Threat Grid Appliance Clustering Overview

Version: 2.4.2

Created: 3/2/2018

Cisco Systems, Inc. www.cisco.com

CLUSTERING

The ability to cluster multiple Threat Grid appliances was introduced in v2.4.0 for early field trials, and becomes a generally available feature with v2.4.2.

Each appliance in a cluster saves data in the shared file system, and will therefore have the same data as the other notes in the cluster.

Goal

The main goal of clustering is to increase the capacity of a single system by joining several appliances together. into a cluster (v2.4.2 supports clusters with 2 to 7 nodes).

The other goal is to support recovery from failure of one or more machines in the cluster, depending on cluster size.

> **Customer Support:** If you have any questions, we ask that you contact customer support for active involvement when installing or reconfiguring clusters to avoid ***mistakes that could destroy your data***.

Features

- **Shared Data:** Every appliance in a cluster can be used as if they were 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 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 (that is, a cluster in which one or more nodes are not operational) of the next size up.
- **Tiebreaker:** As a special case, 2-node clusters have "tiebreaker" support to make availability depend on a specific, designated node rather than becoming entirely unavailable when either node fails.

Limitations

- When building a cluster of existing standalone appliances, only the 1st node (the initial node) can retain its data. The other nodes will have to be reset (i.e. merging existing data into a cluster is not allowed).
- 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. Please contact support@threatgrid.com if you have any questions.

Requirements

- **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 version available.

- **Clust Interface:** Threat Grid appliance clusters require a direct interconnect on the Clust interface. (See figures below.)

- **Clust Interface Configuration:** The Clust interface does not require any configuration: addresses are automatically assigned, and network topologies where the nodes are not on a single physical network segment are not supported.

- **Direct Interconnect:** All of the Clust interfaces ("nodes") in a cluster must be connected to the same layer-2 segment, with no routing required, for example by putting them into the same VLAN.

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

- **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 may contain data.) Moving an existing appliance into a data-free state requires the use of the database reset process that was added in appliance 2.2.4, NOT the destructive Wipe Appliance process that was added in 1.4.3. (Wipe Appliance will not only remove all data, it will make the appliance inoperable until it's returned to Cisco for reimaging.)

- **SSL Certificates:** If the customer is 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

- Threat Grid appliance clusters require an 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 therefore absolutely essential.

Figure 1 - Clustering Network Diagram

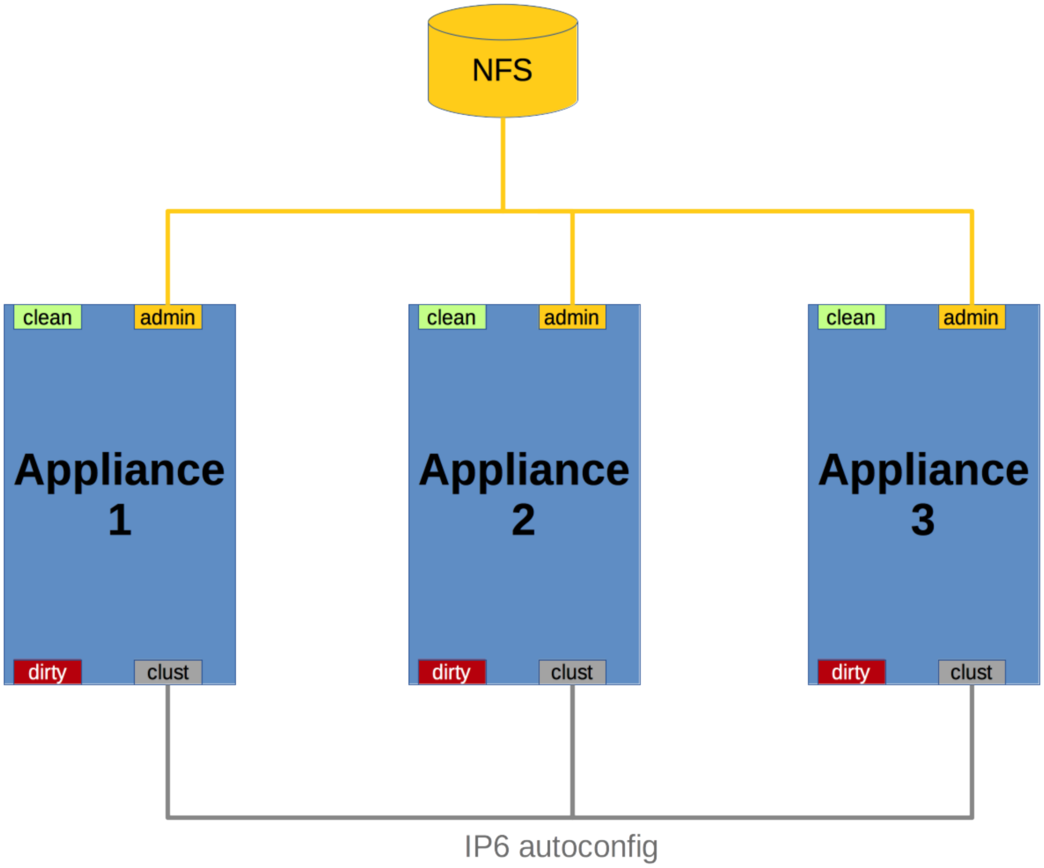

Building a Cluster Overview

Building a cluster in a supported manner requires that all members be on the same version, which should always be the latest available. 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 standalone machines prior, 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 start a new cluster:

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

The Clustering Page

Clusters are configured and managed on the OpAdmin *Clustering* configuration page **(Configuration > Clustering)**. The figure below shows a 3-node, active, healthy cluster.

