Threat Grid Appliance Clustering Overview



Version: 2.5

Updated: 9/14/2018

Cisco Systems, Inc. www.cisco.com

All contents are Copyright © 2015-2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

CLUSTERING

The ability to cluster multiple Threat Grid appliances was introduced in v2.4.0 for early field trials, and became 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 nodes 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 consisting of 2 - 7 nodes).

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

> **Questions? Contact Customer Support:** If you have any questions, we ask that you please 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 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 (that is, 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.

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 manually reset because merging existing data into a cluster is not allowed. Remove existing data with the previously documented `destroy-data` command. (Do NOT use Wipe Appliance, which will make the appliance inoperable until it's returned to Cisco for reimaging.)

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

- 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:** 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" 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 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.

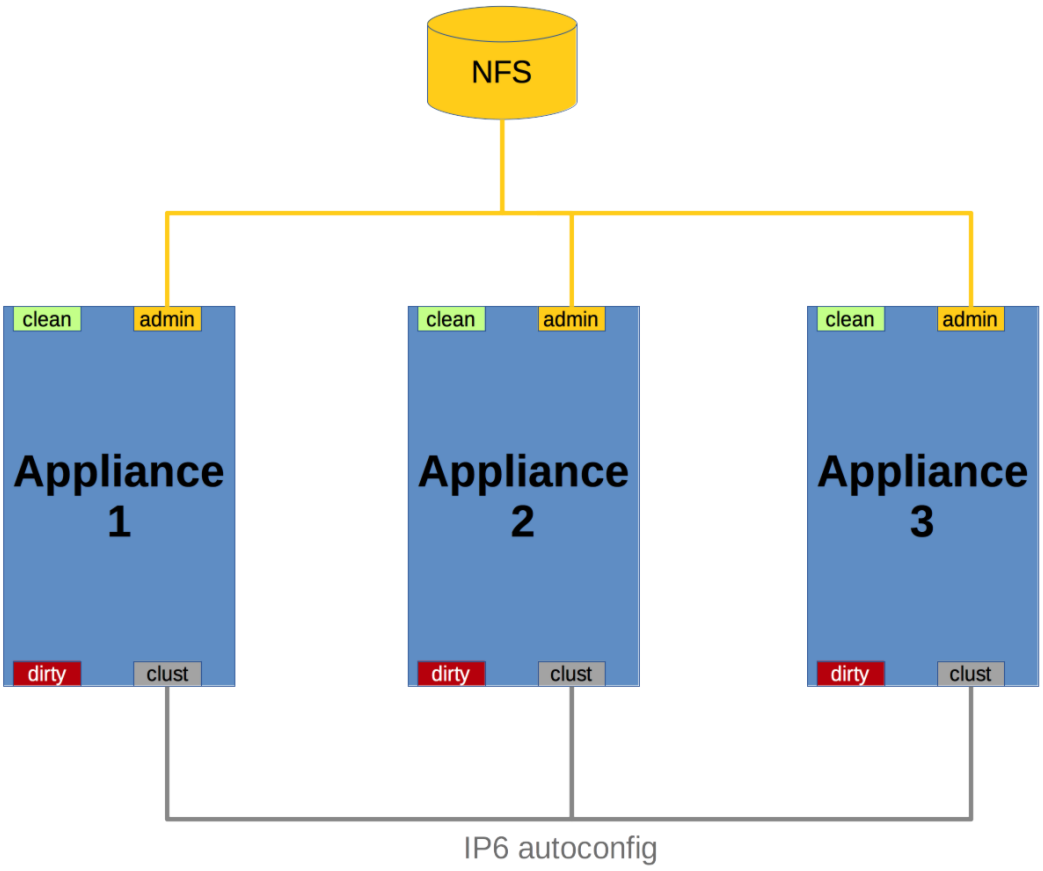
  DO NOT USE the destructive Wipe Appliance process that was added in 1.4.3. (Wipe Appliance will not only remove all data, it will make the appliance inoperable until it's returned to Cisco for reimaging.)
- **SSL Certificates:** If the customer is installing SSL certificates signed by a custom CA on one cluster member, then all other nodes' certificates should be signed by the same CA.

Networking and NFS Storage

- Threat Grid appliance clusters require 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 therefore absolutely essential.

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

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 starting a new cluster:

