

Cisco Threat Grid Appliance Administrator Guide Version 2.7

First Published: 2019-06-01

Last Modified: 2019-08-08

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

?

CHAPTER 1 **Introduction** 1

About Cisco Threat Grid Appliance 1

What's New In This Release 1

Audience 2

Product Documentation 3

Threat Grid Support 3

Enable Support Mode 5

Support Snapshots 6

CHAPTER 2 **Administration** 7

Login Names and Passwords (Default) 7

Reset Administrator Password 7

Install Updates 9

Version Lookup Table 10

Updates Port 10

Troubleshoot Updates 10

Database Schema Updates 10

CHAPTER 3 **Configuration Management** 13

Introduction 13

Network Configuration Using TGSH Dialog 13

Configure Network Using TGSH Dialog 14

Reconnect to TGSH Dialog 14

Configure Networking in Recovery Mode 14

Configuration Using OpAdmin Portal 15
Configure SSH Keys 16
Configure Remote Syslog Server for Notifications 17
Configure LDAP Authentication 17
Configure LDAP Authentication in OpAdmin 18
Configure Third-party Detection and Enrichment Services 19
Apply Configuration Change 20
Using DHCP 21
Explicit DNS for DHCP 21
Network Configuration and DHCP 22

CHAPTER 4 **Manage SSL Certificates** 25
About SSL Certificates and Threat Grid Appliance 25
Configure SSL Certificates for Inbound Connections 26
Validate Common Name in SSL Certificate 27
Replace SSL Certificate 27
Regenerate SSL Certificate 28
Download SSL Certificate 28
Upload SSL Certificate 29
Generate SSL Certificate Using OpenSSL 29
Configure SSL Certificates for Outbound Connections 30
Configure DNS 30
Manage CA Certificates 31
Manage Disposition Update Syndication Service 31
Connect ESA/WSA to Threat Grid Appliance 32
ESA and WSA Integration Process Overview 33
ESA/WSA Integration Process Steps 34
Connect AMP for Endpoints Private Cloud to Threat Grid Appliance 35

CHAPTER 5 **Manage Organizations and Users** 37
Introduction 37
Create New Organization 37
Manage Users 38
Activate New Device User Account 38

CHAPTER 6 **Privacy and Sample Visibility** 41

About Privacy and Sample Visibility 41

Privacy and Visibility for Integrations 41

CHAPTER 7 **Wipe Appliance** 43

About Wipe Appliance Option 43

Wipe Appliance Procedure 43

Wipe and Clusters 45

CHAPTER 8 **Backups** 47

Threat Grid Appliance Backups 47

NFS Requirements 48

File System 48

Backup Storage Requirements 48

Backup Expectations 49

Backup Data Retention 50

Strictly Enforce Retention Period Limits 50

Backup Process Overview 50

Backup Frequency 51

Reset Threat Grid Appliance as Backup Restore Target 51

Data Reset Process 51

Return Target Appliance to Preconfigured State 52

Another Appliance Target Actively Writing to Backup Being Restored 53

Restore Backup Content 53

Backup and Restore Notes 54

Backup-Related Service Notices 54

CHAPTER 9 **Clustering** 57

About Clustering Threat Grid Appliances 57

Clustering Features 57

Clustering Limitations 58

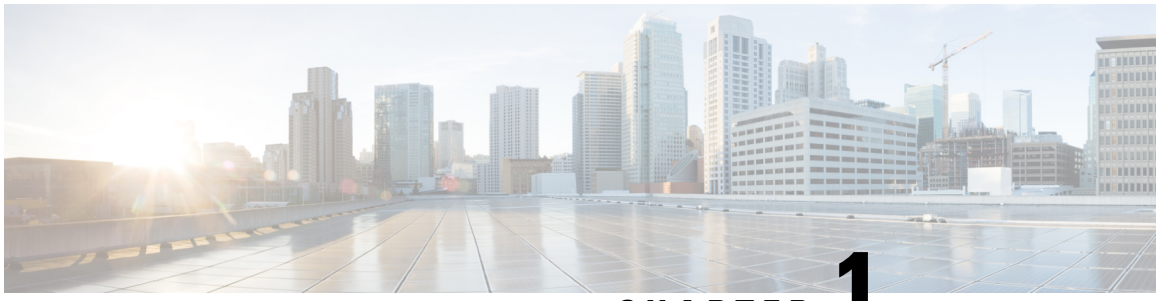
Clustering Requirements 58

Networking and NFS Storage 59

Build Cluster Overview 60
Clust Interface Setup 60
Clustering Configuration 61
Start Cluster of Thread Grid Appliances 63
Start Cluster with Existing Standalone Appliance 63
Start Cluster with New Appliance 69
Join Threat Grid Appliances to Cluster 71
Join Existing Appliance to Cluster 72
Join New Appliance to Cluster 72
Designate Tiebreaker Node 75
Remove Cluster Node 76
Resize Cluster 76
Failure Tolerances 77
Failure Recovery 77
API/Usage Characteristics 78
Operational/Adminstrative Characteristics 78
Sample Deletion 78

CHAPTER 10 **Network Exit Configuration** 79
Configure Network Exit 79

CHAPTER 11 **OpAdmin Menus** 83
Configuration Menu 83
Operations Menu 84
Status Menu 84
Support Menu 85



CHAPTER 1

Introduction

This chapter provide a brief description of the Cisco Threat Grid Appliance, the intended audience and how to access relevant product documentation. It includes the following:

- About Cisco Threat Grid Appliance, on page 1
- What's New In This Release, on page 1
- Audience, on page 2
- Product Documentation, on page 3
- Threat Grid Support, on page 3

About Cisco Threat Grid Appliance

The Cisco Threat Grid appliance provides safe and highly secure on-premises advanced malware analysis, with deep threat analytics and content. A Threat Grid Appliance provides the complete Threat Grid malware analysis platform, installed on a single UCS server: Cisco UCS C220-M3 (TG5000) or Cisco UCS C220 M4 (TG5400). It empowers organizations operating under various compliance and policy restrictions, to submit malware samples to the appliance.

Many organizations that handle sensitive data, such as banks and health services, must follow various regulatory rules and guidelines that do not allow certain types of files, such as malware artifacts, to be sent outside of the network for malware analysis. By maintaining a Cisco Threat Grid Appliance on-premises, organizations are able to send suspicious documents and files to it to be analyzed without leaving the network.

With a Threat Grid Appliance, security teams can analyze all samples using proprietary and highly secure static and dynamic analysis techniques. The appliance correlates the analysis results with hundreds of millions of previously analyzed malware artifacts, to provide a global view of malware attacks and campaigns, and their distributions. A single sample of observed activity and characteristics can quickly be correlated against millions of other samples to fully understand its behaviors within an historical and global context. This ability helps security teams to effectively defend the organization against threats and attacks from advanced malware.

What's New In This Release

The following changes have been implemented in this guide in Version 2.7:

Table 1: Changes in Version 2.7.2

Feature or Update	Section
Background Elasticsearch index migration to ES6-native indexes enabled.	Database Schema Updates
Database base backups are only retained until a new base backup has been successfully created.	Backup Data Retention

Table 2: Changes in Version 2.7.1

Feature or Update	Section
Simulated Only option for Network Exit Mode	Configure Network Exit
Network Simulation option for sample analysis.	See note in Configure Network Exit

Table 3: Changes in Version 2.7

Feature or Update	Section
Network configuration in recovery mode now mirrow the full system.	Configure Networking in Recovery Mode
Threat Grid Appliance serial number is now used as hostname.	Configuration Using OpAdmin Portal
Configuring SSH public key for access to the Threat Grid Appliance disables password-based authentication via SSH.	Configure SSH Keys
TLS v1.0 and v1.1 disabled in Admin interface.	About SSL Certificates and Threat Grid Appliance
XFS is primary file system.	File System
Enhanced data reset process.	Data Reset Process
Sample deletion now includes artifacts, which matches the behavior of the cloud product.	Sample Deletion

Audience

This guide is intended to be used by the Threat Grid Appliance administrator after the appliance has been set up and configured, and an initial test malware sample has been successfully submitted and analyzed. It describes how to manage organizations and users for the Threat Grid malware analysis tool, appliance updates, backups, and other server administration tasks.

This guide also provides information for administrators who are integrating the Threat Grid Appliance with other Cisco products and services, such as Cisco Email Security Appliance, Cisco Web Security Appliance, and AMP for Endpoints Private Cloud devices.

Note For information about Secure Malware Analytics Appliance setup and configuration, see the *Cisco Threat Grid Appliance Getting Started Guide*.

Product Documentation

The latest versions of Cisco Threat Grid Appliance product documentation can be found on Cisco.com:

- Cisco Threat Grid Appliance Release Notes
- *Cisco Threat Grid Version Lookup Table*

Prior version of Cisco Threat Grid Appliance product documentation can be found on the Threat Grid Install and Upgrade page.

Threat Grid Portal UI Online Help

Threat Grid Portal user documentation, including *Release Notes*, *Using Threat Grid* Online Help, API documentation, and other information is available from the **Help** menu located in the navigation bar at the top of the user interface.

Email Security Appliance and Web Security Appliance Documentation

For information on connecting an Email Security Appliance or Web Security Appliance, see Connecting ESA/WSA Appliances to a Threat Grid Appliance in this guide.

See the instructions for *Enabling and Configuring File Reputation and Analysis Services* in the online help or user guide for your ESA/WSA:

- *Cisco Secure Email Gateway User Guide*
- *Cisco Secure Web Appliance User Guide*

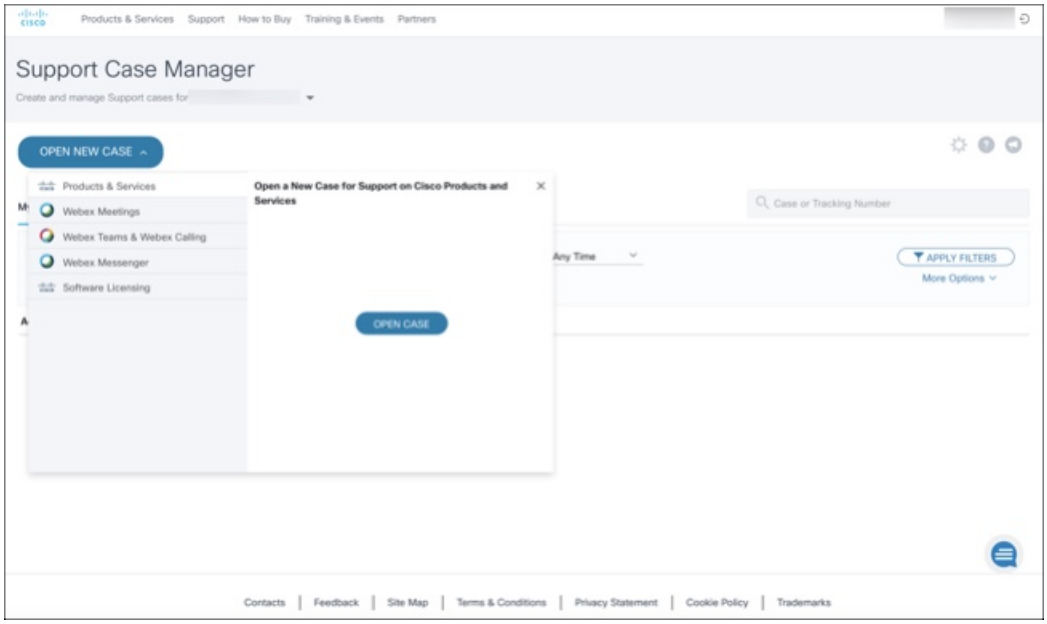
Threat Grid Support

If you have questions or require assistance with Threat Grid, open a Support Case at https://mycase.cloudapps.cisco.com/case.

Procedure

Step 1 In Support Case Manager, click **Open New Case** > **Open Case**.

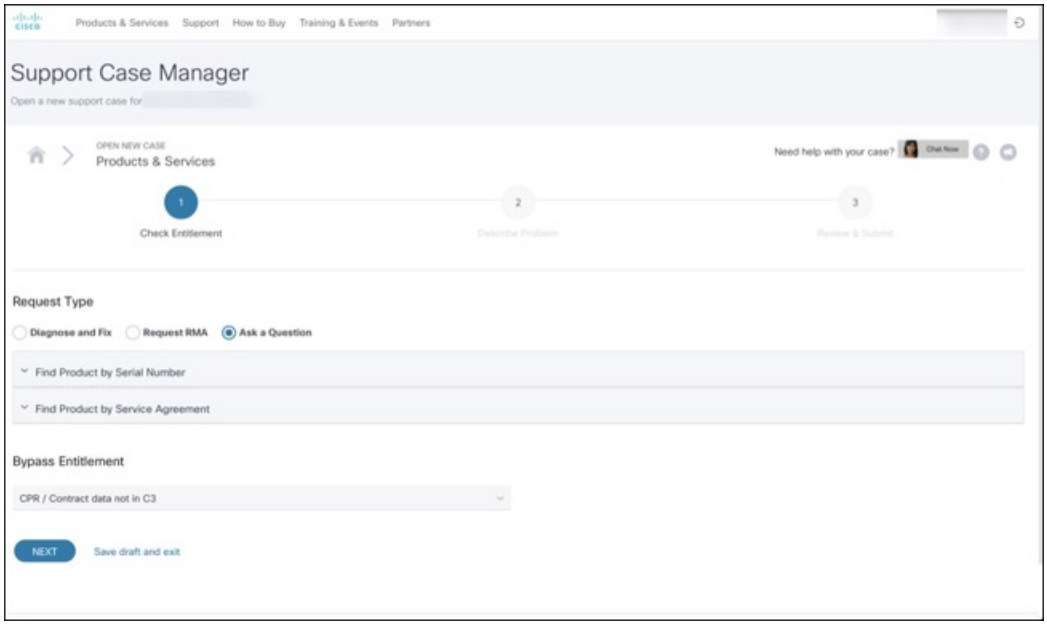
Figure 1: Open New Case



Step 2 Click the **Ask a Question** radio button and search for your Cisco Security **Product Serial Number** or **Product Service Contract**. This should be the serial number or service contract for Threat Grid.

Step 3 If you want to bypass entitlement, choose **Contract Data not in C3** and click **Next**.

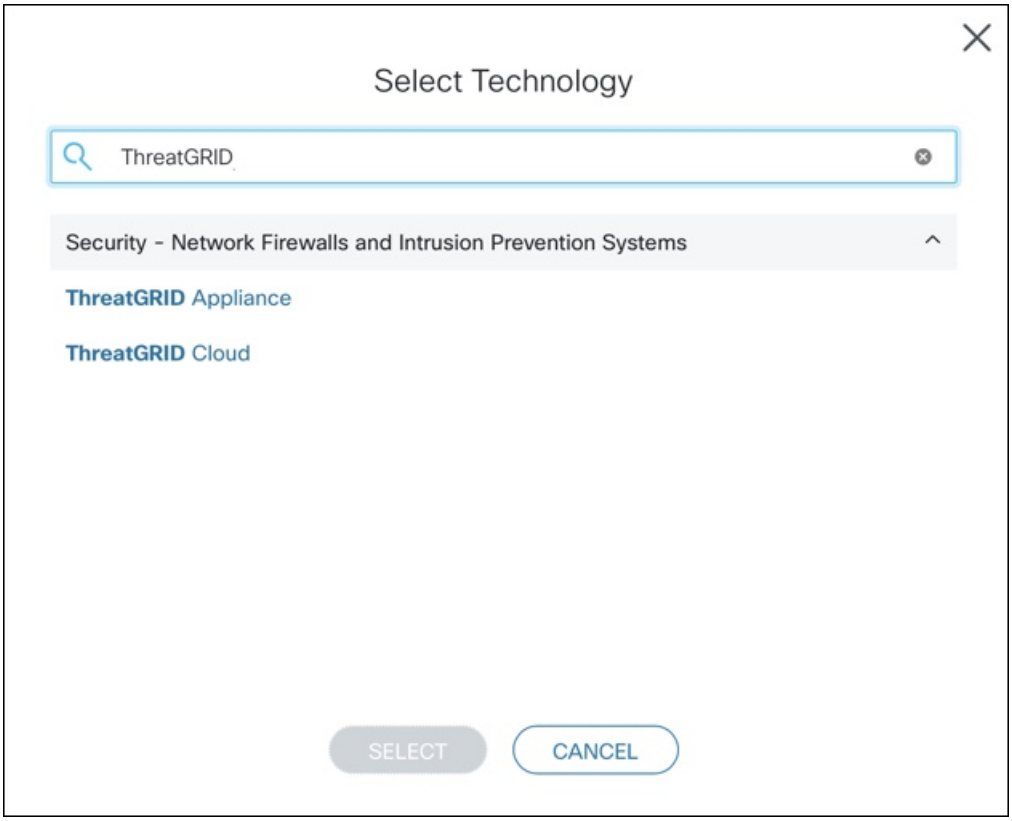
Figure 2: Check Entitlement



Step 4 On the **Describe Problem** page, enter a **Title** and **Description** of the problem (mention Threat Grid in the title).

Step 5 Click **Manually select a Technology** and search for **ThreatGRID**.

Figure 3: Select Technology



Step 6 Choose **ThreatGRID Appliance** from the list and click **Select**.

Step 7 Complete the remainder of the form and click **Submit**.

If you are unable to open a case online, contact Cisco Support:

- **US and Canada**: 1-800-553-2447
- **Worldwide Contacts**: https://www.cisco.com/c/en/us/support/web/tsd-cisco-worldwide-contacts.html

For additional information on how to request support:

- See the blog post: **Changes to the Cisco Threat Grid Support Experience** at https://community.cisco.com/t5/security-blogs/changes-to-the-cisco-threat-grid-support-experience/ba-p/3911407
- See the main **Cisco Support & Downloads** page at: https://www.cisco.com/c/en/us/support/index.html

Enable Support Mode

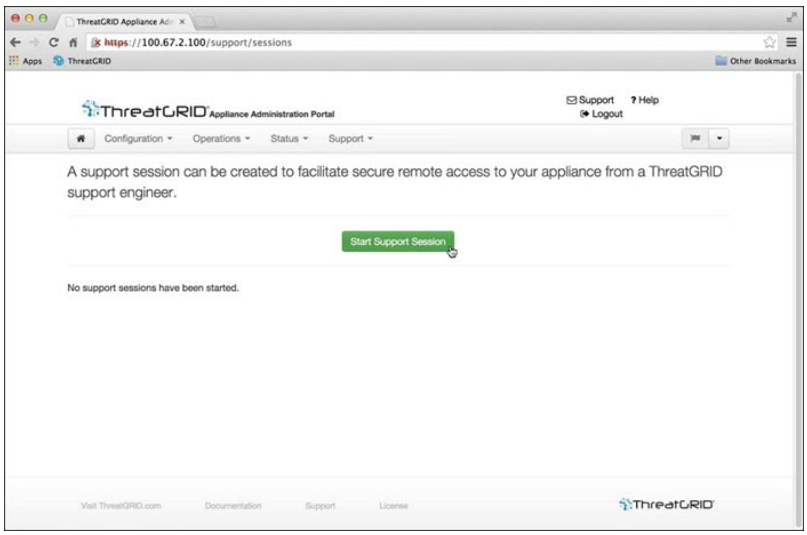
If you require support from a Threat Grid engineer, they may ask you to enable Support Mode, which is a live support session that gives Threat Grid support engineers remote access to the appliance. Normal operations of the appliance will not be affected.

You can enable Support Mode from the OpAdmin portal **Support** menu. You can also enable it from the TGSH Dialog, the legacy Face Portal UI, and when booting up in Recovery Mode.

Procedure

Step 1 In the OpAdmin portal, click the **Support** menu and choose **Live Support Session**.

Figure 4: OpAdmin Start a Live Support Session



Step 2 Click **Start Support Session**.

Note
You can exit the OpAdmin configuration wizard to enable Support Mode prior to licensing.

Support Snapshots

A support snapshot is basically a snapshot of the running system, which contains logs, psoutput, etc., to help Support staff troubleshoot any issues.

