



# Clustering

---

This chapter describes clustering Threat Grid Appliances. It includes the following topics:

- [About Clustering Threat Grid Appliances, on page 1](#)
- [Build Cluster Overview, on page 4](#)
- [Start Cluster of Thread Grid Appliances, on page 7](#)
- [Join Threat Grid Appliances to Cluster, on page 15](#)
- [Designate Tiebreaker Node, on page 19](#)
- [Remove Cluster Node, on page 20](#)
- [Resize Cluster, on page 20](#)
- [Failure Tolerances, on page 21](#)
- [Failure Recovery, on page 21](#)
- [API/Usage Characteristics, on page 22](#)
- [Operational/Adminstrative Characteristics, on page 22](#)
- [Sample Deletion, on page 22](#)

## About Clustering Threat Grid Appliances

The ability to cluster multiple Threat Grid Appliances is available in v2.4.2 or later. Each appliance in a cluster saves data in the shared file system, and will therefore have the same data as the other nodes in the cluster.

The main goal of clustering is to increase the capacity of a single system by joining several appliances together into a cluster (consisting of 2 to 7 nodes). Clustering also helps support recovery from failure of one or more machines in the cluster, depending on the cluster size.

If you have questions about installing or reconfiguring clusters, contact Cisco Support for assistance to avoid possible destruction of data.

## Clustering Features

Clustering Threat Grid Appliances offers the following features:

- **Shared Data** - Every appliance in a cluster can be used as if they are standalone; each is accessing and presenting the same data.
- **Sample Submissions Processing** - Submitted samples are processed on any one of the cluster members, with any other member able to see the analysis results.

- **Rate Limits** - The submission rate limits of each member are added up to become the cluster's limit.
- **Cluster Size** - The preferred cluster sizes are 3, 5, or 7 members; 2-, 4- and 6-node clusters are supported, but with availability characteristics similar to a degraded cluster (a cluster in which one or more nodes are not operational) of the next size up.
- **Tiebreaker** - When a cluster is configured to contain an even number of nodes, the one designated as the tiebreaker gets a second vote in the event of an election to decide which node has the primary database.

Each node in a cluster contains a database, but only the database on the primary node is actually used; the others just have to be able to take over if and when the primary node goes down. Having a tiebreaker can prevent the cluster from being down when exactly half the nodes have failed, but only when the tiebreaker is not among the failed nodes.

Odd-numbered clusters won't have a tied vote. In an odd-numbered cluster, the tiebreaker role will only become relevant if a node (not the tiebreaker) is dropped from the cluster, which would then become even-numbered.




---

**Note** This feature is fully tested only for 2-node clusters.

---

## Clustering Limitations

Clustering Threat Grid Appliances has the following limitations:

- When building a cluster of existing standalone appliances, only the first node (the initial node) can retain its data. The other nodes must be manually reset because merging existing data into a cluster is not allowed.

Remove existing data with the destroy-data command, as documented in [Reset Threat Grid Appliance as Backup Restore Target](#)




---

**Important** Do not use the Wipe Appliance feature as it will render the appliance inoperable until it's returned to Cisco for reimaging.

---

- Adding or removing nodes can result in brief outages, depending on cluster size and the role of the member nodes.
- Clustering on the M3 server is not supported. Contact [support@threatgrid.com](mailto:support@threatgrid.com) if you have any questions.

## Clustering Requirements

The following requirements must be met when clustering Threat Grid Appliances:

- **Version** - All appliances must be running the same version to set up a cluster in a supported configuration, and it should always be the latest available version.
- **Clust Interface** - Each Threat Grid Appliance requires a direct interconnect to the other appliances in that cluster, with a SFP+ (not included with the standalone appliance) installed into the Clust interface slot on each one. Direct interconnect, in this context, means that all appliances must be on the same

layer-2 network segment, with no routing required to reach other nodes, and without significant latency or jitter. Network topologies where the nodes are not on a single physical network segment are not supported.

- **Airgapped Deployments Discouraged** - Due to the increased complexity of debugging, appliance clustering is strongly discouraged in airgapped deployments or other scenarios where a customer is unable or unwilling to provide L3 support access to debug.
- **Data** - An appliance may only be joined to a cluster when it contains no data (only the initial node can contain data). Moving an existing appliance into a data-free state requires the use of the database reset process (available in v2.2.4 or later).



---

**Important** Do not use the destructive Wipe Appliance process, which removes all data and renders the application inoperable until it's returned to Cisco for reimaging.

