THALES

Data Security Manager

DSM Administration Guide
6.4.2

DSM Administration Guide
6.4.2
Document Version 1
5/22/2020

Contents

Preface

The *DSM Administration Guide* describes how to install and configure the DSM. This document is intended for system administrators who install the appliance and connect it to a network.

Documentation Version History

The following table describes the documentation changes made for each document version.

Document Version	Date	Changes
6.1.0 v2	08/06/2018	Removed extraneous text.
6.2.0 v1	4/11/2019	GA release; HA is now active/active, new CLI commands, new API calls expanded KMIP features. This guide contains new troubleshooting and IPMI information. Added rules for hostnames.
6.2.0 v2	3/26/2019	Added rules for hostnames.
6.3.0 v1	08/09/2019	Supports Efficient Storage with VTE 6.2.0, Excluding files from encryption, fixed security vulnerabilities.
6.3.0 v2	9/16/2019	Made clearer the concept that with 6.2.x and 6.3.x user is creating a new cluster, they cannot upgrade a 6.1.x (or lower) cluster.
6.4.0 v1	12/19/2019	DSM now compatible with the Luna HSM and Smart cards, users can create GuardPoints for Cloud Object Storage devices, System admins can prevent domain admins from deleting other admins, LDAP limits raised.
6.4.2 v1	5/22/2020	Supports IDT (In-place data transformation - device) policies and GuardPoints. Luna now supports PED access and Password access. Enabled Secure Start for ESG. Added Key Governance for DSM keys, and registration tokens.

Assumptions

This documentation assumes that you have knowledge of your computer network as well as network configuration concepts. For the hardware appliance, you'll also need access to the data center where your DSM hardware appliance will be racked and cabled.

For more information about what's new in this release, refer to the *DSM Document Version 1 Release Notes*.

DSM Hypervisor Support

The virtual DSM can be installed on the following hypervisors:

- VMware 6.0 or higher
- Microsoft Hyper-V 2012R2, 2016 and 2019
- KVM
- SLES 12 SP3 with Xen v4.9

Document Conventions

The document conventions describe common typographical conventions and important notice and warning formats used in Thales technical publications.

Typographical Conventions

This section lists the common typographical conventions for Thales technical publications.

Table 2-1: Typographical Conventions

Convention	Usage	Example
bold regular font	GUI labels and options.	Click the **System** tab and select **General** Preferences.
`**bold italic monospaced font**`	variables or text to be replaced	https://`<Token Server name>`/admin/ Enter password: `<Password>`
`regular monospacedfont`	Command and code examples XML examples	Example: `session start iptarget=192.168.253.102`
italic regular font	GUI dialog box titles	The *General Preferences* window opens.
	File names, paths, and directories	*/usr/bin/*
	Emphasis	*Do not* resize the page.
	New terminology	*Key Management Interoperability Protocol (KMIP)*
	Document titles	Refer to the *DSM Administrators Guide* for information on how to administer your DSM Appliance.
quotes	File extensions Attribute valuesTerms used in special senses	".js", ".ext" "true" "false", "0" "1+1" hot standby failover

Notes, tips, cautions, and warnings

Notes, tips, cautions, and warning statements may be used in this document.

A Note provides guidance or a recommendation, emphasizes important information, or provides a reference to related information. For example:

> **Note**
> It is recommended to keep tokenization keys separate from the other encryption/decryption keys.

A tip is used to highlight information that helps you complete a task more efficiently, such as a best practice or an alternate method of performing the task.

> **Tip**
> You can also use Ctrl+C to copy and Ctrl+P to paste.

Caution statements are used to alert you to important information that may help prevent unexpected results or data loss. For example:

> ⚠️ **CAUTION**
> **Make a note of this passphrase. If you lose it, the card will be unusable.**

A warning statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data. For example:

> ⚠️ **WARNING**
> **Do not delete keys without first backing them up. All data that has been encrypted with deleted keys cannot be restored or accessed once the keys are gone.**

Hardware-Related Warnings

The following warning statement is used to indicate the risk of electrostatic discharge of equipment:

> ⚠️ **ELECTROSTATIC DISCHARGE**
> **If this warning label is affixed to any part of the equipment, it indicates the risk of electrostatic damage to the module. To prevent equipment damage, follow suitable grounding techniques.**

The following warning statement is used to indicate the risk of hazardous voltages of equipment:

> ⚠️ **HAZARDOUS VOLTAGES**
> **The warnings in this section indicate voltages that could cause serious danger to personnel.**

Sales and Support

If you encounter a problem while installing, registering, or operating this product, please refer to the documentation before contacting support. If you cannot resolve the issue, contact your supplier or Thales Customer Support.

Thales Customer Support operates 24 hours a day, 7 days a week. Your level of access to this service is governed by the support plan arrangements made between Thales and your organization. Please consult this support plan for further information about your entitlements, including the hours when telephone support is available to you.

For support and troubleshooting issues:

- https://supportportal.thalesgroup.com
- (800) 545-6608

For Thales Sales:

- https://enterprise-encryption.vormetric.com/contact-sales.html
- sales@thalesesecurity.com
- (888) 267-3732

Chapter 1: Data Security Manager (DSM)

The Data Security Manager (DSM) creates, stores and manages the policies that protect data. It is available as physical or a virtual appliance. This document describes the work flow needed to set up the DSM to protect your data. Refer to the corresponding version of the release notes for information about new features and updates.

Overview

The DSM lets you create, store, and manage policies that protect data residing on host servers (referred to as 'hosts' from here on throughout the document, unless otherwise specified). The DSM is managed by DSM administrators who access the DSM through a browser-based user interface called the Management Console.

DSM administrators manage VTE and VAE Agents that reside on host servers and protect the data on those hosts. DSM administrators specify data access policies that are sent to these agents. Policies are created, stored and managed by DSM administrators. DSM administrators specify data access policies, create new administrators and administrative domains, generate usage reports, register new hosts, and access security logs.

For high availability (HA), DSMs can be configured together in a cluster. The nodes run in parallel in an active-active format. When one node fails, the other nodes continue to run. When the failed node is back up, it resynchronizes with the other HA nodes. The HA nodes are also used for load balancing the VTE agents.

The DSM generates log entries for all configuration changes, system events, access attempts, and file system agent communications. These log entries can be sent to standard Syslog servers in several formats.

Separation of Duties

Although the main focus of the DSM is the security of your data through encryption, segregation of data, and policy-based access enforcement, a key feature of the DSM incorporates the critical IT security concept of separation of duties with regard to administration of the DSM and the VTE agents and with the overall data center operation. It is this separation of duties that enhances compliance with regulatory requirements.

The DSM allows for the creation of domains to separate administrators and the data they access from other administrators. A domain is a self-contained environment composed of keys, policies, hosts, administrators, and audit records. There are three types of administrators, each with specific roles and permissions. Segmenting administrative functions by type ensures that one administrator cannot control the entire data security process.

Domains

A DSM administrative domain is a logical entity used to separate administrators, and the data they manage, from other administrators. Administrative tasks are done in each domain based upon each administrator's assigned type. The benefits of administrative domains are:

- Segregation of data for increased security

- Separation of responsibilities

- No single administrator has complete control over the DSM and the data it protects

Two types of domains can be created; global domains and restricted or local domains.

Global domains are created at the system level and can share DSM Domain Administrators and DSM Security Administrators. DSM global domains enable different business units, application teams, or geographical locations to share the DSM's protection without having access to each other's security configuration.

Restricted or local domains are domains in which administration is restricted to Domain Administrators and Security Administrators assigned to that domain and configuration data in one domain is invisible to administrators in other domains. DSM Domain administrators in restricted domains cannot be assigned to multiple domains. Once the first Domain Administrator is created and assigned to a restricted domain, that Domain Administrator creates additional Domain Administrators and Security Administrators as required. Domain Administrators created within a restricted domain are not visible outside of that domain, and can only be created and deleted by a Domain Administrator from that restricted domain.

Administrators

The DSM is administered by a DSM System Administrator. DSM System Administrators are different from regular data center system administrators—a DSM administrator's primary responsibility is to provide data access to those who need it and block data access to those who don't need it, including other DSM Administrators and data center system administrators.

To enforce separation of duties for strict adherence to good IT security practices and standards, we recommend creating customized administrator roles for individual users such that no one user has complete access to all data and encryption keys in all domains, see "Separation of Duties" on the previous page for more information.

DSM administrators protect data by establishing data access policies, encrypting data, and auditing data access attempts.

Administrator Types

There are three primary types of administrators, each with specific roles and permissions. Segmenting administrative functions by type ensures that one administrator cannot control the entire data security process. Each Administrator is allowed to do specific administrative tasks.

The DSM provides the following three primary types of administrators:

- **System Administrators**

  System administrators create domains and administrators, and assign a domain's first administrator.

- **Domain Administrators**

  A Domain administrator, once assigned to a domain, can assign more domain administrators and security administrators to the domain or remove them from the domain.

- **Security Administrators**

  A Security administrator, once assigned to a domain (with appropriate roles), can manage hosts, keys and policies in the domain.

However, under a "relaxed security mode", combined administrator type assignments can also be configured:

- **All Administrator**

  Such administrators can operate both inside and outside of global domains. When an All administrator enters a domain, that administrator can perform Domain Administrator and Security Administrator tasks. When this administrator exits a domain, that administrator can perform System Administrator tasks.

- **Domain and Security Administrator**

  A Domain and Security Administrator can do every task that is permitted inside a domain. For example, they can add administrators to the domains of which they are a member, but they cannot create new administrators.

Additionally, any of these administrators can be created as read-only users. A read-only administrator inherits all of the privileges of the designated administrator, but without the ability to modify any settings. A read-only administrator can only view all of the configuration information available to it.

System Administrator

System Administrators operate outside of domains. They create domains and assign Domain Administrator to the domains. They create domains but do not operate within them. Domain and Security Administrators operate within those domains created by the System Administrator. The default Administrator, `admin`, is a System Administrator. This administrator creates additional administrators and domains, and then assigns one or more Domain Administrator to each domain.

Domain Administrator

Domain Administrators operate within domains. They can add additional Domain Administrators and Security Administrators to each domain. There are two types of administrative domains; global domains and restricted domains. Domain Administrators assigned to a global domain can operate within their assigned domains, but can also be assigned to multiple global domains. Global Domain Administrators who are members of multiple global domains can switch between the domains. Global Domain Administrators who are members of multiple global domains must always know what domain they are in before performing any tasks. If you log in as a Domain Administrator or a Security Administrator, and you notice that the administrator, host, or log data is wrong, you are most likely in the wrong domain.

Domain Administrators assigned to a restricted domain are restricted to that particular domain—they cannot be assigned to multiple domains.

The Domain Administrator also adds Security Administrators to a domain and assigns them roles (i.e., Audit, Key, Policy, Host, Challenge & Response, and/or Client Identity) that are applied within that domain.

Security Administrator

All tasks performed by the Security Administrator occur within domains. Security Administrators assigned to a global domain are restricted to their assigned domains, but can be assigned to multiple domains. Security Administrators that are assigned to multiple global domains have only the roles that were assigned when they were made a member of that domain. Meaning, the same administrator can have different roles in different domains.

Security Administrators assigned to a restricted domain are restricted to that particular domain, they cannot be assigned to multiple domains.

Roles are assigned by Domain Administrators when they assign a Security Administrator to a domain. A brief description of the roles is described below.

- **Audit**

  Allows the Security Administrator to view log data.

- **Key**

  Allows the Security Administrator to create, edit, and delete local key-pairs, public keys only, and key groups. Can also view log data.

- **Policy**

  Allows the Security Administrator to create, edit, and delete policies. (A policy is a set of rules that specify who can access which files with what executable during what times. Policies are described in more detail later.) Can also view log data.

- **Host**

Allows the Security Administrator to configure, modify, and delete hosts and host groups. Can also view log data. The **Challenge & Response** role is automatically selected when the Host role is selected.

- **Challenge & Response**

  Allows a Security Administrator to generate a temporary passphrase to give to a host administrator to decrypt data on the host when there is no connection to the DSM.

- **Client Identity**

  Allows a Security Administrator to create a client identity profile. A client identity is used to restrict access to encryption keys for VAE/VKM host users. See "Identity-Based Key Access" on page 198 for details about this feature.

All administrator

All Administrators can operate inside and outside of domains. To operate inside a domain, All administrators must be assigned to that domain first. When the administrator enters a domain, they can perform Domain Administrator and Security Administrator tasks. When an All administrator exits the domain, they can perform System Administrator tasks.

Domain and Security Administrator

The Domain and Security Administrator can perform every task that is permitted for an administrator from inside a domain. For example, the Domain and Security Administrator can add administrators to the domains of which they are a member, but they cannot create new administrators.

The administrator types are outlined in the following table.

Table 1-1: Administrator Types

Type	Permissions
System Administrators This administrator cannot do any security procedures in any domain.	Configure High Availability (HA) Upgrade DSM software Backup and restore DSM Add and delete all administrators Reset passwords for all administrators Add and delete all domains Assign one Domain Administrator to each domain Configure syslog server for system-level messages Install license file Configure DSM preferences View logs
Domain Administrators This administrator cannot remove domains and cannot do any of the domain security roles.	Add and remove administrators (Domain, Security, All) to and from domains Backup and restore DSM Configure Security Administrator roles (Audit, Key, Policy, Host, Challenge & Response, Client Identity) Configure Syslog server for application-level messages View DSM preferences View logs

Table 1-1: Administrator Types (continued)

Type	Permissions
Security Administrators Do the data protection work specified by their roles. Different roles allow them to create policies, configure hosts, audit data usage patterns, apply GuardPoints, and do other duties.	Configure signature sets Configure keys and key groups Configure online and offline policies Configure hosts and host groups Assign host passwords (manually or generated) Apply GuardPoints Share a host with another domain Export the DSM public key Import symmetric keys View DSM preferences View logs
Domain and Security Administrators	Domain Administrator and Security Administrators capabilities combined. These Administrators are deleted from the DSM database upon switching from relaxed to strict domain mode.
All	System, Domain, and Security Administrators capabilities combined. All Administrators are deleted from the DSM database upon switching from relaxed to strict domain mode.

> **Note**
> The person who performs the initial setup and configuration using the DSM CLI can also be thought of as another type of administrator. They are system users or data center system administrators with UNIX login accounts. Although they access the DSM through the CLI, for strict security practices, they should not have access to the Management Console. Conversely, the administrators listed above can access the Management Console, but should not have access to the CLI.

Read-Only Administrators

Administrators can also be created as 'read-only'. A System Administrator can create other administrator as read-only users—except for Domain administrators that are restricted to a domain. The first administrator of a domain must have privileges to create and administer other users within that domain, therefore a restricted Domain administrator cannot be created as read-only by a System or All administrator.

A read-only user inherits all of the privileges of the type of administrator and the associated roles being created however, they can only view all of the information available to that user. A read-only administrator does not have the ability to modify any settings. Read-only administrators can only change their passwords.

CLI Administrators

CLI administrators perform tasks related to setting up and operating the DSM installation—they do not administer the DSM from the browser-based Management Console. CLI administrators are system users with login accounts. Meaning, they are entered in `/etc/passwd` and they have directories under `/home`. The password requirements for both CLI and Management Console administrators are set by the password policy in the Management Console.

Chapter 2: Multitenancy

Multitenancy enables the creation of multiple **restricted** or **local domains** within a single DSM. A restricted or local domain is a DSM domain in which DSM administration is restricted to Domain Administrators or Security Administrators assigned to that domain. Multitenancy is particularly useful for Cloud Service Providers.

Overview

With multitenancy, the DSM platform supports the creation of restricted domains. Restricted or local domains are different from global domains in that Domain Administrators not assigned to that local domain, cannot modify or administer that domain in any way. Unlike global domains, local domain administrator accounts cannot be assigned to any other domains. DSM administration tasks are restricted to local Domain Administrators or local Security Administrators within that domain.

The DSM System administrator creates the first Domain administrator for a restricted domain, all subsequent administrators are created by the Domain administrator of that restricted domain. All other administrative tasks within a restricted domain are done by the local Domain administrator of that domain.

Table 2-1: "Differences between global and local domains" below lists some differences between the two types of domains/administrators.

Table 2-1: Differences between global and local domains

Global Domains and Administrators	Local Domains and Administrators
Administrator names must be unique within all global domains.	Administrator names must be unique within a local domain, but can be identical if they are in different local domains.
Domain and Security Administrators can be assigned to multiple global domains	Local Domain and Security Administrators can only function within their local domain.
DSM System Administrators can:	DSM System Administrators:
Create and assign the first global Domain Administrator to a global domain. That same global administrator can be assigned to other global domains as well. After that the DSM System Administrators do no tasks within global domains.	Create the first local Domain Administrator for a restricted or local domain. After that the DSM System Administrators do no tasks within local domains.
Change the password of any global administrator	Cannot change the password of a local administrator
Delete any global administrator.	Cannot delete local administrators
Add or delete a global domain.	Cannot access log files in a local domain.
Disable all administrators in a global domain.	Can add or delete local domains.
	Can disable all administrators in a local domain.

Implementing Multitenancy

To create a local domain, the DSM System Administrator creates a single Domain Administrator for a domain. After that, complete control of the domain is maintained by that domain's Domain Administrator and any Domain or Security Administrators created by that Domain Administrator.

Administrators in a local domain do DSM duties in **exactly** the same way as in global domains. The only differences are as follows:

- They are restricted to doing DSM work only in their own local domain

- Administrators not in their local domain (including DSM System Administrators) cannot do any domain-related work.

> **Note**
> While DSM System Administrators cannot view the administrators in the local domain, DSM System Administrators can disable all administrators in a local domain.

The Domain Administrator of a local domain can also create 'read-only' administrators. A read-only user inherits all the privileges of the administrator type (and the associated roles in the case of Security administrators), being created. See "Read-Only Administrators" on page 26 for more information about Read-Only administrators.

Creating Local Domain Administrators

This section describes how to create a local domain and its local Domain Administrator.

1. Log on to the Management Console as a DSM System Administrator.

2. Create a domain:
   a. Exit the current domain if necessary.
   b. Go to **Domains > Manage Domains > Add**. Enter domain name (example: **Domain-2**) and click **Ok**.

3. Create a Domain Administrator for this domain.
   a. Go to **Administrators > All > Add**.
   b. Enter **Login** and **Password**.
   c. For **User Type** select **Domain Administrator**.
   d. **Restrict to Domain** field displays. Select the domain to restrict in the pull-down. Click **Ok**.

You have now created a local domain (Domain-2) and a local Administrator (Admin2). When you return to the Administrators window, you will not see the administrator's name listed in the table. The new administrator is in a local domain, and does not appear in the list of global administrators.

Logging in to a local domain

1. Go to the log in screen of the Management Console.

2. Enter the login and password of the local Domain or Security Administrator.

3. Check the **I am a local domain administrator** checkbox and enter the domain name.

4. Click **Ok**. The Dashboard displays the administrator and the current domain on the top right of the console.

Creating a local Security Administrator

Like a global Domain Administrator, the local Domain Administrator cannot do any of the standard security roles (Audit, Key, Policy, Host, Challenge & Response, and/or Client Identity) unless the administrator has been created as a Domain and Security Administrator. If the Domain administrator is a separate role, the local Domain Administrator must create local Security Administrators to do tasks associated with the different security roles.

> **Note**
> DSM System Administrators cannot create DSM Security Administrators for a restricted domain.

1. Go to the log in screen of the Management Console and log in as a local Domain Administrator.

2. Click **Administrators > Manage Administrators > New**.

3. In the **Add Administrator** window, enter a login and password. Select User Type as **Security Administrator**.

4. Select the Roles for this administrator account and click **Ok**.

5. A new local Security Administrator is created.

Creating a local Domain or Security Administrator as Read-Only

1. Go to the log in screen of the Management Console and log in as a local Domain Administrator.

2. Click **Administrators > Manage Administrators > New**.

3. In the **Add Administrator** window, enter a login and password.

4. Select a **User Type** from the drop down list.

5. Select the **Read-Only User** check box to create an administrator with read-only privileges. An administrator with read-only access will not be able to add, delete, or modify any settings in the domain. Read-only administrators will only be able to change their passwords and view the different settings per their type and the roles assigned to them.

Chapter 3: Licensing

System Administrators operate outside of domains. They create domains and assign Domain Administrator to the domains. Domain Administrators and Security Administrator operate within those domains. All Administrators can operate both inside and outside of domains. When an All administrator enters a domain, that administrator can do Domain Administrator and Security Administrator tasks. When an All administrator exits the domain, that administrator can perform System Administrator tasks.

The default administrator, admin, is a System Administrator. The admin administrator creates additional administrators and domains, and then assigns one Domain Administrator to each domain.

> **Note**
> A System Administrator cannot perform any security procedures in a domain.

System Administrators can create licenses. Licensing applies to VTE (FS), and VAE (Key) agents, and KMIP. You must have a valid license to configure hosts and register agents with a DSM. Once your license expires because it reaches its expiration date (term license), or you use all of your hours (hourly license), the agent under that license will continue to operate until the protected host is rebooted or unregistered. However, new agents cannot be registered without a valid license, and the existing agent will need to be re-registered with a valid license.

Overview

Licenses are provided by Customer Support and are uploaded to the DSM. The term of a license begins when you register the agent using that license. The first time you log on to a DSMDSM, the Dashboard displays "License file not found". You must upload a license file to add and register hosts with the DSM.

In an HA environment, the license file only needs to be uploaded to the initial DSM, HA node 1. Once installed on the initial DSM, it automatically replicates to the other nodes.

> **Note**
> Do NOT change the number of cores after the agent(s) has (have) been registered.

License Types

Perpetual licenses

- No expiration date

- Licensed per core

- Usage by agents not accounted for

For a perpetual license to operate on a host, set the following:

- On the Edit Host page under the **General** tab, **License Type** must be set to **Perpetual**. (**Management Console > Domains > Switch Domains > Select Domain > Hosts > Select Host > Host Name > Edit Host > License Type**)

Term licenses

- Subscription based

- Monthly, annual, or multi-year term

- Licensed per core

- Usage by agents accounted for charge back

- There is no business disruption due to expired term licenses until the host is rebooted or unregistered. Agents will not be allowed to register or re-register when the term has expired until a renewed or new license is applied.

For a term license to operate on a host, set the following:

- On the Edit Host page under the **General** tab, **License Type** must be set to **Term**. (**Management Console > Domains > Switch Domains > Select Domain > Hosts > Select Host > Host Name > Edit Host > License Type**)

Hourly licenses

- Usage based

- Purchased in blocks of core hours

- Licensed per core hour

- No term or expiration of license as long as there are unused hours remaining

- There is no business disruption due to expired hourly licenses until the host is rebooted or deregistered. Agents will not be allowed to register or re-register when the term has expired until a renewed or new license is applied.

For an hourly license to operate on a host, set the following:

- On the Edit Host page under the **General** tab, **License Type** must be set to **Hourly**. (**Management Console > Domains > Switch Domains > Select Domain > Hosts > Select Host > Host Name > Edit Host > License Type**)

Logical core hour usage

License usage metering begins whenever an agent registers with a DSM and stops when it unregisters. Each report is generated on a monthly basis starting at the beginning of the month and ending at the end of the month. The logical core hours value is determined as follows:

> `cpuCoreCount` **X** `hoursUsed`

where,

- `cpuCoreCount` refers to the number of cores divided by 16 with a minimum value of 1, on the system where the agent is installed. The DSM gets this information from the agent. To learn more about the core count on the host system, contact your Thales eSecurity sales representative.

- `hoursUsed` starts from the beginning of the month or when an agent is registered with a DSM, whichever is later, and ends at the end of the month or when an agent is deregistered from a DSM, whichever is earlier.

  If the agent deregisters during the month and re-registers within the month, the time between deregistration and reregistration is not counted.

The License Page

On the License page on the Management Console, view information about the licenses on the system, and upload license files.

To access the license page, click **System > License**. The following information appears:

Table 3-1: License Page Information

Column Name	Description
Agent Type	FS (VTE), Key (VAE), KMIP
Term License	**Agents Licensed:** Total number of term licenses per agent type **Cores Licensed**: Total number of cores licensed per license agreement **Expiration Date**: Expiration date of the license
Perpetual License	**Agents Licensed:** Total number of perpetual licenses per agent type **Cores Licensed:** Total number of cores licensed per license agreement
Hourly License **Core-Hours Licensed**	Number of hours licensed per agent type.
LDT License	Enables Live Data Transformation. See the LDT guide for more information.
Docker License	Enables the Docker container. See the VTE Agent Installation and Configuration guide for more information.
Efficient Storage License	Enables Efficient Storage. Efficient Storage encrypts data on an Efficient storage array. See the Efficient Storage chapter in the *VTE Agent Installation and Configuration Guide* for more information.

Uploading a License File

The first time you log on to a DSM, the dashboard displays "`License file not found`". You must upload a DSM license file to remove the message and so you can add and register hosts with the DSM.

> **Note**
> In an HA environment, you only need to upload the license file to one HA node. Once installed on an HA node, the license automatically replicates to the other nodes.

To upload a license file:

1. Log on to a DSM HA node as a System/All administrator. Other administrator types can view the license in the Management Console, but they cannot upload a license file.

2. Get a license file from Customer Support.

3. Select **System > License** in the menu bar. The License window opens.

4. Click **Upload License File**. The Upload License File window opens.

   > **Note**
   > If you are in a domain, the **Upload License File** button is disabled. Click **Domain > Exit Domain**.

5. In the **License File** box, enter the full path of the license file or click **Browse** to locate and select the license file.

6. Click **Ok**.

Allocating Licenses and Hours to a Domain

You can control how many licenses (perpetual or term) or license hours (hourly) can be used in a domain under the License tab in the Edit Domain window.

1. Click **Domains > Manage Domains.** The Domains window lists all the domains available to the current administrator.

2. Click the domain link in the **Name** column. The Edit Domain window opens to the **General** tab.

3. Click the **License** tab

The fields under the **License** tab operate as follows:

- Leave a field blank — agents can be registered in the domain according to the number of licenses available on the system.

- Enter a zero — no agents can be registered in this domain.

- Enter a number in an Agent (Term) or Agent (Perpetual) field — the domain is restricted to that number of hosts registered with that type of license.

- Enter a date in the Expiration Date (Term) field — new hosts cannot register after the expiration date. Active hosts continue to function until they are unregistered or rebooted.

- Enter a number in the Core Hours (Hourly) field — the domain is restricted to that number of core CPU hours with that agent. Active hosts continue to function until they are unregistered or rebooted.

Chapter 4: Creating, Adding, and Deleting DSM Administrators

A default System Administrator, called `admin,` already exists on the DSM. The first time you log on, you do so using the default administrator credentials. Additional administrators must be created to do tasks that a System administrator cannot perform.

Creating DSM Administrators

Thales recommends that you create backup administrators for each administrator type, as a precaution. This way, if a particular administrator is compromised, that administrator can be deleted and their administrative tasks can be assumed by a different administrator.

1. Log on to the Management Console as a System/All Administrator.

   If this is the first time you are logging in, you must log in with the credentials of the default administrator admin; with the default password `admin123`. You will be redirected to the reset password page. You must reset the password. This is true for any administrator logging in for the first time. Note that you cannot delete the default administrator.

2. After resetting your password, the Management Console Dashboard displays. Click **Administrators** on the main menu bar.

   The **Administrators** window opens listing all of the administrators for this DSM.

3. Click **Add**. The Add Administrator window displays.

4. In the Add Administrator window, enter the following information:

   ○ **Login**

   Type a name. Only one instance of an administrator name is allowed.

   ○ **Description** (Optional)

   Enter a phrase or string that helps you to identify the administrator. The maximum number of characters is 256.

   ○ **RSA User ID** (Optional)

   Required:

   ■ RSA Authentication Manager software application deployment

   ■ RSA SecurID device

   The RSA SecurID device and RSA user name are bound together in the RSA Authentication Manager software application by a security administrator. Enter the RSA user name that was configured by the security administrator in the RSA user name text-entry box. The value entered is displayed in the RSA User Name column of the Administrator window.

   ○ **Password**

   Enter a password. The password must conform to the attributes defined in the password preferences. The maximum password length is 256 characters.

If you have enabled and configured multi-factor authentication, an administrator may have two passwords to log on to the Management Console: one for a DSM administrator and one for an RSA user.

The DSM administrator password is used to log on to the Management Console, if multi-factor authentication is disabled or not configured. When multi-factor authentication is enabled and the administrator is configured for multi-factor authentication, the DSM administrator logs into the Management Console with the RSA SecurID password and the Token Code displayed on the RSA SecurID device.

> **Note**
> The first time an administrator logs on to the Management Console with a newly created DSM Administrator account, they are prompted to change the password. Administrators cannot reuse the same password to create the account.

- **Confirm Password**

    Retype the password.

- **User Type**

    Select a type of Administrator from the drop-down menu:
    - System Administrator
    - Domain Administrator
    - Security Administrator
    - Domain and Security Administrator
    - All

> **Note**
> If you enabled Separation of Duties in the **System > General Preferences > System** tab, then the **Domain and Security Administrator** and **All** options are not available.

- **Read-Only User**

    Select this option to create an administrator with read-only privileges. You can assign read-only privileges to any type of administrator—except for Local Domain administrators that are the first administrators to be assigned to a domain. If the first administrator added to a local domain is read-only, that administrator will not be able to create any more administrators for that domain.
    An administrator with read-only access will not be able to add, delete, or modify any settings on the DSM. Read-only administrators will only be able to change their passwords and view the different settings per their type and the roles assigned to them.

- **Smart Card Login Enabled**

    Requires user to use a PIV (Personal Identity Verification) smart card for multifactor authentication. See Chapter 5: "Smart Card Access" on page 39 for more information.

5. Click **Ok**. A new Administrator is created. The Administrators page displays a table with the name and type of the new administrator.

Importing Administrators

The **Import** function imports data from an LDAP server such as Active Directory (AD). Once an LDAP server has been identified and configured, the administrator can import the desired values. See Chapter 9: "LDAP Configuration" on page 66 for more about configuring an LDAP server.

You need the LDAP login ID and password to import values from an LDAP directory.

1. Select **Administrators > All**. Click **Import**.

2. Enter the Login ID and Password on the **Connect to AD/LDAP Server Details** page. If the login and password were entered under **LDAP Server Settings** on the **System > LDAP** page, these values will be populated by those values and do not need to be re-entered. You may also enter a different login and password in place of these stored values when you import administrators.

3. Click **Connect**. The LDAP Users window displays LDAP user names.

4. Search options:

    a. Use the LDAP Query field to filter searches using the LDAP query language. Results depend on how the LDAP service is set up. See RFC2307 for full details on syntax.

    b. Select a Group from the **Group** drop down list.

    c. Enter a User name in the **User** field.

    d. The **Maximum number of entries to return** field lets you limit the maximum number of records to import or display. The default value is 300. The minimum value is 1 and the maximum value is 1 million. A high integer value may result in a delay depending on the database size.

5. Click **Go**.

Selecting LDAP administrators

The Management Console provides a GUI interface to the mapped LDAP directory values such as login and user description. As an Administrator logged into the LDAP directory, you can provide input to the following fields in order to select and manage LDAP users. See Chapter 9: "LDAP Configuration" on page 66 for more information about adding LDAP users.

- **Group Object Class**

    Select a value from the drop down menu to filter by group type.

- **User Object Class**

    Enter a value or partial value to filter on specific users. Entering a partial value acts as a "wild-card" returning all values matching what was entered.

- **Go**

    Click to refresh the screen.

- **Select All**

    Click to select all values on this page.

- **View**

    Select a value from this dropdown box to control how many values appear on any page.

- **Selected**

    Click to select individual values.

- **User Type**

    Select a value from this dropdown box to define the type of Administrator or role of the values you import.

- **Add/Cancel**

    Select to add or cancel your selections.

Deleting DSM Administrators

System or All administrators can delete other administrators, except for the default **admin** administrator and themselves.

If the administrators to be deleted are members of a domain, they must first be removed from that domain (even if the domain has been deleted), before they can be deleted.

To remove an administrator from a domain

1. Log in as a Domain Administrator, Domain and Security Administrator, or All.

2. Remove the administrator you want to delete from every domain from which they are a member.

To delete an administrator

1. Log in as an All/System Administrator.

2. Select **Administrators > All**.

3. In the Administrators window, enable the **Selected** check box of the administrator(s) to be deleted.

4. Click **Delete**.

5. You are prompted to verify that you want to proceed with this operation.

6. Click **Ok**. The selected administrators are deleted from the Management Console and cannot access the DSM.

Resetting DSM Administrator Passwords

Administrator passwords cannot be viewed. If an administrator forgets their password, the System Administrator can assign a new temporary password. The Administrator informs the administrator about their new temporary password. The next time the administrator logs on, they are directed to enter a new password.

If a DSM administrator is currently running an active Management Console session when the System Administrator changes their password, the Management Console session is immediately terminated and the administrator must log on again.

When a System Administrator changes the password for a Domain Administrator, Security Administrator, or All, their account is disabled in every domain of which they are a member. They must be enabled by a different Domain Administrator, Domain and Security Administrator, or All before they can again enter a domain. A disabled administrator can log on to the DSM, but the domain selection radio buttons are opaque and cannot be selected, so the administrator cannot enter any domain and cannot modify the DSM configuration.

The Domain Administrator, Security Administrator, or All account must be enabled in every domain of which they are a member at the time the password is changed. Enabling an administrator in one domain does not enable them for all the domains of which they are a member.

To change another administrator's password:

1. Log in as a **System Administrator** or **All**.

2. Check that the administrator is not currently logged into the Management Console because their login session becomes inactive when the password changes.
   - If you are changing the password of another System Administrator, you can check the log.
   - If you are changing the password of a Domain Administrator or Security Administrator, have a Domain Administrator switch to each domain in which the administrator is a member and check the log of each domain.

3. Select **Administrators > All**. The Administrators window opens.

4. Select an administrator in the **Login** column. The Edit Administrator window opens.

5. Enter the password and then click **Ok**.

6. For Domain Administrator, Security Administrator, or All, have a different Domain Administrator, Domain and Security Administrator, or All re-enable that administrator's domains.

Chapter 5: Smart Card Access

Multifactor authentication adds additional access control to the DSM Management Console by requiring DSM administrators to use a smart card each time the administrator logs into the Management Console.

Prerequisites

Multifactor authentication with a smart card requires the following to be set up prior to use with the DSM:

- The smart card administrator must provision the smart card according to the smart card instructions.

- Install a mini-port driver for the smart card (mini-port driver is provided by either the smart card vendor or Microsoft).

- Program the smart card with a PIN number.

> **Note**
> This feature was tested with a Gemalto IDPrime MD smart card.

Port Configuration

Open the following ports for smart card communication:

- RSA with smart card login: 8449

- EC with smart card login: 8450

- Compatible with smart card login: 8450

Refer to the **Ports to Configure** table in 'Appendix D' of the *DSM Installation and Configuration* Guide.

Setting up and Configuring the Smart Card

Once the smart card is provisioned and the administrator has the certificate chain, an administrator can setup the smart card on the DSM.

> **Note**
> Currently, Smart Cards are compatible with Internet Explorer v11 and Microsoft Edge v10.

Enabling Smart card usage in the DSM

1. Click **System > General Preferences > System.**

2. Go to Smart Card Login and select **Smart Card Login Enabled**.

3. Click **Apply**. Smart Card CA certificates now displays in the System menu.

Import Smart Card CA Certificate

An administrator must import a CA certificate into the DSM.

1. Click **System > Smart Card CA certificates**.

2. Click **Browse** and select the **CA certificate** file.

3. Click **Import Certificate.**

   > **Note**
   > If you need to import more than one CA certificate, you must import the root CA certificate first, followed by the intermediate CA certificate that is signed by the root CA.

4. Click **OK** when the warning dialog opens.

5. Restart all nodes in the DSM cluster through the CLI (0001:system$ `server restart`) after all of the CA certificates have been imported and all of the changes have been replicated successfully to all the other nodes in the cluster.

Export Smart Card CA Certificate

1. Click **System > Smart Card CA certificates**.

2. Select the certificate to export.

3. Click **Export CA Certificate**.

4. In the browser question bar, select **Save As** to export the file to your system.

Delete Smart Card Certificate

To delete a CA certificate:

1. Click **System > Smart Card CA certificates**.

2. Select the certificate to export.

3. Click **Delete**.

4. Restart the browser to refresh the new certificate list.

Enabling LDAP Users to use a Smart Card

You can import Administrators from an LDAP server (System > LDAP) and then enable them for smart card access like any other DSM Administrator. To be successful, you must do the following:

- In the AD/LDAP page, you must enter a Login name and Password.

- Do not delete the contents of the Login name and Password field.

- **Do not** click **Clear** to clear any LDAP server settings.

- In the DSM, in the (Administrators > Import) Connect to AD/LDAP Server page, the Login /Password fields are automatically populated from the LDAP Server settings page. Do not change the contents of those fields.

  > **Note**
  > If you clear any of the settings, the user will not be able to login because the DSM will not be able to bind the user with the LDAP server.

Enabling DSM Administrators to use a Smart Card

A DSM Administrator must enable the smart card feature for other administrators and select the user mapping criteria for the certificate.

To enable an Administrator to use a smart card:

1. In the Edit/Add Administrator window, select **Smart Card Login Enabled.**

   The smart card user mapping criteria list opens. These criteria are used to map a certificate to this user. You must select at least one.

2. Select appropriate criteria for the smart card administrator:

   Table 5-1: Smart Card Options

Name	Description
Subject CN	The Common Name part of the Subject DN
Subject DN	Subject Distinguished Name of the smart card user certificate
Issuer DN	Issuer Distinguished Name of the smart card user certificate
Serial Number	Serial Number of the smart card user certificate
Subject Key Identifier	Subject Key Identifier of the smart card user certificate
Public Key Hash	The SHA-384 hash of the public key of the smart card user certificate
User Principal Name	Used by Microsoft to identify a user in a Windows domain. It is stored in the otherName field of the Subject Alternative Name extension using OID 1.3.6.1.4.1.311.20.2.3
RFC822 Name	The rfc822 Name field of the Subject Alternative Name extension, this is usually the email address of the user

3. Click **OK**.

Initial Login when using a Smart Card User

When an administrator logs in for the first time with smart card login privileges enabled, they will log in as follows:

1. In the user access dialog, enter your admin name and password.

2. Click **OK**. The Upload client certificate dialog opens.
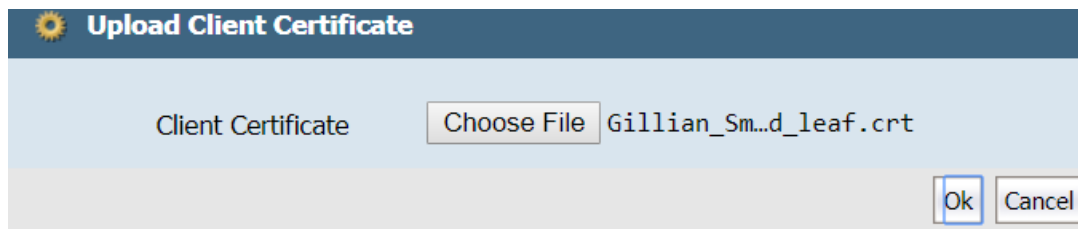
3. Click **Choose File** and select the **user** certificate from the certificate chain

   This is the certificate that was imported into the smart card during the smart card provisioning.

4. Click **OK** to load the **user** certificate for the administrator.

   ○ The certificate file is located on the file system where the web browser runs. It may contain multiple certificates.

5. After you successfully upload the certificate, plug in the smart card reader, insert the smart card into the reader, and restart the web browser.

6. Enter your login name and password in the DSM Login dialog. Click **Login**.

7. The confirm certificate dialog opens. Select the certificate that was uploaded to the DSM in a previous step and press **OK**.

8. Enter the smart card PIN number.
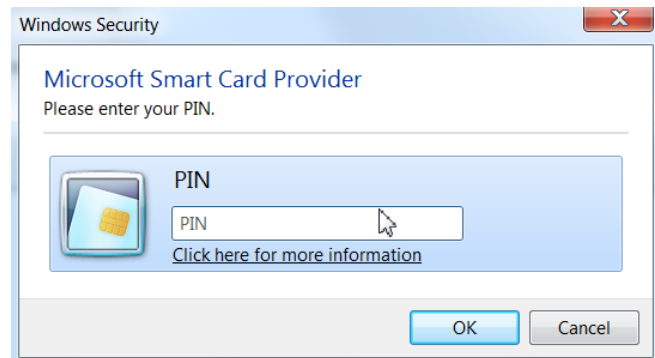
9. Confirm the certificate again, if required.

Logging on after Initial Logon

After the administrator is set up on the DSM and with the smart card, logging in becomes much simpler. When you start the browser with the smart card reader already plugged in and containing a smart card, the Windows Security Confirm Certificate dialog opens.

1. Click **Login** on the Smart Card Login dialog.

2. Select an appropriate certificate and click **OK**.

3. Enter the Smart card PIN in the Smart Card Provider dialog and click **OK**.

Note
If your system is idle for more than a minute, the browser prompts you for the PIN.

Chapter 6: Configuring DSM Preferences and Viewing Logs

As a DSM System Administrator (or type All), you can set many preferences in the Management Console.

Overview

In the **System > General Preferences** page, you can set the following preferences:

- **Display**

  The number of DSM objects displayed based on the object type. For example, you can set a preference that displays all configured policies on one Web page, rather than just 20 per page.

- **System**

  Enable Syslog messaging, enable super administrators, and shorten the update interval when pushing changes to the same policy to hosts on different servers.

- **Password**

  How long a password must be, the types of characters that a password must contain, and password duration. Password preferences can also configure the DSM response to repeated failed login attempts.

In the **System > Log Preferences** page, you can set the following:

- **Server preferences**

  Log maintenance parameters on the DSM. For example, you can set the interval to wait before moving agent log entries from temporary buffers on the DSM to the DSM log database, and consequently to the log viewer.

- **VTE Agent log**

  How the DSM maintains agent-specific log data. For example, you can set the interval at which the agent uploads log data to the DSM.

Configuring DSM Preferences

The General Preferences tab lets you specify display settings, system settings, password settings, and lets you configure the login banner message on the log in screen.

Setting Display preferences

Display preferences are administrator-configurable parameters that control the number of objects to display and set the Management Console expiration time.

To set DSM display preferences:

1. Log on to the Management Console as a System/All administrator.

2. Select **System > General Preferences** in the menu bar.

   The General Preference window opens to the Display tab.

3. Change the values displayed in the attribute text-entry boxes or scroll-list.

   The following table lists and describes attributes and their values.

Table 6-1: General Preferences Display tab attributes and their values

Category	Parameter	Description
Domain Page Settings	Number of Domains Per Page	Sets the maximum number of administrators in the Domains window to display on one page. Navigation buttons are displayed in the Domains window to move between the pages. The default is 20.
Administrator Page Settings	Number of Administrators Per Page	Sets the maximum number of administrators in the Administrators window to display on one page. Navigation buttons are displayed in the Administrators window to move between the pages. The default is 20.
Host Page Settings	Number of Hosts Per Page	Sets the maximum number of hosts in the Hosts window to display on one page. Navigation buttons are displayed in the Hosts window to move between the pages. The default is 20.
	Number of Host Groups Per Page	Sets the maximum number of host groups in the Host Groups window to display on one page. Navigation buttons are displayed in the Host Groups window to move between the pages. The default is 20.
Policy Page Settings	Number of Policies Per Page	Sets the maximum number of policies in the Policies window to display on one page. Navigation buttons are displayed in the Policies window to move between the pages. The default is 20.
Key/Certificate Page Settings	Number of Keys Per Page	Sets the maximum number of keys in the Keys window to display on one page. Navigation buttons are displayed in the Keys window to move between the pages. The default is 20.
	Number of Key Groups Per Page	Sets the maximum number of key groups in the Key Groups window to display on one page. Navigation buttons are displayed in the Key Groups window to move between the pages. The default is 200.
Signature Page Settings	Number of Signature Sets Per Page	Sets the maximum number of signature sets to display on one page. Navigation buttons are displayed in the Signature Sets window to move between the pages. The default is 20.
Log Page Settings	Number of Log Messages Per Page	Sets the maximum number of log entries to display on one page. Navigation buttons are displayed in the Logs window to move between the pages. The default is 20.
Management Console Timeout	Management Console Timeout	Sets the interval of inactivity allowed before automatically logging administrators out of the Management Console Web session. Unsaved changes are discarded. Choices are 5 minutes, 20 minutes, 1 hour, 2 hours, and 8 hours. The default is 1 hour.

4. Click **Apply** to set the changes.

Setting System preferences

You can configure attributes that:

- Enable or disable all Administrator accounts for Domain and Security administrators. By enabling the **Separation of Duties** option, all Domain and Security Administrator accounts are deleted from the database, and only System, Domain Administrator, and Security Administrators remain.

- Speed up DSM updates when policy changes are pushed to VTE Agents that are administered by HA nodes. By enabling **Without Replication Confirmation**, the HA node no longer waits for other HA nodes to synchronize before it begins pushing changes to its own agent hosts.

- Enable syslog logging. Once enabled and configured, a Syslog server can transmit/receive logging data.

To set system preferences:

1. Log on to the Management Console as a System/All administrator.

2. Select **System > General Preferences** in the menu bar.

   The General Preference window opens to the Display tab.

3. Click the **System** tab.

4. Change the values displayed in the attribute check boxes.

   The following table lists and describes the attributes and their values:

Table 6-2: General Preferences System tab attribute values and use

Category	Parameter	Description
Organization	Name	Enter the name of the organization (company, department, or function) responsible for or managed by this DSM. This is useful for reporting and auditing purposes.
Separation of Duties	Enforce separation of duties	Check box to operate in relaxed domain mode or strict domain mode. When enabled, strict domain mode is applied. Administrators are assigned a single administrative type that can do a specific set of tasks. This means that at least three administrators must be configured, each with a specific type, in order to do all DSM tasks. When disabled, the domain mode rules are relaxed, and two additional compound administrative types (Domain and Security, and All) can be configured. When switching from strict to relaxed domain mode, all currently configured administrators are left intact. When switching from relaxed to strict domain mode, all of the primary administrator types are left intact and all of the compound Domain and Security Administrator and All administrators are deleted immediately. The checkbox is disabled by default, indicating relaxed domain mode.
Push Host Configuration	Without replication confirmation	Enable this check box if you want the HA nodes to immediately update the locally administered hosts that are affected by a policy change, even if the same policy is also used for GuardPoints on remotely administered hosts. Disable the check box if you want the HA node to delay pushing policy changes to locally administered hosts until after it successfully synchronizes with the other HA nodes that apply the same policy. If the checkbox is disabled, the HA node can wait up to 15 minutes for all of the other HA nodes to synchronize before it pushes the policy changes to locally administered hosts. The checkbox is enabled by default.

Table 6-2: General Preferences System tab attribute values and use (continued)

Category	Parameter	Description
Agent Keys	Key refreshing period (in minutes)	Defines the refresh period for Agent keys stored on the host. The refresh period value ranges from 1 to 44640 minutes (31 days). The default value is 10080 minutes (7 days). When set outside of a domain under General Preferences, the refresh period is applied globally, for all new keys. The refresh period is not reset for existing keys.
Key Template	Enforce Using Key Template to Define Key	When enabled, administrators creating keys must select a key template to define the key attributes.
Policy	Maximum Number of Policy History Show Validation Warnings	Sets the maximum number of policy history versions stored in the database. The default value is 10. User selectable values are 0, 5, 10, 50, 100. Changing this value does not delete any older versions until the next time a policy is changed and saved. When saved, the XML data of the older version is deleted and cannot be recovered (unless restored from a prior backup). Policy metadata such as who and when is not deleted. If you want to see validation warnings, enable the **Show Validation Warnings** check box. This is disabled by default.
Syslog Setting	Syslog Enabled	When enabled, properly configured syslog servers can receive logging data. Domain, Domain and Security, and All administrators can configure syslog servers. Syslog messaging is domain-specific. Only the events that occur in the local domain are sent to the syslog server. If the administrator is not in a domain when configuring a syslog server, local DSM and appliance system messages are sent to the syslog server. This checkbox can be enabled and disabled by System/All administrators. The checkbox is disabled by default.
Automatic Backup Settings	Automatic Backup Enabled	When enabled, allows System/All administrators, or from within a domain, Domain or Domain and Security administrators, to schedule automatic backups of the DSM or DSM domain configuration. This setting must be disabled to comply with Common Criteria standards. This setting is enabled by default.
Connection Timeout	Max Agent Connection Timeout	Distance and unreliable networks can cause configuration pushes and pulls between DSMs and hosts to timeout. If, because of a slow connection, policy updates are not being pushed to a host, or a host is unable to pull the latest configuration changes, increase the timeout interval. Preference changes are not automatically pushed to hosts. To push a new timeout value to a host, change something in the host configuration, such as Host Settings, and the DSM will push the change, including the new timeout interval, to the host. You can also pull the new timeout onto the host. To pull the change onto a host, log onto the host, either via SSH or a Remote Desktop Connection, and kill the `vmd` process. Wait a moment and the `vmd` process will automatically restart. As `vmd` restarts, it queries the DSM for updates, including policy changes and the connection timeout value. The allowed range is `1` to `600` seconds. The default is `20` seconds.

Table 6-2: General Preferences System tab attribute values and use (continued)

Category	Parameter	Description
Multi-Factor Authentication	Multi-factor authentication Configured	If selected, indicates that multi-factor authentication has been enabled on the HA cluster. This checkbox is a display indicator only. Multi-factor authentication is enabled via the CLI, not the Management Console.
Smart Card Login	Enabled Smart Card login	System enabled for use with a smart card for user access.
Backup Requests Management	Number of Processes to Handle Backup Requests	The DSM may contain several local domains that domain administrators want to backup. If multiple backup requests are made at the same time, this could cause the DSM to hang. Therefore the number of processes set aside to handle backup request is set to 10 by default. This means, if there are more than 10 backup requests, they will remain in the queue until a process is free to perform a backup. Select the number of processes from the drop down list.
HDFS Browse Connections	HDFS connection Time Out (not less than 15 seconds)	Set a time in seconds for connection timeout when browsing HDFS directories from the DSM. When the timeout limit is reached, the DSM aborts the attempt and tries to reconnect.
Active Key Deletion	Allow deletion of active key	Select this check to allow deletion of VAE keys that are in the 'ACTIVE' state without requiring those keys to be deactivated first.
Strict Host Group Deletion	Host groups with GuardPoints	Select this option to prevent Host Groups from being deleted if they contain hosts or GuardPoints. If not selected, then the Host Group will be deleted even if it contains hosts or GuardPoints.
Disable Administrator	Disable domain administrators	Prevent System/All administrators from disabling all of the administrators in a domain by hiding the Disable Administrator button. This setting will be reflected in all newly created domains. It does not affect existing domains.
LDAP Timeout	LDAP Connection Timeout	Timeout duration is now configurable. Set as needed. [minimum time: 1 s | maximum time: 600 s | default time: 30 s]
Key Governance	Key Governance	Enabling this feature prevents a VAE developer from changing a key. Specifically, the developer cannot: create a key; add a key to a key group; add an Attribute; modify, delete or destroy an attribute; change a key state; delete a key. **Note** that this feature does not affect the GUI or REST API. Security administrators can still create, delete and modify keys in the GUI or REST API.

5. Click **Apply** to set the changes.

Setting DSM password preferences

Administrator passwords are a vital part of a good security system. A Management Console administrator password can contain standard ASCII alphabet characters `(a-z, A-Z)`, integers `(0-9)`, and a limited set of special characters (`!@#$%^&*(){}[]`). The individual elements in this combination of characters cannot occur in sequential order. That is, a password cannot contain two instances of the same element if they are next to each other. For example, `mississippi` will not be accepted, but `misSisSipPi` will.

Additional restraints can be applied that require all new passwords to contain at least one uppercase alphabet character, at least one special character, and the minimum number of characters that must be used.

Password preferences are applied to both administrator passwords and host system passwords.

To set DSM password preferences:

1. Log on to the Management Console as an System administrator with `Host` role permissions or type All.

2. Select **System > General Preferences**.

   The General Preferences window opens.

3. Select the **Password** tab.

4. Change the values displayed in the attribute text-entry boxes or scroll-list.

The following is a list of attributes you can configure, and their values:

Password Characteristics

- **Password Duration**

  Passwords expire after the number of days set by an administrator. The password expiration interval is applied globally to each administrator account. If the administrator does not change the password prior to the expiration, the administrator must reset the password immediately the next time the administrator logs in. The expiration interval is an integer between 6 and 365. The default is 90. 'Password Duration' must be set to a value greater than 'Password Expiration Notification'.

- **Password History**

  The DSM maintains a password history. You cannot use the same password more than once per the set limit. The default is 4, and the maximum value that can be set is 12. You can set this value to '0' to permit reuse of the current password.

- **Minimum Password Length**

  Sets the minimum number of characters, including blank spaces that must be in a password. The minimum password length is an integer between 8 and the limit of the operating system. The default is 8.

- **Minimum Number of Character Changes**

  Sets the minimum number of characters, including blank spaces, that constitute a password change.

- **Disallow Password Change Within (Days)**

  Sets the number of days you must wait before you can change the password again.

- **Password Expiration Notification (Days)**

  Sets the number of days prior to the password expiration at which to begin telling the administrator that their password is about to expire. Administrators are notified of the impending expiration at Management Console Login. The notification interval is an integer between 6 and 31. The default is 6.

Password Complexity

- **Require Uppercase**

  When enabled, requires at least one uppercase alphabet character in the administrator password. This is enabled by default.

- **Require Numbers**

  When enabled, requires at least one integer in the administrator password. This is enabled by default.

- **Require Special Characters**

  When enabled, requires at least one special character. (i.e., !@#$%^&*(){}[])

- **Ignore Login Username Case**

  When enabled, user names are case insensitive. Hence, lower and upper case letters are treated the same. Disable this control to make login names case sensitive.

Account Lockout

- **Maximum Number of Login Tries**

  Sets the maximum number of unsuccessful login attempts before disabling access for a set interval of time. The Management Console becomes inoperable and ignores further login attempts by an administrator for the specified interval. The range is between 1 and 10 and the default number of tries allowed is 3.

- **User Lockout Time**

  The interval to wait before re-enabling the Management Console Web interface and allowing administrators to login. The default is 30 minutes.

  The Account Lockout settings also apply to the registration shared secret, that is, if you set the maximum number of unsuccessful login attempts to 4 and the lockout time to 1 hour, then you have 4 attempts to use the registration password before you are locked out for an hour. You can attempt to register an agent again with the correct registration secret after the hour has elapsed.

Log Preferences

The entries displayed in the Message Log depend on the administrator type (System, Domain, Security, All), the domain in which the administrator is working, and, for Security Administrators, the administrator role (Audit, Key, Policy, Host, Challenge & Response, Client Identity).

An System Administrator cannot view the log entries of a Domain or Security Administrator (and vice versa) can view. By design, entries exported to a Syslog log file will have gaps in the number sequence depending on the domains and roles of the DSM administrators who are actively logged on.

Log entries are displayed in the Management Console based on the current administrator type and the domain in which the administrator is working. However, all this log information combined is available in the server.log file on the DSM.

As a System Administrator, you will see log entries such as the administrators that have logged into the Management Console, the administrators created, and policy evaluation.

Additionally, you can view log files from the DSM CLI. For details, see "diag" on page 369.

Setting Log Preferences

The Log Preferences page lets you set logging preferences for the DSM and the encryption agents. Navigate to **System > Log Preferences** on the main menu bar to access the page.

Server Log Preferences

The Server tab displays information about the current DSM logging and communication configuration. You can configure the following attributes:

Logging Settings

- **Logging Level**

  Sets the severity level at which entries are sent to cgss.log. This information is displayed in the 'Logs' window. The choices are DEBUG, INFO, WARN, ERROR, and FATAL. Each level includes the levels below it. For example, FATAL logs only fatal errors, whereas WARN logs warnings, ERROR and FATAL conditions. The default is INFO.

- **Log Upload DB Retry (secs)**

  The interval before resuming the transfer of agent log data that had been uploaded, and is stored in system files, into the log viewer database after a failure, such as after losing the connection to the database. The default is 30 seconds.

- **Log Buffer Size (messages)**

  The maximum number of entries to place in the DSM log. When this limit is reached, or when Log Buffer Flush Time has elapsed, the entries are moved to the log viewer database. The default is 100 entries.

- **Log Buffer Flush Time (secs)**

  The interval to wait before moving log entries in the server log buffers to the log viewer database. The default is 15 seconds.

- **Audit Log File Queue Size (files)**

  The maximum number of audit log files queued for processing by the DSM. This is the number of files that can be queued while the DSM processes files to move them from temporary buffers on the DSM to the DSM log database or remote Syslog servers, or to email depending on the settings. If the queued log files exceed this number, they will be rejected until the DSM can process the ones in the queue. The default is 100. Use this setting with caution as you do not want this number to become so large that it slows the DSM performance.

- **Maximum Stored Audit Log Entries (rows)**

  Previously, combined agent and server logs were purged when they exceeded 10,000 rows. Now, they are purged separately, and you can separately set the maximum number of server and agent log entries to store. The DSM will purge the logs when it reaches the maximum number of rows set in the Server logs field and in the corresponding Agent logs field.

  The minimum number of log entries is 5000; maximum number is 100,000. Note that the maximum limit may be restricted by storage capacity.

- **Communication Settings**

  ○ **Update Host Frequency**

    The interval between scans of the queue to see if any changes have been made to the host configuration on the DSM. Any changes are pushed to the host. The default interval between scans is 30 seconds.

  ○ **Default Host Communication Port**

    The port number on the DSM and on the file agent through which they communicate. When you change this port number, it is applied to all new hosts that are added after the configuration change is made. Existing file agent hosts are unaffected. The change is visible in the Communication Port field in the General tab of each new host. If you change the Communication Port number for an existing host, you must restart the file agent process that runs on that host.

Agent Log Preferences

Depending on the type of agent licenses that you have installed on your DSM, you will see an FS Agent Log tab and a Key Agent Log tab. You can configure logging preferences for the VTE (FS) and Key Agents from the respective tabs.

You can configure the file agent process information that is entered in the Management Console log. You can configure the process information globally, in which all the file system processes running on hosts systems are added after the configuration change inherit the log attributes, but all current file system configurations remain intact. Or, you can configure log attributes for individual file system installations. This section describes global file agent log configuration:

1. Message Type

   ○ **Management Service**

Logs messages that are related to the agent and VMD process server interaction in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is INFO.

- **Policy Evaluation**

  Logs messages that are related to policy evaluation in the agent log. Set the log message level to desired setting. The default log message level is ERROR.

- **System Administration**

  Logs messages that are related to system level events. The default log message level is ERROR.

- **Security Administration**

  Logs messages that are related to security related events. The default log message level is INFO.

2. Message Destination

   Log Messages can be stored in several locations.

- **Log to File**

  Send log messages to the /var/log/vormetric/vorvmd_root.log file of a UNIX host, or a Windows equivalent, such as \Documents and Settings\All Users.WINDOWS\Application\ Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log.

- **Log to Syslog**

  Send log messages to the syslog server for a UNIX host. If a syslog server is not configured, it is sent to the host 'messages' file, such as /var/adm/messages. On a Windows host, the messages are sent to the Event Viewer (Application events).

- **Upload to Server**

  Upload to the DSM and display in the Management Console Logs window.

  Level: Sets the level of error messages to be sent.

  Duplicates: Allow or suppress duplicate messages:

- **Allow**

  All duplicate messages of the corresponding Message Type are captured and displayed in the log.

- **Suppress**

  Messages of the corresponding Message Type will follow the configured Threshold as to how many times duplicate messages are sent to the DSM during the given Interval.

3. File Logging Settings

- **Maximum File Size (bytes)**

  The agent starts a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.

- **Delete Old Log Files**

  Select this check box to delete old FS agent logs. This check box works in conjunction with the Number of Old Log Files to Keep text-entry box. For example, Select this check box and enter 3 as the Number of Old Log Files to Keep value. After 3 logs are generated, the first log, log1, is deleted and a new log, log4, is created. If you do not Select this check box, log files will continue to accumulate in the server database and you will have to remove them manually.

- **Number of Old Log Files to Keep**

  Appears only when you select Delete Old Log Files. Specifies the maximum number of agent log files to leave in the server database. This text-entry box is only displayed when the Delete Old Log Files check box is enabled. The default is 5.Maximum Stored Audit Log Entries (rows)

Previously, combined agent and server logs were purged when they exceeded 10,000 rows. Now, they are purged separately, and you can separately set the maximum number of server and agent log entries to store. The DSM will purge the logs when it reaches the maximum number of rows set in the Server logs field and in the corresponding Agent logs field.

The minimum number of log entries is 5000; maximum number is 100,000. Note that the maximum limit may be restricted by storage capacity.

4. Syslog Settings

- ° Local: Send syslog messages to the local machine.
- ° Server (1, 2, 3, 4): Enter the hostname of the syslog server
- ° Protocol: UDP or TCP
- ° Message Format: Specifies the format of the message; Plain Message, CEF, or RFC5424.

5. Upload Logging Settings

- ° **Maximum Number of Messages to Upload At Once**

    Limits the number of messages sent to the DSM at one time. When the specified number of log entries is reached, those entries are uploaded to the DSM. The default is 1000.

- ° **Upload Messages At Least Every (seconds)**

    The maximum interval to wait before the agent is to upload messages to the DSM. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.

- ° **Upload Messages At Most Every (seconds)**

    The minimum interval to wait before the agent is to upload messages to the DSM. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.

- ° **Normal Time Out (seconds)**

    The maximum interval of time the agent is to wait for the DSM to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the DSM within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.

- ° **Shutdown Time Out (seconds)**

    The maximum interval of time the agent is to wait for the DSM to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.

- ° **Drop If Busy**

    Select to slow log message generation and drop log files during periods of extreme logging.

6. Duplicate Message Suppression Settings

- ° **Enable Concise Logging**

    When enabled, the number of audit log messages are reduced. This option is disabled by default. Instead of logging messages for each file system operation, only the following types of audit messages are logged:

    - Only one audit message for each read or write activity is logged at the start of that activity.
    - Audit messages for reading file status information and setting file attributes (and extended attributes) are not logged.
    - Audit messages for directory open, close and read attributes are not logged.

When this setting is enabled at the system level, it applies to all hosts that are added to the DSM, but will not apply to any existing hosts. Hosts added to the DSM after this setting is enabled will inherit this setting. These settings can be customized on each host and the host setting will override the system level settings. Note that this feature is not available for VTE versions prior to v6.0.

- **Threshold**

  Used when the Duplicates value is set to Suppress. Specifies the maximum number of duplicate messages the agent is to send to the DSM within the amount of time specified by the Interval parameter. The default is 5 messages. The maximum is 100.

- **Interval**

  Used when the Duplicates value is set to Suppress. Specifies the time period in which the number of duplicate messages, specified by Threshold, can be uploaded to the DSM. Once Interval is exceeded, the count specified by the Threshold parameter starts again. The default is 600 seconds (10 minutes). The maximum is 3600.

Network Diagnostics

The **System > Network Diagnostics** page provides a set of tools for diagnosing network related issues. This page can be accessed by every administrator. The available diagnostic tools are:

- **Ping**

  Checks if a system is up or available on the current subnet. It sends ICMP (Internet Control Message Protocol) echo request packets (ECHO_REQUEST) to the specified network host. The ping function sends six packets to the network host and reports the results.

- **IP_Address**

  Shows the current IP address and related information.

- **Arping**

  Sends Address Resolution Protocol (ARP) requests to a neighbor host, pings the address on the device interface with ARP packets, and reports the current number of users with that IP address.

- **Arp**

  Displays the kernel's ARP cache.

- **Traceroute**

  Utilizes the IP protocol time-to-live field to elicit an ICMP time exceeded (TIME_EXCEEDED) response from each gateway along the path to a specified host.

- **Checkport**

  Scans a port on a network-accessible system to verify that a TCP connection can be made to the system using the specified port.

- **Nslookup**

  Returns the IP address associated with a given host name or conversely, the host name associated with a given IP address by querying the DNS.

Chapter 7: Domain Management

System administrators can add and delete domains. However, they are not members of domains. A domain is a group of one or more VTE-protected hosts under the control of an assigned DSM Domain Administrator. Before a protected host can be administered, it must be placed in a domain.

Adding Domains

To add a domain:

1. If you are already logged into the Management Console, log out and log in again as the DSM System Administrator `admin`. Otherwise, just log on as `admin`.

2. Click **Domains > Manage Domains** to bring up the Manage Domains window.

   If you are in a domain click **Exit Domain** to exit the domain and then click **Manage Domains**.

3. Click **Add**. The Add Domain window opens.

4. Under the **General** tab, provide a name for the domain.

   a. **Name**: Enter a name of up to 64 characters for the new domain.

   b. **Organization**: (Optional) Enter the name of the organization responsible for or administered by this domain.

   c. **Description**: (Optional) Enter a phrase or string of up to 256 characters to help identify the domain.

   d. **Help Desk Information**: (Optional) Enter the phone number to call to get the response string for challenge-response authentication. If you leave this box empty, the default message is "`Please contact a Security Server administrator for a response.`" (Note: The term "Security Server" refers to the DSM.)

5. Click **Apply** to save the domain information.

6. Click the **Assign Admin** tab to assign an administrator. If you do not assign an administrator when you add the domain, you can edit the domain later to add an administrator. However, you cannot switch to the domain until you assign an administrator.

7. (Optional) Click the **License** tab to allocate licenses or license hours per agent on this domain.

8. Click **Ok**. The **Domains** window opens with the name and description of the new domain.

After the domain is created and has an assigned DSM Domain Administrator, hosts can be added to it.

Deleting Domains

> **Note**
> Back up security objects such as keys, policies, and logs, before you delete them. Without the keys, you cannot restore or access encrypted data. When you delete a domain, all the log data for that domain is also removed from the DSM database.

To delete a domain:

1. Log in as a Security/All administrator.

2. Switch to the domain to be deleted.

3. Delete all the policy, key, and host configurations.

4. Logout.

5. Log in as a Domain Administrator, Domain and Security Administrator, or All.

6. Switch to the domain to be deleted.

7. Delete all administrators that are assigned to that domain.

   You can delete all but one Domain Administrator; which is the administrator that you are currently logged in as.

8. Log out.

9. Log in as a System/All Administrator.

10. Select **Domains > Manage Domains**.

    The Domains windows is displayed.

11. Enable the **Selected** check boxes for the domains to be deleted.

12. Click **Delete**.

    You are prompted to verify that you want to proceed with this operation.

13. Click **Ok**.

    The deleted domain(s) will no longer appear in the domains table in the Domains window

Assigning Domains to Domain Administrators

A DSM System Administrator creates other DSM administrators but can assign only one Domain or Domain and Security Administrator to a domain. After the first administrator has been assigned to a domain, all subsequent administrators must be assigned or added (depending on the type of domain), from within the domain. The DSM Domain administrators that first are assigned to a domain can log into the domain from the Management Console and add additional Domain Administrators or Domain and Security Administrators to the domain. A global Domain Administrator can add only existing Domain Administrators, Security Administrators, and Domain and Security Administrators to the domain, listed in the global administrators table on the Administrators page to the domain. Restricted Domain administrators can create administrators within their domains and these administrators are not visible outside of the domain and cannot be shared. See "Assigning Domain Administrators or Security Administrators to Domains" on page 122

Chapter 8: External Certificate Authority

You can configure the Data Security Manager (DSM) to have certificates signed by an external Certificate Authority (CA).

You can configure an external CA on a single node or high availability (HA) deployments. You can set up the DSM to have certificates signed by an external Certificate Authority when the system is set up for the first time, when the system is upgraded, or when the system is in production.

Overview

To configure the DSM to work with an external CA, you must have:

- A valid account with an external CA that is network accessible

- Instructions from the CA explaining how to transfer a certificate request file and a signed certificate file to and from the DSM.

The high-level steps for signing the DSM's Web server certificate with an external Certificate Authority are as follows:

1. Use the CLI `genca` command to generate the DSM's self-signed internal certificate authority and Web server certificates.

   This enables access to the Web-based Management Console.

2. Install the license. In HA systems, install the license only on the initial DSM.

3. Generate the Certificate Signing Request (CSR) file, and save it as a Privacy Enhanced Mail (PEM) file.

   The PEM file contains the information you must submit to the external CA to obtain an approved and signed certificate.

4. Import the signed certificate and the signer's certificate(s).

5. Allow the DSM to restart.

6. If the CA is to be used in an HA environment, repeat steps 4 through 6 for each server.

Installing an External Certificate Authority

Installing an External CA on a Single Node

You can create a new single node system or modify an existing single node system to work with an External Certificate Authority.

Generate a self-signed certificate (genca)

1. Log on to the DSM CLI.

2. Generate the self-signed Certificate Authority certificate. Type:

   ```
   dsm$ system
   system$ security genca
   ```

This command regenerates the CA on the DSM. For more information about the `genca` command, see .

3. Log on to the Management Console as a System/All Administrator. Do not enter a domain.

4. Click **System > License > Upload License File** to upload the license file. This step is required only if this is a new installation or a DSM software upgrade.

Web Server Certificate Information

The Web Server Certificate Info tab displays status information about the existing Web server certificate. It can be used to determine if the certificate has been self or externally signed. It also shows the DSM operating mode with respect to Suite B and consists of the following three fields:

- **Issued To**

  Displays a summary of the data required to generate a CSR including Common Name (CN). CN in this field represents the host name of the device requesting the CSR.

- **Issued By**

  Displays the CN of the Certificate Authority issuing the certificate

- **Valid From**

  Displays the certificate's start and expiration date

Figure 8-1: Web Server Certificate
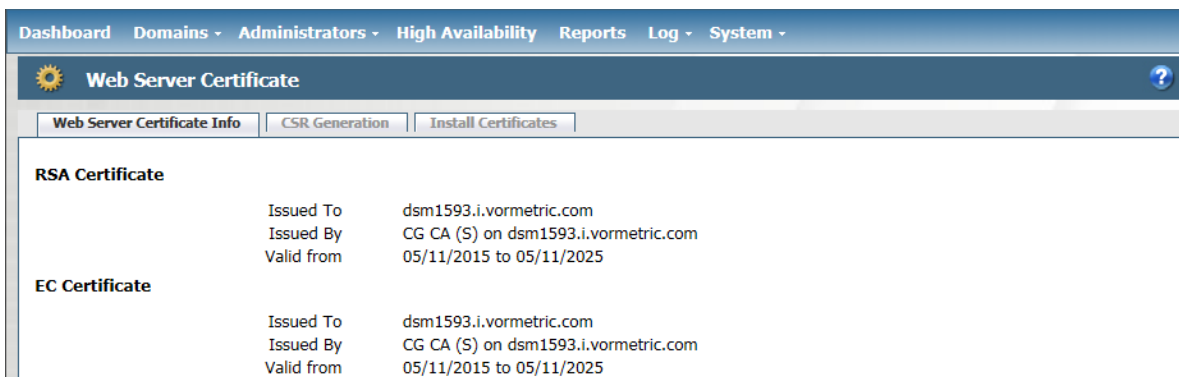
Generate a CSR

1. Select **System > Web Server Certificate** from the Management Console. The Web Server Certificate window opens.

2. Click the **CSR Generation** tab. Enter the information in the fields. If you entered this information while running the `genca` command, the fields (other than the hostname which is updated automatically, but can also be changed) on this tab will contain that same information. You can modify this information if required. Verify that the following pre-populated entries are consistent with the requirements of your external CA. For instance, some CAs will not accept an abbreviation for the name of the city or state.

   ○ **Host Name**

     Network name of the DSM (up to 64 characters). It is possible to edit this field, however it is recommended that you do not change this name.

   ○ **Organizational Unit**

     Typically a department or group name (up to 64 characters)

- ○ **Organization**

  Typically, this is the company name (up to 64 characters)

- ○ **City or locality**

  Location of the Organization (up to 128 characters)

- ○ **State or province**

  Location of the Organization. Refer to external CA for format requirements. Some CAs will not accept an abbreviation for the name of the city or state. (up to 128 characters)

- ○ **Country Code**

  Abbreviation for the country where the Organizational Unit is located (up to 2 characters).

> **Note**
>
> Strings that contain a comma (,) are permitted; however, the use of single or double-quotes in any field on the CSR Generation tab is not allowed.

- ○ **Email address**

  Your valid email address.

Figure 8-2: Certificate Signing Request Tab Information

If you are running the DSM in Suite B or Compatibility mode, when you click Generate CSR, the DSM generates a .zip file containing two PEM files:

- ○ tserver-csr.pem
- ○ EC_tserver-csr.pem

If you are operating in compatibility mode, you will need both PEM files signed. After you receive the signed Web server certificates, install both on the DSM.

3. Click **Generate CSR**. The **File Download** window opens.

4. Click **Save**. The **Save As** window opens.

5. Enter the name and path for the certificate request file. The default file name is *servercsr_<hostname_YYYY_ MM_DD_HHMM>.pem*.

6. Click **Save**. The **Download Complete** window opens. It displays statistical information about the exported PEM file, such as its location and size.

7. Know where the PEM file is saved so you can find it later. Click **Open Folder** to verify the location.

8. Click **X** to close the window.

9. Submit the new CSR to a Certificate Authority for signing/approval.

> **Note**
> Be sure to follow the procedures of the CA to obtain valid certificates. Each CA may have different procedures to obtain the Root certificate, Intermediate certificate, and signed CSR certificate.

Install certificates

1. In the Management Console, on the Web Server Certificate page, click the **Install Certificates** tab.

2. Click **Browse** for the **Root CA Certificate** field and load the Root CA Certificate. The Root CA Certificate is required.

