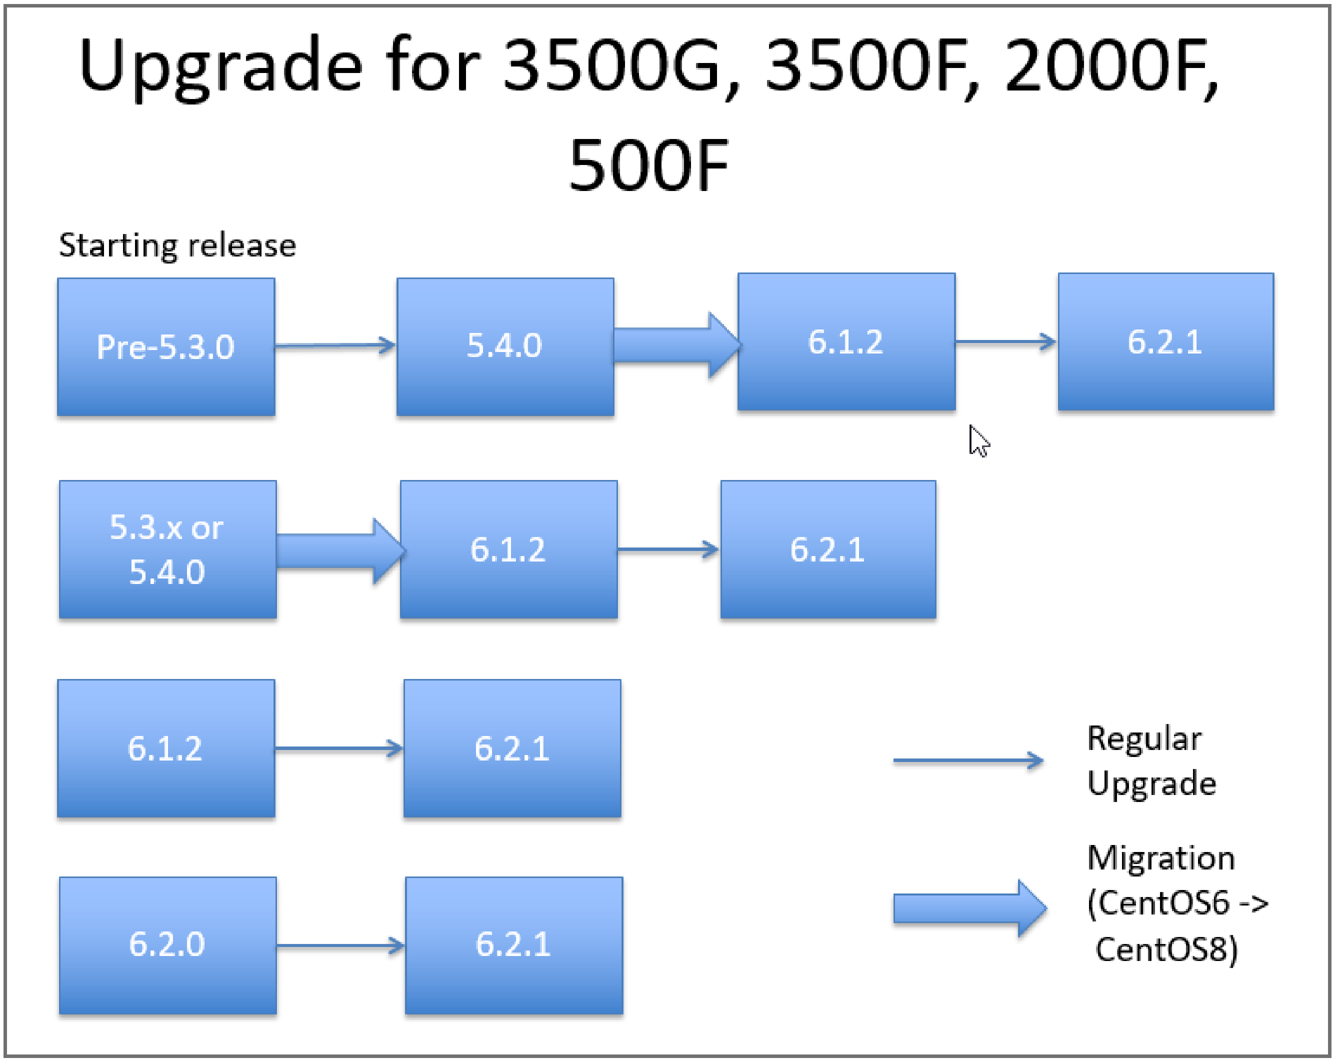

Upgrading From Pre-5.3.0

If you are running FortiSIEM that is pre-5.3.0, take the following steps:

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 6.2.x by following the steps in Upgrading From 6.1.x.

Upgrading From 5.3.x or 5.4.0

Start at step 3 from Upgrading From Pre-5.3.0, and follow the progressive steps.

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.

Upgrading From 6.1.x or 6.2.0 to 6.2.1

Note: Prior to your 6.1.x to 6.2.x upgrade, or 6.2.0 to 6.2.1 upgrade, ensure that the Supervisor, and all Workers, and Collectors are running 6.1.x.

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

After completion of your upgrade, follow the appropriate steps in Post Upgrade Health Check.

There are two possible upgrades after upgrading to 6.1.x. Follow the steps for your appropriate FortiSIEM setup for single node deployment or cluster deployment.

- Upgrade via Proxy
- Post Upgrade Health Check
- Upgrade Single Node Deployment
- Upgrade Cluster Deployment

Upgrade via Proxy

During upgrade, the FortiSIEM Supervisor, Worker or Hardware appliances (FSM-2000F, 3500F, and 3500G) must be able to communicate with CentOS OS repositories (`os-pkgs-cdn.fortisiem.fortinet.com` and `os-pkgs.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.

Post Upgrade Health Check

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 `grep Auth_PQ_dbpass /etc/httpd/conf/httpd.conf`
4. Elasticsearch case: check the Elasticsearch health
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.

Upgrade Single Node Deployment

Upgrading a single node deployment requires upgrading your supervisor. If you have any collectors, make sure to upgrade your supervisor first before upgrading them.

- Upgrade Supervisor
- Upgrade Collectors

Upgrade Supervisor

To upgrade your Supervisor, take the following steps.

1. Make sure Workers are shut down. Collectors can remain up and running.
2. Login to the Supervisor via SSH.
3. Create the `/upgrade` directory.
   ```
   mkdir -p /opt/upgrade
   ```
4. Download the upgrade zip package `FSM_Upgrade_All_6.2.1_build0223.zip`, then upload it to your Supervisor node under the `/opt/upgrade/` folder.
5. Go to `/opt/upgrade`.
   ```
   cd /opt/upgrade
   ```
6. Unzip the upgrade zip package.
   ```
   unzip FSM_Upgrade_All_6.2.1_build0223.zip
   ```
7. Go to the FSM_Upgrade_All_6.2.1_build0223 directory.
   ```
   cd FSM_Upgrade_All_6.2.1_build0223
   ```
   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.

Upgrade Collectors

To upgrade your Collectors, take the following steps.

Extra Upgrade Steps from 6.1.x to 6.2.1

If you are upgrading from 6.1.x to 6.2.1, then take the following steps first. If you are already on 6.2.1, these steps are not required.

1. Login to the Collector via SSH as root.
2. Copy `/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py` from your Supervisor by running the following command. (**Note**: This is copied from the 6.2.1 upgraded 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

1. Login to the Supervisor via SSH as root.
2. Prepare the Collector upgrade image by running the following command on the Supervisor.
   ```
   phSetupCollectorUpgrade.sh /opt/upgrade/FSM_Upgrade_All_6.2.1_build0223.zip
   <SupervisorFQDN>
   ```
   Note: Replace <*SupervisorFQDN*> with the fully qualified domain name of the Supervisor.

   This is what the output is in shell:
   ```
   # phSetupCollectorUpgrade.sh
   ```
   Usage: `/opt/phoenix/phscripts/bin/phSetupCollectorUpgrade.sh coImageZipPath superFQDN`
3. Go to **ADMIN > Health > Collector Health**.
4. Select a Collector.
   a. Download the image by clicking **Download Image**.
   b. Upgrade the image by clicking **Install Image**.
5. Make sure the Collector and all its processes are up.
6. Repeat steps 3 through 5 for all Collectors.