---

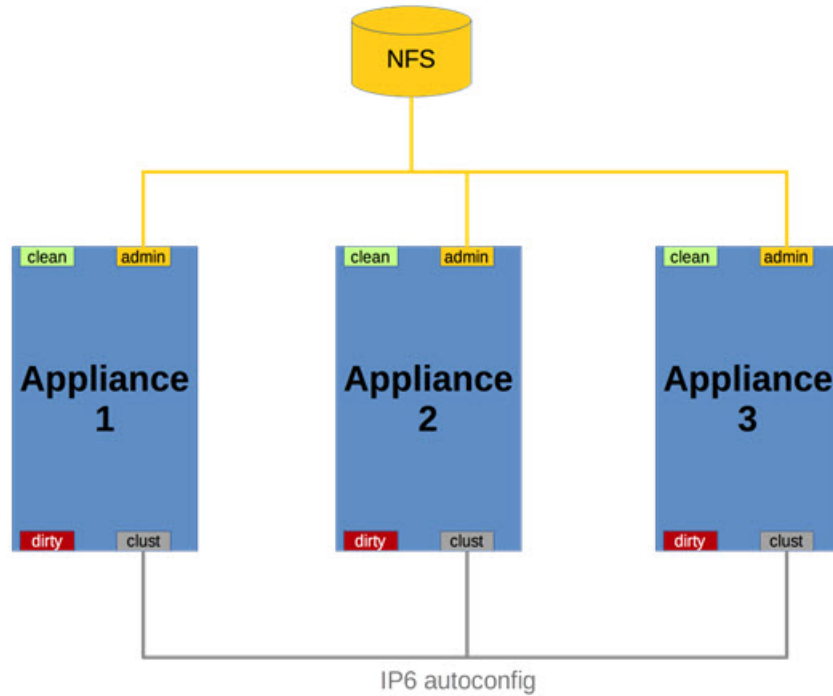
- **SSL Certificates** - If you are installing SSL certificates signed by a custom CA on one cluster member, then all other nodes' certificates should be signed by the same CA.

## Networking and NFS Storage

Clustering Threat Grid Appliances requires the following networking and NFS storage considerations:

- Threat Grid appliance clusters require a NFS store to be enabled and configured. It must be available via the Admin interface, and must be accessible from all cluster nodes.
- Each cluster must be backed by a single NFS store with a single key. While that NFS store may be initialized with data from a preexisting appliance, it **MUST NOT** be accessed by any system which is not a member of the cluster while the cluster is in operation.
- The NFS store is a single point of failure, and the use of redundant, highly reliable equipment for that role is essential.

Figure 1: Clustering Network Diagram



## Build Cluster Overview

Building a cluster in a supported manner requires that all members be on the same version, which should always be the latest available version. This may mean that all of the members have to be built standalone first to get fully updated.

If the appliance has been in use as a standalone machine prior to clustering, only the data of the first member can be preserved. The others need to be reset as part of the build.

Start a new cluster with an initial node, and then join other appliances to it. There are two distinct paths that are available to starting a new cluster:


- Start a new cluster using an existing standalone appliance.
- Start a new cluster using a new appliance.

## Clust Interface Setup

Each appliance in the cluster requires an additional SFP+ for the Clust interface.

Install a SFP+ module in the fourth (non-Admin) SFP port that was previously labeled **Reserved**; it is now used for the **Clust** interface.

Figure 2: Clust Interface Setup for Cisco UCS M4 C220



## Clustering Configuration

Clusters are configured and managed in the OpAdmin portal on the Clustering configuration page (**Configuration > Clustering**). This section describes the fields on the Clustering configuration page to gain an understanding of an active and healthy cluster (the screenshot shows a 3-node cluster).

Figure 3: Clustering Configuration for Active Cluster

ThreatGRID Appliance Administration Portal

Configure your Threat Grid Appliance to use Clustering.

**Clustering Prerequisites Status**

Installation Status	<input checked="" type="radio"/> Complete
Interface Status	<input checked="" type="checkbox"/> Available
NFS Status	<input checked="" type="checkbox"/> Active
Clustering Status	<input checked="" type="radio"/> Clustered

Start Cluster Join Cluster Make Tiebreaker

**Clustering Components Status**

ES	<input checked="" type="checkbox"/> replicated	PG	<input checked="" type="checkbox"/> replicated
----	--	----	--

**Cluster Nodes Status**

Appliance ID	Pulse	Ping	Consul	Tiebreaker	PG Master	Action
FCH1831V0F2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="button" value="x"/>
FCH1832V319	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="button" value="x"/>
FCH1831V0JQ	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="button" value="x"/>

### Clustering Prerequisites Status

- **Installation Status** The installation status of the appliance; must be **Complete**. The appliance must fully set up and configured.
- **Interface Status** - The status of the clustering network interface, "Clust".