3. If needed, click **Browse** for the **Intermediate CA Certificate** field and load the Intermediate CA Certificate.

4. If needed, click **More** to browse for additional Intermediate CA Certificates. You can select up to ten Intermediate CA Certificates.

5. Click **Browse** for the **Signed Certificate** field and load the Signed Certificate. This is required.

Figure 8-3: Install certificates
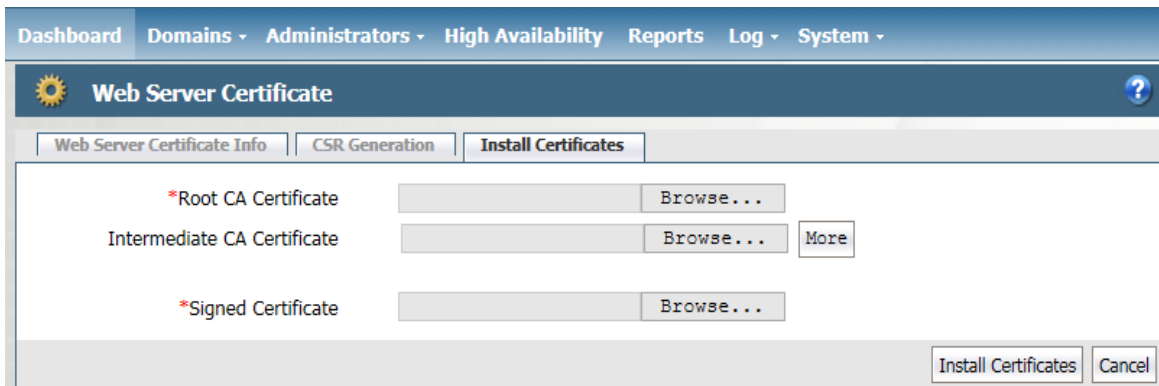
> **Note**
> When you copy a certificate, be certain to copy and paste the certificate just as it appeared originally. Make sure that there are no extra characters or leading spaces as this will invalidate the certificate.

6. Click **Install Certificates**, and then click **OK** to install the certificate and restart the server. The restart takes several minutes.

> **Note**
> During restart, do not close the browser. Do not select Back, Refresh, or the browser Stop buttons.

7. After the server restarts, log on again.

8. To verify the certificate status, click the **Web Server Certificate Info** tab.
   ○ If the Common Name (CN) entry in the **Issued To** and **Issued By** fields shows the same information, the current certificate has been self-signed.
   ○ If the CN entry in the **Issued To** and **Issued By** fields shows different values the current certificate was not self-signed.

If you are running the DSM in Suite B or Compatibility mode, you need to send both the RSA and EC CSRs to be signed by your CA, in which case you need to import both signed certificates; an RSA certificate and an EC certificate.

Installing an external CA in a high availability system

The external Certificate Authority DSM HA system is similar to that of a self-signed DSM HA cluster. To register the agents assigned to a server, the HA nodes must initially have an active connection to the HA cluster.

Once you disable the HA configuration the hosts assigned to a different HA node must re-register before they can operate again. If you anticipate an extended delay in configuring servers, you should reassign the agents to a different HA node before you reconfigure the server.

You can create a new HA environment or modify an existing HA environment to work with an External Certificate Authority.

Generate a self-signed certificate (genca)

> **Note**
> This needs to be done only on the initial node.

1. Log on to the DSM CLI.

2. Generate the self-signed Certificate Authority certificate. Type:

   ```
   system
   security genca
   ```

3. Log on to the Management Console as a System/All Administrator. Do not enter a domain.

4. Click **System > License > Upload License File** to upload the license file. This step is required only if you are doing a new installation, or if you are upgrading the DSM software.

Create a new HA configuration

1. Log on to the DSM as a System Administrator/All. Do not enter a domain.

> **Note**
> The first DSM software version to support an external Certificate Authority is 5.1.1. Be sure the HA nodes are already running the same software version that supports external Certificate Authority.

2. Click **High Availability**, and then click **Add**, the **Add High Availability Server Details** screen opens.

3. Enter the host name or the FQDN in the **Server Name** field, and then click **OK**.

4. Convert to an HA server. Repeat this step for all servers as appropriate.

Administrative Tasks

Tasks in this section are done as required to administer or maintain an external CA environment.

Changing to another external CA

To reconfigure the DSM to use a different CA:

1. Generate a new Certificate Signing Request, see "Generate a CSR" on page 58.

2. Submit the new CSR to a Certificate Authority for signing/approval.

3. Install the new signed certificates from the **Install Certificates** window, see "Install certificates" on the previous page.

Restoring the DSM to a self-signed Certificate Authority

You can restore the DSM to a self-signed Certificate Authority at any time. To revert a DSM to a self-signed Certificate Authority, run the CLI `genca` command. For example:

```
system$ security genca
```

> **Note**
> Reverting to a self-signed Certificate Authority invalidates all configured certificates and they will all have to be regenerated.

The CLI `genca` command overwrites the current server certificate, and must be run to generate a new signer certificate.

Intermediate Certificate Authority

Use the Intermediate Certificate Authority (ICA) page to configure the DSM to have the internal DSM CA signed by an external Certificate Authority (CA).

- "Intermediate CA Info" below
- "CSR Generation" on the facing page
- "Install Certificates" on page 64

> **Note**
> Use of the Intermediate CA is optional.

Intermediate CA Info

The Intermediate Certificate Authority tab displays the following information about the Server/Agent RSA Certificate and the Server/Agent EC Certificate:

- **Issued To**: Displays the host name of the DSM to which the certificate has been issued
- **Issued By**: Displays the name of the CA that has signed this intermediate certificate
- **Fingerprint**: Displays the SHA-256 digest of the certificate
- **Valid From**: Displays the period for which the certificate is valid

Setting up ICA

To correctly set up ICA, DSM must be in compatible mode before the ICA setup. Once ICA is set up successfully, you can set DSM to either RSA or suiteb mode, if desired.

To set the DSM to compatible mode, type:

1. Login to the CLI.
2. Change to the security menu, type:

   ```
   # system
   ```
3. Set the system to compatible mode, type:

   ```
   # security suiteb set compatible
   ```
4. Confirm the change, type:

   ```
   # yes
   ```

CSR Generation

> **Note**
> Consider the certificate validation period when getting the DSM CA certificates signed. Each agent and DSM in a cluster needs to re-register every time the DSM CA certificates are updated. Signing the DSM CA certificates with shorter validation periods requires re-registering every entity in the cluster more often. This also extends to all certificates in the CA chain, as an expired certificate from any of them will require the DSM CA certificates to be re-signed. Thales recommends signing the DSM CA certificates with a validation period of 10 years.

Use the CSR Generation tab to generate a certificate signing request.

1. Navigate to **System > Intermediate CA**.

2. Click **CSR Generation** tab.

3. If certificate information was filled in during the initial configuration of the DSM when running the `genca` command, the form on this page is pre-populated with that information. Verify or enter the information in the fields.

> **Note**
> If the DSM is in compatibility mode, generate both types of certificates. If using suiteb mode, use the EC certificates. If using RSA mode, use the RSA certificates.

  ○ **Host Name**

    Network name of the DSM (up to 64 characters). It is possible to edit this field, however it is recommended that you do not change this name.

  ○ **Organizational Unit**

    Typically a department or group name (up to 64 characters)

  ○ **Organization**

    Typically, this is the company name (up to 64 characters)

  ○ **City or locality**

    Location of the Organization (up to 128 characters)

  ○ **State or province**

    Location of the Organization. Refer to external CA for format requirements. Some CAs will not accept an abbreviation for the name of the city or state. (up to 128 characters)

  ○ **Country Code**

    Abbreviation for the country where the Organizational Unit is located (up to 2 characters).

> **Note**
> Strings that contain a comma (,) are permitted; however, the use of single or double-quotes in any field on the CSR Generation tab is not allowed.

  ○ **Email address**

    Your valid email address.

4. Click **Generate RSA CSR** or **Generate EC CSR**.

5. After generating the CSR, a file download dialog box displays, prompting you to select a location to save the .zip file that contains the CSR. The file format is `hostname-<YYYY_MM_DD_HHMM>-ec-csr.zip` or `hostname-<YYYY_MM_DD_HHMM>-rsa-csr.zip`, depending on the type of CSR generated. Each zip file contains two CSRs, each of which must be signed by the external CA.

6. Download the resulting zip file and extract the two CSRs inside.

7. For each CSR, open it in a text editor and copy the contents.

PKI infrastructure

When you are setting up your PKI infrastructure and creating certificates, note that the DSM **requires** the ICA certificate to include the x509v3 extensions. To create an intermediate certificate, set the CA attribute as follows:

```
X509v3 extensions
X509v3 Basic Constraints: critical
CA: TRUE
```

Obtaining an external certificate

> **Note**
> The following example uses Microsoft Active Directory Certificate Services through Certification Authority Web Enrollment. Other certificate services will differ slightly in their methods.

1. In the DSM, navigate to your web enrollment URL and login.

2. Click **Request a certificate** which takes you to the Request a Certificate page.

3. Click **advanced certificate request**, which opens the Submit a Certificate Request or Renewal page.

4. Paste your CSR into the certificate request box.

5. Select **Subordinate Certification Authority** in the Certificate Template pull-down menu.

6. Click **Submit** to request your certificate.

7. In the Certificate Issued screen, select the **Base 64 encoded** option.

8. Click **Download certificate chain** to download your new certificate chain.

9. Repeat the previous steps for any remaining CSRs.

10. Return to the DSM and click **System > Intermediate CA**.

Install Certificates

The file containing the DSM CA signed certificates for installation must also contain the entire certificate chain of CAs back to a root CA. The certificates must be in PEM format, must have keyCertSign and CRLSign key usages, and must also be in the correct signing order, with the DSM CA certificate at the top, followed by its signer certificate and so on until the root CA certificate which must be the last certificate at the end of the file.

> **Note**
> Installing a new, or renewed, intermediate CA certificate causes the DSM to restart. Any registered agents must be re-registered.

Install Certificate Chain

> **Note**
> This certificate chain format is for Unix only.

For example, for a DSM CA certificate signed by CA1, where CA1 is signed by CA2, which in turn is signed by CA3, which is signed by the root CA, then the order of certificates in the file must be the following:

```
-----BEGIN CERTIFICATE-----
```

```
        (DSM CA cert)
    -----END CERTIFICATE-----
    -----BEGIN CERTIFICATE-----
                (CA1 cert)
    -----END CERTIFICATE-----
    -----BEGIN CERTIFICATE-----
                (CA2 cert)
    -----END CERTIFICATE-----
    -----BEGIN CERTIFICATE-----
                (CA3 cert)
    -----END CERTIFICATE-----
    -----BEGIN CERTIFICATE-----
                (RootCA cert)
    -----END CERTIFICATE----
```

To install the certificate:

1. On the DSM, click the **Install Certificates** tab.

2. Click **Choose File/Browse** to select the certificate chains (RSA or EC) to upload.

   > **Note**
   > You can upload both pairs at once, or only the RSA pair or only the EC pair.

3. Click **Install Certificates** and wait for the DSM to restart before logging in again.

4. Click **Intermediate CA Info tab**. The content should look like the following:

   > **Note**
   > The **Issued To** and **Issued By** fields are different which indicates the certificates are no longer self-signed.

Importing the Root Certificate and Verifying a Secure Connection

To obtain a secure connection (green lock status) in your browser, import the "root CA" certificate into either your browser's certificate store, or the Windows certificate store.

> **Note**
> Some browsers have their own certificate store like Firefox. Chrome and IE/Edge use the Windows certificate store.

The following example is from Firefox.

1. Select **options** from the menu and enter **cert** in the search field.

2. Click **View Certificates > Authorities** tab.

3. Click **Import** and import the root ca certificate: *_root_ca.cer

4. Click **Ok** and then browse to your DSM using its hostname, which must match the hostname in the certificate.

5. In the URL field, you should see a green lock icon next to the URL. This indicates a secure connection.

Chapter 9: LDAP Configuration

The DSM allows for integration with Lightweight Directory Access Protocol (LDAP) directory services such as Active Directory (AD) and OpenLDAP. This feature allows the DSM Administrator to import user criteria instead of recreating it from scratch.

Configuring LDAP

An LDAP server must be configured and authenticated before any information can be imported.

Configuring LDAP Timeout

You can now configure the duration for which the DSM tries to connect to the LDAP server.

- **Minimum time**: 1 s

- **Maximum time**: 600 s

- **Default time**: 30 s

To change the timeout from the default value:

1. Click **System > General Preferences > System**.

2. Scroll down to the LDAP timeout setting.

3. Set the desired value.

4. Click **Apply**.

Configure LDAP server settings

1. Log in and select **System > LDAP**.

2. Enter the hostname of the LDAP server in the **Directory URL** field. If a secure LDAP hostname is specified here, then its LDAPS Server certificate in PEM format must also be entered in LDAPS Server Certificate.

> ⚠️ **WARNING**
> **The hostname in the Directory URL field must match the hostname in the LDAPS Server certificate or the connection fails.**

Examples:

```
ldap://ldapserver.mycorp.com:389
ldaps://ldapserver.mycorp.com:636
```

> **Note**
> The default LDAP port is 389. The default LDAPS port is 636.

3. (Optional) Enter the URL of an alternate LDAP server, in the **Secondary URL** field. This alternate LDAP server will be used if the initial LDAP server is unreachable. If you enter a secure LDAP path, you should browse to the location of a certificate in the **LDAPS Server Certificate** field and upload the certificate for that server.

4. Enter a **Base Distinguished Name**. For example, if you use Active Directory with a domain name such as "mycorp", your base DN would be DC=mycorp, DC=com

5. (Optional) Enter up to a 256 character string to filter searches, in the **LDAP Query** field.

6. (Optional) Enter the LDAP user login name in the **Login** field.

   Example: If your domain name is "mycorp" and using Active Directory with a domain controller your login name might be:

   jsmith@mycorp.com

7. Enter the LDAP password in the **Password** field. Enter it again in **Confirm Password**.

   > **Note**
   >
   > The LDAP user name and password details entered here are cached, so that you do not need to enter them every time you import an administrator or an email address for email notifications.
   > You may also enter a different Login and Password in place of these stored values when you import administrators.

8. **LDAPS Server Certificate**: If a secure LDAP path was entered in the Directory URL field, click **Browse** and navigate to the location of the Root CA Certificate. The CA certificate must be in PEM format. This field does not allow direct user input to avoid typographic errors.

   > **Note**
   >
   > If LDAPS is used for the Directory URL or Secondary URL fields, you must upload a certificate in the LDAPS Server Certificate field. The certificate must be in PEM format.

9. The **CA Certificate Exists** box will be checked if the LDAPS Server Certificate has been uploaded to the DSM.

User Schema Settings

1. Enter the "Object Class" attribute in the **User Object Class** field.

   For example: user or person

2. Enter the user attribute containing the unique user ID in the **Login Name Attribute** field. This is the AD/LDAP schema attribute to be used as the LDAP user login name.

   For example: sAMAccountName or commonName

   > **Note**
   >
   > If a Login Name already exists in the DSM database, the **Import** function will not overwrite existing users with the same login name.

3. (Optional) Enter the user attributes desired in the **User Description Attribute(s)** field. To enter multiple attributes, separate values with a semicolon.

   For example: name or description

4. (Optional) **Email Attribute**: This is the AD/LDAP schema attribute to be used as LDAP user email. For example: userPrincipalName or mail

Group Schema Settings

1. Enter the group "Object Class" attribute in the **Group Object Class** field.

   For example: group or posixGroup

2. Click **OK** to save the settings on the page, or click **Clear** to clear the form. You can also click **Clear** any time later to delete the AD/LDAP settings.

Importing LDAP Administrators

The Import function allows Administrators to import data from an LDAP server such as Active Directory (AD) or OpenLDAP. Once an LDAP server has been identified and configured, the DSM Administrator can import the desired values. To set up access to an AD/LDAP repository, see, "Configuring LDAP" on page 66. You will need an LDAP login ID and password.

To import values from an LDAP directory:

1. Select the **Administrators > All** tab. Click **Import**.

2. Enter the **Login ID** and **Password**. If the Login and Password were entered under **LDAP Server Settings** on the AD/LDAP Details window, these values will be populated and do not need to be re-entered. You may also enter a different Login and Password in place of these stored values when you import administrators.

3. Click **Connect**.

4. The LDAP Users window displays LDAP user names.

Selecting LDAP administrators

The following search option are available on the LDAP Users window:

- **LDAP Query**—Use the field to filter searches using the LDAP query language. Results depend on how the LDAP service is set up. See RFC2307 for full details on syntax.

- **Group**—Select a group from the drop down list.

- **User**—Enter a user name.

- **Maximum number of entries to return**— Limits the maximum number of records to import or display. The default value is 300. The minimum value is 1 and the maximum value is 1 million. A high integer value may result in a delay depending on the database size.

- **Go**—Click to refresh the screen

- **Select All**—Click to select all values on that page

- **View**—Select a value from this drop down box to control how many values appear on any page

- **Selected**—Click to select individual values.

- **User Type**—Select a value from this drop down box to define the type of Administrator or role of the values you import.

- **Add/Cancel**—Select to add or cancel your selections.

> **Note**
> The introduction of the multi-tenancy feature allows the creation of local domains. Each local domain can have its own specific LDAP server. The LDAP server can be configured by the local domain administrator, or a local Domain and Security administrator.

Selecting LDAP users for email notifications

The DSM can be configured to send email notifications about fatal and error conditions on the DSM. This can be done at a system level outside a domain, at a global domain level, or at a local domain level. System level and global domain level administrators use system level LDAP servers and local domain level administrators use local domain level LDAP servers.

To select LDAP users to receive email notifications, do the following:

1. Select **System > Email Notification**.

    > **Note**
    > If an SMTP server has not been configured, the following message is displayed "SMTP is not set." Click the **SMTP Server** tab to configure an SMTP server. A warning will also be displayed if the SMTP server is not correctly configured.

2. Click **Add** to add a group of users who will receive an email notification.

    The Add Email Notification Group window has the following fields that must be configured to enable notifications:

    ◦ **Email Group Name**

    Name of the email group that will receive the email notification.

    ◦ **Email Address List**

    Email addresses that will receive this email notification. Separate addresses with commas.

    If LDAP is configured, you can select addresses from your LDAP address book by clicking **Select**. The Connect to AD/LDAP Server window opens. If the Login and Password were entered under **LDAP Server Settings** on the AD/LDAP Details window, these values will be populated and do not need to be re-entered. If it's not configured, you can enter your login and password to access it. Select the check boxes for those users who are to receive the notifications and click **Add**.

    ◦ **Email Subject**

    Text you want on the subject line.

    ◦ **Notification Type**—

    Generic, Key Expiration, Certificate Expiration. The Generic option is visible both inside and outside a domain, the Key Expiration and Certificate Expiration options are visible only inside a domain.

    ◦ **Email Threshold Level**

    Select either ERROR or FATAL. If the DSM generates a log message with a severity of this specified threshold level, then an email notification is generated. The ERROR threshold option sends log messages about errors and fatal errors since fatal is a subset of error. The FATAL threshold option sends only log messages about fatal errors.

    ◦ **Message Contains**

    This is a string filter that works with the Email Threshold Level. Only messages containing this string will be sent as an email notification. If left blank, then all messages meeting the threshold criteria will be sent.

    ◦ **Enabled**

    A check box that enables or disables email notification to the group.

3. Click **OK**.

Chapter 10: Configuring Syslog Servers for System-Level Messages

This section describes how to add a remote Syslog server to your system, and how to control the severity level and format of the messages that the DSM sends to the Syslog server.

You can configure a Syslog server to receive the same messages that are sent to the Logs window of the Management Console. Use the **System > Log Preferences** menu to create templates that apply to logging configurations for all of the Agents.

Overview

Agent log data is generated on agent hosts. The log data is placed in `/var/log/vormetric` on a UNIX system or in `C:\Documents` or `Settings\All Users\Application Data\Vormetric\DataSecurityExpert\agent\log` on a Windows system, when the **Log to File** logging preference is enabled. The log data can also be forwarded to a Syslog or Event Log server when the **Log to Syslog/Event Log** logging preference is enabled.

> **Note**
> Ensure that the `/var` directory in your system has 256KB to 1MB available for logging to ensure proper DSM logging behavior.

When **Log to Syslog/Event Log** is enabled, log data is placed into a local `/var/log/messages` or `/var/adm/messages` file, or into the local Windows Event Log. The host administrator can choose to configure the agent to forward log data to a remote Syslog server or Event Log server. The host administrator can upload the log data to a remote server with whichever transport protocol is preferred. The DSM is not used to configure the remote log servers for host systems.

The **Syslog Server** window in the Management Console lets you configure the remote Syslog servers to which to send DSM log data. The log data sent to remote Syslog servers consists of log data that is generated on the DSM and, when **Upload to Server** is enabled in the **Log Preferences** window, log data that is generated on hosts. The administrator then configures the DSM to forward log data to a Syslog server using either UDP protocol or TCP protocol.

Figure 10-1: Handling log messages
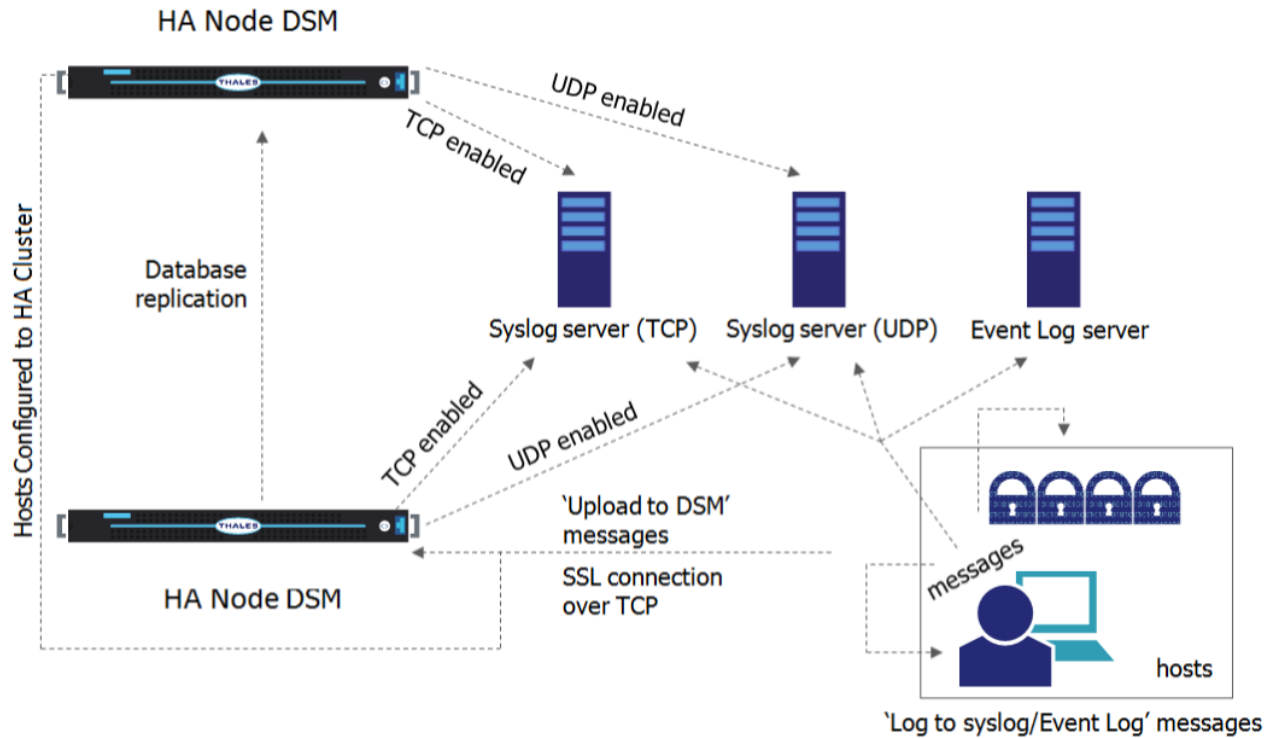
Items to consider before configuring Syslog logging include:

- Only System/All Administrator can enable Syslog messaging.

- Only System/Domain or All Administrators can configure Syslog messaging.

- If Syslog servers are configured in a domain, only events that take place in that domain are logged to the Syslog servers.

- If Syslog servers are configured outside of a domain, only events that take place at the system level are logged to the Syslog servers.

- A default Syslog port number is not provided. The usual industry standard port number for Syslog over UDP is 514. Port 1468 has been used successfully for TCP.

- Configuring a Syslog server is an effective way to consolidate the logs of all the DSMs in an HA configuration in one central repository. The HA nodes in an HA cluster deployment all have the same configuration. The nodes forward log data to HA node 1. Therefore, each node must have network access to the Syslog servers configured on HA node.

Supported Syslog Formats

The DSM supports the following log formats:

- Plain Message

- Common Event Format (CEF)

- RFC5424

- Log Event Extended Format (LEEF)

Plain Message

Originally, DSMSyslog support included only Plain Message format. While simple and efficient, this format did not allow for user enhanced reporting or customization.

The following is an example of a Plain Message formatted log message. The table following the message describes the components of the message.

```
12-07-201216:53:02Local7.Debug10.3.32.2312012-12-08 01:01:58.709 vormetric:SOURCE[linux64-
32231.qa.com]:DAO0445I:Administrator voradmin added SysLog Host 10.3.25.168.
```

Table 10-1: Syslog message parameters and descriptions

Parameter	Description
`12-07-201216:53:02`	Date and time
`Local7.Debug`	Message priority
`10.3.32.231`	Sending machine's IP address
`2012-12-08 01:01:58.709`	Date and time of logged event
`vormetric`	Originator tag
`SOURCE[linux64-32231.qa.com]`	Source of message
`DAO0445I`	Unique message ID
`Administrator voradmin added SysLog Host 10.3.25.168`	Plain text message of the logged event

Common Event Format (CEF) log format

The DSM Syslog supports Common Event Format (CEF) log format. The CEF format is specified in the Arcsight "Common Event Format" standard.

The following is an example of a CEF formatted log message.

```
<27> 2012-10-16T16:01:44.030Z centos-6-0 CEF:0|Vormetric, Inc.|vee-fs|5.1.0.9026|CGP2604E|
Reject access|7|logger=CGP spid=6362 cat=[ALARM] pol=AuditAllExceptLp
uinfo=lp,uid\=4,gid\=7\\lp\\ sproc=/bin/ls act=read_dir_attr gp=/Guard
filePath=/datafiles/file.dat denyStr=DENIED showStr= Code (1M)
```

Table 10-2: CEF Log Format parameters and descriptions

Parameter	Description
`<27>`	A standard syslog facility/priority code
`2012-10-16T16:01:44.030Z`	Date and time
`centos-6-0`	The host name of the machine sending the message.
`CEF:0`	Version of the CEF
`ThalesGroup, Inc.`	Sending device vendor
`vee-fs`	Sending device product
`5.1.0.9026`	Sending device version

Table 10-2: CEF Log Format parameters and descriptions (continued)

Parameter	Description
CGP2604E	Unique message ID
Reject access	Name: A human-readable and understandable description of the event.
7	Severity: An integer that reflects the importance of the event. Only numbers from 0 to 10 are allowed, where 10 indicates the most important event.
logger=CGP spid=6362 cat=[ALARM] pol=AuditAllExceptLp uinfo=lp,uid\=4,gid\=7\\lp\\ sproc=/bin/ls act=read_dir_attr gp=/Guard filePath=/datafiles/file.dat denyStr=DENIED showStr= Code (1M)	Extension: A collection of key-value pairs. The keys are part of a predefined set. The standard allows for including additional keys. An event can contain any number of key-value pairs in any order, separated by delimiting characters.

RFC5424

The DSM Syslog support includes the RFC5424 log format.

An example of an RFC5424 formatted log message follows. Components of the message are described in the table following the message example:

```
<30>1 2012-12-07T21:44:04.875Z t3-normaluser.i.vormetric.com vee-FS 0 CGP2603I [CGP@21513
sev="INFO" msg="Audit access" cat="\[AUDIT\]" pol="normaluser-only-aes256"
uinfo="normaluser,uid=2001,gid=1\\other\\" sproc="/usr/bin/cat" act="read_attr"
gp="/export/home/normaluser/test" filePath="/test.txt" denyStr="PERMIT" showStr="Code
(1M)"]
```

Table 10-3: CEF Log Format parameters and descriptions

Parameter	Description
<30>1	A standard syslog facility and priority code
2012-12-07T21:44:04.875Z	Date and time
t3-normaluser.i.vormetric.com	The host name of the machine sending the message.
vee-FS	Sending device product
0	Process ID field having no interoperable meaning, except that a change in t he value indicates that there has been a discontinuity in syslog reporting.
CGP2603I	Unique message ID
[CGP@21513 sev="INFO" msg="Audit access" cat="\[AUDIT\]" pol="normaluser-only-aes256" uinfo="normaluser,uid=2001,gid=1\\other\\" sproc="/usr/bin/cat" act="read_attr" gp="/export/home/normaluser/test" filePath="/test.txt" denyStr="PERMIT" showStr="Code (1M)"]	Structured data field: Provides a mechanism to express information in a well-defined, easily parsable and interpretable data format. This field consists of the Structured Data (SD) Element, SD-ID, and SD-Parameter.

Log Event Extended Format (LEEF)

The DSM Syslog support includes Log Event Extended Format (LEEF). The LEEF header is pipe ("|") separated and attributes are tab separated.

Adding a Syslog Server

To add a syslog server:

1. Verify that one or more Syslog servers are accessible from the DSM. It is usually enough to ping the Syslog server and run `ps` to check the Syslog process on the Syslog server system.

   If you are going to send the messages to the local host, verify that the syslogd process on the local host is accepting messages. You may need to restart syslogd with the "-r" argument.

   > **Note**
   > Record the Syslog transport protocols and port numbers of the Syslog server(s). You will need this information later.

2. Set the severity level at which to send messages to the Syslog server in the `/etc/syslog.conf` file on the agent host.

   Severity levels in the Log Preferences window are DEBUG, INFO, WARN, ERROR, and FATAL. Severity levels are cumulative, so each level includes the levels below it. For example, FATAL logs only FATAL messages, whereas WARN logs WARN, ERROR, and FATAL messages. To ensure that the syslog server gets the messages set in the Log Preferences window, set the level in the syslog.conf file to debug and direct the output to the local messages file. For example, on a Solaris system, set the output file path to /var/adm/messages.

   ```
   user.debug /var/adm/messages
   ```

3. Log on to the Management Console as a System/All Administrator.

4. Select **System > General Preferences**. The **General Preferences** window opens to the **General** tab.

5. Click the **System** tab, and then select **Syslog Enabled**.

   This enables communication between the DSM and the Syslog server.

   > **Note**
   > You must have the **Syslog Enabled** box selected from outside a domain; otherwise, the **Apply** button will not be selectable from within a domain.

6. Click **Apply**.

7. Select **System > Log Preferences**. The Log Preferences window opens to the Server tab.

8. Set the **Logging Level** property.

   The level you select affects the number of messages that are displayed in the Logs window, and these messages are also sent to the Syslog server.

   Redundant Syslog failure messages are filtered so that only one out of every fifty redundant messages is sent to `/var/log/messages` and the **Logs** window. All the redundant Syslog failure messages are sent when the level is set to DEBUG.

9. Click **Apply**.

   ° If you are configuring a Syslog server to receive system-level log data, remain logged in (for example, 'system-level' is when you are not in a domain).

- ○ If you are configuring a Syslog server to receive domain-level log data, and are logged in as an All administrator, remain logged in and enter the domain to be configured.

- ○ If you are configuring a Syslog server to receive domain-level log data, and are logged in as an System Administrator, log out and log back in as a Domain or All administrator, and enter the domain to be configured.

10. Select **Log > Syslog**. The **Syslog Server** window opens.

11. Click **Add** and enter the following information:

   a. **Server Name:** The host name or FQDN of a Syslog server. Use the network name of a Syslog server which is accessible to all of the nodes in the HA cluster.

   b. **Transport Protocol:** Select UDP, TCP or TLS from the drop down. If you select TLS, a field appears for you to browse to add a **Root Certificate**.

   In the interests of security, we recommend that you use a root certificate rather than a non-root certificate.

> **Note**
> For syslog servers configured with the UDP transport protocol, ensure that UDP packets are not blocked by a firewall or switch rules. Also, verify that the Syslog server is logging messages as expected.
> If you add a Syslog certificate when using TLS protocol, you may need to restart the server. To this you need to do a `system > server restart` from the CLI. After restart, verify that the Syslog server is logging messages as expected.

   c. **Port Number:** The port number the transport protocol uses to connect to the Syslog server. Enter a value between 1 and 65535. There is no default.

   d. **Message Format:** Select Plain Message, CEF, or RFC5424.

   You may configure multiple Syslog servers per DSM however, each Syslog server must have a unique hostname or IP address.

12. Click **Ok**.

13. Do a task on an agent system that normally generates a Syslog entry, such as accessing a GuardPoint.

14. Check the `/var/log/messages` file on the Syslog server for DSM log entries.

Using Syslog to Troubleshoot the DSM and Agents

Syslog entries for DSM activity indicate the source of the Syslog message (system name after the timestamp), the source of the message itself (SOURCE), the log level (AUDIT, ALARM, and so on), and much more.

Analyzing log entries

The format and content of log entries for VTE Agents are described below.

Figure 10-2: Message Log entries

Analyzing VTE Agent log entries

The general format of a VTE Agent log entry is:

```
CGP2602I: [SecFS, 0] Level: Policy[policyName?] User[userID?] Process[command?] Access
[whatIsItDoing?] Res[whatIsItDoingItTo?] Effect[allowOrDeny? Code (whatMatched?)]
```

where:

- SECFS indicates that the message was generated by a VTE Agent. You can enter `secfs` in the **Search Message** text-entry box in the **Logs** window to display VTE Agent policy evaluation and GuardPoint activity for all configured hosts.

- Level indicates the importance of the message. For example, AUDIT indicates an informational message, whereas ALARM indicates a critical failure that should not go ignored.

- Policy[] indicates the name of the policy that is being used to evaluate the access attempt.

- User[] identifies the system user attempting to access data in the GuardPoint. It typically displays the user name, user ID, and group ID.

- Process[] indicates the command, script, or utility being executed.

- Access[] indicates what is being attempted. Access may be read_dir, remove_file, write_file_attr, write_app, create_file, etc. These correspond to the Access methods that you configure in the policy. read_dir corresponds to d_rd. remove_file corresponds to f_rm. And so on.

- Res[] indicates the object being accessed by Process[].

- EFFECT[] indicates the rule that matched and, based upon that rule, whether or not the DSM grants access. Access states may be either PERMIT or DENIED.

For example:

```
CGP2606E: [SecFS, 0] [ALARM] Policy[allowAllRootUsers_fs] User
[hpotter,uid=1111,gid=10\wheel\] Process[/usr/bin/vim] Action[create_file] Res
[/opt/apps/apps1/lib/file1.txt] Effect[DENIED Code (1M)]
```

The format of a rule match is:

```
intchar
```

where:

- int is an integer representing the security rule being used or violated. Security rules are numbered sequentially from top to bottom in the Online Policy Composer window.

- $char$ is an uppercase letter indicating the item that is using or violating the policy.

Table 10-4: Character Codes and Their Descriptions

Character Code	Description
A	The **Action** component of a security rule failed to match.
M	All security rule components match and, unless overridden, the **Effect** for that security rule is applied.
P	The **Process** component of a security rule failed to match.
R	The **Resource** component of a security rule failed to match.
T	The time specified in the **When** component of a security rule failed to match.
U	The **User** component of a security rule failed to match.

For example, the following match codes indicate:

- **1R** – Mismatch in Resource for Security Rule 1.

- **3U** – Mismatch in User for Security Rule 3.

- **4A** – Mismatch in Action for Security Rule 4.

- **2M** – All components matched for Security Rule 2. Since all the rules matched, Security Rule 2 will be used and no other rules will be evaluated.

Log message levels

The detail and extent of information logged is determined by the selected log level. The agent supports five log levels as listed in Table 10-5: "The Agent-Supported 5 Log Levels" below.

Table 10-5: The Agent-Supported 5 Log Levels

Severity	Description
DEBUG	The DEBUG level provides detailed information about events that are intended for support engineers and developers.

Table 10-5: The Agent-Supported 5 Log Levels (continued)

Severity	Description
INFO	The INFO level provides general information that highlights the progress of the application.
WARN	The WARN level designates potentially harmful situations.
ERROR	The ERROR level designates error events that might still allow the application to continue running.
FATAL	The FATAL level designates very severe error events that will presumably lead the application to quit.

Log levels are cumulative. The level that you select not only generates log entries for events that occur at that level, but all the levels below. For example, the WARN level also includes events that occur on the ERROR and FATAL levels.

Using log files

Check the log files to verify the successful installation and configuration of the DSM software, to determine why a backup or restore operation failed, or to monitor DSM activity.

A logged event falls into one of the following categories:

- **Operational status.** The result of any significant action performed by an VTE Agent or DSM is logged.

- **Administrative activity.** The result of any maintenance or administrative activity on the DSM is logged (for example, a key has been created or exported).

- **System status.** The result of any system errors are logged (for example, if the database connection is interrupted).

- **Policy-specified audit.** If the result of a policy evaluation specifies that it should be audited, then a suitable message is logged.

Several logs files are provided. Each serves a different purpose.

(Windows only) The `\ProgramData` folder on Windows Vista and Windows Server 2008, and the `\Documents and Settings\All Users\Application Data` folder for all other supported Windows platforms, are hidden by default. VTE Agent logs, configuration data, and certificates are stored under that folder. If you cannot browse the folder for your platform, enable the **Show hidden files and folders** radio button in the **Folder Options** menu to view the folder and its contents.

Active logs are log files that being currently written to and updated by DSM processes. Inactive logs are logs that have been filled to capacity and then closed. The name of the closed log file is the original name usually appended with the date and some random numbers. For example, the name of an active agent log is vordb2_usr.log. When it reaches the configured capacity, it is made inactive and usually renamed to `vordb2_usr.log.YYYY-MM-DD-MM-SS.tar.gz`. For example, the archive file for `vordb2_db2inst1.log` can be `vordb2_db2inst1.log.2011-01-19-12-25-32`.

Do not try to manually modify or remove active logs. Use the Management Console interface to configure server and VTE Agent logs. Regularly back up and delete inactive logs to maximize available hard disk space.

The Windows system event log can fill quickly. If a Windows host runs out of system event log space, the `vmd` service does not start and issues an error: "`The service did not respond to the start or control request in a timely fashion.`" To prevent the system event log from running out of space, the current event log is archived to a file when it reaches 20MB, all archived entries are then purged from the event log, and logging continues as usual. Archive files are placed in `%SystemRoot%\System32\Config`. The archive file is named `Archive-Vormetric`

`Encryption Expert-timestamp.evt`. For example, `Archive-Vormetric Encryption Expert-2010-05-14-18-14-30-171.evt`. The file is archived in a binary format that you can open in the Event Viewer. Check disk space availability during periods of heavy load and extensive logging. Back up and delete the archive files.

VTE Agent Log Files

The agent logs are the first places to check when communication between the DSM and VTE agent system fails. Also, you may want to check these logs after setting up a new agent or changing the agent configuration.

Sample logging formats include the following:

vorvmd.log (Windows)/vorvmd_root.log (UNIX)

(UNIX)

```
/var/log/vormetric/vorvmd_root.log
```

(Windows)

```
\Documents and Settings\All Users\Application
Data\Vormetric\DataSecurityExpert\Agent\log\vorvmd.log
```

(Windows XP)

```
\Documents and Settings\All Users.WINDOWS\Application
Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log
```

(Windows Vista and Windows Server 2008)

```
\ProgramData\Vormetric\DataSecurityExpert\Agent\log\vorvmd_root.log
```

(Windows) The same information that is sent to `vorvmd.log` can also be sent to the Windows Event Viewer. Enable **Log to Syslog/Event Log** logging options for the agents and open **Event Viewer > Vormetric Encryption Export** to view log events on the host system.

`vorvmd_root.log` contains the VTE Agent transactions for the root user. Transactions consist of a record of vmd actions, such as starting the vmd daemon and setting up communication links with the DSM.

messages (UNIX only)

```
/var/log/messages
```

messages is a Syslog-generated file. It contains standard Syslog entries. It contains kernel entries for enabling/disabling the log service, memory usage, CPU usage, system calls, device initialization, etc. It also contains log entries that are also displayed in the Message Log.

secfs.log (AIX only)

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. The `secfs.log` file is generated only on AIX systems. The secfs.log file is maintained in the `./agent/secfs/tmp` directory. It is used instead of Syslog to log kernel messages. The same log messages are placed in both `/var/log/messages` and `secfs.log`. The `secfs.log` file is archived at 32MB and renamed to `secfs.log.archive`. Only one archive file is maintained.

secfsd.log

(UNIX)

```
/opt/vormetric/DataSecurityExpert/agent/secfs/tmp/secfsd.log
```

(Windows Server 2003)

```
C:\Documents and Settings\All Users\Application
Data\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

(Windows Vista and Windows Server 2008)

```
C:\ProgramData\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

(Windows XP)

```
C:\Documents and Settings\All Users.WINDOWS\Application
Data\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. secfsd.log contains a record of GuardPoint mounts and GuardPoint dismounts (GuardPoints are mounted file systems). Entries are added to this file when you add and remove GuardPoints, as well as when you reboot the agent system.

statusfile

```
/opt/vormetric/DataSecurityExpert/agent/secfs/tmp/statusfile
\Program Files\Vormetric\DataSecurityExpert\agent\secfs\tmp\statusfile
```

`statusfile` is a current record of the local VTE Agent configuration. View this file after updating the VTE Agent configuration on the DSM to verify that the changes have actually been applied. This file should always be checked when the configuration of the VTE Agent is in question. This file lists:

- Each GuardPoint and GuardPoint properties, such as the lock status, protection status, and GuardPoint directory

- The names of applied policies

- The logging information that is captured

- Where captured log information is sent

- Hosts settings

You can also display the file timestamp to see when the agent was last updated.

This file is deleted each time the VTE Agent configuration is updated. You must manually regenerate it using the "`secfsd -status`" command. If you want to keep records of VTE Agent configuration changes, either copy the `statusfile` to a different name, or run "`vmsec status`" and assign the output to a different file.

(Windows) The `secfsd` command has limited support on Windows platforms. You can use the `secfsd -status lockstat` command or use the Vormetric Data Security tray to open the status window. Look for strings like `coreguard_locked=true` and `system_locked=true`. (`false` indicates that a lock is not applied. `true` indicates that a lock is applied.).

You may view the file contents using an ASCII display command, such as `cat`.

DSM Log Files

DSM logs are logs on the DSM system. The primary log is viewed in the **Logs** window of the Management Console. This log is generally the first log that you check to diagnose DSM problems. Check the DSM log after making or restoring a database backup. Look for entries like "`Backup Request for SAMPLE from host vmSSA06 is allowed.`" and "`Backup/Restore completed successfully.`" Messages like "`Backup data request failed: access denied or a related cause.`" indicate a problem has occurred and some debugging on your part is required.

Appliance-based DSM installations must use the `diag` CLI command to list and view the log files. However, the log files can be exported from an appliance using the various export features in the Logs window. Appliance-based server administrators cannot delete log files.

The DSM server creates three log files:

- `boot.log` contains JBoss startup information

- `cgss.log` contains server information

- `server.log` contains system-level information

`boot.log` is managed as a single file. It is not expected to ever become a large file nor is it rotated. The `cgss.log` and `server.log` files can become large and are rotated.

The `cgss.log` and `server.log` files are important log files that can grow quickly under heavy load. Because these logs are vital to analyzing DSM behavior, they should be monitored and backed up regularly.

The names of the active files are `cgss.log` and `server.log`. When either file contains 10MB of log data it is made inactive and renamed to `cgss.log.1` or `server.log.1`, respectively. And a new active `cgss.log` or `server.log` file is opened. When the new active log file reaches 10MB it is made inactive and renamed to `cgss.log.2` or `server.log.2`. And a new active log file is opened. This process continues until there are a total of 10 inactive log files. When there are 10 inactive log files, and the active log file reaches its full 10MB capacity, the first inactive file is discarded, all the other log file names are decremented by one, and the former active log becomes the 10th inactive log file. Using `cgss.log` as an example, when `cgss.log` fills, `cgss.log.1` is thrown away, all the other log file names are decremented by one, and `cgss.log` becomes `cgss.log.10`. Depending on how much load you place on the server, and if your policies audit a lot of data, these files can grow and rotate quickly.

badlog.log

Log files with unparsable data are "bad logs". A badlog.log file contains log data from an agent that is intended for display in the Logs window but which cannot be displayed because the log data cannot be parsed due to format irregularities. Each attempt by an agent to upload an unparsable log file to the server is placed in the badlogs directory as a unique file. Regardless of the number of failed attempts to parse incoming log files, the DSM will continue to accept uploaded logs from the agent.

cgss.log

The cgss.log file contains a record of the events that make up the BEK generation process for an agent requesting to make a backup, as well as the names of uploaded audit files. This file does not contain events that pertain to restore operations. Check this file if the agent fails to back up a database, even though agent/server authentication is correctly configured and the policy for this agent permits the backup operation.

jboss.log

The `jboss.log` file contains information that is related to starting and stopping the JBoss Web application server. This file is generated when the `/etc/init.d/cgss` command is used to start and stop JBoss. Check this log file for problems that are related to JBoss, such as when you are unable to initiate a Management Console session.

This file is located in `/tmp`.

server.log

The server.log file contains details about agent backup and restore requests, connection status, Management Console interaction, Java exceptions, JBoss start and stop processes, and more. This file contains diverse information and should be checked for almost any problem that is related to the DSM. Sometimes it is easier to `grep` a specific error level, like WARN, INFO, or DEBUG, than it is to view the entire file.

Pruning the DSM Logs

After about 10,000 entries in the **Message Log**, the existing logs are automatically pruned (removed) from the database and written to the backup directory, `/opt/vormetric/coreguard/server/jboss-5.1.0.GA/server/default/backup_logs`. (`/opt/vormetric/coreguard/server/appsvr/backup_logs` is a symbolic link to this directory).

The output file name is `CGSS_LOG_VIEW_UNTIL_YYYY-MM-DD-NN.NN.NN.NNNNNN.csv`. For example, `CGSS_LOG_VIEW_UNTIL_2011-06-06-23.16.22.109000.csv`.

Figure 10-3: A pruning entry in the Message Log

| 32995 | 2011-06-06 16:16:22.883 PDT | I | SSA666 | LOG0141I: The Security Server has pruned and exported 10008 log messages into file "/opt/vormetric/coreguard/server/appsvr/backup_logs/CGSS_LOG_VIEW_UNTIL_2011-06-06-23.16.22.109000.csv". These log messages have been deleted from the Security Server. |

Each output file averages 10, 000 lines and 4.3 MB disk space. Each is owned by `db2fenc1`, with a mod of 644 (`rw-r--r--`). The output file is a comma-separated list comprising the entries in the Logs window and is saved as a `.csv` file.

Up to ten log files can reside in the `backup_logs` directory at one time. The first log file is deleted when the eleventh log file is generated.

Pay attention to this directory. If you are generating a massive amount of log data, as can occur when running a lot of `dataxform` sessions or when GuardPoints are under heavy loads, the log files can come and go quickly. Once gone, there is no record of the activity that had occurred.

The output file column organization is the same as the output of the **Export Logs** button on the **Logs** window.

Exporting Logs

You can export the log entries that are displayed in the **Logs** window to maintain a separate record of server and agent activity at the application level. System Administrators can also export log files that track the internal operations of the DSM at the system level.

The data displayed in the **Logs** window can be exported to a file for archival or analysis. Only the entries in the **Logs** window that are appropriate for the administrator type and domain can be saved to a text file. The output file is formatted as a comma-separated list and is usually viewed in a spreadsheet application.

The following example is an excerpt of a .csv file generated by an All administrator that is inside a domain.

Figure 10-4: Figure 158: Excerpt of a log .csv file

	A	B	C	D	E	F	G	H	I	J	K
4	1426712	1000	FS	3009937978	195375299	2011-01-21-20.40.24.563000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
5	1426713	1000	FS	3009937978	195375299	2011-01-21-20.40.37.173000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
6	1426714	1000	FS	3009937978	195375299	2011-01-21-20.40.49.053000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
7	1426715	1000	S			2011-01-21-20.38.15.336000	I	vmSSA05	DAO0235I	DAO0235I: Administr	-480
8	1426716	1000	S			2011-01-21-20.38.15.340000	I	vmSSA05	DAO0235I	DAO0235I: Administr	-480
9	1426717	1000	FS	3009937978	195375299	2011-01-21-20.40.58.303000	I	solaris120	VMD3781I	VMD3781I: [vmd, 244	-480
10	1426718	1000	FS	3262128009	236323513	2011-01-21-20.40.58.255000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
11	1426719	1000	FS	3262128009	236323513	2011-01-21-20.40.58.749000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
12	1426720	1000	FS	3262128009	236323513	2011-01-21-20.40.58.846000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
13	1426721	1000	FS	3262128009	236323513	2011-01-21-20.41.09.216000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
14	1426722	1000	FS	3262128009	236323513	2011-01-21-20.41.09.217000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
15	1426723	1000	FS	3262128009	236323513	2011-01-21-20.41.09.217000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
16	1426724	1000	S			2011-01-21-20.38.30.337000	I	vmSSA05	COM0314I	COM0314I: The Secu	-480
17	1426725	1000	FS	3009937978	195375299	2011-01-21-20.41.44.104000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
18	1426726	1000	FS	3009937978	195375299	2011-01-21-20.42.00.102000	I	solaris120	DXF4344I	DXF4344I: [vmd, 1307	-480
19	1426727	1000	S			2011-01-21-20.39.38.570000	I	vmSSA05	LOG0140I	LOG0140I: Administr	-480
20	1426728	1000	FS	3009937978	195375299	2011-01-21-20.43.53.294000	I	solaris120	DXF4366I	DXF4366I: [vmd, 1308	-480
21	1426729	1000	FS	3009937978	195375299	2011-01-21-20.43.54.302000	I	solaris120	DXF4371I	DXF4371I: [vmd, 1308	-480
22	1426730	1000	FS	3009937978	195375299	2011-01-21-20.43.54.311000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
23	1426731	1000	FS	3009937978	195375299	2011-01-21-20.43.54.315000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
24	1426732	1000	FS	3009937978	195375299	2011-01-21-20.43.54.319000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
25	1426733	1000	FS	3009937978	195375299	2011-01-21-20.43.54.324000	I	solaris120	DXF4200I	DXF4200I: [vmd, 1308	-480
26	1426734	1000	FS	3009937978	195375299	2011-01-21-20.43.54.327000	I	solaris120	DXF4200I	DXF4200I: [vmd, 1308	-480

The format of this table is subject to change. At this time, the columns indicate:

Table 10-6: Exported Message Log Headings and Description

Column Heading	Description
A	ID number in the Management Console (LOG_ID)
B	Internal domain identifier. If you are not in a domain this is zero. (DOMAIN_ID)
C	Entity that generated the message. For example, S (**DSM**), FS (VTE Agent) (SOURCE)
D	Internal tag (TAG)
E	Internal subtag (SUBTAG)
F	Time of action in UTC (LOG_TIMESTAMP)
G	Severity in the Management Console (SEVERITY)
H	Source in the Management Console (HOST_NAME)
I	DSM or agent Message ID. For example, DAO0239I or CGP2603I.The Message ID also identifies the log service. For example, CGP2603I was generated by the CGP (Policy Evaluation Service) log service. (MESSAGE_ID)
J	Message in the Management Console (MESSAGE)
K	Time offset in minutes. Subtract this number from the time in column F to determine local time. F - K = local time. For example, 22:40:19 UTC - 420 offset = 15:40:19 PDT = 3:40 PM PDT. (TIMEZONE_OFFSET)

Exporting the Message Log

1. Log on to the Management Console as an administrator of the appropriate type for the data you want to export.

2. Enter a domain if you want to export domain-related log entries.

3. Open the **Logs** window.

4. Click **Export Logs**. The **File Download** window opens.

   The options are:

   ○ **Open** to display the log entries to be exported in the default spreadsheet application. Usually this is Excel.

   ○ **Save** to export the log to a file on the system running the Management Console Web session or on another network accessible system. The default output file name is `log.csv`.

   ○ **Cancel** to close the window and stop the export operation.

5. Click **Save**. The **Save As** window opens.

6. Enter the name and path for the export file. The default file name is `log.csv`.

7. Click **Save**. The **Download Complete** window opens. It displays statistical information about the exported log, such as its location and size.

   The options are:

   ○ **Open** to open the exported log file in the default spreadsheet application used to process CSV format files.

   ○ **Open Folder** to open a Windows Explorer window in which to select and view the exported log in an application of your choice.

   ○ **Close** to close the window.

8. Click an option to open the exported log in the default spreadsheet application, open the exported log file in a different application, or to close the window and continue other Management Console operations.

Exporting system logs

The Management Console enables System administrators to export a collection of log files that track the DSM's installation, configuration, and internal operations at the system level.

> **Note**
> If there is a major application or server failure, the Management Console graphic interface can stop working and you will be unable to use this feature to export the system.

Periodically export the server log files, and archive them. Later, the exported files may be useful to Thales Customer Support for diagnosing and resolving system related problems. You may also want to use this as an alternative to the CLI `diag` log view command because here you can download all the server and `cgss` logs at one time in one file, including other files that aren't viewable from the CLI. You can unzip the exported file and view the individual log files in your favorite editor, rather than "`more`" through them in the CLI.

The contents and analysis of these files are not described in this document. Should a major problem occur, analyze these files with Thales Customer Support.

This function exports just a subset of the total log files that are on the system. Included in the export file are log files such as:

- `alters.log`
- `boot.log`
- `cgss.log`
- `cgssdb_start_replication_2009-10-30.log`
- `cgssdb_stop_replication_2009-11-15.log`
- `db2setup.log`
- `delver.log`
- `jboss.log`

- `security_server_install.log`

- `security_server_uninstall.log`

- `security_server_upgrade.log`

- `server.log`

- `server_replication_2009-10-30.log`

- `vor_cert.log`

- `vor_est_trust.log`

More and diverse log files are generated on the server during the course of normal usage and maintenance. System administrators on software-only installations can view the additional log files located in `/tmp` and `/var/log`.

Exporting the DSM system log files

1. Log on to the Management Console as a System/All Administrator.

   Note
   This export system logs feature is not available to Domain and Security Administrators.

   It does not matter if you enter a domain or not. The same log files are exported.

2. Select **Log > Logs**. The Logs window opens.

3. Click **Download Logs**. The **File Download** window opens.

   The options are:

   - **Open** to place the individual log files in a cached archive file without saving the archive file. The files can then be extracted and saved as desired.

   - **Save** to export a diverse collection of internal log files to a single zip file. The file may be saved on the system running the Management Console Web session or on another network accessible system.

   - **Cancel** to close the window and stop the export operation.

4. Click **Save**. The **Save As** window opens.

5. Enter the name and path for exporting the file. The default file name is `logs.zip`.

6. Click **Save**. The **Download Complete** window opens. It displays statistical information about the exported log, such as its location and size.

   The options are:

   - **Open** to open the exported log file in the default archive utility used to process zip format files.

   - **Open Folder** to open a Windows Explorer window in which to select and view the exported log in an application of your choice.

   - **Close** to close the window.

7. Click an option to open the exported log in the default archive application, open the exported log file in a different application, or to close the window.

Adding an Email Notification Group

Email Groups are per domain. You can set up email groups for domains of System, Security, Domain, Domain/Security and All Administrators.

Enabling email notification for log messages

You can automatically send email notifications to a set of administrators if the DSM generates a serious log message.

You need to configure an SMTP server first. Navigate to **System > Email Notification** and click the **SMTP Server** tab. Enter the information for the following tabs:

- **SMTP Server**—SMTP server that will send the email notification. SMTP Servers are per appliance and you must be signed in with System Administrator privileges to modify this setting. If you don't have these privileges, the SMTP server setting is grayed out. Note that the appliance does not come with a default SMTP server and that the SMTP server settings are initially empty.

- **SMTP Server Port**—Port used by the SMTP server.

To bring up the Email Notification interface, select **System > Email Notification** when outside a domain. The attributes and interface information for the Email Notification are as follows:

- **Email Group Name**—Name of the email group which will receive the email notification. Email Groups are per domain. You can set up email groups for domains of System, Security, Domain, Domain/Security and All Administrators.

- **Email Threshold Level**—If the DSM generates a log message with a severity of this specified threshold level, then an email notification is generated. Can be ERROR or FATAL.

- **Email Address List**—Email addresses that will receive this email notification. Separate addresses with commas. If LDAP is configured, you can select addresses from your LDAP address book by pressing **Select**. If it's not configured, you can enter your login and password to access it.

- **Email Subject**—Text you want on the subject line.

- **Message Contains**—This is a string filter that works with the Email Threshold Level. Only messages containing this string will be sent as an email notification. If blank, then all messages meeting the threshold criteria will be sent.

- **Enabled**—A checkbox that enables or disables email notification to the group.

To add an email notification group:

1. Select **System > Email Notification**. The **Email Notification** window displays.

2. Under the **Email Notification List** tab, click **Add**. The **Add Email Notification Group** window displays.

3. Enter the information and click **Ok**.

Changing the SMTP server and port for email notification

You must be signed in with System Administrator privileges to modify this setting.

To change the SMTP server and port for email notification:

1. Select **System > Email Notification**. The **Email Notification** window displays.

2. Click the **SMTP Server** tab

3. Enter the SMTP server and server port and click **Ok**.

Chapter 11: Configuring SNMP

Simple Network Management Protocol (SNMP) is a full-featured protocol that is used to manage and monitor network nodes like hosts, routers, and appliances. The specific attributes of network nodes that can be managed and monitored by SNMP are configured as objects in a Management Information Base (MIB). The DSM can be enabled as an SNMP agent and then monitored by SNMP servers using the set of MIB objects described below.

Overview

The DSM supports SNMP version 1 or 2. SNMP is not used to manage DSMs. A small set of MIB objects are provided with which to query DSM configuration and status information. The HA nodes run in parallel, so each HA node contains the same SNMP configuration. Therefore, SNMP servers that can query one HA node can also query every HA node with the same community string.

When the DSM receives an SNMP GET request from an SNMP server, the DSM locates the Object IDentifier (OID) entry in the MIB and returns its value to the SNMP server.

If SNMP is enabled on an HA Cluster, the HA node 1 polls each HA node using an SNMP GET request at five-minute intervals. The response time for each HA node is displayed in the High Availability Servers window in milliseconds. If SNMP is disabled, the Response Time column will display "SNMP Disabled". If an HA node is not reachable, the Response Time column will display "Not Reachable".

SNMP traps are not supported at this time and cannot be configured on the DSM.

Enabling SNMP on the DSM

SNMP is enabled via the **System > SNMP** page on the Configuration tab. You can define the SNMP community string with which to query the DSM.

If the SNMP Access Control List (ACL) is empty, SNMP requests from any IP address will be acknowledged. If the SNMP ACL is defined to allow only certain IP addresses (for example, 10.1.2.3) or IP address blocks (for example, 10.1.2.*) to go through, the DSM will only acknowledge requests from IP addresses specified in the SNMP ACL. The community string and IP address are the only credentials used to verify the legitimacy of the SNMP request. The community string is typically set to a factory default, value of "public". This string must be the same for all devices in the same group for SNMP monitoring to function. For security reasons, the Network Administrator should change the community string from "public" to a custom value.

> **Note**
> Thales recommends that you do not enable SNMP on the DSM unless it is required, as this could pose a security risk. If you do enable SNMP on the DSM, we recommend that you use an SNMP ACL to restrict access to this service, and change the default community string from 'public' to a custom value.

The nodes in an HA cluster share the same SNMP configuration as the all other HA nodes. Enable SNMP listening on one node, and SNMP listening is enabled on all of the HA cluster nodes. The community string that you enter is applied to the all of the nodes in the HA cluster. This means that an SNMP server can query all of the nodes in the HA cluster.

> **Note**
> If a node in an HA configuration does not respond to SNMP requests, restart the node to resolve the issue.

GET requests can be sent to port 161 or port 7025.

Figure 11-1: SNMP
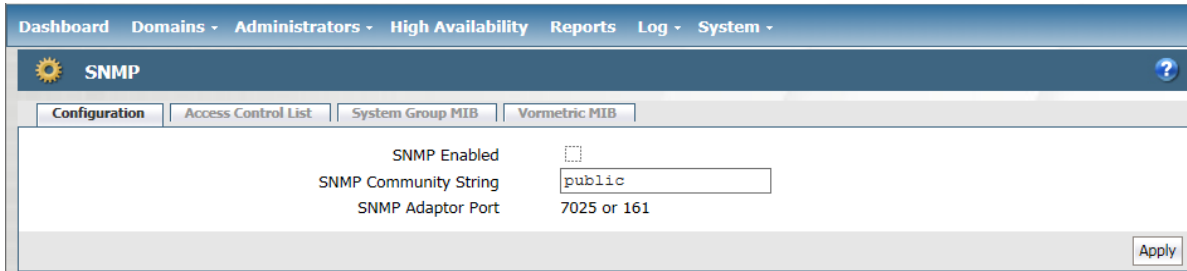
To enable the DSM to listen for SNMP queries and to configure the SNMP community string:

1. Log on as a System/All Administrator.

2. Do not enter a domain.

3. Select **System > SNMP**.

4. The SNMP window opens to the **Configuration** tab.

5. Check **SNMP Enabled** to make the DSM listen for SNMP queries.

6. Enter the community string, or password, with which all SNMP servers will query the DSM in the **SNMP Community String** field.

7. Click **Apply**.

> **Note**
> Once SNMP is enabled, the DSM will respond to requests from any SNMP server unless a preferred SNMP server is specified in the Access Control List. Once the IP address of a SNMP Server is specified in the Access Control List, the DSM will only respond to that SNMP Server.

Adding SNMP Servers

Configure the SNMP servers that are allowed to query the DSM in the SNMP window, Access Control List tab.

SNMP servers can access the DSM using TCP or UDP.

Figure 11-2: SNMP Servers, Access Control List
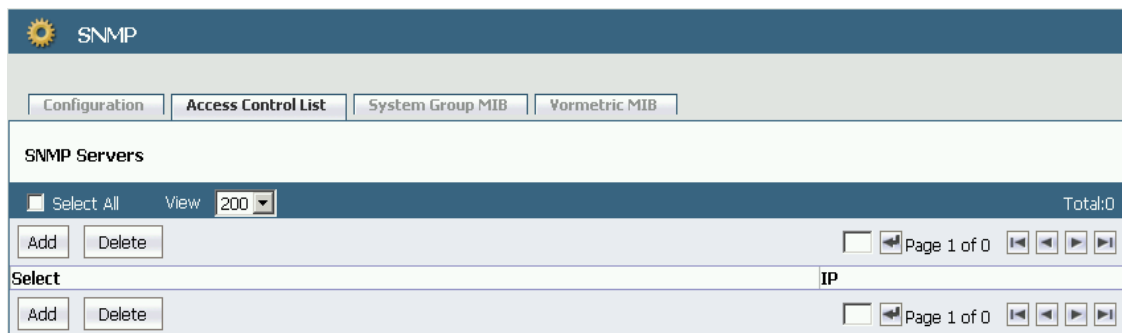
To add a system to the list of SNMP servers that may submit SNMP queries to a DSM:

1. Click **Add**, the Add SNMP Server window opens.

2. Enter the IP address of the SMNP server to be granted access in the **IP Address** field.

   Host names and Fully Qualified Domain Names (FQDN) are not supported at this time.

3. Click **Ok**.

> **Note**
> The IP Address field currently supports the use of a "wild-card" in the 4th octet. For example: 10.1.2.*

Once an SNMP server has been added to the list of allowed servers, a corresponding log entry is created indicating an SNMP server has been added to the ACL.

Figure 11-3: Log entry indicating an SNMP server has been added to the ACL

There is no record of a failed status query in the Logs window; however, a record is entered in the server.log file. For example:

```
2011-09-23 17:41:13,267 ERROR [STDERR] Sep 23, 2011 5:41:13 PM
org.snmp4j.log.JavaLogAdapter log
WARNING: 10.3.244.200 not in ACL
```

The log entry indicates that an SNMP query was attempted from a system that is not configured in the Access Control List (ACL). Such a query is ignored by the DSM and, after the timeout interval has elapsed, the SNMP query is terminated and timeout message is returned. For example:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.2.1.1.4.0
    Timeout: No Response from 10.3.48.1:7025
    #
```
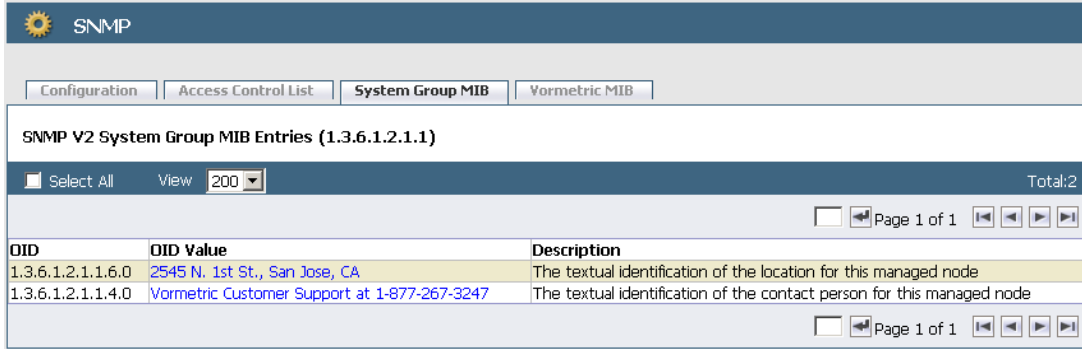
The following example from the server.log file indicates that an SNMP query had been submitted from a configured system. It indicates only that the system submitting the query is configured. It is no indication of the success or failure of the SNMP query itself; only that the SNMP server is allowed to query the DSM.

```
2011-09-23 17:41:49,964 ERROR [STDERR] Sep 23, 2011 5:41:49 PM
org.snmp4j.log.JavaLogAdapter log
WARNING: 10.3.244.200 passed ACL
```

Changing OID Values

The SNMP Object IDentifier (OID) values that can be changed are sysContact (1.3.6.1.2.1.1.4.0) and sysLocation (1.3.6.1.2.1.1.6.0). Customize the OID values so that the information collected by the SNMP server can include the contact for DSM questions and issues, plus the physical location of the DSM. These OIDs are part of the 1.3.6.1.2.1.1 MIB group defined in RFC 1213.

Figure 11-4: Customized contact and location information

To configure the DSM contact and location information:

1. Open the System Group MIB tab.

2. Click a string in the **OID Value** column.

3. The **Edit OID Value** window opens.

Figure 11-5: Editing the OID value

4. Select and delete the text string in the **OID Value** field.

5. Enter a new string in the **OID Value** field.

6. Click **Ok**.

The text in the **Description** column is hard-coded and cannot be changed.

A log entry indicating the OID number and value change is entered in the Logs window.

Displaying Vormetric-specific SNMP Information

The Vormetric MIB tab displays the Vormetric-specific OIDs that can be queried by an SNMP server. The OIDs cannot be manually changed. The OID values are dynamic and change based upon the DSM state and configuration.

Figure 11-6: Vormetric-specific OIDs
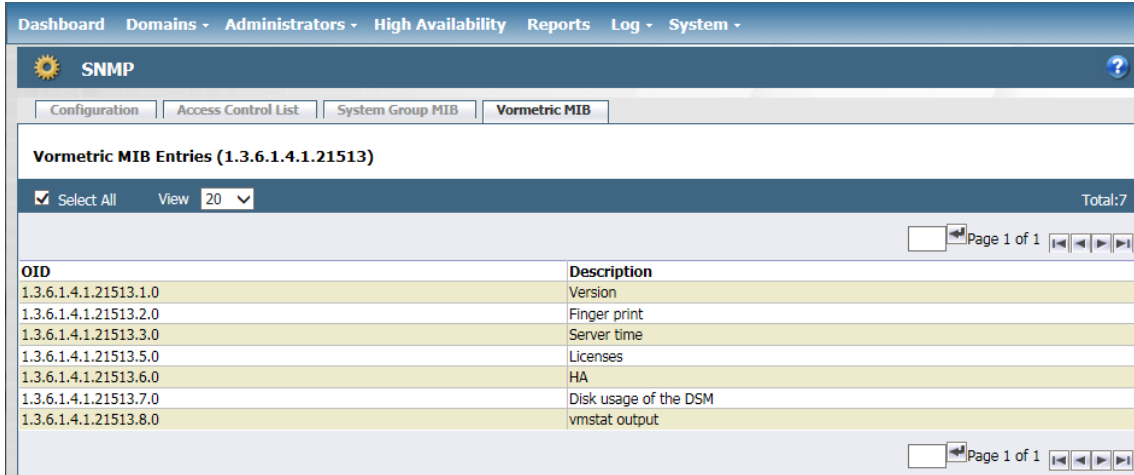
The OIDs in the Vormetric group MIB begin with 1.3.6.1.4.1.21513. The following table lists the Vormetric OIDs and their purpose.

Table 11-1: OID Descriptions

OID	SNMP Object Type	Description
1.3.6.1.2.1.1.4.0	sysContact	The textual identification of the contact person for this managed node, together with information on how to contact this person. If no contact information is known, the value is a zero-length string. Max. length 256 characters.
1.3.6.1.2.1.1.6.0	sysLocation	The physical location of this node (e.g., 'telephone closet, 3rd floor'). If the location is unknown, the value is a zero-length string. Max. length 256 characters.
1.3.6.1.4.1.21513.2.0		Returns the fingerprint of the current DSM deployment. The fingerprint is also displayed in the Management Console Dashboard window.
1.3.6.1.4.1.21513.3.0		Returns the time and date at the time of the SNMP query.
1.3.6.1.4.1.21513.5.0		Returns the agent type (FS, or Key agent), the license installation state (true or false) of each agent type, and, for each installed license, the license expiration date. This information is also displayed in the Management Console License window.
1.3.6.1.4.1.21513.6.0		Returns the name of each node in a DSMHA cluster configuration.
1.3.6.1.4.1.21513.7.0		Returns disk usage information for each file system mounted on the DSM. This is the equivalent of running df -hk -B 1024K on the DSM command line.
1.3.6.1.4.1.21513.8.0		Return DSMs process, memory, paging, I/O, and CPU usage information. This is the equivalent of running vmstat on the DSM command line.

Example SNMP Queries

The following SNMP queries were made on Red Hat Enterprise Linux Server, release 6.0, using SNMPv2.