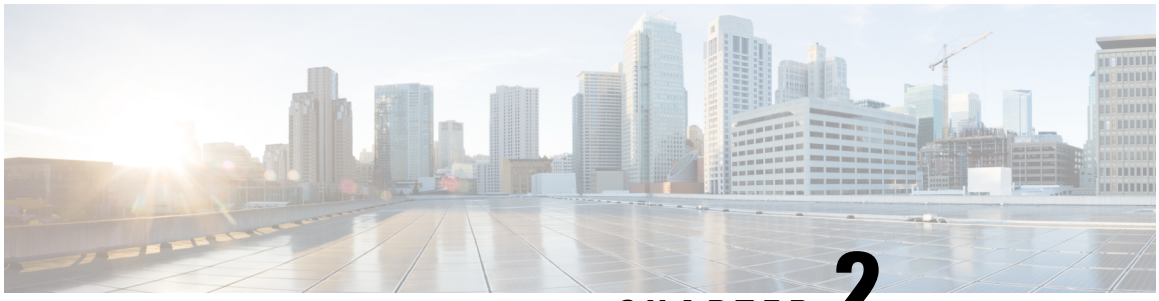
Procedure

Step 1 Verify that SSH is specified for Support Snapshot services.

Step 2 From the **Support** menu, choose **Support Snapshots**.

Step 3 Take the snapshot.

Step 4 Once you take the snapshot, download it as a **.tar** or **.gz** file, or click **Submit**, to automatically upload the snapshot to the Threat Grid snapshot server.



CHAPTER 2

Administration

This chapter provides general information that is useful for Administrators. It includes the following topics:

- Login Names and Passwords (Default), on page 7
- Reset Administrator Password, on page 7
- Install Updates, on page 9

Login Names and Passwords (Default)

The default login names and passwords are listed in the following table:

User	Login/Password
OpAdmin and Shell User	Use the initial Threat Grid/TGSH Dialog randomly generated password, and then the new password entered during the first step of the OpAdmin configuration workflow. If you lose the password, follow the instructions in Reset Administrator Password, on page 7.
Threat Grid Web portal UI Administrator	Login: **admin** Password: Initialize with the first OpAdmin password, and then it becomes independent.
CIMC	Login: **admin** Password: **password**

Reset Administrator Password

The default administrator password is only visible in the TGSH Dialog during the initial appliance setup and configuration. Once the initial configuration is completed, the password is no longer displayed in visible text.

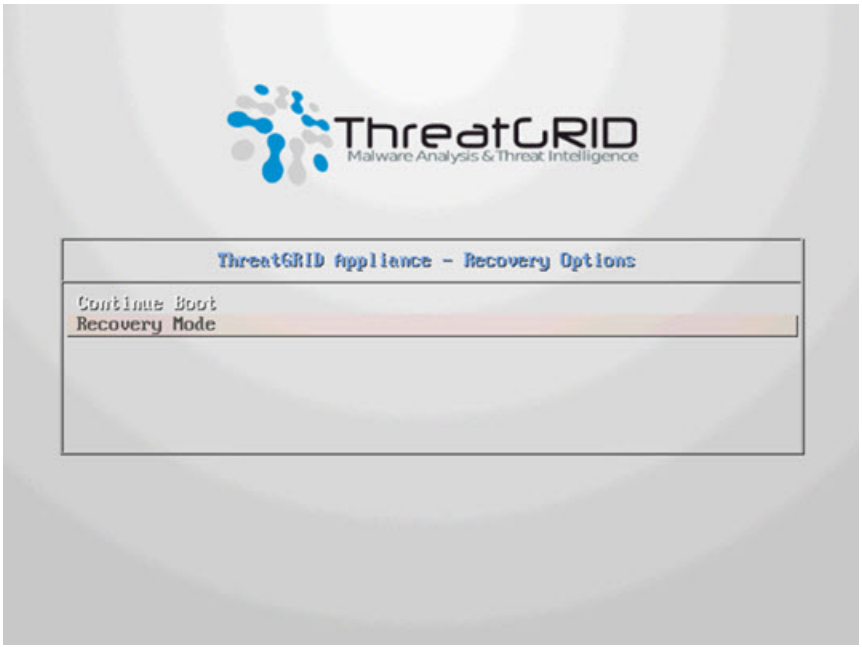
Note LDAP authentication is available for TGSH Dialog and OpAdmin login when you have multiple administrators. If the appliance is configured for LDAP authentication only, resetting the password in recovery mode will reconfigure the authentication mode to allow login with system password as well.

If you lose the administrator password and are unable to login to OpAdmin, complete the following steps to reset the password.

Procedure

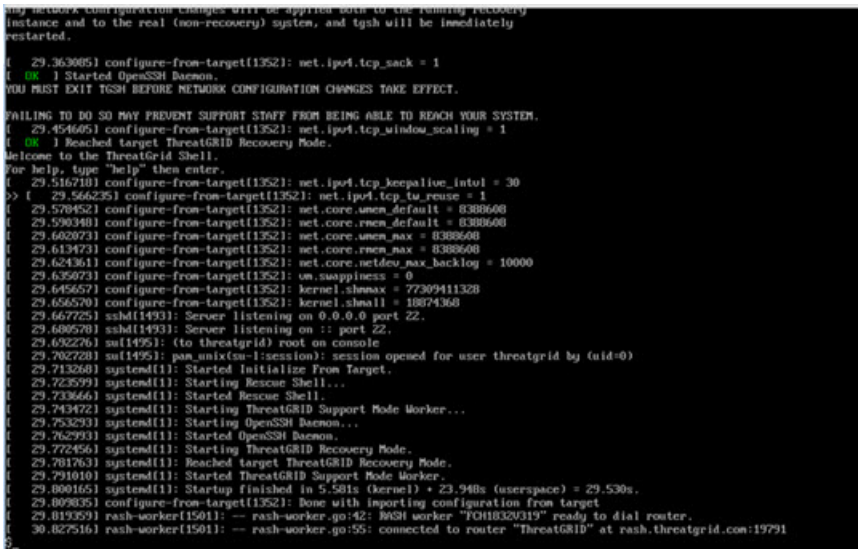
Step 1 Reboot the Threat Grid Appliance and immediately select **Recovery Mode** from the Recovery Options.

Figure 5: Boot Menu - Recovery Mode



The Threat Grid Shell opens.

Figure 6: Threat Grid Shell in Recovery Mode



Step 2 Run `passwd` to change the password.

Figure 7: Enter New Password



Note
The command prompt is not always visible in this mode and logging output may be displayed at any point on top of your input. This does not affect input; you can keep typing *blindly*. Ignore the two lines of logging output.

Step 3 Enter (blindly) the password and press **Enter**.

Step 4 Re-type the password and press **Enter**.

Note
The password will not be displayed.

Step 5 Type **exit** from the command line to save.

Important
You must exit before rebooting to save the new password. If you do not exit, although everything appears to be OK, the password change will be quietly discarded.

Step 6 Type **reboot** and press **Enter** to start the appliance in normal mode.

Install Updates

Before you can update the Threat Grid Appliance with newer versions, you must have completed the initial setup and configuration steps as described in the Cisco Threat Grid Appliance Setup and Configuration Guide.



Note If you have a new appliance that shipped with an older version of software and want to install updates, you must first complete the initial configuration. Threat Grid Appliance updates will not download unless the license is installed, and may not apply correctly if the Appliance has not been fully configured, including the database.

The following considerations should be observed when installing updates:

- Threat Grid Appliance updates are applied through the OpAdmin Portal.
- If the update server sends an update, the client moves all the way forward to that version. It's not always possible to skip interim releases; when not possible, the update server will require the appliance to install the release before it can download the next update.
- If the server allows you to download a version, you are eligible to move to that version directly; that is, with no intervening reboots beyond those needed for a single upgrade.
- Updates are one-directional: you cannot revert to a previous version after you upgrade to a more recent version.

For instructions on installing updates, see the Install Updates section in the Cisco Threat Grid Appliance Setup and Configuration Guide.

Version Lookup Table

To identify the correct build number and corresponding release version, see the Cisco Threat Grid Appliance Version Lookup Table.

Updates Port

The Threat Grid Appliance downloads release updates over SSH, port 22.

- Release updates can also be applied from the textual (curses) interface (Threat Grid Appliance v1.1 and later), not just from the web-based administrative interface (OpAdmin).
- Systems using DHCP need to explicitly specify DNS (v1.3 and later). An upgrade of a system without a DNS server explicitly specified to 1.3 will fail.

Troubleshoot Updates

A *database upgrade not successful* message means that a new Threat Grid Appliance is running an older version of PostgreSQL and the automated database migration process failed. It is critical that this be fixed prior to any upgrade to v2.0.

See Threat Grid Appliance Release Notes v2.0.1 for more information.

Database Schema Updates

Historically, on standalone appliances, database migrations associated with updates would occur while the system was offline in single-user mode, except in a cluster, where the updates would occur after the first

upgraded node came back online. (The exception to this was for unusually long updates that could be run in the background, which we handled on a case-by-case basis.)

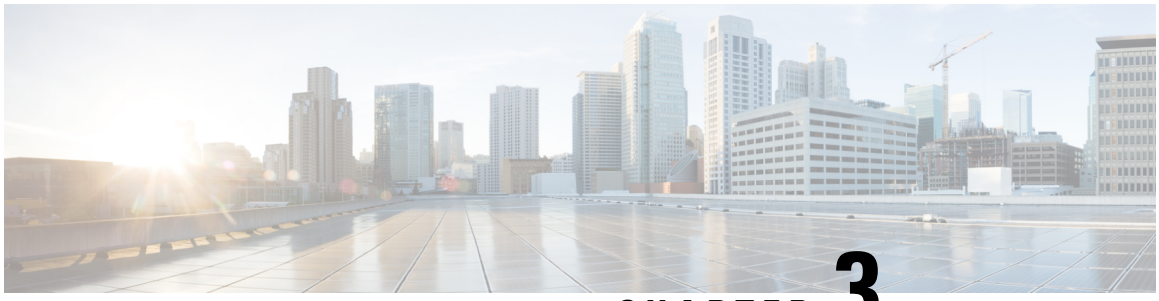
Threat Grid Appliance (v2.5.0 and later) now updates the database schema after the system finishes reboot. This may potentially cause the boot process to take slightly longer. (Very long ones will continue to be handled case-by-case.)

In prior releases, non-clustered systems with backup support enabled would make a best-effort attempt to operate correctly when their NFS server was down. Due to changes in ElasticSearch functionality, we can no longer guarantee this behavior.

Background Elasticsearch index migration to ES6-native indexes is enabled in v2.7.2 and later. This migration must successfully complete before any version of the Threat Grid Appliance which requires Elasticsearch 7.0 or newer is installed.



Note Elasticsearch index migration may cause substantial delays in the NFS backup process, causing related warnings. These warnings should be disregarded, as service notices indicate that index migration is actively ongoing. You should only raise a ticket with Support if the index migration process fails to make progress over an extended period.



CHAPTER 3

Configuration Management

This chapter describes additional information about configuring the Threat Grid Appliance after the initial configuration.

Note The initial configuration is described in the *Cisco Threat Grid Appliance Setup and Configuration Guide*. Refer to this guide for configuring the TGSH Dialog Interface and using the OpAdmin Configuration Wizard.

It includes the following topics:

- Introduction, on page 13
- Network Configuration Using TGSH Dialog, on page 13
- Configuration Using OpAdmin Portal, on page 15
- Configure LDAP Authentication, on page 17
- Configure Third-party Detection and Enrichment Services, on page 19
- Apply Configuration Change, on page 20
- Using DHCP, on page 21

Introduction

The initial Threat Grid Appliance configuration is performed during the appliance setup, as documented in the Cisco Threat Grid Appliance Setup and Configuration Guide, using the TGSH Dialog and the OpAdmin portal.

Note Threat Grid organizations and user accounts are managed in the Threat Grid Portal UI (from the drop-down arrow next to your login name in the navigation bar).

Network Configuration Using TGSH Dialog

The initial network configuration is completed using the TGSH Dialog (see Cisco Threat Grid Appliance Setup and Configuration Guide). This section provides some additional information about using the TGSH Dialog.

Configure Network Using TGSH Dialog

If you want to make changes to your initial network configuration, perform the following steps.

Note If you are using DHCP to obtain IPs, see Network Configuration and DHCP.

Procedure

Step 1 Login to TGSH Dialog.

Note
If you are configured for **LDAP Only** authentication, you can only log into TGSH Dialog using LDAP. If authentication mode is set to **System Password or LDAP**, the TGSH Dialog login only allows the **System** login.

Step 2 In the TGSH Dialog interface, select **CONFIG_NETWORK**.

The Network Configuration console opens and displays the current network settings.

Step 3 Make any necessary changes (you need to backspace over the old entry before you can enter the new one).

Step 4 Leave the Dirty network **DNS Name** blank.

Step 5 After you finish updating the network settings, tab down and select **Validate** to validate your entries.

If invalid values have been entered, you may see errors. Correct the entries and then select **Validate** again.

After validation, the Network Configuration Confirmation displays the entered values.

Step 6 Select **Apply** to apply your configuration settings.

Detailed information about the configuration changes that have been made are displayed.

Step 7 Select **OK**.

The Network Configuration Console refreshes again and displays the IP addresses. Network configuration is now complete.

Reconnect to TGSH Dialog

TGSH Dialog remains open on the console and can be accessed either by attaching a monitor to the appliance or, if CIMC is configured, via remote KVM.

To reconnect to the TGSH Dialog, SSH into the Admin IP address as the user '**threatgrid**'. The required password is either the initial randomly generated password, which is visible initially in the TGSH Dialog, or the new Admin password you create during the first step of the OpAdmin Configuration (see the *Cisco Threat Grid Appliance Getting Started Guide*).

Configure Networking in Recovery Mode

Network configuration in recovery mode mirrors the full system (v2.7 and later):

- All interfaces are brought up.
- Firewall rules and policy routing restricts which processes can communicate on which interfaces.

Note Support mode traffic on port 19791 is allow-listed across all three interfaces.

Perform the following steps to set up networking in recovery mode.

Procedure

Step 1 Reboot the appliance and then choose **Recovery Mode** in the boot menu.

Step 2 Once the system is up, press Enter several times to get a clean command prompt.

Step 3 Enter **netctl clean** and provide the following information:

- **Configuration type** - static
- **IP Address** - <Clean IP Address>/<Netmask>
- **Gateway Address** - <Clean network gateway>
- **Routes** - <leave blank>
- Final Question - Enter **y**

Step 4 Enter **Exit** to apply the configuration.

The appliance will attempt to open an outbound support connection on the Clean interface on port 19791/tcp.

Configuration Using OpAdmin Portal

The initial setup and configuration wizard is described in the Cisco Threat Grid Appliance Setup and Configuration Guide. New Threat Grid Appliances may require the administrator to complete additional configuration, and OpAdmin settings may require updates over time.

The OpAdmin Portal is the Threat Grid Appliance administrator's main configuration interface. It is a Web portal that can be used once an IP address has been configured on the Threat Grid Appliance Admin interface.

OpAdmin is the recommended tool for configuring your appliance, and in fact, much of the appliance configuration can only be done via OpAdmin. OpAdmin is used to configure and manage a number of important Threat Grid Appliance configuration settings, including:

- The administrator's passwords (for OpAdmin and the **threatgrid** user)
- Threat Grid License
- Rate Limits
- SMTP

- SSH
- SSL Certificates
- DNS servers (including DNS configuration for AMP for Endpoints Private Cloud integrations)
- NTP servers
- Server Notifications
- Syslog messages and Threat Grid Notifications remote server setup
- CA Certificate Management (for AMP for Endpoints Private Cloud integrations)
- LDAP Authentication
- Third-party Detection and Enrichment Services (including ClamAV, OpenDNS, Titanium Cloud, and VirusTotal)



Note

- Configuration updates in OpAdmin should be completed in one session to reduce the chance of an interruption to the IP address during configuration.
- OpAdmin does not validate the gateway entries. If you enter the wrong gateway and save it, the OpAdmin interface will not be accessible. You will need to use the console to fix the networking configuration if that was done on the Admin interface. If Admin is still valid, you can fix it in OpAdmin and reboot.
- Threat Grid Appliances (v2.7 or later) use the serial number as the hostname to improve interoperability with some NFS v4 servers.



Important OpAdmin uses HTTPS. Pointing a browser at the Admin IP is not sufficient; you must point to:

https://adminIP/

OR

https://adminHostname/

Configure SSH Keys

Setting up SSH keys provides the Threat Grid appliance administrator with access to the TGSH Dialog via SSH (`threatgrid@<host>`).

It does NOT provide root access or a command shell. Multiple keys may be added in **Configuration > SSH**.

On Threat Grid Appliances (v2.7 and later), configuring a SSH public key for access to the appliance disables password-based authentication via SSH; this makes SSH authentication methods one or the other, not both. After a successful SSH connection using key-based authentication, tgsh-dialog prompts for a password, such that both tokens are required.

Configure Remote Syslog Server for Notifications

In addition to the periodic notifications that can be set up to deliver system notifications via email (in OpAdmin under **Configuration** > **Notifications**), you can configure a remote syslog server to receive syslog messages and Threat Grid notifications.

Procedure

Step 1 Log in to the OpAdmin portal and click **Configuration > Syslog**.

Step 2 Enter the **server DNS** and then choose a protocol from the drop-down list (TCP is the default, the other is UDP).

Step 3 Check the **Verification** check box to perform a DNS lookup after you save the configuration.

If the host cannot resolve the name, it will print an error and will not save (until you enter a valid hostname). If you do not check the **Verification** check box, the appliance will accept any name, whether valid in DNS or not.

Step 4 Click **Save**.

To edit or delete the Syslog DNS, open **Configuration > Syslog** and make your modifications, and then click **Save**.

Configure LDAP Authentication

The Threat Grid Applicance supports LDAP authentication and authorization for OpAdmin and TGSH Dialog login.

You can authenticate multiple appliance administrators with different credentials that are managed on the domain controller or the LDAP server. Authentication modes include: System Password Only, System Password or LDAP, and LDAP Only.

There are three LDAP Protocol options: LDAP, LDAPS, and LDAP with STARTLS.

The following considerations should be reviewed:

- The dual authentication mode (**System Password or LDAP**) is required to avoid accidentally locking yourself out of the appliance when setting up LDAP. Selecting **LDAP Only** is not allowed initially; you must first go through dual mode to make sure it works. You will need to log out of OpAdmin after the initial configuration, and then log back in using LDAP credentials in order to toggle to **LDAP Only**.
- You can only log into TGSH Dialog using LDAP if you are configured for LDAP Only authentication. If authentication mode is set to System Password or LDAP, then the TGSH Dialog login will only allow the System login.
- If the appliance is configured for LDAP authentication only (**LDAP Only**), then resetting the password in recovery mode will reconfigure the authentication mode to allow login with system password as well.
- Make sure that the authentication filter is set up to restrict membership.
- TGSH Dialog and OpAdmin require LDAP credentials only in **LDAP Only** mode: if **LDAP only** is configured, TGSH Dialog will not ask for the system password but for an LDAP user/password.
- If authentication is configured for **System Password or LDAP**, TGSH Dialog will continue to ask for the system pw only, it'll not have both.

- Troubleshooting LDAP: If it breaks, disable it by doing a password reset in Recovery Mode.
- TGSH Dialog access via SSH: A system password or a configured SSH key is required **in addition to** LDAP credentials for tgsh-dialog access via ssh when in LDAP Only mode.
- LDAP is outbound from the Clean interface.