- **NFS Status** - NFS must be available.
- **Clustering Status** - Indicates whether the appliance is a cluster node or standalone
  - **Standalone (unsaved)** - The appliance is not yet configured as either explicitly part of a cluster or a standalone unit. If in the initial setup wizard and the prerequisites for clustering are met, it's possible to make the selection of whether this system will be clustered or not.
  - **Standalone** - Configured as a standalone node. Cannot be configured as part of a cluster without a reset.
  - **Clustered** - The appliance is clustered with one or more other appliances.

### Clustering Components Status

- **ES** - Elasticsearch, the service used for queries that require search functionality.
- **PG** - PostgreSQL, the service used for queries that require up-to-date, definitive data (such as account lookups).

Both services are described with one of the following status values:

- **Replicated** - Everything is working properly. Additionally, everything required to take over on failure is also in place. The appliance is able to tolerate failure and continue working. Being in a replicated state does not mean that a failure will have zero downtime. Rather, a failure should entail zero data loss and constrained downtime (less than a minute in normal circumstances, with the exception of any active analysis on the specific cluster node that fails).

Maintenance operations that bring down nodes should only be performed when the cluster is in the replicated state.

For a fully replicated cluster, recovery should be automatic and require less than a minute to complete in any normal scenario.

- **Available** - Everything is working properly and the referenced service is available for use (that is, it can service API and user requests), but it is not replicated.
- **Unavailable** - The service is known to be non-functional.

### Status Colors:

- **Green** - Replicated
- **Yellow** - Available
- **Red** - Unavailable
- **Grey** - Unknown

For more information, see the [Threat Grid Appliance Clustering FAQ](#) on Cisco.com.

### Cluster Nodes Status

A green checkmark indicates the node is running and healthy.

A red X indicates that something is either not running yet or it's not healthy.

- **Pulse** - Indicates whether the node is actively connected to and using the NFS store (not during initial setup, but while running services).
- **Ping** - Describes whether the cluster node can be seen over the Clust interface
- **Consul** - Indicates whether the node is participating in the consensus store. This requires both a network connection over Clust and a compatible encryption key.
- **Tiebreaker** - Designates the node as the tiebreaker, which will cast the deciding vote in an election to decide the cluster's primary node. See [Designate Tiebreaker Node](#).
- **Keep Standalone** - Indicates that the appliance should not be configured as a node in a cluster. Selecting this option allows the user to complete the rest of the OpAdmin configuration Wizard process for a non-clustered appliance.

## Start Cluster of Thread Grid Appliances

When you build a cluster of Threat Grid Appliances, you must start the cluster with the first node being either an existing standalone Threat Grid Appliance or a new appliance. Refer to the appropriate section for starting the cluster based on your environment.

### Start Cluster with Existing Standalone Appliance

You can start building a cluster from an existing standalone Threat Grid Appliance. This method allows you to preserve existing data from one machine and use it to start a new cluster. An existing backup must be available on tNFS from which the cluster is started.



---

**Note** All other nodes to be joined to the cluster must have data removed before joining; the data from additional nodes cannot be merged into the cluster.

---



---

**Note** In releases prior to v2.4.3.3, standalone appliances with data backed up to NFS no longer require a database reset and restore-from-backup to become the initial node of a new cluster. We suggest that you upgrade to v2.4.3.3 or later and then perform a reset operation prior to initializing a new cluster (if you have an appliance with an earlier version).

---

Perform the following steps to start a cluster for the first node:

---

**Step 1** Fully update the appliance to the latest version. Depending on which version is currently running, this may require more than one update cycle to reach the latest.

**Step 2** If not already done, set up backup of the machine to NFS:

**Note** This step describes the default Linux NFS server implementation, which you may need to adjust depending on your own server setup.

a) In the OpAdmin portal, click Configuration > NFS to open the NFS Configuration page.

Figure 4: NFS Configuration

ThreatGRID Appliance Administration Portal

Support Help  
Logout

Configuration Operations Status Support

**Configuration**

- > Network ✓
- > License ✓
- > **NFS**
- > Clustering ✓
- > Email
- > Notifications
- > Date and Time
- > Syslog ✓

**Other**

- > Review and Install

Start Installation

## NFS

NFS Configuration

Host

Path

Opts

Status

Next >

b) Complete the following fields:

- **Host** - The NFSv4 host server. We recommend using the IP address.
- **Path** - The absolute path to the location on the NFS host server where files will be stored. This does not include the Key ID suffix, which will be added automatically.
- **Opts** - NFS mount options to be used, if this server requires any deviations from standard Linux defaults for NFSv4.
- **Status** - Choose **Enabled (Pending Key)** from the drop-down menu.

c) Click **Next**.

Figure 5: NFS Configuration Enabled (Pending Key)

ThreatGRID Appliance Administration Portal

Support Help  
Logout

Configuration Operations Status Support

**Configuration**

- > Network ✓
- > License ✓
- > **NFS**
- > Clustering ✓
- > Email
- > Notifications
- > Date and Time
- > Syslog ✓

**Other**

- > Review and Install

Start Installation

## NFS

NFS Configuration

Host

Path

Opts

Status

Next >

The page refreshes and **Generate** button becomes available.