To display DSM contact information:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.2.1.1.4.0
```

```
SNMPv2-MIB::sysContact.0 = STRING: Vormetric Customer Support at 1-877- 267-3247
```

To display the physical location of the DSM:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.2.1.1.6.0
```

```
SNMPv2-MIB::sysLocation.0 = STRING: 2545 N. 1st St., San Jose, CA
```

To display the DSM version number:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.1.0
```

```
SNMPv2-SMI::enterprises.21513.1.0 = STRING: "5.3.0.1616"
```

To display the DSM fingerprint:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.2.0
```

```
SNMPv2-SMI::enterprises.21513.2.0 = STRING:
"D2:48:EF:E4:A2:B0:59:8C:5F:DB:9D:3B:30:41:0B:EE:BD:07:8D:67"
```

To display the current date and time on the DSM:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.3.0
```

```
SNMPv2-SMI::enterprises.21513.3.0 = STRING: "2015-08-18 20:56:53.135 PDT"
```

To display the DSM license configuration:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.5.0
```

```
SNMPv2-SMI::enterprises.21513.5.0 = STRING: "FS max # of agents: 30000; Expires: Dec-31-2015;
Key max # of agents: 30000; Expires: Dec-31-2015; FS max # of agents: 30000; ; Key max # of
agents: 30000; FS max # of agents: 30000; Max hours: 1000000; Key max # of agents: 30000; Max
hours: 1000000;Key vault enabled: true; Multi-domain enabled: true; max # of domains: 20000;
Issued to: DSM522-Performance-2015-12-31"
```

To display the DSM HA configuration:

```
# snmpget -c public -v 2c 10.3.48.239:7025 1.3.6.1.4.1.21513.6.0
```

```
SNMPv2-SMI::enterprises.21513.6.0 = STRING: " sys15123.com; : sys48239.com; "
```

To display the mounted file systems and their disk usage:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.7.0
```

```
SNMPv2-SMI::enterprises.21513.7.0 = STRING: "
Filesystem            1M-blocks  Used Available Use% Mounted on
/dev/mapper/vg_sys48001-lv_root
                      50269 3006    44703    7% /
tmpfs                  1917    1     1917    1% /dev/shm
/dev/sda1               477   38      414    9% /boot
/dev/mapper/vg_sys48001-lv_home
                      45867 15185    28346   35% /home
```

To display DSM system usage information:

```
# snmpget -c public -v 2c 10.3.48.1:7025 1.3.6.1.4.1.21513.8.0
```

```
SNMPv2-SMI::enterprises.21513.8.0 = STRING: "
```

```
procs ----------memory---------- ---swap-- -----io---- --system-- -----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa st
 0  0  51040 130248 228572 1777640   0    0     1    12   11    4  0  0 100  0  0
"
```

Chapter 12: Backing Up and Restoring

A backup of the DSM is a snapshot of the configuration at a point in time. When a backup is restored, the Management Console displays the same information captured at the time the backup was originally made, any changes made after the last backup will not be restored.

Overview

You can use a backup to restore the hosts, encryption keys, policies, as well as other configuration information of a DSM in the event of a software crash recovery or system changes. A System/All Administrator creates a system-level backup, and a Domain, Domain and Security, or All Administrator creates a domain-level backup via the Management Console.

> **Note**
> Administrators must be logged into the domain that is to be backed up or restored to perform these operations.
> An All administrator can perform a domain backup and restore operation, as long as that administrator is added to the domain.

System-level configuration such as network and timezone settings are **not** backed up—those remain unchanged after a restore operation.

Each backup is encrypted with a wrapper key. A wrapper key must be created before the DSM can be backed up. The same wrapper key is also required to restore the backup.

DSM backups can be restored at the system-level or at the domain-level.

- A system-level backup can only be restored to the same DSM or another DSM.

- A domain-level backup can only be restored to a domain—the same domain, or another domain on the same DSM, or a domain on another DSM.

A DSM backup can also be restored to another DSM with a higher security level i.e., a backup of a DSM with no HSM can be restored to a DSM with an HSM, but not the other way around.

The following restore scenarios are allowed:

Table 12-1: Backup and Restore compatibility matrix

Backup from	Restore to		
	DSM w/o HSM	**DSM with nCipher Solo or Connect HSM**	**DSM with Thales Luna HSM**
DSM w/o HSM	✔	✔	✔

Table 12-1: Backup and Restore compatibility matrix (continued)

	Restore to		
Backup from	**DSM w/o HSM**	**DSM with nCipher Solo or Connect HSM**	**DSM with Thales Luna HSM**
DSM with nCipher Solo	✘	✔ User must restore to DSM v6.1.0.9229 and then upgrade to 6.4.0	✔
DSM with Thales Luna HSM	✘	✘	✔

Backing Up the DSM Configuration

When a backup is restored, the Management Console displays the same information captured at the time the backup was originally made.

You can create a backup of the DSM configuration at the system level or at the domain level. To create a backup of a domain, you must be logged into that domain.

Differences between System-level and Domain-level Backups

Table 12-2: System-level vs domain-level backups

System-level backup	Domain-level backup
System/All Administrators create the backup.	Domain, Domain and Security, or All Administrators create the backup.
Backs up the configuration information for the complete DSM including; web server certificate, certificates, system preferences, log preferences, users, domains, hosts, encryption keys, signatures, policies, GuardPoints, and license information including all the configuration information in all the domains.	Backs up domain specific information including; web server certificate, certificates, system preferences, log preferences, domains, hosts, encryption keys, signatures, policies, GuardPoints, and license information.
DSM users can be backed up.	Domain level users cannot be backed up, they will need to be recreated or added back to the domain after a restore operation.
GuardPoints and host-sharing information are backed up.	GuardPoints and host-sharing information are not backed up. Host sharing will have to be re-established and GuardPoints recreated after the restore operation.

The procedures to create a wrapper key, create a backup, and restore a backup are the same at the domain level and at the system level.

Backup Encryption Wrapper Key

DSM backup files are encrypted with a wrapper key to keep them secure. This wrapper key must be created, or imported from a previous create operation, before creating a backup. The same wrapper key used to encrypt a backup is also required to restore that DSM backup.

For additional security, wrapper keys can be broken up into key shares—pieces of a wrapper key. These key shares can are then be divided amongst two or more custodians, such that each custodian must contribute their key share in order to assemble a complete wrapper key. This is also referred to as split key knowledge or M of N configuration.

For example, you can break up the wrapper key amongst a total of up to 10 custodians. The minimum value for required custodians is two. When the wrapper key is needed, at least two of the custodians must contribute their key share in order to assemble a complete wrapper key.

To backup a system-level configuration, the wrapper key must be created at the system-level by a System/All Administrator. To create a backup at the domain-level, a wrapper key must be created from within the domain to be backed up by a Domain or Security/All Administrator.

Create a wrapper key

1. Log on to the Management Console as a System/All administrator.

   Or if you are creating a wrapper key at the domain level

   Log on or switch to a domain on the Management Console as a Domain, Domain and Security, or All administrator.

2. Select **System > Wrapper Keys** from the menu bar.

3. In the **Wrapper Keys** window, select **Operation > Create**, then click **Apply** to create the wrapper key.

   You will see a confirmation message stating that the key exists.

4. Select **System > Backup and Restore > Manual Backup and Restore** from the menu bar. A confirmation message is also displayed on this tab, stating that the wrapper key exists. You can now proceed with creating a backup.

5. Return to the **System > Wrapper Keys** menu option and select **Export** from the **Operation** menu to export key shares.

6. Set a number for both the **Minimum Custodians Needed** and the **Total Number of Custodians.** This setting splits the wrapper key value among multiple custodians.

   ○ Minimum value required for Minimum Custodians Needed: 2

   > **Note**
   > If you want the minimum value for Minimum Custodians Needed to be set to 1, you can only set this in the **System > General Preferences > System > Wrapper Keys Export** section.

   ○ Maximum value required for Minimum Custodians Needed: 10
   ○ Minimum value required for Total Number of Custodians: 3
   ○ Maximum value required for Total Number of Custodians: 10

7. Select the check box next to the System/All Administrators who will serve as custodians for the wrapper key shares. Any of these administrators, with the exception of the default initial log-on administrator admin, can be selected as a custodian.

   If more than one custodian has been selected, each of them is given a share of the wrapper key. The wrapper key share is displayed on their Dashboard window when they log into the Management Console, see Figure 12-1: "Management Console Dashboard showing the wrapper key share toggle" on the next page. Each administrator must see a unique wrapper key share displayed on the dashboard beneath the fingerprint for the CA.

8. Click **Apply** on the bottom right hand corner.

   The generated wrapper key or key shares are exported and is visible on the Dashboard, beneath the fingerprint for the CA. The **Wrapper Key Share** displayed in the Dashboard window is a toggle. Click **Show** to display the wrapper key share value. Click **Wrapper Key Share** value to display the string **Show**.

9. Ask each administrator to securely store a copy of this key share. They must provide this as part of their role in a DSM restore operation.

Figure 12-1: Management Console Dashboard showing the wrapper key share toggle
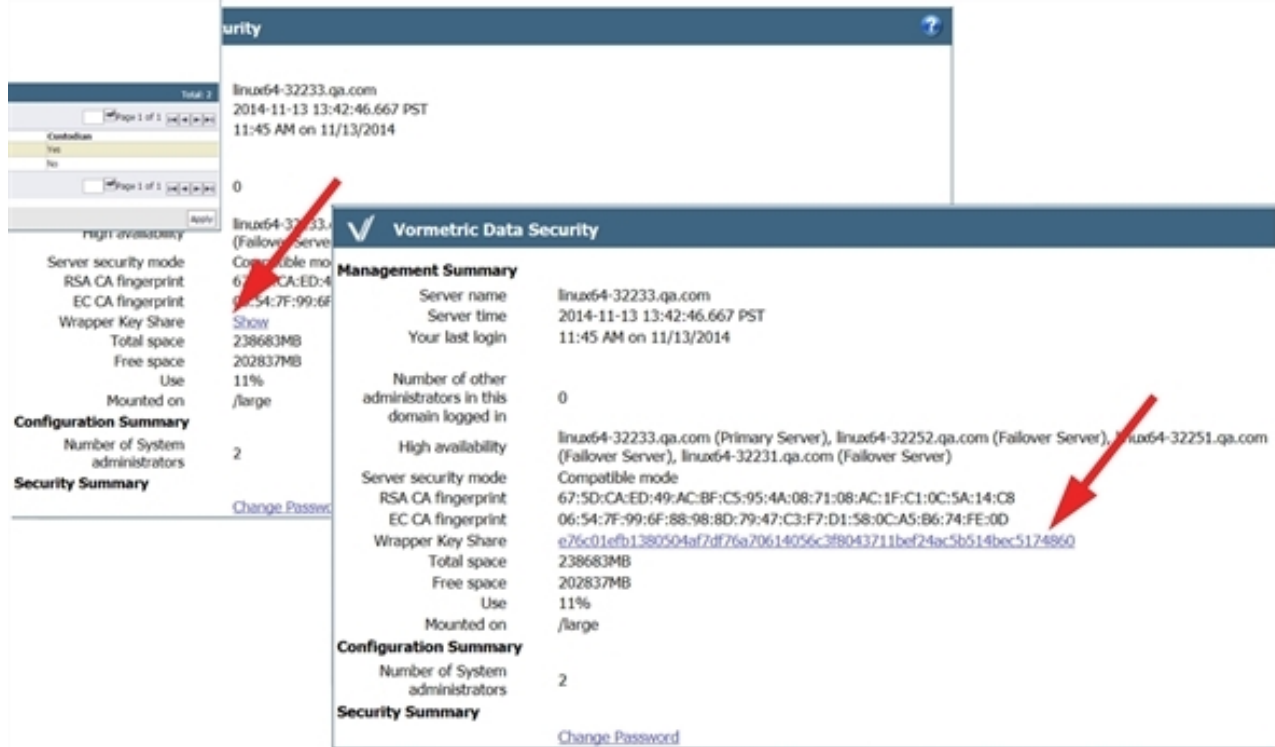
A backup of the DSM can be created after the wrapper key has been created. The procedure to create a backup at the system level or at the domain level is the same.

System-level Backup

1. Log on to the Management Console as a System/All administrator.

2. Select the **System > Backup and Restore** menu option. The **Manual Backup and Restore** page opens.

3. Click the **Backup** tab and then select **Ok**.

Figure 12-2: Manual Backup and Restore dialog with File Download dialog displayed
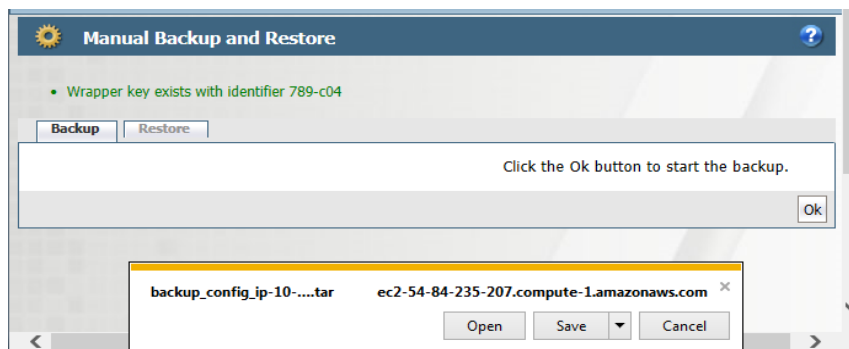
4. Click **Save**. Save the file to a secure location that you are sure will still be accessible if the server fails. By default, the file name will be in the format: `backup_config_<`*dsm server name*`>_yyyy_mm_dd_hhmm.tar`

Where *<dsm server name>* is the FQDN of the DSM that is being backed up.

5. Save the backup to a secure location. Access to the backup should be limited to only a few employees and should be audited.

Per Domain Backup and Restore

In addition to a creating a backup of the DSM, you can also back up and restore the configuration information for a single domain. A domain backup can be restored to:

- The same domain

- To a different domain on the same DSM

  If a domain backup is restored to a different domain on the same DSM, there may be a host name conflict, in which case the host names must be changed.

- To a different domain on another DSM

> ⚠️ **WARNING**
> **A backup containing a KMIP domain must be restored with the same name. You cannot restore it with a different name, unlike a standard domain.**

To create a backup of a domain and to restore that backup, a wrapper key must be created for the domain, and the domain must have an assigned Domain Administrator. The backup and restore operations are done by a Domain, Domain and Security Administrator, or an All administrator from within the domain to be backed up or restored.

> **Note**
> If you have existing domains on your DSM, created with a non-KMIP license, and then you upload a KMIP-enabled license, any domain created after uploading this license is automatically KMIP-enabled.

Domain-level Backup

1. Log on to the Management Console as a Domain, Domain and Security or All Administrator, or switch to the domain that you want to backup.

2. Select the **System > Backup and Restore** menu option. The **Manual Backup and Restore** page opens.

3. Click the **Backup** tab and click **Ok** to start the backup.

4. Click **Save** in the **File Download** dialog box. Save the file to a secure location that you are sure will still be accessible if the server fails. By default, the file name will be in the format: `backup_config_<domain name>_` *<dsm server name>*`_yyyy_mm_dd_hhmm.tar`

   Where `<domain_name>` is the name of the domain being backed up and *<dsm server name>* is the FQDN of the DSM that is being backed up.

5. Save the backup to a secure location. Access to the backup should be limited to only a few employees and should be audited.

Restoring a DSM Backup

A backup of the DSM can be used to restore the hosts, encryption keys, and policies, as well as other configuration information of a DSM, after a software crash recovery or system change. A DSM backup can be restored at the system level or at the domain level.

The procedure to restore a domain-level backup is the same as the procedure to restore a system-level backup. To restore a domain level backup, you must be logged into that domain.

The DSM backup is restored via the Management Console.

If the DSM backup is to be restored in either of the following cases:

- On the same DSM appliance (virtual, V6000, or V6100) on which it was created

- To another appliance, if the current DSMs are lost (virtual appliance and V6000) the procedure is the same, and is described in "System-level restore" on the facing page.

In the event of a disaster involving the V6100 DSM hardware appliance where all the DSMs are lost, and the backup image must be restored to a new DSM appliance (a V6100 backup has to be restored to a V6100 appliance), the procedure is different. The V6100 DSM hardware appliance contains a built-in Hardware Security Module (HSM). This appliance requires the use of a set of smart cards, referred to as the Administrative Card Set (ACS) that contain the keys that let you perform administrative tasks on the DSM. Each ACS is configured to manage a specific group of DSMs and is used to manage system availability and recovery for that group. A group of DSMs managed by a particular ACS is called a "security world". All DSMs in a security world require the same ACS to perform administrative tasks. You can define a quorum of cards from the ACS that are required to perform administrative tasks such as restoring a backup image on a new V6100 DSM hardware appliance. The procedure to restore a V6100 backup image to a new hardware appliance in the section "System-level restore of a V6100 configuration to a new security world" on page 101. For information about how to configure the V6100, refer to the *DSM Installation and Configuration Guide* for procedures describing how to do this.

> ⚠️ **WARNING**
> **A backup containing a KMIP domain must be restored with the same name. You cannot restore it with a different name, unlike a standard domain.**

Disaster Recovery

To restore a V6100 DSM appliance from a backup image, you need to do the following:

- **Create a new security world**. The new security world will be used only until the previous one is restored. Once the previous security world is restored, the new one is overwritten and the ACS released. This ACS is then available for reuse.

- **Upload the backup file**. Ensure that the Management Console is up and running on the new DSM and upload the backup file via the Management Console to the new DSM.

- **Restore the DSM**. Restore the backup file using the ACS of the lost security world, thereby restoring the lost security world. This process releases the security world and the associated ACS that were created for the restoration procedure. This procedure is described in "Restoring the DSM from a backup" below.

Restoring the DSM from a backup

The following procedures describe:

- How to do a system-level restore of a DSM from a backup

- How to do a domain-level restore of a DSM from a backup

- How to restore a DSM V6100 hardware appliance backup image from another security world.

> **Notes**
> - Following a restore operation, the DSM configuration in the Management Console is replaced by the configuration stored in the backup copy. Any new encryption keys, policies, hosts, or GuardPoints added since the last backup will be overwritten and lost.
> - Unless this is a disaster recovery scenario where all DSMs have been lost, always backup the current configuration before running a restore operation.

System-level restore

1. Locate the backup that is to be restored

2. Log on to the Management Console as a System Administrator/All administrator.

   > **Note**
   > If you already have the proper Wrapper Key imported, skip to Step 8.

3. Import wrapper keys. Select **System > Wrapper Keys** from the menu bar.

4. Select **Import** from the Operation pull-down menu.

5. Click the **Add** button.

6. If key shares have created from the wrapper key, paste a Key Share value from one previously stored with a custodian into the **Key Share** text field and click **Ok**.

   Repeat steps 5 and 6 for each administrator selected as a key custodian if you have chosen to have more than one custodian for the wrapper key. A key share must be imported for at least as many as were specified by the **Minimum Number of Custodians** value when the wrapper key was exported.

7. Click **Apply** to finish importing the wrapper key.

8. Restore the backup file. Select **System > Backup and Restore** from the menu bar.

9. Select the **Restore** tab.

10. Click **Browse**. Locate and select the backup file to restore.

11. If this is a disaster recovery, enable the **Include User(s)** check box.

12. Click the **Ok** button. The restored file uploads and the DSM disconnects from the Management Console.

13. Log back on to the Management Console as an Security or All administrator. Verify that the configuration is restored correctly.

Domain-level restore

When restoring a domain-level backup, all host sharing and GuardPoints on shared hosts are removed and users are not restored.

1. Locate the backup that is to be restored

2. Log on to the Management Console as a Domain, Domain and Security, or All administrator.

   > **Note**
   > If you already have the proper Wrapper Key imported, skip to Step 8

3. Import wrapper keys. Select **System > Wrapper Keys** from the menu bar.

4. Select **Import** from the Operation pull-down menu.

5. Click the **Add** button.

6. If key shares have created from the wrapper key, paste a Key Share value from one previously stored with a custodian into the **Key Share** text field and click **Ok**.

   Repeat steps 5 and 6 for each administrator selected as a key custodian if you have chosen to have more than one custodian for the wrapper key. A key share must be imported for at least as many as were specified by the **Minimum Number of Custodians** value when the wrapper key was exported.

7. Click **Apply** to finish importing the wrapper key.

8. Restore the backup file. Select **System > Backup and Restore** from the menu bar.

9. Select the **Restore** tab.

10. Click **Browse**. Locate and select the backup file to restore.

> **Note**
> In the case of a domain-level restore, you will not be able to restore users and this option will not be available.

11. Click **Ok**.

Once the restore operation is complete, verify that the configuration is restored correctly.

> ⚠ **WARNING**
> **Following a restore operation, the DSM configuration in the Management Console is replaced by the configuration stored in the backup copy. Any new encryption keys, policies, hosts, or GuardPoints added since the date/time of the backup file being used for the restore operation, will be overwritten and lost. If there is a reason to do a selective restore from backup, then the following procedure is recommended:**
> **1. Export the keys created since the date/time of the backup file being used for restore operation. Refer to the section on exporting/importing keys in the chapter on "Configuring Keys and Key Groups".**
> **2. Restore from the backup file (note that this operation will replace the current DSM configuration).**
> **3. Import the keys created in step 1.**

System-level restore of a V6100 configuration to a new security world

1. Locate the backup that is to be restored

2. Log on to the Management Console as a System or All administrator.

3. Import wrapper keys. Select **System > Wrapper Keys** from the menu bar.

> **Note**
> If you already have the proper Wrapper Key imported, skip to Step 8.

4. Select **Import** from the **Operation** menu.

5. Click **Add**.

6. Paste a Key Share value from one previously stored with a custodian into the **Key Share** text field and click **Ok**.

7. Repeat steps 5 and 6 for each administrator selected as a key custodian if you have chosen to have more than one custodian for the wrapper key. A key share must be imported for at least as many as were specified by the **Minimum Number of Custodians** value when the wrapper key was exported.

8. Click **Apply** to finish importing the wrapper key.

9. Restore the Backup File. Select **System > Backup and Restore** from the menu bar.

10. Select the **Restore** tab.

11. Click **Choose File**. Locate and select the backup file to restore.

12. If this is a disaster recovery, enable the **Include User(s)** check box.

13. Click **Ok**. The following message displays:

"The configuration file has been successfully uploaded to the server. Since the backup is from a different Security World, you must run

```
maintenance$ config restore <filename>_yyyy_mm_dd_hhmm.tar [withuser | withoutuser]"
```

from the DSM CLI to continue the restore process."

14. Open a DSM CLI window and type the command from the message on the Management Console.

```
maintenance$ config restore <filename>_yyyy_mm_dd_hhmm.tar [withuser | withoutuser]"
```

This command restores the specified backup file. If you enabled the **Include User(s)** check box you need add withuser to the command, or else type the command and add withoutuser.

The restore procedure begins and a series of messages informing you of the status of the restore procedure are displayed.

```
Security Server Configuration Restore

Version 5.3.1.0 Build uni_1716v

Using /var/tmp as restore staging directory.

The Security Server needs to shutdown before the restore can proceed. All the current
data will be over-written with the backup data.

Restoring configuration backup from file /var/tmp/Backup_yyyy_mm_dd_hhmm.tar

Backup is from build uni_1716v

Stopping Security Server

Stopping data store                                          [ OK ]
```

15. You are prompted to have a quorum of cards from the ACS of the Security World to which the backup file belongs, ready to be read by the card reader.

```
HSM detected.

Restoring from a backup config file with Thales Security World data.

Restoring config from another Security World.

Have a quorum of cards from the ACS of the Security World which the given backup config
file belongs to ready.
```

16. Switch the DSM into **Initialization** mode by moving the **Mode** switch on the back panel to the Initialization (I) position. If you have not done this already, you will be prompted to do so.

```
The hardware security module must be in pre-initialization mode in order to be
reprogrammed.

Move the Mode switch on the back panel to the initialization (I) position.

Once that's done, press Enter to continue...

17:31:03 WARNING: Module #1: preemptively erasing module to see its slots!
```

17. You will be prompted to insert the card(s) into the reader. If the card(s) have been configured with a passphrase, you will be prompted to enter it.

```
Indoctrinating Module

 Module 1: 0 cards of 1 read
```

```
 Module 1 slot 0: empty
Card reading complete.
```

> **Note**
>
> If you have more than one card in the quorum, you will be prompted to enter the cards until the quorum of cards has been read.

18. A hash is generated for the security world.

    ```
    security world loaded on 1 module; hknso = f7387fed7f52625bc06b79607bb4b0afdd93a6b1
    ```

19. You will be prompted to move the mode switch on the back panel of the DSM to the operational position (O) and press enter

    ```
    Put the hardware security module into operational mode by moving the Mode switch on the
    back panel to the operational (O) position.
    Once that's done, press Enter to continue...
    ```

20. A series of messages describing the post-restore operations are displayed.

    ```
    Cleaning up staging area...
    Starting Security Server...
    Starting data store
    Starting Security Server
    Restore config complete
    SUCCESS: restore configuration.
    ```

21. The DSM configuration backup has been successfully restored.

22. Log back on to the Management Console as a System or All administrator. Verify that the configuration is restored correctly.

Backing up and Restoring an nShield Connect HSM-enabled DSM

Backup an nShield Connect HSM-enabled DSM

A network HSM-enabled DSM is backed up the same way as any other DSM appliance.

You should also take a backup of the nShield Connect RFS when you backup the DSM and keep the two backup files together. The RFS backup is done separately as part of the nShield Connect administration. Refer to the nShield user documentation for details and procedures.

> **Note**
>
> If you are restoring the backup to the same DSM appliance, then the nShield RFS backup is not required.

If you are restoring the backup to another network HSM-enabled DSM in another Security World, you will need to restore the RFS first and then the DSM backup. The associated ACS is also required.

The high-level steps to backup a network HSM-enabled DSM and restore the backup to another DSM in another Security World are as follows:

1. Backup the network HSM-enabled DSM. The backup procedure for a network HSM-enabled DSM or cluster is the same as for any other DSM deployment.

2. Backup the nShield Connect RFS. Refer to the nShield user documentation for the procedure.

3. Restore the Security World data on the Connect device connected to the new network HSM-enabled DSM that belongs to another Security World.

4. Run the `connect secworldupdate` command on the DSM to update the Security World on that DSM.

5. Restore the network HSM-enabled DSM backup. The restore procedure for a network HSM-enabled DSM or cluster is the same as for any other DSM deployment.

Restore a network HSM-enabled DSM

A backup of an nShield Connect network HSM-enabled DSM can be restored as follows:

- On the same DSM appliance

- On another nShield Connect DSM appliance in the same Security World

- On a DSM V6100 appliance: Restoring a network HSM-enabled DSM backup to a V6100 appliance requires the ACS of the nShield Connect device that was configured with the network HSM-enabled DSM.

Additionally, you can restore the following types of backups to an nShield Connect network HSM-enabled DSM:

- A backup of a non-HSM DSM (V6000 or virtual DSM)

Domain level backups can be restored as follows:

- A domain backup from an nShield Connect network HSM-enabled DSM to a domain on a non- network HSM-enabled DSM (V6000 or virtual appliance) and vice versa

- A domain backup from a network HSM-enabled DSM to a domain on a V6100 appliance and vice versa

Backing up when connected to a Luna-enabled DSM

Backing up a Luna-enabled DSM

Backing up a network Luna-enabled DSM works the same way as backing up to any other DSM appliance.

Restoring to a Luna HSM-enabled DSM

Restoring to a network HSM-enabled DSM works the same way as restoring any other DSM appliance with some limitations.

Prerequisite

Before restoring a configuration, you must remove the HA nodes from the HA cluster and run HA cleanup on each node.

> **Note**
> You do not have to remove the Luna from the DSM.

Restoring a configuration

- After restoring your backup from a Luna-configured DSM to a standard DSM, you **MUST** go to the CLI and run the `HA > Luna add` command. When asked for the Luna partition number, you **MUST** restore to one of the original partitions.

- You can restore a backup from a DSM to a Luna-configured DSM.

Limitations for restoring are as follows:

- You **cannot** restore a backup from an nShield Connect-configured DSM to a Luna-configured DSM.

- You **cannot** restore a backup from an nCipher-configured DSM to a Luna-configured DSM.

Automatic Backup

The DSM system configuration information can be scheduled to be automatically backed up on a daily or weekly basis using the Automatic Backup feature.

Automatic backups can also be configured at the domain level. To schedule an automatic backup at the domain level, you must be logged into the domain for which the backup is to be scheduled.

In addition to scheduling a backup, there is also an option to run a scheduled backup immediately and push the backup file to a configured external file server. To do this, you must access a File Server (a UNIX or Windows host) that is network accessible by the DSM to store the backup files.

The procedure to schedule an automatic backup is the same at the system level and at the domain level.

> **Note**
> Thales recommends using automatic backup for the Luna because the backup includes the metadata text file that lists the partition ID and Luna hostnames.

Schedule an Automatic Backup

1. Select **System > Backup and Restore > Automatic Backup** in the Management Console, to open the **Automatic Backup** page.

2. Enter the settings for the **Automatic Backup Schedule** and the **External File Server** where the backup files will be stored.

   Enter the following information in the **Automatic Backup Schedule** section:

   a. **Active Schedule**: Choose either Daily or Weekly, the default is Weekly.

   b. **Time**: Based on a 12-hour clock and the A.M./P.M. modifiers. Time is relative to the DSM system clock.

   c. **Weekday**: Select the day of the week on which to backup the DSM.

      Enter the following information in the **External File Server Settings** section:

   d. **Active Settings**: Select SCP or Windows Share. This configures the mode in which to copy the generated backup file to the remote system. SSH must be configured on the destination system to use the SCP mode. The selected mode—SCP or Windows—determines the subsequent configuration parameters that must be entered

Figure 12-3: Automatic Backup Schedule for SCP

SCP

If you select **SCP**, enter the following information (all fields marked with a red asterisk are required):

- **This Server Security's Credential**: **Click to Export:** Click this to download the DSM server's public key. Copy the public key onto the destination system and into ~`/user/.ssh/authorized_keys`. The public key is required to use SCP to copy the backup file to the external file server.

- **Target Host**: Enter the host name, IP address, or FQDN of the destination system. If the destination system has a File System Agent, you do not have to use the same host name as configured in the Hosts window. You can use any recognized means of addressing the destination system, just as long as it is recognized on your network.

- **Target Host Fingerprint**: The fingerprint value displayed is the fingerprint of the DSM public key that is currently on the destination system. The fingerprint is retrieved from the destination system and displayed in the Automatic Backup page during a backup. You can verify if the public key on the destination system is current by comparing the key in ~`user/.ssh/authorized_keys` on the destination system with the key generated by **Click to Export**.

- **Target Directory**: Enter the full path of the directory in which to copy the backup file.

- **User Name**: Enter the name of the user to perform the copy operation. The name entered must be a valid user on the destination system. Also, copy the public key into the ~`/.ssh/authorized_keys` file in the home directory of the user you specify in this text-entry box. A password is not required for the SCP user because a public key is used to authenticate the SCP user.

Windows Share

If you select **Windows Share**, enter the following information (all fields marked with a red asterisk are required):

- **Network Host**

  Host name, IP address, or FQDN of the destination system.

- **Network Directory**

  The shared folder path to which to copy the backup file.

- **User Name**

  The name of the user to perform the copy operation. The name entered must be a valid user on the destination system.

- **Password**

  The password for User Name. Sometimes a domain is required for user authentication. To include the user domain, append the domain to the user name in the form user @domain. For example, woodford@thalesgroup.com.

  > **Note**
  > Special characters like vertical bar (|), single quote ('), double quote (") and space ( ) are not supported

- **Confirm Password**

  Re-enter the password for User Name.

- Click **Ok** to save the configuration settings currently displayed on the Automatic Backup page, changes to the settings are stored in cache until you click **Ok**.

Figure 12-4: Automatic Backup schedule for Windows Share
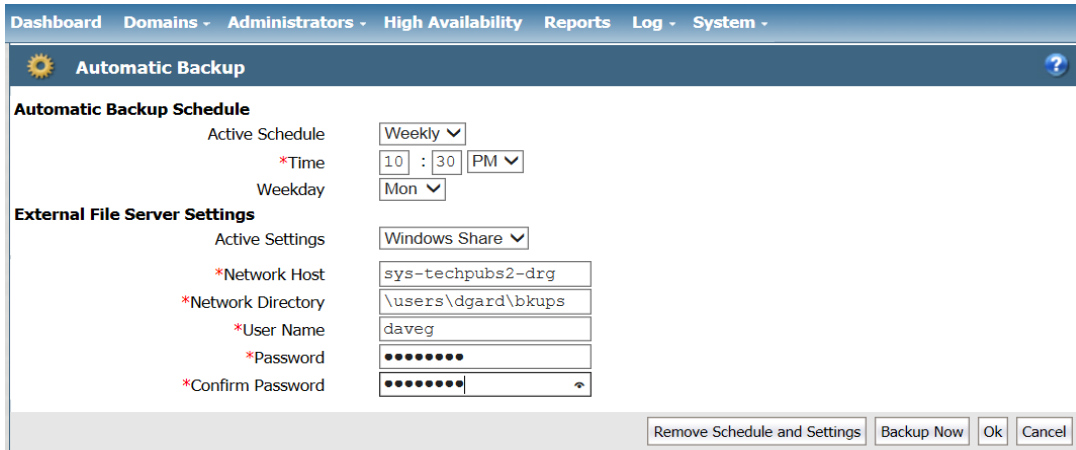
1. Click **Ok** to save the configuration settings or click **Backup Now** to immediately create a backup using the current configuration. This is an easy way to the test network connection and login credentials of the configuration settings you just made.

2. After a successful backup, look in the specified Target Directory on the Target Host to see the backup tar file.

Schedule an immediate backup

You can also schedule an immediate backup once you have made all your selections:

• Click **Backup Now** to create a backup immediately using the current configuration.

This is an easy way to the test network connection and login credentials of the configuration settings you just made.

Remove schedule and settings

Click **Remove Schedule and Settings** to clear all the fields in both the Daily and Weekly configurations. For SCP mode backups, this means the public key is removed, and a new one has to be generated. This new public key has to be copied to the destination system.

A new public key is automatically downloaded the next time you click **Click to Export**. If you create a new key this way, you must also update the `~/.ssh/authorized_keys` file on the destination system because the SSH credentials have changed and will no longer be valid.

Chapter 13: High Availability (HA)

High Availability (HA) is the configuration of multiple DSMs in an HA cluster. Each node runs in parallel. This means there is no longer a primary nor a failover node. All nodes are peers.

High Availability Overview

High Availability (HA) is now configured as Active-Active. This means that all nodes are equal and running in parallel. For example, you can assign GuardPoints and rekey data on any node.

Only DSM System/All Administrators are permitted to configure HA for DSMs.

When one DSM HA node fails, the other HA nodes in the HA cluster continue to run. When the downed DSM HA node is up and running again, it synchronizes with the other DSMHA nodes.

HA configuration and maintenance is done in both the DSM CLI and the UI Management Console.

HA node 1 acts as the Certificate Authority (CA) and creates signing certificates. Once that HA node is configured with its own certificates, it pushes the certificates to other HA nodes when they join the cluster.

Multi-way communication exists between the HA nodes. When a node changes, it synchronizes with all of the other HA nodes.

Because the HA nodes run in parallel, all HA nodes have the same RSA CA and EC CA fingerprints (for example, `5X:5A:51:93:ED:53:B9:8A:1Z:FG:72:3A:BG:` `60:FV:3Q:CE:F7:69:95`).

> **Note**
>
> For configuration information, see the **HA for V6X00 and Virtual Appliances** chapter in the *DSM Installation and Configuration Guide.*

Synchronization Status on the Dashboard

The Management Console **Dashboard** page on all HA nodes displays the high availability synchronization status and the Host assignments.

Figure 13-1: Dashboard HA Status
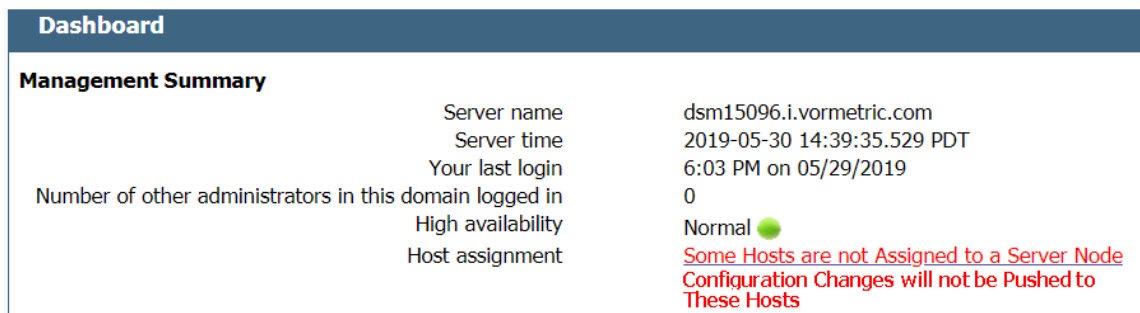
High availability synchronization status

On the Dashboard, next to the term High Availability, the DSM displays the status and a parallel icon. In Figure 13-1: "Dashboard HA Status" above the High Availability status displays as Normal, with a green icon. Clicking the green icon takes you to the High Availability Servers page.

> **Note**
> If you are logged into a domain, you will not see this because domains do not have access to HA.

If high availability has been configured, the FQDN of the HA node(s) display with a synchronization status icon next to the node(s). The icon indicates the status of the HA node.

Clicking the icon next to the HA node link takes you to the Server Node page of the HA node that you clicked on.

The Server Node page on displays the following:

- **From**

  Displays the name of the HA node trying to synchronize.

- **To**

  Displays the name of the HA node with which it is trying to synchronize.

- **Status Collection Time**

  The time shown in this column indicates the last successful synchronization between the HA nodes.

- **HA Replication Status**

  Shows the synchronization status between the nodes.

Dashboard Host Assignment

The Dashboard also displays the status of VTE Agent hosts assigned to the HA cluster.

- If all of the hosts are assigned to nodes, then the message displayed is: **All Hosts are Assigned**.

- If some of the hosts are not assigned to nodes, then the message displayed is: "Some Hosts are not Assigned to a Server Node, Configuration Changes will not be Pushed to These Hosts." When you click on that message, it displays the names of the hosts that are not assigned.

Figure 13-2: Hosts not assigned to nodes

If you have hosts that are not assigned, you can assign them to any nodes in the HA cluster. See the HA for V6X00 and Virtual Appliances chapter in the DSM Installation and Configuration Guide for more information.

Display HA configuration status

1. Log on to the Management Console as a System/All Administrator.

2. Select **High Availability** in the menu bar. The High Availability Servers window opens. It displays the HA nodes.

Table 13-1: High Availability Servers window fields information

Column Header	Description
Selected	Select an HA node to delete the node from the HA cluster.
Name	Displays the fully qualified domain name (FQDN) of the HA node.
Response Time (ms)	Displays the response time in milliseconds, if SNMP is enabled. If SNMP is enabled, an HA node polls the other HA nodes using an SNMP GET request. If SNMP is disabled, the Response Time column displays "SNMP Disabled". If the connection is lost, the Response Time column displays "Not Reachable".
Configured	A check mark in this column indicates that the node is configured and can be accessed by any registered VTE Agents for policy and/or key changes.
Synchronization Status	Shows the synchronization status between the HA nodes. A green circle indicates synchronization between the HA nodes. A red rectangle indicates a synchronization error.

The buttons on the High Availability Servers window are:

- **Add**

  Opens the Add Server window in which to add the host name or FQDN of another HA node.

- **Delete**

  Removes the selected node from the HA cluster.

- **Notify All Hosts**

  Pushes the latest VTE Agent host configurations to every VTE Agent host in the HA cluster.

  If policy changes are not being applied to the hosts that are assigned to an HA node, check the **High Availability** window.

  ○ A green circle should be displayed for all of the HA nodes.

  ○ A red rectangle in the **Synchronization Status** column indicates that an error has occurred.

  ○ Check the network connection between the HA nodes and check that the software is running (for instance, open a Web browser to another HA node). When a host is assigned to another node, policy configuration changes are synchronized on the HA node, and then pushed to the host. If the HA node is going to be down

for an extended period, reassign the hosts to another HA node in the cluster. You can also click **Notify All Hosts** to push policy changes to all the hosts assigned to the DSM, regardless of which DSM they are assigned.

HA Cluster Status

You can monitor the communication status between nodes in the cluster. Communication is now bi-directional and may not be the same between all nodes.

From the High Availability Servers page, if you click on the name of an HA server, the link takes you to the Server Node page for the selected HA node.

The Server Node page displays the following:

- **From**

  Displays the name of the HA node trying to synchronize.

- **To**

  Displays the name of the HA node with which it is trying to synchronize.

- **Status Collection Time**

  Indicates the time of the last successful synchronization between the HA nodes.

- **HA Replication Status**

  Shows the synchronization status between the nodes. Communication is bi-directional.

Server Node HA Status

To find more details on the HA Replication Status, click on the icon to open the Server Node HA Status dialog. The status information listed is:

- **Warning message**

  Error (Red), Warning (yellow), Normal (green)

- **From**

  Node from which the status query originates.

- **To**

  Node queried.

  > **Note**
  > The HA status view no longer shows bi-directional status. It only lists the direction from the node to other nodes. Multi-directional status is illustrated in "HA Topology" on the facing page.

- **Node Status**

  Describes the node status:

  - **Ready**

    Node is fully functional.

  - **Joining**

    Node is bootstrapping. This state occurs when the join has been initiated but synchronizing has not yet been begun.

  - **Joining**

    Node is creating an initial replication slot so that it can communicate with the cluster.

- ○ **Joining**

  Data is transferring to the joining node. The transfer is almost complete.

- ○ **Joining**

  Data is transferred. Node is preparing to join the cluster as a peer.

- ○ **Parting/Parted**

  Node is removed from the cluster by the user.

- **Slot Status**

  Status of channel used to communicate with the cluster.

- **Byte Lag**

  Amount of data available for a node to consume when another node changes, measured in number of bytes.

- **Time Lag (sec)**

  Difference between the current time on a node and the observed time reported by another node.

- **Ping Response Time (ms)**

  Amount of time the ping query took between nodes.

- **Status Collection Time**

  Last time cluster nodes were queried for status.

Setting up HA monitoring for Server Node HA Status

To set up the values for the HA monitoring parameters that display in the Server Node HA status dialog:

1. Click **System > General Preferences** > **HA Monitoring**.

2. Set the following parameters:

   - ○ **HA Node Status Update Interval**

     Set the interval in minutes, (min: 5; max: 59).

   - ○ **HA Status Time Lag Warning**

     Set the interval in seconds, (min: 60 + update-interval)

   - ○ **HA Status Time Lag Error**

     Set the interval in seconds, (> time-lag-warning)

   - ○ **HA Status Byte Lag Warning**

     Set the interval in bytes, (min: 16384)

   - ○ **HA Status Byte Lag Error**

     Set the interval in bytes, (> byte-lag-warn)

HA Topology

HA Topology provides a graphical representation of the HA Cluster node topology. The diagram links in all directions, as indicated by the arrows, from each node to every other node. The edges are color coded to reflect their status: green (normal), yellow (warning), and red (error).

> **Note**
> The HA topology map updates after each node synchronization.

In the following topology map, the yellow lines indicate that system Sys66084.qa.com is having a connection problem with the cluster and the rest of the nodes cannot communicate with it at all. The remaining nodes are communicating properly with the other nodes in the cluster.

Figure 13-3: HA Topology for 8 Nodes
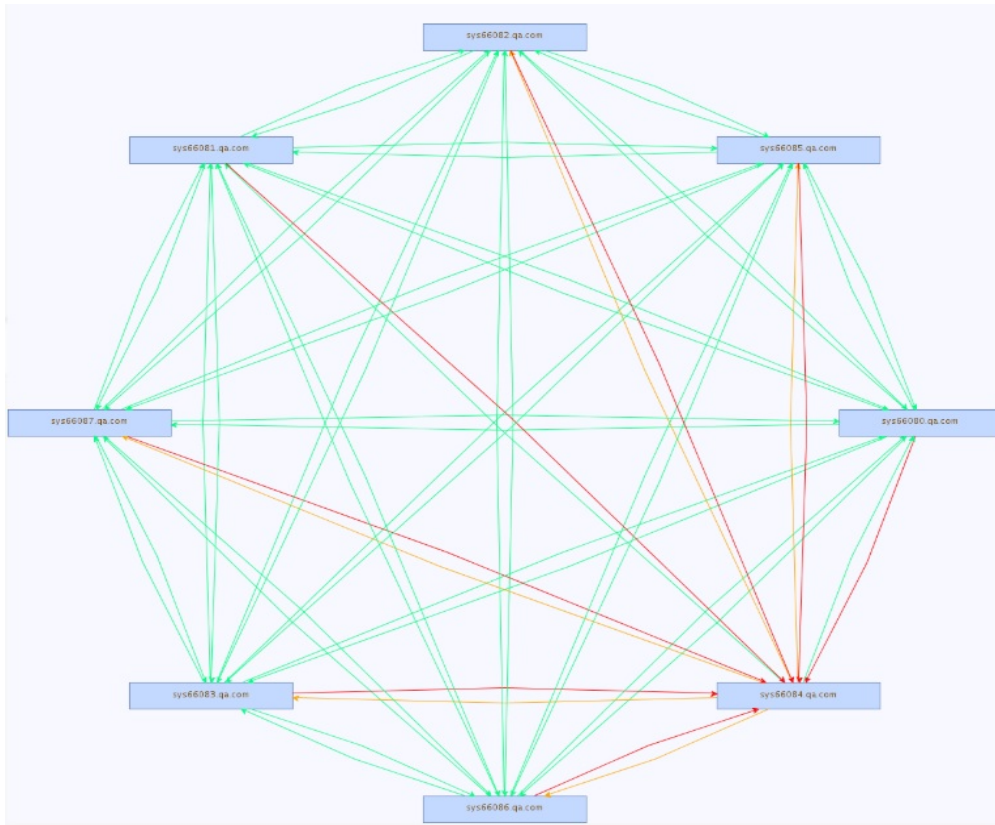
Recovering from incomplete node synchronizations

If the HA nodes do not synchronize completely after replication, do the following to re-initialize an HA node:

1. Remove all nodes from the HA cluster.

2. Run cleanup on each node.

3. Add them back into the HA cluster.

4. Join the HA cluster again.

If the nodes still do not synchronize:

1. In the CLI for the HA node that is not synchronizing properly, go to the Maintenance menu.

2. Reset the HA node, type:

```
0001:maintenance$ config reset
```

3. Type **yes** to continue with the reset.

4. Reconfigure the node.

See the HA for V6X00 and Virtual Appliances chapter in the DSM Installation and Configuration Guide for more information.

Assigning VTE Agent Hosts to HA Nodes in an HA cluster

You can assign VTE agents to any node in an HA cluster. Because the HA nodes are constantly synchronizing, within a few seconds, the VTE agent is available on all nodes in the HA cluster.

Every DSM can check the status of every host in an HA configuration. Each host must have network access to an HA node. It does not have to be an uninterrupted connection, but is required to register the agent and for the node to query the host status. Each host must be on the same network as every node so that it can check the host status.

The Notify All Hosts button in the *High Availability Servers* window pushes the latest host configurations directly from any DSM HA node, to every host in the HA cluster, regardless of which the HA node they are assigned. If you are unsure of how many hosts may be out of sync with their assigned DSM, click the Notify All Hosts button in the High Availability Servers window. This is a convenient way to push the latest host configuration changes to every host, including orphan hosts.

Depending on the number of hosts in the HA cluster and network performance, this can take between a few minutes to a few hours. Check the push status in the Logs window. Messages are placed in the log at intervals to indicate the percentage of completion.

> **Note**
> Do not click Notify All Hosts more than once. Each time you click this button you create a new process and each new process slows the DSM.

VTE Agent Availability

The following describes the high-level process for how a VTE agent joins the HA cluster and becomes available to all of the HA nodes in the cluster:

1. A VTE agent registers to a DSM node that is part of an HA cluster.

2. User adds the VTE host to the HA cluster through the DSM Management Console.

3. The HA node synchronizes with all of the HA nodes.

4. The VTE agent is now available to all HA nodes in the cluster.

Best Practices for HA Clusters

In an HA cluster, the architecture has changed to peer nodes. Therefore, if people are working on the same object on different HA nodes simultaneously, there is a possibility for conflicts.

For example, if two administrators modify the same policy on different HA nodes at the exact same time, the initial modifications will be lost. Only the last operation, meaning the one with the latest timestamp, is saved.

To avoid data conflicts, Thales recommends that you do the following:

- Manage the HA cluster from one dedicated DSM node

- Alternatively, partition the data and manage one partition data in one dedicated node

    For example, you could create domains for different geographic regions, and then manage the domains with the DSM HA node for that region

- If you use the RESTful API, and VMSSC, to manage the DSM HA nodes, do not put any load balancer in front of the DSM HA cluster nodes

- Use the same DSM HA node for data management for a particular job.

For example, if you create and register a host to one DSM HA node, then perform all host operations, such as adding GuardPoints, on that same DSM HA node.

In general, avoid managing the same objects from different DSM HA nodes.

Chapter 14: Multifactor Authentication with RSA

Multifactor authentication increases access control to the DSM Management Console by requiring DSM administrators to enter the value or token code displayed on an RSA SecurID token, along with the administrator name, each time the administrator logs into the Management Console.

Overview

Multifactor authentication on the DSM comprises the DSM, the RSA Authentication Manager, the RSA Authentication Agent, and an RSA SecurID token. The usual sequence is:

1. Configure the DSM HA nodes in the same cluster as the RSA Authentication Agents in the RSA Security Console

2. Create an RSA user ID in the RSA Security Console window

3. Assign the SecurID token to the RSA User ID

4. Test the SecurID token in the RSA self-service console to make certain it is working properly

5. Import the RSA Authentication Agent file into the DSM

6. Associate the RSA user ID to a DSM administrator

7. Enable multifactor authentication

From this point on, a DSM administrator must enter the DSM administrator name, the RSA static PIN (if the DSM administrator ID requires the use of one), and the value displayed on the SecurID token known as a token code, to log into the DSM Management Console.

Figure 14-1: Configuring multifactor authentication
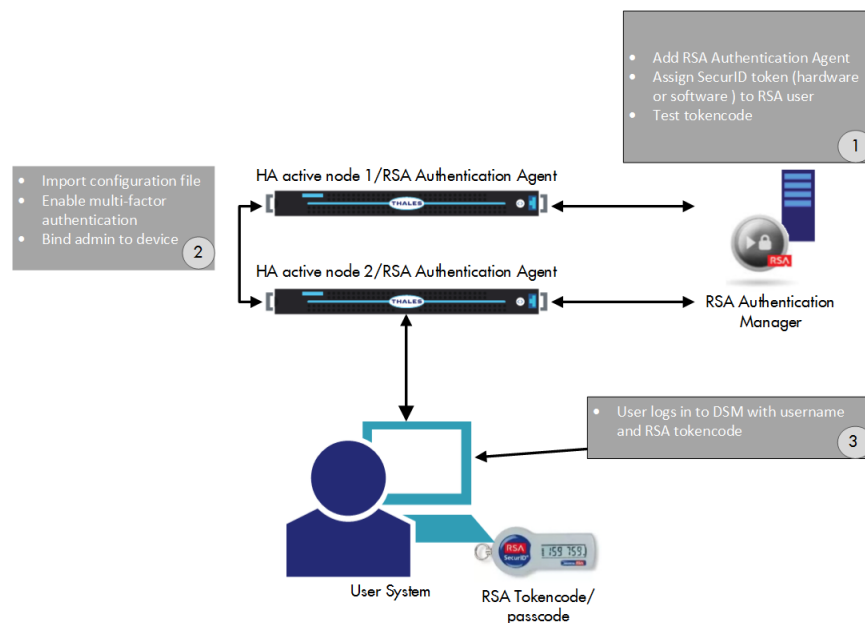
Configuring RSA Authentication

The RSA Authentication Agent is the intermediary between DSM and the RSA Authentication Manager. The RSA Authentication Agent intercepts an access request from the DSM and directs the request to the RSA Authentication Manager server for authentication.

An initial link between DSM and the RSA Authentication Manager is created when an RSA configuration file, `sdconf.rec`, is imported into the DSM Management Console. The first time the RSA Authentication Agent authenticates an administrator with the RSA Authentication Manager, the RSA Authentication Agent node secret is copied and embedded in the DSM.

You can delete the node secret using the `mfauth clean` CLI command. If you delete the node secret, also delete it from the RSA Security Console, and vice versa. This command removes the SecurID file from DSM. A replacement node secret file is automatically downloaded to DSM the next time a DSM administrator logs in with an RSA token code

You must regenerate the node secret file in the RSA Authentication Manager if the DSM installation is destroyed and rebuilt because the DSM authentication credentials are no longer valid.

Multifactor authentication status information is displayed on the DSM Logs page.

Once multifactor authentication is configured, RSA Authentication Manager and DSM startup/shutdown sequence is important.

- Start the RSA Authentication Manager before the DSM.

- Shutdown the DSM before the RSA Authentication Manager.

This sequence is required to ensure that the RSA Authentication Agent can reliably access the RSA Authentication Manager.

Additional RSA configuration notes

The following are additional points to consider when configuring or troubleshooting an RSA configuration.

1. If you are configuring an existing DSM Authentication Agent setup, go to **Access > Authentication Agents > Manage Existing**.

2. Choose Authentication **Agent > Manage Node Secret**.

3. Select **Clear the node secret**.

4. Save your changes.
   - If you are configuring a new DSM Authentication Agent setup, you do not need to go to Manage Node Secret.

The RSA Authentication Manager server requires the IP address and hostname of the DSM. If DNS is configured, the IP address or hostname must be able to be resolved from the RSA Authentication Manager Security Console.

Alternatively, you can edit the `/etc/hosts` file on the DSM and add the DSM IP address and host name. Use the host name to configure the RSA Authentication Agent.

Chapter 15: Viewing and Downloading System-Level Reports

The DSM comes with pre-configured reports that display system information. All reports can be downloaded and saved locally in CSV format.

Overview

All reports are under the **Reports** tab. The availability of reports depends on the DSM administrator type and privileges, and whether the administrator is currently in or out of a domain.

- All Administrators can access all system reports and global domain reports. See Chapter 18: "Viewing and Downloading Domain-Level Reports" on page 143 for more information about domain-level reports.
- System Administrators can access system reports outside global domains.

Viewing and Downloading Reports

To view a report, click the **Report** tab on the Management Console, and then click the name of the report.

To save the report as a CVS file to your local machine, click **Download** at the top left of the report table.

System-Level Reports

The following reports are available to System/All administrators:

- "Administrators" below
- "Servers" on the next page
- "Security Domains" on the next page
- "License Usage by Domain" on the next page
- "System License Usage Summary" on page 120
- "Executive Summary" on page 121

You must be outside of a domain to access system-level reports.

Administrators

The Administrators report is a table of administrators with access to the DSM.

Table 15-1: Administrators

Column Name	Description
User Name	Login

Table 15-1: Administrators (continued)

Column Name	Description
User Type	System Administrator, Security Administrator, Domain Administrator, Domain and Security Administrator, and All
LDAP User ID	The UserID of a user imported from LDAP
Last Login Time	The time of this user's last login. Timestamps are in the form *YYYY-MM-DD HH:MM:SS* where *Y*=year, *M*=month, *D*=day, *H*=hour, *M*=minute, *S*=second

Servers

The Servers report is a table of HA clusters of servers.

Table 15-2: Server Report

Column Name	Description
Server Name	FQDN of the DSM
Up Time	How long the DSM has been active

Security Domains

The Security Domains report is a table of Security domains managed by this DSM.

Table 15-3: Security Domains

Column Name	Description
Organization	Name of the organization responsible for this domain
Domain Name	Name of the domain set when the domain was created.
Description	User added information
Help Desk Information	Phone number of tech support
Domain Administrators	Names of the Domain Administrators assigned to this domain.
Security Administrators	Names of the Security Administrators assigned to this domain.
Domain and Security Administrators	Names of the Domain and Security administrators assigned to this domain.
All Administrators	Names of the All administrators who can access this domain.

License Usage by Domain

The License Usage by Domain report provides a tabular listing of licenses available in the DSM domains and related usage statistics

Table 15-4: License Usage by Domain

Column Name	Description
Organization	Name of the organization responsible for the domain
Domain	Name of the domain
Agent Type	Type of agent license; VTE (FS), VAE (Key), KMIP
License Type	Type of license used in a domain; Perpetual, Term, or Hourly
Agents Licensed to Domain	Total number of agent licenses available to a domain
Agent Licenses Used	Total number of agent licenses used by a domain
Logical Cores Used	Total number of logical cores used in a domain
Logical Core Hours Licensed to Domain	Total number of core CPU hours licenses available to a domain
Logical Core Hour Licenses Used	Total number of core CPU hour licenses used by a domain
LDT Enabled Hosts	Total number of hosts in the domain using VTE Agent licenses with LDT enabled
Docker Enabled Hosts	Total number of hosts in the domain using VTE Agent licenses with Docker enabled
Term Expiration Date	Expiration date of a term license used in a domain, if applicable

System License Usage Summary

The System License Usage Summary report provides a system level overview of all licenses used by the DSM.

Table 15-5: System License Usage Summary

Column Name	Description
Agent Type	Type of agent license; VTE (FS), VAE (Key), KMIP
License Type	Type of license used in a domain; Perpetual, Term, or Hourly
Total Agents Licensed	Total number of agent licenses available on the system
Agents Licensed to Domains	Total number of agent licenses available to domains
Agent Licenses Used	Total number of agent licenses used
Logical Cores Used	Total number of logical CPU cores used by the DSM
Total Logical Core Hour Licenses	Total number of logical CPU core hour licenses available
Logical Core Hours Licensed to Domains	Total number of logical CPU core hour licenses allocated to domains
Logical Core Hour Licenses Used	Total number of logical CPU core licenses used by the DSM
LDT Enabled Hosts	Total number of hosts using VTE Agent licenses with LDT enabled

Table 15-5: System License Usage Summary (continued)

Column Name	Description
Docker Enabled Hosts	Total number of hosts using VE Agent licenses with Docker enabled
Term Expiration Date	Expiration date of a term license used in a domain, if applicable

Executive Summary

The Executive Summary Report shows the totals for the following entities:

- Asymmetric Keys
- GuardPoints
- Hosts (Encryption Expert agents)
- Policies
- Security Domains
- Security Server Administrators
- Security Servers
- Symmetric Keys

Chapter 16: Domain and Security Administrators

Overview

System Administrators create domains but do not operate within them. However, all tasks performed by the Domain Administrators and Security Administrators occur within domains. The Domain Administrators and Security Administrators must always know what domain they are in before executing any task. If you log in as a Domain Administrator or a Security Administrator, and you notice that the administrator, host, or log data is unexpected, you are most likely in the wrong domain.

Domain Administrators can add additional Domain Administrators to each domain. A Global Domain Administrator can be a member of multiple domains. Domain Administrators who are members of multiple domains can easily switch between the domains. They can also add Security Administrators to a domain and assign roles to these Security Administrators (for example, Audit, Key, Policy, Host, Challenge & Response, and/or Client Identity) that are applied only within that domain.

Local or restricted Domain Administrators are restricted to a particular domain. The first Domain Administrator is added to a restricted domain by the System Administrator. After that, the local Domain Administrator creates and adds other Domain or Security Administrators to the Domain as required. Local Domain Administrators and Security Administrators are members of the local domain they are created within, they cannot be members of any other domain. Once created and assigned to a local domain, they are not visible to administrators of other domains.

Domain Administrators cannot remove domains or perform in any of the domain security roles.

A Domain Administrator can:

- Enable and disable Domain Administrator and Security Administrator accounts in the current domain (global Domain Administrators).

- Create, delete, import, enable or disable Domain Administrator and Security Administrator accounts in the current domain (local or restricted Domain Administrators).

- Configure Security Administrator roles (Audit, Key, Policy, Host, Challenge & Response, Client Identity).

- Prevent System Administrators from disabling Domain administrators from within a domain

> **Note**
> When a Domain Administrator changes the configuration of a Security Administrator or another Domain Administrator, the current Management Console session for that administrator is terminated and that administrator must log back in. If a Domain Administrator is removed from a domain, the Domain Administrator cannot switch to or do any work in that domain.

Assigning Domain Administrators or Security Administrators to Domains

There are two types of Domain Administrators that can be created on the DSM:

- Global Domain Administrators

- Local (or restricted) Domain Administrators

A System Administrator adds the first Domain Administrator to a global domain. A global Domain Administrator can add and remove additional administrators (Domain, and Domain and Security) to domains. But, a global Domain Administrator cannot delete administrator accounts.

A System Administrator adds the first Domain Administrator to the restricted domain. The local Domain Administrator can then create new administrators (Domain, Security, or Domain and Security), in the domain, or import LDAP users and make them Domain, Security, or Domain and Security Administrators within the local domain. A local Domain Administrator can also delete those administrator accounts.

A System Administrator can delete global Domain Administrators, but cannot delete local Domain Administrators as they are not visible to the System Administrator. The System Administrator can, however, disable the local Domain Administrator that they added to a local domain.

Add Global Domain or Security Administrators to a domain

1. Log in as a Domain, Domain and Security, or All administrator. The Dashboard window opens.

2. Switch to the domain to which you want to add Domain or Security Administrators.

    a. Select **Domains > Switch Domains**.

    The Domains window opens. All the domains in which the current Administrator is a member, are displayed. The **Selected** radio button of the current domain is grayed out and cannot be selected.

    b. Select the radio button of the domain.

    If the domain is not listed, ask the Domain, Security, Domain and Security, or All Administrator for that domain to add you to it.

    c. Click **Switch to Domain**. The Domains window displays.

3. Select **Administrators > Domain**. The Administrators window opens and displays all of the Administrators who are members of the current domain.

4. Click **Add to Domain**. The Available Administrators window opens. This window lists all the Administrators who are not already assigned to the current domain.

5. Enable the **Selected** check box of the Administrator you want to add to the current domain.

6. If you are adding a Security Administrator, you also need to assign roles to that Security Administrator. Select the administrator role check boxes (Audit, Key, Policy, Host, Challenge & Response, and/or Client Identity) to enable these features.

    Domain and Security Administrator are automatically assigned the following roles: Key, Policy, Host, Challenge & Response and Client Identity. If you want a Domain and Security Administrator to have the Audit role, you must enable that role by selecting the checkbox for Audit.

7. Click **Ok**. The Administrators added to the domain are now active.

Add Local Domain or Security Administrators to a restricted domain

1. Log on to the DSM as a Domain, Domain and Security, or All administrator with your local Domain Administrator credentials. You must select the **local domain** option, and then enter the local domain name in the **Domain Name** field. The Dashboard window opens.

2. Navigate to the Administrators window.

3. Click **New** to create a new Domain, Security, or Domain and Security Administrator.

4. In the Add Administrators window, enter the following information:

- **Login**

  Type a user name for the administrator, it must contain at least 5 characters with an upper limit of 36 characters. Only one instance of an administrator name is allowed.

- **Description (Optional)**

  Enter a description that helps you identify the administrator. The maximum number of characters for this field is 256.

- **Password**

  Enter a password for the administrator. The password must conform to the attributes defined in the password preferences, the maximum password length is 256 characters. The newly created administrators will have to change this password the first time they log on to the DSM.

- **Confirm Password**

  Re-type the password to confirm.

- **User Type**

  Select the type of administrator to create: Domain Administrator, Security Administrator, or Domain and Security Administrator.

- **Read-Only User**

  Select this check box to create an administrator with read-only privileges. You can assign read-only privileges to any type of administrator, except for Local Domain administrators that are the first administrators to be assigned to a domain. If the first administrator added to a local domain is read-only, that administrator will not be able to create any more administrators for that domain.

5. Click **Ok**. The new administrator is displayed in the table on the Administrators page.

6. Click **Import** to import LDAP users to assign as DSM Administrators. You have to have an LDAP Server configured in order to import these users, see Chapter 9: "LDAP Configuration" on page 66 for details.

7. Enter the Login and Password for the LDAP server. If the Login and Password were entered under **LDAP Server Settings** on the AD/LDAP Details window, these values will be populated and do not need to be re-entered. You may also enter a different Login and Password in place of these stored values when you import administrators. Click **Connect**. The LDAP Users window displays LDAP user names.

8. Select LDAP Users:

   The following search option are available on the LDAP Users window:

   - **LDAP Query**

     Use the field to filter searches using the LDAP query language. Results depend on how the LDAP service is set up. See RFC2307 for full details on syntax.

   - **Group**

     Select a group from the drop down list.

   - **User**

     Enter a user name.

   - **Maximum number of entries to return**

     Limits the maximum number of records to import or display. The default value is 300. The minimum value is 1 and the maximum value is 10,000. A high integer value may result in a delay depending on the database size.

   - **Go**

     Click to refresh the screen.

   - **Select All**

     Click to select all values on that page.

- **View**

  Select a value from this drop down box to control how many values appear on any page.

- **Selected**

  Click to select individual values.

- **User Type**

  Select a value from this drop down box to define the type of Administrator or role of the values you import.

- **Add/Cancel**

  Select to add or cancel your selections.

Security Administrator Roles

A Security Administrator can be configured with one or more roles. Domain Administrators assign roles when they assign a Security Administrator to a domain. The roles are applicable only in the current domain. A Security Administrator can be assigned different roles in different domains.

Table 16-1: Security Administrator roles and permitted tasks

Role	Description
Audit	The audit role can only view log data.
Key	The key role can create, edit, and delete local key-pairs, public keys, and key groups. Administrators with this role can also view log data.
Policy	The policy role can create, edit, and delete policies. Administrators with this role can also view log data.
Host	The Host role can configure, modify, and delete hosts and host groups. Administrators with this role can also view log data. The Challenge & Response role is automatically selected when the Host role is selected.
Challenge & Response	The Challenge & Response role must be enabled for a Security Administrator to view the Host Password Challenge & Response window. The window is used to enter a challenge string and display the response string. The response string is a temporary password that a system user enters to decrypt cached encryption keys when there is no connection to the DSM.
	The Challenge & Response role is automatically enabled when the Host role is enabled. You may disable the Host role afterward to leave just the Challenge & Response role enabled. With just this role enabled, the Security Administrator has access to the Dashboard, **Domains > Switch Domains**, and **Hosts > Host Password Challenge & Response** menus only.
	A Security Administrator can open both the **Hosts > Host Password Challenge & Response** window and the **Hosts > Hosts > Challenge Response** tab with the Host and Challenge & Response roles assigned. With just the Challenge & Response role assigned, the Security Administrator can open only the **Hosts > Host Password Challenge & Response** window.
Client Identity	The Client Identity role must be enabled for a Security Administrator to create Identity-Based Key Access. A client identity is used to control access to encryption keys on the DSM by VAE host administrators.

Assigning Security Administrator Roles

> **Note**
> If a Security Administrator is logged in when you assign or change their role(s), that administrator's Management Console session is terminated and they must log on again.

To assign role(s) to a Security Administrator:

1. Log on as a Domain Administrator. The Dashboard window opens.

2. If you are not already in it, switch to the desired domain.

   a. Select **Domains > Switch Domains**. The Domains window opens. All domains in which the current Domain Administrator is a member are displayed. The **Selected** radio button of the current domain is opaque and cannot be selected.

   b. Select the radio button of the desired domain. If the desired domain is not listed, ask the DSM Domain Administrator to add you to that domain.

   c. Click **Switch to Domain**. The **Domains** window opens.

3. Select **Administrators > Domain**. The Administrators window opens and displays all the Domain Administrators and Security Administrators who are members of the current domain.

4. Select an administrator in the **Login** column. The Assign Roles window opens.

5. Enable or disable the **Selected** check boxes for the roles that you want to assign the current administrator. Click **Ok**.

Preventing Security Admins from Disabling Domain Admins

Previously, a DSM System administrator could disable all administrators within a domain, create a new All/security/domain+security administrator and add that administrator to that same domain. This could allow a rogue system administrator to abuse their privileges and misuse the Disable Administrators feature, which would allow them to gain access to a domain to which they should not have access.

The Domain, or Domain & Security administrator, now has the ability to control the behavior of the Disable Administrators feature within the domain in order to prevent other administrators from disabling all of them. They can:

- Override another administrator who set the option globally in System Preferences

> **Note**
> This puts **ALL** of the responsibility within the domain. It means that the Domain administrator would be required to follow practices to ensure that they maintain control over the Domain administrator accounts. If, for example, a Domain administrator were to forget their password, no one could reset it for them.

Globally Preventing Disabling Administrators

An administrator can globally prevent access to the Disable Administrator option to prevent other administrators from disabling all of the administrators in a newly created domain.

> **Note**
> For backward compatibility, the default behavior is to allow the System administrator to disable all of the administrators. This means that the Disable Administrator feature will display in the Assign Admin tab unless the domain administrator checks the option to hide the button.

To prevent a System Administrator from disabling administrators:

1. Click **System > General Preferences > System**.

2. Scroll down to the bottom of the page to the Disable Administrator section.

3. Select **Prevent System administrators from disabling all of the administrators in a domain** option**.**

Locally Preventing Disabling Administrators

DSM added an additional method for controlling the disabling of all of the administrators when creating a Domain, or assigning a Domain administrator. For a newly created domain, the local domain inherits the setting from the General Preferences setting. Hence, if the **Prevent System administrators from disabling all of the administrators in a domain** option is selected in the General Preferences, it will display as checked in the Domain Administrators page for newly created domains. For an existing domain, the default setting is not checked. However, from within the domain, you can override the global setting.

To set the option from within a domain:

1. Click **Administrators**.

2. In the Manage System User field, select **Prevent System administrators from disabling all of the administrators in this domain** option. Click **OK.**

   If you select the option, then the System Administrator CANNOT disable all of the Administrators assigned to that domain. If you DO NOT select it, then the System Administrator CAN disable all of the Administrators assigned to that domain.

3. Click **OK** on the dialog. The Disable Administrators button is hidden in the domain.

Chapter 17: Configuring Syslog Server for Application-Level Messages

This section describes how to add a remote Syslog server to your system, and how to control the severity level and format of the messages that the DSM sends to the Syslog server.

You can configure a Syslog server to receive the same messages that are sent to the Logs window of the Management Console. Use the **System> Log Preferences** menu to create templates that apply to logging configurations for all of the Agents.

Overview

Agent log data is generated on agent hosts. The log data is placed in `/var/log/vormetric` on a UNIX system or in `C:\Documents` or `Settings\All Users\Application Data\Vormetric\DataSecurityExpert\agent\log` on a Windows system, when the **Log to File** logging preference is enabled. The log data can also be forwarded to a Syslog or Event Log server when the Log to **Syslog/Event Log** logging preference is enabled.

> **Note**
> Ensure that the `/var` directory in your system has 256KB to 1MB available for logging to ensure proper DSM behavior.

When **Log to Syslog/Event Log** is enabled, the host administrator can choose to do nothing, which causes log data to be placed into a local `/var/log/messages` or `/var/adm/messages` file, or into the local Windows Event Log, or the host administrator can configure the agent to forward log data to a remote Syslog server or Event Log server. The host administrator can upload the log data to a remote server using the preferred transport protocol. The DSM is not used to configure the remote log servers for host systems.

The **Syslog Server** window in the Management Console lets you configure the remote Syslog servers to which to send DSM log data. The log data sent to remote Syslog servers consists of log data that is generated on the DSM and, when **Upload to Server** is enabled in the **Log Preferences** window, log data that is generated on hosts. The DSM administrator then configures the DSM to forward log data to a Syslog server using either UDP protocol or TCP protocol.

Figure 17-1: Handling log messages
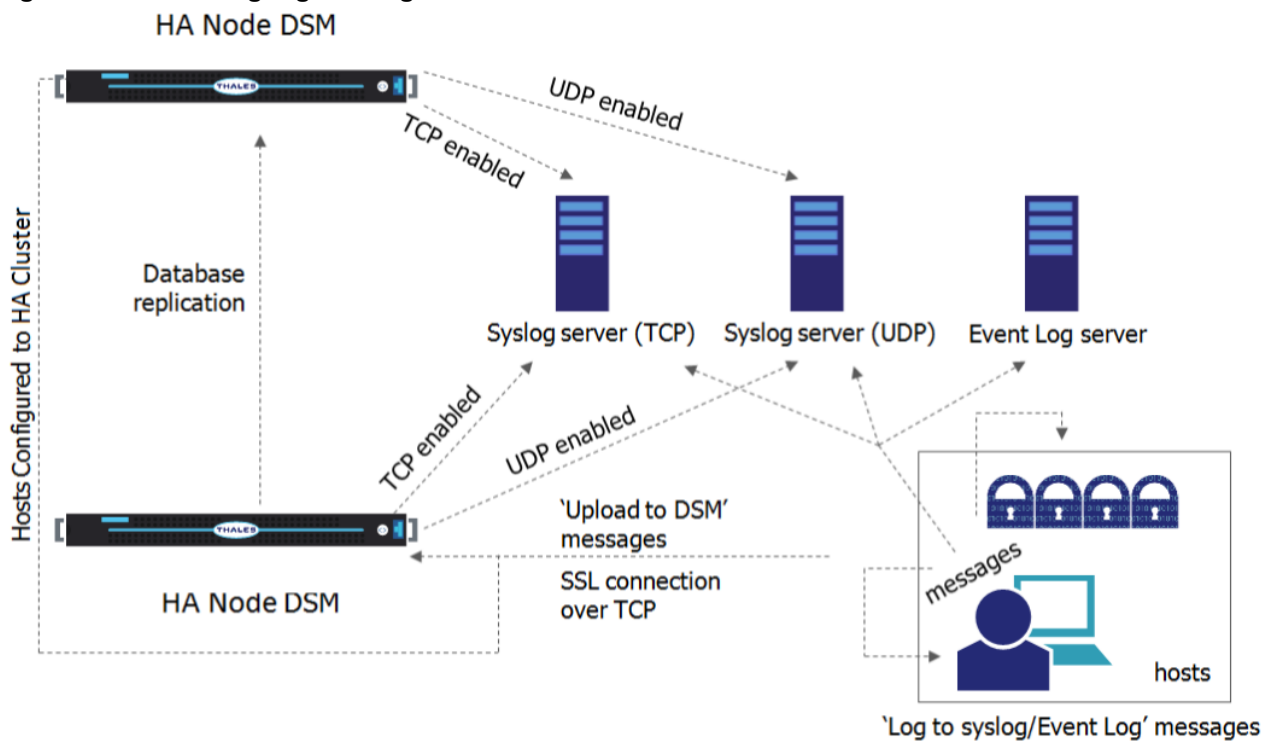
Items to consider before configuring Syslog logging include:

- Only System Domain or All administrators can configure Syslog messaging within a domain.

- If Syslog servers are configured in a domain, only events that take place in that domain are logged to the Syslog servers.

- A default Syslog port number is not provided. The usual industry standard port number for Syslog over UDP is 514. Port 1468 has been used successfully for TCP.

- Configuring a Syslog server is an effective way to consolidate the logs of all of the DSMs in an HA cluster configuration in one central repository.

Supported Syslog Formats

The DSM supports the following log formats:

- Plain Message

- Common Event Format (CEF)

- RFC5424

- Log Event Extended Format (LEEF)

Plain Message

Originally, DSM Syslog supported only Plain Message format. While simple and efficient, this format did not allow for user enhanced reporting or customization.

The following is an example of a Plain Message formatted log message. The table following the message describes the components of the message.

```
12-07-201216:53:02Local7.Debug10.3.32.2312012-12-08 01:01:58.709 vormetric:SOURCE[linux64-
32231.qa.com]:DAO0445I:Administrator voradmin added SysLog Host 10.3.25.168.
```

Table 17-1: Syslog message parameters and descriptions

Parameter	Description
`12-07-201216:53:02`	Date and time
`Local7.Debug`	Message priority
`10.3.32.231`	Sending machine's IP address
`2012-12-08 01:01:58.709`	Date and time of logged event
`vormetric`	Originator tag
`SOURCE[linux64-32231.qa.com]`	Source of message
`DAO0445I`	Unique message ID
`Administrator voradmin added SysLog Host 10.3.25.168`	Plain text message of the logged event

Common Event Format (CEF) log format

DSM Syslog supports Common Event Format (CEF) log format. The Vormetric CEF format is specified in the Arcsight "Common Event Format" standard.

The following is an example of a CEF formatted log message.