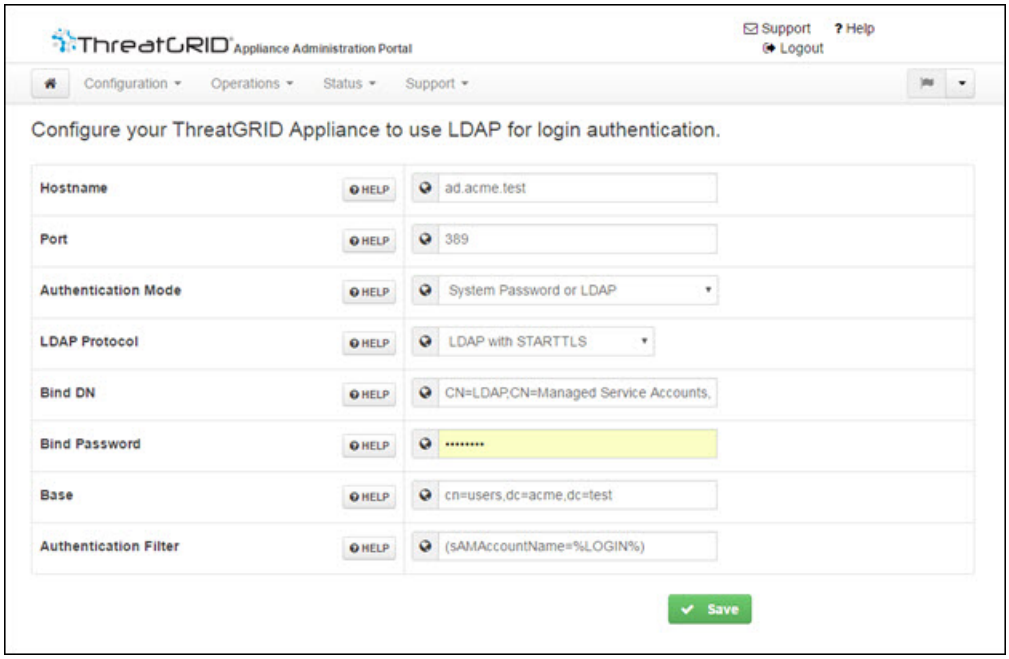
Configure LDAP Authentication in OpAdmin

Perform the following steps to configure LDAP authentication in the OpAdmin portal.

Procedure

Step 1 Log in to the OpAdmin portal and choose **Configuration > LDAP** to open the LDAP configuration page.

Figure 8: LDAP Authentication Configuration



Step 2 Complete the fields on the page. You can click the **Help** icon next to each field for a detailed description and more information.

Note
The first time you configure LDAP authentication, you must select **System Password** or **LDAP**, log out of OpAdmin, and then log back in using your LDAP credentials. You can then change the setting to implement **LDAP Only**.

Step 3 Click **Save**.

When users log in to OpAdmin or TGSH Dialog, they will now see one of the following screens:

Figure 9: LDAP Only

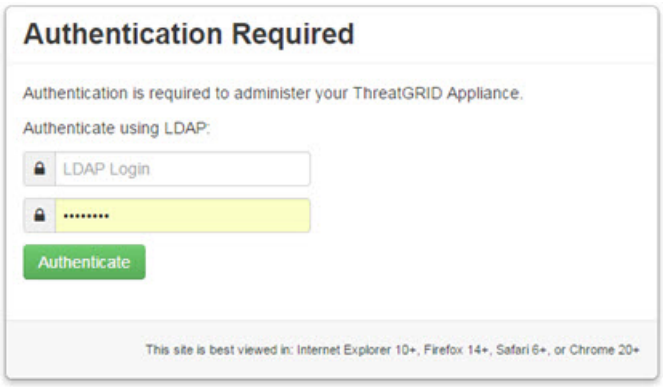
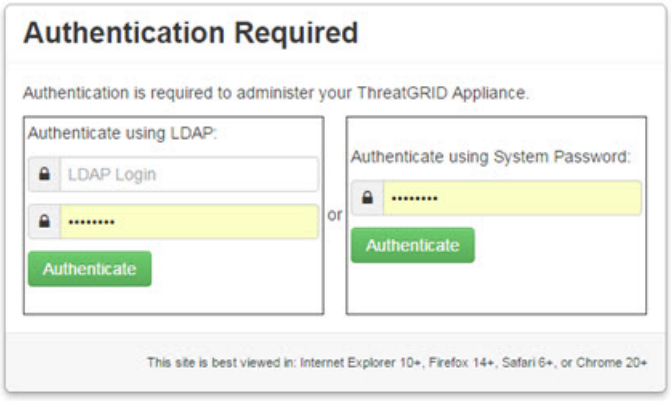


Figure 10: System Password or LDAP



Configure Third-party Detection and Enrichment Services

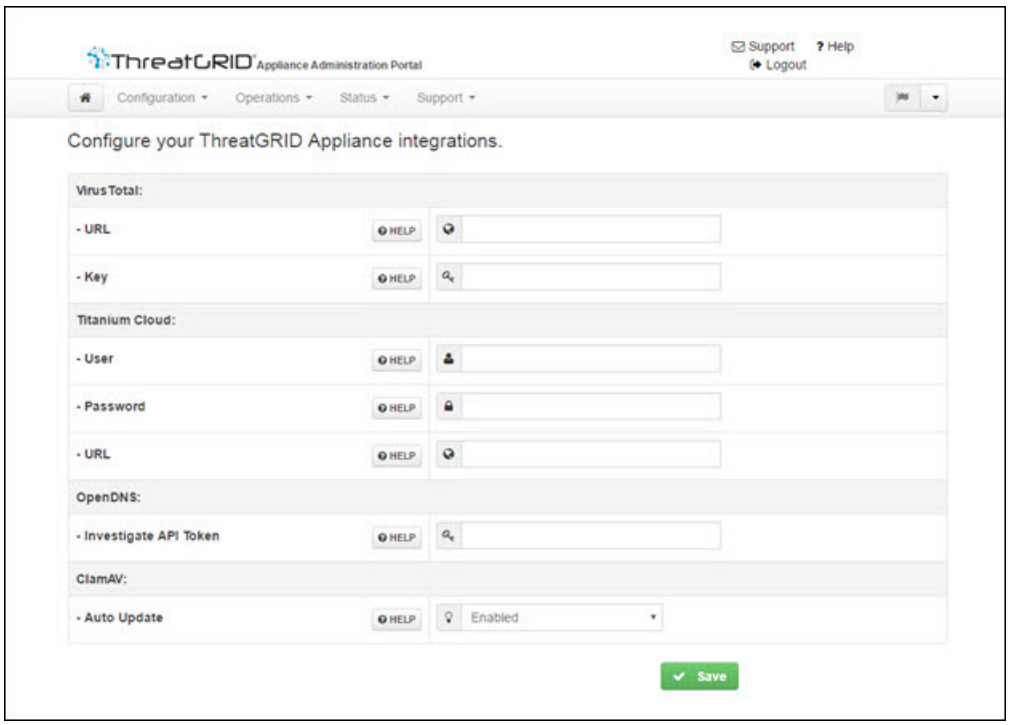
Integrations with several third-party detection and enrichment services, including OpenDNS, TitaniumCloud, and VirusTotal, can be configured on the appliance using the Integration Configuration page. This applies to Version 2.2 and later.

Note If OpenDNS is not configured, the **whois** information on the Domains entity page in the analysis report in the portal (in the Mask version of the UI) will not be rendered.

Procedure

Step 1 Log in to the OpAdmin portal and choose **Configuration > Integrations** to open the Integrations configuration page.

Figure 11: Integrations Configuration



Step 2 Enter the authentication or other values required.

Note
ClamAV signatures can be automatically updated on a daily basis. This is enabled by default, and can be disabled on the Integrations Configuration page (**ClaimAV** field).

Step 3 Click **Save**.

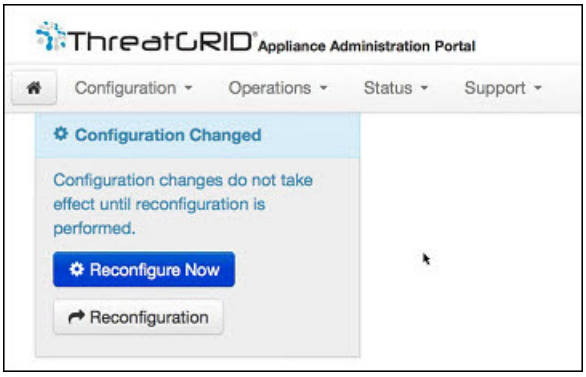
Apply Configuration Change

Any time changes are made to configuration settings, a light blue **Configuration Changed** alert appears below the **Configuration** menu. When you are done updating any OpAdmin configuration settings, you must save the new configuration in a separate step.

Procedure

Step 1 Click **Configuration Changed** to open the dialog.

Figure 12: Configuration Changed Dialog



Step 2 Click **Reconfigure Now** to apply your changes to the appliance.

Using DHCP

Most Threat Grid Appliance users do not use a network configured with DHCP. However, if you are connected to a network configured to use DHCP, then read this section.

Note If the initial appliance network configuration used DHCP and you now need to switch to static IP addresses, see Network Configuration and DHCP.

TGSH Dialog displays the information you will need to access and configure the OpAdmin portal interface. It may take some time for the IP addresses for DHCP to display after your appliance boots.

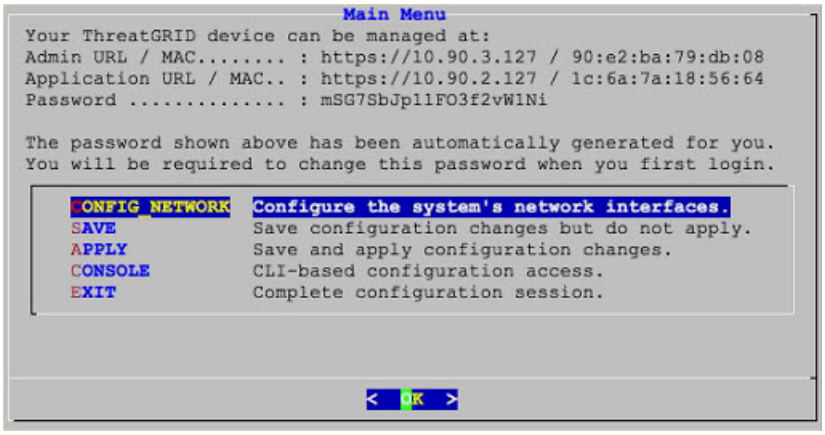
Explicit DNS for DHCP

Threat Grid Appliances (v1.3 and later) that use DHCP need to explicitly specify DNS.

Warning An upgrade of a system without a DNS server explicitly specified to v1.3 will fail.

Open the TGSH Dialog and note the following information.

Figure 13: TGSH Dialog (Connected to a Network Configured to Use DHCP)



- **Admin URL** - The Admin network. You will need this address in order to continue the remaining configuration tasks with OpAdmin.
- **Application URL** - The Clean network. Note: This is the address to use after completing the configuration with OpAdmin, in order to access the Threat Grid application.

  The Dirty network is not shown.
- **Password** - The initial administrator's password, which is randomly generated during the appliance installation. You will need to change this password later as the first step the OpAdmin configuration process.

If you plan to use DHCP on a permanent basis, then no additional network configuration is necessary, unless you need to change the Admin IP address to static.

Network Configuration and DHCP

If you used DHCP for initial configuration, and you now need to adjust the IP assignment from DHCP to your permanent static IP addresses for all three networks, perform the following steps.

Note OpAdmin will not validate the gateway entries. If you enter the wrong gateway and save it, the OpAdmin interface will be inaccessible. You will need to use the console to fix the networking configuration if that was done on the admin interface. If Admin is still valid, you can fix it in OpAdmin and reboot.

Procedure

Step 1 In the OpAdmin portal, click **Network** in the navigation pane (although **Configuration** > **Network** is checked in the License window, the DHCP network configuration has not yet been done).

The Network configuration page opens.

Step 2 Complete the following fields:

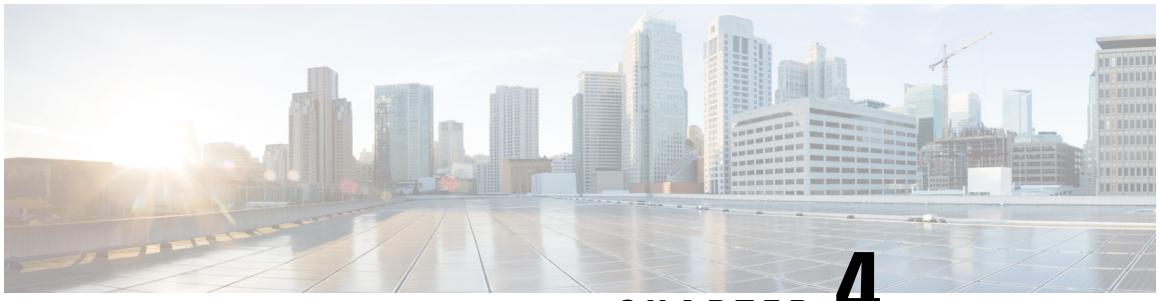
Note
The Admin network settings were configured using the **TGSH Dialog** during the initial appliance setup and configuration.

- **IP Assignment** - Choose **Static** from the drop-down menu for both Clean and Dirty network interface.
- **IP Address** - Enter a static IP address for the Clean or Dirty network interface.
- **Subnet mask** - Complete as appropriate for the type of network interface.
- **Validate DNS Name** - For Clean network interface, check the **Validate DNS Name** check box to verify that the DNS resolves to the IP address.
- **Primary and Secondary DNS** - Enter the primary and secondary DNS server information.

Step 3 Click **Next (Applies Configuration)** to save your network configuration settings.

Note
Email configuration is managed from the Email page and Time NTP servers are managed on the Date and Time page.

Step 4 Click **Configuration Changed** and choose **Reconfigure Now** to apply your DHCP configuration settings.



CHAPTER 4

Manage SSL Certificates

This chapter provides information about managing SSL certificates for your Threat Grid Appliance and integrated appliances and devices. It includes the following topics:

- About SSL Certificates and Threat Grid Appliance, on page 25
- Configure SSL Certificates for Inbound Connections, on page 26
- Configure SSL Certificates for Outbound Connections, on page 30
- Connect ESA/WSA to Threat Grid Appliance, on page 32
- Connect AMP for Endpoints Private Cloud to Threat Grid Appliance, on page 35

About SSL Certificates and Threat Grid Appliance

All network traffic passing to and from the Threat Grid Appliance is encrypted using SSL. The following information is provided to assist you through the steps for setting up SSL certificates to support Threat Grid Appliance connections with Email Security Appliance (ESA), Web Security Appliance (WSA), AMP for Endpoints Private Cloud, and other integrations.

Note A full description of how to administer SSL certificates is beyond the scope of this guide.

Interfaces Using SSL

There are two interfaces on the Threat Grid Appliance that use SSL:

- **Clean** interface for the Threat Grid Portal UI and API, and integrations (ESA/WSA appliances, AMP for Endpoints Private Cloud Disposition Update Service).
- **Admin** interface for the OpAdmin Portal.

Supported SSL/TLS Version

The following versions of SSL/TLS are supported on the Threat Grid Appliance:

- TLS v1.0 - Disabled in the Admin interface (v2.7 and later)
- TLS v1.1 - Disabled in the Admin interface (v2.7 and later)
- TLS v1.2

Note TLS v1.0 and TLS v1.1 are disabled in the Admin interface (v2.7 and later), and disabled by default for the main application. If one of these protocols is required for integration compatibility purposes, they can be re-enabled (for the main application only) from the tgsh.

Supported Customer-Provided CA Certificates

Customer-provided CA certificates are supported (v2.0.3 and later) to allow customers to import their own trusted certificates or CA certificates.

Self-Signed Default SSL Certificates

The Threat Grid Appliance is shipped with a set of self-signed SSL certificates and keys already installed. One set is for the **Clean** interface and the other is for the **Admin** interface. The appliance SSL certificates can be replaced by an administrator.

The default Threat Grid Appliance SSL certificate hostname (Common Name) is **pandem**, and is valid for 10 years. If a different hostname was assigned to the Threat Grid appliance during configuration, then the hostname and the Common Name in the certificate will no longer match.

The hostname in the certificate must also match the hostname expected by a connecting ESA or WSA appliance, or other integrating Cisco device or service, as many client applications require SSL certificates where the Common Name used in the certificate matches the hostname of the appliance.

Configure SSL Certificates for Inbound Connections

Cisco security products, such as Email Security Appliance, Web Security Appliance, and AMP for Endpoints Private Clouds, can integrate with a Threat Grid Appliance and submit samples to it. These integrations are *Inbound* connections from the perspective of the Threat Grid Appliance.

The integrating appliance or other device must be able to trust the Threat Grid Appliance SSL certificate. You must first validate that the hostname matches the Common Name; if it doesn't match, you must regenerate or replace it. You then must export the SSL certificate from the Threat Grid Appliance, and then import it into the integrating appliance or service.

The certificates on the Threat Grid Appliance that are used for inbound SSL connections are configured in the SSL Certificate Configuration page. The SSL certificates for the Clean and Admin interfaces can be configured independently.

Procedure

Step 1 In the OpAdmin portal, click **Configuration > SSL** to open the SSL Certificate configuration page.

Figure 14: SSL Certificate Configuration Page



In this example, there are two SSL certificates: **ThreatGRID Application** is the Clean interface, and **Administration Portal** is the Admin interface.

Step 2 Validate that the hostname matches the Common Name used in the SSL certificate (green padlock icon). See Validate Common Name in SSL Certificate.

Validate Common Name in SSL Certificate

The hostname must match the Common Name used in the SSL certificate on the Threat Grid Appliance.

On the SSL Certificate Configuration page, the padlock icon in the column to the left of the interface name indicates the status of the SSL certificates:

- **Green** - Indicates the interface hostname matches the Common Name used in the SSL certificate.
- **Yellow** - Indicates that the interface hostname does not match the Common Name in the SSL certificate. You must replace the certificate with one that uses the current hostname.

If the hostname and Common Name do not match, you must replace the certificate with one that uses the current hostname (see Replace SSL Certificate).

Replace SSL Certificate

SSL certificates usually need to be replaced at some point for various reasons, such as certificate has expired, the hostname has changed, or to support integrations with other Cisco devices and services.

Cisco Email Security Appliance, Web Security Appliance, and other CSA Cisco integrating devices may require an SSL certificate in which the Common Name matches the Threat Grid appliance hostname. You must replace the default SSL certificate with a newly generated certificate that uses the same hostname to access the Threat Grid Appliance.

If integrating a Threat Grid Appliance with an AMP for Endpoints Private Cloud to use its Disposition Update Service, you must install the AMP for Endpoints Private Cloud SSL Certificate so the Threat Grid Appliance can trust the connection.

There are several ways to replace an SSL certificate on a Threat Grid Appliance:

- Regenerate SSL Certificate that uses the current hostname for the Common Name.
- Download SSL Certificate.
- Upload SSL Certificate; this can be a commercial or enterprise SSL, or one you create using OpenSSL.
- Generate SSL Certificate Using OpenSSL

Regenerate SSL Certificate

You can regenerate a SSL certificate from the SSL Certificate Configuration page (v1.4.2 or later required) if your hostname does not match the Common Name in the certificate.

In the OpAdmin SSL Certificate Configuration page, click Regenerate.

Procedure

Step 1 In the OpAdmin portal, click **Configuration > SSL** to open the SSL Certificate configuration page.

Step 2 In the Operations column, click **Regenerate** for the interface that needs a new certificate.

A new self-signed SSL certificate is generated on the Threat Grid Appliance that uses the current hostname of the appliance in the Common Name field of the certificate. The Common Name validation padlock icon next to the interface name changes to green.

You can now Download SSL Certificate the regenerated certificate (.cert file) and install it on the integrating appliance.

Download SSL Certificate

The Threat Grid Appliance SSL certificate can be downloaded and installed on your integrating device so it can trust connections from the Threat Grid Appliance.

Procedure

Step 1 In the OpAdmin portal, click **Configuration > SSL** to open the SSL Certificate configuration page.

Step 2 In the Operations column, click **Download** for the interface certificate. The SSL certificate **.cert** file is downloaded.

Step 3 Install the downloaded SSL certificate (**.cert** file) on the Email Security Appliance, Web Security Appliance, AMP for Endpoints Private Cloud, or other integrating Cisco products per the product documentation.

Upload SSL Certificate

If you already have a commercial or corporate SSL certificate in place within your organization, you can use that to generate a new SSL certificate for the Threat Grid Appliance and use the CA cert on the integrating device.

Procedure

Step 1 In the OpAdmin portal, click **Configuration > SSL** to open the SSL Certificate configuration page.

Step 2 In the Operations column, click **Upload** for the appropriate interface.

The Common Name validation padlock icon next to the interface name changes to green.