The first time you configure this page, the **Remove** and **Download** the encryption key buttons become visible.

The **Upload** button is available if you have NFS enabled but no key created. Once you create a key, the **Upload** button changes to **Download**. If you delete the key, the **Download** button becomes **Upload** again.

**Note** If the key correctly matches the one used to create a backup, the **Key ID** displayed in OpAdmin after upload should match the name of a directory in the configured path. Backups cannot be restored without the encryption key. The configuration process includes the process of mounting the NFS store, mounting the encrypted data, and initializing the appliance's local datastores from the NFS store's contents.


- d) Click **Generate** to generate a new NFS encryption key.

**Figure 6: Generate a New NFS Encryption Key**




- e) Click **Next**. The page refreshes and the **Key ID** is displayed, and the **Activate** and **Download** buttons become available.

Figure 7: Activate the NFS Configuration



- f) Click **Activate**. This will take a few seconds (the status indicator is located in the lower left corner). The **Status** becomes **Active**.

Figure 8: NFS Active



- g) Click **Download** to download the backup encryption key. Save the generated file in a secure location. You will need the key for joining additional nodes to the cluster.

**Important** If this step is missed, all data will be lost in the following steps.

**Step 3** Complete the configuration, as needed, and reboot the appliance to apply the NFS backup configuration.

**Step 4** Perform a backup.

**Note** If you do the backup at least 48 hours in advance, as recommended, and there are no service notices indicating problems with the backup, then the following manual steps are unnecessary.

Backup and other service notices are available in the Threat Grid portal UI from the icon in the upper-right corner. If you see a service notice **There is no PostgreSQL backup yet**, then DO NOT PROCEED.

If you do the backup immediately after reboot, then you will need to manually initiate a backup of all data to NFS to ensure it's complete. Performing the manual backup commands is only necessary if you are setting up backup immediately before rebuilding the standalone box into a cluster.

a) Open tgsh and enter the following commands:

```
service start tg-database-backup.service
service start freezer-backup-bulk.service
service start elasticsearch-backup.service
```

**Figure 9: Initiating a Backup of All Data to NFS**

```

:: [I]string([I]string("CONSOLE"))
Welcome to the ThreatGrid Shell.
For help, type "help" then enter.
>> help
COMMANDS:
  configure -- show|set: View or modify configuration variables
  conns -- listening|open|all: Show open connections
  destroy-data -- Reset appliance to be a target for the restore process
  exit -- Exit tgsh.
  halt -- Halt appliance
  help -- List available commands, or 'help COMMAND' for details.
  netctl -- Configure the network
  netinfo -- routes|firewall|address|stats: Show network configuration and status
  opadmin -- import|check: Sync from, or validate, new configuration format
  passud -- Change password for this account
  ping -- ping [-c count] [-I interface] host: ping a remote host
  poweroff -- Power off appliance
  queues -- Show status of various application queues
  reboot -- Reboot appliance
  service -- {status|start|stop|restart} [svc-name]: Toggle ThreatGRID services
  support-node -- status|start|stop|enable|disable: Toggle support node
  traceroute -- Determine the path used to a network location
  version -- Shows appliance version
>> service start tg-database-backup.service
>> service start freezer-backup-bulk.service
>> service start elasticsearch-backup.service
>>

```

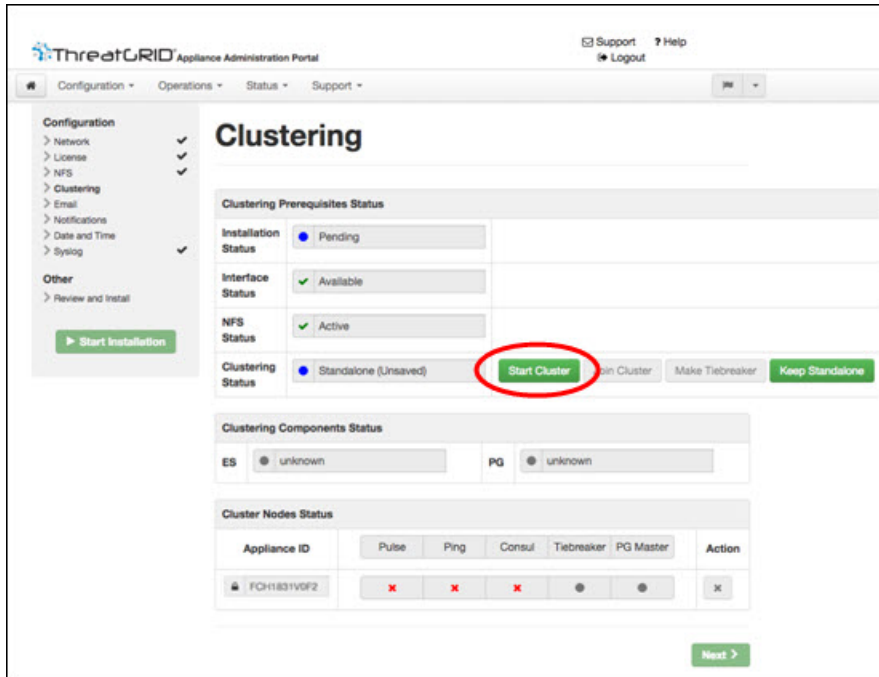
b) Wait about 5 minutes after the last command returns.