Figure 2 - The Clustering Page of an Active Cluster

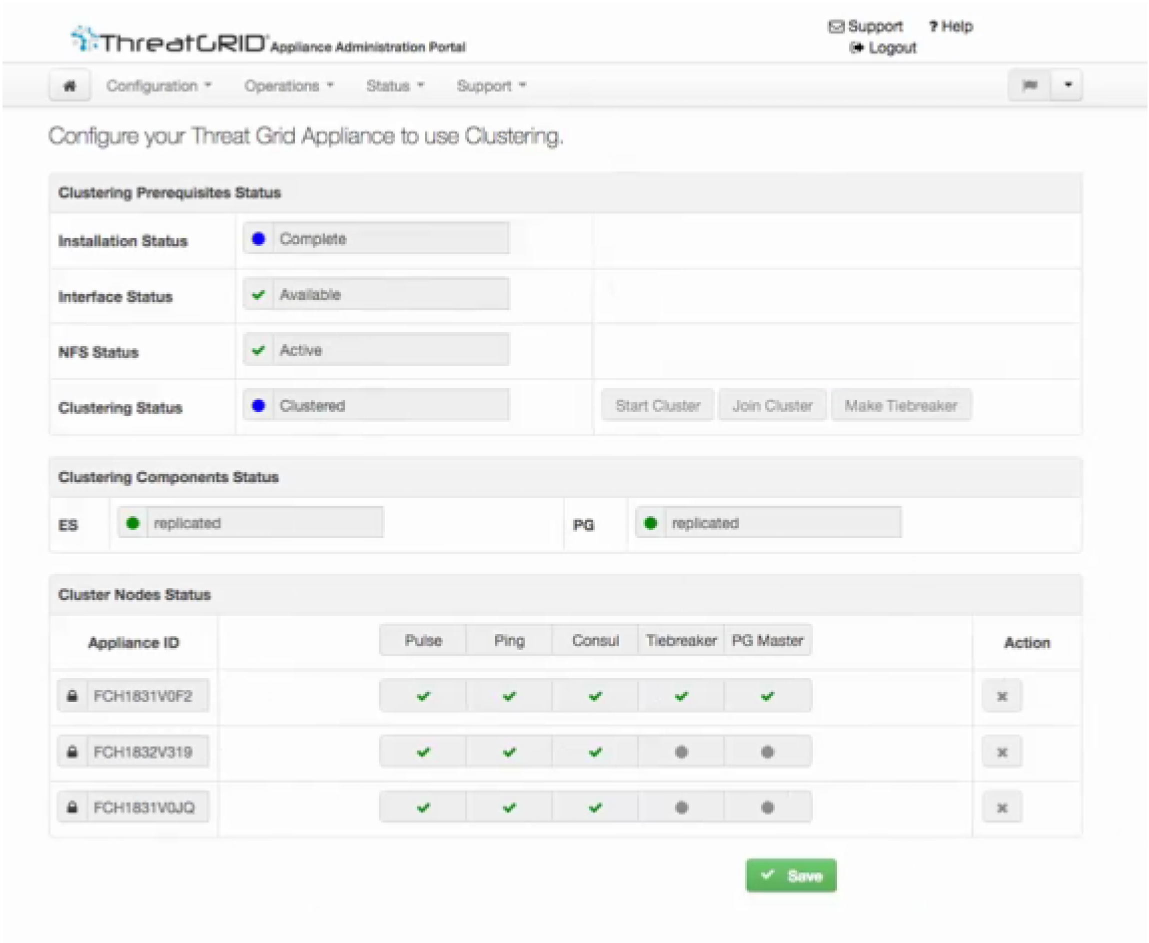

Configure your Threat Grid Appliance to use Clustering.

Clustering Prerequisites Status

Installation Status	● Complete	
Interface Status	✓ Available	
NFS Status	✓ Active	
Clustering Status	● Clustered	Start Cluster · Join Cluster · Make Tiebreaker

Clustering Components Status

ES	● replicated	PG	● replicated

Cluster Nodes Status

Appliance ID	Pulse	Ping	Consul	Tiebreaker	PG Master	Action
FCH1831V0F2	✓	✓	✓	✓	✓	×
FCH1832V319	✓	✓	✓	●	●	×
FCH1831V0JQ	✓	✓	✓	●	●	×

Clustering Status

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.

(See the *Clustering FAQ* published on the *Threat Grid Appliance product documentation page*, for more information.)

-

Status Colors:

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

Cluster Node Status

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.

A green checkmark indicates running and healthy.

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

Tiebreaker - Designates the node as the "tiebreaker".

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.

More Information

For detailed instructions on creating and managing clusters, please see the *Threat Grid Appliance Administrator's Guide v2.4.2* and other documentation available on the *Threat Grid Appliance product documentation page* located on the cisco.com website.