Generate SSL Certificate Using OpenSSL

You can manually generate a SSL certificate using OpenSSL when there is no SSL certificate infrastructure already in place on your premises and upload it to the Threat Grid Appliance (as described in Upload SSL Certificate). OpenSSL is a standard open-source SSL tool for creating and managing OpenSSL certificates, keys, and other files.



Note OpenSSL is not a Cisco product, therefore does not provide technical support for it. It is recommended that you search the Web for additional information on using OpenSSL. Cisco does offer a SSL library, *Cisco SSL*, for generating SSL certificates.

Procedure

Step 1 Run the following command to generate a new self-signed SSL certificate:

```
openssl req -x509 -days 3650 -newkey rsa:4096 -keyout tgapp.key -nodes -out
tgapp.cert -subj "/C=US/ST=New York/L=Brooklyn/O=Acme Co/CN=tgapp.acmeco.com"
```

- **openssl** - OpenSSL
- **req** - Specifies to use X.509 certificate signing request (CSR) management. X.509 is a public key infrastructure standard that SSL and TLS use for key and certificate management. In this example, this parameter is used to create a new X.509 cert.
- **-x509** - This modifies the req parameter X.509 to make a self-signed certificate instead of generating a certificate signing request.
- **-days 3650** - This option sets the length of time for which the certificate will be considered valid. In this example, it is set for 10 years.
- **-newkey rsa:4096** - This specifies to generate a new certificate and a new key at the same time. Because the required key was not previously created, it must be created with the certificate. The **rsa:4096** parameter indicates to make an RSA key that is 4096 bits long.

- **-keyout** - This parameter indicates where OpenSSl should save the generated private key file that is being created.
- **-nodes** - This parameters indicates that OpenSSL should skip the option to secure the certificate with a passphrase. The appliance needs to be able to read the file, without user intervention, when the server starts up. A certificate that is secured with a passphrase requires that the user enter the passphrase every time the server is restarted.
- **-out** - This parameter indicates where OpenSSL should save the certificate that is being created.
- **-subj:** (Example):
  - **C=US** - Country
  - **ST=New York** - State
  - **L=Brooklyn** - Location
  - **O=Acme Co** - Owner's name
  - **CN=tgapp.acmeco.com** - Enter the Threat Grid Appliance FQDN (Fully Qualified Domain Name). This includes the HOSTNAME of the Threat Grid Appliance (in this example, **tgapp**) and the associated domain name (in this example, **acmeco.com**).

    Important
    You must at least change the Common Name to match the FQDN of the Threat Grid Appliance Clean interface.

Step 2 Once the new SSL certificate is generated, upload the certificate to the Threat Grid Appliance from the SSL Certificate Configuration page (see Upload SSL Certificate). You must also upload the certificate (**.cert** file only) to the Email Security Appliance or Web Security Appliance.

Configure SSL Certificates for Outbound Connections

The Threat Grid Appliance (v2.0.3 and later) supports integration with Cisco AMP for Endpoints Private Cloud for the Disposition Update Service. This integration is a *Outbound* connection from the perspective of the Threat Grid Appliance.

Configure DNS

By default, DNS uses the Dirty interface. If the hostname of an integrating appliance or service, such as a AMP for Endpoints Private Cloud, cannot be resolved over the Dirty interface because the Clean interface is used for the integration, a separate DNS server that uses the Clean interface can be configured in OpAdmin.

Procedure

Step 1 In OpAdmin, click **Configuration** > **Network**.

Step 2 Complete the **DNS** fields for the Dirty and Clean networks.

Step 3 Click **Save**.

Manage CA Certificates

The CA Certificate page (v2.0.3 or later) in the OpAdmin portal is used to manage the CA Certificate trust store for outbound SSL connections so that the Threat Grid Appliance can trust the Cisco AMP for Endpoints Private Cloud to notify it about analyzed samples that are considered malicious.

Procedure

Step 1 In the OpAdmin portal, click **Configuration > CA Certificates**.

Step 2 Choose one of the following import options:

- **Import from Host** to retrieve the certificate from the server. Enter the **Host** and **Port** for the AMP for Endpoints Private Cloud and then click **Retrieve**.
- **Import from Clipboard**, paste the PEM from the clipboard, and then click **Add Certificate**.

Step 3 Click **Import**.

Manage Disposition Update Syndication Service

You can manage the Disposition Update Syndication Service for AMP for Endpoints Private Cloud appliance integrations in the Threat Grid portal user interface (v2.2 or later). URLs can be added, edited, and deleted from the Disposition Update Syndication Service page.

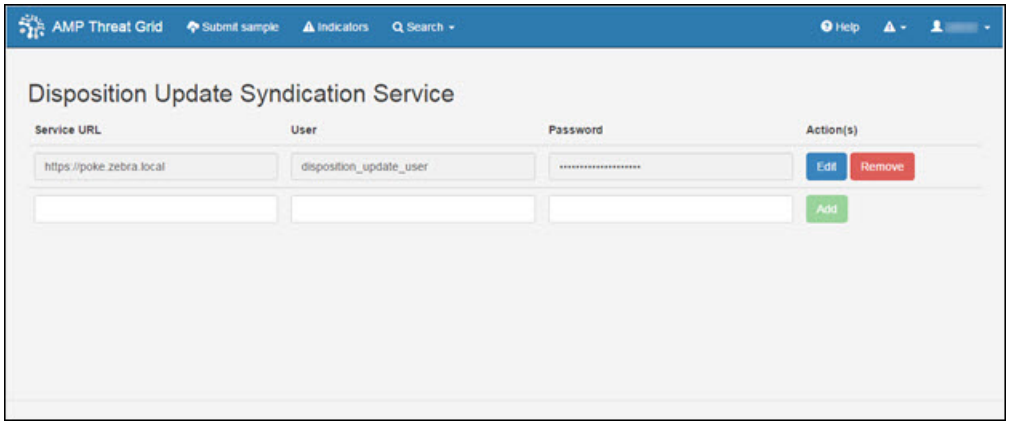


Note For more information about AMP for Endpoints Private Cloud appliance integrations, see Connect AMP for Endpoints Private Cloud to Threat Grid Appliance.

Procedure

Step 1 In the Threat Grid portal, click the drop-down on the navigation bar next to your login name and choose **Manage FireAMP Integration** to open the Disposition Update Syndication Service page.

Figure 15: Disposition Update Syndication Service



Step 2 Enter the following information:

- **Service URL** - The AMP for Endpoints Private Cloud URL.
- **User** - The admin user name.
- **Password** - The password provided by the AMP for Endpoints configuration portal.

Step 3 Click **Config**.

Connect ESA/WSA to Threat Grid Appliance

Cisco products such as Email Security Appliance (ESA), Web Security Appliance (WSA), and other appliances, devices, and services, can integrate with Threat Grid Appliances through connections encrypted with SSL to submit possible malware samples for analysis.

Integration between ESA/WSA and Threat Grid Appliance are enabled by the Cisco Sandbox API (CSA API) and are often referred to as CSA Integrations.

An integrating ESA/WSA must be registered with the Threat Grid Appliance before it can submit samples for analysis. Before the integrating ESA/WSA can be registered with the Threat Grid Appliance, the ESA/WSA administrator must first set up the SSL certificate connection as appropriate for their appliance and their network environment.

This section describes the steps necessary for setting up ESA, WSA and other Cisco products to communicate with the Threat Grid Appliance.

ESA and WSA Documentation

See the instructions for *Enabling and Configuring File Reputation and Analysis Services* in the online help or user guide for your ESA/WSA.



Note The Threat Grid appliance is often referred to as an analysis service, or private cloud file analysis server in these guides.

- Cisco Email Security Appliance User Guides
- Cisco Web Security Appliance User Guides

ESA and WSA Integration Process Overview

This section provides an overview of the steps in setting up a connection between an ESA/WSA or other CSA integration (inbound) with a Threat Grid Appliance. See ESA/WSA Integration Process Steps for more detailed descriptions and steps.

Threat Grid Appliance SSL Certificate SAN or CN Must Match its Current Hostname and ESA/WSA Expectations

The Threat Grid Appliance SSL certificate SAN (Subject Alternative Name), or the CN (Common Name) needs to match the hostname and the ESA/WSA expectations; for a successful connection with an integrating ESA/WSA, this must be the same hostname by which the integrating ESA/WSA identifies the Threat Grid Appliance.

Depending on your requirements, you may need to regenerate the self-signed SSL certificate on the Threat Grid Appliance so it uses the current hostname in the SAN/CN field, then download it to your working environment and upload and install it onto the integrating ESA/WSA.

Alternatively, you may need to replace the current Threat Grid Appliance SSL certificate by uploading an enterprise or commercial SSL certificate (or a manually generated certificate). For detailed instructions, see Configure SSL Certificates for Inbound Connections.

Verify Connectivity

Once the SSL certificate setup is complete, the next step is to verify that the ESA/WSA can communicate with the Threat Grid Appliance.

The ESA/WSA must be able to connect to the Clean interface of the Threat Grid Appliance over your network. Follow the instructions in the appropriate guide for your product to verify that the Threat Grid Appliance and ESA/WSA can communicate with each other (see Connect ESA/WSA to Threat Grid Appliance).

Complete the ESA/WSA File Analysis Configuration

Enable the **File Analysis Security** service and configure the advanced settings.

Register ESA/WSA with Threat Grid Appliance

An ESA/WSA that is configured according to the product documentation, registers itself automatically with the Threat Grid Appliance. Upon registration of the connecting device, a new Threat Grid user is automatically created with the Device ID as the login ID, and a new organization is created with a name based on the same ID. An administrator must activate the new Device user account.

Activate the New ESA/WSA Account on the Threat Grid Appliance

When the ESA/WSA or other integration connects and registers itself with the Threat Grid Appliance, a new Threat Grid user account is automatically created. The initial status of the user account is *deactivated*. A Threat Grid Appliance administrator must manually activate the device user account before it can be used for submitting malware samples for analysis.

ESA/WSA Integration Process Steps

The connection between the ESA/WSA is *incoming* from the perspective of the Threat Grid Appliance. The integration uses the CSA API.



Note Refer to the ESA and WSA User Guides for more detailed information on the tasks that must be performed on that side.

Procedure

Step 1 Set up and configure the Threat Grid Appliance as normal (no integration yet). Check for updates and install, if necessary.

Step 2 Set up and configure the ESA/WSA as normal (no integration yet).

Step 3 The TGA SSL Certificate SAN or CN Must Match its Current Hostname and ESA/WSA Expectations. If you are deploying a self-signed SSL certificate, generate a new SSL Certificate (on the Threat Grid Application the Clean interface), to replace the default if needed, and download it to install in the ESA/WSA appliance device (see Replace SSL Certificate).

Note
Be sure to generate a certificate that has the hostname of your Threat Grid Appliance as the SAN or CN. The default certificate from the Threat Grid Appliance does not work. Use the hostname, not the IP address.

Step 4 Verify that the ESA/WSA can connect to the Clean interface of the Threat Grid Appliance over your network.

Step 5 Configure the ESA/WSA for Threat Grid Appliance integration. See the ESA/WSA guides for complete instructions. The following steps are specific to the ESA, as this is currently the most common type of integration:

a) Select **Security Services > File Reputation and Analysis**.
b) Click **Enable**.
c) Click **Edit Global Settings**.
d) In the File Analysis section, File Analysis is enabled by default. If you do not uncheck the **Enable File Analysis** check box, the File Analysis feature key will be activated after the next commit. Select the file types to send to the Cloud for analysis.
e) Configure the **Advanced Settings** for File Analysis as needed, according to the ESA or WSA guides:

- **File Analysis Server URL** - Choose Private Cloud.
- **Server** - URL of the on-premises Cisco Threat Grid appliance. Use the hostname, not the IP address, for this value and for the certificate.
- **SSL Certificate** - Upload a self-signed certificate that you have generated from your on-premises Cisco Threat Grid Appliance. The most recently uploaded self-signed certificate is used. It is not possible to access a certificate uploaded prior to the most recent certificate; if needed, upload the desired certificate again.

Step 6 Submit and commit your changes.

Note the **File Analysis Client ID** that appears at the bottom of the page. This identifies the user that will be activated.

Registration of your ESA/WSA with the Threat Grid Appliance occurs automatically when you submit the configuration for File Analysis.

Step 7 Activate the new device user account on the Threat Grid Appliance:

a) Log into the Threat Grid Portal UI as Admin.
b) From the navigation bar drop-down menu next to your login name, choose Manage Users to open the Threat Grid Users page.
c) Open the User Details page for the device user account (you may need to use Search to find it).
d) The user status is currently *de-activated*. Click **Re-Activate User**.
e) On the confirmation dialog, click **Re-Activate** to confirm the action.

The ESA/WSA or other integrating appliance or device can now initiate connections with the Threat Grid Appliance.

Connect AMP for Endpoints Private Cloud to Threat Grid Appliance

The Threat Grid appliance supports integration with AMP for Endpoints Private Cloud for the Disposition Update Service as an outbound connection.

Note The Threat Grid Appliance Disposition Update Service and AMP for Endpoints Private Cloud integration set-up tasks must be performed on the devices in the specified order, particularly if you are setting up new appliances. If you are integrating appliances that are already set up and configured, the order is not as critical.

Refer to the AMP for Endpoints Private Cloud documentation for more detailed information on the tasks which must be performed in the product.

Procedure

Step 1 Set up and configure the Threat Grid Appliance as normal (no integration yet). Check for updates and install, if necessary.

Step 2 Set up and configure the AMP for Endpoints Private Cloud as normal (no integration yet).

Step 3 In the Threat Grid Appliance OpAdmin interface, Regenerate SSL Certificate on the Clean interface), to replace the default if needed, and download it to install on the AMP for Endpoints Private Cloud device.

Obtain the following information, which is needed to configure the integration in AMP for Endpoints Private Cloud device:

- **Hostname** - Click **Configuration** > **Hostname** and note the hostname.
- **API Key** - Copy the **API Key** from the User Details page in the Threat Grid Portal (click the drop-down next to your login name and choose Manage Users, then navigate to the integration user account).

Note
This does not need to be the admin user, but can be another user that was specifically created for this purpose on the Threat Grid Appliance.

Step 4 Configure the AMP for Endpoints Private Cloud device for Threat Grid Appliance integration:

a) Select **Integrations > Threat Grid** and go to the **Connection to Threat Grid** section.
b) Complete the following fields:

- **Hostname** - Enter the Threat Grid Appliance hostname (obtained in previous step).
- **API Key** - Enter the Threat Grid API Key for the account to be used for integrations (obtained in previous step).
- **SSL Certificate** - Choose the Threat Grid Appliance SSL Certificate file.

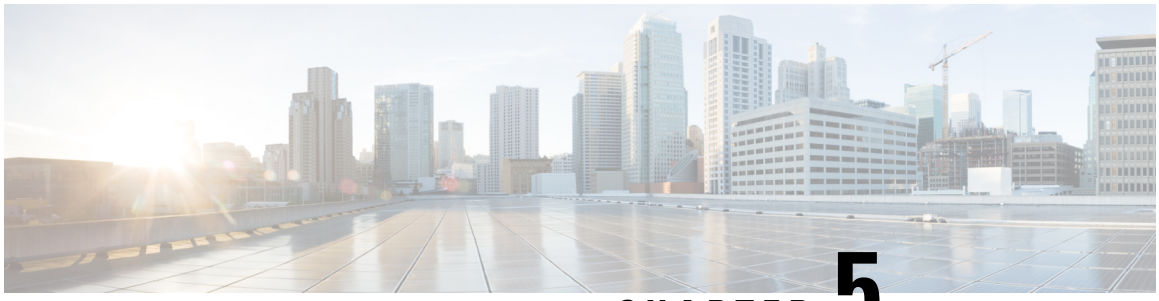
c) Click **Save Configuration**.
d) Click **Test Connection**.

Once the connection test passes, you must run the Reconfiguration on the AMP for Endpoints Private Cloud to apply the changes. This allows AMP to talk to the Threat Grid Appliance, and you can now submit samples to Threat Grid.

However, you must complete the remaining steps to set up the Disposition Update Service to communicate disposition results to the Threat Grid Appliance. (For more information, see the user documentation for AMP for Endpoints Private Cloud.)

Step 5 In the Threat Grid Appliance, set up the Disposition Update Syndication Service:

a) Configure DNS, if needed.
b) Download or copy and paste the AMP for Endpoints Private Cloud SSL certificate to the Threat Grid Appliance so it can trust the integrating device. See Manage CA Certificates.
c) From the upper-right menu, choose **Manage FireAMP Integration** and specify the AMP Disposition Update Service URL and credentials (see Manage Disposition Update Syndication Service).
d) Click **Config**.



CHAPTER 5

Manage Organizations and Users

This chapter describes how to manage organizations and users in Threat Grid. It includes the following topics:

- Introduction, on page 37
- Create New Organization, on page 37
- Manage Users, on page 38
- Activate New Device User Account, on page 38

Introduction

Threat Grid is installed on the appliance with a default organization and Admin user. Once the appliance is set up and the network configuration is completed, you can create additional organization and user accounts, so people can login and begin submitting malware samples for analysis.

Adding organizations, users, and administrators may require planning and coordination among multiple users and teams, depending on your organization.

Create New Organization

Users are always affiliated with an organization; before you can add users, you must first create the organization so you can add them to it.



Important You cannot delete an organization from this interface once it has been created so plan this task carefully.

Procedure

Step 1 Log into the Threat Grid portal as Admin.

Step 2 Click the **Administration** menu and choose **Manage Organization**. The Organizations page opens shows all the organizations on the appliance.

Step 3 Click **New Organization** in the upper-right corner of the page to open the New Organization dialog.

Step 4 Complete the following information:

- **Name** - Add a name for the organization (there is currently no size limit to the name).
- **Industry** - Select the type of business from the Industry drop-down menu. If none of the industries on the list are applicable, then leave it set to Unknown, and contact Threat Grid support (support@threatgrid.com) to request that an option be added.
- **ATS Id** - Enter the Advanced Threat Services ID.

Step 5 Click **Submit**. The new organization is created and is now visible in the list of Organizations.

Step 6 Edit the newly created organization and complete the following information:

- **Options** - Complete as appropriate.
- **Rate Limit** - Set the default user submission rate limit.

  The API rate limit is global for the appliance under the terms of the license agreement. This affects API submissions only, not manual sample submissions. The rate limit in the license applies to the organization.

  You can also set sample submission rates on individual users, as documented in *Using Threat Grid* in the online Help (from the navigation bar click **Help > Using Threat Grid Online Help**).

  Rate limits are based on a 24-hour window of rolling time, not to a calendar day. When the submission limit is exhausted, the next API submission will return a 429 error, plus a message about how long to wait before retrying.

Once the organization is created, the Admin or Organization Admin can manage it (see *Managing Organizations* in the online Help.

Manage Users

For instructions and documentation on creating and managing user accounts, including how to add users, see the Threat Grid Portal UI online help:

In the navigation bar, click **Help > Using Threat Grid Online Help > Managing Threat Grid Users**.

Note Users can only be removed via the API, and only if they have submitted no samples.

Managing device user accounts for integrating Cisco Email Security Appliance, Web Security Appliance, and other devices is described in Activate New Device User Account.

Activate New Device User Account

When the Cisco Email Security Appliance, Web Security Appliance, or other Cisco Sandbox API integration connects and registers itself with a Threat Grid Appliance, a new Threat Grid user account is automatically created. The initial status of the user account is *de-activated*. The device user account must be manually activated by a Threat Grid Appliance administrator before it can be used for submitting malware samples for analysis.

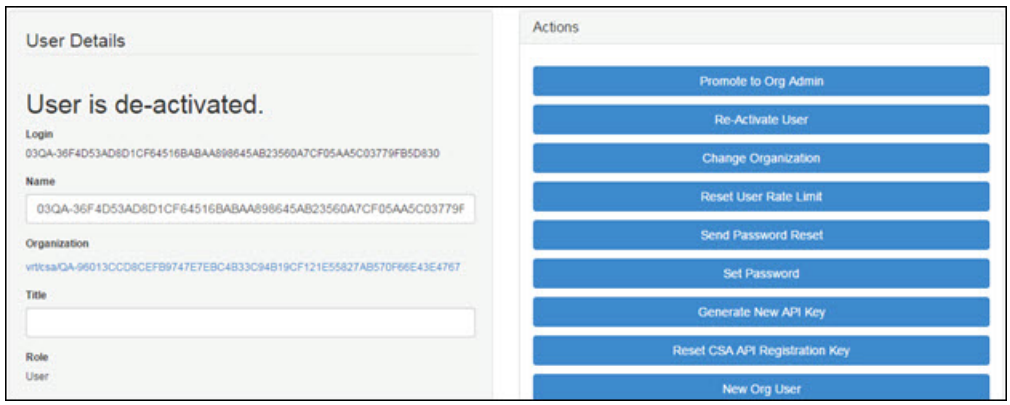
Procedure

Step 1 Log into the Threat Grid Portal UI as Admin.