**Step 5** In the Threat Grid portal UI, check for service notices. If any notices indicate a backup process failure, such as a warning that there is no PostgreSQL backup yet, then DO NOT PROCEED.

**Important** Do not continue unless these processes have completed successfully.

**Step 6** Navigate to **Configuration > Clustering**.

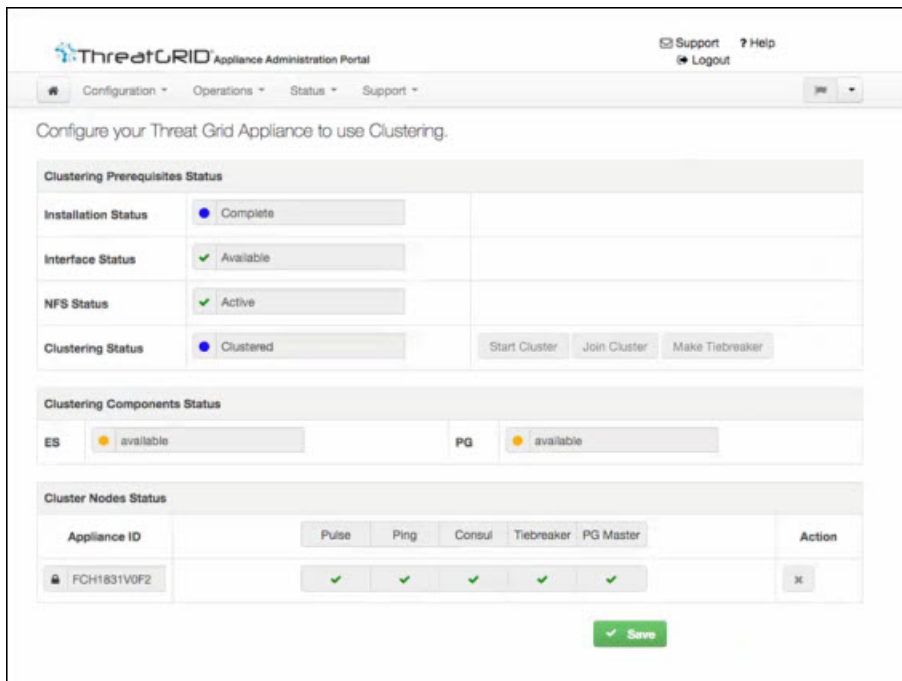
Figure 10: Start Cluster



**Step 7** Click **Start Cluster**.

**Step 8** On the confirmation dialog, click **OK**. The **Clustering Status** changes to **Clustered**.

Figure 11: Clustering Status - Clustered



Once the data restore is complete, return to the Clustering configuration page to check the health of the new cluster.

**Step 9** Finish the installation. This initiates a restore of the data in cluster mode.

---

#### What to do next

Now you can begin joining other appliances to the new cluster, as described in [Join Threat Grid Appliances to Cluster](#).

## Start Cluster with New Appliance

This method of starting a cluster can be used for new appliances that are shipped with cluster-capable versions of the appliance software, or for existing appliances that have had their data reset.



---

**Note** Remove existing data with the `destroy-data` command, as documented in [Reset Threat Grid Appliance as Backup Restore Target](#). Do not use the Wipe Appliance feature.

---

**Step 1** Set up and begin the OpAdmin configuration as normal.

**Step 2** In OpAdmin, click **Configuration > NFS**.

**Note** See the figures in [Start Cluster with Existing Standalone Appliance](#).

**Step 3** Configure the **Network** and **License**.

**Step 4** On the NFS configuration page, complete the following fields:

- **Host** - The NFSv4 host server. We recommend using the IP address.
- **Path** - The absolute path to the location on the NFS host server where the files will be stored. This does not include the Key ID suffix, which will be added automatically.
- **Opts** - NFS mount options to be used, if this server requires any deviations from standard Linux defaults for NFSv4.
- **Status** - Choose **Enabled (Pending Key)** from the drop-down menu.

