Threat Grid Appliance Clustering Overview and FAQ

Version: 2.4

Created: 12/14/2017

Cisco Systems, Inc. www.cisco.com

All contents are Copyright © 2015-2017 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

Contents

Introduction

The Threat Grid appliance release 2.4 introduces pre-release support for clustering multiple appliances. This requires additional hardware, and is presently only available for customers who are eligible for, and have opted into, an early field trial.

The main goal of clustering is to increase the capacity of a single system by joining several appliances into a cluster. This initial release supports joining clusters of 2-3 appliances. The ability to support more (up to 5) is expected as part of the next production release, which is scheduled for February, 2018.

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

Technical Overview

NFS Requirements

- Threat Grid appliance clusters require an NFS store be enabled and configured: it must be available via the admin interface, 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.

Other Requirements

- **Clust Interface:** Threat Grid appliance clusters require a direct interconnect on the Clust interface. (See figures below.)
- 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, for example by putting them into the same VLAN.
- **Hardware:** Enabling the direct interconnect requires an additional SFP+ module installed in each appliance, in order to connect the "Clust" interface to a switch.
- **Airgapped 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 added to a cluster when it contains no data. (Only the initial node may contain data.)
- **License:** Prior to the February release of clustering as a fully-supported feature, use of clustering functionality requires that all appliances participating in a cluster have a license installed with a flag indicating that the customer is participating in an early field trial.
- **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.

Setup and Configuration

Installations that meet all of the following conditions for the 2.4.0 release are configured as clustered:

1. Install a license that indicates the appliance in question is enrolled in the clustering EFT.
2. Install a SFP+ module in the 4th (non-Admin) SFP port that was previously labeled **Reserved**; it is now used for the **Clust** interface. Each appliance in the cluster requires an additional SFP for the Clust interface.

Figure 1 - Clust Interface - Cisco UCS M3 C220

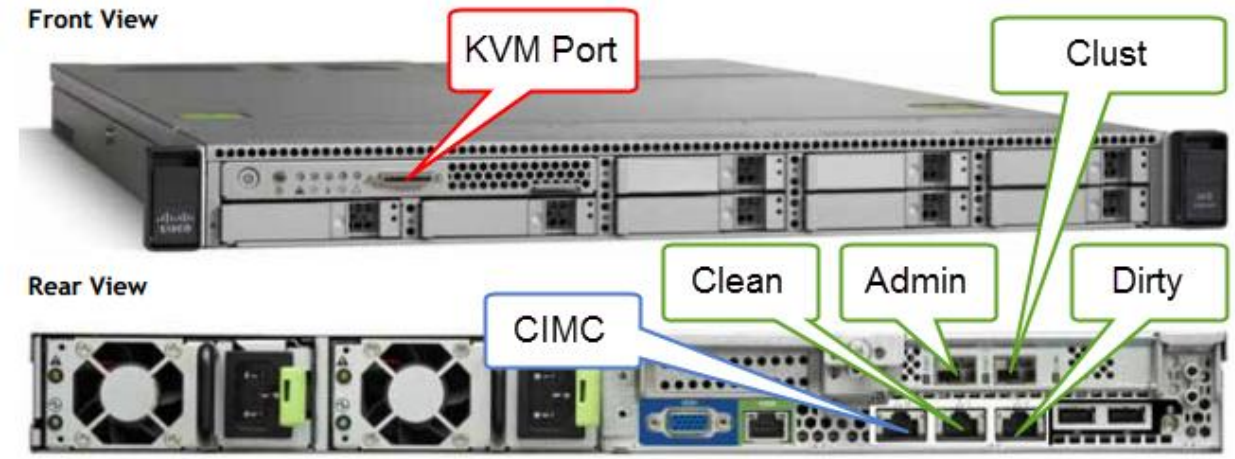

Figure 24 - Clust Interface Setup for Cisco UCS M4 C220

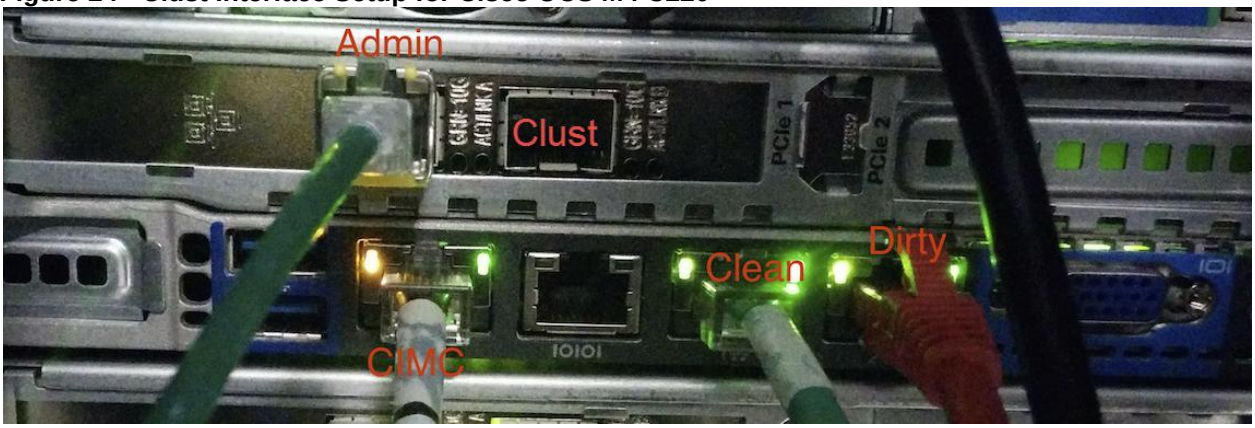

3. Enable and configure a single NFS store and generate the key.
   - If the NFS store is empty, installing the 2.4 updates will create a new cluster.
   - If the NFS store is concurrently used by other appliances sharing the *clust* network segment with a given key, installing the 2.4 release will join to that existing cluster.

Number of Appliances

QA coverage presently covers clusters of 2 nodes (capacity only, *not* HA); or 3 nodes (for both HA and capacity purposes). 5-node clusters may be supported as of the February release.

Creating a New Cluster

There are several ways to create a new cluster:

- Start "from scratch" on a new appliance.
- A new cluster may be created by restoring from a backup that is not being used concurrently as a backup by any preexisting appliance or cluster.
- Start a new cluster with data from a single preexisting appliance.

Adding Cluster Nodes

An appliance may be added to an existing cluster *only when it contains no data*. (Unlike the starting ("pre-existing") appliance, which may contain data.)

Moving an existing appliance into a data-free state requires the use of the database reset process added in appliance 2.2.4, NOT the destructive wipe process available since 1.4.3.

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 “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 3 5000-sample appliances – with a total 15,000-samples/day rate limit – will, when used at full capacity, have retention minimums 33% shorter than the 10,000-sample/day estimates provided in the **Threat Grid Appliance Data Retention Notes**, located with other appliance documentation on cisco.com.

Clustering Frequently Asked Questions

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 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.

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.