Step 2 Click the **Administration** menu and choose **Manage Users**. The Threat Grid User Details page opens.

Step 3 Locate the device user account and open the User Details page. The user status is currently *de-activated*.

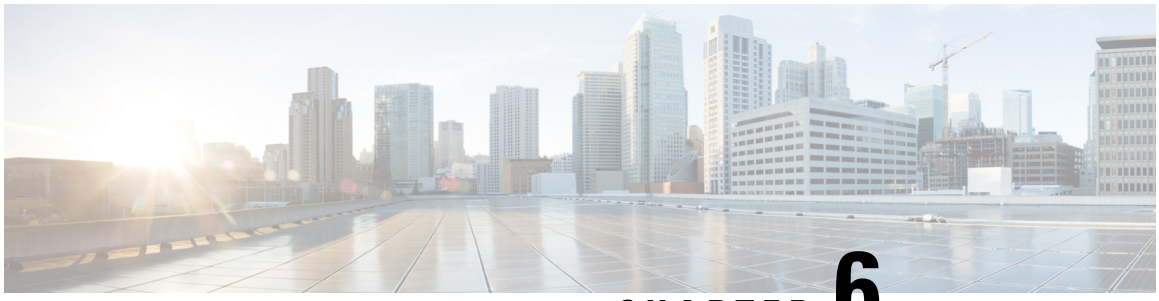
Figure 16: User Details



Step 4 Click **Re-Activate User**. A dialog opens asking you to confirm.

Step 5 On the confirmation dialog, click **Re-Activate User** to confirm the action.

The ESA/WSA or other integrating appliance or device can now communicate with the Threat Grid Appliance.



CHAPTER 6

Privacy and Sample Visibility

This chapter provides information about the privacy and sample visibility model for sample submissions to Threat Grid. It includes the following topics:

- About Privacy and Sample Visibility, on page 41
- Privacy and Visibility for Integrations, on page 41

About Privacy and Sample Visibility

When submitting samples to a Threat Grid appliance for analysis, an important consideration is the privacy of their contents. Privacy is a particularly important consideration if sensitive documents or archive types are submitted for analysis, because locating sensitive material could be relatively easy for those with access to the Threat Grid Appliance, especially with the search API.

The privacy and sample visibility model for sample submissions to Threat Grid is as follows:

- Unless samples are designated as Private, they are visible to users who are outside the submitter's organization.
- Private samples can only be seen by Threat Grid users within the same organization as the user who submitted the sample.

Privacy and Visibility for Integrations

The privacy and sample visibility model is modified on Threat Grid Appliances for samples that are submitted by integrations. Integrations are Cisco products such as Email Security Appliance (ESA), Web Security Appliance (WSA), and other devices or third-party services (you may see the term CSA Integrations, which refers to ESA/WSA and other Cisco appliances, devices, and services that are integrated; for example, registered, with Threat Grid Appliance via the Cisco Sandbox API.)

All sample submissions on Threat Grid Appliances are Public by default, and can be viewed by any other appliance user, including integrations, regardless of the organization to which they belong. All appliance users can see all details of samples submitted by all other users.

Threat Grid users may also submit Private samples to the Threat Grid Appliance, which are only visible to other Threat Grid Appliance users, including integrations, from the same organization as the sample submitter.

Privacy and sample visibility model on Threat Grid Appliances are illustrated in the table.

Figure 17: Privacy and Visibility on a Threat Grid Appliance

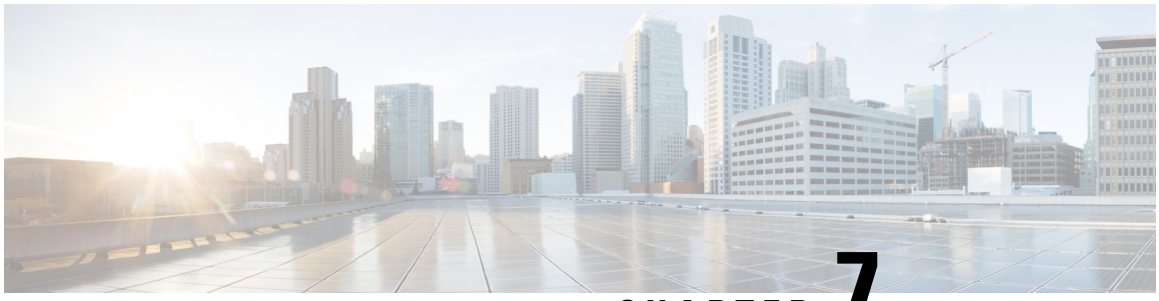
Sample and Analysis Results are visible to:	Public Submissions (Default)	Private Submissions	CSA Integration Submissions (Public by Default)
Users from the Same Organization	✔	✔	✔
Users from a Different Organization	✔	✔	✔
CSA Integrations from the Same Organization	✔	✔	✔
CSA Integrations from a Different Organization	✔	✖	✔

- **Full Access** - The green check mark indicates that users have full access to the sample and the analysis results.
- **Scrubbed Reports** - The grey check mark indicates that the Private submission results are scrubbed. Users have partial access to the sample and analysis results, but all potentially sensitive information about the sample is removed. There are no filenames, process names, screenshots, or even specifics about its activity in the glovebox.

  We omit details from the Metadata section, such as the sample submitter's login information. If you encounter a hash from a private sample in the course of doing business, this will let alert you to known threats, and if you need more details, submit your own copy of the sample for full analysis.

  Private samples may not be downloaded. Scrubbed reports include Artifacts (with filename removed), Behavioral Indicators, Domains, and IPs.
- **No Access** - The red X indicates that users have no access to the sample or the analysis results.

The same basic privacy rules apply to Threat Grid appliance integrations with AMP for Endpoints Private Cloud.



CHAPTER 7

Wipe Appliance

This chapter describes how to use the Wipe Appliance boot option. It includes the following topics:

- About Wipe Appliance Option, on page 43
- Wipe Appliance Procedure, on page 43
- Wipe and Clusters, on page 45

About Wipe Appliance Option

The Wipe Appliance boot option (v1.4.4 or later) enables you to wipe the disks on a Threat Grid Appliance to remove all data from the appliance prior to decommissioning or returning it to the Cisco Demo Loan Program.

Important After performing the wipe appliance procedure, the appliance will no longer operate without being returned to Cisco for reimaging.

Wipe Appliance Procedure

Perform the following steps to wipe the appliance:

Procedure

Step 1 Reboot your appliance.

During the boot, there will be a 4-second window in which you can select **Wipe Appliance**.

Figure 18: Wipe Appliance Option



Step 2 Enter the following information:

- **Username** - wipe
- **Password** - I ACCEPT ALL RESPONSIBILITY FOR THIS ACTION

Step 3 Select a Wipe option:

Figure 19: Wipe Options

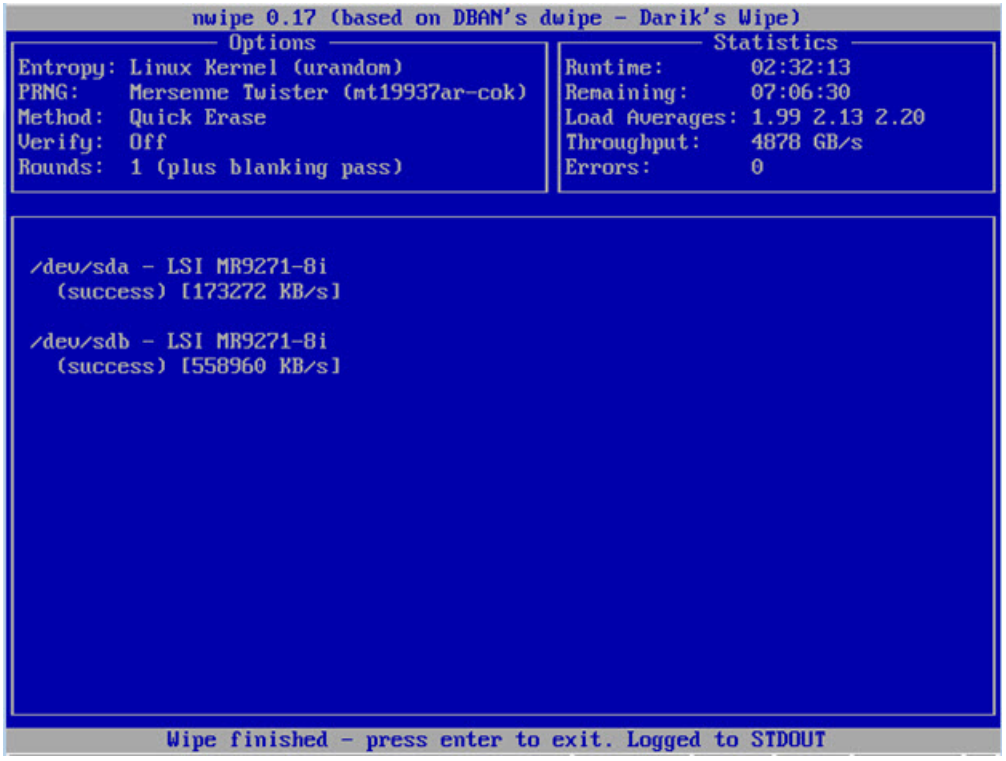


- Wipe (Fast: Zero Disks) - 2.5 hours approximate run time.

- Wipe (3-pass DOD method) - 16 hours approximate run time.
- Wipe (Random Overwrite) - 12 hours approximate run time.

The Wipe Finished screen is displayed when the wipe operation is complete.

Figure 20: Wipe Finished

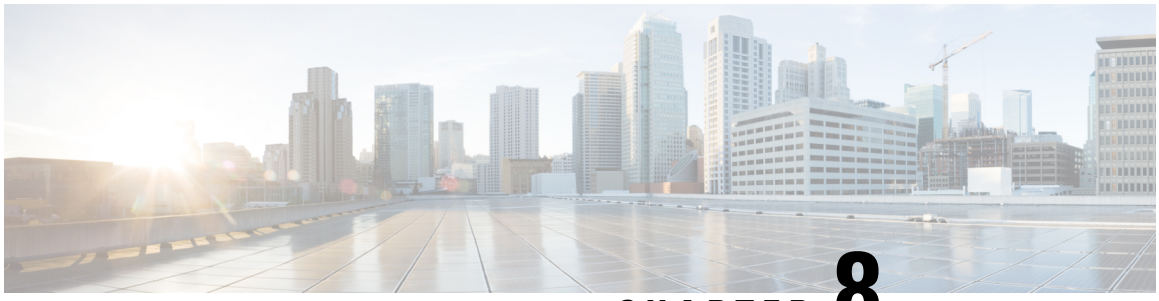


Step 4 Press **Enter** to exit.

Wipe and Clusters

After performing a wipe operation, the appliance will no longer operate without being returned to Cisco for reimaging. Wipe should only 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; instead, wipe it and then 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 removed, not just inactive.



CHAPTER 8

Backups

This chapter describes Threat Grid Appliance requirements, expectations, data retention policy, and procedures for backups and restore. It includes the following topics:

- Threat Grid Appliance Backups, on page 47
- NFS Requirements, on page 48
- File System, on page 48
- Backup Storage Requirements, on page 48
- Backup Expectations, on page 49
- Backup Data Retention, on page 50
- Backup Process Overview, on page 50
- Reset Threat Grid Appliance as Backup Restore Target, on page 51
- Restore Backup Content, on page 53
- Backup-Related Service Notices, on page 54

Threat Grid Appliance Backups

Threat Grid Appliances (v2.2.4 or later) support encrypted backups to NFS-backed storage, initialization of data from such storage, and reset to an empty-database state into which such a backup can be loaded.

Note Reset is different from the Wipe Appliance process; it is used to allow an appliance to be shipped off customer premises without information leakage, and is for backup preparation. The wipe process appropriate for that purpose already exists in the recovery bootloader, but is not suitable for preparing a system to restore a backup.

Content is encrypted with gocryptfs, a third-party open source product.

Note Filename encryption is disabled for performance reasons. Samples and other content in Threat Grid are not stored with their original names under any circumstances so this does not leak customer-owned data.

We strongly encourage consulting the documentation prior to use. Extended documentation regarding the backup functionality is available, and we strongly encourage consulting it prior to use. For additional technical information and instructions see the *Backup Notes and FAQ*.

NFS Requirements

The following NFS requirements must be met for encryped backups to NFS-backed storage:

- Must be running the NFSv4 protocol over TCP, accessible from the appliance admin interface.
- Configured directory, must be writable by nfsnobody (UID 65534).
- The NFSv4 server must be accessible via the Admin 10-Gb interface.
- Sufficient storage (see Backup Storage Requirements).
- The following mount parameters are unconditionally used: `rw`, `sync`, `nfsvers=4`, `nofail`



Note These parameters do not need to be entered manually, and manually entering any parameters that conflict with them is explicitly unsupported and may result in undefined behavior.

- Invalid NFS configuration (or configuration pointing the service at an incorrectly-configured NFS server) will generally cause the process of applying configuration to fail. Correcting this configuration in OpAdmin and reapplying should result in success.
- Exposing files for write by **nfsnobody** is secure. The only processes on the Threat Grid Appliance running as **nfsnobody** or with write to **nfsnobody**, are those responsible for encryption of data. Plain text data is exposed under distinct user accounts for different subtrees according to principal of least privilege; the PostgreSQL service on the appliance cannot access Elasticsearch data or the freezer; the Elasticsearch service cannot access PostgreSQL or freezer data.
- Using the **nfsnobody** account simplifies configuration, preventing the need to build an **idmap.conf** for each customer site, mapping local and remote account names together.

File System

Threat Grid Appliance (manufactured with v2.7 and later) use XFS as the primary file system, instead of the ZFS file system that was used on older appliances that have not been reset. This change does not affect pre-existing appliances except as otherwise noted (see Data Reset Process for more information).

Backup Storage Requirements

Total storage required for a backup store should not require more than 5.6 TB. A backup store consists of the following components:

- **Object Store** - This is normally the bulk of the storage in use. Data retention for the bulk component of a backup store follows the same policies and limits documented for the appliance release in use and places maximum storage use for this component as 4.1 TB. See the Threat Grid Appliance Data Retention Notes.

- **PostgreSQL database store** - This contains two full backups of the PostgreSQL store, and a chain of WAL logs sufficient to allow replay from the oldest of the retained full backups. This should be less than 500 GB in total.
- **Elasticsearch snapshot store** - This should be less than 1 TB in total.

Backup Expectations

The following backup expectations should be considered.

Included in Backup

The initial release of the Threat Grid Appliance backup process includes the following customer-owned bulk data:

- Samples
- Analysis results, artifacts, flagging
- Application-layer (not OpAdmin) organization and user account data.
- Databases (including users and organizations)
- Configuration done within the Face or Mask portal UI

Not Included in Backup

The following is not included in the initial release of the Threat Grid Appliance backup process:

- System logs
- Previously downloaded and installed updates
- Configuration inside the appliance OpAdmin interface, including SSL keys and CA certificates

Other Expectations

Other considerations about the backup process include:

- PostgreSQL - PostgreSQL base backup takes place on a 24-hour cycle. Database backup cannot be restored, and a warning will be displayed, until this has successfully completed at least once.
- Elasticsearch - Elasticsearch backup takes place incrementally, once every 5 minutes.
- Freezer - Freezer backup takes place on an ongoing basis, with a job following behind every 24 hours to handle any objects which were missed from the ongoing backup.
- New Key Generation - Generating a new key creates a new, independent backup store. Like the original, this new store is not valid until a base backup has taken place on a 24-hour cycle.

Backup Data Retention

During a backup, data is retained as follows:

- **PostgreSQL** - The last two successful backups and all WAL segments since those backups are retained.
- **Elasticsearch** - The last two 5-minute snapshots are retained.
- **Bulk Storage** - The same retention policy followed and documented for a single appliance is used for the shared store.

For customers who wish to retain historical data for longer periods, making use of a NFS server with filesystem- or block-layer snapshot support is strongly recommended.

For Threat Grid Appliances, v2.7.2 and later, database base backups are only retained until a new base backup has been successfully created.

Strictly Enforce Retention Period Limits

A **tgsh** configuration option, **strict_retention**, is available (v2.6 or later) that allows you to strictly enforce the retention period limit by not storing artifacts from analysis for more than fifteen (15) days. When set, files will be deleted during the first nightly pruning on which they are more than 15 days old.



Note The time period of 15 days cannot be configured or changed.

Artifacts refers to the samples themselves and other things generated from them. Artifacts do *not* include the analysis report HTML, which is subject to its original limits as otherwise documented. Artifacts also do not include database entries and search indexes.

The **strict_retention** option is disabled (false) by default. To enable the hard-pruning of artifacts after 15 days, in tgsh, set the option to *true*:

configure set strict_retention true

Backup Process Overview

The backup process on Threat Grid Appliance consists of the following steps:

Procedure

Step 1 Create the backup target directory according to the NFS Requirements.

Step 2 Complete the NFS Configuration page of the setup wizard in OpAdmin (**Configuration > NFS**), as described in the Cisco Threat Grid Appliance Setup and Configuration Guide.

Step 3 Download the encryption key that is generated once you complete the NFS configuration. You need this key to restore the backup data.

Important

The customer is responsible for backing up the encryption key and securely storing it. Threat Grid does not retain a copy. Backup cannot be completed without this key.

Step 4 Reset the backup restore target as described in Reset Threat Grid Appliance as Backup Restore Target.

Step 5 Restore the backup data as described in Restore Backup Content, on page 53.

Backup Frequency

The backup frequency of data is as follows:

- For bulk storage of samples, artifacts and reports, content is continuously backed up. Additionally, a pass is performed to look for and transfer missing content on a 24-hour cycle.
- For the PostgreSQL database, a base backup is created on a 24-hour cycle, and incremental content is continually added thereafter, either as soon as a 16-MB threshold of newly-written database content is reached, or not less than once every 5 minutes.
- For the Elasticsearch database, content is incrementally added to the backup store on a 5-minute cycle.

Backup frequency cannot be controlled or tuned because doing so would make estimates regarding storage usage, restore-process time, and performance overhead invalid.

Reset Threat Grid Appliance as Backup Restore Target

Caution Leveraging this process will destroy customer-owned data. Be very careful, and very certain. Read through all of the documentation before working any tasks.

Before an appliance can be used as a restore target, it must be in a preconfigured state. Appliances ship in this state. However, getting one back to the preconfigured state once it has been configured requires explicit administrative action.

Note Reset is not the same as the secure wipe that is available in recovery mode; only the recovery-mode secure wipe is appropriate to completely remove customer-owned data from a machine before shipping it to a DLP reimaging center. However, the secure wipe in recovery mode is not a replacement for this reset: secure wipe renders an appliance unusable until reimaged, while this reset prepares an appliance to restore a backup.

Data Reset Process

The data reset process was updated in Threat Grid Appliance v2.7 and later and is now more comprehensive. While the Wipe process (in the recovery bootloader menu) is still required for a firm guarantee of the destruction of all customer-related data, the reset process now clears operating system logs and other state which was previously left in place.

A successfully reset Threat Grid Appliance now has a new randomly-generated password displayed on its console (identical to behavior in newly-installed state). This improved process now reboots multiple times, and can be invoked from recovery mode (as opposed to the prior process, which could only be successfully invoked when booted into regular operation).

If a Threat Grid Appliance has its data reset, the datastore will be changed from a ZFS file system to a XFS file system. This improves forward compatibility and provides OS-level support for I/O usage monitoring on a per-service basis.