**Step 5** Click **Next**.

The page refreshes; the **Generate** and **Activate** buttons become available.


**Step 6** Click **Generate** to generate a new NFS encryption key.

**Step 7** Click **Activate**.

The **Status** changes to **Active**.

**Step 8** Click **Download** to download a copy of the encryption key for safekeeping. You will need the key for joining additional nodes to the cluster.

Figure 12: Clustering Configuration Page



**Step 9** Browse to the Clustering configuration page (**Configuration > Clustering**).

**Step 10** Click **Start Cluster**, and then click **OK** on the confirmation dialog.

The **Clustering Status** changes to **Clustered**.

**Step 11** Complete the remaining steps in the wizard and click **Start Installation**. This initiates a restore of the data in cluster mode.

**Step 12** Open the Clustering configuration page to check the health of the new cluster.

Figure 13: Clustering Status: Clustered

The screenshot shows the Threat Grid Appliance Administration Portal with the following details:

- Clustering Prerequisites Status:**
  - Installation Status: Complete (blue dot)
  - Interface Status: Available (green checkmark)
  - NFS Status: Active (green checkmark)
  - Clustering Status: Clustered (blue dot). Buttons: Start Cluster, Join Cluster, Make Tiebreaker.
- Clustering Components Status:**
  - ES: available (yellow dot)
  - PG: available (yellow dot)
- Cluster Nodes Status:**

Appliance ID	Pulse	Ping	Consul	Tiebreaker	PG Master	Action
FCH1831V0F2	✓	✓	✓	✓	✓	X

A green 'Save' button is located at the bottom right of the Cluster Nodes Status section.

**What to do next**

Proceed to [Join Threat Grid Appliances to Cluster](#).

## Join Threat Grid Appliances to Cluster

This section describes how to join new and existing Threat Grid Appliances to a cluster.



**Note** A Threat Grid Appliance can be joined to an existing cluster only when it contains no data; unlike the initial appliance, which may contain data.

Also, it is critically important that the appliance being joined to the cluster is the latest version. All appliances in a cluster must be running the same version. This may require setting up the appliance first, then update it, reset the data, and join it to the cluster.

Add one node at a time, and wait for Elasticsearch (ES) and Postgres (PG) to reach the state of **Replicated** before adding the next node. The Replicated status is expected in clusters of two or more nodes.



**Note** The wait for the state change for ES and PG to reach **Replicated** does not apply to the single-node case. If you are initializing a single-node cluster from a backup, you should wait for the restore to be completed and the application to be visible in the UI before adding the second node.

When joining an appliance to a cluster, the NFS and clustering must be configured during the initial setup run.

## Join Existing Appliance to Cluster

Perform the following steps to join an existing Threat Grid Appliance to a cluster:

- 
- Step 1** Update the appliance to the latest version. This may require several update cycles depending on the current version on the appliance. All nodes in a cluster must be the same version.
- Step 2** Run the `destroy-data` command in `tgsh` to remove all data; when joining an existing appliance to a cluster, all data must be removed prior to being merged into the cluster. See [Reset Threat Grid Appliance As Backup Restore Target](#).

After running the `destroy-data` command on an existing appliance, it basically becomes a new node, and joining it to a cluster follows the same steps as joining a new appliance.

---

### What to do next

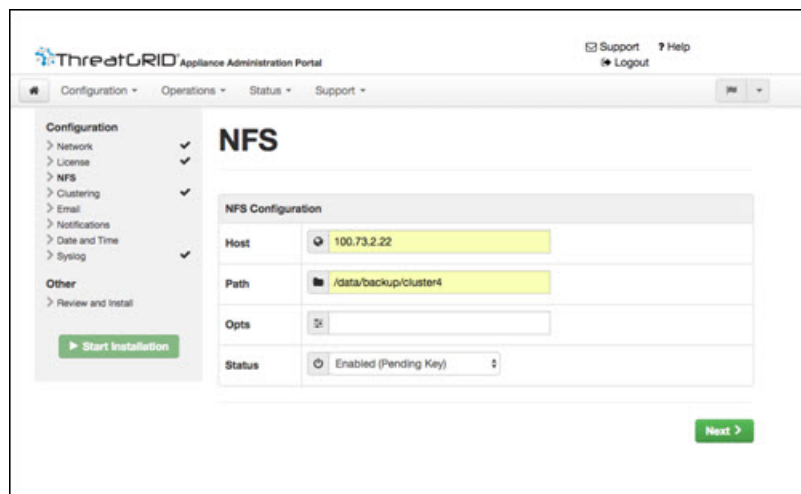
Proceed to [Join New Appliance to Cluster](#).