Upgrade Cluster Deployment

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

- Overview
- Detailed Steps
- Upgrade Supervisor
- Upgrade Workers
- Upgrade Collectors

Overview

1. Shut down all Workers.
   - Collectors can be up and running.
2. Upgrade the Supervisor first, while all Workers are shut down.
3. After the Supervisor upgrade is complete, verify it is up and running.
4. Upgrade each Worker individually.
5. 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.
6. Upgrade each Collector individually.

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.

Note: 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 your 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. Shutdown all Worker nodes.

```
# shutdown now
```

2. Upgrade your 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 are running Elasticsearch, and upgrading from 6.1.x to 6.2.1, then take the following steps, else skip this step and proceed to Step 4.
   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.
4. Upgrade each Worker one by one, using the procedure in Upgrade Workers.
5. 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.
6. Upgrade Collectors using the steps in Upgrade Collectors.

Upgrade Supervisor

To upgrade your Supervisor, take the following steps.

1. Make sure Workers are shut down. Collectors can remain up and running.
2. Login to the Supervisor via SSH.
3. Create the `/upgrade` directory.
   ```
   mkdir -p /opt/upgrade
   ```
4. Download the upgrade zip package `FSM_Upgrade_All_6.2.1_build0223.zip`, then upload it to your Supervisor node under the `/opt/upgrade/` folder.
5. Go to `/opt/upgrade`.
   ```
   cd /opt/upgrade
   ```
6. Unzip the upgrade zip package.
   ```
   unzip FSM_Upgrade_All_6.2.1_build0223.zip
   ```
7. Go to the FSM_Upgrade_All_6.2.1_build0223 directory.
   ```
   cd FSM_Upgrade_All_6.2.1_build0223
   ```
   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.

Upgrade Workers

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

1. Login to a worker via SSH.
2. Create the `/upgrade` directory.
   ```
   mkdir -p /opt/upgrade
   ```
3. Download the upgrade zip package `FSM_Upgrade_All_6.2.1_build0223.zip` to `/opt/upgrade`.
4. Go to `/opt/upgrade`.
   ```
   cd /opt/upgrade
   ```
5. Unzip the upgrade zip package.
   ```
   unzip FSM_Upgrade_All_6.2.1_build0223.zip
   ```
6. Go to the FSM_Upgrade_All_6.2.1_build0223 directory.
   ```
   cd FSM_Upgrade_All_6.2.1_build0223
   ```
   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.

Upgrade Collectors

Extra Upgrade Steps from 6.1.x to 6.2.1

If you are upgrading from 6.1.x to 6.2.1, then take the following steps first. If you are already on 6.2.1, these steps are not required.

1. Login to the Collector via SSH as root.
2. Copy `/opt/phoenix/phscripts/bin/phcollectorimageinstaller.py` from your Supervisor by running the following command. (**Note**: This is copied from the 6.2.1 upgraded 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 Collectors, take the following steps.

1. Login to the Supervisor via SSH as root.
2. Prepare the Collector upgrade image by running the following command on the Supervisor.

```
phSetupCollectorUpgrade.sh /opt/upgrade/FSM_Upgrade_All_6.2.1_build0223.zip
<SupervisorFQDN>
```

Note: Replace <*SupervisorFQDN*> with the fully qualified domain name of the Supervisor.

This is what the output is in shell:

```
# phSetupCollectorUpgrade.sh
Usage: /opt/phoenix/phscripts/bin/phSetupCollectorUpgrade.sh coImageZipPath superFQDN
```

3. Go to **ADMIN > Health > Collector Health**.
4. Select a Collector.
   a. Download the image by clicking **Download Image**.
   b. Upgrade the image by clicking **Install Image**.
5. Make sure the Collector and all its processes are up.
6. Repeat steps 3 through 5 for all Collectors.

Upgrade Log

The 6.2.1_build0223 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 `/opt` disk by 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, or 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
```

FORTINET®

www.fortinet.com