The data reset process now also requires sufficient storage to contain all content necessary for a fresh install on the system SSDs. Any pre-existing data is only deleted after the presence and validity of this content has been ensured. It is possible that systems that have been in use for an extended period (particularly first-generation hardware), may not have sufficient space immediately available. If this is the case, customer support can assist, if needed.

Return Target Appliance to Preconfigured State

If you are not restoring to a system fresh from manufacturing, the restore target appliance must be returned to the preconfigured state by clearing pre-existing data and NFS-related configuration from the system.

Procedure

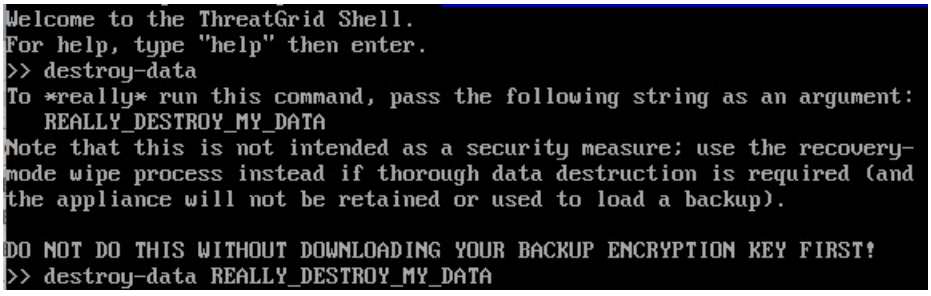
Step 1 Access the tgsh-dialog configuration interface, either via the appliance TTY or via SSH.

Step 2 Select the CONSOLE option to enter tgsh. (Note that entering tgsh via recovery mode is not suitable for this use case.)

Step 3 At the tgsh prompt, enter the command `destroy-data`. Carefully read and follow the instructions provided with the prompt.

Caution
There is no *Undo* from this command. All data will be destroyed.

Figure 21: The destroy-data REALLY_DESTROY_MY_DATA Command and Argument



The following data is destroyed:

- Samples
- Analysis results, artifacts, flagging
- Application-layer (not OpAdmin) organization and user account data
- Databases (including users and organizations)

- Configuration done within the Face or Mask portal UI
- NFS configuration and credentials
- The local copy of the encryption key used for NFS

Another Appliance Target Actively Writing to Backup Being Restored

If another system or appliance is actively writing to the backup that is being restored, for example, a test restore of content being written by a second master appliance actively used in production, return that appliance to the preconfigured state.

Procedure

Step 1 Generate a consistent, writable copy of the datastore.

Step 2 Point your appliance that is doing the test restore to this writable copy instead of to the store which is being continuously written.

Once the appliance is in a preconfigured state, it can function as the target for the backup store as described in Restore Backup Content.

Restore Backup Content

Important The system is unavailable for sample submission during the restore process.

Perform the following steps to restore the backup content:

Note The system is unavailable for sample submission during the restore process.

Procedure

Step 1 In the OpAdmin portal, click **Configuration** > **NFS** to open the NFS Configuration page.

Step 2 Click **Upload** to retrieve the backup key previously generated when configuring the server on which the backup was created.

If the key correctly matches the one used to create a backup, the **Key ID** displayed in OpAdmin portal should match the name of a directory in the configured path. The install wizard checks for a directory matching the backup key, and if it finds one, begins restoring the data to that location.

Note
There is no progress bar. The amount of time required to restore data depends on the size of the backup and other factors. In testing, a 1.2-GB restore is quick, while a 1.2-TB restore required over 16 hours. For large restores it may appear that the install has hung so be patient. OpAdmin will report that the restore has succeeded, and the appliance will start up.

Step 3 Confirm that the restored data looks the same as the original data.

Backup and Restore Notes

- Sample submission is unavailable during the restore process.
- Backups can only be restored from the setup wizard.
- Set up the same NFS store as used previously, and the same encryption key as used previously, with a process identical to the original.
- The act of setting up an appliance with a prior NFS store and encryption key will trigger a restore.
- To test the restore process on a different Threat Grid appliance while your primary appliance is still operational, make a copy of a consistent snapshot of the backup store, and point a new appliance (with the encryption key uploaded) at that copy.



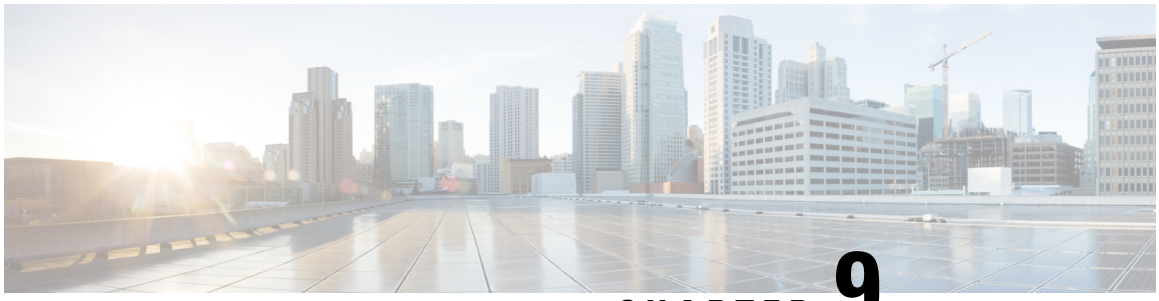
Important Only one server can be running with data from a given backup store active at a time.

Backup-Related Service Notices

The following service notices may be displayed during the backup process:

- **Network storage not mounted** - Check that the network file system being used as a backend is fully operational, and try reapplying configuration through OpAdmin or rebooting your appliance.
- **Network storage not working** - Check that the network file system being used as a backend is fully operational; if the system does not recover within 15 minutes of correcting any problems with the NFS server, try rebooting your appliance.
- **Backup file system access failure** - Contact customer support.
- **No PostgreSQL backup found** - This is a normal condition between the point in time when a backup store has been configured and the point in time when the first base backup (run automatically on a 24-hour cycle) takes place. Note that until this is complete, a backup is not considered complete and cannot be restored. If and only if this message persists for more than 48 hours, contact customer support.
- **Newest PostgreSQL base backup more than two days old** - This indicates that the system has not been successful in generating a new base backup for PostgreSQL. If unremediated, this can result in unbounded usage on the backup store (to retain a full chain of writes necessary to restore from an increasingly-old backup point), and unacceptably long processing time needed for a restore to take place. Contact customer support.
- **Backup Creation Messages** - These reflect errors detected when starting or triggering a backup.

- **ES Backup (Creation) Inactive** - Indicates that when Elasticsearch was started, the backup store was unavailable. This can be remediated by rebooting the appliance, or (if NFS and the encryption service are now functional) by logging into tgsh and running the command service restart elasticsearch.service.
- **Backup Maintenance Messages** - These reflect errors detected when checking status of previously-created backups.
- **ES Backup (Maintenance) snapshot (...) status FAILED** - This indicates that in the most recent attempt to update the backup of the Elasticsearch database, no indices could be successfully written. Check that the NFS server is functional and has free space; if no issue can be identified and the issue persists, contact customer support.
- **ES Backup (Maintenance) snapshot (...) status INCOMPATIBLE** - Should only occur immediately after an appliance upgrade installing a new version of Elasticsearch; will be displayed until the backup store has been upgraded to be compatible with this new release. Restoring from an INCOMPATIBLE backup may require customer service assistance, should a failure occur while in this state.
- **ES Backup (Maintenance) snapshot (...) status PARTIAL** - Contains one of two messages in the body: No prior successful backups seen, so retaining. (if we're keeping a partial backup as better than none at all); or Prior successful backups exist, so removing. (if we're discarding that partial backup with the intent to retry later).
- **ES Backup (Maintenance) - Backup required (...) ms** - Occurs if a backup requires more than 60 seconds. This is not necessarily an error: Elasticsearch performs periodic maintenance which can cause significant write load even on idle systems. However, if it takes place consistently when under periods of low load, investigate storage performance or contact customer service for assistance.
- **ES Backup (Maintenance)** - Unable to query Elasticsearch snapshot status - Elasticsearch could not be contacted; and this failure took place after a backup creation was successfully started. Generally, this will occur in conjunction with other appliance failures, and remediation should focus on those issues. If this error is seen when the appliance is otherwise fully functional and does not go away of its own accord, contact customer support.



CHAPTER 9

Clustering

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

- About Clustering Threat Grid Appliances, on page 57
- Build Cluster Overview, on page 60
- Start Cluster of Thread Grid Appliances, on page 63
- Join Threat Grid Appliances to Cluster, on page 71
- Designate Tiebreaker Node, on page 75
- Remove Cluster Node, on page 76
- Resize Cluster, on page 76
- Failure Tolerances, on page 77
- Failure Recovery, on page 77
- API/Usage Characteristics, on page 78
- Operational/Adminstrative Characteristics, on page 78
- Sample Deletion, on page 78

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 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.

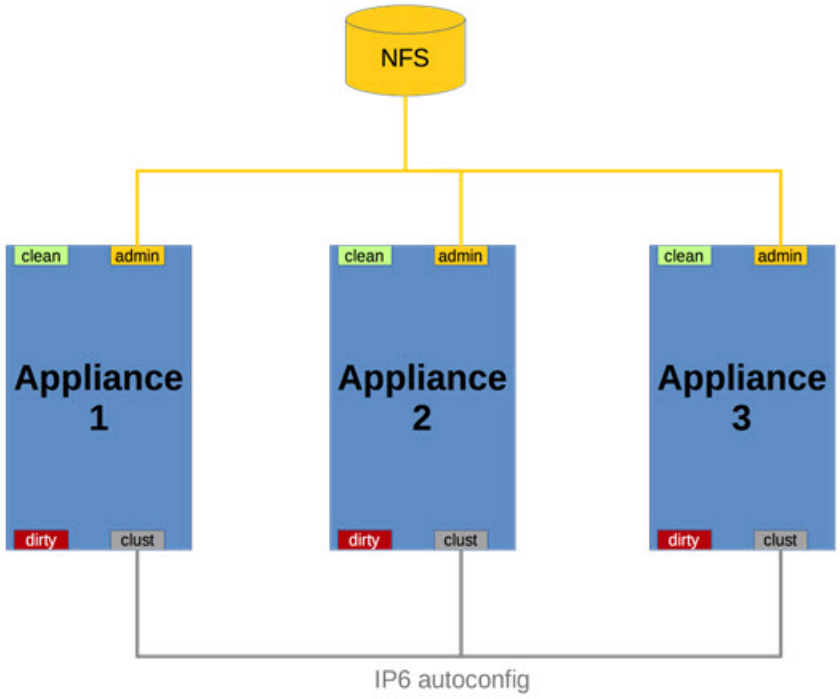
- **SSL Certificates** - If you are installing SSL certificates signed by a custom CA on one cluster member, then all other nodes' certificates should be signed by the same CA.

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 pre-existing 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 22: 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:

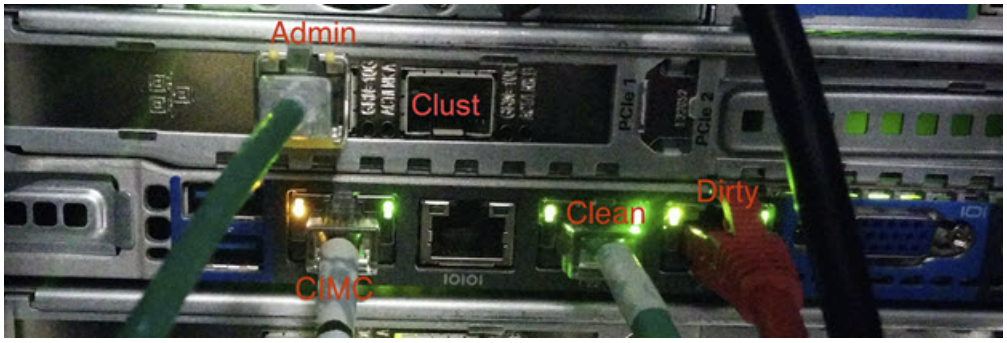
- Start a new cluster using an existing standalone appliance.
- Start a new cluster using a new appliance.

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.

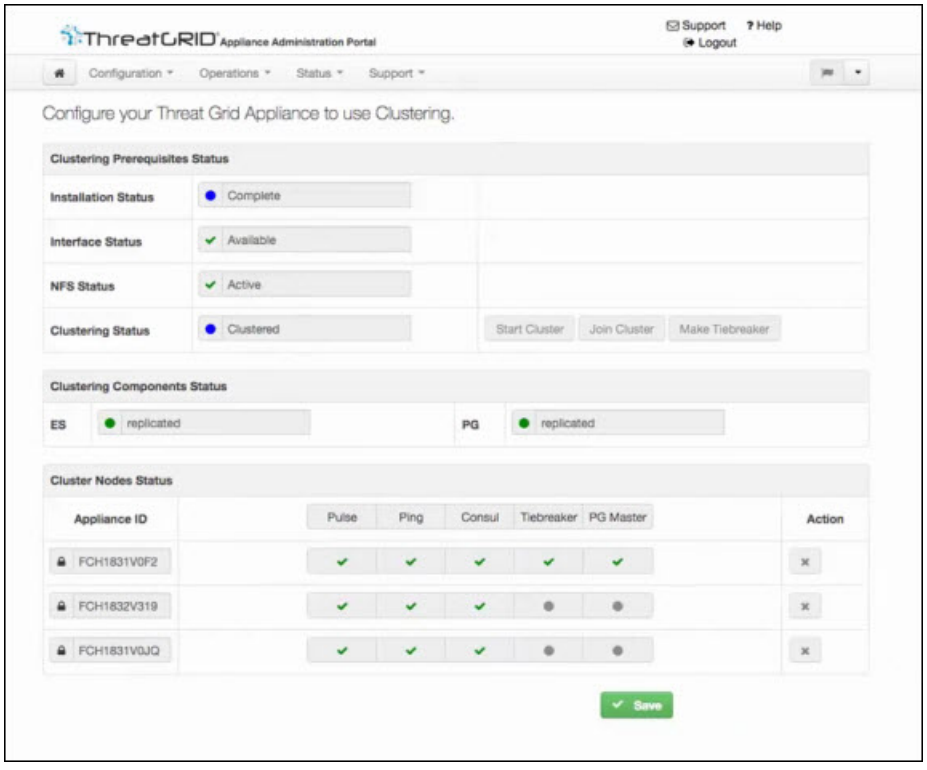
Figure 23: Clust Interface Setup for Cisco UCS M4 C220



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 24: Clustering Configuration for Active Cluster



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:

Procedure

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.

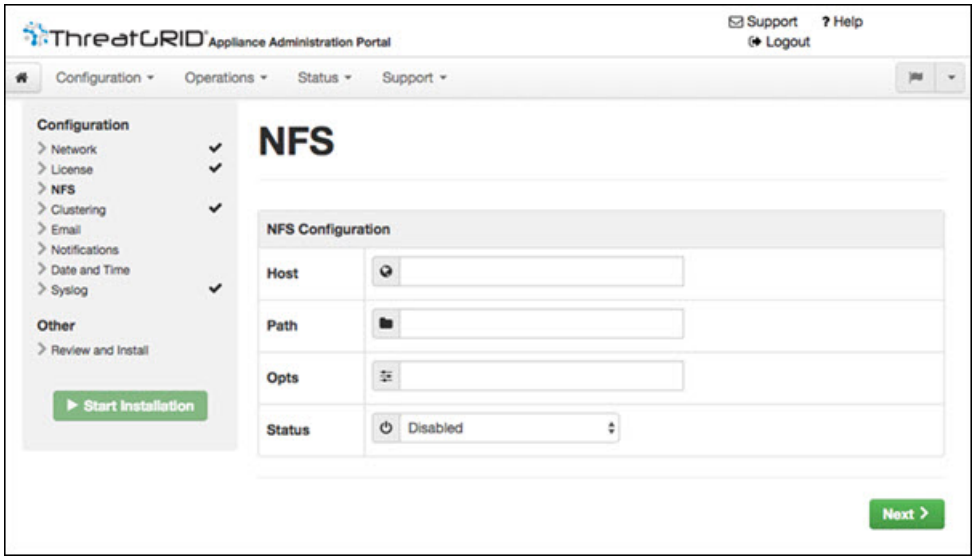
Step 2 If not already done, set up backup of the machine to NFS:

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 25: NFS Configuration

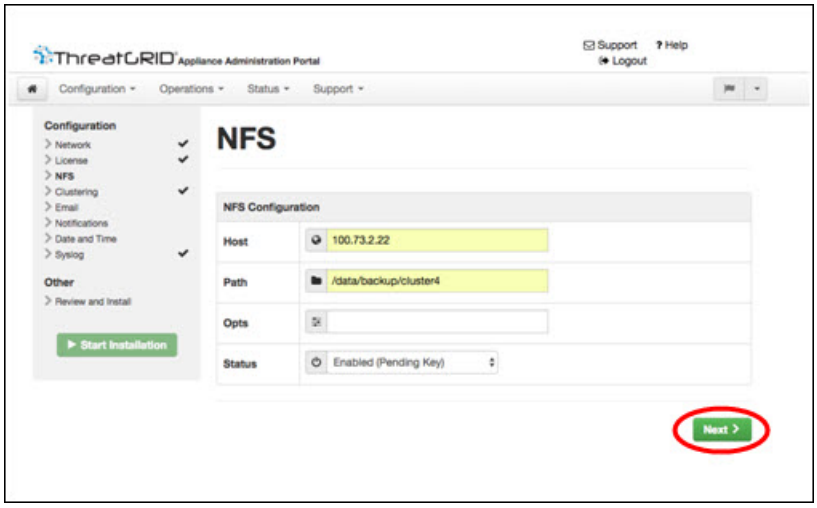


b) Complete the following fields:

- **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 26: NFS Configuration Enabled (Pending Key)



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

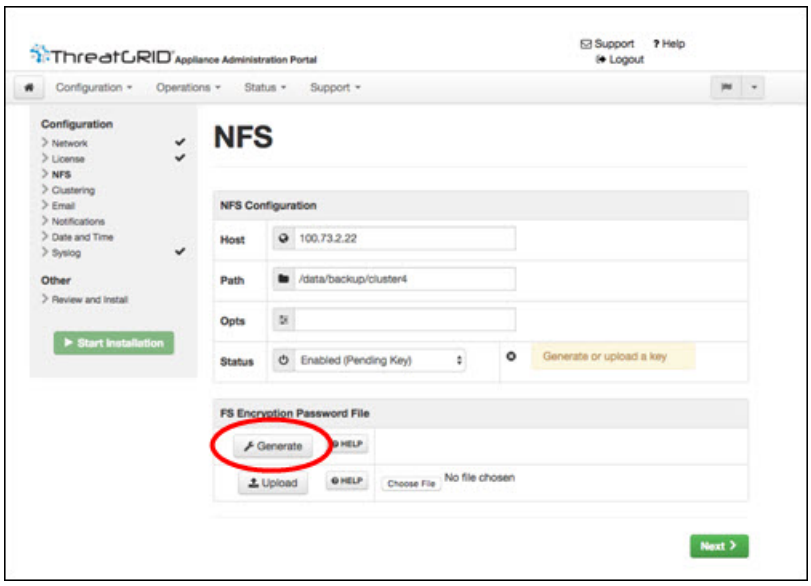
The first time you configure this page, the **Remove** and **Download** the encryption key buttons become visible.

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.