## Join New Appliance to Cluster

Perform the following steps to join a new Threat Grid Appliance to a cluster:

- 
- Step 1** Set up and begin the OpAdmin configuration as normal.
- Step 2** Browse to the OpAdmin NFS configuration page (**Configuration > NFS**) and specify the **Host** and **Path** to match what was set in the first (initial node in the cluster).
- Step 3** In the **Status** drop-down menu, choose **Enabled (Pending Key)**.

*Figure 14: NFS for Joining a Cluster*

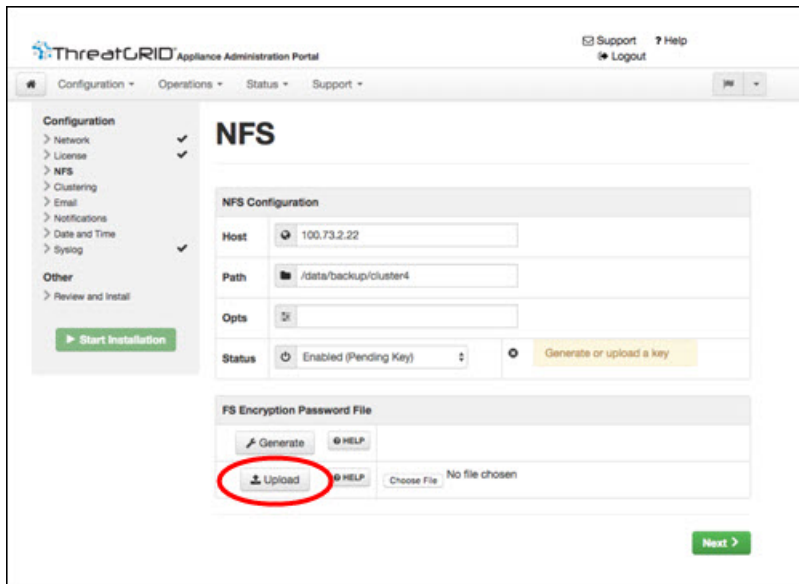




**Step 4** Click **Next**. The page refreshes and Upload becomes available.

**Note** If the key correctly matches the one used to create a backup, the **Key ID** displayed in OpAdmin after upload should match the name of a directory in the configured path. Backups cannot be restored without the encryption key. The configuration process includes the process of mounting the NFS store, mounting the encrypted data, and initializing the appliance's local datastores from the NFS store's contents.

*Figure 15: Upload the NFS Encryption Key*




**Step 5** Click **Upload** and choose the NFS encryption key you downloaded from the first node when you started the new cluster.

**Step 6** Click **Next**.

The page refreshes; the **Key ID** is displayed and the **Activate** button is enabled.


Figure 16: Activate the NFS Encryption Key of the Joining Appliance



**Step 7** Click **Activate**. This **Status** becomes **Active** after a few seconds (lower left corner).

**Step 8** Click **Next** to continue to the Clustering configuration page.

Figure 17: Join Cluster




**Step 9** Click **Join Cluster**, and then click **OK** on the confirmation dialog.

The **Clustering Status** changes to **Clustered**.

**Step 10** Finish the installation. This will initiate a restore of the data in cluster mode.

**Figure 18: Active and Healthy 3-Node Cluster**



**Step 11** Repeat the Step 1 through Step 10 for each node you want to join to the cluster.

## Designate Tiebreaker Node

When a cluster is configured to contain an even number of nodes, the one designated as the tiebreaker gets a second vote in the event of an election to decide which node has the primary database.

Each node in a cluster contains a database, but only the database on the primary node is actually used; the others just have to be able to take over if and when the primary node goes down. Having a tiebreaker can prevent the cluster from being down when exactly half the nodes have failed, but only when the tiebreaker is not among the failed nodes.

We recommend that clusters contain three, five, or seven nodes. Having tiebreaker support is part of an ongoing effort to mitigate the loss of reliability in moving from a standalone appliance to a two-node cluster.

When a cluster is completely healthy and the current node is not the tiebreaker, the **Make Tiebreaker** button is active on the Clustering configuration page.

To designate a node as the tiebreaker, click **Make Tiebreaker**. There will be a brief service disruption, after which the current node will be the one which is not allowed to fail, and the other node can be shut down without breaking the cluster.

In the event of a permanent failure of the tiebreaker node where you are unable to modify the designation ahead of time, either reset the surviving node and restore from backup, or contact [support@threatgrid.com](mailto:support@threatgrid.com) for assistance.

## Remove Cluster Node

To remove an appliance node from a cluster, click the **Remove** icon (X) in the **Action** column in the **Cluster Nodes Status** pane on the Clustering configuration page.

- Removing an appliance from the cluster indicates that it should no longer be considered part of the cluster, rather than a node that is temporarily down. You should remove an appliance when it is being decommissioned; either being replaced with different hardware or will be rejoined to a cluster only after its data has been reset.
- Removing an appliance indicates to the system that you are not going to re-add a node, or if you do re-add it, it will have been reset.
- An appliance is not marked as having been permanently removed from a cluster if it has pulse (is actively writing to NFS), or is active on consul (part of the consensus store).

