Threat Grid Appliance Clustering FAQ

Version: 2.5

Updated: 9/14/2018

Cisco Systems, Inc. www.cisco.com

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

Contents

What is the basic process of creating a cluster?

Answer: 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.

For a complete discussion of how to create and join a cluster, please see the Clustering section in the *Threat Grid Administrator's Guide v2.4.2* on the appliance documentation page on cisco.com.

How many appliances can be joined in a cluster?

Answer: Preferred cluster sizes are 3, 5, or 7 members. 2-, 4-, and 6-member 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.

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. If the tiebreaker node fails in a 2-node cluster, then the cluster is down until that node is recovered.

When joining a cluster, when should the NFS and clustering be configured?

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

Should we use Wipe to remove a node from a cluster?

Answer: Probably not. After performing a wipe operation, an appliance will no longer operate without being returned to Cisco for reimaging. Wipe should only ever be used on a cluster node after that node has been flagged in OpAdmin as permanently removed. Do not remove a node from a cluster, wipe it, and re-add it. Otherwise, if that node ever becomes master after being re-added, undesirable outcomes may result.

Use the **Remove** button in OpAdmin to inform the system that the node is not just inactive but removed.

What is a "tiebreaker"?

Answer: In a 2-node cluster, one node may be designated as the "tiebreaker", which is not allowed to fail, so the other node can be shut down without breaking the cluster. If the tiebreaker node does fail in a 2-node cluster, then the cluster is down until that node is recovered.

We recommend 3-, 5-, or 7-node clusters. Having tiebreaker support is part of an ongoing effort to mitigate the loss of reliability in moving from a standalone appliance to a 2-node cluster.

How do we keep track of which samples were submitted to which appliance in the cluster?

Answer: The 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.

How are custom CAs handled for appliances in a cluster?

Answer: 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.

Do I need a load balancer?

Answer: 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.

An appliance cluster does not present a "virtual IP" -- a single IP address which will be responded to by whichever appliance(s) are currently operational. If such single-endpoint functionality is desired, using a load balancer to provide failover support is suggested. However, programmatic endpoints are encouraged to support falling back between appliances in a cluster without needing to have such a single endpoint; and Cisco ESA appliances explicitly will have this support.

If a system is temporarily down from a cluster, does the cluster still count that node for purposes of the cluster's combined rate limit?

Answer: Only systems that are actively joined to and capable of running samples for the cluster count towards the cluster's rate limit.

Can we split nodes of the same cluster between two different datacenters?

Answer: The Clust interface needs to be physically on the same segment (connected to the same switch, no hops, no extra latency); so no, you can't split across datacenters.

Can cluster nodes be in different racks or different locations?

Can an Admin use a switch as the connection point? That is, from the SFP+ to a switch and then out to the next machine in the cluster and a SFP+?

We were wondering if we can rig things in such a way as to connect in a cluster machines that sit in different racks, or even in a different location if they're on the same LAN. Can systems in the cluster be connected directly to each other using crossover cables as opposed to straight through cables that would be use with a switch?

Answer: The nodes need to be in the same broadcast segment, *period*, such that there are no routing hops. How that's done in physical topology doesn't really matter -- for a 2-node cluster, yes, you *could* use a crossover cable.

For clusters with 3 or more nodes, yes, you could have multiple switches, so long as they're bridged together and not routed (and their connection doesn't involve latency, as would be the case if there were something like a routed VPN). The cluster nodes can't be kept in different locations because of latency, even if it's bridged.

Why isn't Pulse showing green?

Answer: a system won't show up green for pulse when the NFS Activate button is first pressed but only after it actually finishes configuring itself to use that datastore. Pulse won't be green until the you finish the wizard and launch the configuration process

What is the difference between "clustered" and "replicated"?

Answer: "Clustered" indicates that a node is not a standalone. "Replicated" indicates that the data involved in a storage service has been backed up such that the number of failures indicated in the table can be tolerated, and no failures are presently ongoing. (If there are currently either ongoing failures, or data is not yet distributed across all node, but failures are not so severe as to cause a service outage, then status will be "Available" rather than "Replicated").

Service status will be "Replicated" when all nodes in the cluster are healthy, and sufficient backup copies of data in the cluster have been distributed to permit the number of failures described in the table. If backup copies have not yet been distributed, or some of the nodes have failed, then status will be "Available".

How are rate limits calculated for clustered appliances?

Answer: The submission rate limits of each cluster member are added up to become the cluster's limit.

What happens to the cluster rate limit if a node fails?

Answer: Appliances that have failed do not contribute towards the combined rate limit. For example, if a 5-node cluster is healthy, the rate limit will is the combination of all 5 nodes; but if two nodes have failed, your rate limit will be that of the 3 remaining nodes combined.

Do cluster nodes analyze samples, or are they on standby?

Answer: All systems joined to a cluster actively analyze submitted samples and assist in running searches.

What is the meaning of the Elasticsearch green or yellow cluster status?

Answer:

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

An ES index is considered "green" if all shards in it are green; the ES service as a whole is "green" only if all indexes are green.

Similarly, you can go from green to yellow if you add a new node but haven't had a chance to copy data out to it yet, because adding a new node can increase the number of copies that are expected to exist for each piece of data.

Different nodes store the master and read-only copies for each shard, to spread out the workload. For example, in a 5-node cluster: for each Elasticsearch shard, one node has the master copy and two nodes keep read-only copies, which they also use to help answer queries.

So each shard has a replication status. Therefore, if you lose a node and you see Elasticsearch change from green to yellow in the Clustering page, it means that you lost backups of some shards, and there hasn't yet been a chance to get another system to take over the role of hosting those backups.

The red ES status will only occur if at least one shard has no copy at all on a live node that can take over the master role, and thus the data in that shard is unsearchable; in that case, you need to either 1) bring up any dead nodes that previously hosted that content; 2) contact customer support (who can potentially restore the specific indexes impacted by the problem from backup, or discard them if the data is

old content the customer doesn't care about; or 3) rebuild the cluster from the backup on NFS. (The latter should only occur if the documented constraints around the number of required nodes are broken.)

If you go from 4 nodes to 5, for example, then instead of one read-only backup copy of each piece of data, there are now expected to be 2. Until there are 2 read-only backups of each piece of data, a cluster that is recently resized from 4 to 5 nodes will be have a status of yellow.

Note: The same color status indicators apply to the PostgreSQL service.

What are the failure tolerances?

In the event of a failure, clustered 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 so long as the number of nodes that are available is not smaller than the number in the **Nodes Required** column; or will recover after the number of available nodes increases to meet that count; so long as the cluster was in a healthy state prior to those failures (as indicated by services listed as "Replicated" in the *Clustering* page).

The number of failures a cluster of a given size is expected to tolerate:

Figure 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

*Non-Tiebreaker Only

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 and 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.

Something else to consider: 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 as 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 changed your hardware fault rate to once every 5 years).

What is the recommended storage size of NFS for clustering?

The ideal is 6TB. This number provides headroom over the maximum possible storage size (which is limited by the amount of storage on an appliance). The number does not depend on the number of appliances in the cluster.

Can samples be deleted from a cluster?

Yes, starting with the 2.5.0 release, which updates the portal software to 3.5.11, it is possible to delete samples from a cluster.

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

Don't see your question listed? Contact support@threatgrid.com for additional assistance.