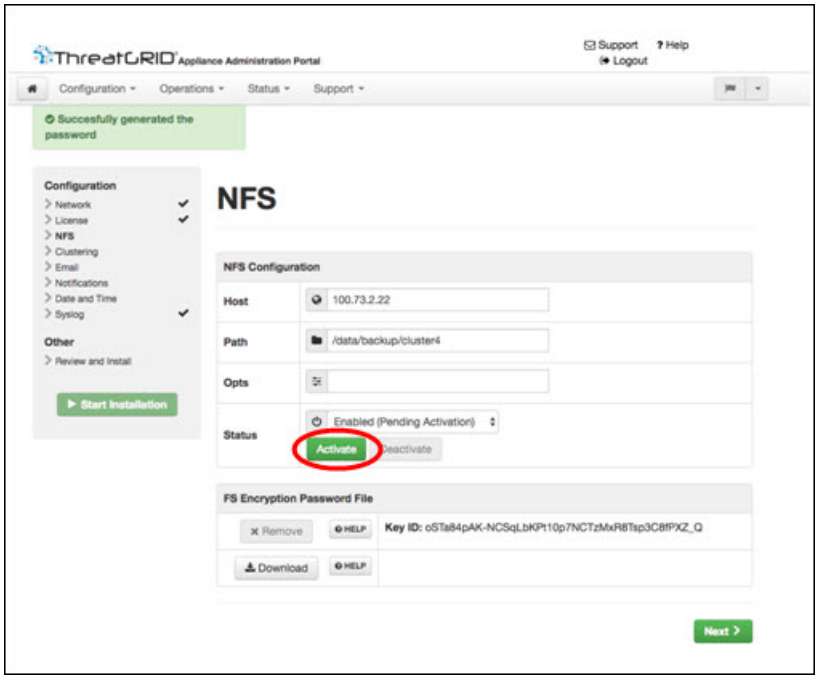
d) Click **Generate** to generate a new NFS encryption key.

Figure 27: Generate a New NFS Encryption Key



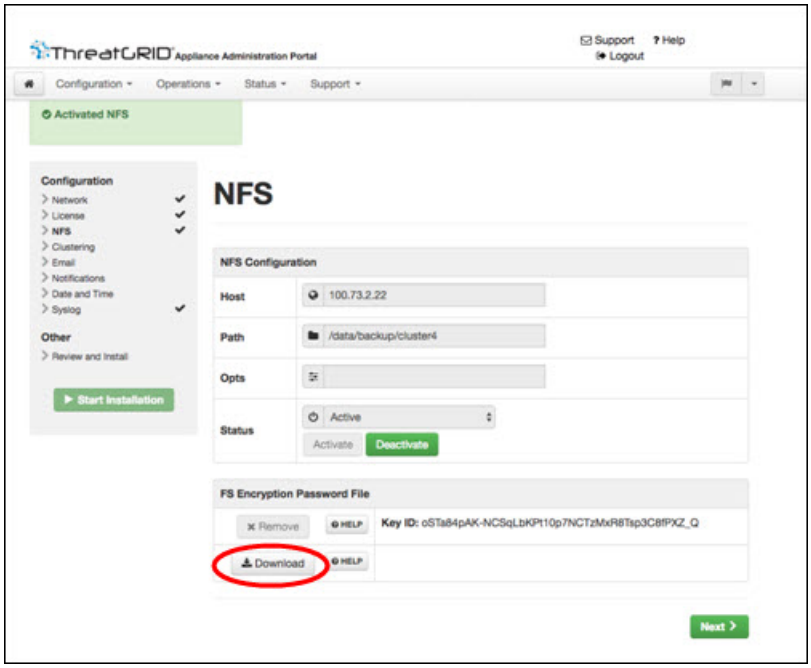
e) Click **Next.** The page refreshes and the **Key ID** is displayed, and the **Activate** and **Download** buttons become available.

Figure 28: 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 29: 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.

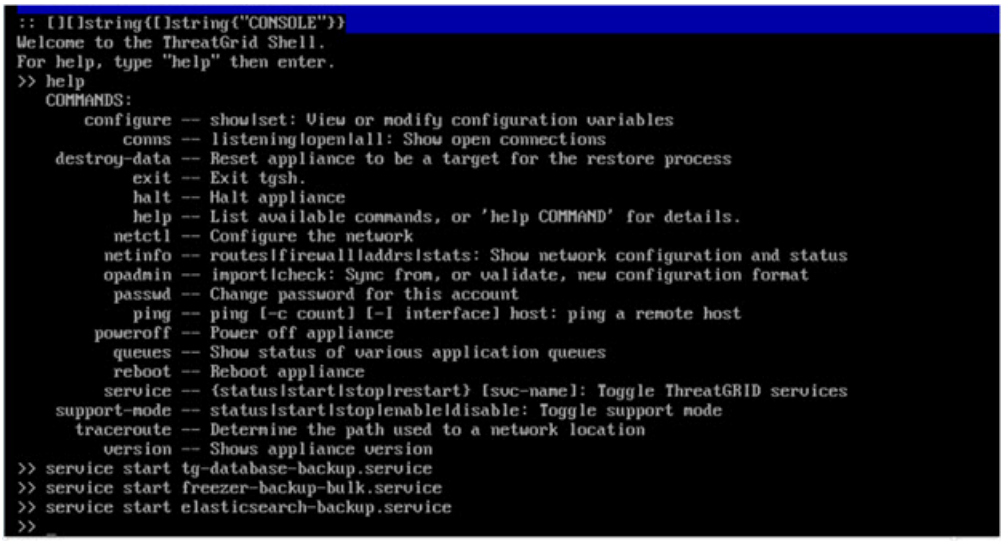
Backup and other service notices are available in the Threat Grid portal UI from the icon in the upper-right corner. If you see a service notice **There is no PostgreSQL backup yet**, then DO NOT PROCEED.

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 30: Initiating a Backup of All Data to NFS



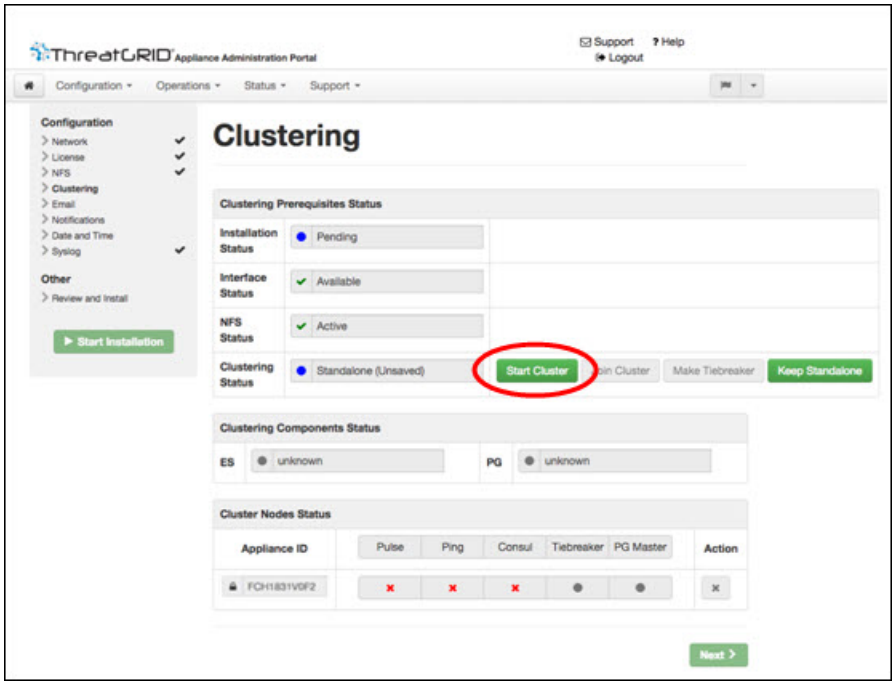
b) Wait about 5 minutes after the last command returns.

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**.

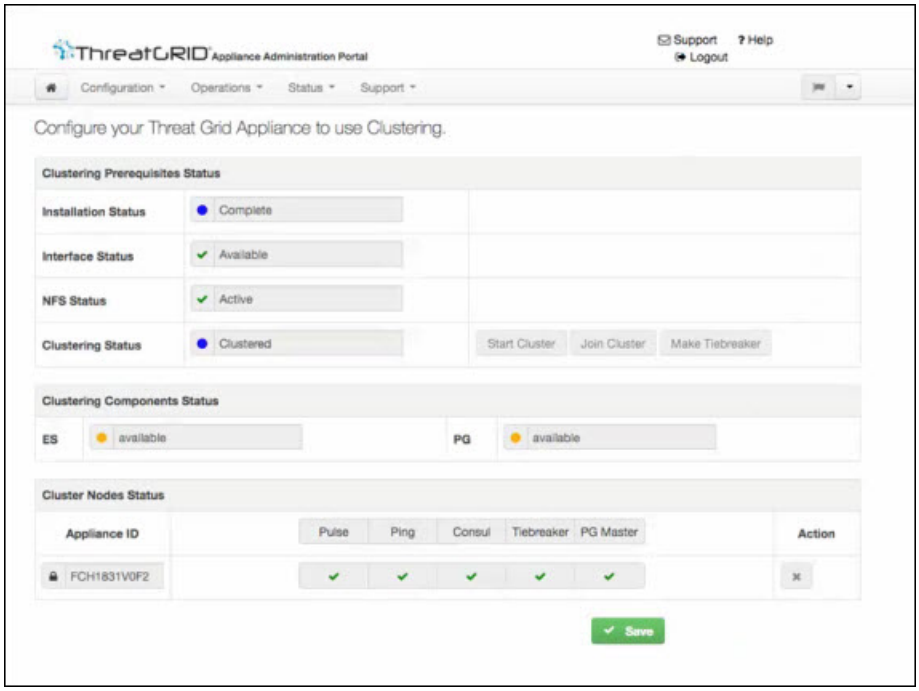
Figure 31: Start Cluster



Step 7 Click **Start Cluster**.

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

Figure 32: 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.

Procedure

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.

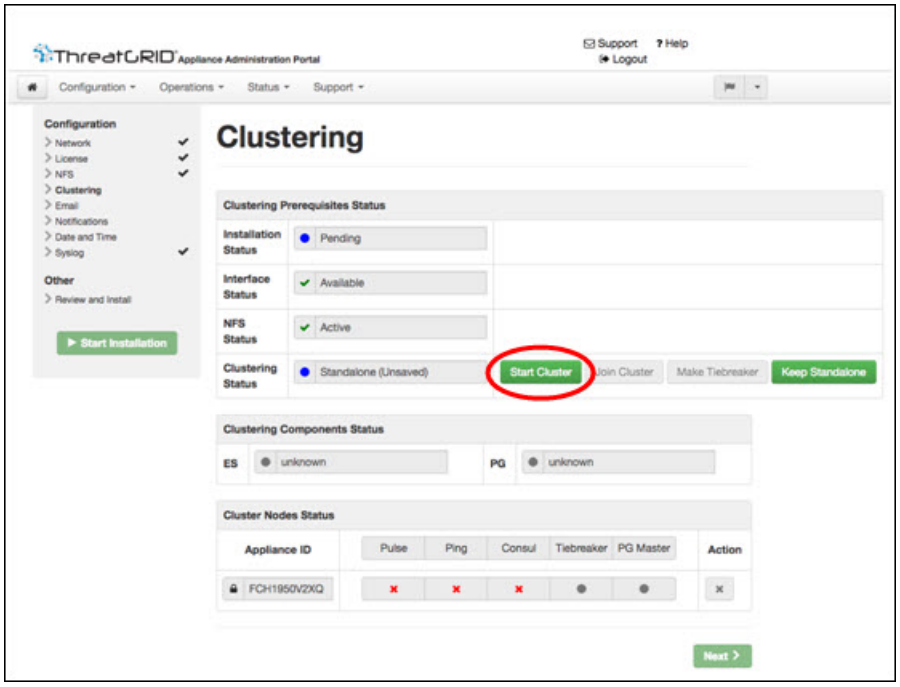
Step 6 Click **Generate** to generate a new NFS encryption key.

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 33: Clustering Configuration Page



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

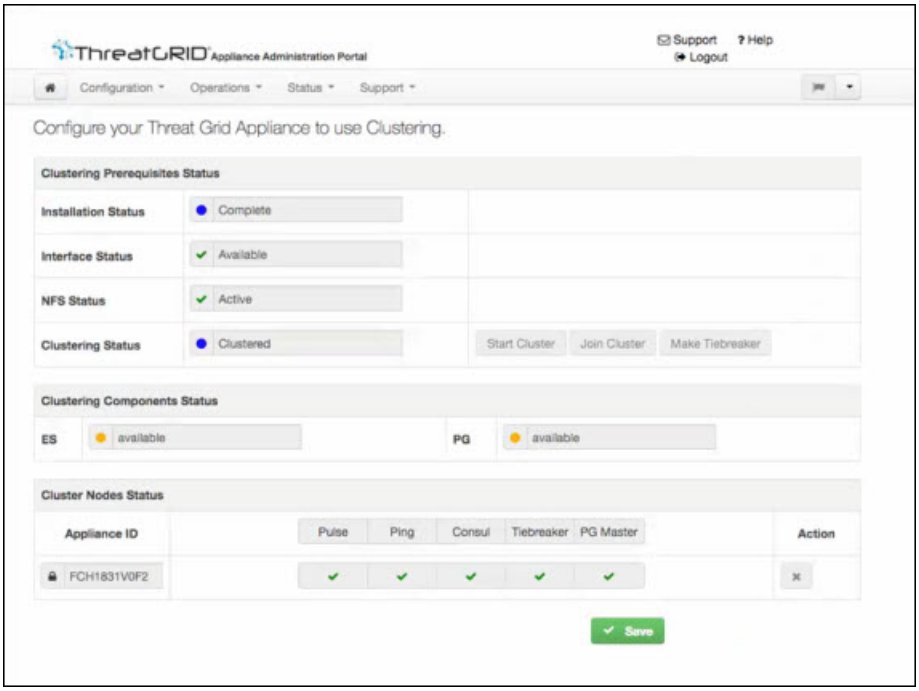
Step 10 Click **Start Cluster**, and then click **OK** on the confirmation dialog.

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 34: Clustering Status: Clustered



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:

Procedure

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:

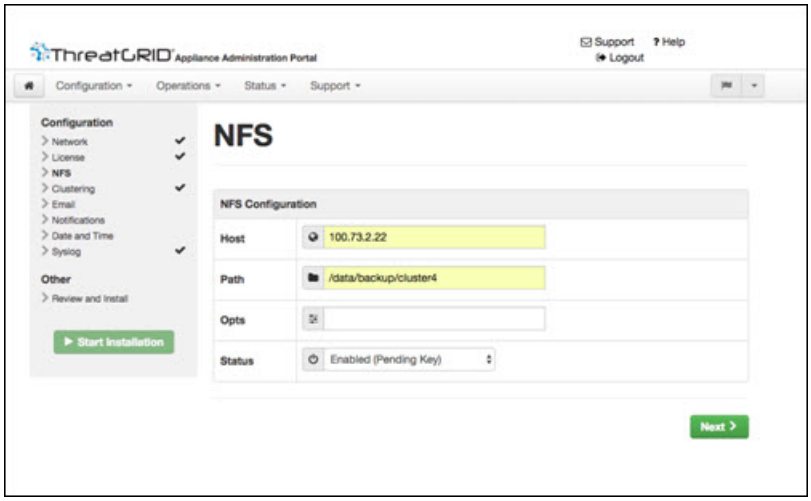
Procedure

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)**.

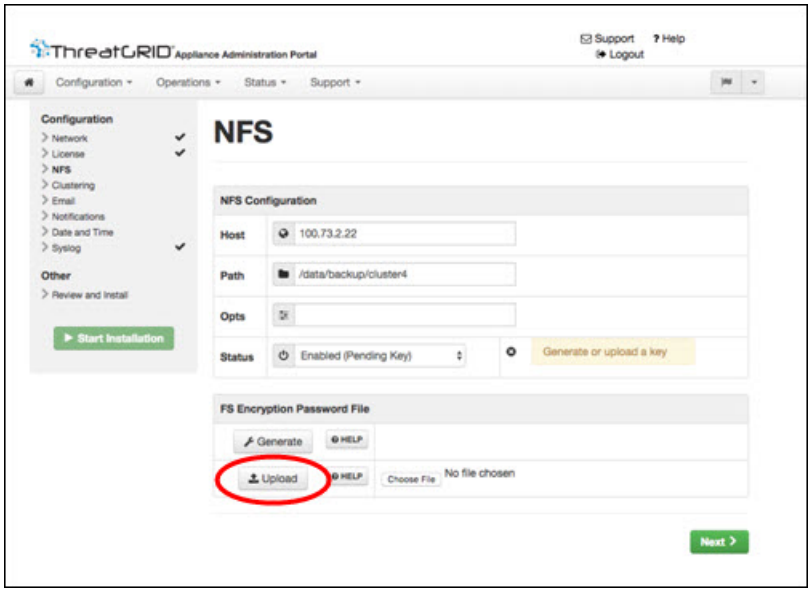
Figure 35: NFS for Joining a Cluster



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 36: Upload the NFS Encryption Key

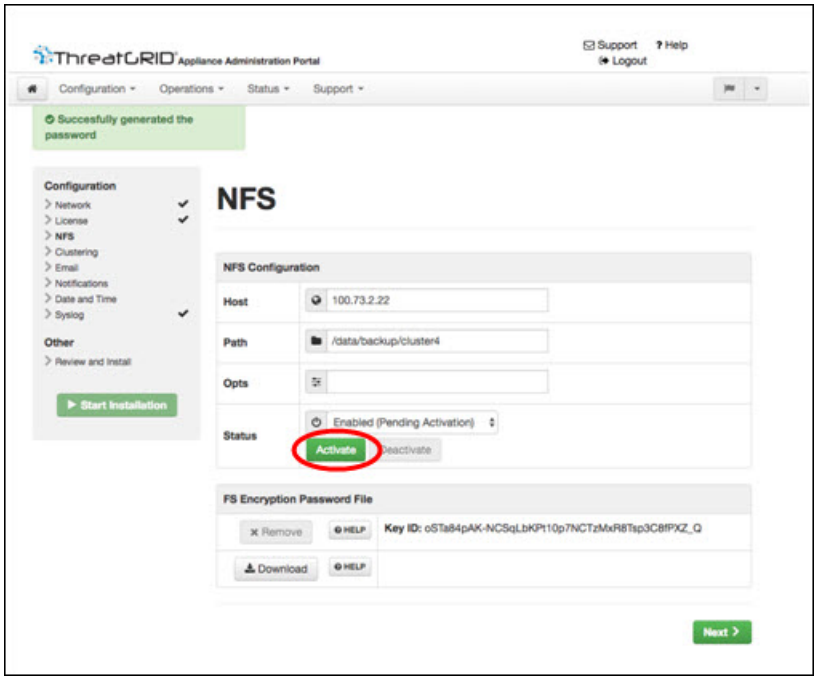


Step 5 Click **Upload** and choose the NFS encryption key you downloaded from the first node when you started the new cluster.

Step 6 Click **Next**.