To replace a still-live node (in a cluster with less than seven nodes), add the new node, wait for the cluster to go green, then remove the old one offline using the **Remove** button. This alerts the system that it's not coming back.

When you first take the node offline, the cluster status changes to yellow. After you click **Remove**, the status reverts back to green (since the cluster will resize such that it no longer expects the now-removed node to be present).

## Resize Cluster

When a node is removed from a cluster using the **Remove** icon, the cluster resizes; this may affect the number of failures it is expected to tolerate. If a cluster is resized in such a way as to change the number of expected failure tolerances (as defined in [Failure Tolerances](#)), it will force an Elasticsearch restart, which will cause a brief service interruption.

**Exception:** This does not include a system other than the PostgreSQL master being rebooted or having a transient failure. Disruption should be minimal in that case except for clients actively using that node, or if samples are running on it. If you add an appliance that was not part of the cluster already, or if you click **Remove**, and this changes the cluster size such that the number of tolerated failures is changed, then there will be a brief interruption as the rest of the cluster reconfigures.

## Failure Tolerances

In the event of a failure, clustered Threat Grid Appliances will not lose any data, with the exception of any analysis being actively run by the failed node, and will recover service with a minimal (less than one minute) service disruption period and no user involvement.

Most failures will recover in less than a minute if the number of available nodes is not smaller than the number shown in the **Nodes Required** column in the Failure Tolerances Table; or will recover after the number of available nodes increases to meet that count. This is true if the cluster was in a healthy state prior to failures (as indicated by services listed as Replicated on the Clustering page).

The number of failures a cluster of a given size is expected to tolerate is shown in the Failure Tolerances Table.

*Table 1: Failure Tolerances Table*

Cluster Size	Failures Tolerated	Nodes Required
1	0	1
2	1*	1*
3	1	2
4	1	3
5	2	3
6	2	4
7	3	4

These figures represent best-case scenarios. If the cluster is not showing green across the board when all nodes are up, then it may not be able to tolerate the full failure count indicated.

For example, if you have a 5-node cluster size with 2 failures tolerated, 3 nodes required, and all 5 appliances are actively processing data, the cluster will be able to reconfigure itself and continue operation without human administrative action if up to 2 failures take place.

Another consideration, in a 5-, 6-, or 7-node cluster, the +1 in the number of failures tolerated means that the percentage of nodes that can fail is higher, which is particularly important because the number of nodes acts as a multiplier to the failure rate. (If you have two nodes, and each has a hardware fault once every 10 years, then you just change your hardware fault rate to once every 5 years.)

## Failure Recovery

Most failures recover automatically. If not, you should contact Threat Grid Support ([support@threatgrid.com](mailto:support@threatgrid.com)), or restore the data from backups. See [Restore Backup Content](#) for more information.

## API/Usage Characteristics

Status of samples submitted to any node in a cluster may be queried from any other node in the cluster; there is no need to track to which individual node a submission took place.

Processing of sample submissions made to one node will be split across all nodes in the cluster; there is no need to actively load-balance from the client side.

## Operational/Administrative Characteristics

In a 2-node cluster, one of the nodes is the tiebreaker, and acts as a single-point-of-failure. However, the other node may be removed from the cluster without ill effect (beyond transient failures during cutover). When a 2-node cluster is healthy (both nodes are fully operational), the tiebreaker designation may be modified by the user, to alter which of the nodes is a single point of failure.

Service may be temporarily disrupted during a failover event; samples which were actively running during a failover will not be automatically rerun.

Inasmuch as capacity is referred to in the context of clustering, this refers to throughput, not storage. A 3-node cluster prunes data to the same maximum storage levels as a single appliance. Consequently, a cluster of three 5000-sample appliances, with a total 15,000-samples/day rate limit, will (when used at full capacity), have retention minimums of 33 percent shorter than the 10,000-sample/day estimates provided in the [Threat Grid Appliance Data Retention Notes](#) on Cisco.com.

## Sample Deletion

Support for deleting samples is available on Threat Grid Appliances (v2.5.0 or later):

- The **Delete** option is available in the **Actions** menu in the samples list.
- The **Delete** button is available in the upper-right corner of the sample analysis report.



---

**Note** It may take up to 24 hours for backup copies of deleted samples to be removed from all nodes.

---

Deleted samples are removed from the shared NFS store immediately; removed from the node processing the deletion request immediately, but the other nodes will lag until the nightly cron job is run. In clustered mode, the NFS store is considered the primary source for samples, so even if the sample is not physically removed from other nodes, it should no longer be retrievable from any of them.