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

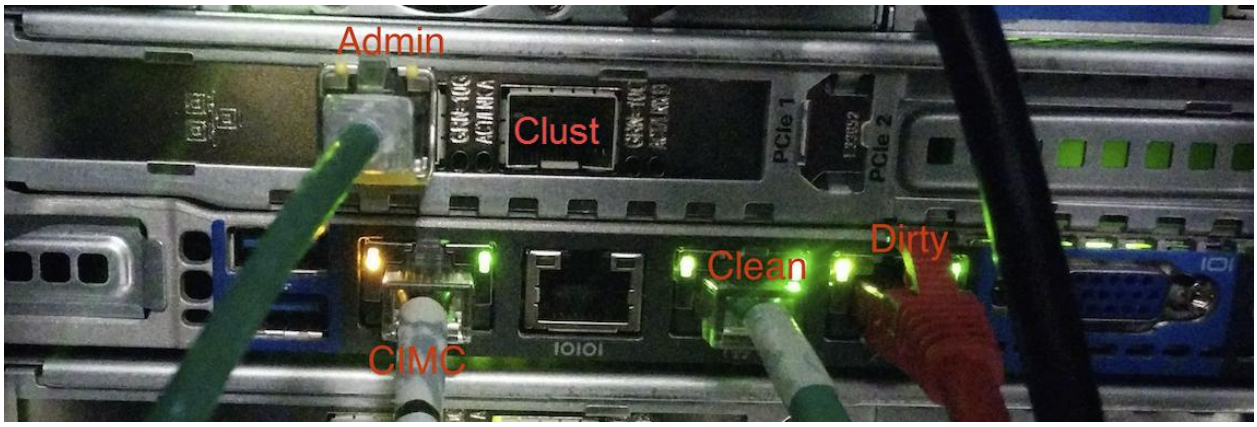
Clust Interface Setup

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

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

Install a SFP+ module in the 4th (non-Admin) SFP port that was previously labeled **Reserved**; it is now used for the **Clust** interface, as illustrated in the next two figures:

Figure 2 - Clust Interface Setup for Cisco UCS M4 C220

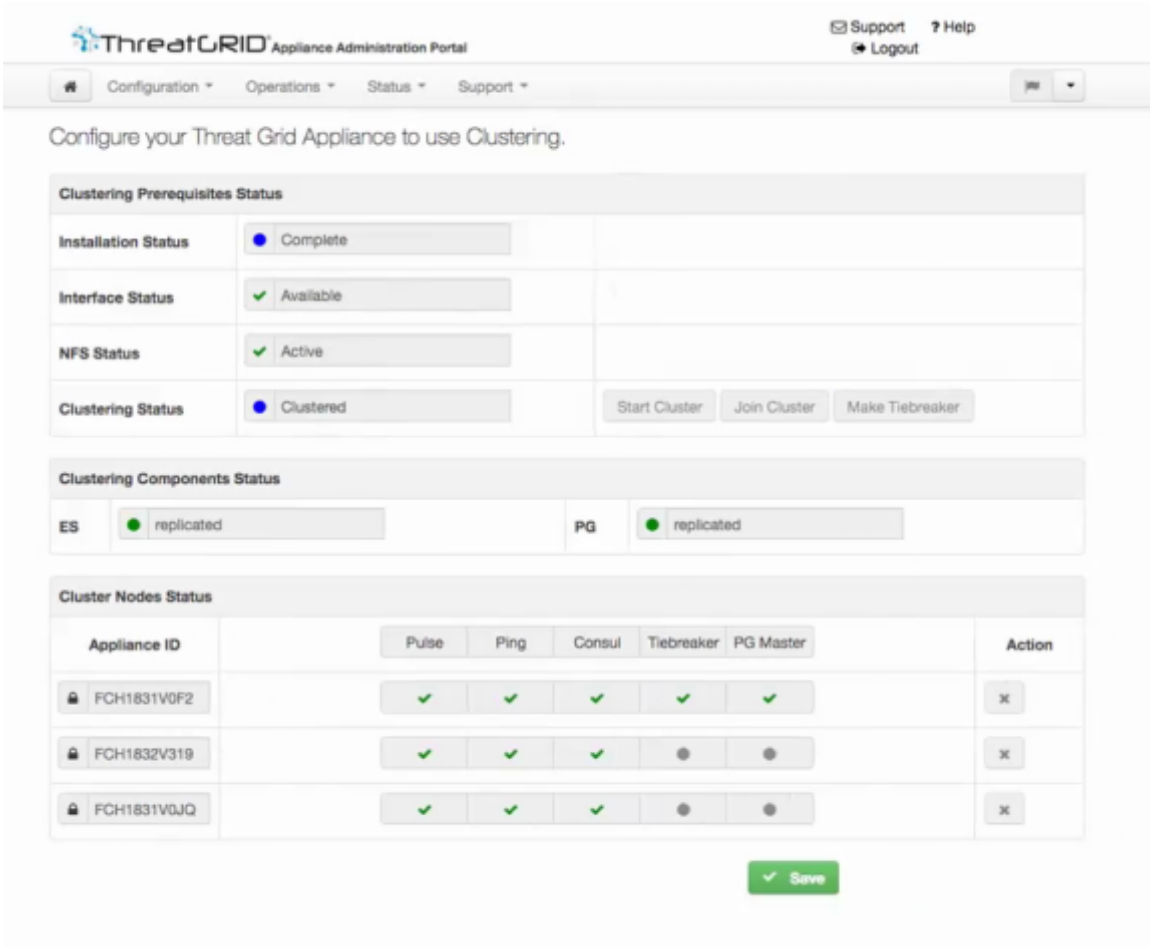


The Clustering Page

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

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

Figure 3 - The Clustering Page of an Active Cluster



More Information

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

All contents are Copyright © 2018 Cisco Systems, Inc. and/or its affiliates. All rights reserved.