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

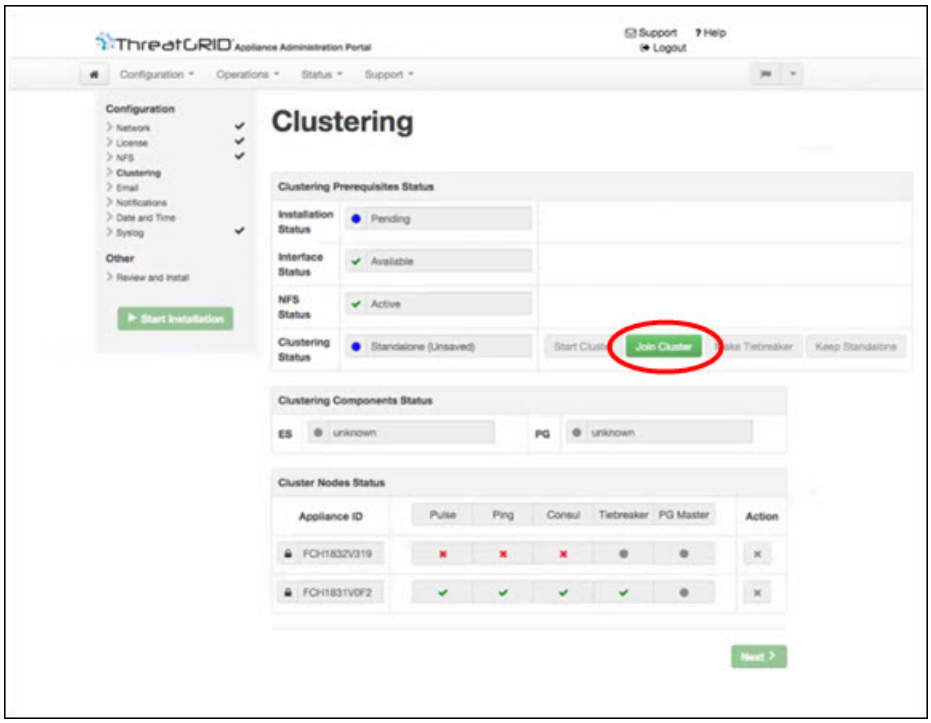
Figure 37: Activate the NFS Encryption Key of the Joining Appliance



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 38: Join Cluster

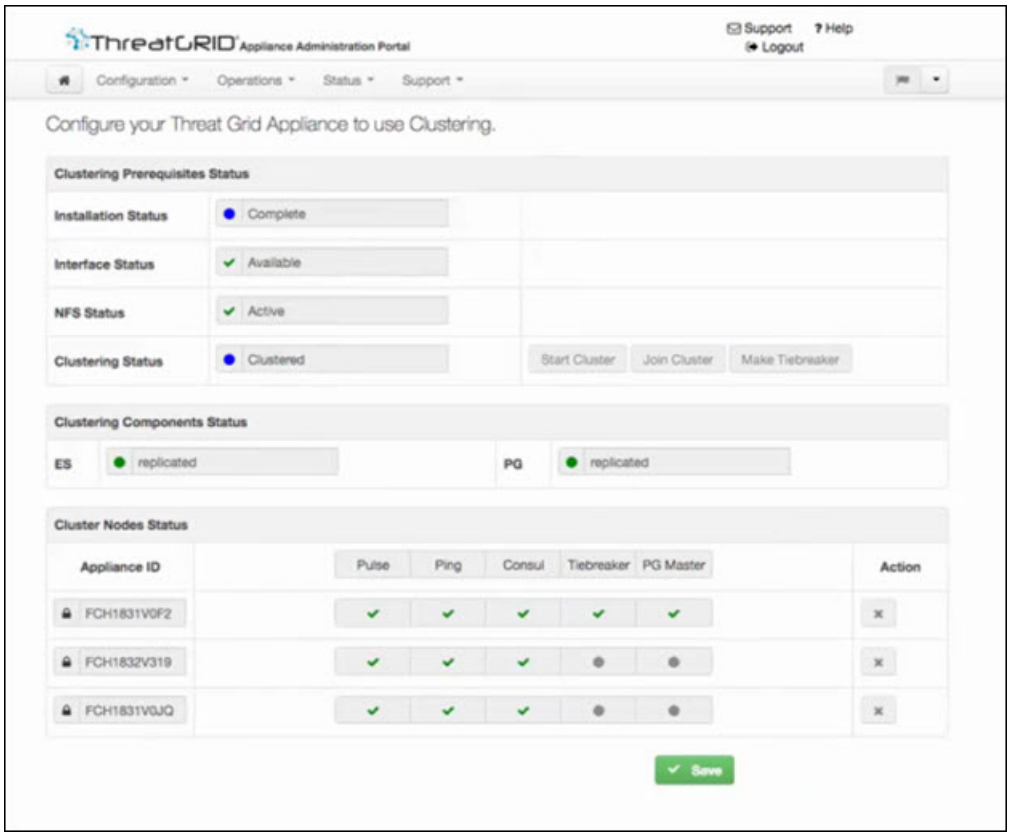


Step 9 Click **Join Cluster**. and then click **OK** on the confirmation dialog.

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

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

Figure 39: 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 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 4: 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.ccom), 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/Adminstrative 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):

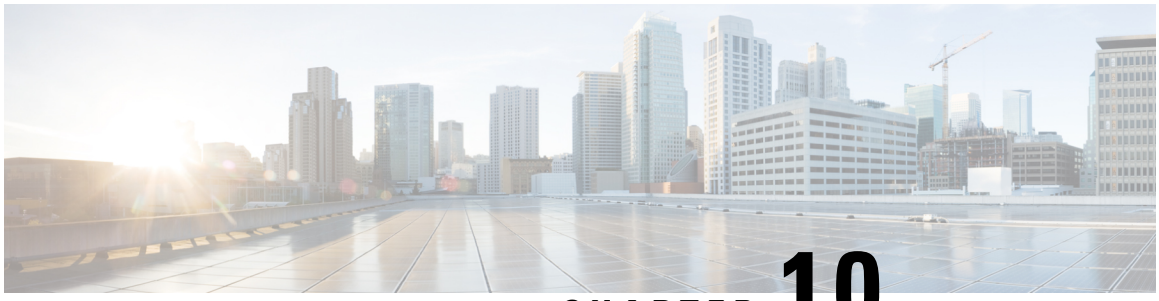
- The **Delete** option is available in the **Actions** menu in the samples list.
- The **Delete** button is available in the upper-right corner of the sample analysis report.



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.

In Threat Grid Appliance v2.7 and later, sample deletion is extended to include artifacts, which matches the behavior of the cloud product.



CHAPTER 10

Network Exit Configuration

This chapter describes the Network Exit feature and how to configure it.

- Configure Network Exit, on page 79

Configure Network Exit

Geographic location is often an important issue for malware analysis. Some types of malware behave differently depending on geographic location, while other types may target a specific area. Similar in concept to VPN, the Network Exit setting makes any outgoing network that is generated during sample analysis appear to exit from that location.

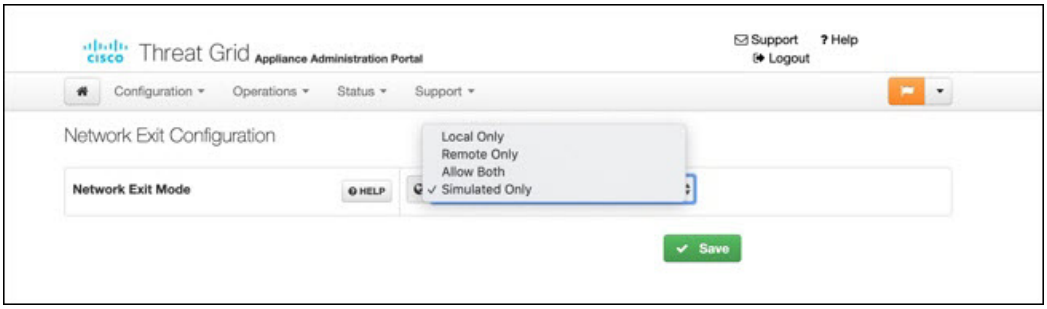
The Network Exit setting is available in v2.4.3 or later, and replaces the tg-tunnel solution. Configuration files are automatically distributed and there is no need for suport staff to manually install or update them.

Note If you were previously using tg-tunnel, you must allow outbound traffic to 4.14.36.142:21413 and 63.97.201.68:21413 before installing v2.4.3. Otherwise, that traffic only needs to be permitted before enabling remote exit use.

Procedure

Step 1 In the OpAdmin portal, click **Configuration** > **Network Exit**.

Figure 40: Network Exit Configuration



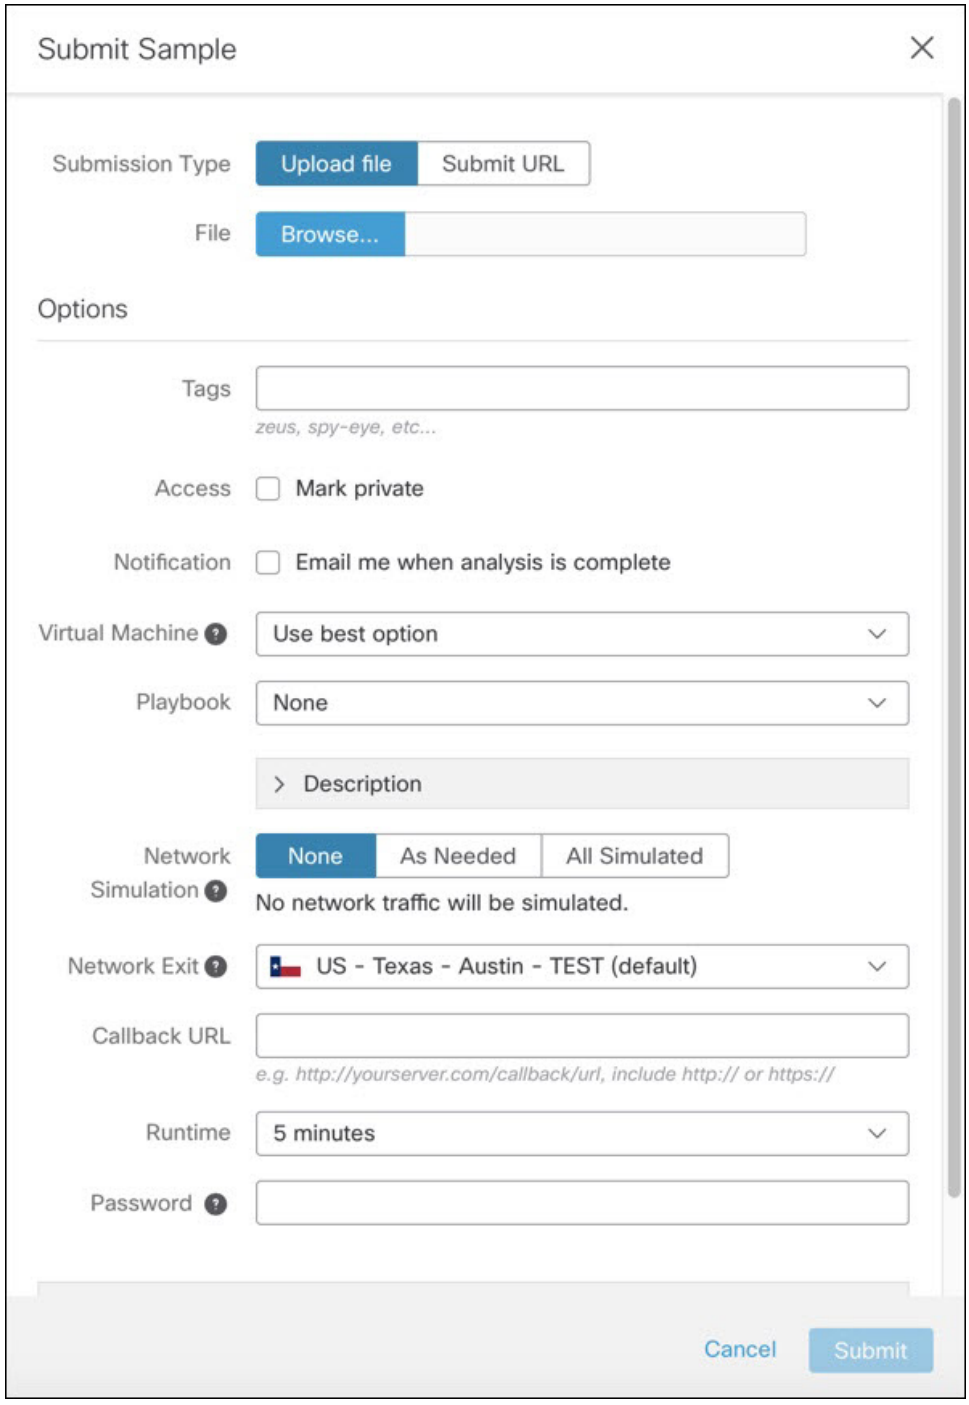
Step 2 In the **Network Exit Mode** field, choose **Local Only**, **Remote Only**, **Allow Both**, or **Simulated Only**. This field determines which Network Exit options will be available in the application, such as when submitting samples in the UI.

If you select **Local Only** or **Remote Only**, then the application will only make those options available to users.

If you select **Simulated Only**, then API and UI users cannot select any option that would send network traffic from virtual machines to destinations outside of the local appliance.

Accessing private networks, even for DNS lookup, is not allowed even for Network Exit. All malware traffic comes out of the Dirty interface, using the Dirty DNS server configured.

Figure 41: Submit Sample



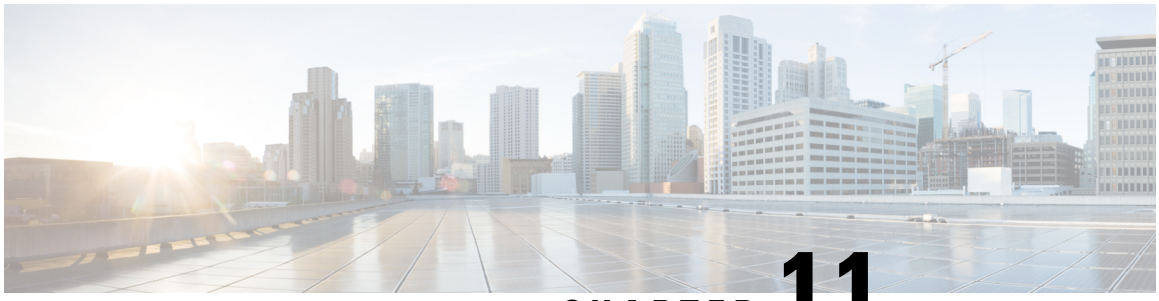
Note

Sometimes it may be necessary to simulate network connections during analysis. Network simulation provides analysts with a way to present network resources to malware samples that may otherwise be unavailable, and for other reasons. For example, you may want to select a network simulation option to simulate network connections when the upstream

servers are not accessible; when they have been taken down; when their DNS records are gone; or if other restrictions on outbound connectivity apply in order to improve sample execution and convictions.

In addition, network simulation can provide at least some connectivity to air-gapped appliances and improve sample execution on them.

The **Network Simulation** option for sample analysis is available on Threat Grid Appliances v2.7.1 and later. See the Threat Grid portal UI online help topic for additional information.



CHAPTER 11

OpAdmin Menus

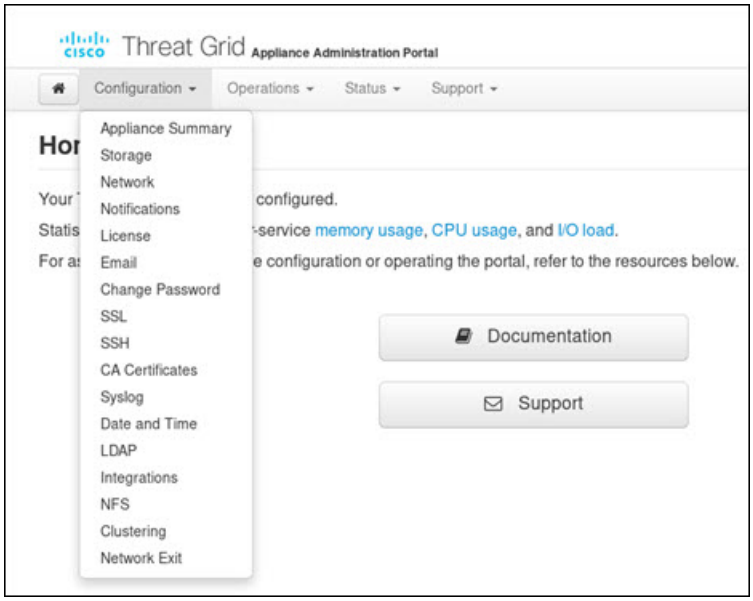
This chapter provides an overview and screenshot of each of the menus in the OpAdmin portal to illustrate the various menu options that are available for performing tasks. It includes the following topics:

- Configuration Menu, on page 83
- Operations Menu, on page 84
- Status Menu, on page 84
- Support Menu, on page 85

Configuration Menu

The Configuration menu in the OpAdmin portal provides options for configuring your Threat Grid Appliance. Any changes that need to be made to your configuration, must be done using this menu to be in edit mode.

Figure 42: OpAdmin Portal Configuration Menu

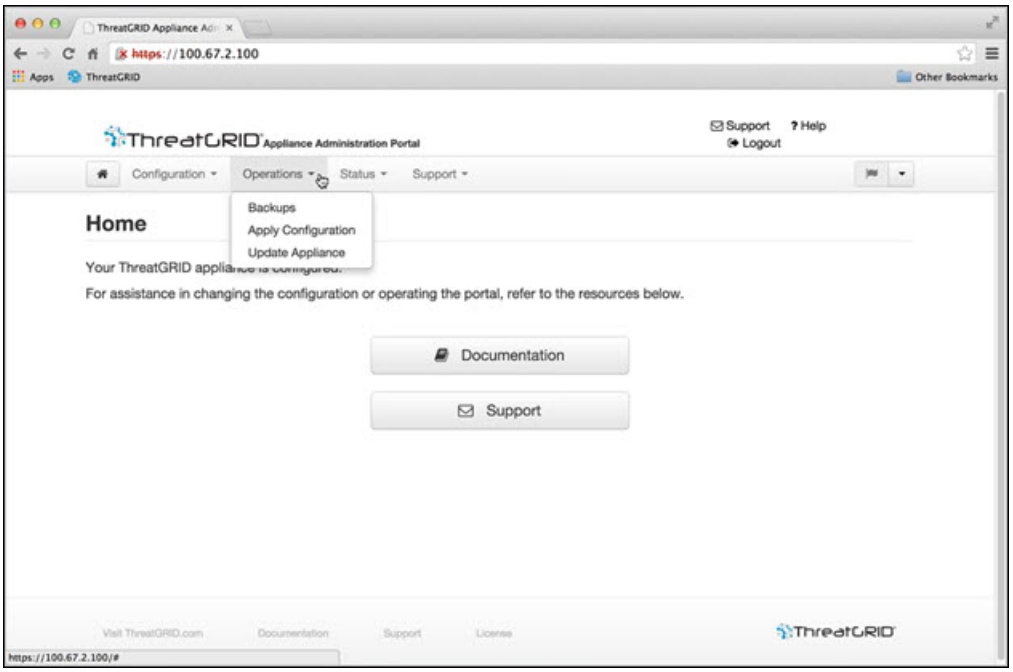


Operations Menu

The Operations menu in the OpAdmin portal provides options for backups, applying configuration, and updating the appliance.

Note Choose **Update Appliance** in the Operations menu to view the Release Notes.

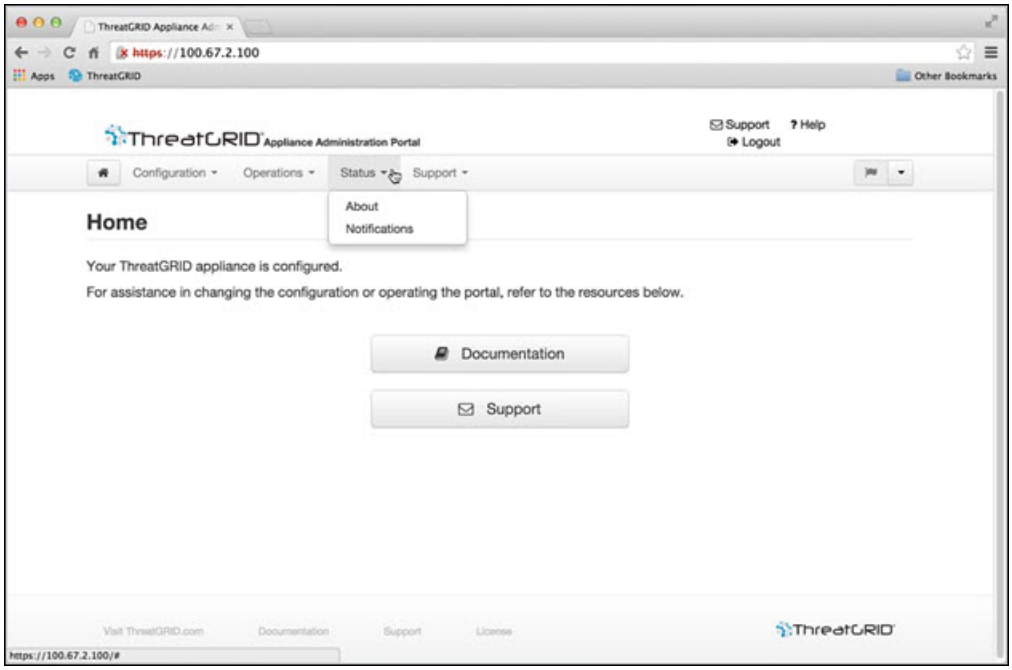
Figure 43: OpAdmin Portal Operations Menu



Status Menu

The Status menu in the OpAdmin portal is used to view the installed version and information about notifications.

Figure 44: OpAdmin Portal Status Menu



Support Menu

The Support menu in the OpAdmin portal is used to start a live support session, take snapshots of your system, and access the Threat Grid Cloud Platform.

Figure 45: OpAdmin Portal Support Menu