```
<27> 2012-10-16T16:01:44.030Z centos-6-0 CEF:0|Vormetric, Inc.|vee-fs|5.1.0.9026|CGP2604E|
Reject access|7|logger=CGP spid=6362 cat=[ALARM] pol=AuditAllExceptLp
uinfo=lp,uid\=4,gid\=7\\lp\\ sproc=/bin/ls act=read_dir_attr gp=/Guard
filePath=/datafiles/file.dat denyStr=DENIED showStr= Code (1M)
```

Table 17-2: CEF Log Format parameters and descriptions

Parameter	Description
`<27>`	A standard Syslog facility/priority code
`2012-10-16T16:01:44.030Z`	Date and time
`centos-6-0`	The host name of the machine sending the message.
`CEF:0`	Version of the CEF
`Vormetric, Inc.`	Sending device vendor
`vee-fs`	Sending device product
`5.1.0.9026`	Sending device version
`CGP2604E`	Unique message ID
`Reject access`	Name: A human-readable and understandable description of the event.
`7`	Severity: An integer that reflects the importance of the event. Only numbers from 0 to 10 are allowed, where 10 indicates the most important event.

Table 17-2: CEF Log Format parameters and descriptions (continued)

Parameter	Description

`logger=CGP spid=6362 cat=[ALARM]`
`pol=AuditAllExceptLp`
`uinfo=lp,uid\=4,gid\=7\\lp\\ sproc=/bin/ls`
`act=read_dir_attr gp=/Guard`
`filePath=/datafiles/file.dat denyStr=DENIED`
`showStr= Code (1M)`
	Extension: A collection of key-value pairs. The keys are part of a predefined set. The standard allows for including additional keys. An event can contain any number of key-value pairs in any order, separated by delimiting characters.

RFC5424

Vormetric Syslog supports RFC5424 log format.

An example of an RFC5424 formatted log message follows. Components of the message are described in the table following the message example:

```
<30>1 2012-12-07T21:44:04.875Z t3-normaluser.i.vormetric.com vee-FS 0 CGP2603I [CGP@21513
sev="INFO" msg="Audit access" cat="\[AUDIT\]" pol="normaluser-only-aes256"
uinfo="normaluser,uid=2001,gid=1\\other\\" sproc="/usr/bin/cat" act="read_attr"
gp="/export/home/normaluser/test" filePath="/test.txt" denyStr="PERMIT" showStr="Code
(1M)"]
```

Table 17-3: CEF Log Format parameters and descriptions

Parameter	Description
`<30>1`	A standard Syslog facility and priority code
`2012-12-07T21:44:04.875Z`	Date and time
`t3-normaluser.i.vormetric.com`	The host name of the machine sending the message.
`vee-FS`	Sending device product
`0`	Process ID field having no interoperable meaning, except that a change in t he value indicates that there has been a discontinuity in Syslog reporting.
`CGP2603I`	Unique message ID

`[CGP@21513 sev="INFO" msg="Audit access"`
`cat="\[AUDIT\]" pol="normaluser-only-aes256"`
`uinfo="normaluser,uid=2001,gid=1\\other\\"`
`sproc="/usr/bin/cat" act="read_attr"`
`gp="/export/home/normaluser/test"`
`filePath="/test.txt" denyStr="PERMIT"`
`showStr="Code (1M)"]`
	Structured data field: Provides a mechanism to express information in a well-defined, easily parsable and interpretable data format. This field consists of the Structured Data (SD) Element, SD-ID, and SD-Parameter.

Log Event Extended Format (LEEF)

The DSM supports Log Event Extended Format (LEEF). The LEEF header is pipe ("|") separated and attributes are tab separated.

Adding a Syslog Server

To add a Syslog server:

1. Verify that one or more Syslog servers are accessible from the DSM system. It is usually enough to ping the Syslog server and run ps to check the Syslog process on the Syslog server system.

   If you are going to send the messages to the local host, verify that the syslogd process on the local host is accepting messages. You may need to restart syslogd with the "-r" argument.

   > **Note**
   > Record the Syslog transport protocols and port numbers of the Syslog server(s). You will need this information later.

2. Set the severity level at which to send messages to the Syslog server in the `/etc/syslog.conf` file on the agent host.

   Severity levels in the Log Preferences window are DEBUG, INFO, WARN, ERROR, and FATAL. Severity levels are cumulative, so each level includes the levels below it. For example, FATAL logs only FATAL messages, whereas WARN logs WARN, ERROR, and FATAL messages. To ensure that the Syslog server gets the messages set in the Log Preferences window, set the level in the syslog.conf file to debug and direct the output to the local messages file. For example, on a Solaris system, set the output file path to /var/adm/messages.

   ```
   user.debug /var/adm/messages
   ```

3. Log on to the Management Console as a System/All Administrator.

4. Select **System > General Preferences**. The **General Preferences** window opens to the **General** tab.

5. Click the **System** tab, and then select **Syslog Enabled**.

   This enables communication between the DSM and the Syslog server.

   > **Note**
   > You must have the **Syslog Enabled** box selected from outside a domain; otherwise, the **Apply** button will not be selectable from within a domain.

6. Click **Apply**.

7. Select **System > Log Preferences**. The Log Preferences window opens to the Server tab.

8. Set the **Logging Level** property.

   The level you select affects the number of messages that are displayed in the Logs window, and these messages are also sent to the Syslog server.

   Redundant Syslog failure messages are filtered so that only one out of every fifty redundant messages is sent to `/var/log/messages` and the Logs window. All the redundant Syslog failure messages are sent when the level is set to DEBUG.

9. Click **Apply**.

   ○ If you are configuring a Syslog server to receive domain-level log data, and are logged in as an All administrator, remain logged in and enter the domain to be configured.

   ○ If you are configuring a Syslog server to receive domain-level log data, and are logged in as a System Administrator, log out and log back in as Domain/All Administrator and enter the domain to be configured.

10. Select **Log > Syslog**. The **Syslog Server** window opens.

11. Click Add and enter the following information:

    a. **Server Name:** The host name or FQDN of a Syslog server. Use the network name of a Syslog server which is accessible to the HA cluster.

    b. **Transport Protocol:** Select UDP, TCP or TLS from the drop down. If you select TLS, a field appears for you to browse to a Root Certificate.

In the interests of security, Vormetric recommends that you use a root certificate rather than a non-root certificate.

> **Note**
>
> For Syslog servers configured with the UDP transport protocol, ensure that UDP packets are not blocked by a firewall or switch rules. Also, verify that the Syslog server is logging messages as expected.
>
> If you add a Syslog certificate when using TLS protocol, you may need to restart the server. To this you need to do a `system > server restart` from the CLI. After restart, verify that the Syslog server is logging messages as expected.

    c. **Port Number:** The port number the transport protocol uses to connect to the Syslog server. Enter a value between 1 and 65535. There is no default.

    d. **Message Format:** Select Plain Message, CEF, or RFC5424.

      You may configure multiple Syslog servers but only one instance of a Syslog server name is allowed in the DSM database.

12. Click **Ok**. Perform a task on an agent system that normally generates a Syslog entry, such as accessing a GuardPoint.

13. Check the `/var/log/messages` file on the Syslog server for DSM log entries.

Using Syslog to Troubleshoot the DSM and Agents

Syslog entries for DSM activity indicate the source of the Syslog message (system name after the timestamp), the source of the message itself (SOURCE), the log level (AUDIT, ALARM, and so on), and much more.

Analyzing log entries

The format and content of log entries for File System Agents are described below.

Figure 17-2: Message Log entries

Analyzing VTE Agent log entries

The general format of a VTE Agent log entry is:

```
CGP2602I: [SecFS, 0] Level: Policy[policyName?] User[userID?] Process[command?] Access
[whatIsItDoing?] Res[whatIsItDoingItTo?] Effect[allowOrDeny? Code (whatMatched?)]
```

where:

- SECFS indicates that the message was generated by a VTE Agent. You can enter `secfs` in the **Search Message** text-entry box in the **Logs** window to display VTE Agent policy evaluation and GuardPoint activity for all configured hosts.

- Level indicates the importance of the message. For example, AUDIT indicates an informational message, whereas ALARM indicates a critical failure that should not go ignored.

- Policy[] indicates the name of the policy that is being used to evaluate the access attempt.

- User[] identifies the system user attempting to access data in the GuardPoint. It typically displays the user name, user ID, and group ID.

- Process[] indicates the command, script, or utility being executed.

- Access[] indicates what is being attempted. Access may be read_dir, remove_file, write_file_attr, write_app, create_file, etc. These correspond to the Access methods that you configure in the policy. read_dir corresponds to d_rd. remove_file corresponds to f_rm. And so on.

- Res[] indicates the object being accessed by Process[].

- EFFECT[] indicates the rule that matched and, based upon that rule, whether or not the DSM grants access. Access states may be either PERMIT or DENIED.

For example:

```
CGP2606E: [SecFS, 0] [ALARM] Policy[allowAllRootUsers_fs] User
[hpotter,uid=1111,gid=10\wheel\] Process[/usr/bin/vim] Action[create_file] Res
[/opt/apps/apps1/lib/file1.txt] Effect[DENIED Code (1M)]
```

The format of a rule match is:

```
intchar
```

where:

- `int` is an integer representing the security rule being used or violated. Security rules are numbered sequentially from top to bottom in the Online Policy Composer window.

- `char` is an uppercase letter indicating the item that is using or violating the policy.

Table 17-4: Character Codes and Their Descriptions

Character Code	Description
A	The Action component of a security rule failed to match.
M	All security rule components match and, unless overridden, the **Effect** for that security rule is applied.
P	The **Process** component of a security rule failed to match.
R	The **Resource** component of a security rule failed to match.
T	The time specified in the **When** component of a security rule failed to match.
U	The **User** component of a security rule failed to match.

For example, the following match codes indicate:

- **1R** – Mismatch in Resource for Security Rule 1.

- **3U** – Mismatch in User for Security Rule 3.

- **4A** – Mismatch in Action for Security Rule 4.

- **2M** – All components matched for Security Rule 2. Since all the rules matched, Security Rule 2 will be used and no other rules will be evaluated.

Log message levels

The detail and extent of information logged is determined by the selected log level. The agent supports five log levels as listed in Table 17-5: "The Agent-Supported 5 Log Levels" below.

Table 17-5: The Agent-Supported 5 Log Levels

Severity	Description
DEBUG	The DEBUG level provides detailed information about events that are intended for support engineers and developers.
INFO	The INFO level provides general information that highlights the progress of the application.

Table 17-5: The Agent-Supported 5 Log Levels (continued)

Severity	Description
WARN	The WARN level designates potentially harmful situations.
ERROR	The ERROR level designates error events that might still allow the application to continue running.
FATAL	The FATAL level designates very severe error events that will presumably lead the application to quit.

Log levels are cumulative. The level that you select not only generates log entries for events that occur at that level, but all the levels below. For example, the WARN level also includes events that occur on the ERROR and FATAL levels.

Using log files

Check the log files to verify the successful installation and configuration of the Vormetric Data Security software, to determine why a backup or restore operation failed, or to monitor Vormetric Data Security activity.

A logged event falls into one of the following categories:

- **Operational status**

    The result of any significant action performed by an VTE Agent or DSM is logged.

- **Administrative activity**

    The result of any maintenance or administrative activity on the DSM is logged (for example, a key has been created or exported).

- **System status**

    The result of any system errors are logged (for example, if the database connection is interrupted).

- **Policy-specified audit**

    If the result of a policy evaluation specifies that it should be audited, then a suitable message is logged.

Several logs files are provided. Each serves a different purpose. The log files are:

(Windows only) The `\ProgramData` folder on Windows Vista and Windows Server 2008, and the `\Documents and Settings\All Users\Application Data` folder for all other supported Windows platforms, are hidden by default. VTE Agent logs, configuration data, and certificates are stored under that folder. If you cannot browse the folder for your platform, enable the **Show hidden files and folders** radio button in the **Folder Options** menu to view the folder and its contents.

Active logs are log files that being currently written to and updated by DSM processes. Inactive logs are logs that have been filled to capacity and then closed. The name of the closed log file is the original name usually appended with the date and some random numbers. For example, the name of an active agent log is vordb2_usr.log. When it reaches the configured capacity, it is made inactive and usually renamed to `vordb2_usr.log.YYYY-MM-DD-MM-SS.tar.gz`. For example, the archive file for `vordb2_db2inst1.log` can be `vordb2_db2inst1.log.2011-01-19-12-25-32`.

Do not try to manually modify or remove active logs. Use the Management Console interface to configure server and VTE Agent logs. Regularly back up and delete inactive logs to maximize available hard disk space.

The Windows system event log can fill quickly. If a Windows host runs out of system event log space, the `vmd` service does not start and issues an error: "`The service did not respond to the start or control request in a timely fashion.`"

To prevent the system event log from running out of space, the current event log is archived to a file when it reaches 20MB, all archived entries are then purged from the event log, and logging continues as usual. Archive files are placed in `%SystemRoot%\System32\Config`.

The archive file is named `Archive-Vormetric Encryption Expert-timestamp.evt`. For example, `Archive-Vormetric Encryption Expert-2010-05-14-18-14-30-171.evt`. The file is archived in a binary format that you can open in the Event Viewer. Check disk space availability during periods of heavy load and extensive logging. Back up and delete the archive files.

VTE Agent Log Files

The agent logs are the first places to check when communication between the DSM and VTE Agent system fails. Also, you may want to check these logs after setting up a new agent or changing the agent configuration.

vorvmd.log (Windows)/vorvmd_root.log (UNIX)

(UNIX)

> `/var/log/vormetric/vorvmd_root.log`

(Windows)

> `\Documents and Settings\All Users\Application Data\Vormetric\DataSecurityExpert\Agent\log\vorvmd.log`

(Windows XP)

> `\Documents and Settings\All Users.WINDOWS\Application Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log`

(Windows Vista and Windows Server 2008)

> `\ProgramData\Vormetric\DataSecurityExpert\Agent\log\vorvmd_root.log`

(Windows) The same information that is sent to `vorvmd.log` can also be sent to the Windows Event Viewer. Enable **Log to Syslog/Event Log** logging options for the agents and open **Event Viewer > Vormetric Encryption Export** to view log events on the host system.

`vorvmd_root.log` contains the VTE Agent transactions for the root user. Transactions consist of a record of vmd actions, such as starting the vmd daemon and setting up communication links with the DSM.

messages (UNIX only)

> `/var/log/messages`

Messages is a syslog-generated file. It contains standard syslog entries. It contains kernel entries for enabling/disabling the log service, memory usage, CPU usage, system calls, device initialization, etc. It also contains log entries that are also displayed in the Message Log.

secfs.log (AIX only)

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. The `secfs.log` file is generated only on AIX systems. The secfs.log file is maintained in the `./agent/secfs/tmp` directory. It is used instead of syslog to log kernel messages. The same log messages are placed in both `/var/log/messages` and `secfs.log`. The `secfs.log` file is archived at 32MB and renamed to `secfs.log.archive`. Only one archive file is maintained.

secfsd.log

(UNIX)

```
/opt/vormetric/DataSecurityExpert/agent/secfs/tmp/secfsd.log
```

(Windows Server 2003)

```
C:\Documents and Settings\All Users\Application
Data\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

(Windows Vista and Windows Server 2008)

```
C:\ProgramData\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

(Windows XP)

```
C:\Documents and Settings\All Users.WINDOWS\Application
Data\Vormetric\DataSecurityExpert\agent\log\secfsd.log
```

The `secfs.log` file contains kernel-related messages, and the `secfsd.log` file contains process-related messages. secfsd.log contains a record of GuardPoint mounts and GuardPoint dismounts (GuardPoints are mounted file systems). Entries are added to this file when you add and remove GuardPoints, as well as when you reboot the agent system.

statusfile

```
/opt/vormetric/DataSecurityExpert/agent/secfs/tmp/statusfile
```

```
\Program Files\Vormetric\DataSecurityExpert\agent\secfs\tmp\statusfile
```

`statusfile` is a current record of the local VTE Agent configuration. View this file after updating the VTE Agent configuration on the DSM to verify that the changes have actually been applied. This file should always be checked when the configuration of the VTE Agent is in question. This file lists:

- Each GuardPoint and GuardPoint properties, such as the lock status, protection status, and GuardPoint directory

- The names of applied policies

- The logging information that is captured

- Where captured log information is sent

- Hosts settings

You can also display the file timestamp to see when the agent was last updated.

This file is deleted each time the VTE Agent configuration is updated. You must manually regenerate it using the "`secfsd -status`" command. If you want to keep records of VTE Agent configuration changes, either copy the `statusfile` to a different name, or run "`vmsec status`" and tee the output to a different file.

(Windows) The `secfsd` command has limited support on Windows platforms. You can use the `secfsd -status lockstat` command or use the Vormetric Data Security tray to open the status window. Look for strings like `coreguard_locked=true` and `system_locked=true`. (`false` indicates that a lock is not applied. `true` indicates that a lock is applied.).

You may view the file contents using an ASCII display command, such as `cat`.

DSM Log Files

The initial DSM log is viewed in the Logs window of the Management Console. This log is generally the first log that you check to diagnose server problems. Check the DSM log after making or restoring a database backup. Look for

entries like "`Backup Request for SAMPLE from host vmSSA06 is allowed.`" and "`Backup/Restore completed successfully.`" Messages like "`Backup data request failed: access denied or a related cause.`" indicate a problem has occurred and some debugging on your part is required.

Viewing the log files is easier on a software-only server than an appliance-based server. On a software-only server you can use a favorite editor to search a log or copy logs nightly as part of a batch process. Use the `diag` CLI command to list and view the log files. However, the log files can be exported from an appliance using the various export features in the Logs window. DSM administrators cannot delete log files.

The JBoss application server creates three log files in `/opt/vormetric/coreguard/server/jboss-5.1.0.GA/server/default/log`.

- `boot.log` contains JBoss startup information.

- `cgss.log` contains DSM information.

- `server.log` contains system-level information.

`boot.log` is managed as a single file. It is not expected to ever become a large file nor is it rotated. The `cgss.log` and `server.log` files can become large and are rotated.

The three log files are physically stored in `/opt/vormetric/coreguard/server/jboss-5.1.0.GA/server/default/log`. Alternate access is provided through the symbolic link, `/opt/vormetric/coreguard/server/log`.

The `cgss.log` and `server.log` files are important log files that can grow quickly under heavy load. Because these logs are vital to analyzing DSM behavior, they should be monitored and backed up regularly.

The names of the active files are `cgss.log` and `server.log`. When either file contains 10MB of log data it is made inactive and renamed to `cgss.log.1` or `server.log.1`, respectively. And a new active `cgss.log` or `server.log` file is opened. When the new active log file reaches 10MB it is made inactive and renamed to `cgss.log.2` or `server.log.2`. And a new active log file is opened. This process continues until there are a total of 10 inactive log files. When there are 10 inactive log files, and the active log file reaches its full 10MB capacity, the first inactive file is discarded, all the other log file names are decremented by one, and the former active log becomes the 10th inactive log file. Using `cgss.log` as an example, when `cgss.log` fills, `cgss.log.1` is discarded, all the other log file names are decremented by one, and `cgss.log` becomes `cgss.log.10`. Depending on the load you place on the server, and if your policies audit a lot of data, these files can grow and rotate quickly.

badlog.log

Log files with unparsable data are "bad logs". A badlog.log file contains log data from an agent that is intended for display in the Logs window but which cannot be displayed because the log data cannot be parsed due to format irregularities. Each attempt by an agent to upload an unparsable log file to the server is placed in the badlogs directory as a unique file. Regardless of the number of failed attempts to parse incoming log files, the DSM will continue to accept uploaded logs from the agent.

Log files are in an XML format. Log files originate on the UNIX agent in the `/var/log/vormetric` directory and they are removed from the agent after they are successfully uploaded to the DSM. If the DSM cannot parse the file, it is placed in `/opt/vormetric/coreguard/server/jboss-5.1.0.GA/server/default/auditlog/badlogs/vmd_upload_hostName.num`. For example, `vmd_upload_vmlinux101.374`.

cgss.log

The cgss.log file contains a record of the events that make up the BEK generation process for an agent requesting to make a backup, as well as the names of uploaded audit files. This file does not contain events that pertain to restore operations. Check this file if the agent fails to back up a database, even though agent/server authentication is correctly

configured and the policy for this agent permits the backup operation.

This file is located in:
`/opt/vormetric/coreguard/server/jboss-5.1.0.GA/server/default/log`.

server.log

The server.log file contains details about agent backup and restore requests, connection status, Management Console interaction, Java exceptions, JBoss start and stop processes, and more. This file contains diverse information and should be checked for almost any problem that is related to the DSM. Sometimes it is easier to `grep` a specific error level, like WARN, INFO, or DEBUG, than it is to view the entire file.

This file is located in: `/opt/vormetric/coreguard/server/jboss-as/standalone/log`.

Exporting Logs

You can export the log entries that are displayed in the Logs window to maintain a separate record of server and agent activity at the application level.

The data displayed in the Logs window can be exported to a file for archival or analysis. Only the entries in the Logs window that are appropriate for the administrator type and domain can be saved to a text file. The output file is formatted as a comma-separated list and is usually viewed in a spreadsheet application.

The following example is an excerpt of a .csv file generated by an All administrator that is inside a domain.

Figure 17-3: Excerpt of a log .csv file

	A	B	C	D	E	F	G	H	I	J	K
4	1426712	1000	FS	3009937978	195375299	2011-01-21-20.40.24.563000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
5	1426713	1000	FS	3009937978	195375299	2011-01-21-20.40.37.173000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
6	1426714	1000	FS	3009937978	195375299	2011-01-21-20.40.49.053000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
7	1426715	1000	S			2011-01-21-20.38.15.336000	I	vmSSA05	DAO0235I	DAO0235I: Administr	-480
8	1426716	1000	S			2011-01-21-20.38.15.340000	I	vmSSA05	DAO0235I	DAO0235I: Administr	-480
9	1426717	1000	FS	3009937978	195375299	2011-01-21-20.40.58.303000	I	solaris120	VMD3781I	VMD3781I: [vmd, 244	-480
10	1426718	1000	FS	3262128009	236323513	2011-01-21-20.40.58.255000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
11	1426719	1000	FS	3262128009	236323513	2011-01-21-20.40.58.749000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
12	1426720	1000	FS	3262128009	236323513	2011-01-21-20.40.58.846000	I	solaris120	CGA3001I	CGA3001I: [SecFS, 0	0
13	1426721	1000	FS	3262128009	236323513	2011-01-21-20.41.09.216000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
14	1426722	1000	FS	3262128009	236323513	2011-01-21-20.41.09.217000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
15	1426723	1000	FS	3262128009	236323513	2011-01-21-20.41.09.217000	I	solaris120	CGP2603I	CGP2603I: [SecFS, 0	0
16	1426724	1000	S			2011-01-21-20.38.30.337000	I	vmSSA05	COM0314I	COM0314I: The Secu	-480
17	1426725	1000	FS	3009937978	195375299	2011-01-21-20.41.44.104000	I	solaris120	VMD3794I	VMD3794I: [vmd, 244	-480
18	1426726	1000	FS	3009937978	195375299	2011-01-21-20.42.00.102000	I	solaris120	DXF4344I	DXF4344I: [vmd, 1307	-480
19	1426727	1000	S			2011-01-21-20.39.38.570000	I	vmSSA05	LOG0140I	LOG0140I: Administr	-480
20	1426728	1000	FS	3009937978	195375299	2011-01-21-20.43.53.294000	I	solaris120	DXF4366I	DXF4366I: [vmd, 1308	-480
21	1426729	1000	FS	3009937978	195375299	2011-01-21-20.43.54.302000	I	solaris120	DXF4371I	DXF4371I: [vmd, 1308	-480
22	1426730	1000	FS	3009937978	195375299	2011-01-21-20.43.54.311000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
23	1426731	1000	FS	3009937978	195375299	2011-01-21-20.43.54.315000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
24	1426732	1000	FS	3009937978	195375299	2011-01-21-20.43.54.319000	I	solaris120	DXF4201I	DXF4201I: [vmd, 1308	-480
25	1426733	1000	FS	3009937978	195375299	2011-01-21-20.43.54.324000	I	solaris120	DXF4200I	DXF4200I: [vmd, 1308	-480
26	1426734	1000	FS	3009937978	195375299	2011-01-21-20.43.54.327000	I	solaris120	DXF4200I	DXF4200I: [vmd, 1308	-480

The format of this table is subject to change. At this time, the columns indicate:

Table 17-6: Exported Message Log Headings and Description

Column Heading	Description
A	ID number in the Management Console (LOG_ID)
B	Internal domain identifier. If you are not in a domain this is zero. (DOMAIN_ID)
C	Entity that generated the message. It can be S (DSM), FS (VTE Agent). (SOURCE)

Table 17-6: Exported Message Log Headings and Description (continued)

Column Heading	Description
D	Internal tag (TAG)
E	Internal subtag (SUBTAG)
F	Time of action in UTC (LOG_TIMESTAMP)
G	Severity in the Management Console (SEVERITY)
H	Source in the Management Console (HOST_NAME)
I	DSM or agent Message ID. For example, DAO0239I or CGP2603I.The Message ID also identifies the log service. For example, CGP2603I was generated by the CGP (Policy Evaluation Service) log service. (MESSAGE_ID)
J	Message in the Management Console (MESSAGE)
K	Time offset in minutes. Subtract this number from the time in column F to determine local time. F - K = local time. For example, 22:40:19 UTC - 420 offset = 15:40:19 PDT = 3:40 PM PDT. (TIMEZONE_OFFSET)

Exporting the Message Log

1. Log on to the Management Console as an administrator of the appropriate type for the data you want to export.

2. Enter a domain if you want to export domain-related log entries.

3. Open the Logs window.

4. Click **Export Logs**. The File Download window opens.

   The options are:

   ○ **Open** to display the log entries to be exported in the default spreadsheet application. Usually this is Excel.

   ○ **Save** to export the log to a file on the system running the Management Console Web session or on another network accessible system. The default output file name is `log.csv`.

   ○ **Cancel** to close the window and stop the export operation.

5. Click **Save**. The Save As window opens.

6. Enter the name and path for the export file. The default file name is `log.csv`.

7. Click **Save**. The Download Complete window opens. It displays statistical information about the exported log, such as its location and size.

   The options are:

   ○ **Open** to open the exported log file in the default spreadsheet application used to process CSV format files.

   ○ **Open Folder** to open a Windows Explorer window in which to select and view the exported log in an application of your choice.

   ○ **Close** to close the window.

8. Click an option to open the exported log in the default spreadsheet application, open the exported log file in a different application, or to close the window and continue other Management Console operations.

Adding an email Notification Group

Email Groups are per domain. You can set up email groups for domains of System, Security, Domain, Domain/Security and All Administrators.

Enabling email notification for log messages

You can automatically send email notifications to a set of administrators if the DSM generates a serious log message.

You need to configure an SMTP server first. Navigate to **System > Email Notification** and click the SMTP Server tab. Enter the information for the following tabs:

- **SMTP Server**

  SMTP server that will send the email notification. SMTP Servers are per appliance and you must be signed in with System Administrator privileges to modify this setting. If you don't have these privileges, the SMTP server setting is grayed out. Note that the appliance does not come with a default SMTP server and that the SMTP server settings are initially empty.

- **SMTP Server Port**

  Port used by the SMTP server.

To bring up the Email Notification interface, select **System > Email Notification** when outside a domain. The attributes and interface information for the Email Notification are as follows:

- **Email Group Name**

  Name of the email group which will receive the email notification. Email Groups are per domain. You can set up email groups for domains of System, Security, Domain, Domain/Security and All Administrators.

- **Email Threshold Level**

  If the DSM generates a log message with a severity of this specified threshold level, then an email notification is generated. Can be ERROR or FATAL.

- **Email Address List**

  Email addresses that will receive this email notification. Separate addresses with commas. If LDAP is configured, you can select addresses from your LDAP address book by pressing Select. If it's not configured, you can enter your login and password to access it.

- **Email Subject**

  Text you want on the subject line.

- **Message Contains**

  This is a string filter that works with the Email Threshold Level. Only messages containing this string will be sent as an email notification. If blank, then all messages meeting the threshold criteria will be sent.

- **Enabled**

  A checkbox that enables or disables email notification to the group.

To add an email notification group:

1. Select **System > Email Notification**. The **Email Notification** window displays.

2. Under the **Email Notification List** tab, click **Add**. The **Add Email Notification Group** window displays.

3. Enter the information and click **Ok**.

Chapter 18: Viewing and Downloading Domain-Level Reports

The DSM comes with pre-configured reports that display system information. All reports can be downloaded and saved locally in CSV format.

Overview

All reports are under the **Reports** tab. The availability of reports depends on administrator type and privileges, and whether the administrator is currently in or out of a domain. You must be in a domain to access the domain-level reports.

Security and Domain and Security administrators must have AUDIT privileges to access the reports inside domains (for both global and restricted domains).

Viewing and Downloading Reports

To view a report:

1. Click the **Report** tab on the Management Console.

2. Click the name of the report.

To save the report as a CVS file to your local machine:

- Click **Download** at the top left of the report table.

Domain-Level Reports

The following security reports are available inside any domain:

- "Keys" on the next page

- "Key-Policy" on page 145

- "Policies" on page 145

- "Policy-Key" on page 146

- "Policy-Host" on page 146

- "Hosts" on page 147

- "GuardPoints" on page 147

- "Certificates" on page 148

- "Domain License Usage Summary" on page 149

- "License Usage by Hosts" on page 150

- "Host Registration Activities" on page 150

- "Hosts with GuardPoint Status" on page 151

The reports displayed depend on the roles assigned to a Security or Domain and Security Administrator, in addition to the 'Audit' role.

Keys

The Keys report is a table of keys available in the current domain.

The Keys report can generate more specific views through use of the following **Search** fields in the report's top panel.

- **Key Name**

  Enter a specific key name.

- **Source**

  Enter the source, either IP address or FQDN.

- **Key Flavor**

  Symmetric or Asymmetric from the drop-down menu.

- **Key Algorithm**

  AES128, AES256, ARIA128, ARIA256, RSA1024, RSA2048, RSA3072 or RSA4096 from the drop-down menu.

  > **Note**
  > 3DES is no longer available for new key creation. However, DSM will continue to support legacy keys created with 3DES.

- **Key Usage**

  Select All, Vault-Keys or Non-Vault-Keys from the drop-down menu.

- After completing the fields, click **Go**.

Table 18-1: Keys Report

Column Name	Description
Name	Name of the key
Source	The origin of the key, e.g., from DSM
Description	User defined description for the key
Algorithm	Algorithm used to create the key. Available options are: AES128, AES256, ARIA128, ARIA256, RSA1024, RSA2048, RSA3072, RSA4096
Key type	Indicates whether the keys are: **Stored on server:** Each time the key is needed, it retrieved from the DSM and downloaded to non-persistent memory on the host **Cached on Host:** Downloads and stores the key in persistent memory on the host. **Cached, Unique to Host**: Generated key is unique to the host and downloaded and stored in persistent memory
Flavor	Whether the key is Symmetric or Asymmetric
Creation Time	Time stamp of when the key was created. Format of the time stamp is YYYY-MM-DD HH:MM:SS:ms

Table 18-1: Keys Report (continued)

Column Name	Description
Time Expired	Time and date when the key expires or has already expired
Number of Policies	Number of policies that use this key
State	NIST Key State. The key can be in one of the following states: **Pre-activation**: Key is generated, but is not yet authorized for use **Active**: Key authorized to encrypt/decrypt information **Suspended**: Key has been suspended. It is currently not in use. **Deactivated**: A key whose active encryption period has expired, but may still be needed to perform cryptographic processing, is deactivated until it is destroyed. **Compromised**: The integrity, or secrecy, of the key is suspect, therefore, the key will not be used to apply cryptographic protection to information and it will be revoked. **Destroyed**: Key has been destroyed.

Key-Policy

The Key-Policy report lists keys and the policies that use the key. The Key-Policy report can generate a more specific view through use of the **Search** field in the report's top panel.

In the **Key Name** field, enter the specific key name you want to search for. Click **Go**.

> **Note**
> 3DES is no longer available for new key creation. However, DSM will continue to support legacy keys created with 3DES.

Key-Policy report

Table 18-2: Key-Policy Report

Column Name	Description
Key Name	Name of the key
Algorithm	Algorithm used to create the key; 3DES, AES128, AES256, ARIA128, ARIA256, RSA1024, RSA2048, RSA3072, RSA4096
Policy Name	Name of the policy that uses the key
Policy Type	The type of policy that uses the key, e.g., FS

Policies

The Policies report is a table of available policies. This report can generate more specific views through use of the following **Search** fields in the report's top panel.

- **Policy Name**

  Enter a policy name.

- **Policy Type**

  Select from available options, click **Go**.

Table 18-3: Policies Report

Column Name	Description
Policy Name	Name of the policy
Policy Type	The type of policy that uses the key, e.g., FS
Creation Time	Time stamp of when the key was created. Format of the time stamp is YYYY-MM-DD HH:MM:SS:ms
Keys Used	Number of keys used by this policy
Total GuardPoints	Number of GuardPoints using this policy
GuardPoints Enabled	Number of GuardPoints where this policy is enabled
GuardPoints Disabled	Number of GuardPoints where this policy is disabled

Policy-Key

The Policy-Key Report is a table of Key Names associated with Policy Names. This report can generate a more specific view through use of the **Search** field in the report's top panel.

> **Note**
> 3DES is no longer available for new key creation. However, DSM will continue to support legacy keys created with 3DES.

In the **Policy Name** field, enter the specific policy name you want to search for, click **Go**.

Table 18-4: Policy Key Report

Column Name	Description
Policy Name	Name of the Policy
Policy Type	The type of policy that uses the key, e.g., FS
Key Name	Name of the key
Algorithm	Algorithm used to create the key: AES128, AES256, ARIA128, ARIA256, RSA1024, RSA2048, RSA3072, RSA4096

Policy-Host

The Policy-Host report is a table of Host Names associated with Policy Names. This report can generate a more specific view through use of the **Search** field in the report's top panel.

In the **Policy Name** field, enter the specific policy name you want to search for, click **Go**.

Table 18-5: Policy Host Report

Column Name	Description
Policy Name	Name of the policy
Policy Type	The type of policy that uses the key, e.g., FS

Table 18-5: Policy Host Report (continued)

Column Name	Description
Host Name	IP address or FQDN of the host on which the policy is applied
OS Type	Operating System running on the host: for example, UNIX, Windows, Linux
GuardPoint Enabled	Yes (Y) or No (N)

Hosts

The Hosts report is a table of Hosts and the registration status of the agents installed on them. This report can generate a more specific view through use of the **Search** field in the report's top panel.

In the **Host Name** field, enter the specific Host name (IP address or FQDN) you want to search for, click **Go**.

Table 18-6: Hosts Report

Column Name	Description
Host Name	IP address or FQDN of the host on which the policy is applied
OS Type	Operating System running on the host: for example, UNIX, Windows, Linux, or OFFLINE if the host OS cannot be detected.
FS Agent Registration Status	Registration status of the agent on the host. Possible values are; Registered, Registration Allowed, Registration Not Allowed.
FS Agent Version	Version of the VTE (FS) Agent installed on the host
Key Agent Registration	Registration status of the agent on the host. Possible values are; Registered, Registration Allowed, Registration Not Allowed.
Key Agent Version	Version of the Key (VAE) Agent installed on the host
Last Policy Update	Time of the last policy update. Format of the time stamp is YYYY-MM-DD HH:MM:SS:ms
Number of Policies	Total number of policies on the host
Number of Enabled Policies	Total number of enabled policies on the host

GuardPoints

The **GuardPoints** report is a table of GuardPoints associated with each host. This report can generate more specific views through use of the following **Search** fields in the report's top panel.

- **Host Name**

  Enter a host name (IP Address or FQDN).

- **Guard Path**

  Enter the path for the folder location where the GuardPoint is installed.

- Click **Go**.

Table 18-7: GuardPoints Report

Column Name	Description
Host Name	IP address or FQDN of the host on which the GuardPoint has been created
GuardPoint Type	Type of GuardPoint created. Options are DIR, RAWDEVICE
Guard Path	GuardPoint location
Guard Enabled	Yes (Y) or No (N)
Policy Name	Name of the policy that applies to the GuardPoint

Certificates

The Certificates report is a table that describes certificates in the Certificate Vault. This report can generate more specific views through use of the following **Search** fields in the report's top panel.

- **Subject**

  Enter the certificate name.

- **Issuer**

  Enter the entity that issued the certificate.

- **Valid from**

  Enter the date in the following format MM/DD/YY or select a day from the calendar button at the right of the field.

- **Valid to**

  Enter the date in the following format MM/DD/YY or select a day from the calendar button at the right of the field.

- **Public Key Algorithm**

  Enter the type of key you want to search for: RSA or DSA

- Click **Go**.

Table 18-8: Certificates Report

Column Name	Description
ID	Identification number of the certificate
Subject	Name of the certificate
Issuer	Entity that issued the certificate
Valid From	Issuing date of the certificate
Valid To	Expiration date of the certificate
Public Key Algorithm	Algorithm used to sign the certificate, e.g., RSA or DSA
Public Key Size	512, 1024, 2048, 3072, 4096, etc.
Signature Algorithm	Algorithm for creating digital signatures, such as Sha1 with RSA, MD2 with RSA
CA	Certificate authority such as VeriSign, AT&T, GoDaddy, etc.

Table 18-8: Certificates Report (continued)

Column Name	Description
Type	Root, intermediate, leaf, unknown
Extended Key Usage	Server authentication and/or client authentication, code signing, email, time stamping, OCSP signing, unknown
CRL	Certificate revocation list
OCSP	Online Conflict Status Protocol
Format	Format of the vaulted certificate, such as PEM, DER, PKCS#7, PKCS#11, PKCS#12
With Private Key	Contains a private key? Y or N
File Name	Name of file
Description	Information that more fully describes the certificate

Domain License Usage Summary

The Domain License Usage Summary report summarizes all license usage in this domain. For more information about core hour licensing, see "Logical core hour usage" on page 31.

> **Note**
> Although other administrator types can view the DSM license in the Management Console, only DSM System administrators can upload a license file. If you are in a domain, the **Upload License File** button is disabled.

Table 18-9: Domain License Usage Summary Report information

Column Name	Description
Organization	Name of the organization responsible for this domain
Domain	Name of the domain set when the domain was created
Agent Type	FS (VTE), Key (VAE), KMIP
License Type	Perpetual, Term, or Hourly
Agents Licensed to Domain	Total number of licenses purchased in this domain
Agent Licenses Used	Total number of agents used
Logical Cores Used	Total number of core CPU hours used
Logical Core Hours Licensed to Domain	Total number of core CPU hours purchased
Logical Core Hour Licenses Used	Total number of cores in use
LDT Enabled Hosts	Total number of hosts using VTE Agent licenses with LDT enabled
Docker Enabled Hosts	Total number of hosts using VTE Agent licenses with Docker enabled
Term Expiration Date	Expiration date of a term license used in a domain, if applicable

License Usage by Hosts

The License Usage by Hosts report is a table of license usage by hosts in this domain. For more information about core hour licensing, see "Logical core hour usage" on page 31.

Table 18-10: License Usage by Hosts report information

Column Name	Description
Organization	Name of the organization responsible for this domain
Domain	Name of the domain set when the domain was created
Agent Type	FS (VTE), Key (VAE), KMIP
License Type	Perpetual, Term, or Hourly
Host Name	Name of the host
Agent Licenses Used	Number of agents used
Logical Cores Used	Total number of cores used
Logical Core Hour Licenses Used	Total number of core hours used
LDT Enabled Hosts	Total number of hosts using VTE Agent licenses with LDT enabled
Docker Enabled Hosts	Total number of hosts using VTE Agent licenses with Docker enabled

Host Registration Activities

The Host Registration Activities report is a table of host registrations and deregistrations for hosts under hourly licenses.

Table 18-11: Host Registration Activities report information

Column Name	Description
Organization	Name of the organization responsible for this domain
Domain	Name of the domain set when the domain was created
Agent Type	VTE (FS), VAE (Key), KMIP
License Type	Perpetual, Term, or Hourly
Host Name	Name of the host
Agent Licenses Used	Number of agent licenses used
Logical Cores	Number of logical cores
Logical Core Hour Licenses Used	Number of logical core hours used within this registration period
Registration Start (UTC)	Date host was registered
Registration End (UTC)	End date of the registration

Hosts with GuardPoint Status

The Hosts with GuardPoint Status report is a table of the total number of hosts with the status of their GuardPoints. This report helps with facilitating audits and other compliance metrics.

The Hosts with GuardPoint Status report features an overall status panel at the top that has the following fields:

- **Report ID:** The ID of the current report being generated

- **State:**

  - **QUEUED**

    When the request has been submitted and is waiting in the queue.

  - **STARTED**

    When the request is being processed.

  - **COMPLETED**

    When the report has been generated without errors found.

  - **CANCELLED**

    When the request in the queue has been canceled before completion by the administrator.

  - **ERROR**

    When errors have been found on the hosts or when a timeout or internal error occurred during processing.

> **Note**
> If State indicates an ERROR, check the fields for Total Hosts and Completed Hosts. If Total Hosts and Completed Hosts are equal, then the report table will show those host(s) indicating an ERROR condition. When Total Hosts and Completed Hosts are NOT equal, this indicates a problem with the DSM and host communication connections or another issue that should be troubleshooted.

- **Total Hosts**

  The total number of hosts being queried by the report.

- **Completed Hosts**

  The total number of hosts that have been queried to create the report.

- **Creation Time**

  The month, day, year, and hour and minute that the query was started.

- **Completion Time**

  The month, day, year, and hour and minute that the query ended with results.

- **Position in Queue**

  If multiple requests for reports have come in to the DSM for processing, this is the number this request is that reflects how many other requests are already lined up.

- **Total Tasks**

  If multiple requests for reports have come in to the DSM for processing, this is the total number of requests that are in the queue.

- **Remaining Tasks**

  If multiple requests for reports have come in to the DSM for processing, this is the number of requests that are yet to be worked on.

- **Last Update Time**

This field indicates the month, day, year, hour, minute (and AM/PM) that the DSM report database was last updated.

Report Tasks

- To start a query and generate a report, click **Generate Report**.

- To update the DSM database, click on **Refresh**.

- To cancel a query report generation request, click **Cancel**.

- To download the results of this report, click **Download** to produce a CSV text file.

Table 18-12: Hosts with GuardPoint Status Report

Column Name	Description
Report ID	Unique identifier for this report view
Host Name	IP address or FQDN of the host
Host Description	(Optional) User-entry field for more clearly defining the host during provisioning
OS Type	Operating System installed on the host
Port	Port number used for DSM <-> Agent communication
License Type	Type of license used by the host; Perpetual, Term, or Hourly
One-Way Enabled	The agent was registered with One-Way Communication enabled between the Agent and the DSM (Yes—Y) or not (N).
FS Agent Registration Status	Registered, Allowed, Not Allowed
FS Agent Version	Version of the VTE (FS) Agent installed on the host
Key Agent Registration Status	Registered, Allowed, Not Allowed
Key Agent Version	Version of the VAE (Key) Agent installed on the host
Last Policy Update	Timestamps are in the form *YYYY-MM-DD HH:MM:SS:ms*
Docker Image ID	Unique identifier for the Docker image. This column is displayed only if you have a VTE Agent license with a Docker Extension
Docker Container ID	Unique identifier for the Docker container. This column is displayed only if you have a VTE Agent license with a Docker extension
GuardPoint Type	DIR, RAWDEVICE
Guard Path	GuardPoint location
Policy Name	Name of the GuardPoint policy
Guard Enabled	Indicates whether or not the GuardPoint is enabled. Values; Y or N
GuardPoint Status	Up, Down, Server Pending, Agent Pending, Error, or Unavailable
Transformation Status	Indicates the transformation status of a rekey operation. Displayed only if you have a VTE Agent license with an LDT extension and an LDT policy applied to a GuardPoint

Table 18-12: Hosts with GuardPoint Status Report (continued)

Column Name	Description
Transformation Progress	If a rekey operation is underway, this indicates the percentage of the operation completed. Displayed only if you have a VTE Agent license with an LDT extension and an LDT policy applied to a GuardPoint
Estimated Rekey Completion Time	Provides an estimate of the time it will take to transform the data in the GuardPoint based on the available resources and the size of the data.
Transformation Error	Indicates whether there was an error in the transformation operation.
Last Transformation Start Time	Date and time the last data transformation started.
Last Transformation Completion Time	Date and time when the last data transformation was done.
Total Files to be Transformed	The total number of files in that GuardPoint to be transformed by the policy.
Totals Files Transformed	Total number of files in that GuardPoint transformed by the policy. If the rekey operation is successful, this should match the number in the Total Files to be Transformed field.
Total Bytes to be Transformed	Total number of Bytes to be transformed.
Total Bytes Transformed	Total number of bytes transformed. If the rekey operation is successful, this should match the number in the Total Bytes to be Transformed field.
Total Files Skipped	Indicates the number of files skipped during data transformation. If any files were skipped, use the `voradmin ldt skip` command from the CLI on the host to see why these files were skipped.
Error Reason	If the Guard Status reads "Error" and Guarded reads "N", the reason will be indicated here

Chapter 19: Viewing DSM Preferences and Logs

Viewing DSM Preferences

Preferences for viewing the various windows and panels on the Management Console are set by the DSM System Administrator, as a DSM Domain Administrator you can still set some viewing preferences within the domains you are authorized to access.

You can set Log viewing preferences from the Edit Host page for the available agent log tabs. You can also configure Docker log settings from the Docker Log tab. Docker support is a separately licensed feature, see "Enabling Docker Support" on page 239 for more information about this feature.

Setting Log Preferences on a Host

Log settings for the VTE Agent (FS Agent Log) are configured at the System level on the DSM. These settings are inherited by all the hosts on the DSM. However, you can fine those tune log settings for a specific host, and those settings will override the system settings.

Navigate to the Hosts page and click on the name of the host in the **Host Name** column for which you want to set log viewing preferences. Click the agent log that you want to configure (e.g., FS Agent Log, Key Agent Log, Docker Log). From this page, you can set the following parameters for the host:

1. **Message Type**

   ° **Management Service**

   Logs messages that are related to the agent and VMD process server interaction in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is INFO.

   ° **Policy Evaluation**

   Logs messages that are related to policy evaluation in the agent log. Set the log message level to desired setting. The default log message level is ERROR.

   ° **System Administration**

   Logs messages that are related to system level events. The default log message level is ERROR.

   ° **Security Administration**

   Logs messages that are related to security related events. The default log message level is INFO.

2. **Message Destination**

   Log Messages can be stored in several locations.

   ° **Log to File**

   Send log messages to the /var/log/vormetric/vorvmd_root.log file of a UNIX host, or a Windows equivalent, such as \Documents and Settings\All Users.WINDOWS\Application\ Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log.

   ° **Log to Syslog**

   Send log messages to the syslog server for a UNIX host. If a syslog server is not configured, it is sent to the host 'messages' file, such as /var/adm/messages. On a Windows host, the messages are sent to the Event Viewer (Application events).

   ° **Upload to Server**

Upload to the DSM and display in the Management Console Logs window.

- **Level**

  Sets the level of error messages to be sent.

- **Duplicates**

  Allow or suppress duplicate messages:

- **Allow**

  All duplicate messages of the corresponding Message Type are captured and displayed in the log.

- **Suppress**

  Messages of the corresponding Message Type will follow the configured Threshold as to how many times duplicate messages are sent to the DSM during the given Interval.

3. **File Logging Settings**

   - **Maximum File Size (bytes)**

     The agent starts a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.

   - **Delete Old Log Files**

     Select this check box to delete old FS agent logs. This check box works in conjunction with the **Number of Old Log Files to Keep** field. For example, Select this check box and enter 3 as the **Number of Old Log Files to Keep** value. After 3 logs are generated, the first log, log1, is deleted and a new log, log4, is created. If you do not Select this check box, log files will continue to accumulate in the server database and you will have to remove them manually.

   - **Number of Old Log Files to Keep**

     Appears only when you select **Delete Old Log Files**. Specifies the maximum number of agent log files to leave in the server database. This text-entry box is only displayed when the **Delete Old Log Files** check box is enabled. The default is 5.

   - **Maximum Stored Audit Log Entries (rows)**

     Previously, combined agent and server logs were purged when they exceeded 10,000 rows. Now, they are purged separately, and you can separately set the maximum number of server and agent log entries to store. The DSM will purge the logs when it reaches the maximum number of rows set in the Server logs field and in the corresponding Agent logs field.

     The minimum number of log entries is 5000; maximum number is 100,000. Note that the maximum limit may be restricted by storage capacity.

4. **Syslog Settings**

   - Local: Send Syslog messages to the local machine.
   - Server (1, 2, 3, 4): Enter the hostname of the Syslog server
   - Protocol: UDP or TCP
   - Message Format: Specifies the format of the message; Plain Message, CEF, or RFC5424.

5. **Upload Logging Settings**

   - **Maximum Number of Messages to Upload At Once**

     Limits the number of messages sent to the DSM at one time. When the specified number of log entries is reached, those entries are uploaded to the DSM. The default is 1000.

   - **Upload Messages At Least Every (seconds)**

     The maximum interval to wait before the agent is to upload messages to the DSM. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.

○ **Upload Messages At Most Every (seconds)**

The minimum interval to wait before the agent is to upload messages to the DSM. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.

○ **Normal Time Out (seconds)**

The maximum interval of time the agent is to wait for the DSM to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the DSM within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.

○ **Shutdown Time Out (seconds)**

The maximum interval of time the agent is to wait for the DSM to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.

○ **Drop If Busy**

Select to slow log message generation and drop log files during periods of extreme logging.

6. **Duplicate Message Suppression Settings**

○ **Enable Concise Logging**

When enabled, audit log messages are reduced. This option is disabled by default. Instead of logging messages for each file system operation, only the following types of audit messages are logged;

- Only one audit message for each read or write activity is logged at the start of that activity.

- Audit messages for reading file status information and setting file attributes (and extended attributes) are not logged.

- Audit messages for directory open, close and read attributes are not logged.

These settings can be customized on each host and the host setting will override the system level settings. Note that this feature is not available for VTE versions prior to v6.0.

○ **Threshold**

Used when the Duplicates value is set to Suppress. Specifies the maximum number of duplicate messages the agent is to send to the DSM within the amount of time specified by the Interval parameter. The default is 5 messages. The maximum is 100.

○ **Interval**

Used when the Duplicates value is set to Suppress. Specifies the time period in which the number of duplicate messages, specified by Threshold, can be uploaded to the DSM. Once Interval is exceeded, the count specified by the Threshold parameter starts again. The default is 600 seconds (10 minutes). The maximum is 3600.

> **Note**
>
> We recommend turning on **Log to File** or **Log to Syslog** instead of **Upload to Server** for INFO and DEBUG levels. And, for general day-to-day operation, we recommend enabling and setting only ERROR Level (so that only ERROR, WARNING, and FATAL log entries are received). Setting Upload to Server to INFO or DEBUG level for policy evaluation can degrade DSM performance.

Configure Docker Log Settings

With the introduction of Docker support, you can now configure log settings for Docker images and containers. Docker logs evaluate GuardPoint policies.

1. Log on to the Management Console and switch to a domain or log in as a local Domain and Security administrator with a Host role.

2. Navigate to the Hosts page.

3. Click the name of your Docker host in the **Host Name** column, the Edit Host page opens.

   Enter the following information in the **Configure Docker Log Setting** panel:

   ○ **Docker Image/Container**

   Click **Browse** to select an image or container from the Docker host. If you select an image the **Docker Image ID** field displays the image ID. If you select a container, the **Docker Image ID** field displays the image from which the container was spawned and the **Docker Container ID** displays the container ID. You can use these IDs to search for Docker specific logs on the Logs page later.

   ○ **Policy Evaluation Level**

   Select a log message level. For more information about log levels, refer to the *Administrators Guide*.

   ○ **Policy Evaluation Duplicated**

   You can choose to suppress or allow duplicate messages. Select SUPPRESS or ALLOW, the default is SUPPRESS.

4. Click **Ok**. The Policy Evaluation settings are saved in a table under the **Configure Docker Log Setting** panel.

Docker log messages are displayed on the Logs page. To search for Docker specific log messages:

1. Navigate to the Logs page.

2. Enter the following information in the **Search** panel:

   ○ **Log Type**

   Select whether you want to display logs from both the DSM and the agents, only the DSM, or only the agents. The default is All, which means from both DSM and agents.

   ○ **Source**

   Enter the hostname of the DSM or agent for which you want to return log files.

   ○ **Last Refreshed**

   Displays the date and time of when the displayed log files were last refreshed. Format is YYYY-MM-DD HH:MM:SS

   ■ Message Contains

   Type in text string that you want to search for in the log messages.

   ■ Docker Host

   Click **Browse** to select the Docker Host for which you want to return log files.

   ■ Docker Image/Container

   Click **Browse** to select an image or container for which you want to display logs.

   ■ Docker Image ID

   Displays the ID for the selected Docker image.

   ■ Docker Container ID

   Displays the ID of the selected Docker container.

3. Click **Go**. The relevant logs are displayed in the table under the **Search** panel.

Viewing Logs

The entries displayed in the Message Log depend on the DSM administrator type (System, Domain, Security, All), the domain in which that administrator is working, and, for Security Administrators, that administrator's role (Audit, Key, Policy, Host, Challenge & Response, Client Identity).

A Domain Administrator cannot view the log entries that can be viewed by a System or Security (and vice versa) administrator. By design, entries exported to a Syslog log file will have gaps in the number sequence depending on which domains and roles the DSM Administrators are actively logging.

The Domain Administrator sees log entries such as Domain Administrator and Security Administrator logins, SSL handshaking, and policy evaluation.

Log entries are displayed in the Management Console based on the current administrator type and the domain in which the administrator is working. However, all this log information combined is available in the server.log file on the DSM.

Chapter 20: Creating and Configuring Signature Sets

File signing checks the authenticity and integrity of executables and applications before they are allowed to access GuardPoint data. When you initiate file signing on the DSM, the VTE Agent calculates the cryptographic signatures of the executables that are eligible to access GuardPoint data. Files are individually signed as part of a set and the set is configured in a policy that defines the processes to allow.

When an executable tries to access a GuardPoint, the `secfs` service checks the fingerprint, a SHA-2 (Secure Hashing Algorithm) message digest, of the executable against the fingerprint stored in the DSM database. If they match, the executable's authenticity is verified and it can be allowed to access protected data. A hostile or compromised executable, such as a Trojan application, malicious code, or rogue process, with a missing or mismatched signature, is denied access.

Once a set of files to be signed is created, the executables are signed on a selected host and a copy of each signature is stored on the DSM. This is done as a background process on the selected host. The time it takes to complete signing depends upon the number of files to be signed, the response time of the host system, and other load factors. The completion status is indicated in the Signature Sets window.

Completed signature sets are configured in a (VTE Agent) policy so that not only are the executables attempting GuardPoint access identified, but their signatures are checked to ensure that they had not been compromised.

Creating Signature Sets

A signature set is a collection of file names and/or directory names. You can enter the full path of files and directories manually or use the browser to locate and select them. Specify a directory to sign all the files in that directory and all the subdirectories that it may contain.

Signing many files can take a while. To shorten processing time, verify that the files and directories in the signature set exist. It takes longer to process non-existent files. If they do not exist, we recommend that you delete them as sources from the signature set.

By default, a generic error message is generated and displayed in the Logs window about a non-existent source being detected; the name of the offending file or directory is not specified. However, it is specified in the agent log on the host. To identify the offending file or directory, open the agent log file `vorvmd.log/vtray > View > File System > Log` on Windows systems. Look for "`Number of failed files`" to determine how many files were affected and "`is invalid for the signature request`" to identify the files that were not signed.

To create a signature set:

1. Log on to the Management Console as an All administrator or one with `Host` role permissions.

2. From the menu bar, select **Signatures**.

    The Signature Sets window opens. All configured signature sets are displayed.

Figure 20-1: Default Signature Sets window
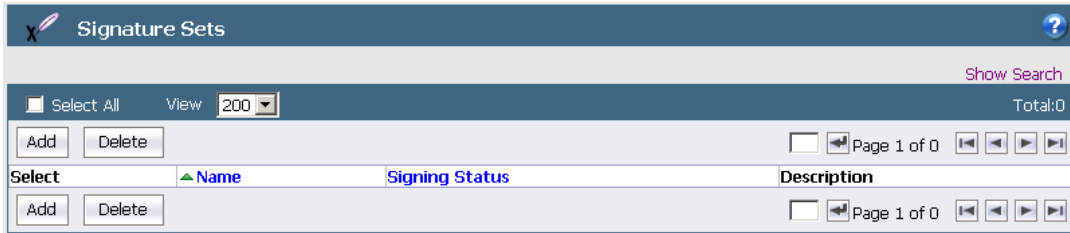
3. (Optional) Display only specific signature sets by entering all or part of a signature set name, and select the completion status, in the **Search** panel to display a subset of all signature sets in the DSM database.

   The **Show Search** label located below the **Signature Sets** banner opens the **Search** panel. You can enter a string and/or limit the search to sets with a specific completion status, and then click **Go** to display only those signature sets that match the search criteria. Click **Hide Search** to conceal the **Search** panel. The **Search** panel is not displayed to reduce graphic size.

4. Click **Add**. The **Add Signature Set** window opens.

5. Enter a name to assign the signature set in the **Name** text-entry box.

   Enter a unique string for the signature set name. The string you enter cannot exist in the current domain nor any other domain.

   This field is mandatory. The name must consist of alpha-numeric characters; starting with an alphabet character. The only non-alpha-numeric characters allowed are underscore (_) and dash (-). The maximum number of characters is 64.

6. (Optional) Enter a brief phrase or string in the **Description** text-entry box to make signature set identification easier. The maximum number of characters is 256.

Figure 20-2: Adding a signature set

7. Click **Ok**.

   The **Signature Sets** window reopens and displays all the signature sets, including the one you just created.

Figure 20-3: Unsigned signature set
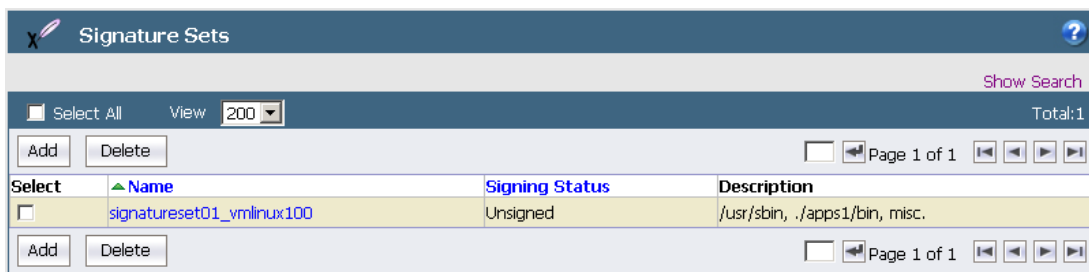
By default, the signature set has an Unsigned status.

8. Add the files to be signed and directories whose files are to be signed to the signature set.

Adding files to a set

You can enter the full path of files or directories manually or use the browser to locate and select the files. It is quicker and easier to manually enter the paths of files; however, manual entry is prone to typographic errors and incorrect paths. Browsing can take longer but it ensures that the files exist and paths are entered correctly.

A cryptographic hash is created for each file in a signature set that meets a specific criteria. It would take longer, bloat the DSM database, and reduce performance to sign all the files in a signature set, especially when the set consists of top-level directories. The criteria is listed below. All other files are skipped.

- On a Windows host, all the files in the signature set that are inside a GuardPoint are signed. Only the compiled Windows executable files in the signature set that are located outside a GuardPoint are signed.

- On UNIX, it makes no difference if the files are inside or outside a GuardPoint. Only the files in the signature set with one or more of the execute bits (for example, `-rwxrwxrw-`) set on a UNIX host are signed.

- File extension has no impact. Files like `.bat` and visual basic programs on Windows, and files that end with `.so` on UNIX, are skipped.

Each instance of a file that has been copied to a different location or to a different name will have the same signature. This can be convenient way to detect duplicate files on your system.

To add files and/or directories to the signature set:

1. Select **Signatures** in the menu bar.

2. Click the name of a signature set in the **Name** column.

3. Click the Source tab in the **Edit Signature Set** window.

4. Select the host that contains the files to be signed.

   You must specify a host before you can browse for sources or initiate the signing process. You cannot enter the host name manually in the **Host** text-entry box.

   a. Click **Select** next to the **Host** test-entry box.

      The **Select a host to continue** window opens. All configured hosts are displayed and available for selection.

   b. Enable the **Select** radio button for the host that contains the files to be signed.

      Do not click the name of a host—that will open the **Edit Host** window.

   c. Click **Select** on the bottom of the window.

      The **Edit Signature Set** window is redisplayed and includes the name of the selected host in the **Host** text-entry box. The files on this host will be signed.

5. Click **Add**.

   The **Add Sources** window opens. Do one of the following:

   ○ Enter file names and directory paths manually in the **Sources** scroll-list

   ○ Select files and directories by browsing the host

   ○ Do a combination of the two

   Adding sources is cumulative.

6. To add sources manually:

   a. Enter the full paths to files and directories in the **Sources** scroll-list (Enter one file or directory per line).

The asterisk can be used in a limited capacity as a wildcard character in file name searches. Place it somewhere in a file name string. Any executable or application file in the specified directory, and in every subdirectory, that matches the string will be located and can be signed. The wildcard is ignored when used in directory names. Directories that would normally match the wildcard are ignored. Check the logs for skipped files and directories.

A trailing slash (/) or backslash (\) at the end of directory paths is optional.

Figure 20-4: Manual source entry

If you plan to add sources using both the manual and browser methods, be sure to click **Ok** before you open the browser. If you do not, all the sources that you had manually entered in the **Sources** scroll-list will be deleted and only the browser-selected sources will be listed. Inversely, you can browse for sources first and then manually add additional sources later without losing browser-selected and manually-entered sources.

b. Click **Ok**.

The **Source** tab displays the added sources.

7. To add sources using the browser:

a. Display the **Source** tab of a signature set.

b. Click **Add**.

The **Add Sources** window opens.

c. Click **Browse**.

The **Remote File Browser** window opens.

The **Type** scroll-list is hardwired to `Directory and File`.

The **Start Directory** text-entry box displays the top-level directory that is appropriate to the platform type: Windows (\) or UNIX (/).

d. (Optional) Enter a start point in the **Start Directory** text-entry box.

You cannot browse above the **Start Directory**. Enter a start point that is higher in the directory hierarchy than all the directories and files that you want to select, or you will have to re-enter start points to locate and select the desired files. The default is the top-level, either slash or backslash.

e. Click **Go** or, with the mouse cursor in the **Start Directory** text-entry box, press the `<Enter>` key.

f. Navigate to and select the desired files.

Click the plus symbol (+) next to a folder to display the next level of the directory hierarchy. Click the minus symbol (-) to collapse the hierarchy. Click a folder or file name to select that directory or file.

Figure 20-5: Adding files/directories to the set, browser method
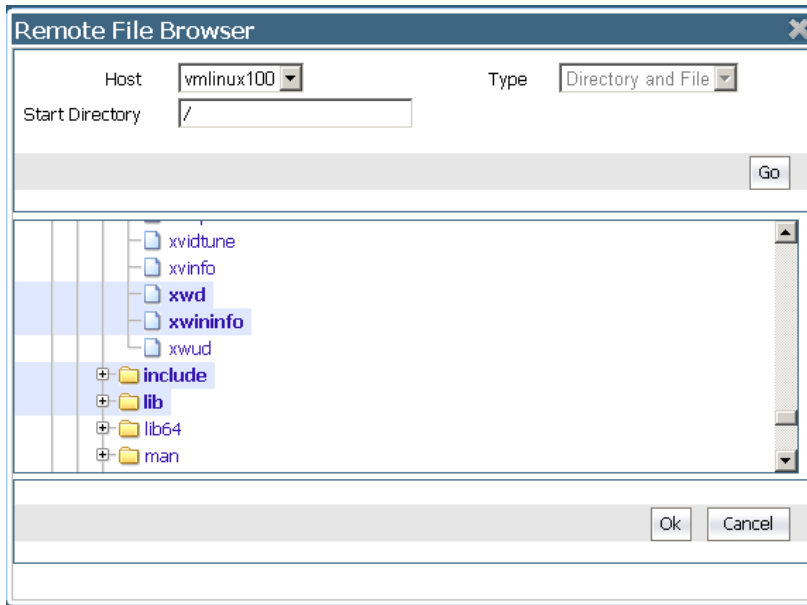
Single-click one or more files and/or directories. When you select a directory, all the files in all the subdirectories are also added to the set.

g. Click **Ok**.

The **Source** scroll-list displays the new additions.

Figure 20-6: Displaying browser-selected sources
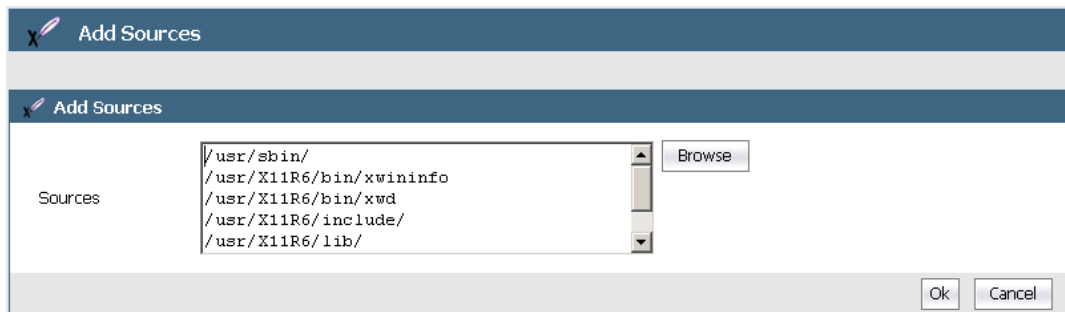
You can make changes by single, double, or triple clicking a source in the **Sources** scroll-list. Single-click to add or delete individual characters. Double-click to select a word. Triple-click to select an entire line.

h. Click **Ok**.

The **Source** tab displays the added sources.

Figure 20-7: Signature set with files
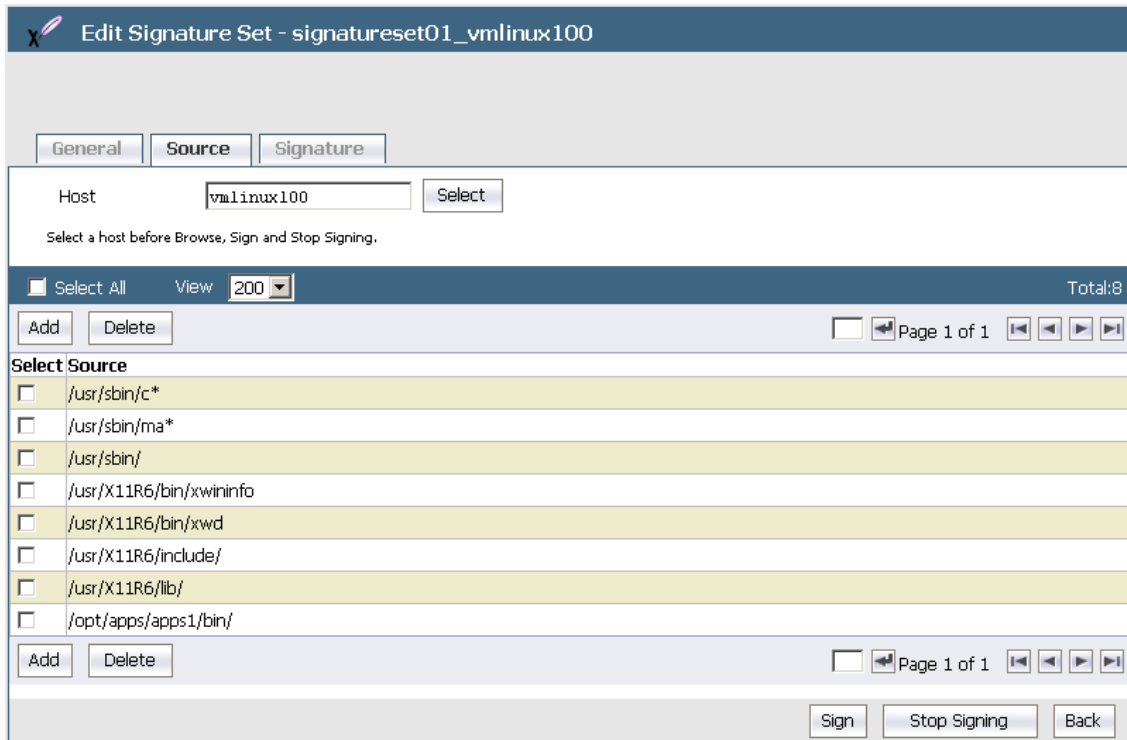
Signing Files in a Signature Set

Signing involves calculating a hash value for a file and storing the value on the DSM. Later, when a policy checks signatures, the signature of the process or executable accessing the GuardPoint is calculated and compared against the value in the DSM. If the two values match, the process or executable satisfies the Process requirement of a policy and may be granted access to the guarded data.

> **Note**
> If the executable itself is volatile, or subject to frequent change, it may not be worthwhile to use a file signature as a criteria in a policy because you have to re-sign the executable after each change. If the volatile executables are members of a large signature set, it can take a while to re-sign the files because the signature of every file in the signature set is recalculated. If the volatile files are few, it might be quicker to add the volatile files to a different signature set and sign that set, rather than re-sign all the files in the original signature set.

To sign the files in a signature set:

1. Log on to the Management Console as an All Administrator or one with `Host` role permissions.

2. Select **Signatures** in the menu bar.

   The **Signature Sets** window opens. Note the completion status of the desired signature set in the **Signing Status** column.

3. Click the name of the signature set in the **Name** column.

   The **Edit Signature Set** window opens to the **General** tab.

4. Click the **Source** tab.

> **Note**
> Do not enable any of the **Select** check boxes. The **Select** check boxes are used only to delete sources from the set.

5. If not already selected, specify a host that contains the files to be signed.

   You must specify a host before you can start signing. You cannot enter the host name manually in the **Host** text-entry box.

   a. Click the **Select** button next to the **Host** test-entry box.

      The **Select a host to continue** window opens. All configured hosts are displayed and available for selection.

   b. Enable the **Select** radio button for the host that contains the files to be signed.

      Do not click the name of a host—that will open the **Edit Host** window.

   c. Click **Select** on the bottom of the window.

      The **Edit Signature Set** window is redisplayed and includes the name of the selected host in the **Host** text-entry box. The files on this host will be signed.

6. Click **Sign**.

   The time for this process to complete depends on how many files are being processed. The percentage of files in the signature set that have been signed is indicated in the status bar on the **Edit Signature Set** window. Also, you can view signing status in the **General** tab.

Figure 20-8: Signing progress indicator
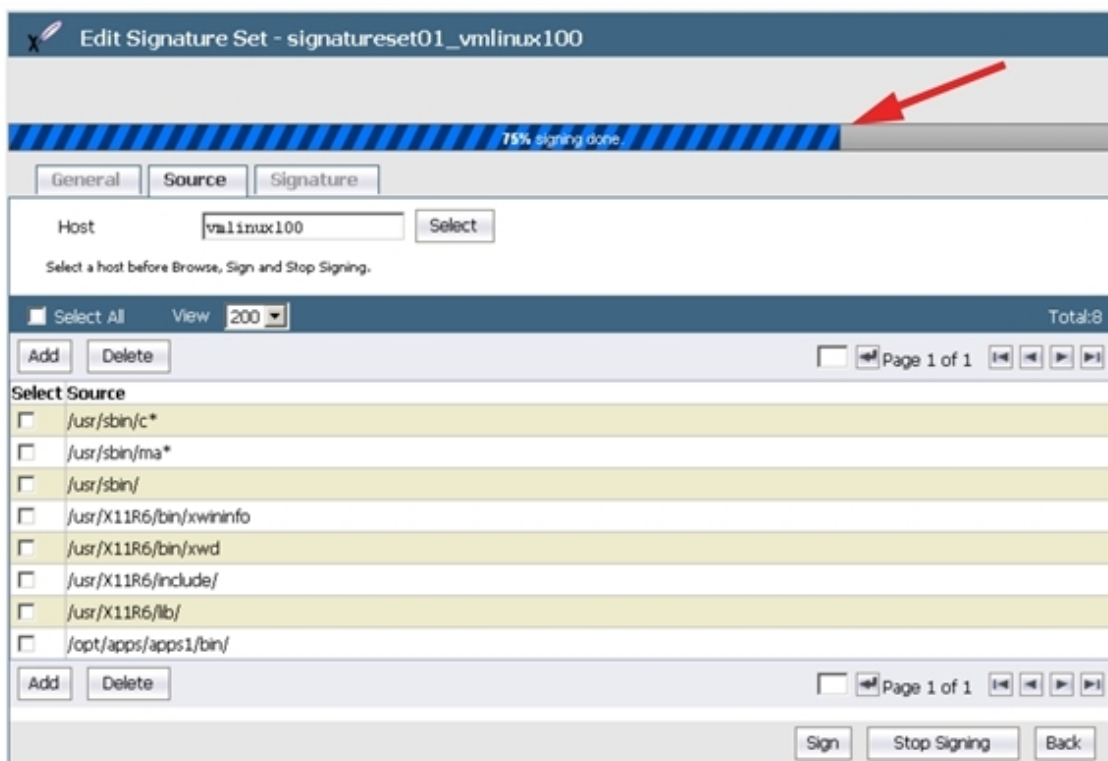
The signing status for the set in the **Signature Sets** window and the **Edit Signature Set** window, **General** tab, is IN_PROGRESS. Signing occurs as a background process, so you can use the Management Console for other administrative functions during this operation.

7. When signing completes, display the **General** tab and note both the signing status and percentage of completion. These should be FINISHED and 100 percent respectively.

8. View the resulting file: signature pairs in the **Signature** tab.

Using signature sets in a policy

Policies can be configured to identify the executables trying to access GuardPoint data and to verify that the executables themselves are unchanged since they were signed. You must, however, anticipate the effect of encryption on file signatures.

When a file inside a GuardPoint is copied to a location outside the GuardPoint, the two files will have different signatures because the file in the GuardPoint is encrypted and the file outside the GuardPoint is not.

Encryption makes the two files different, and the `vmd` process does not decrypt guarded files before checking their signatures. This means that when you rekey guarded files their signatures also change, and you must re-sign the files that use the signatures of those files in Process sets.

If both files, the one inside the GuardPoint and the one outside, must access GuardPoint data, add both files to the signature set and sign them. If encryption is not applied, both files will have the same signature and a signature mismatch should not occur.

Checking the agent logs if signing fails

If signing fails, or you want more information about the signing process such as which files were skipped, check the agent logs.

Messages are logged to:

(UNIX) `/var/log/vormetric/vorvmd_root.log` with **Log to File** enabled, and to `messages` with **Log to Syslog/Event Log** enabled.

(Windows) `\Documents and Settings\ ...\agent\log\vorvmd.log` with **Log to File** enabled.

In the Management Console, look for messages like:

```
COM0591W: The agent at host sys-techpub2 failed to generate the signature on this file
E:\apps\lib\dataxform_auto_config. Please check the VMD log for the cause.
```

In the host `vorvmd_root.log/vorvmd.log` file, look for messages like:

```
[VMD] [WARN ] [3732] [VMD3824W] Failed to create the signature for E:\apps\lib\dx1\aa_
dir\dataxform_auto_config for the signature request 9, error code 3
```

Table 20-1: Error Codes

Code	Description
0	System is okay.
1	Does not exist.
2	DO NOT USE.
3	Invalid argument.
4	Operation not supported.
5	Out of memory.
6	No space left on device.
7	Timeout reached.

Table 20-1: Error Codes (continued)

Code	Description
8	I/O error.
9	Interrupted.
10	Permission denied.
11	Too many keys in key group.
12	Error in soap rpc layer.
13	Returned buffer is partially full.
14	Given target buffer is too small.
15	Unable to compress buffer.
16	Internal test failed.
17	Overflow.
18	Error setting up logging.
19	Overloaded error.
20	Server responded to a file upload with "bad request".
21	Unknown user name or bad password.
22	The directory service is not available.

Checking Signing Completion Status

To display the completion status of a signature set:

1. Select **Signatures** in the menu bar.

2. Click the name of a signature set in the **Name** column.

   The **Edit Signature Set** window opens. The **General** tab is displayed.

Figure 20-9: The Edit Signature Set window, General tab
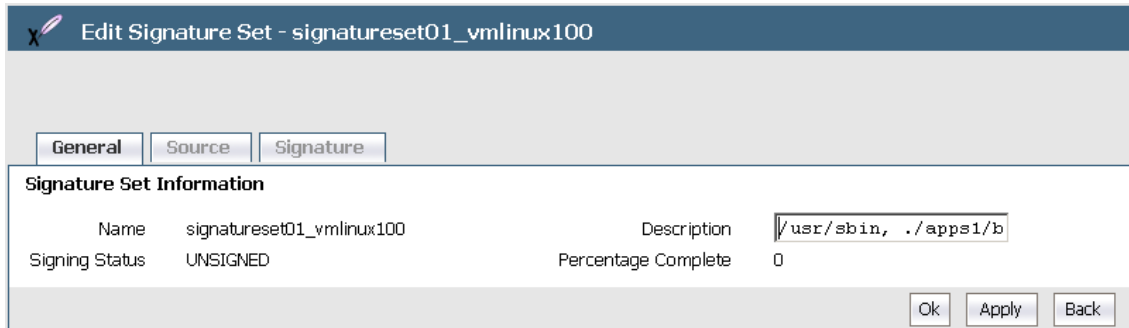
The **General** tab displays the file signing status and the percentage of files in the signature set that have been signed. The parts of the **General** tab are described below.

Table 20-2: General tab information for the Signatures page

Field	Description
Name	The name of the signature set.
Description	(Optional) Descriptive string to simplify set identification.
Signing Status	A signature set can be in one of five states: UNSIGNED—there may or may not be files in this signature set. If there files in the set, no attempt has been made to sign them. IN_PROGRESS—the DSM is actively signing the files in the set. This can take a while depending upon system load, accessibility, and the number of files being signed at one time. FINISHED—all the files in the set have been successfully signed and the set is ready to be used. FINISHED_WITH_WARNING—the VTE Agent was able to process each file in the signature set, but that one or more files in the set could not be signed. Possible causes are missing files or inadequate access permissions. Check the vmd log for details — vorvmd_root.log on UNIX and **Event Viewer > Vormetric Encryption Expert** on Windows. ABORTED—signing had been started but was stopped before completing.
Percentage Complete	Indicates the percentage of files that had been signed relative to the total number of files in the set.

The only field you can modify on the **General** tab is **Description**.

Stopping Signing

You can interrupt signing at any time by clicking **Stop Sign** in the **Source** tab. You are not prompted to verify your choice.

It takes a while for signing to stop. Once it does, the **Signing Status** displayed in the **General** tab is set to ABORTED and the **Percentage Complete** indicates the percentage of files that have been successfully signed. All the files that were signed remain intact and can be viewed in the **Signature** tab.

Re-Signing Files in a Signature Set

Signatures are computed values and, unless the files in a set have been compromised or modified, the signature should always be the same for a given file. If files have been modified, they must be resigned so that their signatures match the signatures in the DSM.

To re-sign files, you can:

- Create and sign a new signature set with the files and the directories to be signed. If there are only a few files and directories in the set, this can be the easiest solution.

- Create and sign a new signature set with the paths of the directories that contain the files to be signed. If there are only a few files in the directories, or most of the files in the directories need re-signing, this can be the easiest solution.

- Open and re-sign the same signature set that was originally used to sign the files. If you have the time, or the signature set is small, this can be the easiest solution.

Signatures are computed values. It takes a long time to compute the signature of every file in a large signature set. It takes longer to re-sign the files in a set than it does to sign the files the first time because of the additional handling

required to update information on the DSM. If you plan to re-sign many files, on the scale of hundreds of thousands, it can be quicker to delete the signature set that was initially used to sign the files, recreate the signature set from scratch, and sign the members of the signature set as if it were a new signature set.

Displaying Signatures, Detecting Duplicate Files

You can display up to 200 files of a signature set on one page, or you can display a subset of the files across multiple pages. You can specify a search criteria to display a subset of the signed files. The search criteria can be a case-insensitive string that is in all or part of a file name, and/or it can be the signature itself. String search is a convenient way to display signed files with a specific extension, such as ".exe", or signed files with names that contain a specific string, such as "lib".

Displaying files based on their signature is one way to locate identical files, regardless of name differences. If you search using a signature, you must enter the entire signature.

Figure 20-10: Displaying identical files based on their signature—Same file in different locations and with different names
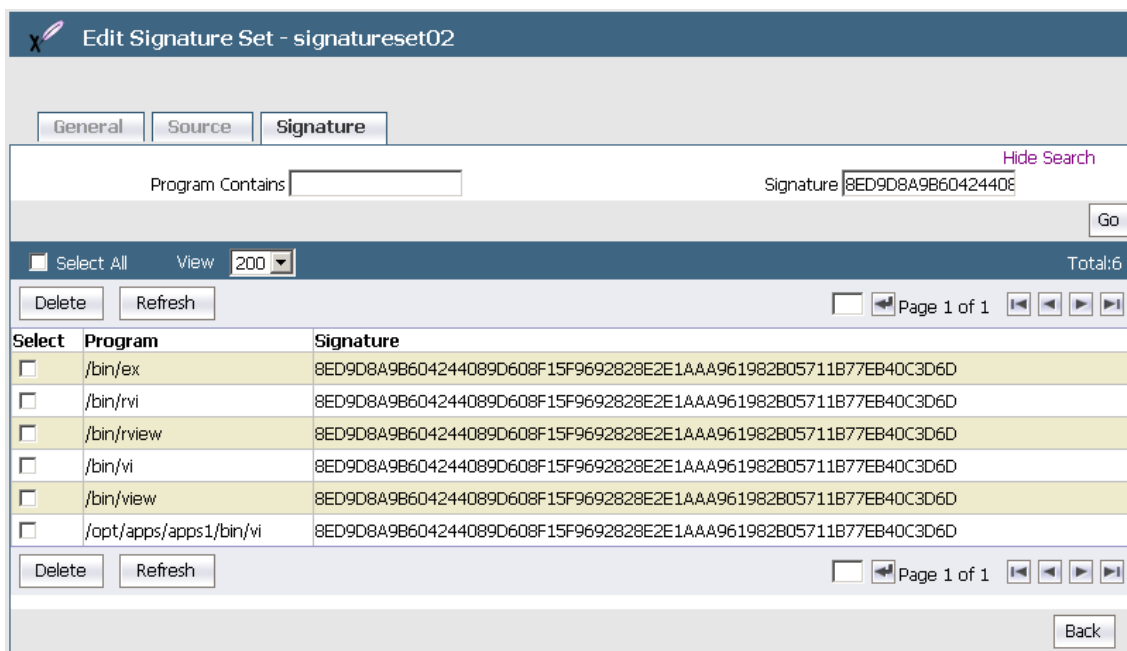
Displaying Specific Signed Files in a Signature Set

To display specific signed files in a signature set:

1. Log on to the Management Console as an All administrator or one with `Host` role permissions.

2. Select **Signatures** in the menu bar.

   The **Signature Sets** window opens.

3. Click the name of a signature set in the **Name** column.

   The **Edit Signature Set** window opens.

4. Click the **Signature** tab.

   By default, all signed files in the set are displayed in alphanumeric order.

5. Click **Show Search** to display the **Search** panel.

6. To search for files with a specific string in their names, enter the string in the **Program Contains** text-entry box.

   Enter all or any part of the desired file name. Case does not matter. All files that contain the string and, if configured, match the signature in the **Signature** text-entry box, will be displayed.

7. To search for files with a specific signature, enter the entire signature in the **Signature** text-entry box.

   Enter the entire signature. Unlike the **Program Contains** text-entry box, the **Signature** text-entry box is case-sensitive. All files that have the same signature and, if configured, match the string in the **Program Contains** text-entry box, will be displayed.

8. Click **Go**.

Deleting Signatures from a Set

Delete signatures from a signature set when you want to continue to use the signature set to authenticate processes, and you no longer want to authenticate the files that you are deleting from the set.

You can delete signatures individually or one page at a time.

To delete individual signatures from a signature set:

1. Log on to the Management Console as an All administrator or one with `Host` role permissions.

2. Select **Signatures** in the menu bar.

   The **Signature Sets** window opens.

3. Click a signature set in the **Name** column.

   The **Edit Signature Set** window opens.

4. Click the **Signature** tab.

   The signed files in the signature set are displayed, along with their signatures.

Figure 20-11: Signed signature set

5. If you plan to delete many file signatures, set the **View** number high because file selection applies only to the files on the current page. Files are deselected when you go to another page.

6. Select the file signatures to be deleted.

   You can click the **Select** check boxes of individual files on the current page or you can enable the **Select All** checkbox to select all the files on the current page.

7. Click **Delete**.

   A dialog box opens and prompts you to verify that you want to delete the selected signatures.

8. Click **OK**

   The signatures are removed from the DSM database.

Deleting Signature Sets

You can delete individual signatures within a set or whole signature sets.

You cannot delete an active signature set. If it is defined in a VTE Agent policy, it must be removed from the policy before the set can be deleted from the DSM.

To delete one or more signature sets:

1. Log on to the Management Console as an All administrator or one with `Host` role permissions.

2. Select **Signatures** in the menu bar.

   The **Signature Sets** window opens. All configured signature sets are displayed.

3. Enable the **Select** checkbox of each signature set to be deleted.

4. Click **Delete**.

   A dialog box opens that prompts you to verify that you want to delete the selected signature sets.

5. Click **OK**.

Chapter 21: Managing Keys

You can use the DSM to create agent keys, as a secure centralized repository for storing and retrieving third-party encryption keys, and to create key templates.

Overview

Encryption keys are required for ensuring data integrity and privacy as well as user authentication.

Types of keys used by the DSM include:

- **Authentication keys**

    Verify the identity of the DSM to the host and the host to the DSM. The DSM and host cannot communicate without valid authentication keys. Authentication keys are referred to as *authentication certificates*.

- **Symmetric encryption keys**

    A symmetric key is one that is a randomly generated AES key used both to encrypt and decrypt information. To decrypt information, one must have the same key that was used to encrypt it. The keys, in practice, represent a shared secret between two or more parties that can be used to maintain a private information link.

- **Asymmetric (public and private) keys**

    Encrypt and decrypt the randomly generated AES symmetric keys that encrypt and decrypt backed-up data. The public RSA key encrypts the symmetric key. The private RSA key decrypts the symmetric key.

    Asymmetric keys require two different keys, one to lock or encrypt the plain text, and one to unlock or decrypt the ciphertext. Neither key can do both functions. One key is published (public key) and the other is kept private (private key). If the lock/encryption key is the one published, the system enables private communication from the public to the unlocking key's owner. If the unlock/decryption key is the one published, then the system serves as a signature verifier of documents locked by the owner of the private key.

- **Imported symmetric encryption keys**

Allows the import of externally generated symmetric keys to encrypt data.

> ⚠️ **WARNING**
>
> **Once encryption is applied, you must keep track of the encryption keys you are using. Encrypted data is unusable without the proper keys. Missing or improperly applied keys are the primary source of data retrieval problems.**
>
> **Back up encryption keys to a secure location without encrypting them. This way, if you must build a new DSM from scratch, you have the keys in a usable form.**

Protecting Keys

A System/All Administrator can protect the DSM keys that they created by preventing an application developer from modifying their existing DSM keys. Application developers may have access to the keys, and may need to create them for testing purposes. However, you can ensure that they do not alter existing keys that were created in DSM. Specifically, you can prevent an application developer from:

- Rotating a key
- Adding a key to a key group
- Adding an attribute to a key
- Modifying, deleting or destroying a key attribute
- Changing the state of a key
- Deleting a key

> **Note**
> Keys created through the Management Console do not have all of the required Key Identifier attributes for certain VAE use cases. Refer to the VAE Guide for details about key usage.

To prevent an application developer from modifying a key:

1. Click **System > General Preferences > System**.

2. Scroll to the **Key Governance** option and select it.

Agent Keys

The DSM creates two types of agent keys: symmetric and asymmetric.

Symmetric keys can be used by:

- VTE agent
- VAE agent
- KMIP client
- Key agent for Oracle TDE

Asymmetric keys can be used by;

- Key agent for Microsoft TDE
- VAE agent

> **Note**
>
> Keys created through the Management Console do not have all of the required Key Identifier attributes for certain VAE use cases. Refer to the *VAE Guide* for details about key usage.

Keys are partitioned into their own DSM domains. In other words, an agent that is registered to Domain-A cannot retrieve, delete, or modify keys stored in Domain-B and conversely, Domain-B cannot retrieve, delete, or modify keys stored in Domain-A. This applies to keys created by the DSM and agents. The key name does not have to be globally unique, but it must be unique within a domain. Therefore, you can have duplicated key names across different domains, but key names must be unique within a domain.

The VTE Agent policies use symmetric keys. Since security policies are only applicable to the VTE agent, the DSM is aware that a symmetric key is used by the VTE Agent, once a key has been assigned to a policy. An attempt to delete a key used by a policy will fail until the key is removed from the policy.

Creating and Adding symmetric keys

The DSM lets you manually create or add symmetric keys and import symmetric keys. (See "Importing Symmetric Keys (BYOK)" on page 192 for more information.)

You can create your own keys or copy third-party keys to the DSM. If you have access to the hex values of AES and ARIA keys, you can enter those values in the Add Vault Key page on Management Console, and then use those keys to encrypt and decrypt files. This is particularly useful if you want to migrate encrypted files from other applications into VTE Agent GuardPoints.

Symmetric keys are based on AES, and ARIA algorithms and are used to encrypt the data in GuardPoints. You can only configure symmetric keys for VTE Agents.

Create a symmetric key

1. Log on to the Management Console as an All/Security Administrator with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the menu bar.

   The Agent Keys page displays.

3. Click **Add**. The Add Agent Key window opens.

4. Select the **Symmetric** tab.

Complete the fields in this window by using the following information.

Table 21-1: Symmetric Key Fields

Field	Description
Name	Enter a name for the key in the **Name** field. This field is mandatory. The maximum number of characters is 64. The name can contain special characters such as: (+ . / _ - =). Note that it **must** start with an alpha character.
Description	(Optional) Enter a phrase or string in the **Description** text-entry box that helps you to identify the key. The maximum number of characters is 256.
Template	(Optional) A key template with a set of pre-defined attributes. Key templates are useful for creating keys of a specific type with specific attributes. Default Microsoft SQL Symmetric and Asymmetric key templates are also provided. If you use this template do NOT modify any of the template attributes. You can create your own templates by selecting Keys > Key Template. To create a Microsoft SQL Server TDE agent symmetric or asymmetric key, choose this template and do not change any of the custom attribute values.

Table 21-1: Symmetric Key Fields (continued)

Field	Description
Expiration Date	Date the key expires. Set a date per your security policies, when the expiration date is reached. It displays in red. Once an expiration date is set for a non-versioned key, when the key expires, you can reset the expiration date. You can also choose to create a new key for your policy, in which case you must rekey your data with the new key. Once an expiration date is set for a versioned key, when the key expires, you can rotate the key and set a new expiration date, or have the expiration date update automatically to a value that equals the date the key is rotated plus the period defined in **Key Version Life Span**. **IMPORTANT!** You must assign an expiration date to a versioned key. Without an expiration date, the key does not contain all of the properties required for versioned keys. As a result, LDT does not recognize the files that need transforming in a GuardPoint when using that key.
Algorithm	Select an encryption algorithm from the **Algorithm** list. Your choices are `AES128`, `AES256`, `ARIA128`, and `ARIA256`. The default is `AES256`.
Encryption Mode	Select an encryption mode from the drop-down list. The options are CBC, CBC-CS1 or XTS. The default is mode is CBC-CS1. **Note 1:** Keys that use the CBC-CS1 mode are only supported by v6.1.0+ VTE agents and DSM. See "Enhanced Encryption Mode" on page 185 for more information. **Note 2:** Keys that use the XTS mode are only supported with Efficient Storage GuardPoints by VTE agents v6.2.0+.
KMIP Accessible	Restricted to encryption keys that use XTS mode. Selecting KMIP Accessible replicates the key to the KMIP database. This makes the key available for retrieval by KMIP clients through KMIP. This is mandatory for keys used with Efficient Storage GuardPoints.
Key Type	Select the location for the generated key from the **Key Type** scroll-list. Your choices are **Stored on Server** and **Cached on Host**. The default is **Stored on Server**. **Stored on Server** keys are downloaded to non-persistent memory on the host. Each time the key is needed, the host retrieves the key from the DSM. Stored on Server requires a constant network connection to the DSM. **Cached on Host** downloads and stores (in an encrypted form) the key in persistent memory on the host. The cached keys are used when there is no network connection between the host and DSM. All hosts using the same encryption key can access encrypted data on other hosts that use the same key. The **Unique to Host** checkbox is displayed when **Cached on Host** is selected.
Unique to Host	This check box is displayed when the **Key Type** is set to **Cached on Host**. When enabled, this check box uses a token that is stored in the host record on the DSM to make the encryption key unique. The unique host encryption key is downloaded to the host and stored in an encrypted manner using the host password. These keys are used for locally attached devices, as files encrypted by them can be read by only one machine. Therefore, do not enable this checkbox for cloned systems, RAID configurations, clustered environments, or any environment that uses host mirroring. The **Unique to Host** checkbox can be enabled only when the **Key Type** is set to **Cached on Host** and the **Key Creation Method** is set to **Generate**.

Table 21-1: Symmetric Key Fields (continued)

Field	Description
Key Creation Method	Select if the key is to be generated automatically using a random seed or if it is to be generated by importing a file. Your choices are **Generate** and **Manual Input**. **Generate** is the default. The **Unique to Host** check box is disabled when **Key Creation Method** is set to **Manual Input**.
Key String	This list is displayed when **Key Creation Method** is set to **Manual Input**. Enter a hex string `[0-9, a-f, A-F]` for the key in the **Key String** text-entry box. - Enter 32 hex characters (128 bits) if the selected algorithm is AES128 or ARIA128. - Enter 64 hex characters (256 bits) if the selected algorithm is AES256 or ARIA256. Re-enter the string in the **Confirm Key String** field.
Key Refresh Period (minutes)	When an Agent Key is cached on host, a DSM administrator can define the refresh period. This setting only applies to VAE keys. Values are from 1 to 44640 minutes, with 10080 minutes as the default value. When set outside of a domain (on the General Preferences page, System tab), the refresh period is applied globally, to all keys.
Automatic Key Rotation	Selecting this option creates a 'versioned' key required for a Live Data Transformation (LDT) policy. The key is automatically rotated based on the expiration date and the period defined in the **Key Version Life Span** option. Refer to the *Live Data Transformation Guide* for more information about using this option with LDT policies.
Key Version Life Span	This field is displayed once you enable the **Automatic Key Rotation** check box. This option specifies the frequency of key rotation in days. Refer to the *Live Data Transformation Guide* for more information about using this option for LDT policies.

5. Click **Ok**.

   The DSM creates new versions of keys which have expired or are about to expire within 24 hours. It computes a new expiration date for the newly created (rotated) key version as follows:

   ```
   (version creation date) + (key version life span)
   ```

For example:

1. Create a versioned key, `TestKey`, on 2/21/2019.

2. Set the key **Expiration Date** to a week after the date you create the key, for our example that would be 2/28/2019.

3. Set the **Key Version Life Span** to 10 days.

4. The key is created with the **Current Version** as '0' to indicate this is the base version of the key.

The DSM creates the first version of the key 24 hours before the expiration on 2/27/2019, with a new expiration date of 03/09/2019. The **Current Version** column for `TestKey`, on the Agent Keys page, displays '1' indicating this is the first version of the key.

The DSM creates a second version of this key on 03/08/2019 (24 hours before expiration) with a new expiration date of 03/18/2019. The **Current Version** column for `TestKey`, on the Agent Keys page, displays '2' indicating that this is the second version of the key.

Creating Asymmetric keys

Asymmetric keys are based on the RSA algorithm and are used to encrypt the symmetric keys. Note that You can configure the symmetric keys for VTE Agents only.

The public half of an RSA key-pair can be imported into other DSMs so that these other DSMs can encrypt data but not decrypt it. A DSM with the private half of the RSA key-pair is required to decrypt data.

Create an asymmetric key

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the menu bar.

   The **Add Agent Key** window opens.

3. Click **Add**.

4. Select the **Asymmetric** tab.

5. Complete the fields displayed using the following information.

Table 21-2: Asymmetric Key Fields

Field	Description
Name	Enter the name for the key in the **Name** text-entry box. This field is mandatory. The maximum number of characters is 64. The name can contain special characters such as: (+ . / _ - =). Note that it **must** start with an alpha character.
Description	(Optional) Enter a phrase or string in the **Description** text-entry box that helps you to identify the key. The maximum number of characters is 256.
Template	A key template with a set of pre-defined attributes. Key templates are useful for creating keys of a specific type with specific attributes. A pre-defined template is provided for Microsoft SQL Server TDE agent asymmetric keys called **Default_SQL_Asymmetric_Key_Template**. To create a Microsoft SQL Server TDE agent asymmetric key, choose this template and do not change any of the custom attribute values.
Key Type	Select the type of RSA key to generate. The choices are **Key Pair** and **Public Key**. **Key Pair** creates a standard RSA key in two parts: a public key and a private key. The **Public Key File** text-entry box is displayed when you select **Public Key**. Use the browser to locate and select a public key file that was generated by another server. The default is **Key Pair**. The key format should be "PEM", which is a base64 encoded format.
Algorithm	Select an encryption algorithm from the **Algorithm** scroll-list. Your choices are RSA1024, RSA2048, RSA3072 and RSA4096. The default is RSA1024.
Public Key File	This text-entry box is displayed when **Key Type** is set to **Public Key**. Click **Browse...** to select the X.509 certificate file that contains the public key.

6. Click **Ok**.

Efficient Storage Encryption Keys

Efficient Storage GuardPoints must be encrypted with a KMIP accessible **XTS-AES256** symmetric keys. An XTS-AES256 key is a 512-bit key composed of two parts:

- The first 256 bits of the key is the AES256 encryption key component

- The second 256 bits is the tweak component

To create an Efficient Storage encryption key:

1. Create a symmetric key. See "Creating and Adding symmetric keys" on page 175.

2. For Algorithm, select **AES256**.

3. For Encryption mode, select **XTS**.

4. Select the **KMIP Accessible** option.

> **Note**
>
> If you do not select KMIP Accessible, the key is not replicated in the KMIP database, and consequently, not available for retrieval through KMIP.

Following is an example of a key created on a DSM with the KMIP Accessible option selected. As depicted, the KMIP key objects corresponding to the AES256/XTS key are available on the DSM through KMIP.

Figure 21-1: Edit Agent Key - Key attributes of ESG_Array1_P1_Key

In-place Data Transformation - Device Encryption Keys

In-place Data Transformation Device (IDT) GuardPoints must be encrypted with an **XTS-AES256** symmetric keys. An XTS-AES256 key is a 512-bit key composed of two parts:

- The first 256 bits of the key is the AES256 encryption key component

- The second 256 bits is the tweak component

> **Note**
>
> The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

To create an IDT encryption key:

1. Create a symmetric key. See "Creating and Adding symmetric keys" on page 175.

2. For Algorithm, select **AES256**.

3. For Encryption mode, select **XTS**.

4. If you are using this for an Efficient Storage device, then select **KMIP accessible**.

Versioned Keys

You can create 'versioned' keys to use with Live Data Transformation policies. The LDT feature enables DSM Security Administrators to encrypt or rekey GuardPoint data without blocking user or application access to that data. Standard (non-LDT) policies require you to associate a non-versioned key with a policy to transform your data, while user and application access to the data is blocked during the transformation process. Transforming this data to use a new key would require a separate policy with this new key being applied to the data.

When you create a 'versioned' key for an LDT policy, you must define a life span for the key. The key is then automatically rotated when it reaches its expiration date. When the key rotates, all its properties including the key name and cryptographic algorithm remain unchanged, except the cryptographic key material that changes the output of the key's cryptographic algorithm. Under LDT policy, the new key material is applied to transform data to the new key version, as part of the same LDT policy that also protects data. You can still manually rotate the key if circumstances require it.

The **Add Agent Key > Symmetric > Automatic Key Rotation** option must be selected in order to create a versioned key. See "Creating and Adding symmetric keys" on page 175 for more about this option.

Refer to the *VTE Live Data Transformation User Guide* for information about implementing LDT and to Chapter 22: "Policies" on page 207 for procedures to create LDT policies.

Storing and Caching Encryption Keys

Encryption keys can be stored exclusively on the DSM, downloaded to the host, or downloaded to the host and stored in non-persistent memory. The keys can also be downloaded and stored (in an encrypted form) in persistent memory on the host for use when there is no network connection between the host and DSM. If a network connection to the DSM is unavailable, and the VTE Agent is configured with persistent keys, enter the host password using the `vmsec passwd` utility. If the host password is configured using challenge-response authentication, run the `vmsec challenge` utility, then contact your DSM administrator with the challenge string, and enter the response string provided by the administrator. Afterwards, you can read and write encrypted data without corrupting it.

Keys are stored or cached in three different ways and have different effects.

Stored on Server keys

- Stored only on the DSM.

- Downloaded to non-persistent memory on the host.

- Remain in effect if they were used before losing the DSM connection.

> **Note**
> A connection to the DSM is required to download keys after a system reboot in order to access encrypted data. Do not apply **Stored on Server** keys to offline host files because, even if you enter the offline password, when there is no network connection and an attempt is made to access the files, the window making the attempt may wait indefinitely for the online keys.

Cached on Host keys

- Stored on the host for offline use.

- A host encryption key, encrypted using the host password, is downloaded to the host. All hosts using the same encryption key can encrypt/decrypt data on other hosts that use the same key.

- When needed, the key is decrypted and cached. When the host is disconnected from the DSM, the current policy remains in effect because the encryption key is locally available.

- If the VTE Agent cannot connect to the DSM after a reboot or `secfs` restart, any attempt to access the contents of an encrypted file on an unconnected host will not complete. The application hangs until the host password is provided to unlock the encryption keys. The application resumes accessing the contents of an encrypted file once the password is provided.

- Connection can be reestablished after a reboot without access to the DSM by entering the host password. You can specify the host password using the "`vmsec passwd`" utility, or, you can display the challenge string in challenge-response host deployments using the "`vmsec challenge`" utility.

Cached on Host with Unique to Host

- Unique key stored on the host for offline use.

- Using a token stored with the host record on the DSM, the encryption key is made unique to each host.

- This unique host encryption key is downloaded to the host and stored and encrypted using the host password.

- When the key is needed, it is decrypted and cached. If the host goes offline and is disconnected from the DSM, the current policy remains in effect because the encryption key is locally available in the system cache.

- If the VTE Agent cannot connect to the DSM after a reboot, any attempt to access the contents of an encrypted file on an unconnected host will not complete. The application will hang until the VTE Agent host password is entered in another terminal window. The application resumes execution once the password is provided.

- The key can also be reestablished after a reboot without access to the DSM through the use of the host password. You can specify the host password using the "`vmsec passwd`" utility, or, you can display the challenge string in challenge-response host deployments using the "`vmsec challenge`" utility.

- These keys are used for locally attached devices, as files encrypted by them can be read by only one host. Therefore, do not use **Cached on Host** with **Unique to Host** keys in any situation where data may be shared by more than one host, such as in clustered environments or any environment that uses host mirroring.

- These keys provide greater security because a key compromised on one host does not compromise the keys on other hosts.

Modifying and Displaying Key Configuration

The following information is displayed in tabular format on the **Keys > Agent Keys > Keys** page, about the keys on the DSM (both symmetric and asymmetric).

- **UUID**

  The Universally Unique Identifier of the generated key.

- **Name**

  Name assigned to the key when created. Names must be unique within a domain but can be repeated across different domains.

- **Versioned Key**

  This column indicates whether a key is a 'versioned' key, which means that it can be automatically rotated and a new version created. It contains two sub-columns:

  ○ **Versioned**

    Indicates if a key can be versioned. If it is a versioned key, a check mark displays.

  ○ **Current Version**

    Indicates the version of the key. When a versioned key is created for the first time, the version number is '0'. When the key is rotated, the version number increments by 1. Refer to the *LDT Guide* for more information about versioned keys.

- **Algorithm**

    The algorithm used to create the key.

- **Key Type**

    'Stored on Server' keys are downloaded to non-persistent memory on the host. Each time the key is needed, the host retrieves the key from the Security Server. 'Cached on Host' downloads and stores (in an encrypted form) the key in persistent memory on the host.

- **Encryption**

    Indicates whether the key is symmetric or asymmetric.

- **Creation Time**

    Date and time the key was created.

- **Expiration Date**

    Date the key expires. This is set when creating a key. Set the date per your security policies, when the expiration date is reached, it displays in red.

    When you set an expiration date for a non-versioned key, when the key expires, you can reset the expiration date. You can also choose to create a new key for your policy, in which case you must rekey your data with the new key. You can create an email notification to alert you when a key is due to expire from the **System > Email Notification** option, at the system level or at the domain level.

    When you set an expiration date for a versioned key, when the key expires, you can rotate the key and set a new expiration date, or have the expiration date update automatically to a value that equals the date the key is rotated plus the period defined in **Key Version Life Span** setting.

    > **Note**
    > You must set an expiration date for a versioned key, if you do not set an expiration date, the key will **not** be rotated.

- **Source**

    The server that requested the key creation. This can be from a key agent host that submits a request to the DSM, or from the DSM itself if the request is generated through the Management Console or `vmssc`. This field is for informational purposes only and is not editable.

- **Description**

    Optional text description of the key.

You can change the following key information of symmetric and asymmetric keys:

- Description (both symmetric and asymmetric keys)

- Expiration date

- Key type

- Key Refresh Period

Modify and display key information

1. Log on to the Management Console as an All or Security Administrator with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the menu bar.

    The **Agent Keys** window displays configured keys and their properties. The table lists a special key; `clear_key` that is available by default and is provided to remove encryption from guarded files and restore the files to their original unencrypted form. This key cannot be deleted.

3. (Optional) Enter all or part of a key name and select a key type in the **Search** panel to display only the keys that match.

   The **Show Search** label located below the **Keys** banner opens the **Search** panel. You can enter a string and/or limit the search to a specific type of key, and click **Go** to display only those keys that match the search criteria. Click **Hide Search** to conceal the **Search** panel.

4. In the **Name** column, click the key that you want to modify.

   The **Edit Agent Key** window displays. The content of this window changes based on the type of key (symmetric or asymmetric) being modified. The Agent Key fields are detailed below.

Table 21-3: Edit Agent Key window field information

Field	Description
UUID	The key's Universally Unique Identifier used to generate a license file.
Name	Name assigned to the key when it was first created. Names must be unique within a domain but can be repeated across different domains.
Source	The machine that requested the key creation. This can be from a key agent host that submits a request to the DSM, or from the DSM itself if the request is generated through the Management Console or vmssc. This field is for informational purposes only and is not editable.
Description	(Optional) Text description of the key. The maximum number of characters is `256`.
Creation Date	Date the key was created.
Expiration Date	Date the key expires. This is the only field on this screen you can modify.
Algorithm	Algorithm used to create the key. The symmetric key algorithms are `AES128, AES256, ARIA128,` and `ARIA256`. The asymmetric key algorithms are `RSA1024, RSA2048, RSA3072` and `RSA4096`.
Encryption Mode	Select an encryption mode from the drop-down list. The options are CBC, CBC-CS1 or XTS. The default is mode is CBC-CS1. **Note 1:** Keys that use the CBC-CS1 mode are only supported by v6.1.0+ VTE agents. See "Enhanced Encryption Mode" on page 185 for more information. **Note 2:** XTS keys are only for use with Efficient Storage and IDT (In-place data transformation - device) GuardPoints.
KMIP Accessible	The XTS key will be replicated to the KMIP DB and you can view it in KMIP objects. This is mandatory for a key for use with Efficient Storage GuardPoints.

Table 21-3: Edit Agent Key window field information (continued)

Field	Description
Key Type	If a symmetric key algorithm is configured, `Stored on Server` or `Cached on Host` can be displayed. `Stored on Server` keys are downloaded to non-persistent memory on the host. Each time the key is needed, the host retrieves the key from the DSM. `Cached on Host` downloads and stores (in an encrypted form) the key in persistent memory on the host. For symmetric keys without `Unique to Host` enabled, you can toggle between `Cached on Host` and `Stored on Server` only. You can enable/disable `Unique to Host` only when configuring a new key. When you switch between `Stored on Server` and `Cached on Host` symmetric keys, the configuration change is pushed to the host. If an asymmetric key algorithm is configured, `Key Pair` or `Public Key` can be displayed. `Key Pair` is a standard RSA key in two parts: a public key and a private key. This type of key can allow an Encryption Agent to back-up and restore data. `Public Key` indicates that key contains only the public key component of a public:private key pair. The DSM with this key can allow an Agent to back-up data only.
Export Key	(Asymmetric keys only) This button opens or saves the key file. Public key only.
Unique to Host	(Symmetric keys only) When activated, unique keys are stored on the host for offline use when there is no connection to the DSM. This option can only be enabled/disabled when configuring a new key.
Key Version	(Versioned keys only) Indicates the version of the key. Any time a versioned key is rotated, the version number increments by 1.
Key Hash	(Versioned keys only) The hash value of the key generated using the key string. If the key is a versioned key, this value changes when the key is rotated. This attribute is not applicable to asymmetric keys.
Automatic Key Rotation	(Versioned keys only) If checked, indicates that the key is a versioned key.
Key Refresh Period (minutes)	(Symmetric keys only) When the Agent Key is cached on host, the administrator can define the refresh period. This setting only applies to VAE keys. Values are from 1 to 44640 minutes with 10080 minutes as the default value. When set outside of a domain under **General Preferences**, the refresh period is applied globally, for all new keys. The refresh period is not reset for existing keys.
Key Version Life Span (days)	This field is displayed once you enable the **Automatic Key Rotation** check box. This option specifies the frequency of key rotation in days. You can edit this field. Refer to the *Live Data Transformation Guide* for more information about using this option for LDT policies.

5. Click **Ok** if you are applying changes.

6. If you are viewing a symmetric key type, click **Back** to return to the **Keys** window.

7. If you are viewing an asymmetric key type, click **Click to Export**.

    The **File Download** window opens.

8. Click **Open** to display the public key component of the asymmetric key in a Web browser.

How the public key is displayed depends on your Web browser and what it does with XML files `.xml`. The key data can be displayed as raw XML code in a Web browser page, or, if your system is configured with an XML editor, the editor can be opened and the formatted XML file displayed.

9. Click **Save** to save the public key component of the asymmetric key.

   The **Save As** dialog box opens.

   a. Specify a path and name for the file.

   b. Click **Save**.

      The **Download Complete** dialog box opens.

   c. Click **Open** to display the public key component of the asymmetric key in a Web browser or click **Close**.

10. Click **Back** to return to the **Keys** window.

Delete keys

> ⚠️ **WARNING**
> **Do not delete keys without first backing them up. All data that has been encrypted with deleted keys cannot be restored or accessed once the keys are gone.**

1. Log on to the Management Console as an All/Security with `Key` role permissions administrator or Domain and Security.

2. Select **Keys > Agent Key > Keys** in the menu bar.

   The **Agent Keys** window opens.

3. Enable the **Selected** checkbox for those keys you want to delete.

4. Click **Delete**.

Enhanced Encryption Mode

A new encryption mode has been introduced for symmetric keys; AES CBC-CS1.

This new encryption mode is supported only by VTE v6.1.0+ and DSM v6.1+. If you have a host group that contains a mix of VTE 6.1.0 and earlier versions of VTE, and you apply a policy containing keys that use the CBC-CS1 encryption mode, the policy does not apply and fails with an error message to the effect that the new encryption mode is not supported by all of the protected hosts in the host group. Similarly if you add a CBC-CS1 encryption key to a host group that has older registered hosts (earlier than v6.1), it fails as the new encryption mode is not supported by all the hosts.

> ⚠️ **WARNING**
> **Once data is encrypted with keys that use a selected encryption mode (either legacy CBC or the new CBC-CS1), the mode is permanent—you cannot switch between encryption modes.**
>
> **To change the encryption mode, i.e. move from using the new CBC-CS1 encryption mode to the legacy CBC mode, then you must transform the data using keys that use the legacy CBC encryption mode using the offline data transform tool or LDT.**

Symmetric keys that use the new CBC-CS1 encryption mode are only supported on DSM v6.1. If you try importing keys that use this new encryption mode to an earlier version of the DSM, the import fails. These keys are not recognized by earlier versions of the DSM. The legacy CBC encryption mode is the default mode when creating a new encryption key.

> **Note**
> For COS policies, only encryption mode CBC-CS1 is supported. CS1 is not supported for cloud object storage.

Encryption Key Protection

In prior versions of this software, the encryption key strings were stored in clear text when they were stored in both key lists and key cache. This made them vulnerable to security attacks. Now DSM can encrypt keys that are cached in kernel memory.

To activate this feature:

- Select the **Encryption Key Protection** option on the Edit Hosts page.

> **Note**
> If you are using a version of VTE Agent that is prior to v6.3.0, this feature does not display.

Exporting and Importing Keys

This section describes exporting and importing symmetric and asymmetric keys for archival, key restoration, or distribution to other DSMs.

You can export, import, and archive the symmetric keys used to encrypt GuardPoint data. You can export and import symmetric keys between DSMs in different HA clusters. You can export the keys of a server to a file in a secure location to ensure that you always have the keys needed to restore encrypted archive data. Without the right keys, encrypted backups are worthless.

> **Notes**
> - Keys that use the new CBC-CS1 encryption mode are only supported by and VTE v6.1+. If you attempt to import keys that use the new encryption mode to versions earlier than 6.1+, the import will fail.
> - You cannot export keys from a later DSM version to an earlier DSM version. For example, you cannot export a key from v6.1.0 and import it into v6.0.0.

The exported key file is itself encrypted. Before you export any keys, create and distribute the key shares of the wrapper key that will be used to encrypt the key file.

> ⚠️ **CAUTION**
> **If you are going to import the keys on another DSM, be sure to import the wrapper key (s) into the other DSM before you import the key file. Do not lose the key shares or you will be unable to decrypt the key file wrapper.**

Importing keys

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. Select **Keys > Agent Keys > Export Import Keys**.

The **Export Import Keys** window opens and displays the **Export** tab.

- If the message "`Export/Import Wrapper Key set.`" is not displayed, set the wrapper key before proceeding. See "Create a wrapper key" on page 96 for more information.

Configure the same wrapper key that you used to create the key file. Copy and paste the same key shares in the **Wrapper Keys** window that you used to make the wrapper file for the exported key file, otherwise you will be unable to import the key file.

3. In the **Import** tab, click **Browse**.

4. Locate and select the key file.

5. Click **Open**.

6. Click **Ok**.

   The **Keys** window opens. If a problem occurs, either real or potential, the **Export Import Keys** window remains open and displays a message. A warning message is displayed if keys in the imported file already exist on the DSM.

   Another typical warning message is "`Wrong path/file name.`" that is displayed when the specified file cannot be accessed. It is also displayed if there is a wrapper key mismatch. If you get this error message, and you are sure that the path and file name are correct, verify that the same key share used to export the key file is also used to import the key file.

   Upon completion without errors or warnings, the **Keys** window is opened. It shows all configured keys, including the imported keys. Similarly named keys are imported with the same name and appended with _*X*, where *X* is an integer. Each time a key with the same name is imported, *X* increments by `1`.

7. (Optional) Check the **Logs** window for additional information about the key import process. A log entry should be generated for each key that is created on the DSM. For example:

   `DAO0239I: Administrator "alladmin" created Symmetric Key "testkeyaes128".`

   Also, a log entry should be generated that identifies the user who initiated the key import process, the number of keys in the file, the SHA hash of the key file, the file size, and the names of the keys in the file. The following example is for successfully importing a small file that contains only three keys:

   ```
   KMG0611I: Administrator "alladmin" imported the following 3 symmetric keys:
   testkey2aes128, testkeyaes128, testkeyaes128_1
   ```

Exporting keys

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. If you are not already in the appropriate domain, switch to it.

3. Select **Keys > Agent Keys > Export Import Keys**.

   The **Export Import Keys** window opens and displays the **Export** tab.

4. If the message "`Export/Import Wrapper Key set.`" is not displayed, create or import a wrapper key before proceeding.

   - Create and distribute the wrapper key. See "Create a wrapper key" on page 96 for more information.

     If the symmetric keys are to be exported to a different server, rather than restored on the originating server, be sure to import the same key shares to the other server to make an identical wrapper key. This way both servers will use the same wrapper key and should be able to successfully encrypt the exported key file on one server and decrypt it on the other.

5. In the **Export** tab, select the check boxes of the keys you want to export.

6. Click **Ok**.

The **File Download** window opens.

The options are:

- **Open** to open the `.dat` file. It will be encrypted so this option is pointless at this time.
- **Save** to save the `.dat` file on the system running the Management Console Web session, or on another network-accessible system.
- **Cancel** to close the window and stop the export operation.

7. Click **Save**.

   The default file name is <*dsm server name*>`_keys_YYYY_MM_DD_HHMM.dat`. For example, `dsmserver1.domain.com_keys_2016_05_11_1252.dat`.

   The **Download Complete** window opens. You can view the file location of the downloaded file.

8. Click **Close**.

9. (Optional) Check the **Logs** window for additional information about the key export process.

   A log entry should be generated that identifies who initiated the key export process, the number of keys in the file, the SHA hash of the key file, the file size, and the names of the keys in the file. The following example is for successfully exporting a small file that contains nine keys:

   ```
   KMG0610I: Administrator "admin1" exported 9 symmetric keys to a file with
   sha1=8c6c3544bd4352f3a8e93a3f478c16489ecd97e5 and size=3524 bytes, containing the following
   keys: aes128, aes128_1, aes128_100901, aes128_StoredOnServer, aes256, aria128,
   testkey2aes128, testkeyaes128, testkeyaes128_1
   ```

Exporting a public key

A public key is the public-key component of a public/private RSA key-pair. The public key of an RSA key-pair is used only to make backups. The private key of an RSA key-pair is used to restore backups. The public key can be imported into other DSMs to enable them as backup-only DSMs. In the figure below, the DSM with copied keys and their SHA-2 hashes (DSM2) can download the keys to another appliance, or system, to do backups only. It does not have the keys to restore those backups. The originating DSM (DSM1) can restore backups because it retains the matching private key. Shared public keys are for environments in which data is backed up in one place with one set of policy constraints, and the backup is restored in another place with a different set of policy constraints.

Figure 21-2: Affiliated DSM configuration
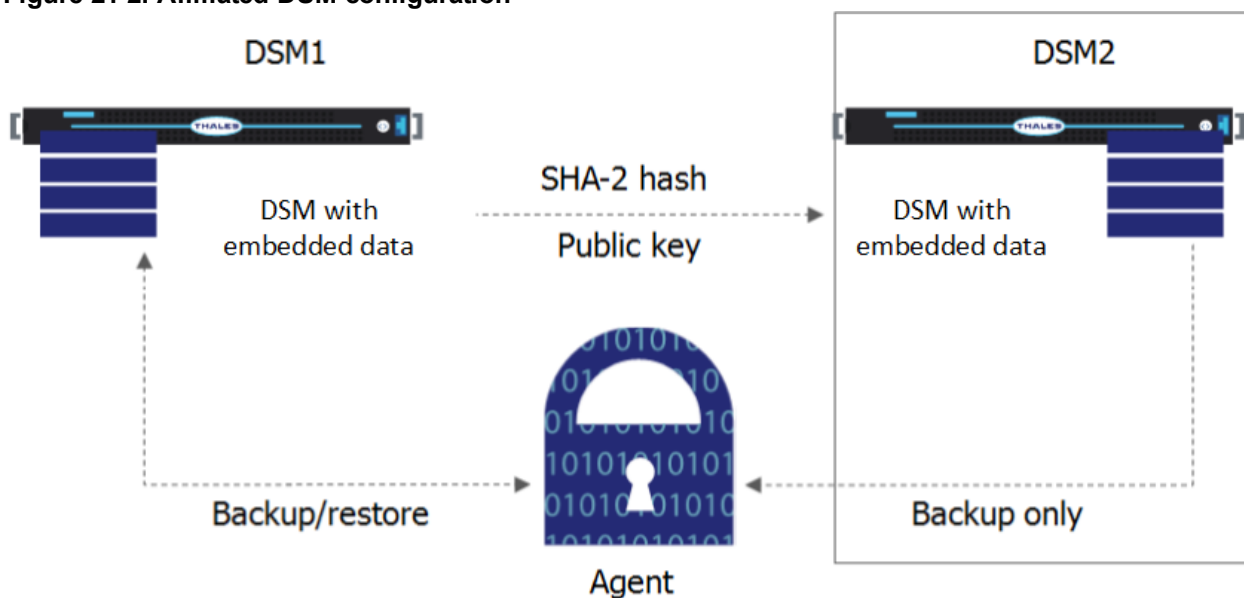
Export the public key of an RSA key pair

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the Management Console menu bar.

   The **Agent Keys** window opens.

3. Click the RSA key-pair or RSA public key in the **Name** column that you want to export.

   The **Edit Agent Key** window opens.

4. Select an export format, PEM or XML. Default is XML.

5. Click **Click to Export**.

   The **File Download** window opens, prompting you to save the public key.

6. (Optional) Click **Open** to display the public key.

   How the public key displays depends on your Web browser and what it does with XML files. The key data can be displayed as raw XML code in a Web browser page, or, if your system is configured with an XML editor, the editor can be opened and the formatted XML file displays.

7. Click **Save to save the public key**.

   The file locater opens.

8. Enter the path and name for the file.

   The default file name is `PublicKey.xml`.

9. Click **Save**.

Importing an Asymmetric public key of an RSA key pair

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the Management Console menu bar.

   The **Agent Keys** window opens.

3. Click **Add**.

4. Select the **Asymmetric** tab.

5. Enter the name to assign the imported public key in the **Name** text-entry box.

6. Select **Public Key** from the **Key Type** scroll-list.

   The **Algorithm** scroll-list is replaced with the **Public Key File** text-entry box.

7. Click Choose File/**Browse...** to open the file locater.

8. Locate and select the public key file.

9. Click **Open**.

10. (Optional) To set an expiration date for the key, enter the date manually in the **expiration Date** text-entry box in the form *MM/DD/YY*, or click the calendar icon, 🗓, and select the expiration date from the graphic interface.

11. Click **Ok**.

   The **Agent Keys** window opens and displays the imported key. The key type is **Public Key**.

Importing Externally Generated Keys (BYOK)

The DSM provides a Bring Your Own Key (BYOK) solution for enterprises that want to use their own keys for encryption operations on the DSM. These externally generated symmetric keys i.e., not generated on the DSM, can be imported to the DSM and used in security policies.

Externally generated symmetric keys can be imported using the Management Console or through the DSM RESTful API. To use the DSM RESTful API, refer to the DSM RESTful API docs located at:

```
https://<dsm_IP_address|FQDN>/app/doc/
```

BYOK

To import an externally generated key, you must first create a key that uses AES encryption, (AES128 / AES256). You will wrap this symmetric key with an Asymmetric RSA public key using either the RSA2048, RSA3072 or the RSA4096 algorithm, from the DSM. The main requirements for importing a key:

- The AES128 or AES256 symmetric key that you wish to import must be converted from hex format to binary

- The symmetric key, in binary form, is encrypted with an RSA 2048, RSA3072 or RSA4096 public key from DSM, using a supported mask generation function and hash padding algorithm

- Convert the encrypted symmetric key to base64 format which can then be pasted into the DSM GUI.

   > **Note**
   > A combination of mask generation function and hash padding algorithm must be used to generate the wrapped bytes, also known as a secret.

The public key encryption mechanism is OAEP (Optimal Asymmetric Encryption Padding). The mask generation function can be SHA256 / SHA384 / SHA512, the hash padding algorithm can be SHA256 / SHA384 / SHA512.

For example:

Encrypted symmetric key:

```
Symmetric AES key (bytes) + Public key (pem format) + mask generation function + hash
algorithm = secret (base64)
```

Decrypted symmetric key:

```
Secret(base64) + Asymmetric private key = Symmetric key imported into DSM
```

Generating a symmetric key file

You can generate a symmetric file using many different methods.

Exporting the public key

In the case of this example, you would export the key with the PEM option selected. See "Export the public key of an RSA key pair" on page 189 for more information.

Convert the key from hex to binary

If the external symmetric key is stored in hex or any other format, you must convert the key to binary.

- **Type**:

  ```
  # xxd -r -p <file_name>.txt <file_name>.bin
  ```

- **Example**

  ```
  # xxd -r -p my_key.txt my_key.bin
  ```

Encrypt symmetric key

The following example uses Openssl to encrypt the key using a previously exported DSM public key (*PublicKey.xml*) with supported hash padding algorithm and mask generation function. In the example below, we use SHA256 for both:

- In OpenSSL, type:

  ```
  openssl pkeyutl -in <path to my_key.bin> -encrypt -pubin -inkey /
  <path to PublicKey.pem> -pkeyopt rsa_padding_mode:oaep -pkeyopt / rsa_oaep_md:<hash type> -
  pkeyopt rsa_mgf1_md:<hash type> -out /
  <path to output.dat>
  ```

  where:

  ° `mykey.bin` is the symmetric key

  ° `PublicKey.pem` is the RSA public asymmetric key

  ° `rsa_oaep_md:`*<hash type>* algorithm can have the following values: sha256, sha384, sha512

  ° `rsa_mgf1_md:` mask generation algo can have the following values: sha256, sha384, sha512

  ° `output.dat` is the secret in binary

  > **Note**
  > The asymmetric public key in the example below must be converted from PKCS#1 format to PKCS#8.
  > Convert it with the following openssl command:
  > ```
  > # openssl rsa -RSAPublicKey_in -in <filename> -pubout
  > ```

  ```
  openssl pkeyutl -in /root/Desktop/mykey.bin -encrypt -pubin -inkey /root/Desktop/pub_
  key.pem -pkeyopt rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256 -pkeyopt rsa_mgf1_
  md:sha256 -out /root/Desktop/output.dat
  ```

Converting the bytes to base64

The DSM requires you to convert the key to base64. You will then paste the generated base64 text information into the **Symmetric Key Material** field in the Import External Symmetric Key page.

To convert the file to base64, type:

```
# base64 -w 0 <filename>
# $ base64 -w 0 output.dat
```

Ncr2PHIaO8S8RfyAnO/XDYHQi26YPWjcyNbqoV4l859Oi/boVOM7g6wE/c+i24x199UyK9V6bN9EVdMpk4O+/vbJpYeTRah
jaDFsHXItwZe3XNBVsJPb65BqlBi8mwa9YwVIR2lFWq/E8+Ve3ITB6lyuDpqgimthm4/Bg22K0/aAP7xuuGGMpy9gCd6pbi
8XeXXkIjkUZKGLdNwTqsRWYJG01HmKKpca4eBzVnNUyutPqiu87O/D/5nQKiqC/vxctP+5dnq5WWzAc7TRe2Db0Kw2zTCp1
XWGcOdX0zPsA1c9SPK2WS3kS8RoIadnBoGCgNfwSrGF8TQJt2m4RLBB3VRMbxKcPIXeCWSEeQgvcxA1xYWeC39k2uU/SwoF
ZzwHVepWSYTZNharT9olE9iMmXdflbEJqwPFfTKdsvvnFskTzdU/FyvnvxD6veRyjpIUrj4TuR+ZJe5O08HqSxCdJRZB0PS
j3o/6xO0YSV1168KgYCrUw1LcBaH0aBxU6qsthmm2m5lRwmP2NEDjm3amAw1XKjwfX+o9ZNI2UiqnBh9139OjG3eYRaivZ5
e0Z5RhGitb65SdrDr32CNY9w2DdbiRb0Vdhub4L+lBAgIdzBDhS8+2DYJ6xkyEkfgmsXhZafCsGkjww1k6XOlR8qlVbRZLM
fFGZzCQ/cT3otE0EsU=

Importing Symmetric Keys (BYOK)

To import an external symmetric key into the DSM:

1. Log on to the Management Console as an All administrator or one with `Key` role permissions.

2. Select **Keys > Agent Keys > Keys** in the Management Console menu bar.

   The **Agent Keys** window opens.

3. Click **Import External Symmetric Key**.

4. Enter the required information in the following fields:

   ◦ **Name**

   Type a name for the key to be imported. (Required)

   ◦ **UUID**

   Enter the unique identifier of the externally generated key. This ID would have been automatically generated with the external key. (Optional)

   ◦ **Description**

   Type a description for the key to be imported. (Optional)

   ◦ **Algorithm**

   Select the algorithm that was used to create the key to be imported. (Required)

   ◦ **Key type**

   Select whether the imported key should be stored on the DSM, or cached on the host. The default option is cached on host. (Required)

   ◦ **Hash Padding Algorithm**

   Select the padding algorithm used to wrap the key. Supported algorithms are: SHA256, SHA384, SHA512. The default value is SHA256. (Required).

   ◦ **Mask Gen Function**

   Select the mask generation function used to wrap the key. Supported algorithms are, SHA256, SHA384, SHA512. The default value is SHA256. (Required).

   > **Note**
   > You must use the same combination of algorithms for the Hash Padding Algorithm and the Mask Gen Function when you wrap the key. For example, if you select SHA512 for Hash Padding Algorithm, and SHA384 for Mask Gen Function, you must select that same combination when you import the wrapped key. It will be unwrapped using the same information.

   ◦ **Encryption Mode**

   The encryption mode used by the key; CBC or the new enhanced encryption mode CBC_CS1.

> **Note**
>
> Only VTE v6.1 agents, and future versions, support the CBC_CS1 encryption mode. Refer to the *VTE Agent Installation and Configuration Guide* for more information.

- ○ **Wrapper Key**

  Click **Select** to select the public key used to wrap the external key. (Required).

- ○ **Symmetric Key Material**

  Paste the generated symmetric key material into the field. (Required). See "Converting the bytes to base64" on page 191 for more information.

5. Click **Ok** to import the key. The imported key is displayed on the **Agent Keys** page, with a **Source** field value of **External**.

Key Groups

Key groups are used to control access to encryption keys by VAE or VKM host administrators. To control access to encryption keys, keys are grouped into key groups, and the key group is then associated with a client identity. The client identity can only access keys in the associated key group.

Only Security Administrators with the 'Key' role, or Administrators for type: All, can create key groups and associate key groups with identities. A key group can be associated with multiple client identities, and a key can be part of multiple key groups. (See "Identity-Based Key Access" on page 198 for more information about creating and managing identities.)

For example, create a key group `keyGrp1` and add two encryption keys to this group. Associate `keyGrp1` with a specific Client Identity. When a user logs on with those Client Identity credentials, that user can only access encryption keys associated with `keyGrp1`.

Key groups can be assigned to more than one client identity, they can be reassigned to different client identities, or can be deleted. Keys in a key group can also belong to more than one key group, and can be removed and reassigned to different key groups.

This functionality is also available through the DSM REST API, refer to the DSM REST API documentation for more information:

```
https://<dsm_IP_address|FQDN>/app/doc/
```

Add Key Group

1. Log on to the DSM as a Security Administrator with the 'Key' role.

2. Navigate to **Keys > Agent Keys > Key Groups** and on the Agent Key Groups page, click **Add**.

3. Type in a name for the key group in the **Key Group Name** field. (Required)

4. Add a description. (Optional)

5. In the Keys section, click **Add**.

6. Select keys to add to this key group from the **Keys** list on the Add Keys to Key group page.

7. Click OK. The DSM returns to the Add Key Group page.

8. In the Assigned Client Identities section, click Add, the Available Client Identities page displays. See "Add Client Identity" on page 198 for more information.

9. Select one or more identities to assign to the key group from the list, click Ok.

10. Click Ok to create the key group and associate it with a client identity or identities.

Edit Key Group

Keys can be added or deleted from a key group.

1. Click the name of the key group on the Name column of the table on the Agent Key Groups page.

2. Add keys to a key group, see "Add Key Group" on the previous page for more information.

3. To delete keys from a key group, select the keys and click **Delete**.

4. To remove assigned client identities, select one or more identities from the **Assigned Client Identities** table and click **Delete**.

5. To add client identities, see "Add Client Identity" on page 198 for more information.

Delete Key Group

1. Select the key group name(s) on the Agent Key Groups page.

2. Click **Delete**.

3. Click **OK**.

Key Templates

Key templates let you quickly add agent keys or third-party vault keys by specifying a template with predefined attributes. You can define specific attributes in a template, then you can call up the template to add a key with those attributes. This is particularly helpful for applications with keys that have customized attributes.

> **Note**
> Pre-defined templates for Microsoft SQL Server TDE agent keys are provided: `Default_SQL_Asymmetric_Key_Template` and `Default_SQL_Asymmetric_Key_Template`.

> ⚠️ **WARNING**
> **DO NOT modify any of the attributes in either of the Microsoft SQL Server TDE key templates or you may prevent access to the database.**

You can also enforce key template usage when creating keys. This means that DSM administrators creating keys must select a key template to define the key's attributes.

The key templates feature allows you to specify common attributes (for example, name, description and algorithm) and custom attributes (attributes specific to certain types of keys such as Microsoft SQL Server TDE keys). The attributes and interface information for key templates are as follows:

Common template attributes:

- **Name**

  Name you assigned the key template when you created it. Names must be unique within a domain but can be repeated across different domains.

- **Description**

  Optional text description of the key template.

- **Algorithm**

  Algorithm used to create the key. The symmetric key algorithms are:

  - AES128

- ° AES256
- ° ARIA128
- ° ARIA256

The asymmetric key algorithms are:

- ° RSA1024
- ° RSA2048
- ° RSA3072
- ° RSA4096

- **Key Type**

  Stored on Server, Cached on Host, Key Pair, Public Key, Vault.

- **Encryption Mode**

  - ° CBC: For DSM versions prior to v6.1.0
  - ° CBC-CS1: For DSM ➔ v6.1.0
  - ° XTS: use with KMIP for Efficient Storage devices

- **Unique to Host**

  Select with `Cached on Host`.

- **Expiration Date**

  Date the key expires.

- **Application Specific Information**

  Optional data that is specific to the application.

- **Contact Information**

  Optional contact information.

- **Attribute Name**

  Name of the added custom attribute.

- **Attribute Value**

  Value of the added custom attribute.

Default Microsoft SQL Server TDE symmetric key template attributes:

- **Attribute Index**

  Value indicating whether a key is `supported` (0x01), `volatile` (0x02), `exportable` (0x04) or `importable` (0x08). The attribute value can be any combination of these bit masks. For example, an attribute that is supported and exportable would be 0x05.

- **Cryptographic Usage Mask**

  A bit mask to define the key cryptographic usage. The first 7 bits indicate Sign, Verify, Encrypt, Decrypt, Wrap Key, Unwrap key, and Export. The SQL server requires that a symmetric key can be used for all of these purposes. The cryptographic usage mask value in binary bits is 1111111 (decimal value is 127).

- **Object Type**

  The type of object. Values can be `SymmetricKey`, `PublicKey` or `PrivateKey`. Since this is the key template for symmetric keys, the value must always be `SymmetricKey`.

- **x-VormCanBePlainText**

  Specifies whether the key value can be revealed in plain text outside the DSM (`true`) or not (`false`). This value is always set to `true`.

- **x-VormCanNeverBeExported**

Specifies whether the DSM can never export key values, and will return an error when a user tries to do an export (`true`) or not (`false`). This attribute protects sensitive key material from being exported outside the server. In general, only public keys can be exported. This value is always set to `true` which means that you cannot export key values.

- **x-VormCanNeverBePlainText**

    Specifies whether the key value can never be revealed in plain text outside the DSM (`true`) or if revealing it in plain text is allowed (`false`). This value is always set to false which means that key values can be revealed in plain text.

- **x-VormCanObjectPersist**

    Specifies whether the DSM can store the key after it creates it (`true`) or does it create the key and return it to Key Agent without storing it (`false`). This value is always set to `true`.

- **x-VormID**

    This is the SQL-server-customized key identifier that the SQL server uses to locate the key. The value must be unique in DSM.

Default Microsoft SQL Server TDE asymmetric key template attributes:

- **Attribute Index**

    Value indicating whether a key is `supported` (0x01), `volatile` (0x02), `exportable` (0x04) or `importable` (0x08). The attribute value can be any combination of these bit masks. For example, an attribute that is supported and exportable would be 0x05.

- **Cryptographic Usage Mask**

    A bit mask to define the key cryptographic usage. The first 7 bits indicate Sign, Verify, Encrypt, Decrypt, Wrap Key, Unwrap key, and Export. The SQL server requires an asymmetric private key be used for signing, decryption, unwrapping and exporting purposes. The cryptographic usage mask value in binary bits is 1101001 (decimal value is 105).

- **Object Type**

    The type of object. Values can be `SymmetricKey`, `PublicKey` or `PrivateKey`. Since this is the key template for asymmetric key, the value here will always be `PrivateKey`.

- **x-VormID**

    This is the SQL server customized key identifier that the SQL server uses to locate the key. The value must be unique in DSM.

Common Key Template Procedures

Use the following procedures to manage your key templates.

Adding a key template

1. Select **Keys > Key Templates.** The **Key Templates** window opens.

2. Click **Add.** The **Add Key Template** window opens.

3. Fill in the appropriate standard attributes.

4. Click **Add** to add customized attributes.

5. Add the attribute name and value and click **Ok**.

6. Add as many customized attributes as needed.

7. Click **Ok** in the **Add Key Template** window. The **Key Templates** window displays with the new template listed.

Deleting a key template

1. Select **Keys > Key Templates**. The **Key Templates** window opens.

2. Select the template that you want to delete.

3. Click **Delete**.

Modifying a key template

1. Select **Keys > Key Templates.** The **Key Templates** window opens.

2. Click on the template name to modify existing attributes or add custom attributes.

> **Note**
> DO NOT modify any of the attributes in either of the Microsoft SQL Server TDE key templates called `Default_SQL_Asymmetric_Key_Template` and `Default_SQL_Symmetric_Key_Template`. Doing so may prevent access to the database.

Using a key template

1. Select **Keys > Vault Keys > Vault Keys** or **Keys > Agent Keys > Keys**.

2. Click **Add**. This displays the **Add Vault/Agent Key** window

3. Click the **Template** pull down menu and select a key template.

Enforcing a key template to define a key

1. Select **System > General Preferences > System Tab.**

2. Select **Enforcing Using Key Template to Define Key**.

3. Click **Apply**.

Agent Objects

The Agent Objects page displays objects created by the VAE agent/VAE libraries in the current domain. Opaque objects are used to import previously unsupported size keys, create keys with unsupported algorithms, and import certificates associated with a key. The object itself is not necessarily a key. The page displays the following information about the objects in a tabular format:

- **Name**

    The name of the object. Click the name of the object to view more details about the object.
- **Algorithm**

    To view details, you must log on to the VAE host and view the object.
- **Key Type**

    Objects created by the VAE agent and stored on the DSM are Opaque Objects.
- **Creation Date**

    The date when the object was created.
- **Expiration Date**

    The date when the object expires, if any.
- **Description**

    A description, if any, of the object.

Identity-Based Key Access

The VAE and VKM agents provide identity-based access control for encryption keys stored on the DSM. The identity of a VAE or VKM user is established using credentials; user name and password, and a corresponding identity profile is created on the DSM. This DSM identity profile is then associated with a key group that contains the keys that client identity is allowed to access. When a VAE or VKM user logs in with a Client Identity profile and tries to access keys, the DSM verifies that identity profile and then grants access to keys in key groups associated with that identity profile. For more about creating and managing key groups see "Key Groups" on page 193.

An identity can be associated with multiple key groups. Only Security Administrators with the 'Client Identity' role, or All Administrator can create client identities.

> **Note**
> Security Administrators with the Client Identity role assigned can only create identities and have a limited view of menu options. They cannot perform any other tasks on the DSM. As a best security practice, Thales recommends that you do not assign both 'Key' and 'Client Identity' roles to a single Security administrator.

Username Requirements

The identity's username requirements are as follows:

- May contain the following non-alpha-numeric characters:
  - at (@)
  - dot (.)
  - underscore (_)
  - dash (-)

Password Requirements

The password restrictions are:

- Cannot allow colon ":"
- Min password length: 1 character
- Maximum password length: 256 characters

For password requirements:

- Your application may contain other requirements for passwords. Follow the requirements provided in the application documentation.

Add Client Identity

1. Log on as an All/Security Administrator with the 'Client Identity' role.
2. Navigate to the **Keys > Identities** page.
3. On the Client Identities page, click **Add**, the Add Client Identity page displays.
4. Enter the following information:
   - **Identity Name**: Name of the VAE/VKM user.
   - **Description**: Description for the identity profile, this is optional.
   - **User password**: User password, enter the password again to confirm

5. Click **Ok** to add that identity profile to the DSM.

Edit Client Identity

Edit a client identity to change the password or description.

1. On the Client Identities page, click the client identity in the **Name** column, the Edit Client Identity page displays.

2. Select the **Update User Credentials** check box.

3. Enter the new password information in the **User Password** and **Confirm User Password** fields, or change the **Description** field.

4. Click **Ok** to confirm the update.

Delete Client Identity

To delete a client identity or identities:

1. Select the identity to remove.

2. Click **Delete**.

Vault Keys

The Key Vault is a secure centralized repository in the DSM for storing and retrieving third-party encryption keys. Agent Keys are not stored in the Key Vault. Here are a few things to note about securing keys in the DSM key vault:

- When an administrator adds a key to the key vault, only another administrator in the same security domain can use that key.

- You cannot create new keys in the key vault. You can only add existing keys created elsewhere.

Symmetric Vault Keys

The Vault Keys page lets you manage your key vault. It displays the following information about the Vault Keys on the DSM for Symmetric keys:

- **Name**

  Key name.

- **UUID**

  This field is only available on the Edit Vault Key page. The Universally Unique Identifier of the generated key.

- **Source**

  This field is only available on the Edit Vault Key page. The server from which the add key request is sent. This will always be this DSM. It is for informational purposes only and is not editable.

- **Description**

  Optional text description of the key.

- **Template**

  You can select a key template with a set of pre-defined attributes. Key templates are used to create keys of a specific type with specific attributes, without having to manually enter those attributes.

- **Algorithm**

  Algorithm used to create the key. The symmetric key algorithms are AES128, AES256, ARIA128, and ARIA256.

- **Key Type**

  All keys stored in the Key Vault are of type VAULT.

- **Key String**

  Only available on the Add Vault Key page, Symmetric tab. Enter the key string value.

- **Confirm Key String**

  Only available on the Add Vault Key page, Symmetric tab. Confirm the key string value by re-entering it in this field.

- **Creation Date**

  Date that the key was created. If you do not enter a date, the date that the key was stored in the vault is used.

- **Expiration Date**

  Date on which the key expires.

- **Validate**

  If selected, the system validates the uploaded key against the specified algorithm. If you disable this option, you can vault a key without validation as long as it does not exceed the maximum file size. This option is enabled by default.

- **Attributes**

  Only available from the Edit Vault Key page, Attributes tab. This tab allows you to add, modify or delete key attributes, such as whether a key can be exported or read in plain text outside of the DSM. Any key created on the DSM (symmetric or asymmetric) has the following NIST compliant attributes:

  - **x-deactivation-date**: Available only if an expiration date is set for that key. The expiration date is in UNIX epoch format.

  - **x-key-state**: Values reflect the current state of the key: ACTIVE, SUSPENDED, COMPROMISED or DEACTIVATED. You can change the state of the x-key-state by selecting a different value from the dropdown menu. If you select DEACTIVATED and click OK, then the dropdown list changes to COMPROMISED, DESTROYED or DEACTIVATED.

  > ⚠️ **WARNING**
  > **Use the DESTROYED function VERY CAREFULLY. Destroyed keys cannot be recovered and all data encrypted by that key will be lost.**

  > **Note**
  > If the base version (version 0) is deleted, the entire key is deleted.

You can **Add**, **Delete** and **Export** keys from this page.

Asymmetric Vault Keys

The following fields display when you edit a Vault key.

- **Name**

  Key name.

- **Description**

  Optional text description of the key.

- **Template**

  This is only available on the Add Vault Key page. You can select a key template with a set of pre-defined attributes. Key templates are useful for creating keys of a specific type with specific attributes without having to manually enter those attributes. To create a key template, select **Keys > Key Template** and select **Key Type** as Vault.

- **Algorithm**

DSM Administration Guide
200

The algorithm used to create the key.

The asymmetric key algorithms are:

- ° `RSA1024`
- ° `RSA2048`
- ° `RSA3072`
- ° `RSA4096`

- **Key Type**

  All keys stored in the Key Vault are of type VAULT.

- **Public Key File**

  Only available on the **Add Vault Key** page, **Asymmetric** tab. This is the public key file to store in the Key Vault.

- **Private Key File**

  Only available on the **Add Vault Key** page, **Asymmetric** tab. This is the private key file to store in the Key Vault.

- **Creation Date**

  Date that the key was created. If you do not enter a date, the system will enter the date that the key was stored in the vault.

- **Expiration Date**

  Date the key expires.

- **Validate**

  If selected, the system validates the uploaded key against the specified algorithm. If you disable this option, you can vault a key without validation as long as it does not exceed the maximum file size. This option is enabled by default.

- **Source**

  The server from which the add key request is sent. This will always be this DSM. This field is only available on the Edit Vault Key page. It is for informational purposes only and is not editable.

- **Creation Date**

  Date the key was created. Available on the Edit Vault Key page.

- **Attributes**

  The asymmetric SQL key attributes are:

  - ° Attribute Index
  - ° Cryptographic Usage Mask
  - ° Object Type
  - ° x-VormID

Adding a Vault key

1. Select **Keys > Vault Keys > Vault keys**.

2. Click **Add**.

3. On the **Add Vault Keys** page, click the **Symmetric** or **Asymmetric** tab depending on the type of key you are adding to the vault.

   - ° To vault symmetric keys, paste in the string.
   - ° To vault asymmetric keys, upload the raw public and private key files.

4. Enter the appropriate key information. (See the attributes and interface information described in "Symmetric Vault Keys" on page 199 or "Asymmetric Vault Keys" on the previous page.)

5. Click **Ok**.

Add an attribute to an Asymmetric Key

To add an attribute:

1. Click **Add** and enter the a name and value for the attribute:

   ○ Attribute Name

     Name of the added custom attribute.

   ○ Attribute Value

     Value of the added custom attribute.

2. Click **OK**.

Deleting a Vault key

1. Select Keys > Vault Keys > Vault keys.

2. Select the key to be deleted.

3. Click **Delete**.

Exporting a Vault key

The export feature allows an administrator to export the selected keys to an external location. The default location for Windows is:

```
C:\Users\<username>\Downloads
```

The file is saved as a .zip file with the following default naming convention:

```
vault_key_<servername_YYYY_MM_DD_HHMM>.zip
```

1. Select **Keys > Vault Keys > Vault keys**.

2. On the Vault Keys page, select the key(s) to be exported.

3. Click **Export**.

Editing a Vault key

1. Select **Keys > Vault Keys > Vault keys**.

2. On the Vault Keys page, click the key to be edited.

3. Click the **General** tab to change the standard attributes.

4. Click the **Attributes** tab to modify or add custom attributes.

Generating a list of vault keys with their expiration dates

1. Click **Reports**. This displays the Reports window

2. On the **Reports** page, click **Keys** under the **Name** column. The View Keys page is displayed.

3. In the **Search** pane, select **Vault Keys** and click **Go**.

   A list of vault keys with the expiration dates displays.

4. Click Download if you want to downloaded the data as a `.csv` file.

Certificate Vault

This section covers the following topics:

-

-

-

-

-

-

The Certificate Vault stores certificates used by applications and devices across the enterprise. It maintains an inventory of certificates with expiration tracking and notification. No agent is associated with the Certificate Vault.

Table of Vaulted Certificates

The Certificates page displays a table that lists all of the certificates in the vault, and information about each certificate. The following table describes each of the columns in the Certificates table.

Table 21-4: Vault Certificates

Column Name	Description
ID	Unique identification number for the certificate
Subject	Name of the certificate
Issuer	Entity that issued the certificate
Valid From	Issuing date of the certificate
Valid To	Expiration date of the certificate
Signature Algorithm	Algorithm for creating digital signatures, such as Sha1 with RSA, MD2 with RSA
Format	Format of the vaulted certificate, such as PEM, DER, PKCS#7, PKCS#12
Type	Root, intermediate, leaf
With Private Key	Contains a private key? Y or N
Extended Key Usage	Server authentication and/or client authentication
Public Key Algorithm	Algorithm used to sign the certificate, e.g., RSA or DSA
Public Key Size	512, 1024, 2048, 3072, 4096, etc.
File Name	Name of the certificate file
Description	Additional information for the certificate

Import Certificate to Certificate Vault

1. On the Certificates page, click **Import**. The Upload Certificates page displays.

2. Click **Choose File**. Navigate to the certificate file, and then click **Open**.

3. Add additional detailed information in the **Description** field, if desired.

4. Click **Next**. The Upload Private Key page displays.

   For PKCS#12 certificates you will get a Fetch Password page instead of the Upload Private Key page.

5. Enter and then confirm the password.

6. If your certificate has an associated key, click **Browse**, navigate to the key file, and then click **Open**.

7. Click **Next**.

8. Click **Finished**.

Delete Certificate from Certificate Vault

1. On the Certificates page, select the certificate(s) to be deleted.

2. Click **Delete**.

Export Certificate

1. On the Certificates page, select the certificate(s) to be exported.

2. Click **Export**.

The file is saved in .zip format with the default naming convention of certYYYY_MM_DD_HHMM.zip

Search Certificate Vault

- If the search panel is not visible at the top of the Certificates page, click **Show Search**.

- To hide the search panel, click **Hide Search**.

Search for certificates using one or more of the following fields:

Table 21-5: Search Fields

Search Field	Description
Subject	Name of the certificate.
Valid From	Select a date from the calendar or enter a date equal to or earlier than the issuing date on the certificate.
Type	Filter by the level in the certificate authority chain of trust. In the drop down box, click one of the following: All Root Intermediate Leaf Unknown
Public Key Min Size	Specify the minimum size of the key in bits, e.g. 512, 1024, 2048, 3072, 4096.
Public Key Algorithm	Enter RSA or DSA.
Description	Optional description

Table 21-5: Search Fields (continued)

Search Field	Description
Valid To	Select a date from the calendar, or enter a date equal to or earlier than the expiration date on the certificate.
Extended Key Usage	Method of server and/or client authentication. In the drop down box, click one of the following: Server Authentication Client Authentication Code Signing Email Time Stamping OCSP Signing Unknown
Public Key Maximum Size	Specify the maximum size of the key in bits, e.g. 512, 1024, 2048, 3072, 4096.
With Private Key	Click if you are searching for a certificate where a Private Key was used when the certificate was imported.

Generate Certificates Report

To generate a report about certificates, you must be logged into a domain as an administrator with AUDIT privileges.

1. On the Management Console, click **Reports**. The Reports page opens.

2. In the Name column, click **Certificates**. The View Certificates page displays showing a list of vault keys with their expiration dates.

The Certificates Report table describes all of the certificates in the Certificate Vault.

Table 21-6: Certificates Report

Column Name	Description
ID	ID of certificate
Subject	Name of the certificate
Issuer	Entity that issued the certificate
Valid From	Issuing date of the certificate
Valid To	Expiration date of the certificate
Public Key Algorithm	Algorithm used to sign the certificate, e.g., RSA or DSA
Public Key Size	512, 1024, 2048, 3072, 4096, etc.
Signature Algorithm	Algorithm for creating digital signatures, such as Sha1 with RSA, MD2 with RSA
CA	Certificate authority such as VeriSign, AT&T, GoDaddy, etc.
Type	Root, intermediate, leaf
Extended Key Usage	Server authentication and/or client authentication, code signing, email, time stamping, OCSP signing, unknown

Table 21-6: Certificates Report (continued)

Column Name	Description
CRL	Certificate Revocation List
OSCP	Online Conflict Status Protocol
Format	Format of the vaulted certificate, such as PEM, DER, PKCS#7, PKCS#12
With Private Key	Contains a private key? Y or N
File Name	File Name
Description	Description

Chapter 22: Policies

The primary job of a DSM Security Administrator is to create policies that protect data. Policies govern access to, and encryption of, the files in VTE-protected directories. VTE-protected directories are called GuardPoints.

Overview

A Data Security policy is a collective set of rules that govern data access and encryption. Think of a policy as an if-then statement. The rules are processed sequentially. If the criteria of rule one are not met, the policy enforcement engine moves on to the second rule and so on. The following criteria are processed by the policy enforcement engine:

- **Order**

  Security rule enforcement sequence.

- **Resource**

  Files and/or directories to which the policy will apply, plus key rules that govern those files and directories.

- **User**

  Users and user groups authorized to access the resources.

- **Process**

  Executables which will access the files.

- **Action**

  Type of user access being made (read, write, copy, move etc.). Before you can define Data Transformation Rules, you must select a Key_op action.

- **Effect**

  When all the other rules match, this describes the type of access granted or denied per the rule.

- **When**

  Time frame within which the action occurs.

- **Browsing**

  Allow browsing is enabled by default, while the Enable Communication check box is enabled on the host. This allows the server to browse the host's file system. This option can be deselected even if host communication is still enabled.

A policy comprises 'Security Rules' and 'Key Rules'. A **security rule** defines the users or user groups authorized to have specified access to specific files or directory paths for a designated period of time. In short, it defines who is accessing data (**User**), what they can do with the data (**Action**), which applications or executables have access to the data (**Process**), where the data is located (**Resource**), the time frame that the 'Security Rule' is applicable (**When**), how the data can be accessed (**Effect**), and if it can be viewed from the DSM (**Browsing**).

A **key rule** defines the encryption key to apply to a specific resource set, or the encryption key to use as the default key, in the event that no other key rule matches. It defines the sequence in which the key rules are to be executed (**Order**), the location of the data to be encrypted (**Resource**), the encryption key to be applied to the resource set (**Key**).

When defining a key rule for a Live Data Transformation (LDT) policy, you can select a key that is applied to the resource set (**Current Key**) and the key to use to rekey that resource set (**Transformation Key**).

When defining a key rule for a Cloud Object Storage policy, you can only select one key. Resource selection is not applicable for COS because the resource will always be the cloud object storage device. For DSM v6.4.0, it is only compatible with Amazon AWS S3 Cloud Object Storage.

> **Note**
> The DSM Cloud Object Storage feature is compatible with VTE v6.3.0 and all subsequent versions.

Figure 22-1: Policies relationship to the GuardPoints, Hosts, and the DSM
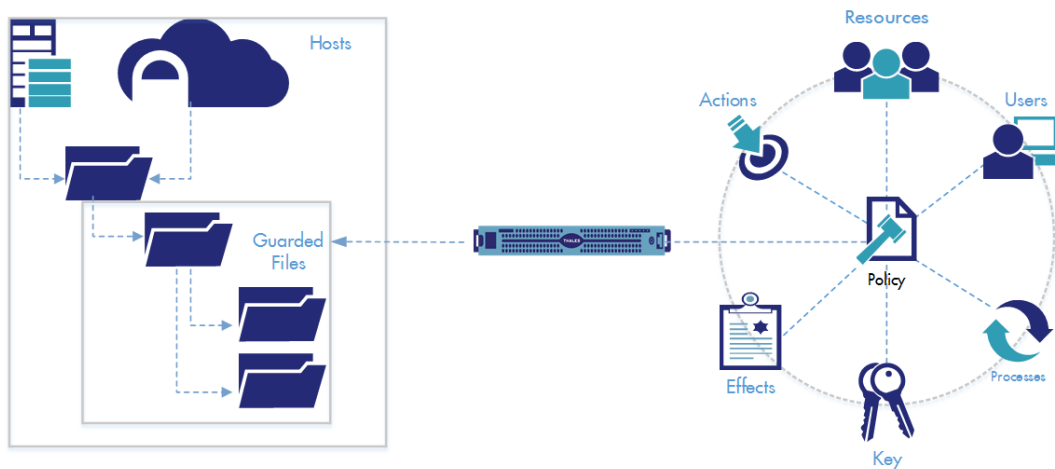
Policy Rule Criteria and Effects

Policy Rules consist of five criteria, which specify the attributes of an access attempt, and effects, which define whether that access is permitted or denied, and whether encryption/decryption is required.

Table 22-1: Policy Rule Criteria

Criteria	Action
Resource	Specifies which files and/or directories in a GuardPoint are to be blocked. Example: `/secure_dir/financials` **Note**: You cannot select a resource set for policies for Cloud Object Storage. The resource set is automatically the cloud object storage.
User	Specifies a which users or groups of users can access the files.
Process	Specifies executables that can operate on the files.
When	Specifies the time range when files can be accessed.
Action	Specifies the allowed file action. Example: read, write, remove, rename, make directory.

Table 22-2: Policy Rule Effects

Effect	Action
Permit	Permit access to the data.
Deny	Deny access to the data.

Table 22-2: Policy Rule Effects (continued)

Effect	Action
Apply Key	Encrypt data written into GuardPoint with the key specified in the Key Selection Rules tab. Decrypt data that is accessed using the same key.
Audit	Creates an entry in the Message Log that describes what is being accessed, when it is being accessed, the security rule being applied.

Every time a user's application tries to access a GuardPoint file, the security policy tests that access attempt against the criteria of each rule. For example, suppose user `Harry` wants to access and modify a file called `secret`, using the command `cp`, at 3AM. For `Harry` to be successful, there must be a rule that allows access to `secret` (resource), by user `Harry` (user), using the command `cp` (process), at 3AM (when), and includes the permission `write` (action).

A blank criteria field specifies a value of All. If User is blank, the rule applies to all users; if When is blank, the rule applies to all times; if Process is blank, the rules applies to all executables, and so on. Effect can never be blank. It must have at least a permit (allow access) or deny (disallow access).

A policy can have multiple rules. Rules are evaluated much like firewall rules; they are evaluated in order, from first to last, and evaluation stops when a rule is found for which all the criteria are met. The effect for that rule is then enforced. Therefore, you must carefully order a policy's rules to achieve the desired result.

Creating and Configuring VTE Policies

Accessing the Domain to be Protected

1. Log on to the Management Console as a Security Administrator in the domain containing your protected host.

   Alternatively, log into the local (restricted) domain to which you belong.

2. Click **Domains > Switch Domains.** The **Switch Domains** window opens.

3. Select the domain containing the protected host and click **Switch to domain**. The active domain name displays in the upper right corner of the Management Console.

Add a Policy

1. Click **Policies > Manage Policies > Manage Policies** to list the policies available to a domain.

Figure 22-2: Policies window
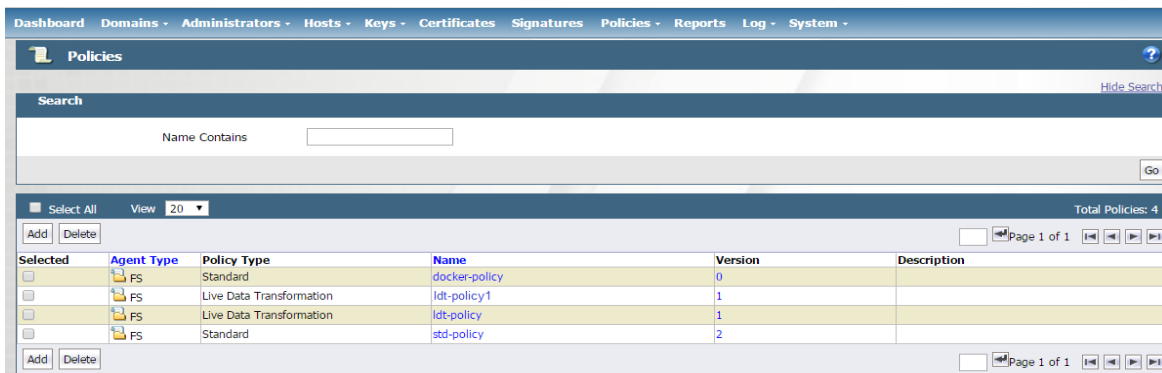
2. Click **Add**. The Add Policy page displays.

Figure 22-3: Add Policy window
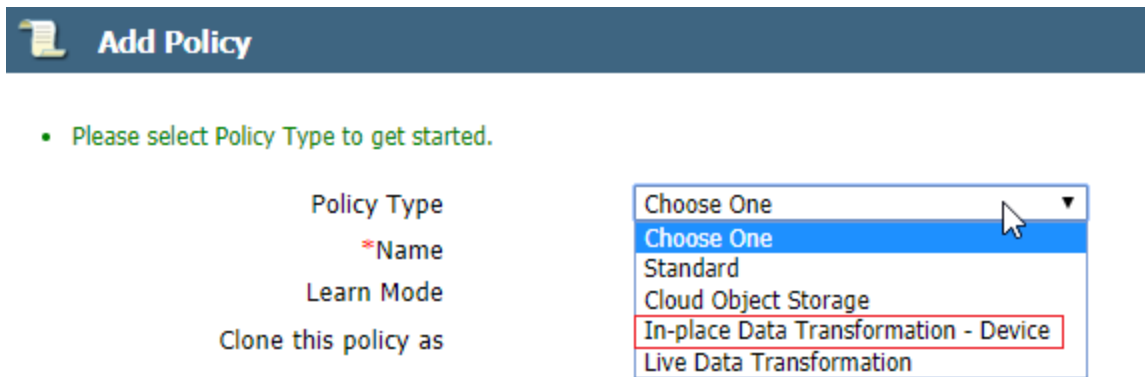
3. Select a policy type from the **Policy Type** drop-down list. The options are:

    ○ **Standard**

    Use for offline Data Transformation policies.

    ○ **Cloud Object Storage**

    A Cloud Object Storage policy is only available for Agents that support Cloud Object Storage. Currently, DSM only supports Amazon AWS S3 Cloud Object Storage.

    ○ **In-place Data Transformation - Device**

    You can use an IDT-Device policy with a raw device, as of VTE v6.3 and DSM v6.4.2, if you have a valid VTE license. You can use an IDT policy for Efficient Storage GuardPoints or In-place data transformation GuardPoints.

> **Note**
> The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

> **Note**
> - For a new key for an In-place data transformation policy for an Efficient Storage device, you must select an AES key with XTS encryption that can be KMIP accessible.
>
> - For a new key for an In-place data transformation device policy, you must select an AES key with XTS encryption
>
> - An IDT policy requires two key rules. However, you can only add one key rule, with two keys, in the UI. Therefore, when you create an IDT policy, DSM converts the key rule into two key rules so that the policy contains the required two key rules. (The API does not have this limitation.)

    ○ **Live Data Transformation**

    An LDT policy is only available if you have a valid LDT license.

    See "Enabling Live Data Transformation" on page 240 for more information about LDT. Refer to the *Live Data Transformation Guide* for information about implementing LDT.

4. Create a name for your policy a **Name** (for example, `basic-access-policy` or `ldt-policy`) and an optional **Description**.

5. **Learn Mode (**Optional). This mode permits a policy to be tested without actually denying access to the GuardPoint. In Learn Mode, all actions that would have been denied are instead permitted, but logged. This allows you to test policies by tracking how rules are evaluated, without enforcing the policy. Monitor the log to determine how data is being accessed, then modify the policy accordingly.

A **deny** statement in **Effect** must include **apply_key** when Learn Mode is enabled. This option generates a warning each time an access attempt is made that matches any security rule in the policy. This warning is sent as a log message and it can be viewed in the Management Console (if it's configured to accept warnings).

Learn Mode is recommended for policies that restrict by application (process), as many applications use multiple binaries that may not be known to the creator of the policy at time of creation.

> **Note**
>
> Learn mode is not supported for Cloud Object Storage or In-place data transformation - device polices.

6. **Clone this policy as** (Optional). Type in a new policy name and click **Clone**. This creates a clone of the original policy.

7. Clicking **Ok,** at this point, creates a blank policy if you are creating a standard or cloud object storage policy. This policy has no rules.

   If you selected an LDT policy, the policy has one security rule added to it by default—key_op—that cannot be deleted, edited or reordered.

8. To add rules to the policy, click **Add**. The **Add Security Rule** window opens.

Adding a Security Rule Resource Set to a Policy

Security rules specify how the DSM will respond to an access request.

To add security rules to a policy:

1. Open the **Add Security Rule** window if it is not displayed. Click the policy name in the **Policies** window if the policy has already been created.

2. Click **Add** in the **Security Rules** panel. The **Add Security Rule** window opens.

   - If you are creating an **Live Data Transformation** policy type, then the first security rule for the policy is created by default. This rule permits key operations on all resources for that policy, without denying user or application access to resources, so that a rekey operation can be done whenever the encryption key is versioned. This rule is **always** the first rule in an LDT policy and cannot be edited.

     For more information about creating policies for Live Data Transformation, refer to the *Live Data Transformation Guide*.

Figure 22-4: Add Security Rule window
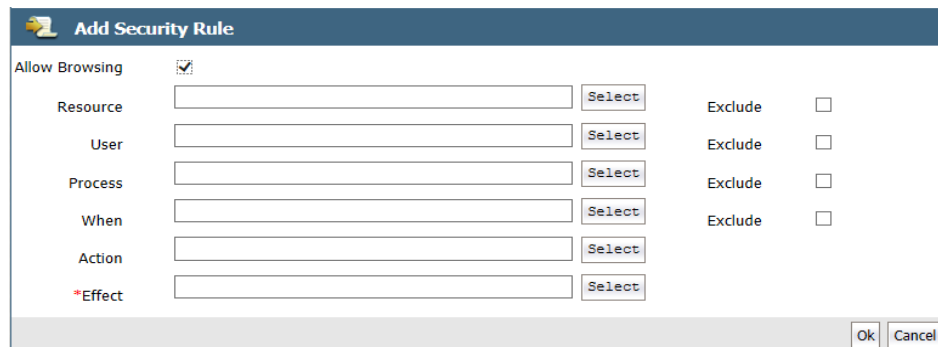
3. Click **Allow Browsing** to enable the user to access and traverse directories below the GuardPoint, leading down to the resources in the rule. Users that match the criteria set by the security rule can access the directories between the GuardPoint and the resource. If you define a resource, then that resource should exist in the GuardPoint. The files in these directories can be listed like any file in a browser, but they cannot be modified, copied, or deleted.

> **Note**
> This section walks you through adding the criteria by clicking the **Select** button. If the criteria has already been defined and you know the names, you can type their names in the text boxes.

> **Note**
> Allow Browsing is not supported for Cloud Object Storage polices.

4. **Resource** (Optional): Specifies the hosts, files and directories to which a user or process will be permitted or denied access. Though not mandatory, if you define a resource, that resource should exist in the GuardPoint.

> **Note**
> Policies for Cloud Object Storage do not contain a resource set. The resource set is automatically the cloud object storage.

a. To specify all resources, leave **Resource** blank.

To define specific resources in a GuardPoint, click **Select**. The **Select Resource Set** window opens.

b. Click **Add** to add a Resource Set. The **Add Resource** window opens.

Figure 22-5: Add Resource window
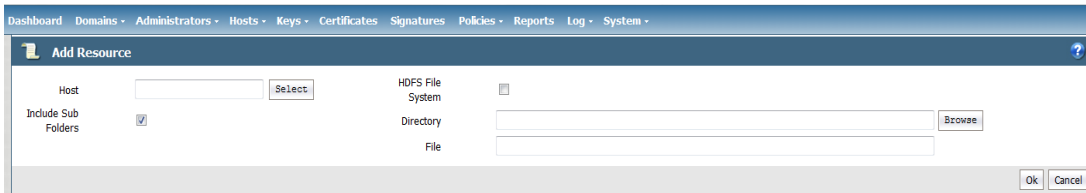
c. Specify the **Host**, **Directory** and **Files** on which to apply the rule.

- **Host** is the hostname containing the directory. Enter the hostname and click **Browse** to browse for the directory.

- If your host is a Docker host, another field displays: **Docker Image/Container**. Click **Browse** to open the **Remote Docker Browser** to select a Docker image or container, from which you can select a resource.

Figure 22-6: Remote Docker Browser
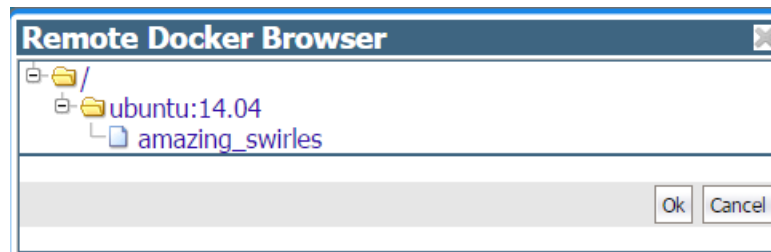
- **HDFS File System**, select this option if the resource is located on a host that is part of an HDFS cluster group.

- **Directory** is appended to the GuardPoint. If the GuardPoint is `/mnt/remote2` and the directory is `/remoteDir`, then the policy applies to the files and directories in `/mnt/remote2/remoteDir`.

- If your host is a Docker host, clicking **Browse** opens the **Remote File Browser**, select a Docker image or container, from which to select a resource.

The asterisk and question mark can be used to indicate one to many characters (*), or exactly one character (?). Directory examples:

- `sales/`: Access limited to `sales` directory under the GuardPoint.

- `*sales/`: Access limited directories that end in "`sales`" under the GuardPoint. If you omit a leading path delimiter (back slash or forward slash) in a directory name before an asterisk e.g., `*sales`, a pop-up will prompt you to select whether the directory is on a Windows or non-Windows OS, and a '`/`' or '`\`' will be added accordingly as follows; for example if you select Windows the result will be `*\sales`.

- `*/sales`: Access is granted to any directories named "`sales`" anywhere.

- The variables `|uname|` (user name) and `|gname|` (group name) can be used. On UNIX systems, `|uid|` and `|gid|` may also be used. When the security rule is applied, the variable is replaced by the actual user name or user group name.

  For example, if **Directory** is set to `/opt/local/ |gname|`, when you later make `/opt/local` the GuardPoint, only the members of the group specified as "engineering" in **Users** are allowed access to `/opt/local/engineering`.

  `uname` and `gname` are like macros. Another example: if you want to define a policy to protect all of the user directories under `/home`, you do not need to enumerate `/home/steve`, `/home/george`, `/home/Richard`, and so on. You only need to define `/home/|uname|`, When the agent evaluates the policy, it replaces `|uname|` with the actual user, so when Steve logs on, the agent evaluates the policy with `/home/steve`, and they will not be able to access `/home/george`.

  When a resource set is defined with a leading asterisk in the directory path, a leading path delimiter (back slash or forward slash) is inserted at the beginning of the string, e.g., `*/sales` results in `/*/sales`. If the **File** field is left blank, a trailing delimiter and asterisk are added to the path, the asterisk indicating that all files under that directory are included in the definition. For example, if a directory path is defined as `*/sales/*` and the **File** field is blank, it results in `/*/sales/*/*`.

  ○ **File** is the filename and can include variables or patterns.

  ○ **Include subfolders** finds all occurrences of the resource pattern under the GuardPoint and applies policy protection to them. For example, if the GuardPoint is at *a/b* and the resource is defined as *c/*.txt* and **Include subfolders** is checked, then every occurrence of **.txt* anywhere under the GuardPoint is protected (example: *a/b/c/d/*.txt*). If **Include subfolders** is not checked, then only *a/b/c/*.txt* is protected.

  d. Click **Ok** to add the resources to the Resource Set.

Figure 22-7: Add Resource Set window
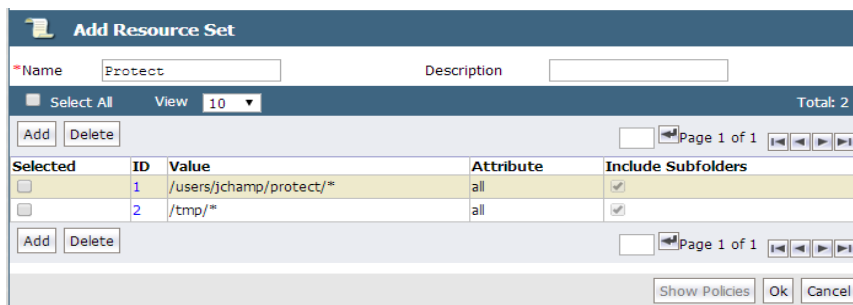
> **Note**
> You can also create or select Resource Sets directly, without first creating a policy, by clicking **Policies > Manage Policies > Resource Sets** to open the **Resource Sets** window.

  e. Once you have added all resources to your resource set, click **Ok**. The **Select Resource Set** window opens with the new resource set added.

Figure 22-8: Select Resource Set window with new resource set added
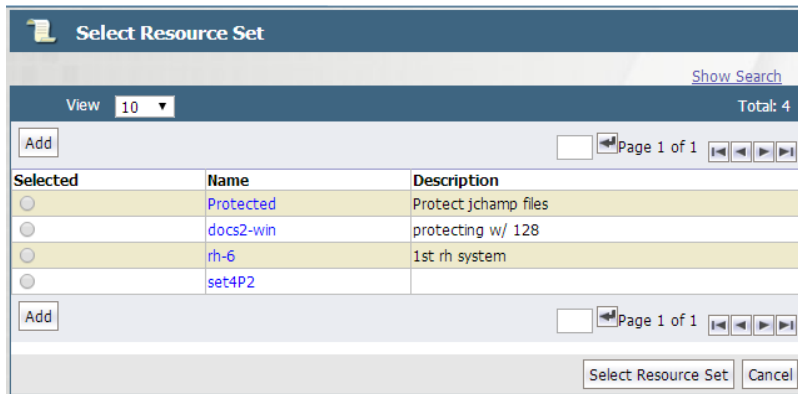
f. Select the resource set for this policy and click **Select Resource Set**. The **Edit Security Rule** window opens with the resource added (in this example, **Protected**).

g. Check the **Exclude** option to the right of the **Resource** text-entry field to include all host resources except those resources in the resource set. Uncheck the option to include just the resources in the resource set.

Adding a Security Rule User Set to a Policy

Specify the **User** criteria. **User** allows you to specify the users that are permitted or denied GuardPoint access.

To specify all users, leave **User** blank.

1. To define specific users, select **User**. The **Select User Set** window opens.

2. Click **Add** to create a User Set.

Figure 22-9: Add User Set window
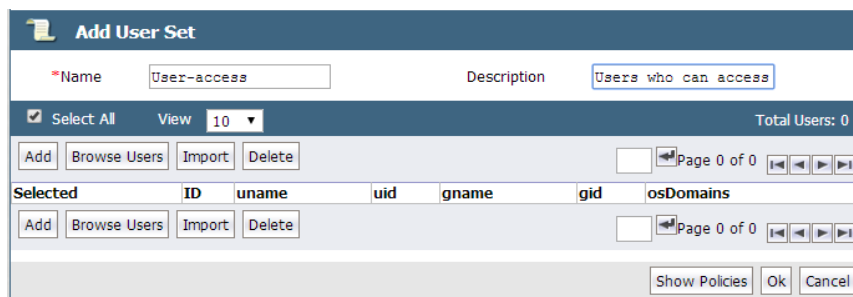
3. Enter a **Name** (for example, `User-access`) for the User Set and optionally, a **Description**.

4. Click **Add**. The Add User window opens. You must specify at least one field.

Figure 22-10: Add User window

- **uname**

  Login name.

- **uid** (*UNIX only*)

  User identification number

- **gid** (*UNIX only*)

  User group number, enter only the primary group ID number of the user

- **gname**

  Comma-separated list of group names

- **osDomain** (*Windows only*)

  Network domain of the user. Multiple domain names, separated by commas, may be entered. Enter the string **localhost** to configure a generic domain.

  If you click **Browse Users**, the Add Users page opens, you can select users from an LDAP server if configured, or from a selected host. To select users from docker images or containers, use the default Agents selection and select the host name (FQDN) of the Docker host from the list. Since this a Docker host, another field is displayed; Docker Image/Container. Click **Browse** to open the Remote Docker Browser to select a Docker image or container from which to select users.

5. Click **Ok** to add this user to the User Set.

6. Add as many users to the User Set as needed by repeating steps **c** through **e**.

   Figure 22-11: Add User Set window
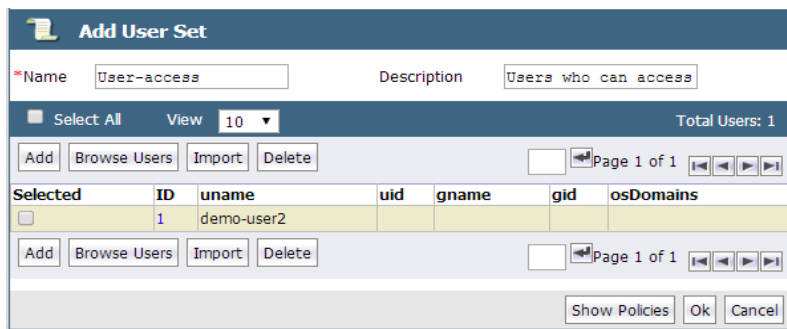
7. Click **Ok**. The **Select User Set** window opens with the new User Set added.

   Figure 22-12: Select User Set window
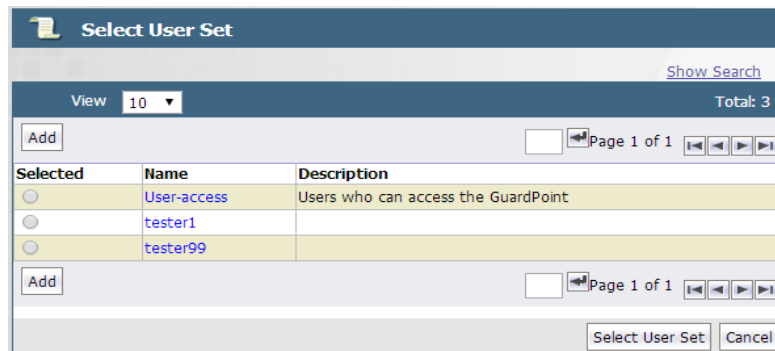
8. Select the User Set for this policy and click **Select User Set**. The **Edit Security Rule** window opens with the new User Set resource added (in this example, **User-access**).

9. Check the **Exclude** box to the right of the **User** text-entry box to include all host users except those users in the User Set. Uncheck the box to include just the users in the User Set.

10. You can also create or select User Sets directly, without creating a policy, by clicking **Policies > Manage Policies > User Sets** to bring up the **User Sets** window.

Adding a Security Rule Process Set to a Policy

Specify the **Process** criteria. A Process Set is a named collection of executable processes that are permitted or denied access to the GuardPoint data.

- To specify all processes, leave **Process** blank.

1. To define specific processes, select **Process**. The **Select Process Set** window opens.

2. Click **Add** to create a Process Set. The **Add Process Set** window opens.

Figure 22-13: Add Process Set window
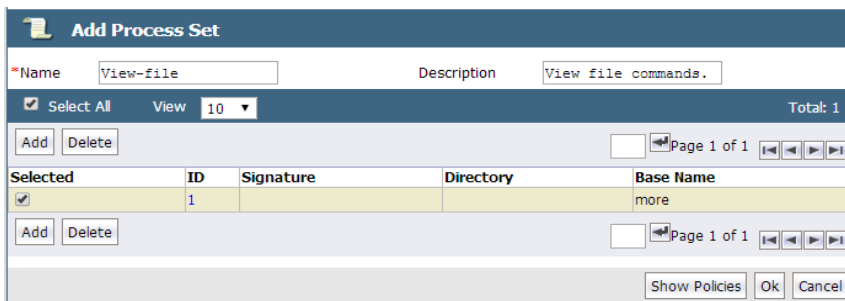
3. Enter a **Name** (for example `View-file`) and optional **Description**.

4. Click **Add**. The Add Process window opens.

Figure 22-14: Add Process window
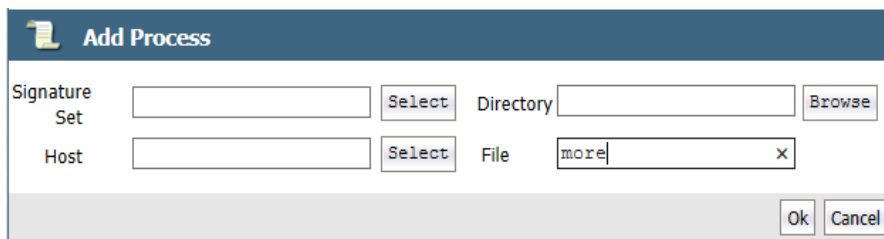
- **Signature Set**

  Collection of signed files and/or directory names. Files that are signed confirm software integrity and guarantee that code has not been altered since it was cryptographically signed. If you created a signature set to use with this policy, select the signature set from the scroll-list. Otherwise, you must first create a signature set.

- **Host**

  Host of the directory or executable and activates the **Browse** function.

- **Directory**

Directory path information. It may be a full path, a relative path, or left blank.

- **File**

    Name of the executable. The **more** command is used in the example.

5. Click **Ok** to add this process to the Process Set.

6. If you selected a Docker host, the **Docker Image/Container** field is displayed, select a Docker image or container, in the **Directory** field, click **Browse** and select a directory from the **Remote File Browser**, fill in the file name field as required. Click **Ok** to return to the Add Process Set page.

7. If you get the pop-up shown below, click **Windows** for Window hosts and **Non-Windows** for Linux hosts.

Figure 22-15: OS Type Pop-Up dialog
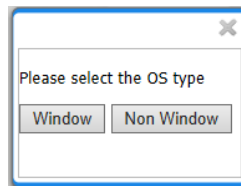
8. Repeat the previous steps to add additional process sets, if needed.

9. Click **Ok**. The **Select Process Set** window opens with the new Process Set(s) added.

Figure 22-16: Select Process Set window
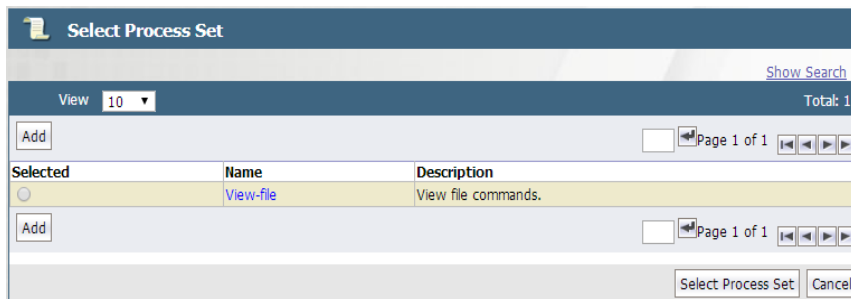
10. Select the Process Set for this policy and click **Select Process Set**. The **Edit Security Rule** window opens with the new Process Set added (in this example, **View-file**).

11. Check the **Exclude** option to the right of the **Process** text-entry option to include all host processes except those users in the Process Set. Uncheck the option to include just the processes in the Process Set.

> **Note**
> You can also create or select Process Sets directly, without creating a policy, by clicking **Policies > Manage Policies > Process Sets** to bring up the **Process Sets** window.

Adding a Security Rule Time Set to a Policy

Specify the **When** criteria. **When** enables you to specify when GuardPoint access is allowed.

> **Note**
> Policies for Cloud Object Storage do not contain a When/Time set.

- To specify all times, i.e. 24-hour access, leave **When** blank.

To define specific allowable times, select **When**. The **Select Time Set** window opens.

1. Click **Add** to create a **Time Set**. A Time Set is a named collection of times when GuardPoint access is permitted or denied. The **Add Time Set** window opens.

    Figure 22-17: Add Time Set

2. Enter a **Name** (for example **Time-1**) and optional **Description**.

3. Click **Add**. The **Add Time** window opens.

    Figure 22-18: Add Time window

- **Week Day From - To** is a range of days of the week during which access is denied or permitted. Values are Sunday through Saturday. Enter a day of the week to begin allowing access and a day of the week to stop access.

- **Date From - To** is a range of dates during which access is denied or permitted. Enter a calendar dates to from when to begin allowing access and when to stop.

- **Start Time - End Time** is a range of times during which access is denied or permitted. Enter a start time in the format hh:mm, select AM/PM, to allow access, and an end time in the format hh:mm, select AM/PM. This defines the exact start time and end time during which access is permitted.

    > **Note**
    > Time is set to the protected host clock, not the DSM clock.

4. Click **Ok** to add this time range to the Time Set.

5. Repeat the previous steps to add additional time sets, if needed. Add as many times to the Time Set as needed.

Figure 22-19: Add Time Set window
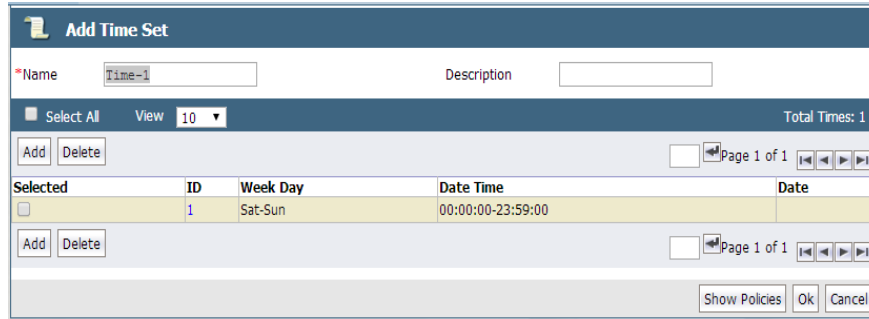
6. Click **Ok**. The **Select Time Set** window opens with the new Time Set added.

7. Select the Time Set for this policy and click **Select Time Set**. The **Edit Security Rule** window opens with the new Time Set resource added (in this example, `Time-1`).

8. Check the **Exclude** box to the right of the **When** text-entry box to include all times except those times in the Time Set. Uncheck the box to include just the times in the Time Set.

9. You can also create or select Time Sets directly, without creating a policy, by clicking **Policies > Manage Policies > Time Sets** to bring up the Time Sets window.

Adding a Security Rule Action Set to a Policy

An **action** is an attempt to access protected data in some way. **Action** allows you to specify the type of file and direct action allowed in a GuardPoint. It can be an attempt to read, delete, rename, etc. A Policy provides a wide range of access methods to specify the action required for the security rule. The action is identified, evaluated, and used as a factor in determining whether or not to grant access. Specifying an access method is optional. If you do not specify one, the Policy defaults to all_ops, or "all operations".

Any action that starts with `d_` applies to directories. Any action that starts with `f_` applies to files. Only a few actions are available for COS (Cloud Object Storage) and IDT (In-place Data Transformation - Device) policies.

Table 22-3: Action Definitions

Action	Description	Policy Types
all_ops	All operations. User can access the data using any method. This is the default if no action is specified.	All
d_chg_att	Change directory attributes (e.g., chown usr dir1).	Standard, LDT
d_chg_sec	Change any security property of a directory.	Standard, LDT
d_mkdir	Make a new directory.	Standard, LDT
d_rd	View directory contents (e.g., ls dir; cd dir).	Standard, LDT, COS
d_rd_att	Read the attributes of a directory (e.g., ls -la dir1).	Standard, LDT
d_rd_sec	View the security properties of a directory.	Standard, LDT
d_ren	Rename a directory (e.g., mv dir1 dir2).	Standard, LDT
d_rmdir	Delete a directory (e.g., rm -r dir1).	Standard, LDT

Table 22-3: Action Definitions (continued)

Action	Description	Policy Types
f_chg_att	Change file attributes (e.g., chmod).	Standard, LDT
f_chg_sec	Change any security property of a file.	Standard, LDT
f_cre	Create a new file.	Standard, LDT
f_link	Link to a file (e.g., ln file_name link_name).	Standard, LDT
f_rd	Read a file.	Standard, LDT, COS
f_rd_att	Read the attributes of a file (e.g., ls -l file).	Standard, LDT
f_rd_sec	View the security properties of a file.	Standard, LDT
f_ren	Rename a file.	Standard, LDT
f_rm	Remove/delete a file.	Standard, LDT, COS
f_wr	Write to an existing file.	Standard, LDT, COS
f_wr_app	Append data to a file.	Standard, LDT
key_op	Key operations. This requires 2 key rules. One key is specified in the Key Selection Rules tab and the other in the Data Transformation Rules tab. Key Selection keys are the encryption keys for the current security rule. Data Transformation keys are used to migrate data from an encrypted form to a non-encrypted form or to change the encryption keys used to access the data.	Standard, LDT
read	Read operations. This method is a collection of the preceding methods that are related to reading files, directories, and their attributes. This method comprises f_rd, f_rd_att, f_rd_sec, d_rd, d_rd_attr, and d_rd_sec.	Standard, IDT, LDT
write	Write operations. This method is a collection of the preceding methods that are related to writing to files, directories, and their attributes. This method comprises f_wr, f_cre, f_ren, f_rm, f_link, f_chg_attr, f_chg_sec, d_ren, d_chg_att, d_chg_sec, d_mkdir, and d_rmdir.	Standard, IDT, LDT

Specify the **Action** criteria:

1. Click **Select** next to the Action field. The Select Action window opens.

2. To specify all actions, select **Select All**.

3. To specify allowable actions, select the option for that action.

4. When finished, click **Select Action**. The **Add Security Rule** window displays with the allowable actions listed in the Action field.

> **Note**
> The action, **key_op**, is used for the `dataxform` policies. If you select `key_op` and return to the Add Policy window, the Data Transformation Rule field displays. You must add a Data Transformation rule.

Adding a Security Rule Effect to a Policy

Specify the **Effect** for each security rule. **Effect** is the action that occurs when the attempted access matches all the criteria in the rule.

1. Click **Select** next to the Effect field. The **Select Effect** window opens.

   Figure 22-20: Select Effect window
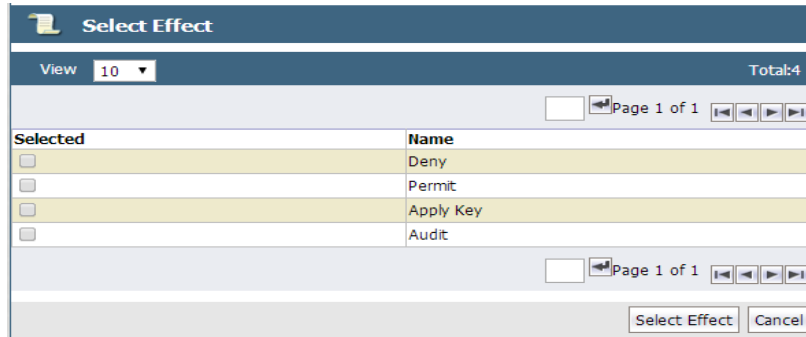
   - **Deny**

     Denies the access attempt to the resource.

   - **Permit**

     Grants the access attempt to the resource.

   - **Apply Key**

     Applies an encryption key to data in a GuardPoint. Data copied into the GuardPoint is encrypted with the key specified in the **Key Selection Rules** panel and data that is accessed in the GuardPoint is decrypted using the same key.

     > **Note**
     > If you select **Apply Key**, you must also specify the key rules to apply for encrypting and decrypting the resources.

   - **Audit**

     Used in conjunction with **permit** or **deny**, **audit** creates an entry in the Message Log that describes what is being accessed, when it is being accessed, and the security rule being applied.

2. Select desired effects and click **Select Effect**. The Edit Security Rule window opens with all criteria and the effects displayed.

Figure 22-21: Edit Security Rule window.

3. Click **Ok**. The **Add Policy** window opens.

Default Security Rule

The last rule of the policy is called a default security rule or a catchall rule. This rule catches any access attempt that is not matched by other security rules.

To create a default security rule, leave all criteria fields blank with the exception of **Action** and **Effect** on the **Security Rules** tab:

- Set **Action** to **all_ops**
- Set **Effect** to **deny audit**

This security rule will match any attempt to access any data on the host. After creating this rule, click **Ok** in the **Edit Online Policy** window.

The **Policies** window opens and the policy you just created can be applied to a GuardPoint.

Add Key Selection Rules

After setting up the Security Rules, set up your Key Selection Rules.

⚠ **WARNING**
　Be careful when changing a VTE policy that has already been implemented. The following practices may cause data corruption in certain scenarios:

　* **Changing the order of key rules**

　* **Modifying a key rule**

　* **Adding a new key rule**

　* **Deleting a key rule**

　* **Modifying a resource set**

　* **Modifying a resource**

1. Click **Add** in the **Key Selection Rules** panel at the bottom of the Add Policy window.

   The **Add Key Rule** window opens.

Figure 22-22: Add Key Rule window

2. Click **Select** and select a **Resource Set**, if desired.

   A key resource set lets you specify files and directories that are to be encrypted. If you do not specify a resource set in the **Key Selection Rules** tab, encryption is applied to the resources specified in the **Security Rules** tab.

3. Select **(Current) Key**. The **Select Symmetric Key** window opens.

   ° For an **LDT** policy, the **current** key specifies the key applied to existing data prior to the application of the LDT policy. The current key can be a clear key or a non-versioned key. When a policy is applied to protect data, LDT uses the current key to transform the data to the current version of the Transformation Key.

   ° For a **Standard** policy, you can enter a key name, or click **Select**, and the **Select Symmetric Key** window opens allowing you to select an existing non-versioned key.

   ° For an **Efficient Storage** policy, you must select an AES key with XTS encryption that is KMIP accessible.

   ° For an **In-place Data Transformation Device (IDT)** policy, you must select an AES key with XTS encryption.

   > **Note**
   > The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

   ▪ **New Key**: Key that replaces the current key. The new key becomes the current key.

   ▪ **Current Key**: Key that is currently in use. Note that you can choose clear-key for the current key.

   > **Note**
   > An IDT policy requires two key rules. However, you can only add one key rule, with two keys, in the UI. Therefore, when you create an IDT policy, DSM converts the key rule into two key rules so that the policy contains the required two key rules. (The API does not have this limitation.)

   ° For a **Cloud Object Storage** policy, it can only contain one key. Enter a key name, or click **Select**, and the **Select Symmetric Key** window opens allowing you to select an existing CBC-CS1 key.

   > **Note**
   > For COS policies, only encryption mode CBC-CS1 is supported. CBC is not supported for cloud object storage.

4. If this is a key rule for an LDT policy, and you are not excluding files, click **Select** to open the **Select Symmetric Key** window, and select a **Transformation Key** on the Add Key Rule page.

   The Transformation Key is the versioned key applied to data for the initial transformation from current key and subsequent rekeying to the next version of Transformation Key.

   (See "Creating and Adding symmetric keys" on page 175, for more information about LDT versioned keys.)

5. Select **Exclusion Rule**, if needed.

   The Exclusion rule allows you to exclude the files or directories, associated with a specific resource set, from being transformed.

   ° You cannot exclude files when using a versioned key for the current key or transformation key.

- If Exclusion Rule is selected, the Resource field becomes **mandatory**. DSM can only exclude files associated with a resource set.

- The two keys must be identical. Therefore, when you select the Exclusion Rule, the Transformation Key and Current key are automatically set to the same value and the Transformation Key field is hidden.

- User cannot add an Exclusion Rule to any policy that was created in a version prior to v6.4.0 of the DSM.

- You can clone a policy created in a previous version of DSM, and then modify the cloned policy to use an LDT exclusion key rule policy.

- The Exclusion rule is only compatible VTE v6.2.0 and subsequent versions.

- Policies for Cloud Object Storage do not support Exclusion rules.

6. Once you've selected the keys, click **Ok** to return to the Edit Policy page.

Displaying Policies

Policies are displayed in the Policies window. Policies displayed can be selected for modification or deletion.

Display Policies

1. Log on to the Management Console as an Administrator with `Policy` role permissions.

2. Select **Policies** in the menu bar. The Policies window opens. Configured policies are displayed.

Policy History

The number of times a policy has been changed is displayed in the **Version** column of the **Policies** window. This number indicates the current revision only. (You cannot use it to roll-back to a previous version. Restore a backup to revert to a previous online policy version.) The version count starts at zero when the online policy is initially created and increments by one each time it is saved thereafter. Click the policy version number in the **Version** column to view the version history of a policy.

Figure 22-23: Policy version history

Customize display in the Policy window

- **Show Search**

  Located below the **Policies** banner; this opens the **Search** panel. You can enter all or part of a policy name and/or limit the search to policies that are used by a specific type of agent (All or FS). Click **Go** to display only those policies that match the search criteria. Click **Hide Search** to conceal the **Search** panel.

- **Select All**

Selects all the policies that are displayed on the current Web browser page. Select this checkbox to select all the policies on the current page at one time. If you have enabled the **Select** checkbox for many individual policies, a quick way to deselect them is to enable and then disable the **Select All** checkbox.

- **View**

A scroll-list from which to select the maximum number of policies to display on the current page. Up to 200 policies can be displayed on one page. Displays up to the specified number of policies on one Web page, regardless of the display number specified in the preferences.

Navigation buttons are displayed in the **Policies** window. Use these buttons to advance between pages. The buttons are shown below:

Table 22-4: Policy Window Panel Navigation Buttons

⏮	First. Display the first page of policies in the Policy window.
◀	Previous. Display the previous page of policies in the Policy window.
▶	Next. Display the next page of policies in the Policy window.
⏭	Last. Display the last page of policies in the Policy window.
⏎	Jump to. Advance to the specified page of information. Enter the page number in the text-entry box that is next to this button.

Exporting and Importing Policies

Policies can be exported and then imported to the same or another DSM where you want to replicate the policies you've already created.

You can choose to export all policies from a DSM or a specific policy. Policies are generated as xml files and then compressed and exported to a .tar file with the following naming convention; *policy_<YYYY_MM_DD_HHMM>.tar*.

> **Note**
> Editing the exported policy file is strongly discouraged. This can cause issues with policy enforcement.

When policies are imported to a domain, all the sets (resource sets, user sets, process sets, and time sets) are imported with the following conditions:

- If a set in the imported policy does not exist in the domain that policy is being imported to, then that set is created.

- If a set in the imported policy exists in the domain to which that policy is being imported, then the existing set is overwritten.

- If a name of a policy being imported conflicts with a name on the domain where it is being imported to, then a number is appended to the name before it is imported. For example, if the imported policy and a policy in the domain to which the policy is being imported both contain a policy named 'secure_file_policy' then the policy will be imported as 'secure_file_policy_1'.

Export a policy

If you choose to export only some specific policies, then only those Resource sets, Process sets, User sets, and Time sets, and associated action and effects used by those policies are exported.

> **Note**
> The second option is only available in DSM v6.0. Earlier versions of the DSM software export all Resource sets, Process, sets, User sets, and Time sets and the associated actions and effects from the domain, even if they are not used by the policies you choose to export.

1. Log on to the Management Console as an All/Security administrator.

2. Click **Policies > Import Export Policies** on the Management Console.

3. On the **Export** tab, select the policies that you want to export.

4. Click **Ok**.

5. The policy export file is exported as `policy_<YYYY_MM_DD_HHMM>.tar`, follow the prompts to save the file to your preferred location.

Import a policy

1. Log on as an All /Security/Domain and Security administrator.

2. Click **Policies > Import Export Policies** on the Management Console.

3. On the **Import** tab, click **Browse** to locate the policy file to import.

4. Click **Import Policy**.

The DSM performs pre-import checks on the policies to make sure that there are no conflicts or missing items.

- If there are no conflicts or missing keys the import proceeds and the Resource sets, User sets, Process sets, and Time sets are imported. A message confirming that the operation was successful is displayed on the **Import** tab.

In the event of conflicts or missing keys are detected any the following could occur:

- If the pre-check process finds that the policy keys are missing, the import is aborted and a message informing you that the operation failed is displayed and the **Messages** text box on the **Import** tab provides the names of the missing key(s).

- If the policy or policies you import contain Resource sets, User sets, Process sets, or Time sets that have names that match existing policies on the DSM to which they are being imported, or the policy or policies being imported have the same names, the **Messages** text box will contain a message listing the imported sets that conflict with existing sets. You can choose to **Continue** or **Abort** the import operation.

  If you choose to continue, the existing policy will retain it's name and the imported policy will have '_1' or the relevant number in sequence appended to the name. For example, if you are importing policies to DSM B and it has an existing policy called 'policy1' and the imported policy has the same name, if you choose to continue the import operation, the imported policy will be rename 'policy1_1'.

- If the pre-check operation detects that there unused sets referenced in the policies, you will be prompted to do either of the following:
  - Select **Policies & associated sets**, which means only sets that are used will be imported.
  - Select **Policies & all sets**, which means all sets, regardless of whether they are referenced by the policy or not will be imported.

    You can choose to **Continue** or **Abort** the import operation.

- If you choose **Policies & all sets** and the pre-check finds there are conflicts with exiting sets, you can again choose to continue or abort the operation.

The **Policies & all sets** option is useful when importing policies from earlier versions of the DSM. Earlier versions of the DSM always exported all sets regardless of whether they were used by a policy or not.

Chapter 23: Configuring Hosts and Host Groups

A "protected host" is a computer system on which Agents (VTE/VAE/VTS) are installed. The agent on a host may protect data on that host, or data on other devices connected to that host.

Overview

The Hosts page on the Management Console displays all hosts protected by encryption Agents. DSM Security Administrators manage hosts via this page; hosts can be added, imported, or deleted.

Refer to the *VTE Agent Installation & Configuration Guide* for information about installing and configuring a VTE Agent.

Viewing Hosts

To see all protected hosts registered with a DSM:

1. Log on to the Management Console as an All /Security/Domain and Security administrator.

2. If you log on as type All, click **Domains > Switch Domains**.

   a. In the **Selected** column, click the radio button for the domain you want, then click switch to domain. Skip to step 4.

3. Navigate to **Hosts > Hosts**.

4. Click **Hosts > Hosts**. The Hosts page has a table listing names of the protected hosts in the DSM and the following details about each protected host:

Table 23-1: Hosts Window Table Details

Column	Description
Select	Select this checkbox to select the host for deletion. Multiple check boxes can be selected at one time.
OS Type	Values may be `Unknown`, `AIX`, `HPUX`, `Linux`, `Solaris`, or `Windows`. `Unknown` indicates that the host has not been registered or is an unsupported type.
Host Name	The name of the host on the DSM.

Table 23-1: Hosts Window Table Details (continued)

Column	Description
VTE/Key/KMIP Agent	This column consists of child columns of check boxes for the VTE Agent and for VAE and KMIP: - **Reg. Allowed**—Registration Allowed, indicates that the host can register and be configured to run VTE (File System) Agent software. - **Comm. Enabled**—Communication Enabled, indicates that a policy can be applied to a host. **Reg. Allowed** must be enabled before you can set **Comm. Enabled**. - **Pushing Status**—Status for pushing policy and configuration changes to locally assigned hosts. Status is specific to the local DSM. Run the Management Console on an HA node DSM to see the push status of the hosts assigned to that server. **> Done**—Host has the latest policy and configuration changes. **> Pending**—Update is in progress or is queued for download to the host. **> N/A**—Local host is disabled or the host is being administered by a different server. **Th**is last column is not available for VAE and KMIP.
One Way Communication	Indicates that the agent was registered with One-Way Communication enabled between the Agent and the DSM.
Delete Pending	Indicates the status of a request to delete a host.
License Type	Perpetual, Term, Hourly. See Chapter 3: "Licensing" on page 30 for information about the different license types.
LDT Enabled	Indicates whether the Live Data Transformation (LDT) feature is enabled on the host. If this feature is enabled, Docker support cannot be enabled.
Docker Enabled	Indicates whether support for Docker feature is enabled on the host. If this feature is enabled, LDT cannot be enabled.
Efficient Storage Enabled	Indicates that Efficient Storage is enabled. This feature requires KMIP and a Pure Storage Flash storage array.
Description	(Optional) Text to help you identify the host.
Sharing	Indicates if the host is shared with another domain. The column may have a value of `Shared`, `External`, or blank. - `Shared` indicates that the host is in the current domain and, if the Security Administrator has the correct roles, it can be fully configured. The shared host is visible only in the domain with which it is being shared. - `External` indicates that host is administered in another domain. You can assign the host a VTE Agent GuardPoint, but you cannot change the host configuration. - A blank value indicates that the host is not shared.

The first time you log in, the list of hosts is empty because you have not yet registered any hosts with the DSM.

The following tasks can be done from this page:

- **Select All**

  Selects all hosts displayed on the current page.

  If you have enabled the **Select** check box for many individual hosts, a quick way to cancel the selection is to enable and then disable the **Select All** check box.

- **View**

  Specifies the number of hosts to display on the current page. Up to 200 hosts can be displayed on one page.

- **Search**

The **Show Search** label located below the **Hosts** banner opens the **Search** panel. You can enter a string and/or search for a specific type of agent, and click **Go** to display the hosts that match the search criteria. Click **Hide Search** to conceal the **Search** panel.

Navigation buttons are displayed in the **Host** window. Use these buttons to advance between pages.

- **Add**

    Click to create a new host record.

- **Delete**

    Enable the check box in the **Select** column for one or more hosts and click **Delete** to remove the selected host from the DSM database. The agent installation is left intact on the host system and needs to be uninstalled from the host if required.

- **Import**

    Click **Import** to select a configuration file to add multiple hosts in a batch operation. When hosts are added in batch files, the **License Type** defaults to Term.

Adding Hosts to the DSM

You can add Hosts to a DSM manually via the Management Console, or automatically through the Shared Secret Registration method. You can also add a registration token to the shared secret for multi-factor authentication. Agents on the host are registered with the DSM using either the Fingerprint Registration method or the Shared Secret Registration method. The Shared Secret method is the default.

This section describes the following:

- "Adding hosts using a shared secret" below

- "Adding hosts using a certificate fingerprint" on page 233

- "Adding hosts using a batch file" on page 235

Adding hosts using a shared secret

The Shared Secret method requires a DSM Administrator to create a registration password for a domain or host group. This password is shared with the Agent Installer, which uses this password to add and register protected hosts with the DSM in a single step. There is no need to manually add hosts to the DSM before registering the agent. Adding a host before registering it using the shared secret method, is optional. Multiple protected hosts can be added with a single shared secret password. As of this release, DSM Administrators can enforce the shared secret registration method for all hosts, by selecting the option on the Web UI.

The DSM only allows hosts that know the secret to register. The agent in turn, knows that it is registering with the correct DSM because it has the same secret. Hosts can be added to a domain, or to a host group within a domain, which means that a shared secret can be defined at the domain level or the host group level.

Add a host to a domain using Shared Secret Registration

When you use the registration secret feature for the first time, the **Current Registration Secret** section does not contain any information. If there is an existing shared secret, then **Show Registration Shared Secret** displays. You can display the secret criteria, if needed.

The first step is for the **Security Administrator** to create a registration shared secret

1. Log on an administrator with Host role permissions.

2. Switch to the domain to which you want to add the host.

3. Select **Hosts > Registration Shared Secret** in the menu bar. The Registration Shared Secret window opens.

4. Enter the following information in the **Create new Registration Shared Secret** section:

   a. **Registration Shared Secret creation method**—The same constraints that apply to password creation, namely uppercase letters, numbers, and special characters required, apply to the shared secret creation.

      - **Manual**—This is the default method. Select this to create the shared secret yourself.

      - **Generate**—Select this option to get an automatically generated password.

   b. **Validity Date**—Enter a date, or select a date by clicking the calendar icon. The date must be in the format MM/DD/YY.

   c. **Require that hosts are added first**—Optional. If you select this option, you need to first add the host to the DSM database with the **Registration Allowed** check box enabled before you install and configure the agent.

   d. **Enforce shared secret during host registration**—Optional. If you select this option, hosts **must** register with the DSM using the shared secret. Any attempt to register a host using the fingerprint method will fail.

5. Click **Ok**.

6. To remove an existing shared secret, click **Expire Registration Shared Secret**. The expiration date turns red to indicate that the shared secret is no longer valid.

The Account Lockout settings defined in **General Preferences > Password > Account Lockout** also apply to the registration shared secret, see "Account Lockout" on page 50 for more information about these settings.

Creating and Using a Registration Token

The shared secret is shared by all of the hosts in a specific domain, or a host group. If the shared secret is somehow compromised, it must be regenerated (or re-entered) and this could disrupt many other agent hosts. Therefore, the DSM now provides a second authentication credential, a token, that is tied to a specific host or a host group. Only that host/hostgroup can use that token to register to a DSM domain. This provides two-factor authentication for agent registration.

> **Note**
> This feature is only compatible with VTE v6.3.0 (or greater).

To create a Registration Token:

1. Log on as an administrator with Host role permissions.

2. Switch to the domain to which you want to add the token.

3. Select **Hosts > Registration Shared Secret > Registration Shared Secret** in the menu bar. The Registration Shared Secret window opens.

4. Select the option for **Registration Token Enabled**.

5. Enter the host name for the host to which to assign the token.
   - If you want to assign the token to a group of hosts, use a wild card. (For example: *.ca.amazon.com)

6. Set the expiration date and time.

7. Click **Generate Registration Token**. The DSM generates a token and downloads it to your downloads drive.

8. Click **OK** to save the configuration.

9. Open the token file with a text editor and save it as a text file called `dsm_reg_token`.

10. Send the token to the VTE administrator. They must install the token on the agent system prior to registration with the DSM.

11. To install, the VTE admin must copy the file into the following location:

- For **Linux/AIX** systems: `/etc/vormetric/dsm_reg_token`

- For **Windows** systems: `C:\Windows\Temp\Vormetric\dsm_reg_token`

  When the agent runs the registration code, the token is read and submitted as part of the registration payload.

12. When finished, click **OK** in the Registration Shared Secret Window.

Select a license type for the host

The DSM Administrator must select a license type for the host:

1. Select **Hosts > Registration Shared Secret > Default License Type**.

2. Select the appropriate license for FS agents (VTE) or Key agents (VAE).

3. Click **Ok**.

In the event that the license type selected for the agent registration does not contain anymore licenses, the registration process will try registering the agent with the next available license type with available licenses.

> **Note**
>
> If you add a host to the DSM before registering it using the shared secret method, and if you change the **License Type** setting on the Add Host page, the setting on the Add Host page takes precedence over the setting on the Default License Type for Registration Shared Secret page.
>
> For example, if you set license type on the Default License Type page to TERM, then on the Add Host page, while adding the host entry, you select License Type as PERPETUAL, then when you register the host using the shared secret method, the license type for that host will be PERPETUAL.

Register the host on the DSM

After the agent has been installed, you will be prompted to register the host to the DSM:.

> **Note**
>
> The exact sequence of steps may differ from agent to agent, for details about how to install specific agent types, refer to the *VTE Agent Installation & Configuration Guide*.

1. You will be prompted to select a method to register the host. Select the shared secret option to register the host. This is the default option.

2. Enter the following information when prompted:

   a. What is the registration shared secret?

      Enter the shared secret of the domain to which you the host is to be added. Or enter the shared secret of the host group to which the host is to be added.

> ⚠ **WARNING**
>
> **Be sure to enter the shared secret correctly, the prompt will not display any entered text, nor does the prompt move until you press enter. If the shared secret was entered incorrectly, an error message is displayed, saying that the certificate signing was unsuccessful. If you exceed the number of tries defined in the Maximum Number of Login Tries setting on the Password Preferences page, you will be locked out of the system for a period defined in the User Lockout Time setting on the same page.**

   b. Domain name

Enter the name of the domain to which the host is to be added.

c. **Host Group** (Optional)

If the host is to be added to a host group, enter the name of the host group to which it is to be added, else click enter or next and continue to the next step.

d. **Host description** (Optional)

Enter a description of the host to be registered.

3. Confirm the information is correct and proceed with the registration.

4. Open the Management Console on the DSM, switch to the domain where the host has been added, the host should be listed in the hosts table.

   If the host was added to a host group, select **Hosts > Host Groups** and click the host group where the host has been added, the host should be visible in the table.

Create a Registration Shared Secret for a HostGroup

When adding a host to a host group in a domain, the Security Administrator must first create a registration shared secret:

1. Log on an administrator with Host role permissions.

2. Switch to the domain to which you want to add the host.

3. Select **Hosts > Host Groups** in the menu bar. The Host Groups window opens.

4. Click a host group name or create a host group where the host is to be added and click the host group name. The Edit Host Group page is displayed. Click the **Registration Shared Secret** tab.

Figure 23-1: Host Group Registration Shared Secret window
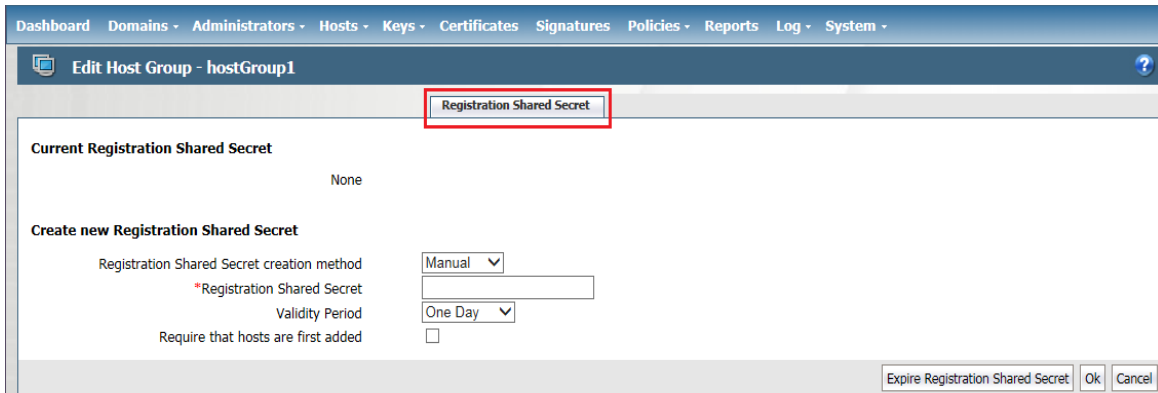
The remaining steps to create a registration shared secret for a hostgroup are the same as "Add a host to a domain using Shared Secret Registration" on page 230.

Adding hosts using a certificate fingerprint

The Fingerprint Registration method requires you to first add the host name or its IP address to the DSM from the Management Console. Once the host is added to the DSM, you can register the host from the Agent Installer on the host.

To add hosts to the DSM:

1. Determine the manner in which you want to address the host. That is, FQDN, host name, or IP number.

- If FQDN, verify that DNS is configured and working on the DSM.

- If host name, use the host CLI command to link IP numbers with host names, or edit `/etc/hosts` directly.

2. Log on to the Management Console as a Security administrator with **Host** role permissions, or a Domain and Security, or All administrator.

3. Switch to the domain where you will add the host.

4. Select **Hosts > Hosts** in the menu bar. The Hosts window opens.

Figure 23-2: Hosts window
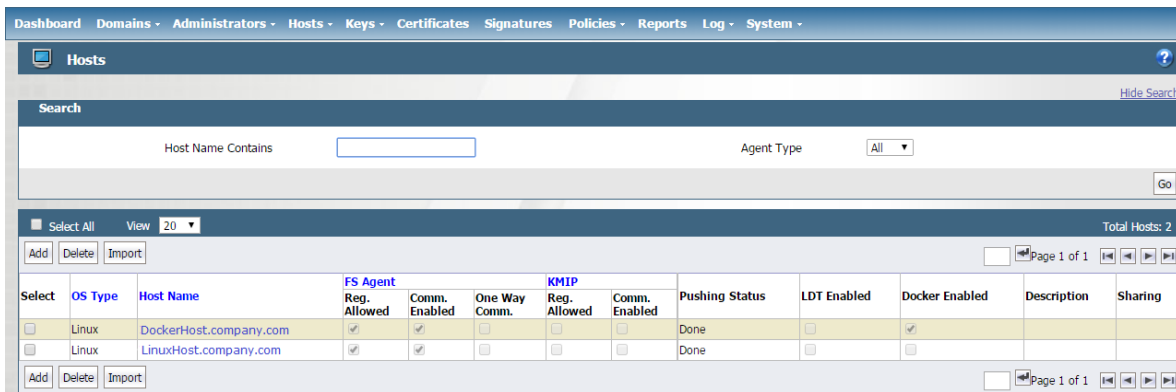
5. Click **Add**. The Add Host window opens.

Figure 23-3: Add Host window
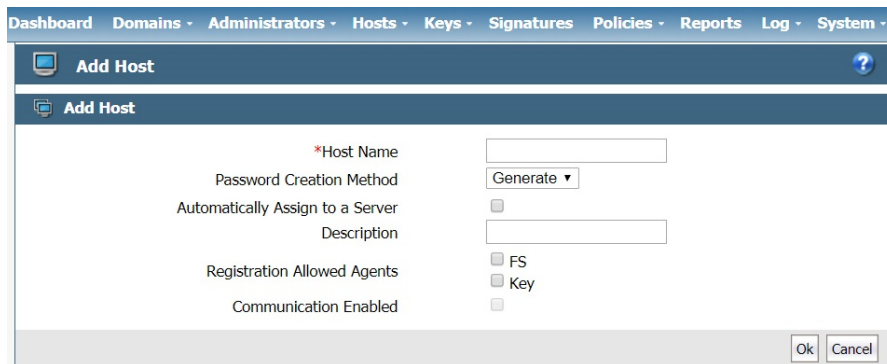
6. Enter the following information:

   a. Enter the IP address, host name or FQDN (253 characters max.).

   > **Note**
   > Host names that include an underscore are rejected by the Management Console.
   > Host names that have a dot ('.') appended to them, prevents the agent configuration log files from being uploaded to the DSM. However, if your hostname does contain a dot appended to it, then you must re-register that host using the host IP address and then upload the log files.

   b. Select a **Password Creation Method.** This is the password you use to unlock a GuardPoint when there is no server connection.

   ○ **Generate** (challenge-response)

Dynamic password. Each time a host password is required, the Security Administrator requests a new password from a DSM Administrator.

- ○ **Manual**

  Static password that is entered each time a host password is required. Select **Manual**, then enter and re-enter the password in the **Password/Confirm Password** fields.

c. Select to **automatically assign the host** to a DSM during host registration. Automatic host assignment is a load-balancing function in an HA cluster. If servers in the HA cluster are physically distributed over great distances, you may not want to use this option because hosts can be assigned to distant servers with slow connections. By default, hosts are assigned to the initial DSM when they are added. If you leave it unchecked, you can specify the server to explicitly assign the host later in the **Hosts for High Availability Server** window.

d. Enter text that helps you to identify the host in the description field. The maximum number of characters is 256.

e. Choose the type of **license** that will run on this host. Options are **Perpetual**, **Term**, and **Hourly**, depending on the system license. For instance, if only a Term license is installed on the system, only 'Term' appears in the box.

f. **S**elect the agents that will run on the host system. Depending on your license, your choices are **FS** (VTE), and **Key** (VAE). The agent must be selected here before you can register that agent with the DSM. Only the agents you have a license for will display here.

g. Select **Communication Enabled** to enable communication between the DSM and the agent. This can also be done later by going to the Edit Host page.

7. Click **Ok**.

8. The host administrator (with root access) installs the agent software on the host.

   You can manually configure agent certificates later if the certificate generation and exchange phase of agent software installation fails.

Adding hosts using a batch file

You can add multiple hosts to the DSM simultaneously. After they are added to the DSM, you can install and register the agents that run on those hosts.

> **Notes**
> - This batch process does not reduce the time it takes to add individual hosts to the DSM. It only makes it easier to add many hosts by reducing key strokes and permitting unattended operation.
> - The batch input file is not verified as it is read. If there are errors in the batch input file, such as malformed passwords or inappropriate characters, this operation can fail and hang and no hosts will be added.

The lines in a batch file are individual host definitions that follow an identical format. Each line is a comma-separated list consisting of six fields. If you want to include a comma as part of the field value, enclose the whole field in double-quotes ("). Do not enclose other special characters, such as the colon (:), in double-quotes.

The format of a batch file line is:

```
hostname,description,password,allow_fs_agent
```

where:

- **Hostname**

  An alphanumeric string that represents the host name or FQDN of the host being added. This is the network identity of the host.

> **Note**
>
> Do not enter a host name that contains the underscore character (_). Host names that include the underscore character are rejected by the Management Console.

- **Description**

  A text string that describes the host.

- **Password**

  This is not a regular login or user password. This is the host password to be used by the host system to decrypt cached keys when the DSM is not accessible. The host must also be configured with **Cached on Host** keys.

- **Allow_fs_agent**

  A boolean string that is either "yes" or "no" to enable or disable VTE Agent registration.

  An example batch file is shown below:

  ```
  host1,This is host 1,onlyMe78,yes,yes
  host2,This is host 2,bobsNum1,yes,yes
  host3,This is host 3,goOd4U678,no,yes
  host4,This is host 4,some1Else,yes,no
  host5,This is host 5,qwerty123,no,yes
  host6,This is host 6,ooPB2AUoo,no,yes
  ```

> **Note**
>
> When hosts are added in batch files, the License Type defaults to Term.

To add hosts using a batch file:

1. Create the batch file as described above.

2. Log on an administrator with Host role permissions.

3. Select **Hosts > Hosts** in the menu bar. The **Hosts** window opens.

4. Click **Import**. The Import Hosts window opens.

5. Click **Browse** next to the **Import Hosts File** text-entry box.

6. Navigate to, and select, the batch file from the Choose File to Upload window.

7. Click **Open**.

8. Click **Ok** in the Import Hosts window.

   Wait until the following message is displayed: `The operation is successful.`

9. Click **Hosts** on the Management Console menu bar to display the Hosts window and the newly added hosts.

Agent software can now be installed on these systems and the agents can be registered with the DSM.

Configuring Hosts

After adding and registering hosts with the DSM, you need to configure the new host.

1. Select **Hosts > Hosts** in the menu bar. The **Hosts** window opens.

2. Click the link in the **Host Name** column of the host you want to modify. The **Edit Host** window opens to the **General** tab.

> **Note**
>
> If the **Edit Host** window displays only two tabs, **General** and **GuardPoints**, it means you are working with a shared host. Check the host status in the **Sharing** column of the Hosts window. If you want to do more than add or remove GuardPoints, switch to the domain in which the host was created.

The following host attributes are displayed, some of the fields can be modified:

- **Name**

  FQDN of the host

- **Description**

  Add or modify a description of the host.

- **OS Type**

  Operating system on the host

- **Communication Port**

  You can change the port number used to exchange policy enforcement data between the DSM and the VTE (FS) Agent. Generally, you change the port number only when the default port number is already in use or if your firewall requires a different port number.

  If you change the port number, click **Ok**. The configuration change is downloaded to the VTE Agent host after the interval set by the **Update Host Frequency** parameter.

  After the update is downloaded, you must manually restart the VTE Agent. Execute one of the following commands on the VTE Agent host to restart the VTE Agent:

  - On Linux, Solaris, and AIX:

    ```
    # /etc/init.d/secfs restart
    ```

  - On Redhat 7.2:

    ```
    # /etc/vormetric/secfs restart
    ```

  - On HP-UX:

    ```
    # /sbin/init.d/secfs restart
    ```

  > **Note**
  >
  > Information about UNIX agents applies to earlier versions of those agents. As of v6.0, UNIX agents such as Solaris and HP-UX are EOL. AIX is still supported. Refer to the DSM Compatibility Matrix for more information about supported agent versions.

- **License Type**

  Choose the type of license that will run on this host. Options are Perpetual, Term, and Hourly.

- **FS Agent Locked**

  Locks the contents of the VTE Agent directories on the host. See "Setting Host Locks" on page 241 for about this setting.

System Locked

  Applies an internal policy to the host to lock host system directories, like `/var`, `/bin`, `/etc`. This can be selected only if **FS Agent Locked** is enabled.

- **Support Challenge & Response**

  Indicates whether this feature is enabled on the host. It becomes enabled when the VTE Agent running on the host registers with the DSM.

- **Password Creation Method**

  If you switch the password method from **Manual** to **Generate**, regenerate the password. Select **Regenerate Password** and click **Apply**. A new generated password is downloaded to the host.

If you switch the password method from **Generate** to **Manual**, enter a new password in the **Password** and **Confirm Password** boxes.

By default, the **Password** and **Confirm Password** text-entry boxes display dots. The dots are just graphic placeholders and do not indicate that a password had been entered. You must enter a password in both text-entry boxes or the **Manual** password method will not be applied to the host.

> **Note**
>
> If you select the generate password creation method for an agent that does not support the challenge-response feature, an ERROR-level audit message is generated and entered in the log after the agent registers with the DSM, plus a red warning message is displayed on the **Edit Host** window for the host. In effect, a randomly generated password is created and downloaded to the host system; however, the `vmsec challenge` command is not available on the host system so a user cannot display a challenge string. The solution is to change the host configuration from **Generate** to **Manual** and manually enter the host password.

- **Regenerate Password**

  Password method to use to unlock the agent. The host user may be prompted to supply a password to decrypt encrypted data when there is no network connection between the host and the DSM. The methods are **Generate** (challenge-response) and **Manual** (static password). When **Generate** is selected, the host user must request a new password from a DSM administrator each time a host password is required.

  Password/Confirm Password—Displayed when **Password Creation Method** is set to **Manual**, enter and re-enter the password to use to unlock a GuardPoint when there is no server connection.

- **Docker Enabled**

  Select this check box to enable docker support.

  > **Note**
  >
  > If you selected the option to enable Docker support during the agent registration procedure, this check box will display as selected, indicating that this feature has been enabled. Refer to the *VTE Agent Installation and Configuration Guide* for more information.

  Once Docker support is enabled, it cannot be disabled. The **Docker Enabled** check box is selectable only if your VTE agent license includes this feature. If your VTE Agent (FS Agent) license includes Live Data Transformation (LDT) and you choose to enable Docker support on a host, then the LDT check box is disabled, as the two features cannot coexist. Similarly if a host has the **Live Data Transformation** check box enabled then the **Docker Enabled** check box is disabled. See "Enabling Docker Support" on the next page for steps to enable this feature.

- **Live Data Transformation**

  Select this check box to enable Live Data Transformation (LDT) on the host.

  > **Note**
  >
  > If you selected the option to enable LDT support during the agent registration procedure, this check box will display as selected, indicating that this feature has already been enabled. Refer to the *VTE Agent Installation and Configuration Guide* for more information.

  Once LDT support is enabled, it cannot be disabled. The **Live Data Transformation** check box is selectable only if your VTE Agent (FS Agent) license includes this feature. If your VTE license includes Docker support and you choose to enable LDT, then the **Docker Enabled** check box is disabled as the two features cannot coexist. Similarly, if a host has the **Docker Enabled** check box enabled then the **Live Data Transformation** check box is disabled. See "Enabling Live Data Transformation" on page 240 for steps to enable this feature.

- **Secure Start GuardPoint**

Select this option if you want to create a Secure Start GuardPoint for Active Directory or MSSQL directories. This feature is only supported on hosts running Windows OS. Refer to the *VTE Agent Installation & Configuration Guide* for more information about using Secure Start GuardPoints.

> **Note**
> The **Docker** and **Live Data Transformation** options are only displayed if you have the relevant license.

3. If your host is a KMIP client then you need to import that host's certificate, see Chapter 25: "Key Management Interoperability Protocol (KMIP)" on page 298, for more about adding KMIP clients.

4. Once your host is registered with the DSM, you can start protecting your data by creating GuardPoints, see Chapter 24: "Managing GuardPoints" on page 275 for more about creating and managing GuardPoints.

5. The **Sharing** tab lets you share the GuardPoints on the host with Security Administrators in other domains, see "Sharing a Host" on page 246.

6. Use the **Host Settings** tab to set authentication options for applications running on the host. See "Host Settings" on page 248 for more information.

7. The **Challenge Response** tab allows a DSM Security Administrator to generate a temporary passphrase to give to a host administrator to decrypt data on the host when there is no connection to the DSM. See "Security Administrator" on page 24 for more information.

8. Use the agent log tabs (FS Agent Log, Key Agent Log, Docker Log) to define log settings. See "Agent Log Settings" on page 252 for more information.

9. You can optionally add the host to a host group using the **Member** tab, see "Configuring Host Groups" on page 262.

Enabling Docker Support

Data protection policies can be set up for Docker images and Docker containers. In addition to data encryption, the DSM also provides Docker container-level access control and container-level audit logging. DSM Security Administrators can create GuardPoints on Docker images and containers via the Management Console.

Docker support is available on the following platforms:

- Docker Host: RHEL 7.0, 7.1, and 7.2

- Docker containers: heterogeneous container support, including but not limited to, RHEL, CentOS, Ubuntu, SUSE

- Docker storage driver: devicemapper

In order to use the Docker support feature on a host, you must have the following:

- VTE Agent 6.0 license with Docker support

- A host with Docker configured and running

- VTE Agent version 6.0 installed on the Docker host

Refer to the *VTE Agent Installation & Configuration Guide* for information about installing and configuring a VTE Agent.

After installing the VTE Agent on the Docker host and registering it with DSM, you must enable Docker support on the DSM:

1. Log on to the Management Console as an All, Domain and Security, or Security administrator.

2. On the main menu bar of the Management Console, click **Hosts**.

3. On the Hosts page, click the name of the Docker host in the **Host Name** column, the Edit Host page opens.

4. In the **Host Information** panel of the Edit Host page, select the **Docker Enabled** check box.

Once you have enabled Docker on a host, you cannot disable it. To disable the feature, you must first unregister and then delete the host and then re-register the host, without enabling the feature. This will let you reclaim the license for use on another host.

Next, edit the **Host Settings**:

1. Log on to the Management Console as an All, Domain and Security, or Security administrator.

2. On the main menu bar of the Management Console, click **Hosts**.

3. On the Hosts page, click the name of the Docker host in the **Host Name** column, the Edit Host page opens.

4. Click the **Host Settings** tab.

- If you are using a Docker engine version earlier than version 1.12.1, add the following entry to the **Host Settings** text box:

  ```
  |authenticator|/usr/bin/docker
  ```

- If you are using a Docker engine version 1.12.1 or later, add the following entry to the **Host Settings** text box

  ```
  |authenticator|/usr/bin/dockerd
  ```

For details about creating Docker GuardPoints, see .

Enabling Live Data Transformation

The Live Data Transformation (LDT) feature enables DSM Security Administrators to encrypt or rekey GuardPoint data without blocking user or application access to that data.

In standard VTE deployments, access to data is blocked during initial encryption or rekeying of data. With Live Data Transformation (LDT), encryption and rekeying of data takes place in the background, without disrupting user or application access.

In order to use LDT you must have the following:

- VTE Agent 6.0 license with LDT.

- VTE Agent version 6.0 installed on a host. Refer to the *VTE Agent Installation & Configuration Guide* for information about installing and configuring a VTE Agent.

  Refer to the *Live Data Transformation Guide* and the *VDS Compatibility Matrix* for information about implementing LDT and the supported platforms.

> **Note**
> The LDT feature uses 'versioned keys', which automatically expire and rotate, as defined by the key's settings. The key rotation and key expiration occur in the background, and it is possible that a DSM backup may not contain the latest versions of the rotated keys. In the event of a DSM failure, all keys that were automatically rotated after the last backup would be lost, making all data encrypted with those keys unusable or unrecoverable. Therefore, we recommend that the LDT feature be used in a high availability deployment.
>
> If LDT must be used in a single DSM configuration, we recommend that you specify 'Cached On Host' for all keys that are created, and to set the password creation method to 'Manual' for all hosts. In the event that the standalone DSM fails and is unavailable, access to the data on the host is still available by entering the known passphrase, and the data is available as the encryption keys are cached on the host.

After installing the VTE Agent on a host and registering it with DSM, you must enable LDT support:

1. Log on to the Management Console as an All, Domain and Security, or Security administrator.

2. On the main menu bar of the Management Console, click **Hosts**.

3. On the Hosts page, click the name of the host on which you want to enable the feature in the **Host Name** column, the Edit Host page displays.

4. In the **Host Information** panel of the Edit Host page, select the **Live Data Transformation Enabled** check box.

Once Live Data Transformation has been enabled, it cannot be disabled. To remove the feature, you must migrate existing data protected under LDT policies, unregister and delete the host and then re-register the host, without enabling the feature. This will let you reclaim the license for use on another host. See "Deleting Hosts" on page 261 for more information.

For details about creating LDT GuardPoints see, Chapter 24: "Managing GuardPoints" on page 275.

For details about how LDT works, guidelines and best practices for using the feature, refer to the *Live Data Transformation Guide*.

Setting Host Locks

FS Agent Locked and **System Locked** are two options used to protect the VTE Agent and certain system files. VTE Agent protection includes preventing some changes to the VTE Agent installation directory and preventing the unauthorized termination of VTE Agent processes. These options appear in **General** tab of the Edit Host and Edit Host Group windows and are disabled by default.

> **Note**
> You might not be able to upgrade or delete agent software if you do not disable the locks first.
> - Disable **FS Agent Locked** before updating or deleting agent software on the host system.
> - Disable **FS Agent Locked** before deleting the host record from the Management Console.
> - Disable **System Locked** before updating, deleting, or modifying protected system files.

To apply locks:

1. Check that no one is currently in or accessing the Agent installation directories; otherwise, the DSM may be unable to lock the Agent software.

2. Log on an administrator with Host role permissions.

3. To set the locks on an individual host:

   a. Select **Hosts > Hosts** in the menu bar.

      The Hosts window opens.

   b. Click a host name in the **Host Name** column.

      The Edit Host window opens to the **General** tab.

   c. To protect VTE Agent files from modification and deletion, enable the **FS Agent Locked** check box.

   d. To protect a set of system files from modification and deletion, enable the **Host > System Locked** check box.

      System Locked is automatically enabled when **FS Agent Locked** is enabled. You can enable and disable **System Locked** only when **FS Agent Locked** is enabled.

   e. Select **Ok** to finalize the changes.

To set locks on hosts in a host group:

1. Select **Hosts > Host Groups** in the menu bar.

   The Host Groups window opens.

2. Click a host group in the **Name** column.

The **Edit Host Group** window opens to the **General** tab.

3. To protect VTE Agent files from modification and deletion, enable the **FS Agent Locked** check box.

4. To protect a set of system files from modification and deletion, select **Host > System Locked**.

    System Locked is automatically enabled when **FS Agent Locked** is enabled. You can enable and disable **System Locked** only when **FS Agent Locked** is enabled.

5. Select **Ok** to finalize the changes.

6. (Optional) As a host administrator with root permissions, verify that the locks have been applied to the agent.

    a. Log onto the host (agent) system.

    b. Execute the `secfsd` command with the `lockstat` argument:

        ```
        # secfsd -status lockstat
          FS Agent Lock: true
          System Lock: true
        #
        ```

> **Note**
>
> Sometimes there is a discrepancy between what the DSM reports as the VTE Agent configuration and the actual VTE Agent configuration. This may be due to the time delay between log uploads to the DSM, or because a GuardPoint is in use when the lock is applied.

If the locks are enabled and the DSM cannot administer the host, such as can occur after changing authentication credentials or removing the certificate fingerprint, the host administrator must unlock the host manually. The certificate fingerprint can be removed if the **Registration Allowed** check box on the **General** tab of the Edit Host page, is not selected

To unlock the host manually, boot the host into single-user mode and edit the `./secfs/.sec/conf/configuration/secfs_config` file.
Set both `coreguard_locked` and `system_locked` to `false`. Save the file. Boot the system into multi-user mode. You should now be able to administer the host again.
On Windows systems, boot in safe mode, rename `C:\Windows\system32\drivers\vmmgmt.sys and .\drivers\vmfiltr.sys` to something else, then boot in regular mode.

The host administrator must inform the Security Administrator of changes to the system hierarchy.

- **Example 1:** The host system administrator can request to have the locks temporarily disabled to do some administrative functions.

- **Example 2:** The host system administrator can remove directories and files, then, later when the lock is reapplied, the DSM is protecting non-existent data.
  Another common administrative issue pertains to mounted GuardPoints. The host system administrator can remove or unmount an unlocked, non-automounted GuardPoint. The DSM Management Console interface is not aware of this change and does not issue a warning when you reapply the lock to the now non-existent mounted GuardPoint.

- **To recover an unmounted GuardPoint:**

    a. Disable the GuardPoint for the file system in the Management Console.

    b. Mount the file system on the host.

    c. Enable the GuardPoint for the file system.

FS Agent locked

FS Agent Locked locks the contents of the VTE Agent directories on the host. These directories are `/<install root>/agent/secfs` and `/<install root>/agent/vmd`.

Files in these directories cannot be modified or removed when **FS Agent Locked** is enabled; however, the DSM can still propagate updates to the host system.

When **FS Agent Locked** is enabled:

- **System Locked** is automatically enabled

- Certificates are exchanged and the host is bound to the DSM

- The VTE Agent installation directory cannot be deleted or overwritten

- The VTE Agent services cannot be stopped

- The VTE Agent GuardPoints cannot be forcefully unmounted

When **FS Agent Locked** is disabled:

- **System Locked** is automatically disabled

- The VTE Agent software on the host is not protected

> **Note**
> Do not unregister or delete the VTE Agent while locks are applied. The locks stay in effect after the agent is unregistered and, without agent credentials, the DSM cannot administer that Agent and it cannot disable the locks. You must boot the host into single-user mode and manually modify the agent configuration to disable the locks.

On Linux systems, all operations are permitted in the following directory when **FS Agent Locked** is enabled.

```
/<install root>/agent/secfs/tmp
```

On Linux systems, the following directories cannot be removed or renamed, and directory and file creation will fail when **FS Agent Locked** is enabled.

```
/<install root>/agent/secfs/bin
/<install root>/agent/vmd
```

On Linux systems, file creations and other operations will work for the following directory, but the directory cannot be removed or renamed when **FS Agent Locked** is enabled.

```
/<install root>/agent/secfs/
```

On AIX systems, the contents of the following directories cannot be changed or moved when **FS Agent Locked** is enabled.

```
/<install root>/agent/vmd
```

On AIX systems, the contents of the following files and directories can be modified, but not removed or renamed when **FS Agent Locked** is enabled.

```
/<install root>/agent/secfs/
/<install root>/agent/secfs/tmp
```

On Windows systems, when **FS Agent Locked** is enabled, the following folder cannot be moved and its contents cannot be modified:

```
C:\Program Files\Vormetric\DataSecurityExpert\Agent\secfs\sec
```

Also, the VTE Agent entries in the registry cannot be modified or deleted when **FS Agent Locked** is enabled on a Windows system.

System locked

System Locked applies an internal policy to the host to lock host system directories, such as `/var`, `/bin`, `/etc`, and so on. When you enable **FS Agent Locked**, **System Locked** is automatically enabled.

> **Note**
>
> To upgrade or install third-party software, add new applications, open an SSH session remotely, or modify system directories, you must disable **System Locked**.

> **Note**
>
> (Windows only) Verify that the volume letter and the path for the Windows system are correct before proceeding. When Windows VTE Agent software is installed, the volume letter defaults to "`C:`" The executables in the **Host Settings** tab may be on a different volume or in a different folder. If the volume or path information is incorrect, the DSM cannot sign the applications, and it will be unable to apply **FS Agent Locked** and **System Locked**.

When **System Locked** is enabled:

- Operating system directories on the host are protected.

- Microsoft Update cannot be run on Windows systems to protect the host. Microsoft update and other installation-related executables are specifically blocked. Executables like `wuacuclt.exe` and `msiexec.exe` cannot be run.

- The installation utility checks if **System Locked** is enabled on the host system. If it is, the utility aborts installation and displays a message telling you to "`unlock system before running install/update program`". Other third-party installation utilities do not check if **System Locked** is enabled and are not prevented from installing software.

- New file or directory creation inside a protected directory is not allowed.

When **System Locked** is disabled:

- The internal policy is disabled.

- **FS Agent Locked** remains enabled.

- You can install or update system software.

The following files, directories, and subdirectories are, by default, automatically protected when **System Locked** is enabled. NB: Asterisks (*) indicate pattern matching.

On Linux systems, the following files and the contents of the following directories cannot be changed or moved when **System Locked** is enabled.

- `/etc/pam.d`
- `/etc/rc*`
- `/etc/security`
- `/usr/lib/security`

On Linux systems, the contents of the following files and directories can be modified, but not removed or renamed when **System Locked** is enabled.

- `/etc`
- `/etc/init.d/secfs`
- `/usr`
- `/usr/bin/vmd`
- `/usr/bin/vmsec`
- `/usr/bin/secfsd`
- `/usr/bin/dataxform`

- `/usr/lib`
- `/usr/lib/*pam*`
- `/usr/lib/security`
- `/var/log/vormetric`

On Solaris systems, the following files and directories cannot be created, edited, or deleted, when **System Locked** is applied.

- `/usr/lib/fs*`
- `/usr/ker*`
- `/usr/pla*`
- `/usr/lib/securi*`
- `/etc/rc*`
- `/etc/ns*`
- `/etc/vfs*`
- `/etc/init.d/secfs`
- `/etc/system`
- `/ker*`
- `/pl*`
- `/sbin`

On AIX systems, the following files and the contents of the following directories cannot be changed or moved when **System Locked** is enabled.

- `/etc/rc.d`
- `/etc/security`
- `/usr/lib/security`
- `/sbin/helpers/mount_secfs`

On AIX systems, the contents of the following files and directories can be modified, but not removed or renamed when **System Locked** is enabled.

- `/var/log/vormetric`

On HP-UX systems, the following files and the contents of the following directories cannot be changed or moved when **System Locked** is enabled:

- `/sbin/rc[0-4].d`
- `/sbin/init.d`
- `/usr/lib/security`
- `/etc/pam.conf`
- `/etc`
- `/usr`
- `/sbin`
- `/sbin/rc`
- `/etc/inittab`
- `/usr/lib`

On HP-UX systems, the contents of the following files and directories should not be modified when **System Locked** is enabled:

- `/sbin/fs/secfs2`
- `/usr/bin/secfs`
- `/usr/bin/vmd`
- `/usr/bin/vmsec`
- `/usr/bin/secfsd`
- `/usr/bin/dataxform`

When **System Locked** is applied, a protected file or path cannot be renamed or deleted; however, if it is a directory, other files may be added to it. For example, `/etc` cannot be deleted nor renamed, though you can add files to it. A file that cannot be modified cannot be opened and edited in any way.

On Windows systems, files with the following extensions in the Windows OS installation folder (for instance: \Windows, \WinNT, and so on) cannot be moved or modified when **System Locked** is enabled:

- `.exe`
- `.dll`
- `.sys`
- `.cmd`
- `.com`

Setting locks on Docker hosts

The **FS Agent Locked** and **System Locked** options are applicable to Docker host systems but, they are not applicable to Docker images and containers. Files and directories that are locked on the Docker host using these options remain locked, even if they are indirectly accessed through a Docker image or container.

Sharing a Host

Security Administrators in other domains may administer GuardPoints on a locally configured host if sharing is enabled. This feature is used to allow Security Administrators in other domains to manage a host or host group in that domain. The domains that are allowed to administer the local host are set and displayed in the **Sharing** tab of the Hosts window.

The shared/not shared status of a host is indicated on several Management Console windows:

- Hosts and Host Groups windows

- Edit Hosts window, **Sharing** tab

- Edit Hosts window, **GuardPoints** tab—indicated by an obscured **Select** check box and italicized host policy name

- A host that is not configured for sharing displays a blank in the **Sharing** status field in the Hosts window.

- A shared host that is being accessed in the same domain in which it was created has a fully functional interface and displays Sharing in the **Sharing** status field. All the tabs in the Edit Host window are displayed and can be used to configure the host.

- A shared host that is being accessed by a Security Administrator in a different domain than the domain in which it was created has a partially functional interface and displays a sharing status of External. Only the **General** and **GuardPoints** tabs in the Edit Host window are displayed and they are used to add and remove GuardPoints.

The Edit Host window in the Management Console normally displays tabs that are used to configure VTE Agents, agent logs, and set other host parameters. The Edit Host window for a shared host displays only the **General** and **GuardPoints** tabs

The current Security Administrator domain will be displayed in the top-right corner of the Management Console window. The domain of a shared VTE Agent is displayed in the **GuardPoints** tab of the Hosts and Edit Hosts windows. Three GuardPoints were applied in datadomain1 and the other in datadomain2. The **Select** check boxes for GuardPoints in datadomain1 are disabled, but enabled for datadomain2, indicating that the Security Administrator is logged into datadomain2. Therefore, the current Security Administrator can delete the GuardPoint made in datadomain2 but not the one made in datadomain1. The current Security Administrator can also add additional GuardPoints.

Configuration attributes are local to the domain in which the Security Administrator is currently working. Primarily, the keys and policies that are in the local domain are used to configure GuardPoints.

The shared host is indicated in the **Sharing** column of the Hosts window. A state of External indicates that the host you are accessing is a remote, shared host, and only a limited set of VTE Agent features are available for configuring it. A state of Shared indicates that the local host is being shared in one or more other domains.

A grayed-out **Select** checkbox and an italicized **Policy** name in the **GuardPoints** tab indicates a GuardPoint that is configured on the same system but in another domain. You cannot determine specifically which domain, other than by switching to each domain and checking configured hosts.

Sharing a Local Host with Another Domain

1. Log on an administrator with Host role permissions.

2. Change to the desired domain, if you are not already in it.

   a. Select **Domains > Switch Domains**.

      The **Domains** window opens. All the domains in which the current Domain Administrator is a member are displayed. The current domain is not selectable.

   b. Enable the radio button of the desired domain.

      If the desired domain is not listed, ask the Domain Administrator for that domain to add you to it.

   c. Click **Switch to domain**.

      The Domains window is redisplayed.

3. Select **Hosts > Hosts** or **Hosts > Host Groups**.

4. Select the host or host group to be shared from the **Host Name** column of the Hosts window or the **Name** column of the Host Groups window.

   The Edit Host or Edit Host Group window opens.

5. Select the **Sharing** tab.

6. Click **Share**.

   The **Sharing** window opens.

7. Enter the name of the domain to be given shared access to the current host in the **Domain Name** text-entry box.

   All configured domains are available, even domains the current Security Administrator is not configured to access. Available domain names are not displayed and a domain browser is not provided. Domain name handling is case-sensitive. Enter the name exactly as it is configured. The Management Console will tell you if you enter an incorrect or non-existent domain name.

8. Click **Ok**.

Shared Host Logging

Shared hosts and shared host groups are administered in the domain in which they were created. All of the VTE Agent log data generated on a shared host is displayed only in the domain in which the host was created. You must be in the domain in which the host was created to view GuardPoint activity in the logs.

Only server-generated messages are displayed in the log of the domain that is being granted shared access. That is, only log messages that indicate that the DSM performed an action are displayed. Host acknowledgment is not displayed. You must enter the domain in which the host or host group was originally created to view host acknowledgment and GuardPoint access activity. If the VTE Agent is assigned to an HA node DSM, agent activity is logged on the HA cluster.

Host Settings

The Host Settings tab allows you to set authentication options for the applications running on the host. Applications such as `su`, `sshd`, and `login` that authenticate a user's identity by requesting a user name and an associated password, are signed applications that identify and authenticate before a child process executes.

GuardPoints may have an associated policy that restricts access to the data contained in those GuardPoints. For a process to be able to access the data, the user's associated identity must be authorized. This authorization can be done by adding an entry in the host settings table that specifies a program, such as mentioned above, along with a keyword that indicates the type of authorization that is applied.

Host Settings on the DSM are pushed to the hosts periodically. In an HA deployment, you can also click **Notify All Hosts** in the High Availability Servers window to push the latest host configurations directly from an HA Node to every host in the HA cluster.

> **Note**
> Do not click **Notify All Hosts** more than once. Each time you click this button you spawn a new process and each new process slows the DSM.

Applications in the Host Settings tab used to be automatically signed when new settings were pushed from the DSM. Therefore you could apply host settings after any of the following tasks:

- Installing VTE Agent software

- Installing VTE Agent software with Docker enabled

- Upgrading VTE Agent software

- Changing any of the files listed in the Host Settings tab

As of DSM v6.0.1 and VTE v6.0.1, any changes to the Host Settings are no longer automatically signed. The signatures of the newly added process or processes are compared against the signatures of the existing settings and, if they differ, an error message is generated. See section "Re-Sign Settings" on page 252 for how to configure this setting, and refer to the v6.0.1 *VTE Agent Installation and Configuration Guide* for details about this feature.

For specific information about HDFS hosts settings information, refer to the *VTE Installation and Configuration Guide*.

As of DSM v6.0.2, Host Settings can also be configured at the host group level, see "Host Group Host Settings" on page 270 for details.

Host settings for Linux and AIX

The text entry box on the Host Settings tab is where you specify what authentication mechanisms are in place for certain binaries on the host machine. Each line has the format:

```
|behavior|/path/to/binary
```

Default settings for Linux

```
|authenticator|/usr/sbin/sshd
|authenticator|/usr/sbin/in.rlogind
|authenticator|/bin/login
|authenticator|/usr/bin/gdm-binary
|authenticator|/usr/bin/kdm
|authenticator_euid|/usr/sbin/vsftpd
|protect|/etc/passwd
|protect|/etc/group
```

Default settings for AIX

```
|authenticator|/usr/sbin/tsm

|authenticator|/usr/sbin/sshd

|authenticator_euid|/usr/sbin/ftpd

|authenticator|/usr/dt/bin/dtlogin
```

Host settings for Windows

For applications running under Wow64 that require some form of user authentication, create entries in the Host Settings tab for Windows. The syswow64 paths are created by default during Windows file agent installation. `\Windows` is for Windows XP and Windows Itanium operating systems.

In Wow64, all file-access to `C:\Windows\System32` is redirected to `C:\Windows\syswow64`, and is implemented using the File System. Redirected syswow64 paths are effective only for 64-bit Windows file agents. This is the path where programs compiled for 32-bits are stored in order to run on a 64-bit system.

Verify that the volume letter and the path for the Windows system are correct before proceeding. When Windows VTE Agent software is installed, the volume letter defaults to "`C:`" It is possible that the executables in the Host Settings tab are on a different volume or in a different folder. If the volume or path information is incorrect, the DSM cannot sign the applications, and it cannot apply **FS Agent Locked** and **System Locked**.

Default settings for Windows

```
C:\WINDOWS\system32\winlogon.exe

|lock|C:\WINDOWS\system32\msiexec.exe

|lock|C:\WINDOWS\system32\wuauclt.exe

|lock|C:\WINDOWS\system32\wupdmgr.exe

|lock|C:\Program Files\Vormetric\DataSecurityExpert\agent\secfs\sec\bin\vminstall.exe

|exempt|C:\WINDOWS\explorer.exe

|exempt|C:\WINDOWS\regedit.exe

|exempt|C:\WINDOWS\system32\regedt32.exe

|exempt|C:\WINDOWS\system32\svchost.exe

|exempt|C:\WINDOWS\system32\services.exe

|exempt|C:\WINDOWS\system32\smss.exe
```

Host settings for a Docker enabled host

1. Log on to the Management Console as an All/Domain and Security administrator.

2. On the main menu of the Management Console, click **Hosts**.

3. Click the host in the **Host Name** column, the Edit Host page opens

4. Click the **Host Settings** tab.

   ○ If you are using a Docker engine earlier than version 1.12.1, add the following entry to the **Host Settings** text box:

   ```
   |authenticator|/usr/bin/docker
   ```

   ○ If you are using a Docker engine version 1.12.1 or later, add the following entry to the **Host Settings** text box:

   ```
   |authenticator|/usr/bin/dockerd
   ```

You can also define host settings for docker containers. It allows all tags for example, `authenticator`, `su`, `protect` etc for containers as well. If you want to tag specific containers, you need to add them as follows:

```
|<tag name><+arg=<+cid=<container ID>>>| path_to_binary
```

For example if you want to add `sshd` authenticator for a Docker container:

1. Log on to your DSM.

2. On the main menu of the Management Console, click **Hosts**.

3. Click the host in the **Host Name** column, the Edit Host page opens

4. Click the **Host Settings** tab and add the following entry:

   ```
   |authenticator+arg=+cid=b4c6a9ca8ce4|/usr/sbin/sshd
   ```

   where `cid` is the 12 character container ID.

Oracle database in a guarded NFS mount on AIX

If you plan to locate your Oracle database in a guarded NFS mount, add the following entries to host settings.

```
|vfsnumber|<path to>/oracle
|vfsnumber|<path to>/dbca
```

Example:

```
|vfsnumber|/u01/app/oracle/dbhome_1/bin/oracle
|vfsnumber|/u01/app/oracle/dbhome_1/bin/dbca
```

Host setting keywords

The following table lists the keywords that you can enter in the Host Settings tab that override different authentication requirements:

Table 23-2: Host Settings tab keywords

Keyword	Description
|authenticator|	(UNIX only) This keyword means that the given binary is trusted to authenticate users. For example, the `sshd` process on UNIX is a good |authenticator| because it takes incoming network connections and authenticates the user that is attempting to log in to the system. All child processes from this session will be trusted as the original user.
|authenticator_euid|	(UNIX only) The |authenticator| keyword authenticates based upon the real user ID (ruid) credentials of a process. The |authenticator_euid| keyword authenticates based upon the effective user ID (euid) credentials of a process. The |authenticator_euid| keyword is used when you want to authenticate the credentials of a setuid process with the euid value rather than the ruid value.
|vfsnumber|	(AIX [all supported]/Oracle 10gR2) Use this host setting in the case that Oracle RMAN backups fail on NFS as a result of not receiving underlying file system identifiers. Apply |vfsnumber| to the Oracle binaries directory.
|realfsid|	(AIX[All supported], HPUX [All supported]) On AIX, use this host setting if the cp operation fails while copying files with extent attributes on guarded Veritas file systems. The failure is due to the underlying file system identifier not being received. The same host setting should also be used on HPUX environments when using the Veritas `vxresize` utility.

Table 23-2: Host Settings tab keywords (continued)

Keyword	Description
|lock|	(Windows only) Specifies an application that cannot be executed on the host. An application defined with `lock` does not go through an internal policy check. It is not allowed to run at all. A default set of applications is locked on the Windows host to prevent their execution and causing potential failure during bootup. The same effect can be achieved by configuring the **Resource** and **Process** security rule attributes in a policy; however, certain default applications are automatically locked in the Host Settings tab as a precautionary measure for when you fail to include these applications in the policy. Sometimes problems occur when installing software on a locked host, such as installation failure or application lockup. Specific processes can be identified where, when they are locked, they cannot be started and the failure goes away. For example: `|lock|c:\winnt\system32\msiexec.exe`
|exempt|	(Windows only) When processes or applications are started, the internal policy and regular policies are checked locally or by the Security Server. When a policy check is performed and exempt is applied to the process, a 6 second timeout is imposed on the check. Without exempt, an application can wait indefinitely for a policy access check to complete, as when the Security Server is required but is not accessible. If the check times-out because the Security Server is unavailable for any reason, access is denied. Exempt host processes are also "exempt" from pop-up messages that describe the occurrence of access violations. An example of what causes such pop-ups is an application that tries to memory map a file for which it does not have encryption permission (for instance, memory map with no view ability key on Windows). The only reasons to include exempt in the configuration are shorter wait periods and blocked pop-ups.

> **Note**
> |trust| and |trustfrom| have been deprecated. Please re-evaluate host settings and replace with |authenticator| or |authenticator_euid| as appropriate. These settings will continue to be supported.

The different results you get when using `authenticator` or `authenticator_euid` to verify user identities is shown in the following table.

Table 23-3: Results from authenticator to verify user identity

Product	Application	Host Setting	User
Oracle	oracle	authenticator_euid	"oracle"
Oracle	oracle	authenticator	*

    * Indicates the real uid of the user who starts the application. This means that if the policy is configured to check user ID, a security rule must be generated for every possible user.

> **Note**
> Apply the `|authenticator_euid|` keyword to the `oracle` binary in the Host Settings tab to authenticate the `oracle` user because regardless of who starts the `oracle` process, the EUID is always `oracle`.

Configuring Application Authentication Credentials

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Select **Hosts > Hosts** in the menu bar. The Hosts window opens.

3. Click the host in the **Host Name** column. The Edit Host window opens.

4. Select the Host Settings tab. This tab displays a default set of system applications that may require authentication entries.

5. Add, modify, or delete entries to control their access permissions. When you add more processes, you must include the entire path.

> **Note**
> You must use a keyword such as `|authenticator|` in front of a process or it will be ignored by the Management Console.

6. Click **Ok**.

7. Any users who are currently logged on to the system must log off, and then log on again to refresh their user authentication credentials.

8. Verify the change by logging on to the host and accessing a GuardPoint, then check the user information in the Message Log.

Re-Sign Settings

If you add another process to the set of trusted applications on the Host Settings tab, check the **Re-Sign Settings** check box to ensure that the new process is signed and authenticated by the host.

The next time host settings are pushed to the VTE Agent, the updated host settings are re-signed and the **Re-Sign Settings** check box on the Management Console is cleared (or reset).

To ensure that the new process is signed and authenticated by the host, do the following:

1. Navigate to the **Hosts > Hosts** option on the Management Console menu.

2. Che **Re-Sign Settings** option. Selecting this option will force a signature update. The next time host settings are pushed to the VTE Agent, the updated host settings are re-signed and the **Re-Sign Settings** check box on the DSM Console is cleared (or reset).

If you do not select this option after adding a new process, the host will ignore the newly added process.

Agent Log Settings

Configure log viewing settings for the various agents from the specific tabs—VTE (FS) Agent on the FS Agent Log tab, Key Agent (VAE/VKM) settings from the Key Agent tab, or Docker settings from the Docker Log tab. If you are outside of a domain, i.e., at the 'system' level, then these configuration settings are applied globally. All host systems added after this change inherit the log settings attributes, but all current hosts configurations remain intact. To configure log settings attributes for a specific host, you need to log into a domain and make the changes on that host. The host level settings take precedence over the system level settings.

FS Agent Log

This section describes VTE (FS) agent log configuration.

The table at the top of the page displays the Message Type and log message destination i.e., where the log files will be stored.

Message Type

- **Management Service**

  Logs messages that are related to the agent and VMD process server interaction in the agent logs. Log to File and Upload to Server are enabled by default. The default log message level is INFO.

- **Policy Evaluation**

  Logs messages that are related to policy evaluation in the agent log. Set the log message level to desired setting. The default log message level is ERROR.

- **System Administration**

  Logs messages that are related to system level events. The default log message level is ERROR.

- **Security Administration**

  Logs messages that are related to security related events. The default log message level is INFO.

The detail and extent of information to be logged by the current agent is determined by the selected error level. The agent supports five log levels. These logs can be logged to a local file, a Syslog server or uploaded to the DSM.

In sequence they are:

- **DEBUG**

  Designates fine-grained informational events that are targeted towards support engineers and developers.

- **INFO**

  Designates informational messages that highlight the progress of the application at coarse-grained level.

- **WARN**

  Designates potentially harmful situations.

- **ERROR**

  Designates error events that might still allow the application to continue running.

- **FATAL**

  The 'FATAL' level designates very severe error events that will presumably lead the application to abort.

Log levels are cumulative. The level that you select not only generates log entries for events that occur at that level, but all the levels below. For example, the 'WARN' level also includes events that occur on the 'ERROR' and 'FATAL' levels.

Message Destination

Log Messages can be stored in several locations.

- **Log to File**

  Send log messages to the /var/log/vormetric/vorvmd_root.log file of a UNIX host, or a Windows equivalent, such as \Documents and Settings\All Users or WINDOWS\Application\ Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log.

- **Log to Syslog**

  Send log messages to the syslog server for a UNIX host. If a syslog server is not configured, it is sent to the host 'messages' file, such as /var/adm/messages. On a Windows host, the messages are sent to the Event Viewer (Application events).

- **Upload to Server**

  Upload to the DSM and display in the Management Console Logs window.

- **Level**

  Sets the level of error messages to be sent.

- **Duplicates**:

  - **Allow**

    All duplicate messages of the corresponding Message Type are captured and displayed in the log.

  - **Suppress**

Messages of the corresponding Message Type will follow the configured Threshold as to how many times duplicate messages are sent to the DSM during the given Interval.

File Logging Settings

- **Maximum File Size (bytes)**

  The agent starts a new, empty log file when the specified limit is exceeded. The default is 1000000 bytes.

- **Delete Old Log Files**

  Select this check box to delete old FS agent logs. This check box works in conjunction with the Number of Old Log Files to Keep text-entry box. For example, Select this check box and enter 3 as the Number of Old Log Files to Keep value. After 3 logs are generated, the first log, log1, is deleted and a new log, log4, is created. If you do not Select this check box, log files will continue to accumulate in the server database and you will have to remove them manually.

Syslog Settings

- **Local**

  Send Syslog messages to the local machine.

- **Server (1, 2, 3, 4)**

  Enter the hostname of the Syslog server.

- **Protocol**

  Select the protocol to connect to the syslog server; UDP or TCP

- **Message Format**

  Specifies the format of the message; Plain Message, CEF, or RFC5424.

Upload Logging Settings

- **Maximum Number of Messages to Upload At Once**

  Limits the number of messages sent to the DSM at one time. When the specified number of log entries is reached, those entries are uploaded to the DSM. The default is 1000.

- **Upload Messages At Least Every (seconds)**

  The maximum interval to wait before the agent is to upload messages to the DSM. Use this attribute to update the log viewer even when the Maximum Number of Messages to Upload At Once has not been reached. You can lower the interval if there is little agent activity. The default is 10 seconds.

- **Upload Messages At Most Every (seconds)**

  The minimum interval to wait before the agent is to upload messages to the DSM. You can increase the interval if there is considerable agent activity, so the agents do not flood the network with log messages. The default is 1.

- **Normal Time Out (seconds)**

  The maximum interval of time the agent is to wait for the DSM to acknowledge a backup or restore request and upload related message data. If the agent cannot connect to the DSM within the specified interval, the agent will try again after the interval configured by the Upload Messages At Least Every attribute. The default is 2 seconds.

- **Shutdown Time Out (seconds)**

  The maximum interval of time the agent is to wait for the DSM to acknowledge job completion and upload related message data. If the agent is unable to upload the log messages within the specified interval, they are left on the agent system. The agent will resend the messages at the beginning of the next job. The default is 30 seconds.

- **Drop If Busy**

  Select to slow log message generation and drop log files during periods of extreme logging.

Duplicate Message Suppression Settings

- **Enable Concise Logging**

  When enabled, audit log messages are reduced. This option is disabled by default. Instead of logging messages for each file system operation, only the following types of audit messages are logged:

  - Only one audit message for each read or write activity is logged at the start of that activity.

  - Audit messages for reading file status information and setting file attributes (and extended attributes) are not logged.

  - Audit messages for directory open, close and read attributes are not logged.

- **Threshold (1-100)**

  Used when the Duplicates value is set to Suppress. Specifies the maximum number of duplicate messages the agent is to send to the DSM within the amount of time specified by the Interval parameter. The default is 5 messages and the maximum is 100 messages.

- **Interval (seconds) 1-1000**

  Used when the Duplicates value is set to Suppress. Specifies the time period in which the number of duplicate messages, specified by Threshold, can be uploaded to the DSM. Once Interval is exceeded, the count specified by the Threshold parameter starts again. The default is 600 seconds (10 minutes). The maximum is 3600.

- **Maximum Space for Caching Log Files (MB)**

  This setting indicates the space available for caching agent log files. Log files are copied from the agent to the DSM soon aster they are created, assuming a good network connection. If the network is a little slow, a backlog will build up and the log files are cached. If the space for caching files fills up, the system slows down and new log messages are dropped. The agent sends warning messages to that effect, which can be viewed on the Logs page on the Management Console.

- **Maximum Number of Cached Log Files**

  This setting indicates the number of files that can be stored in the space for caching log files pending upload to the DSM. If the limit is reached, the agent will drop any new log messages and send warning to the DSM which can be viewed on the Logs page of the Management Console.

> **Note**
> The default values for **Maximum Space for Caching Log Files (MB)** and **Maximum Number of Cached Log Files** are the recommended values. If these values are changed, they should be kept in the same ratio, since each log file can be about 500Kbytes in size. Additionally, users should ensure that the underlying file system can accommodate a larger backlog of files.

Key Agent Log

Configure log viewing settings for the Key Agent on the Key Agent Log tab. If you are outside of a domain i.e., at the 'system' level, these configuration settings are applied globally. All Key Agent host systems added after this change inherit the log settings attributes, but all current hosts configurations remain intact. To configure log settings attributes for a specific host, you need to log into a domain and make the changes on that host. The host level settings take precedence over the system level settings. This section describes global Key agent log configuration.

The table at the top of the page displays the Message Type and log message destination i.e., where the log files will be stored.

Message Type

Key Operation: Enters messages that are related to the key operation. Log to File and Upload to Server are enabled by default. The default log message level is INFO.

Message Destination

- **Log to File**

  Send log messages to the /var/log/vormetric/vorvmd_root.log file of a UNIX host, or a Windows equivalent, such as \Documents and Settings\All Users.WINDOWS\Application\ Data\Vormetric\DataSecurityExpert\agent\log\vorvmd.log.

- **Log to Syslog**

  Send log messages to the syslog server for a UNIX host. If a syslog server is not configured, it is sent to the host 'messages' file, such as /var/adm/messages. On a Windows host, the messages are sent to the Event Viewer (Application events).

- **Upload to Server**

  Upload to the DSM and display in the Management Console Logs window.

- **Level**

  Sets the level of error messages to be sent.

- **Duplicates**:

  - **Allow**

    All duplicate messages of the corresponding Message Type are captured and displayed in the log.

  - **Suppress**

    Messages of the corresponding Message Type will follow the configured Threshold as to how many times duplicate messages are sent to the DSM during the given Interval.

  The rest of the settings; File Log Settings, Syslog Settings, Upload Log Settings, and Duplicate Message Suppression Settings are the same as for the FS(VTE) Agent, see "FS Agent Log" on page 252 for details.

Docker Log

The Docker Log tab lets you configure log settings for a docker image or container. The docker logs record events related to the policy applied to the selected images or containers. If no log settings are defined on this tab, the settings defined on the FS Agent Log tab will apply. However if policy evaluation log settings are defined on the Docker Log tab, they take precedence over any policy evaluation settings defined on the FS Agent Log tab.

To configure Docker Logs:

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Navigate to the Hosts page.

3. Click the name of your Docker host in the **Host Name** column, the Edit Host page opens.

   Enter the following information in the **Configure Docker Log Setting** panel:

   - **Docker Image/Container**

     Click **Browse** to select an image or container from the Docker host. If you select an image the **Docker Image ID** field displays the image ID. If you select a container, the **Docker Image ID** field displays the image from which the container was spawned and the **Docker Container ID** displays the container ID. You can use these IDs to search for Docker specific logs on the Logs page later.

   - **Policy Evaluation Level**

     Select a log message level

   - **Policy Evaluation Duplicated**

     You can choose to suppress or allow duplicate messages. Select SUPPRESS or ALLOW, the default is SUPPRESS.

4. Click **Ok**. The Policy Evaluation settings are saved in a tabular format under the **Configure Docker Log Setting** panel.

Docker log messages are displayed on the Logs page.

To search for Docker specific log messages:

1. Navigate to the Logs page.

2. Enter the following information in the **Search** panel:

   ◦ **Log Type**

   Select whether you want to display logs from both the DSM and the agents, only the DSM, or only the agents. The default is All, which means from both DSM and agents.

   ◦ **Source**

   Enter the hostname of the DSM or agent for which you want to return log files.

   ◦ **Last Refreshed**

   Displays the date and time of when the displayed log files were last refreshed. Format is YYYY-MM-DD HH:MM:SS

   ◦ **Message Contains**

   Type in text string that you want to search for in the log messages.

   ◦ **Docker Host**

   Click **Browse** to select the Docker Host for which you want to return log files.

   ◦ **Docker Image/Container**

   Click **Browse** to select an image or container for which you want to display logs.

   ◦ **Docker Image ID**

   Displays the ID for the selected Docker image.

   ◦ **Docker Container ID**

   Displays the ID of the selected Docker container.

3. Click **Go**. The relevant logs are displayed in the table under the **Search** panel.

Automatic Renewal of Host Certificates

Certificates are used to verify the identity of a remote peer when agents communicate with the DSM. The current lifespan of these certificates is 365 days. For the automatic agent certificate renewal process to work, you must have the following:

- Current (not expired) and valid host certificates installed

- Access to the Management Console as a Security, Domain and Security, or All administrator.

The system prompts the administrator and automatically renews any certificate that is 60 days or closer to expiration. The renewal process is transparent and requires no intervention by the administrator. If multiple host agents require renewal at the same time, the server staggers the renewal process to avoid network congestion. This staggering could introduce a delay of up to 48 hours in the renewal process.

Certificate renewal notification

The DSM automatically renews certificates for the VTE (FS) agent.

Certificate renewal may cause the agent to restart. When an agent restarts or certificate is renewed, the agent sends a system notification and log entry.

VMD restart sends the following notification for all installed products:

- `Certificates for the <agentname> agent expire in <number> days`

Certificate renewal causes the agent to report the following message on restart:

- `The new certificate set has been activated`

For information about the Key Agent, refer to the relevant Key Agent documentation.

Updating host certificates

You must regenerate host certificates when you:

- Configure an agent to access a new initial DSM

- Update agent certificates as part of a scheduled update process

- Delete and reinstall agent software

- Regenerate the CA signer certificate of the DSM

> **Note**
>
> The default host registration timeout is 10 minutes. If the host is unable to reach the DSM within the allotted period because of an extremely slow network connection, set the `REGISTER_HOST_TIMEOUT` environment variable to extend the registration timeout. The variable value is an integer expressed in seconds. You may also have to extend the default TCP timeout. See also "RFC 5482 - TCP User Timeout Option".

Since you are updating host certificates, the host already has certificates and the host is already registered with the initial DSM. The certificates on the local host will be deleted and regenerated automatically. However, you must unregister the host on the DSM before proceeding. This is described below.

If you are upgrading agent certificates with the same DSM, there is no need to disable GuardPoints.

If you are upgrading the agent certificates with a different DSM, disable all configured GuardPoints for the host before proceeding. After certificate upgrade completes, assign the GuardPoints from the new DSM.

To update host certificates:

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Click **Hosts > Hosts**. The **Hosts** window opens.

3. Click the host in the **Host Name** column. The **Edit Host** window opens to the **General** tab.

4. Disable the **Registration Allowed** check box for the agent whose certificate you want to change.

   A dialog box opens warning you that the agent certificates will be removed and DSM > agent communication will be disabled. You will have to re-register the agents. Note that the agent configuration stays in place so you do not have to reconfigure policies, keys, and so on.

5. Click **OK** in the dialog box.

6. Click **Apply** in the Edit Hosts page to finalize the configuration change.

   The **Certificate Fingerprint** for the agents should be gone.

7. Re-enable the **Registration Allowed** and **Communication Enabled** check boxes.

8. Click **OK**.

   The DSM is now ready to re-register the host.

Log on to hosts that needs to be re-registered. Refer to the *VTE Installation and Configuration Guide* for procedures to re-register the host.

Modifying Host Configuration

To modify a host configuration:

1. Select **Hosts > Hosts** in the menu bar. The Hosts window opens.

2. Click the link in the **Host Name** column of the host you want to modify. The Edit Host window opens to the **General** tab.

   > **Note**
   >
   > If the Edit Host page displays only two tabs, **General** and **GuardPoints**, you are working with a shared host. Check the host status in the **Sharing** column of the **Hosts** window. If you want to do more than add or remove GuardPoints, switch to the domain in which the host was created.

3. In the **Host Information Panel** you can modify the following:

   a. **FS Agent Locked**

      Locks the contents of the VTE Agent directories on the host.

   b. **License Type**

      Choose the type of license that will run on this host. Options are Perpetual, Term, and Hourly.

   c. **Password Creation Method**

      `Generate` (dynamic) or `Manual` (static)

   ○ If you switch the password method from `Manual` to `Generate`, regenerate the password. Select **Regenerate Password** and click **Apply**. A new generated password is downloaded to the host.

   ○ If you switch the password method from `Generate` to `Manual`, enter a new password in the **Password** and **Confirm Password** boxes.

   > **Notes**
   >
   > - If you configure a dynamic password for an agent that does not support the challenge-response feature, an ERROR-level audit message is generated and entered in the log after the agent registers with the DSM, plus a red warning message is displayed on the Edit Host window for the host. In effect, a randomly generated password is created and downloaded to the host system; however, the `vmsec challenge` command is not available on the host system so a user cannot display a challenge string. The solution is to change the host configuration from `Generate` to `Manual` and manually enter the host password.
   > - By default, the **Password** and **Confirm Password** text-entry boxes display dots. The dots are just graphic placeholders and do not indicate that a password had been entered. You must enter a password in both text-entry boxes or the `Manual` password method will not be applied to the host.

   d. **Description**

      Add or modify a description of the host.

   e. **Communication Port**

      You can change the port number used to exchange policy enforcement data between the DSM and the VTE Agent. Generally, you only change the port number when the default port number is already in use or if your firewall requires a different port number.

      If you change the port number, click **Ok**. The configuration change is downloaded to the VTE Agent host after the interval set by the **Update Host Frequency** parameter.

f. After the update is downloaded, you must manually restart the VTE Agent. The "host administrator", must execute one of the following commands on the VTE Agent host to restart the VTE Agent:

- On Linux, Solaris, and AIX:

  ```
  # /etc/init.d/secfs restart
  ```

- On HP-UX:

  ```
  # /sbin/init.d/secfs restart
  ```

- **On RedHat 7.2:**

  ```
  # /etc/vormetric/secfs restart
  ```

g. **System Locked**

   Applies an internal policy to the host to lock host system directories, like `/var`, `/bin`, `/etc`. This can be selected only if **FS Agent Locked** is enabled.

h. The **Support Challenge & Response** check box indicates whether this feature is enabled on the host. It becomes enabled when the VTE Agent running on the host registers with the DSM.

i. Enable the **Registration Allowed** check boxes for **FS**.

   Successfully registered agents display a hash value in the **Certificate Fingerprint** column. The **Registration Allowed** check box must be enabled before you can enable the **Communication Enabled** check box. An agent must be registered and the **Communication Enabled** check box enabled before you can apply policies to that agent.

   Configure keys and policies before enabling the host. You can optionally configure the host in a host group.

4. Click **Ok** to finalize the changes.

Changing the VTE Agent host password

The offline password feature is designed to protect the data on a laptop or similar portable system from being accessed by unauthorized users. You must provide a password when there is no connection between the VTE Agent and the DSM in order to decrypt/encrypt files that are encrypted with an offline key (`Cached on Host`). The offline password feature controls access to encryption keys that are stored locally on a particular machine as a way to keep data secure when the DSM is not accessible. Provide the password and the VTE Agent will encrypt/decrypt guarded data per the applied policy.

The host password is initially set when the host is added to the DSM. Passwords can be set on a host-by-host or host group basis.

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Select **Hosts > Hosts** in the menu bar.

   The Hosts window opens.

3. Click the host in the **Host Name** column.

   The Edit Host page opens to the General tab.

4. Select either `Generate` or `Manual` in the **Password Creation Method** scroll-list.

5. If you selected `Manual`, enter the new password in the **Password** and **Confirm Password** text-entry boxes.

   By default, the **Password** and **Confirm Password** text-entry boxes display dots, which makes you think that a password had already been entered or a default password is being used. The dots are just graphic placeholders and do not indicate that a password had been entered. You must enter a password in both text-entry boxes or the `Manual` password method will not be applied to the host group.

6. If you selected `Generate`, enable the **Regenerate Password** check box.

   You must enable the check box or the `Generate` password method will not be applied to the host group.

7. Click **Apply** or **Ok**.

8. When changing a static password, or changing a host from a dynamic password to a static password, tell the host user(s) the new static password or they will be unable to access encrypted data when there is no network connection between the host and DSM. When changing a host from a static password to a dynamic password, tell the host user(s) that challenge-response authentication has been enabled and that they need to run `vmsec challenge` on UNIX/Linux hosts, or select **Password...** on the Windows etray, when the host cannot connect to the DSM.

Deleting Hosts

When a host is deleted from the Management Console, the host record and configuration are deleted from the DSM only. The agent installations on the host continue to run, complete with the applied policies. To completely remove an agent host, run the software removal utility on the host system after you delete the host from the Management Console.

Only All/Security/Domain and Security administrators can delete hosts. If the host is shared with other domains, the DSM administrator must be in the same domain in which the host was first created in order to delete that host.

When a host record is deleted from the DSM, it pushes the configuration change to the VTE Agent running on that host. This change deletes VTE Agent certificates from the host and it deletes the "`URL`" line from the `agent.conf` file. The GuardPoints are removed, the host is no longer recognized by the DSM. If the agent tries to communicate with the DSM, the connection is refused.

Indications that a host has been deleted from the DSM

If there are missing certificates in the `./agent/pem` directory and no URL line in the VTE Agent `agent.conf` file, that is an indication that the host has been deleted. However, if the host is offline when it is deleted from the DSM, and the host identity is changed before the host comes back online, the GuardPoints will not be removed, the certificates will remain intact in the `./agent/pem` directory and the `agent.conf` file will be unchanged, but the agent and server still will not be able to communicate with each other.

The DSM URL is deleted from the VTE Agent `agent.conf` file when the host record is deleted from the Management Console.

Logging for the VTE Agent on the DSM is also affected. If you delete a host from the DSM while the host is offline, when the host comes back online, log messages concerning the denied connection can be viewed only by System or All administrators when not in a domain. This is because the DSM no longer has the host record, and does not know which domain the host belonged to, and cannot send messages to the appropriate log service.

When you delete a host you also delete that host from any host groups of which it may be a member.

Deleting a host

If a host has active GuardPoints, you will be prevented from deleting the host. A warning message is displayed telling you to unguard or disable the GuardPoints associated with the policy before you can delete the host. Make sure your data is accessible before you disable or unguard GuardPoints. This applies to LDT enabled hosts as well. Refer to the *Live Data Transformation Guide* for more information about data recovery and changing data from encrypted data to clear data on GuardPoints and for information about recovering data from LDT GuardPoints.

Deleting hosts with System or FS Agent Locks

Do not unregister or delete the VTE Agent while locks are applied. The locks stay in effect after the agent is unregistered and, without agent credentials, the DSM cannot administer that VTE Agent and it cannot disable the locks. You must boot the host into single-user mode and manually modify the agent configuration to disable the locks.

To remove everything associated with a host, including the agent software that runs on the host:

1. Apply a rekey policy and run `dataxform` on the host files that you want unencrypted.

2. Disable the locks for the host in the Edit Host window, **General** tab.

3. Remove all the GuardPoints for the host in the Edit Host window, **GuardPoints** tab.

4. Select **Hosts > Hosts** from the Management Console menu bar.

5. The Hosts window opens.

6. Enable the **Select** check box of each host to be deleted.

7. Click **Delete**.

    A dialog box opens that asks if you are sure you want to proceed with the operation.

8. Click **Ok**.

The host administrator with root permissions must log on to the host system and delete the agent software.

Deleting One Way communication hosts

In the case of one-way communication hosts, the host is deleted when the host receives the next status push from the DSM. If, for any reason the host cannot communicate with the DSM, a one- way communication host can be deleted manually as follows:

1. Select **Hosts > Hosts** from the Management Console menu bar. The Hosts window displays.

2. Select the host to be deleted. The **Delete Pending** column indicates the host as marked for deletion with a check mark.

3. Click on the host name to view the Edit Host page.

4. Clear the **Registration Allowed** checkbox, click **Ok** to return to the Host page.

5. Select the host again and click **Delete**. The host is removed from the DSM.

Configuring Host Groups

A Host Group is used to group one or more hosts to simplify configuration and administration. GuardPoints created on a host group are applied to all members of the group. Additionally, you can choose to apply host group configuration settings (except for password, FS Agent Lock, System Lock, Registration Allowed, and Communication Enabled settings), to all hosts that are members of that host group. It is important to keep this in mind when adding hosts to a host group. For example if you create an LDT policy in a host group, and then add a Docker enabled host to that host group, the Docker host will not be protected by that GuardPoint. Similarly, if you create a Linux file system GuardPoint in a host group, and then you add a Windows host to that host group, the Windows host will not be protected by that GuardPoint. See for more about creating GuardPoints on host groups.

The DSM supports two types of host groups; non-cluster and cluster. DSM cluster groups contain hosts that are members of a cluster with a cluster file system.

A host can be a member of more than one host group. However, membership in a cluster group is exclusive, so a host that belongs to a cluster, cannot join another cluster group, or host group.

Creating a host group

1. Select **Hosts > Host Groups** in the menu bar. The Host Groups window opens.

2. Click **Add**. The Add Host Groups window opens.

3. In **Host Group Name** field, enter a name for the new host group. This field is mandatory. The maximum number of characters is 64.

4. Select the host group type from the **Cluster Type** drop-down list. The options are **Non-Cluster**, **GPFS** or **HDFS**. If the host group is not a cluster group, select **Non-Cluster**. See "Creating a cluster host group" on the facing page for how to create a cluster group.

5. (Optional) Enter a phrase or string in the **Description** text-entry box that helps you to identify this host group. This field is optional. The maximum number of characters is 256.

6. Click **Ok**. The Host Groups page opens. The newly created host group is visible in the host group table.

Registering hosts to a host group

Add hosts to the host group using either the registration shared secret or the fingerprint method.

1. If using the shared secret method, create the host group and the shared secret to be used by hosts that will be added to that host group.

2. If hosts were added to the DSM using the fingerprint method, create a host group and add the hosts to the host group (we recommend that these steps be scripted for large scale deployments).

Creating a Registration Shared Secret for a host group

You can create a registration shared secret at the same time that you create a host group or, you can create a registration secret later once you have planned your host group creation.

1. Create the host group, click the Registration Shared Secret tab, or if you have already created a host group and you want to register hosts using the a shared secret, click the name of the host group on the Host Groups page and on the Edit Host Groups page, click the Registration Shared Secret tab.

2. When you use the registration secret feature for the first time, the **Current Registration Secret** section will not have any information. If there is an existing shared secret, a message, **Show Registration Shared Secret** is displayed, select **Yes** to view the secret. The default setting is No.

3. Enter the following information in the **Create New Registration Shared Secret** section:

   a. **Registration Shared Secret creation method**—The same constraints that apply to password creation, namely uppercase letters, numbers, and special characters required, apply to the shared secret creation.

      ○ **Manual**—This is the default method. Select this to create the shared secret yourself.

      ○ **Generate**—Select this option to get an automatically generated password.

   b. **Validity period**—Select the period for which the shared secret will be valid. Click the calendar icon to select the dates.

   c. **Require that hosts are added first**—(Optional) If you select this option, you need to first add the host to the DSM database with the **Registration Allowed** check box enabled before you install and configure the agent.

4. Click **Ok**.

   To remove an existing shared secret, click **Expire Registration Shared Secret**. The expiration date turns red to indicate that the shared secret is no longer valid.

Adding Hosts to a Host Group using Fingerprint method

1. Select the Member tab and then click **Add**. The Add Host window displays all configured hosts, with the exception of current host group members.

2. Select the hosts to add to the group.

   Select the hosts to add to the host group based on the policies to be applied. For example, if you want to apply file system protection, then the hosts you select should run the VTE Agent.

3. Indicate if the host is to maintain its current host configuration or if the host group configuration is to be applied to the host.

   At "`Do you want to apply the selected host(s) settings to host group settings?`"

   Select **Yes** to apply the host group settings for **System Locked**, **FS Agent Locked**, communication enabling, and so on to the hosts.

   Select **No** to add the hosts as they are and retain their individual configurations. This choice is not recommended. There is little reason to add a host to a host group and leave the host configuration intact. If you choose this option, you must be especially careful not to introduce configuration conflicts.

   The default is **Yes**.

   Click **Ok**. The Member tab displays the new host group members.

Refer to the *VTE Agent Installation Guide* for procedures to install and register the VTE Agent.

Creating a cluster host group

A cluster host group is a group of hosts that form a cluster.

1. Select **Hosts > Host Groups** in the menu bar. The Host Groups window opens.

2. Click **Add**. The Add Host Groups window opens.

3. In **Host Group Name**, enter the name of the new cluster host group. This field is mandatory. The maximum number of characters is 64.

4. Select the cluster group type from **Cluster Type** drop-down-list. The options are **Non-Cluster**, **GPFS** or **HDFS**. Select GPFS or HDFS depending on the type of file system on the host.

   Add the cluster nodes to the host group.

   > **Note**
   > GPFS is only supported on VTE Agent versions 5.x. See VTE Agent Release Notes for more information.

5. (Optional) Enter a phrase or string in the **Description** text field that helps you to identify this host group. The maximum number of characters is 256.

6. Click **Ok**. The Host Groups window opens. The newly created host group is visible in the host group table.

7. Click the host group in the **Name** column.

8. The Edit Host Group window opens. It has the following tabs: General, **GuardPoints**, Guard Docker (if you have a license for it), Sharing, Member, and Registration Shared Secret.

   If the group is an HDFS cluster group, you will see a tab labeled **HDFS**.

9. Click the **HDFS** tab to complete the HDFS cluster group configuration. Enter the following information:

   a. **Name Node URL**: Enter the URL of the Name Node. If Hadoop authentication is configured as Simple mode, only the NameNode URL information is needed in the URL format `hdfs://<host>:<port>`. By default the port number is 8020, but check the HDFS configuration to make sure this is so. For HDFS HA cluster, the URLs for both active and standby are required.

   b. **Second Name Node URL (HA)**: If this is a high availability configuration, enter the name of an HA Node.

   c. **Required Kerberos Authentication**: Select this check box if Kerberos authentication is required for the HDFS cluster.

- **Kerberos Principal**: Enter the name of the Kerberos principal

- **Kerberos Realm**: Enter the name of the Kerberos realm

- **KDC Host**: Enter the FQDN or IP address of the Kerberos Key Distribution Center (KDC)

- **Keytab File**: Enter the name of the keytab file to be used for authenticating HDFS cluster hosts. Click **Browse** to navigate to the file.

For more information about protecting data on HDFS configurations, see the *VTE Installation and Configuration Guide*.

10. Click **Ok** or **Apply** to save the configuration to the DSM database and then click **Test** to test the connection of the HDFS host to the Kerberos authentication server. The result of the test is displayed in the space above the tabs, if the test is successful, it displays 'Successful'.

Displaying host groups

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Select **Hosts > Host Groups** in the menu bar. The Host Groups window opens. All configured host groups are displayed.

Editing host groups

Once you create a host group and add hosts to the group, you can configure the host groups. The following can be modified or configured from the Edit Host Group page:

- Change a group description, enabling agent communication, locking VTE agent files on the host

- Change the VTE Agent password for the hosts in the host group.

- Enable policy enforcement, editing policies, applying policies

- Define GuardPoints

- Add hosts to a host group

The General tab allows you to enable agent communication for the host group, or enable System Lock or FS Agent Lock to control access to agent or system files.

- **Name**

  Name of the host group.

- **Description**

  Optional. Enter a description for the Host Group.

- **Enable FS Agent Communication**

  Select to enable/disable interactive communications of File System Agents installed on members of the host group.

- **Enable Key Agent Communication**

  Select to enable/disable interactive communications of key agents installed on members of the host group.

- **System Locked**

  Select to lock down the key operating system files of the hosts of members of the host group. (If this is enabled, patches to the operating system of the host will fail due to the protection of these files)

- **FS Agent Locked**

  Select to lock down the configuration of the File System Agent on the members of the host group. This will prevent updates to any policies on the members of the host group.

- **System Locked**

This check box is automatically selected when FS Agent Locked option is selected. It locks down the key operating system files of the host. If this is enabled, patches to the operating system of the host will fail.

- **Password Creation Method**

    Select the password method to use to unlock the agent. The host user may be prompted to supply a password to decrypt encrypted data when there is no network connection between the host and the DSM. The methods are **Generate** (challenge-response) and **Manual** (static password).

    When you select **Generate**, the host user must request a new password from a DSM administrator each time a host password is required. If you select **Generate**, an additional option is displayed; **Regenerate Password**. Enable this toggle to download a new randomly generated password to all hosts in the group.

    When you select **Manual**, the host user must request a new password from a DSM administrator each time a host password is required. Enter the password to apply to the hosts in the host group. The password is applied to each host in the host group and remains in effect when the hosts are removed from the host group or the host group is deleted. If you do not enter a password, the individual host password for each host in the host group remains unchanged. Enter the same password in the **Confirm Password** field to ensure that it had been typed correctly.

Host group password management

The DSM allows for host password management using host groups. For large-scale deployments where the DSM must manage several hundreds or thousands of agents, administering passwords on a per-host basis becomes untenable and administratively burdensome. Using a common password across all the hosts in a host group mitigates the administrative burden.

This feature is also useful for offline agent recovery. If a remote agent reboots (planned or unplanned) and cannot communicate with the DSM in the central office, it will prompt the administrator at the remote site to enter the host password. The remote site administrator typically calls the corporate help desk for the password. Using the password provided by the help desk personnel, the remote site administrator enables offline agent recovery and the resumption of services. Since the password is now known to the remote site administrator and the help desk personnel, it may result in a breach of security and/or render the IT operations to be non-compliant with respect to guaranteeing data privacy. To remedy the compromised situation, the security administrators should change the password—rotate the password—according to existing security practices. The host group password management feature allows changing the password on all the hosts in the host group when the password is compromised.

The use cases for host group password feature can be summarized as follows:

1. Set a common password for all hosts in a host group

2. Reset the common password for all hosts in a host group. If the password is provided to a remote agent administrator for offline agent recovery.

This feature is best used for deployments of scale when many agents are under the management of a DSM cluster.

Resetting a host group password

1. Select the host group whose password must be changed.

2. Apply the new password.

    When the new password is applied, the server pushes the password to all the hosts in the host group. Hosts that are removed from the host group retain the password set for the host group; hosts added to the host group later do not receive the new password.

    Pushing the host group password to thousands of agents is demanding on the DSM. Initiating other transactions while the password push is in progress may result in the server returning the following message: "Server busy please retry".

Protecting a host group

There are two ways to apply host protection. You can apply protection on a host-by-host basis or you can configure multiple hosts into a group and apply the same protection to all hosts in the group. Host groups are a convenient way to assign policies and keys simultaneously to a collection of hosts, rather than configuring each host individually.

You can configure hosts either before or after configuring host groups; however, creating hosts before creating host groups is quicker and requires fewer steps.

Before you apply GuardPoints:

- Create the initial host configuration in the Management Console for each host to be added the host group. See "Configuring Hosts" on page 236.

- Make sure that VTE Agent software is installed on each host system.

- Create encryption keys. See "Creating and Adding symmetric keys" on page 175.

- Configure the policies using the encryption keys your just created, to apply to the hosts in the host group. See "Creating and Configuring VTE Policies" on page 209.

To create and apply protection to a group of hosts running VTE Agents:

1. Create a host group, see "Creating a host group" on page 262.

2. On the Host Groups page, click the host group in the Name column.

3. The Edit Host Group window opens. It has five tabs: **General**, **GuardPoints**, **Sharing**, and **Member**.

4. Add hosts to the host group.

    a. Select the Member tab and then click **Add**. The Add Host window displays all configured hosts, with the exception of current host group members.

    b. Select the hosts to add to the group.

       Select the hosts to add to the host group based on the policies to be applied. For example, if you want to apply file system protection, then the hosts you select should run the VTE Agent.

    c. Indicate if the host is to maintain its current host configuration or if the host group configuration is to be applied to the host.

       At "`Do you want to apply the selected host(s) settings to host group settings?`" Select **Yes** to apply the host group settings for **System Locked**, **FS Agent Locked**, communication enabling, and so on to the hosts.

       Select **No** to add the hosts as they are and retain their individual configurations. This choice is not recommended. There is little reason to add a host to a host group and leave the host configuration intact. If you choose this option, you must be especially careful not to introduce configuration conflicts.

       The default is **Yes**.

    d. Click **Ok**. The Member tab displays the new host group members.

5. Apply GuardPoints.

    a. Select the **GuardPoints** tab.

       This tab displays the applied policies, the host groups to which the policies are being applied, and their enforcement status. Nothing is displayed if this is a new installation or no policies are applied.

    b. Click **Guard**. The Guard Host Group File System window opens to display all VTE Agent policies.

    c. Complete the policy application process.

For more about creating GuardPoints on a host group, see "Creating GuardPoints on a Host Group" on page 289. If a host group contains LDT enabled hosts, see "Creating LDT GuardPoints" on page 283. If the host group contains Docker hosts, see "Creating Docker GuardPoints" on page 283.

6. Select the **Guard Docker** tab

   This tab displays the applied policies, the host groups to which the policies are being applied, and their enforcement status. Nothing is displayed if this is a new installation or no policies are applied.

   a. Click **Guard**. The Guard Host Group File System window displays.

   b. Complete the policy application process.

7. Select the General tab. The **General** tab displays the host group name and its description. It is also used to enable the DSM to begin administering the host group members.

   a. Enable the **Enable FS Agent Communication** check box.

   b. The member hosts are administered as a group when you enable these check boxes.

   c. (Optional) Enable the **FS Agent Locked** and **System Locked** check boxes to apply protection—prevent the deletion or modification of VTE Agent installation files—to system files and VTE Agent files that reside on the host.

   d. (Optional) Set the password method for unlocking GuardPoints when the host cannot communicate with the DSM.

   The password method is applied to each host that is currently a member of the host group. The password method remains in effect until it is changed in the Edit Host Group window or the Edit Host window. If a host is removed from the group, or the group is deleted, the host retains the current password method. You can use the Edit Host window to change the password or password method of an individual host at any time.

   Select either **Generate** or **Manual** from the **Password Creation Method** scroll-list. **Generate** enables the challenge-response feature where the user displays a string on the host system, gives the string to the DSM administrator, and the DSM administrator returns a response string for the host user to enter. The response string is a single-use password that expires within 15 minutes. **Manual** is used to assign a static password to the host. The static password does not expire and can be used repeatedly until the DSM administrator changes it. The default method is **Generate** for non-cluster host groups and HDFS host groups, for GPFS cluster groups the only option is **Manual**.

   > **Note**
   >
   > If you select **Generate**, all the hosts in the host group must support the challenge-response feature. Hosts that do not support the challenge-response feature will still receive the randomly generated password; however, they will be unable to create the challenge string.

   The **Support Challenge & Response** field displays the dynamic password generation status of the host. The **Support Challenge & Response** field is not displayed in the Edit Host Group window. To determine if a host supports dynamic passwords, open the Edit Host window for the host to the **General** tab to display the **Support Challenge & Response** field on that tab.

   The **Password Creation Method** drop-down is used to apply a password creation method to the members of a host group only. It does not indicate the current password method for the host group. By default, the Edit Host Group window always displays the **Generate** password method when it is opened. Also, when the **Manual** password method is displayed, the dots in the password text-entry boxes do not indicate that a default password is provided or that a password had been entered.

8. If you switch the password method from `Manual` to `Generate`, regenerate the password.

   The **Regenerate Password** check box is displayed on the General tab when you change **Password Creation Method** from **Manual** to **Generate**. Select the **Regenerate Password** check box and click **Apply**. A new randomly generated password is created and downloaded to the hosts in the host group.

9. If you switch the password method from **Generate** to **Manual**, enter a new password.

   The **Password** and **Confirm Password** text-entry boxes are displayed. Enter the password to assign the hosts in the **Password** and **Confirm Password** text-entry boxes.

   Ignore the dots in the **Password** and **Confirm Password** text-entry boxes when you open the Edit Host Group window. They do not indicate a default password or that a password had already been entered.

   If you do not enter a password, the hosts in a host group retain their original passwords.

   > **Note**
   >
   > The host group password is not applied when a host is added to a host group. The `Do you want to apply the host group configuration to the selected host(s)?` field does not include the host group password. New host group members retain their original host password. To apply the host group password to the hosts in the group, change the password fields the `Edit Host Group` window and click **Apply**.

10. Click **Ok** to finalize the changes and close the window.

11. Check the configuration of each host in the host group.

    We recommend that you open each host in the Edit Host (not Edit Host Group) window to double-check that no configuration conflicts were introduced by adding the host to the host group. Also, check the status of GuardPoints to ensure that the GuardPoints and policies were applied as expected.

12. For VTE Agents, try accessing a GuardPoint to verify that the DSM and the host in the host group can communicate, as well as to verify the policy itself.

13. Display the DSM log to monitor the backup process.

Protecting a Docker host group

You can manage a group of Docker hosts by adding them to a host group and applying security policies to the host group.

Before you apply GuardPoints:

1. Create the initial host configuration in the Management Console for each host to be added the host group. See "Configuring Hosts" on page 236.

2. Install the VTE Agent software on each host system.

3. Create encryption keys. See "Creating and Adding symmetric keys" on page 175.

4. Configure the policies using the encryption keys your just created, to apply to the hosts in the host group. See "Creating and Configuring VTE Policies" on page 209

Apply a Docker GuardPoint

1. Select the **Guard Docker** tab

   This tab displays the applied policies, the host groups to which the policies are being applied, and their enforcement status. Nothing is displayed if this is a new installation or no policies are applied.

2. Click **Guard**. The Guard Host Group File System window displays.

3. Complete the policy application process; select the Docker host, the policy, the Docker image or container on which you want to apply the GuardPoint, and the path to the image directory or container volume on which to apply the GuardPoint.

If you are creating a Docker Image based GuardPoint, your Docker container stores the data in a Docker volume. You need to enter the path of the folder to be protected manually. This path should match the path that will be seen from inside the Docker container.

Sharing host groups

You can share the members of the host group with other domains. Sharing allows remote Security Administrators in other domains to administer GuardPoints on the local host. Only GuardPoints guarded by File System agents can be shared.

Host sharing example:

Hostgroup_1 in domain_1 is configured with two GuardPoints; gp_A, a manual guard set to /home/manual, and gp_B, an autoguard set to /home/autoguard. Hostgroup_1 has one member; host_1. If hostgroup_1 is now shared with domain_2, it means domain_2 imports hostgroup_1 and any hosts in domain_2 can be added as members of hostgroup_1. GuardPoint configurations defined in hostgroup_1 will now apply to any hosts from domain_2 that are added to that host group.

Share a host group:

1. Select the Sharing tab.

2. Click **Share**.

3. Enter the name of the domain with which to share the members of the host group in the **Domain Name** text-entry box.

4. Click **Ok**.

Remove sharing:

Click **Unshare** to remove sharing and return GuardPoints to the domain in which the host was configured.

Host Group Host Settings

Host Settings can be applied at the Host Group level. The Host Settings tab allows you to set authentication options for the applications running on the hosts in this host group. For a detailed explanation of Host Settings options, see "Host Settings" on page 248.

> ⚠️ **CAUTION**
> **Care must be taken while defining host settings at the host group level. If a host group contains member hosts with different operating systems (e.g., Linux and Windows), or host with Docker and non-Docker hosts, that inherit host settings from the host group, this may result in conflicts and affect file and user access permissions.**

A host that joins a host group has the option to inherit host group configuration, this includes host settings. If host settings have not been defined at the host group level i.e., left blank, then the host retains its own settings. If host settings at the host group level are modified later, then those settings will apply to all members of the group that are set to inherit configuration from that host group. Individual members of that host group will have host settings overwritten by the host group host settings. For example;

- `hostA` has host settings defined and then joins `hostGroup1` and inherits `hostGroup1` configuration. `hostB` also joins `hostGroup1` but, is not set to inherit the host group configuration. `hostGroup1` does not have any Host Settings defined, `hostA` retains it's own Host Settings and so does `hostB`.

- `hostGroup1` modifies its Host Settings, all members set to inherit host group settings will now have their individual settings overwritten by the host group Host Settings. `hostA` inherits the host group Host Settings but, `hostB` does not, as it does not inherit host group configuration.

- `hostB` then changes it's inheritance settings from the Host Settings tab to inherit settings from `hostGroup1`. The next time `hostGroup1` updates Host Settings, the changes will apply to both `hostA` and `hostB`.

A host can be a member of more than one host group. If the host is set to inherit host group configuration from the first host group it joins, and the next group it joins, it inherits the Host Settings of the last host group that it joins. For example;

- `hostC` joins `hostGroup2` and inherits the host group configuration, `hostC` now has `hostGroup2` Host Settings. `hostC` is then added to `hostGroup1` and is set to inherit host group configuration and so it gets `hostGroup1` host settings.

If a host group empties its Host Settings, any member hosts that inherit, retain the last Host Settings that were defined. For example;

- `hostGroup1` then deletes its Host Settings. All member hosts (`hostA`, `hostB`, and `hostC`) retain the last Host Settings defined for `hostGgroup1`—blank Host Settings are not passed on to members of the group. `hostB` leaves `hostGroup1`, and it retains the Host Settings it last inherited from `hostGroup1`.

If the Host Settings of a member of a host group are modified, that host no longer inherits Host Settings from the host group. For example;

- Host Settings on `hostB` are modified. Then the Host Settings for `hostGroup1` are modified, all members except `hostB` will inherit the changes made to the Host Settings for `hostGroup1`.

Configure Host Group Host Settings

1. Navigate to **Hosts > Host Groups**, click the host group for which to modify Host Settings, the Edit Host Group windows displays.

2. Click the Host Settings tab of the Edit Host Group window.

3. In the **Host Group Settings** text box, add `|authenticator|` before the path of the binary. (e.g., `|authenticator|/bin/su` to allow su to be a trusted method of authentication). For further consideration of authentication options, refer to "Host Settings" on page 248.

4. If you add another process to the set of trusted applications in the Host Settings, check the **Re-Sign Settings** check box to ensure that the new process is signed and authenticated by the host. The next time host settings are pushed to the VTE Agent host, the updated host settings are re-signed and the **Re-Sign Settings** check box on the DSM Console is cleared (or reset). If you do not select this option after adding a new process, the host will ignore the newly added process. See "Re-Sign Settings" on page 252 for more information about this setting.

5. Select one of the available choices from the **Apply Settings to Hosts** option:
   - **Only Hosts which currently inherit from this Host Group**: this will propagate changes only to the hosts that have chosen to apply Host group configuration.
   - **All hosts in this host group**: this will apply changes to all hosts that are members of this host group.

6. Click **Apply** after making changes to the host settings.

Change Host Group Host Settings inheritance

Hosts that are members of more than one host group inherit host group configuration (including host settings) from the last host group that they joined with inheritance set to 'Yes'. To change the host group from which to inherit Host Settings:

1. Navigate to **Hosts > Hosts** and click the host for which the host group host settings inheritance is to be changed, the Edit Host window displays.

   The **Host Settings from** field displays the host group from which the shared host inherits Host Settings.

2. From the **Make Host Settings inherit from** drop-down list, select the host group whose Host Settings you want to apply to this host.

3. Check the **Re-Sign Settings** check box to ensure that the new process is signed and authenticated by the host. The next time host settings are pushed to the VTE Agent host, the updated host settings are re-signed and the **Re-Sign Settings** check box on the DSM Console is cleared (or reset). If you do not select this option after adding a new process, the host will ignore the newly added process.

The Member tab of the Edit Host Group window displays where the host inherits its Host Settings, see "Adding hosts to a host group" below.

Adding hosts to a host group

The Member tab on the Edit Host Group window displays the following information about members of the host group:

- **OS Type**

  Indicates the host operating system type, e.g., Linux, Windows.

- **Host Name**

  The fully qualified domain name of the member host.

- **FS Agent**

  Indicates whether a VTE (FS) Agent is installed on the member host.

- **Key Agent**

  Indicates whether a Key (VAE/VKM) Agent is installed on the member host.

- **KMIP Agent**

  Indicates whether a KMIP client is installed on the member host.

- **One Way Comm**

  Indicates whether the installed agent is configured to use one way communication.

- **FS Agent Lock**

  If checked, indicates that the VTE (FS) Agent configuration on that host are locked.

- **System Lock**

  If checked, indicates that the key operating system files on the host are locked. If this is enabled, software patches applied to the operating system will fail.

- **LDT Enabled**

  Indicates whether this feature has been enabled or not.

- **Docker Enabled**

  Indicates whether this feature has been enabled or not.

- **Efficient Storage Enabled**

  Indicates that Efficient Storage is enabled. This feature requires KMIP and a Pure Storage Flash storage array.

- **Secure Start**

  Indicates whether this feature has been enabled or not.

- **Host Settings From**

  Indicates how the host gets its Host Settings. The following are possible:

  - This host - which means the host does not inherit host settings from any host group, they are set on the host.
  - This host group - which means the member host inherits its host settings from the current host group.

- ◦ *<Name of host group>* - which means that the member host inherits Host Settings from another host group of which it is a member.

Add hosts to a host group from the Member tab on the Edit Host Group page.

1. On the Member tab page, click **Add**. The Add Host window displays all configured hosts, with the exception of current host group members.

2. Select the hosts to add to the group based on the policies to be applied. For example, if you want to apply file system protection policies, then the hosts you select should run the VTE Agent.

3. Indicate if the host is to maintain its current host configuration or if the host group configuration is to be applied to the host.

4. The following message is displayed under the table listing the available hosts; "Do you want to apply the host group configuration to the selected host(s)?"

   a. Select **Yes**, to apply the complete host group configuration (except for the host group password) will be applied to this host including:

      - ◦ Host Settings from the selected host group
      - ◦ File System Agent Lock
      - ◦ System Lock
      - ◦ Registration Allowed
      - ◦ Communication Enabled

   b. Select **No** to add the hosts as they are and retain their individual configurations. This choice is not recommended. There is little reason to add a host to a host group and leave the host configuration intact. If you choose this option, you must take care not to introduce configuration conflicts.

      The default is Yes.

5. Click **Ok**. The Member tab displays the new host group members.

Deleting host groups

As part of DSM maintenance, you occasionally must remove host groups from the DSM. Deleting a host group removes only the group; the individual hosts that are members of that group remain intact. You cannot delete host groups that are configured with a policy. You must delete the host group GuardPoints from the host group before you can delete the host group itself. If you configured a host group password, the individual hosts retain the host group password.

To remove a host group:

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Select **Hosts > Host Groups** in the menu bar.

   The Host Groups window opens. All configured host groups are displayed.

3. Enable the selection check boxes of those host groups that you want to delete.

   The selection check boxes are located in the **Select** column of the Host Groups window.

4. Click **Delete**.

   You are prompted to verify the deletion.

> **Note**
>
> If you selected (System > General Preferences > System) **Strict Host Group Deletion** then if there are hosts and/or Guard Points in the host group, the delete fails. A message displays informing you that the host group cannot be deleted because it contains hosts and/or GuardPoints. Remove all remaining hosts and/or Guard Points before deleting the host group. If you did not select **Strict Host Group Deletion**, then the delete succeeds.

5. Click **Ok**.

Chapter 24: Managing GuardPoints

GuardPoints are directories protected by VTE Agent security policies. Access to files and encryption of files in protected directories is controlled by security policies.

Overview

Before you apply GuardPoints you must do the following:

- Add a host to the DSM, see Chapter 23: "Configuring Hosts and Host Groups" on page 228.

- Install and register the VTE Agent on the host system, as described in the VTE Agent Installation and Configuration guide.

- Create encryption keys, see Chapter 21: "Managing Keys" on page 173.

- Configure policies using the encryption keys you created, see Chapter 22: "Policies" on page 207.

- Create a GuardPoint.

> **Note**
> Check that no one is using the directory to be guarded before making it a GuardPoint.

If users are working in the directory when it is made into a GuardPoint, users can continue to use data in memory rather than use the actual data in the GuardPoint. Tell users to:

1. Save their work.

2. Close applications that are running in the directory.

3. Exit the directory before applying the GuardPoint.

   When they re-enter the directory they will use protected data and the VTE Agent will work appropriately.

This chapter also describes how to create Secure Start GuardPoints, LDT GuardPoints on LDT enabled hosts or host groups, on container images, and on containers on hosts.

See "Understanding GuardPoints" on page 277 for how to create GuardPoints on a host and "Creating GuardPoints on a Host Group" on page 289 for how to create GuardPoints on a host group.

Considerations before creating a GuardPoint

1. If a host is to be added to a host group, do not apply a GuardPoint at the host level, rather, apply the GuardPoint at the host group level. You can do both, but it is harder to keep track of GuardPoints applied at the host group level and custom GuardPoints applied at the host level.

2. Certain directories are protected against guarding, plan your GuardPoints accordingly:

   a. The top-level "Program Data" folder on Windows Vista and Windows 2008, and the top-level "Documents and Settings" folder on all other Windows platforms, cannot be guarded because a GuardPoint cannot be applied to a folder that contains open files. The same is true for the "Users" folder. The VTE Agent opens and continually maintains log files in subfolders under "ProgramData" and "Documents and Settings". Other subfolders below "ProgramData" and "Documents and Settings" can be guarded as long as there are no open files in any subfolder at the time the GuardPoint is applied.

   Be especially careful when specifying paths for Windows agents. Cross-guarding the same folder with different policies and encryption keys will give unexpected results and will corrupt the files in that folder.

   GuardPoint paths must use standard Windows path notation and delimiters. Incorrect notation and delimiters are ignored and discarded by the Windows agent. Therefore, it is possible to enter two paths that resolve to the same Windows folder and successfully guard both of them. The DSM reports that it is guarding two unique folders when, in fact, it is guarding the same folder twice.

   Do not use any of the following characters as path delimiters: `|?<>:*"/,`

   For example, both `C:\gp\` and `C:\gp/\` are allowed by the DSM. When the second GuardPoint is applied, the extraneous "/" is discarded by the Windows VTE Agent and the Windows VTE Agent applies a GuardPoint to `C:\gp\` a second time.

   b. On Linux, the following directories cannot be guarded:
   - `<secfs install root>/agent/secfs/`
   - `<install root>/agent/secfs/bin` and all subdirectories
   - `<secfs install root>/agent/vmd` and all subdirectories
   - `/etc/vormetric` and all subdirectories
   - `/etc`
   - `/etc/pam.d` and all subdirectories
   - `/etc/security` and all subdirectories
   - `/usr`
   - `/usr/lib`
   - `/usr/lib/pam`
   - `/usr/lib/security` and all subdirectories
   - `/etc/rc*` and all subdirectories
   - `/var/log/vormetric`

   c. You cannot apply VTE Agent protection to already mounted and guarded directories, nor can you nest GuardPoints.
   The `/opt/vormetric/DataSecurityExpert/agent/secfs/.sec` directory is automatically mounted and guarded by `secfs` when the VTE Agent process starts on the host. You cannot apply a GuardPoint to `/opt` because it contains the existing GuardPoint, `/opt/vormetric/DataSecurityExpert/agent/secfs/.sec`; however, you can guard a directory like `/opt/myapps` because it is in a different hierarchy and has no impact on `/opt/vormetric`.

   Display mounted and guarded directories using the `df` command.

3. As of the v6.1.0 release, both DSM and VTE support a new enhanced encryption mode (CBC-CS1). If your host groups contain v6.1.0 VTE hosts and other hosts with earlier versions of VTE, you **cannot** apply policies containing keys that use this new encryption mode. The action fails with an error message informing you that all

hosts in the host group do not support the key's encryption mode. Only hosts with VTE v6.1.0 support the new encryption mode. Refer to "Creating and Adding symmetric keys" on page 175 and to the *VTE Agent Installation and Configuration Guide* for more about the new encryption mode.

Changing a policy or rekeying a GuardPoint

To change a policy or rekey a GuardPoint, be prepared to temporarily stop access to the GuardPoint. Changing policies for a GuardPoint requires an interruption of service because the transition process entails disabling one policy and then enabling another policy. The GuardPoint must be inactive during the transition period to ensure GuardPoint integrity. The same rule applies to moving a host between host groups when it includes a change in policies. Coordinate policy changes during a maintenance outage window.

If Live Data Transformation (LDT) is enabled on your hosts, encryption and rekeying of GuardPoint data is done without blocking user or application access to the data. LDT is a separately licensed feature, refer to "Enabling Live Data Transformation" on page 240 and the *VTE Live Data Transformation User Guide* for more information about implementing LDT.

Understanding GuardPoints

> **Note**
> Information about UNIX agents applies to earlier versions of those agents, since as of v6.0 UNIX agents are EOL. Refer to the DSM and VTE Release Notes for more information about supported agent versions.

1. Log on to the Management Console as an administrator with `Host` role permissions.

2. Select **Hosts > Hosts** on the menu bar.

   The Hosts window opens.

3. Click the target host in the **Host Name** column. The Edit Host window opens to the General tab for the selected host.

   ○ Select the **Registration Allowed** option so that the VTE Agent running on the target host can register with the DSM.

   ○ Select the **Communication Enabled** option so that the DSM can push policy and configuration changes to the host, and for the DSM to accept VTE Agent policy evaluation requests.

   ○ To create LDT GuardPoints, select the **Live Data Transformation**option, see "Creating LDT GuardPoints" on page 283.

   ○ To create ESG GuardPoints for an ESG storage array, select **Efficient Storage**.

   ○ To create IDT GuardPoints for an IDT device, such as a hard drive, select In-place data transformation - device.

   > **Note**
   > The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

   ○ To create Docker GuardPoints, select the **Docker Enabled** option, see "Creating Docker GuardPoints" on page 283.

   ○ To create a Secure Start GuardPoint, select the **Secure Start GuardPoint** option. This feature is only supported on Windows hosts. See .

4. Click **GuardPoints** tab.

The panel displays applied policies in a tabular format. The table consists of the following columns:

   ○ **Policy**

Name of the policy applied to the GuardPoint.

- ° **Host group**

  Name of the host group of which the current host is a member

- ° **Protected Path**

  GuardPoint path that is protected

- ° **Disk/Disk Group**

  If a raw partition is a member of an Oracle ASM disk group, it is displayed in the form, group_name/disk_name.

- ° **Type**

Type of GuardPoint being applied on a **Cloud Object Storage** host:

- Cloud Object Storage (Auto Guard)
- Cloud Object Storage (Manual Guard)

Type of GuardPoint being applied on a **UNIX** host:

- Directory (Auto Guard)
- Directory (Manual Guard)
- Raw or Block Device (Auto Guard)
- Raw or Block Device (Manual Guard)

Type of GuardPoint being applied on a **Windows** host

- Directory (Auto Guard)
- Raw or Block Device (Auto Guard)

- ° **Domain**

  Domain in which the host is administered

- ° **Auto Mount**

  Indicator of the file system mount type, whether a regular mount or an automount

- ° **Enabled**

  Displays the policy enforcement status, can be either enabled or disabled.

- ° **Secure Start**

  Indicates whether the GuardPoint is a Secure Start GuardPoint. This can be enabled or disabled by selecting the GuardPoint and clicking **Secure Start On**, or **Secure Start Off** as applicable.

- ° **Transform Sparse Regions**

  Indicates whether transform sparse regions is enabled or not. If this was set when creating the GuardPoint, you can disable it by unchecking the option in the column. Once disabled, it cannot be re-enabled. This column is displayed only if LDT is enabled for that host.

- ° **Status**

  Connection status to the host

- ° **Rekey Status**

  Indicates the transformation status of the data rekey operation.

You can perform any of the following functions from this tab:

- ° **Refresh**

  Update the Edit Host page.

- ° **Suspend Rekey**

  Click to suspend rekey or data transformation operations, for all GuardPoints on the selected host.

- ○ **Re-Push Policies**

  Click to push a policy update to a host. For example, if a rekey operation is underway on your host and you rotate the encryption key, the agent will not accept the policy push. You can re-push the policy until the agent accepts it and performs the rekey operation again.

- ○ **Transform Sparse Regions**

  This is only applicable for LDT policies. If you selected this option while creating a GuardPoint, it means that sparse file regions will be transformed. Once selected, this option cannot be disabled. If you did not select this option while creating a GuardPoint, sparse regions will not be transformed. You have one opportunity to disable this option from the **GuardPoints** tab. Once you change the setting, you cannot roll it back. It is a one-time change.

- ○ **Secure Start On**

  This button is displayed only if the Secure Start feature has been enabled on the host. Select this option to create a Secure Start GuardPoint.

- ○ **Secure Start Off**

  This button is displayed only if the Secure Start feature has been enabled on the host. Select this option to turn off Secure Start for the GuardPoint.

Creating and Applying a GuardPoint

1. Select the GuardPoints that you want to change or click **Guard** to add a new one.

2. Select the type of policy to apply from the **Policy** drop-down menu. You must select a policy before you can browse the agent file system.

   If **LDT** is enabled on your host, then the Live Data Transformation policy type is available, see "Enabling Live Data Transformation" on page 240 for more information.

   If your agent supports **Cloud Object Storage**, then the COS policy type is available.

   Later, when you select the directories to configure as GuardPoints, if you select multiple directories, they are configured with the currently selected policy.

3. Select the type of GuardPoint to apply in the **Type** drop-down menu.

   UNIX/AIX options:
   - ○ **Directory (Auto Guard)**
   - ○ **Directory (Manual Guard)**
   - ○ **Raw or Block Device (Auto Guard)**
   - ○ **Raw or Block Device (Manual Guard)**

   Windows options:
   - ○ **Directory (Auto Guard)**
   - ○ **Raw or Block Device (Auto Guard)**

   Cloud Object Storage options:
   - ○ **Cloud Object Storage (Auto Guard)**
   - ○ **Cloud Object Storage (Manual Guard)**

   IDT/ESG options:
   - ○ **Raw or Block Device (Auto Guard)**
   - ○ **Raw or Block Device (Manual Guard)**

   > **Note**
   > The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

4. In the **Path** text box:

   ○ Enter the full paths of one or more directories in the **Path** text-entry box and click **Ok** to apply the policy to the target GuardPoint. Enter one path per line in the **Path** text-entry box.

   > **Note**
   >
   > DSM cannot discern between the two styles used for the Cloud Object Storage GuardPoints. Therefore, use only one GuardPoint style:
   >
   > **Path style**: `https://s3.amazonaws.com/vte-repository`
   > **Virtual host style**: `https://vte-repository.s3.amazonaws.com`

   ○ Enter part of a directory path in the **Path** text-entry box and click **Browse** to jump to the specified point. From there, you can use the browser to descend further into the directory hierarchy and select one or more directories to be guarded.

   ○ Click the **Browse** button to locate the host directory to guard. The Remote File Browser window opens.

   If a target GuardPoint exists, use the browser to select the GuardPoint path. If it does not exist, be sure to enter the GuardPoint path correctly. The DSM does not parse manually entered paths for correct syntax.

   See, for what to be aware of before creating a GuardPoint.

   > **Note**
   >
   > When browsing a Docker image on a host, volumes created on a container run off that image are not visible if that container has been removed. If you want to create a GuardPoint on a container volume that container must exist (e.g. running or stopped), in order for the volume to be visible. Or, you can manually enter a path for a volume you want to guard, and then when a container instance is run off that image, you must remember to create those volumes in order for the GuardPoint to apply.

   ○ Find target GuardPoints. Click the plus symbol (+) next to a folder to display the next level of the directory hierarchy. Click the minus symbol (-) to collapse the hierarchy. Click a folder or file name to select that directory or file.

Figure 24-1: Browsing for GuardPoints
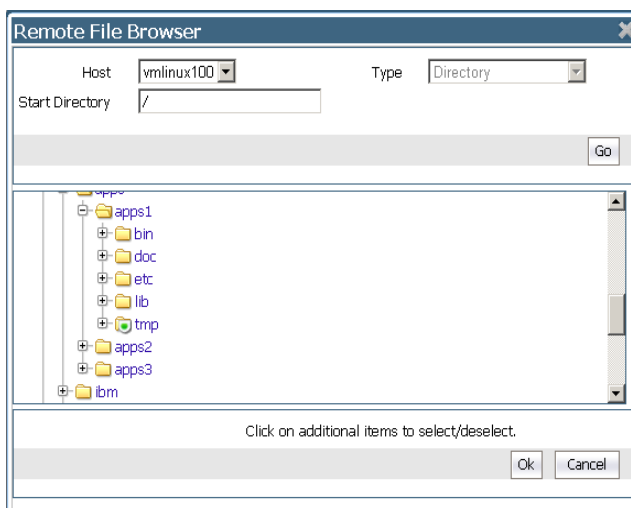
Configured GuardPoints are displayed as folders overlaid with a shield icon. If you suspect that the GuardPoint status is incorrectly indicated, note that the agent status displayed in the window shows the status as it is configured on the DSM. It is not a real-time indication of the actual status. For actual status, log onto the agent system and run VTE Agent utilities, like "`vmsec status`" and "`secfsd -status guard`". Compare the two to

ensure that the GuardPoint status on the DSM and VTE Agent match. If the two do not match, go with what you see on the agent. The shield indicates a configured GuardPoint only. The GuardPoint can be enabled or disabled and the shield will still be displayed. The shield remains displayed until the GuardPoint is unguarded (deleted).

To quickly traverse different directory hierarchies, you can enter part of the path to the GuardPoint in the **Start Directory** text-entry box and click **Go**, or press `<Enter>`, to display and select the rest of the path.

5. Select one or more directories to be configured as GuardPoints.

   Single-click a directory in the scroll-list to select an individual directory. Hold the `<Ctrl>` key down to select multiple directories. Hold the `<Ctrl>` and `<Shift>` keys down to select a range of directories.

   You should check that no file or directory below a selected GuardPoint is being accessed. If something under a GuardPoint is being used or accessed, the DSM may not be able to take control of the directory and apply protection.

   Keep the following in mind while selecting a GuardPoint path:

   a. The maximum number of characters allowed in a GuardPoint path is determined by your operating system. You can specify a GuardPoint path up to the restriction imposed by the host operating system. However, we recommend that you keep it below 1,000. Beyond 1,000 characters, the path information for the **Resource** field in the Message Log and host messages file (for instance, `/var/log/messages`) is truncated, and the Key and Effect fields that normally follow the **Resource** field are not displayed.

   b. The directory (or directory path) specified in a resource set is appended to the GuardPoint. This means if the GuardPoint is `/mnt/remote2` and the resource set directory path is `/remoteDir`, then the policy is applied to the files and directories in `/mnt/remote2/remoteDir`.

6. The **Auto Mount** check box disappears when **Directory (Manual Guard)** or **Raw or Block Device (Manual Guard)** is selected because only regular mounts are supported by these types.

   > **Note**
   > The **Auto Mount** option is not supported for GuardPoints on Docker hosts and on Cloud Object Storage devices. Cloud Object Storage GuardPoints do not support Transform Sparse Regions.

7. **Secure Start**: This option is only available if the Secure Start feature has been enabled. Select this option to create a Secure Start GuardPoint for a service or raw or block device that starts before VMD starts.

8. **Transform Sparse Regions**: This option is only available if LDT is enabled. Select this option to encrypt any sparse file regions in the GuardPoint during LDT. A sparse file is one that contains data that consists of mainly zeros.

9. **Efficient Storage Device**: Provides encryption for a Pure Storage Flash storage array. Efficient Storage GuardPoints protect Efficient Storage devices. These features are only available if you have KMIP and a Pure Storage Flash storage array.

10. **In-place Data Transformation - Device**: IDT GuardPoints are GuardPoints that protect the offline transformation of ESG or IDT devices that already contain user data. Use an IDT GuardPoint to protect raw devices.

    > **Note**
    > The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

Figure 24-2: GuardPoint selection
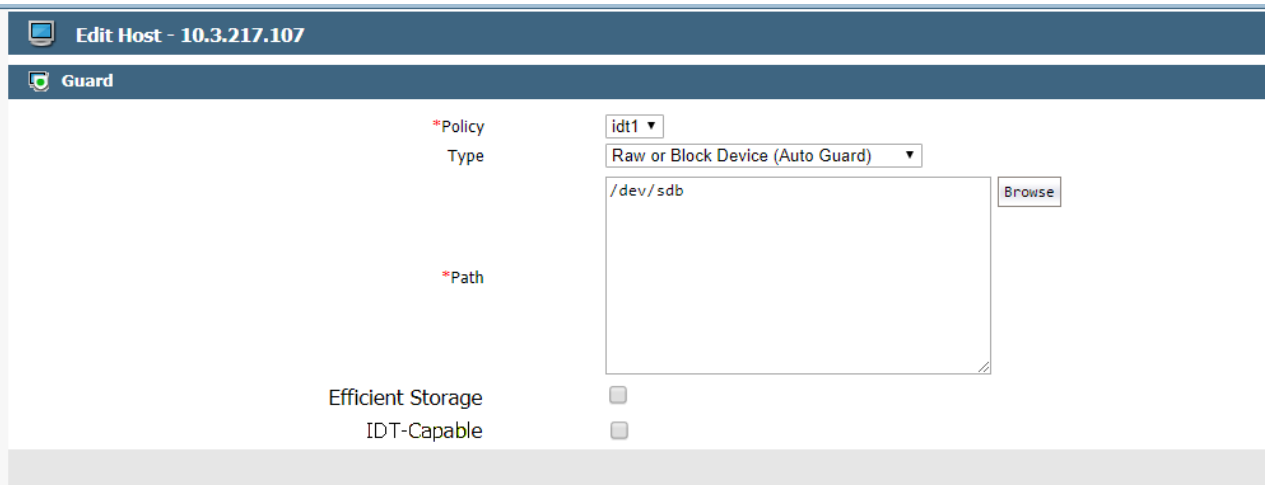
11. Click **OK**.

    The Edit Host page is updated to display the new GuardPoint or GuardPoints.

    Note the GuardPoint status:

    ° A green circle indicates an active and healthy connection to the agent system.

    ° A red square indicates that a policy has been configured but not applied on the agent system; that a GuardPoint is disabled or is in the process of being disabled; or that a communication error has occurred between the DSM and agent systems.

    ° A yellow triangle indicates that an attempt to delete a GuardPoint is still pending. The DSM awaits confirmation from the agent before it deletes the GuardPoint from the DSM. A yellow triangle also indicates a GuardPoint that is `Directory (Manual Guard)` or `Raw or Block Device (Manual Guard)` that is not mounted on the host system.

12. Wait a moment then click the **Refresh** button to update the display.

    The red square should change to a green circle.

    It may be easier to execute the `df` command repeatedly on the host system until you notice a `secfs` mount for the new GuardPoint, or, execute `tail -f /var/log/vormetric/vorvmd_root.log` and wait until a message like the following is displayed:

    ```
    Successfully received and implemented a new security configuration.
    ```

LDT Quality of Service

If you have a Live Data Transformation (LDT) license and the LDT feature enabled on your host, this tab displays **Quality of Service** in the top panel of the **GuardPoints** tab. The QoS feature allows administrators to maintain operational efficiencies in their systems in conjunction with LDT operations. QoS lets administrators specify percentage of CPU usage or a rekey rate and schedules for LDT operations. (See the *Live Data Transformation Guide* for best practices about using LDT and QoS.) The following options are available:

- **Rekey Option**: Choose to use Rekey Rate to define QOS or CPU/IO utilization.

  > **Note**
  > This feature is supported in VTE v6.1.2+. If you do not have a supported version, you will not see these fields.

- **Schedule**: Select a schedule to run LDT. The options are:

  ° ANY_TIME: LDT runs any day at any time of the week

- ○ WEEKENDS: LDT runs between 9:00 PM Friday to 7:00 AM on Monday

- ○ WEEKNIGHTS: LDT runs between midnight to 7:00 AM from Monday to Friday

You can also create custom QoS schedules:

1. Navigate to **Hosts > QoS Schedules**, click **Add**.

2. The Add/Edit QoS Schedule page displays. Enter a name for the schedule and a description (optional). Click **Add** again.

3. The scheduling options are displayed. You can make the following selections:

   - ○ **Starting Day**: Day of the week to start the LDT process

   - ○ **Ending Day**: Day of the week to end the LDT process

   - ○ **Start Time**: Time at which to start the LDT process.

   - ○ **Ending Time**: Time at which to stop the LDT process

4. Click **OK**, then click **OK** again, to go back to the QoS Schedules page.

   The new schedule is listed on the table and is also available in the **Schedule** drop down list in the LDT Quality of Service panel on the **GuardPoints** tab.

   - ○ **Set % of available CPU usage for rekey**: If you chose to use the CPU/IO for rekey, then define what percentage of the host servers CPU should be reserved for LDT rekey operations. Refer to the LDT Guide for more information about rekey operations.

   - ○ **Cap CPU Allocation**: If you chose to use the CPU/IO for rekey, then, select this option to cap CPU usage to the percentage defined in Set % of available CPU usage for rekey. If you do not select this option, LDT operations will utilize all of the available CPU memory.

   - ○ **MB/s rate**: If you chose to use the rekey rate for QOS, then define the rekey rate in MB/s. For rate guidance, refer to the LDT Guide.

Creating LDT GuardPoints

To create an LDT GuardPoint:

1. Create an LDT policy.

2. Set the Quality of Service, see "LDT Quality of Service " on the previous page.

3. Click **Guard** on the **GuardPoints** tab to apply an LDT policy to a directory or file and create a GuardPoint.

See the *VTE Live Data Transformation User Guide* for more information about creating policies, creating QoS schedules, and creating LDT GuardPoints. See "Creating and Configuring VTE Policies" on page 209 for procedures to create LDT policies.

Creating Docker GuardPoints

GuardPoints can be created for Docker images or for docker containers. Before creating GuardPoints on Docker images and containers, the following must be taken into consideration:

- In order to use Vormetric data security protection, you must add the Docker engine process to the Host Settings, see "Host settings for a Docker enabled host" on page 249.

- When applying GuardPoint policies to Docker containers, users must ensure that the root user has at least 'permit' effect on the GuardPoint, or else the GuardPoint will be completely inaccessible to all users, even for users with 'apply_key', and 'permit' effects.

- If you create a Docker image-based GuardPoint, that GuardPoint is pushed to any container that is run off that image. A Docker container started from that protected image, stores data in a Docker volume. To protect volumes used by the container, you need to enter the path of the folder to be protected manually. This path should match the path that will be seen from inside the Docker container.

1. Log on to your DSM as an All/Security/Domain and Security administrator.

2. Navigate to **Hosts**.

Figure 24-3: Guard Docker tab
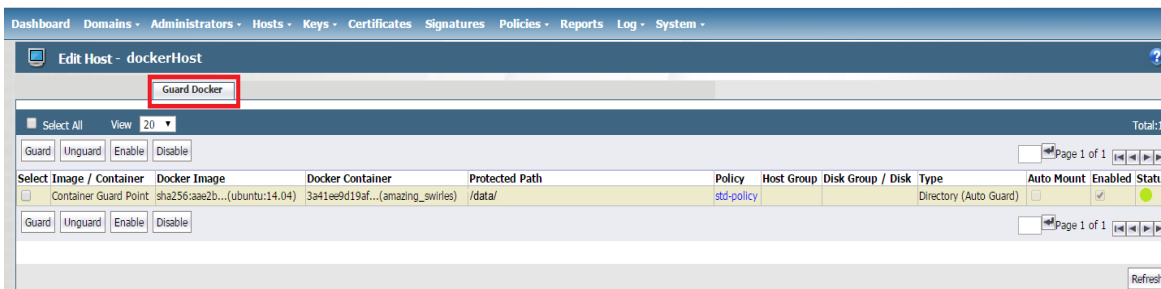
3. On the Hosts page, click the name of the host in the **Host Name** column, the Edit Host page opens.

4. Click the **Guard Docker** tab.

5. Click **Guard** to open the Guard File System page, from where you can select a policy to apply to a Docker image or container on your docker host.

6. Select a policy to apply to the GuardPoint you are about to create.

7. Click **Browse** next to the **Docker Image/Container** field to browse the Docker host for an image or container to which to apply the policy.

8. Select the type of directory to guard.

9. Click **Browse** next to the **Path** text box to browse the image or container for a file path to add the GuardPoint.

Figure 24-4: Guard File System: Select a Docker image or container
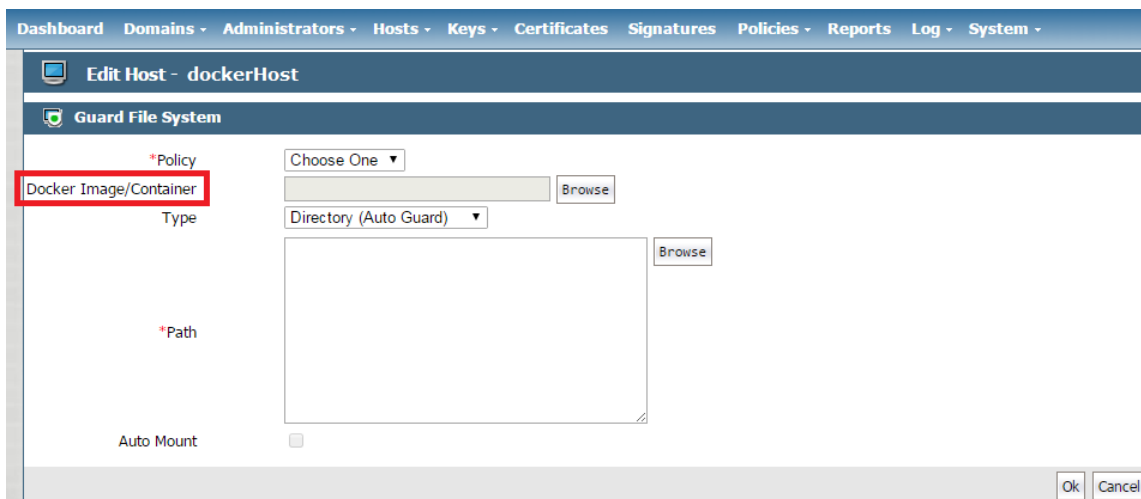
10. Click **Ok**, the Edit Host page opens with the newly created GuardPoint listed in the table.

> **Note**
> Auto Mount is not supported in a Docker environment.

Refer to the *VTE Agent Installation and Configuration Guide* for details about transforming data on Docker image and container GuardPoints.

Creating Efficient Storage GuardPoints

Efficient Storage is a feature that enables GuardPoint protection for Efficient Storage devices.

The initial Efficient Storage GuardPoint must guard the Efficient Storage device as a raw or block device on the DSM. When guarding a device as a raw or block device, the Guard File System page shows the Efficient Storage Device option. You must select this option. If the Efficient Storage Device option is not selected, the VTE host will not enable the device as an Efficient Storage GuardPoint. Consequently, it does not share the key with the storage array. The following example shows that the LUN (/dev/sdb) is guarded as a block device:
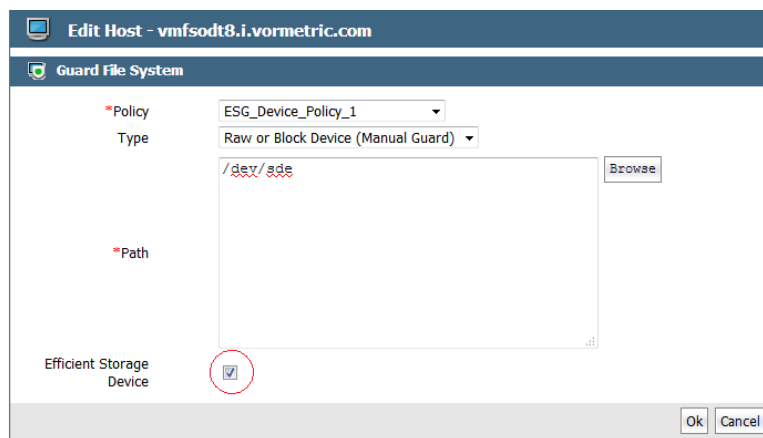
Refer to the chapter 'Using VTE with Efficient Storage Devices" in the *VTE Agent Installation and Configuration Guide* for information on setting up and initializing the Efficient Storage device or the IDT device.

Creating IDT Device GuardPoints

IDT GuardPoints are GuardPoints that protect the offline transformation of ESG or IDT devices that already contain user data. You can use an IDT GuardPoint to protect raw and block devices. You can also use it anywhere that you would use a standard GuardPoint.

> **Note**
> The IDT feature is only compatible with VTE v6.3.0 or succeeding versions.

When guarding a device as a raw or block device, the Guard File System page shows the IDT-Capable Device option. You must select this option. If it is not selected, the VTE host will not enable the device as an IDT GuardPoint.

Enabling GuardPoints for Raw Devices Using Secure Start

Secure Start is a feature that enables GuardPoint protection for raw devices which start earlier in the boot sequence than VMD (VTE agent daemon). Using Secure Start, those raw devices can now be protected.

If a user enables Efficient Storage, then DSM automatically enables SecureStart on the host, if it is not already enabled. When you apply a GuardPoint for a raw device on an ESG-enabled host, DSM enables Secure Start for that GuardPoint.

> **Note**
> This feature is only supported on hosts running Windows OS.

For more information about protecting raw devices using Secure Start, refer to the Secure Start chapter in the *VTE Agent Installation & Configuration Guide*. To determine if another application qualifies for Secure Start, contact Thales technical support.

Access to a Secure Start GuardPoint is only permitted during the boot sequence and for a short period of time. Once the VMD is up and running, it performs the normal agent initialization and communicates with the DSM to access files within a GuardPoint location.

To apply Secure Start protection to servers:

1. Click **Hosts > Hosts > <hostName>** on the Management Console.

2. In the **General** host information section, select the option: **Secure Start GuardPoint**.

3. Click **GuardPoints** tab.

4. Click **Guard** to set a new GuardPoint.

5. In the **Policy** field, select a Standard or LDT Production policy.

6. Set **Type** to Raw or Block Device (Auto Guard).

7. Click **Browse** and navigate to the device.

8. Select the option: **Secure Start**.

9. Select the option: **Efficient Storage Device**.

10. Click **OK**.

11. Select the GuardPoint and click **Secure Start On**.

Creating COS GuardPoints

Cloud Object Storage GuardPoints encrypt the contents of the Cloud Object Storage (COS). For version 6.4.0, only the AWS S3 bucket is supported.

> **Note**
> The DSM Cloud Object Storage feature is compatible with VTE v6.3.0 and all subsequent versions.

Security/All administrators can create new GuardPoints or edit existing host GuardPoints.

> **Note**
> DSM allows **exactly one** COS GuardPoint per S3 bucket.

The following features are NOT supported in COS GuardPoints:

- Transform Sparse Regions
- Secure Start
- Browsing to a directory to be guarded
- Host to Browse
- Auto Mount
- Efficient Storage

> **Note**
> In the Edit Host window, you will see that the **Guard FS** tab has been renamed to **GuardPoints**.

To create a COS GuardPoint:

1. In the Hosts window, click on the host for which you want to set GuardPoints.

2. In the Edit Host window, click **GuardPoints**.

3. In the GuardPoints tab, click **Guard**.

4. In the Policy dropdown menu, select an appropriate policy.

5. For type, click Cloud Object Storage (Auto Guard or Manual Guard).

6. In the Path field, enter the path for the GuardPoint.

7. Click **OK**. COS GuardPoints display on the GuardPoints tab of the Host Detail page.

Creating Cloud Object Storage GuardPoints in a Host Group

Security/All administrators can create new COS GuardPoints or edit existing GuardPoints for Host Groups.

> **Note**
> In the Edit Host window, you will see that the **Guard FS** tab has been renamed to **GuardPoints**.

The following features are NOT supported in COS GuardPoints:

- Transform Sparse Regions
- Secure Start
- Browsing to a directory to be guarded

- Host to Browse

- Auto Mount

- Efficient Storage

> **Note**
> DSM allows **exactly one** COS GuardPoint per S3 bucket.

To create a COS GuardPoint:

1. In the Host Group window, click on the host group for which you want to set GuardPoints.

2. In the Edit Host Group window, click **GuardPoints**.

3. In the GuardPoints tab, click **Guard**.

4. In the Policy dropdown menu, select an appropriate policy.

5. For type, click **Cloud Object Storage (Auto Guard or Manual Guard)**

6. In the Path field, manually enter the path for the GuardPoint.

7. Click **OK**. COS GuardPoints display on the GuardPoints tab of the Host Detail page.

Valid Styles for Cloud Object Storage GuardPoints

You can create COS GuardPoints for S3 buckets with either the path style, or the virtual host style:

- Path style: https://s3.amazonaws.com/vte-repository

- Virtual host style: https://vte-repository.s3.amazonaws.com/

Both styles are supported. Note that both styles represent the SAME GuardPoint. You can only guard a GuardPoint once. So if you successfully guard https://vte-.s3.amazonaws.com/ and then try to guard, https://s3.amazonaws.com/vte-repositrepositoryory, DSM generates the following error:

```
GuardPoint "https://vte-s3-dev.s3.amazonaws.com/" with policy "thales-s3-pol" was created on
host "192.168.255.63". It was disabled because the same path is already guarded.
```

Valid Directory Levels for Cloud Object Storage GuardPoints

DSM only supports guarding the entire cloud bucket. It does not support guarding directories within the bucket. Therefore, the following GuardPoints are invalid because they contain directories below the bucket level:

- https://s3.amazonaws.com/vte-repository/directory1

- https://vte-repository.s3.amazonaws.com/directory1/directory2

If you try to create a GuardPoint for a bucket with a sub-directory, DSM generates an error message that states that it is an invalid path.

The following GuardPoints are valid because they are guarding the entire S3 bucket:

- https://s3.amazonaws.com/vte-repository

- https://vte-repository.s3.amazonaws.com

> **Note**
> For COS GuardPoints, you cannot click Browse to locate the host directory to guard. You must enter the path manually.

Creating GuardPoints on a Host Group

GuardPoints created on a host group are applied to all members of the group. Additionally, you can choose to apply host group configuration settings (except for password, FS Agent Lock, System Lock, Registration Allowed, and Communication Enabled settings), to all hosts that are members of that host group. It is important to keep this in mind when adding hosts to a host group. For example if you create an LDT policy in a host group, and then add a Docker enabled host to that host group, the Docker host will not be protected by that GuardPoint. Similarly, if you create a Linux file system GuardPoint in a host group, and then you add a Windows host to that host group, the Windows host will not be protected by that GuardPoint.

If you create a host group and add a host to that group that does not have LDT enabled but you create LDT GuardPoints on the host group, those GuardPoints will not be propagated to that host. However if you subsequently enable LDT on that host in that host group (assuming you have a license for this feature), the LDT GuardPoint is now propagated to the LDT enabled host.

Similarly in the case of a Docker enabled host, if you later enable Docker on a host and the host contains the same Docker image as the host group GuardPoint, then that Docker GuardPoint is propagated to the Docker enabled host.

Create a host group GuardPoint

Applying a GuardPoint to a host group uses the same methodology as applying a GuardPoint to a Host.

1. Log on to the Management Console as an administrator with Host role permissions.

2. Select **Hosts > Host Groups** on the menu bar.

   The Host Groups page opens.

3. Click the target host in the **Host Name** column. The Edit Host Group window opens to the General tab for the selected host. The options are the same as for a host.

4. Select the **GuardPoints** tab.

5. Click **Guard** to add a shared GuardPoint to all members within the host group.

6. Select a host in **the Host to Browse** field to apply the GuardPoint. It is important to note that for this GuardPoint to be applicable to all hosts in the host group, they must all have the same file system type as the host selected here.

7. Select the type of policy to apply from the **Policy** drop-down menu as you would for a host.

8. Select the type of GuardPoint to apply in the **Type** drop-down menu as you would for a host.

9. In the **Path** text box, select an appropriate path.

10. Select the remaining options as appropriate. See "Understanding GuardPoints" on page 277 and "Creating and Applying a GuardPoint" on page 279 for more information.

11. Click **Ok** to create the GuardPoint and go back to the Edit Host Group page. The new GuardPoint will be listed in the table.

Creating LDT GuardPoints on a host group

The steps to create a an host group LDT GuardPoint are the same as for a creating a host group GuardPoint except that you must select an LDT policy to apply to the GuardPoint.

To create an LDT GuardPoint on a host group:

1. Create an LDT policy.

2. Set the Quality of Service, see "Understanding GuardPoints" on page 277.

3. Click **Guard** on the **GuardPoints** tab to apply an LDT policy to a directory or file and create a GuardPoint. See "Creating GuardPoints on a Host Group" on the previous page for procedures.

See the *Live Data Transformation Guide* for more information about LDT policies, QoS schedules, and LDT GuardPoints. See "Creating and Configuring VTE Policies" on page 209 for procedures to create LDT policies.

## Creating Docker GuardPoints on a host group

The steps to create a an host group LDT GuardPoint are the same as for a creating a host group GuardPoint except that select the Guard Docker tab and select a Docker host on which to apply the GuardPoint. When you a create a GuardPoint on a Docker image, for a Docker image-based GuardPoint to apply to all the Docker hosts in a Docker host group, that same Docker image must also be available on all the Docker hosts.

To create a Docker GuardPoint on a host group:

1. Create a policy.

2. Click Guard on the Guard Docker tab to apply a policy to a Docker image or container. See "Creating GuardPoints on a Host Group" on the previous page and for Docker specific information see, "Creating Docker GuardPoints" on page 283.

Using Efficient Storage for GuardPoints on a host group

When you enable Efficient Storage to protect ESG-enabled devices on a host group for the first time, or when you want to rekey the device, you must apply the Efficient Storage GuardPoints on the initial host first. The initial host is the Agent that you used to configure the ESG device.

See the *VTE Agent Installation and Configuration Guide* for more information.

Creating Cloud Object Storage GuardPoints in a Host Group

Security/All administrators can create new COS GuardPoints or edit existing GuardPoints for Host Groups.

> **Note**
> In the Edit Host window, you will see that the **Guard FS** tab has been renamed to **GuardPoints**.

The following features are NOT supported in COS GuardPoints:

- Transform Sparse Regions
- Secure Start
- Browsing to a directory to be guarded
- Host to Browse
- Auto Mount
- Efficient Storage

> **Note**
> DSM allows **exactly one** COS GuardPoint per S3 bucket.

To create a COS GuardPoint:

1. In the Host Group window, click on the host group for which you want to set GuardPoints.

2. In the Edit Host Group window, click **GuardPoints**.

3. In the GuardPoints tab, click **Guard**.

4. In the Policy dropdown menu, select an appropriate policy.

5. For type, click **Cloud Object Storage (Auto Guard or Manual Guard)**

6. In the Path field, manually enter the path for the GuardPoint.

7. Click **OK**. COS GuardPoints display on the GuardPoints tab of the Host Detail page.

Automatic and Manual GuardPoints

> **Note**
> Manual GuardPoints supported by UNIX platforms only.
>
> Information about UNIX agents applies to earlier versions of those agents, since as of v6.0, UNIX agents are EOL. Refer to the DSM v6.0 Release Notes for more information about supported agent versions.

A GuardPoint is usually applied immediately after it is configured in the Management Console; however, it can be applied later on the host system.

Generally, when you get error messages, check that only active nodes are properly guarded.

Automatic and manual GuardPoint application is set in the Edit Host window, Guard File System sub-window.

The GuardPoint type is usually set to `Directory (Auto Guard)` for file-system based directories and to `Raw or Block Device (Auto Guard)` when applying GuardPoint protection to raw or block devices. When an auto GuardPoint is applied, regardless if it is a file system directory or a raw device, the change is pushed to the host system, and the GuardPoint is applied immediately.

Use the `df` command to display `secfs` mounts (for example, GuardPoints) or `secfsd` to display the GuardPoints themselves. The `secfsd` output shows a guard type of `local` for directories configured with `Directory (Auto Guard)`.

For example:

```
# df
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/mapper/VolGroup00-LogVol00
                     40123784  11352236  26733380  30% /
/dev/sda1              101086     14590     81277  16% /boot
none                  254492         0    254492   0% /dev/shm
/opt/vormetric/DataSecurityExpert/agent/secfs/.sec
                     40123784  11352236  26733380  30%
/opt/vormetric/DataSecurityExpert/agent/secfs/.sec
/opt/apps/apps1/tmp  40123784  11352236  26733380  30% /opt/apps/apps1/tmp
/opt/apps/apps1/lib  40123784  11352236  26733380  30% /opt/apps/apps1/lib
/opt/apps/apps1/doc  40123784  11352236  26733380  30% /opt/apps/apps1/doc
# secfsd -status guard
GuardPoint          Policy                      Type    ConfigState  Status   Reason
----------          ------                      ----    --------     ------   ---
/opt/apps/apps1/tmp allowAllOps_fs              local   guarded      guarded  N/A
```

```
/opt/apps/apps1/lib  allowAllRootUsers_fs     local   guarded   guarded   N/A
/opt/apps/apps1/doc  allowAllOps-winusers1_fs local   guarded   guarded   N/A
#
```

When a manual GuardPoint is applied, regardless if it is a file system directory or a raw device, the change is pushed to the host system only. The host is aware of the GuardPoint but the host does not mount it. This is indicated in the `Type` column of the "`secfsd -status guard`" output. For example, the GuardPoint `/opt/apps/apps2/bin` has been configured with `Directory (Manual Guard)` so the guard type is set to "`manual`".

```
# secfsd -status guard

GuardPoint          Policy                    Type     ConfigState  Status       Reason
----------          ------                    ----     --------     ------       ---

/opt/apps/apps1/tmp  allowAllOps_fs           local    guarded      guarded      N/A

/opt/apps/apps1/lib  allowAllRootUsers_fs     local    guarded      guarded      N/A

/opt/apps/apps1/doc  allowAllOps-winusers1_fs local    guarded      guarded      N/A

/opt/apps/apps2/bin  HR_policy01              manual   unguarded    not guarded  Inactive
#
```

Note the `Type` value. A `Type` of `manual` indicates a manual GuardPoint. A `Type` of `local` indicates an automatic GuardPoint.

A manually applied GuardPoint retains a yellow triangle status (Pending) until the GuardPoint is applied on the host. After the GuardPoint is applied on the host, and the host communicates the change to the server, the status changes to a green ball (Normal). It returns to the yellow triangle when the GuardPoint is manually unguarded.

Use the `secfsd` command to guard and unguard `Directory (Manual Guard)` and `Raw or Block Device (Manual Guard)` GuardPoints. The `secfsd` syntax is:

```
secfsd -guard path
secfsd -unguard path
```

> **Note**
> In zone-based VTE Agent deployments, such as Solaris Zones, always specify paths relative to the global zone, never the local zone. Also, you must guard and unguard manual GuardPoints in the global zone.

For example, to manually guard and unguard a file system directory:

1. Configure a GuardPoint with the type `Directory (Manual Guard)`.

2. The host administrator with root permissions must log on to the agent system as a root user.

3. Wait until the configuration change is downloaded to the agent system.

    The status command is run until the manual GuardPoint displays.

    For example:

```
# secfsd -status guard

GuardPoint          Policy           Type     ConfigState  Status       Reason
----------          ------           ----     -----------  ------       ------

/opt/apps/etc       allowAllOps_fs   manual   unguarded    not guarded  N/A

/opt/apps/lib/dx3   allowAllOps_fs   local    guarded      guarded      N/A
#
```

4. Enable the GuardPoint:

```
# secfsd -guard /opt/apps/apps2/bin

  secfsd: Guard initiated
```

```
#
```

The GuardPoint is active and the policy is enforced.

5. Disable the GuardPoint:

```
# secfsd -unguard /opt/apps/apps2/bin
  secfsd: Unguard initiated
#
```

Selecting a GuardPoint mount type

Under random circumstances, NFS file systems can be mounted before the VTE Agent drivers are loaded. When this occurs, the VTE Agent is unable to protect GuardPoints on the file system. The **Auto Mount** feature prevents this from occurring. Select the **Auto Mount** toggle in the Edit Host window when the GuardPoint is in an automounted file system.

When applying file system protection to an automounted file system, do not apply the GuardPoint to the link-target directory. Rather, apply the GuardPoint to the full path to the directory underneath it. For example, if the automounted directory "/Auto" mounts a link-target directory named "/documents", do not set the GuardPoint to "/documents". Instead, set the GuardPoint to "/Auto/documents".

Do not configure Linux 64-bit hosts to automount directories with the "/net" option. The automounter uses the automount map associated with each mount point to locate each file system as it is accessed. The VTE Agent cannot resolve file system selections for GuardPoints, including any directories below a GuardPoint, that are configured with the "/net" option.

Displaying VTE Agent GuardPoint Status

The VTE Agent GuardPoint status can be displayed on the DSM and on the host running the VTE Agent. The agent status displayed in the Management Console shows the status as it is configured on the DSM. It is not a real-time indication of the actual status.

For actual status, the host administrator with root permissions must log on to the agent system and run VTE Agent utilities. Compare the two to ensure that the GuardPoint status on the DSM and Encryption Agent match. If the two do not match, go with what you see on the host (agent) system.

Viewing VTE Agent GuardPoint status on the DSM

1. Log on to the Management Console as a All/Security Administrator with `Host` role permissions.

2. Select **Hosts > Hosts** in the menu bar.

    The Hosts window opens.

3. Click the host in the **Host Name** column.

    The Edit Host window opens to the General tab.

4. Click **GuardPoints** tab to view GuardPoints on the host. Click the status indicator of a GuardPoint.

    The status indicator is a green circle, a yellow triangle, or a red square in the **Status** column.

    The GuardPoint Status pop-up displays.

Figure 24-5: GuardPoint Status summary
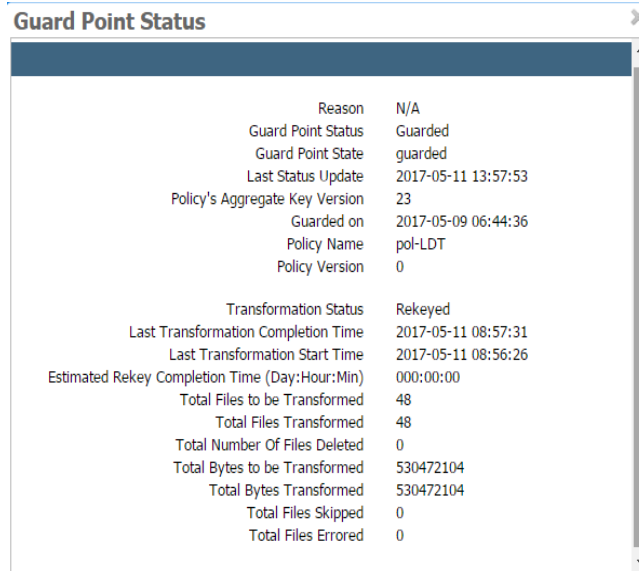
Do not click a GuardPoint with a red square status indicator. The Guard Point Status window will not display any configuration or status data when a red square is displayed.

The window is not automatically updated. You must close and reopen the window after the DSM and VTE Agent synchronize and the status indicator turns green.

5. Click the "**X**" on the Guard Point Status window to close it.

Viewing Docker GuardPoint Status

To view Docker GuardPoint status information:

1. Log on to the Management Console as a All/Security Administrator with `Host` role permissions.

2. Select **Hosts > Hosts** in the menu bar.

    The Hosts window opens.

3. Click the host in the **Host Name** column.

    The Edit Host window opens to the General tab.

4. Click Guard Docker tab to view GuardPoints on a Docker host.

Figure 24-6: Docker GuardPoints

A Docker image-based GuardPoint does not display any information in the **Status** column of the table. However, if there are containers running off that image, then the image-based GuardPoint applies to those containers and the **Docker Container** column displays the number of containers that are running.

---

5. Click the number in the **Docker Container** column, a pop-up dialog displays the Docker container GuardPoints. Click the status indicator in the **Status** column to view Docker GuardPoint Status.

Figure 24-7: Docker GuardPoint Status
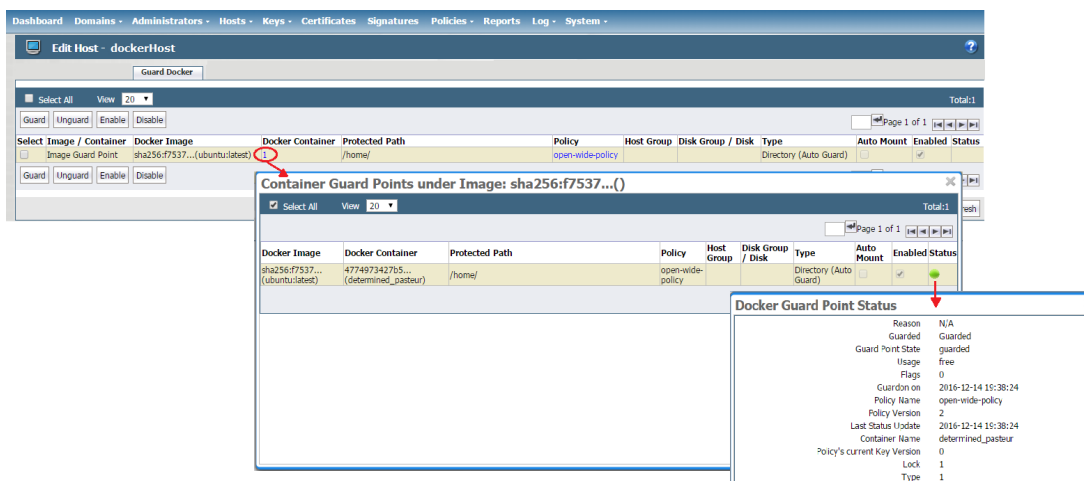
If there are no containers running off the Docker image, the **Docker Container** column displays '0' and no pop-up is available.

Configuring Windows Network Drives

Windows network drives may need user credentials and domain information for the DSM to configure GuardPoints and to push configuration changes to the VTE Agent. The Remote File Browser window enables you to automatically supply the user credentials.

Guard network mapped drives on a Windows host using the complete Universal Naming Convention (UNC) name for each file path. For example:

- `\\1.2.3.4\ShareName\dirpath`
- `\\ServerName.DomainName.com\ShareName\dirpath`
- `\\ServerName\ShareName\dirpath`

We recommend that you use the DSM IP address instead of the DNS name. GuardPoint protection is still enforced even when the DSM name is used.

> **Note**
> The **Auto Mount** check box is displayed but not selectable for Windows platforms. Auto Mount is for UNIX platforms only.

To configure a network drive:

1. Open Guard File System window.

2. Click **Browse**.

   The Remote File Browser window opens.

3. Enable **Network Drive**.

   Three text-entry boxes are displayed. They are **Username**, **Password**, and **Windows Domain**.

4. Enter the network name of the user who has access permission to the network drive in the **Username** text-entry box.

5. Enter the password for the specified user in the **Password** text-entry box.

6. Enter the domain name of the system hosting the network drive in the **Windows Domain** text-entry box.

7. Select the GuardPoint and apply the policy as you would a non-network resource.

Deleting GuardPoints

Take the following preliminary steps before deleting a GuardPoint:

- Encrypted data in a GuardPoint will still be encrypted when the GuardPoint is removed. If you are not going to reuse the GuardPoint for any reason, such as uninstalling the VTE Agent software from a host, either copy the encrypted files out of the GuardPoint so that they are saved as unencrypted files or rekey the encrypted files while the GuardPoint is still applied.

- If the GuardPoint is an LDT GuardPoint, make sure you run through the procedures described in the *VTE Live Data Transformation User Guide* to ensure that the data in those GuardPoints remains available.

- Take the GuardPoint out of service so that no user or application is accessing the directories and files in the GuardPoint. A GuardPoint is a mounted file system. Removing a GuardPoint involves unmounting the file system. File systems cannot be unmounted when in use.

- Delete all the GuardPoints and disable the locks for a host before deleting the host from the DSM. This ensures that there are no residual GuardPoints in effect on the host.

1. Log on to the Management Console as a All/Security administrator with `Host` role permissions, or a Domain and Security administrator.

2. Select **Hosts > Hosts** in the menu bar.

Figure 24-8: Hosts window

3. Select a host in the **Host Name** column of the Hosts page.

Figure 24-9: Edit Host window
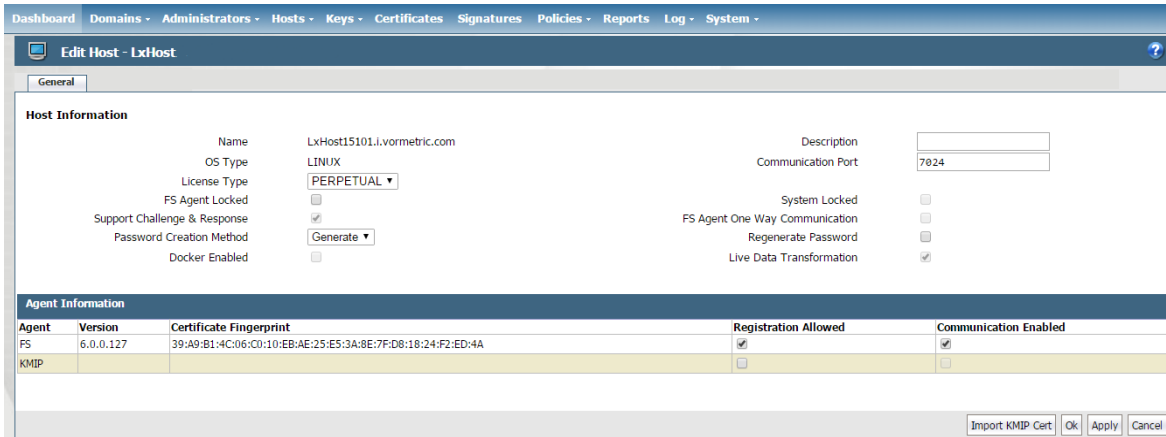
4. Select the **GuardPoints** tab.

   The GuardPoints are displayed.

5. Select the radio button in the **Select** column for the GuardPoint to be deleted.

   Only one GuardPoint at a time can be selected at a time.

6. Click **Unguard**.

7. Note the GuardPoint status:

   ° A green circle indicates an active and healthy connection to the agent system.

   ° A red square indicates that a policy has been configured but not applied on the agent system; that a GuardPoint is disabled or is in the process of being disabled; or that a communication error has occurred between the DSM and Agent host systems.

   ° A yellow triangle indicates that an attempt to delete a GuardPoint is still pending. The DSM awaits confirmation from the agent before it deletes the GuardPoint. A yellow triangle also indicates a GuardPoint that is `Directory (Manual Guard)` or `Raw or Block Device (Manual Guard)` that is not mounted on the host system.

8. Click **Refresh** to update the tab.

   After the VTE Agent acknowledges that the GuardPoint has been removed from the host, it is removed from the Management Console **GuardPoints** tab.

9. Check the mount points on the VTE Agent host to ensure that the GuardPoint has been removed.

   On UNIX, you can run the `df` command or the `secfsd -status guard` command. On Windows, you can select the Vormetric icon and **View > File System > Guardpoints**.

Chapter 25: Key Management Interoperability Protocol (KMIP)

This section describes how to enable the DSM as a key management server (key manager) supporting the Key Management Interoperability Protocol (KMIP) protocol.

Overview

The Key Management Interoperability Protocol (KMIP) is an open comprehensive protocol for communication between enterprise key management systems and encryption systems. The KMIP standard is governed by the open standards consortium known as OASIS. The DSM uses this standard protocol to simplify key management, reduce operational costs, use different key management servers, and avoid vendor lock-in.

KMIP is a wire protocol. Any device or client software that is KMIP-enabled can communicate with the DSM to manage encrypted keys. Examples of KMIP clients include storage devices, switches, tape drives, and virtual machines.

Support for KMIP client certificate validation through an external Certificate Authority (CA) is available as of DSM release v6.0.2, and later.

DSM Administrators now have the option to create Non-KMIP domains, after installing a DSM KMIP license.

Note the following restrictions:

- The DSM cannot be a client to another key manager.

- You cannot run a VAE or VKM Agent and KMIP on the DSM at the same time (unless you use an IP for the host name).

- You cannot run the KMIP client with a DSM configured in "Suite-B" mode, it must be configured in "Compatibility" mode.

KMIP Data Sheet

Table 25-1: KMIP Information for DSM

Item	Description	Comments
Protocol support	Key Management Interoperability Protocol	KMIP Version 1.9.0j, which supports IPv6 and KMIP protocol 1.0-1.4
Licensing	KMIP is a licensed feature of the DSM	The KMIP license is either enabled or disabled, and the licenses count toward the Key Agent license.
High Availability	Up to eight DSM HA nodes	All DSM HA nodes respond to KMIP client requests.

Table 25-1: KMIP Information for DSM (continued)

Item	Description	Comments
Backup/Restore	KMIP keys can be backed up/restored along with the rest of the DSM configuration using M of N shares	No changes to DSM backup/restore procedures.
DSM Upgrade	The DSM must be at v6.0 to be able to upgrade to v6.0.3. If you are upgrading from an earlier version (v5.3.0), the upgrade path is as follows: 5.3.0 > 5.3.1-patch > 6.0 > 6.0.x Contact Thales Support for more information.	When upgrading from a version of DSM that does not include KMIP functionality to one that does(5.3.1 or later), the keys already created and in use for non KMIP clients, cannot be used by KMIP clients. You must create a new domain after installing a KMIP-enabled license to register KMIP clients.
Concurrent KMIP client connections	Up to 32 concurrent clients can be used with the DSM	Performance with more concurrent connections may be slower.
Number of KMIP keys	Up to 10,000 keys	There is no known limitation to creating more than 10,000 keys but, performance may be slower.
KMIP client certificate validation	The DSM requires the KMIP certificate to authenticate the identity of the client. Obtain this from a CA.	The DSM validates the KMIP certificate when it is imported. Existing KMIP client certificates will not work in DSM v6.2.0 and subsequent versions unless they contain client identity authentication. To remedy the situation, obtain new KMIP certificates with client identity authentication and import them into the DSM. Microsoft IIS (Internet Information Services) for the Windows CA, is using the key to determine the purpose of the certificate. This key is listed as the Enhanced Key Usage or EKU in the Windows certificate dialog. The DSM needs the Client Authentication in the EKU of the KMIP client certificate to authenticate itself with the KMIP server. For Linux CA, the certificate must contain the CN, (common name). The DSM needs to match the KMIP client name to the certificate CN to have a valid KMIP client certificate that can be successfully imported and authenticated with the KMIP server.

Enable KMIP Support

To enable KMIP support on the DSM, you must install a KMIP license on the DSM, add and register KMIP clients with the DSM with the appropriate certificates, and then create a domain with KMIP enabled.

> **Note**
> You cannot run the KMIP client with a DSM configured in "**Suite-B only**" mode. The DSM must be configured in "**Compatibility**" mode.

After installing the KMIP license, you will be prompted to restart the DSM. To restart the server:

1. Start a CLI session on the initial DSM.

2. At the prompt, type the following:

```
# system
# server restart
```

3. Confirm the restart when prompted to do so.

> ⚠️ **WARNING**
>
> **After you upload a KMIP license to the DSM, domain level backup is not available for KMIP-enabled domains. However, domain level backup will continue to work for any non-KMIP domains.**

To register a KMIP client with the DSM, you must upload a signed certificate, with client authentication, which you can obtain from a third party.

> **Note**
>
> The KMIP client's externally signed certificate is separate from the DSM Web server certificate, which can also be signed by an external CA. The DSM Web server certificate is used by the browser to communicate with the DSM and can be signed by an external CA, see Chapter 8: "External Certificate Authority" on page 57 for more information about the Web Server certificate.

Requirements

- DSM appliance: Either a physical or virtual DSM appliance

- DSM **KMIP license**: Contact Support to obtain a license

- **Host:** VM or physical appliance running a KMIP client

- **Signed Certificate**: From a trusted provider with client authentication. DSM supports the following types of certificates:

  - **PEM**: Privacy Enhanced Mail, X.509v3 file which contains ASCII (Base64) encoded DER certificate

  - **DER**: Distinguished Encoding Rules, binary DER encoded certificates

    The following certificate format is not supported by KMIP:

  - **PKCS#12:** Binary format, with a protected password

KMIP Client Registration

The following high-level steps describe how to register a KMIP client with the DSM, as shown in Figure 25-1: "KMIP System Diagram" on the next page:

1. On the DSM:

   a. Install a KMIP license and restart the server as prompted.

   b. Add a KMIP client to the DSM. Navigate to **Hosts > Hosts** on the Management Console and click **Add** to add a host.

> **Note**
> Only domains created **after** a KMIP license is uploaded and enabled, will be KMIP enabled. Any domains created before the KMIP license was added will not support KMIP.
> To check whether a domain is KMIP enabled, navigate to **Domains > Manage Domains** on the top navigation bar. The table displays all the available domains on your DSM. The **KMIP Supported** column will be selected against the domain if it is available.

2. Obtain the proper certificate to register the KMIP client with the DSM—import a properly signed third-party certificate to the DSM.

3. Copy the certificate to the proper location on the KMIP client.

4. Create a KMIP-enabled domain.

Figure 25-1: KMIP System Diagram
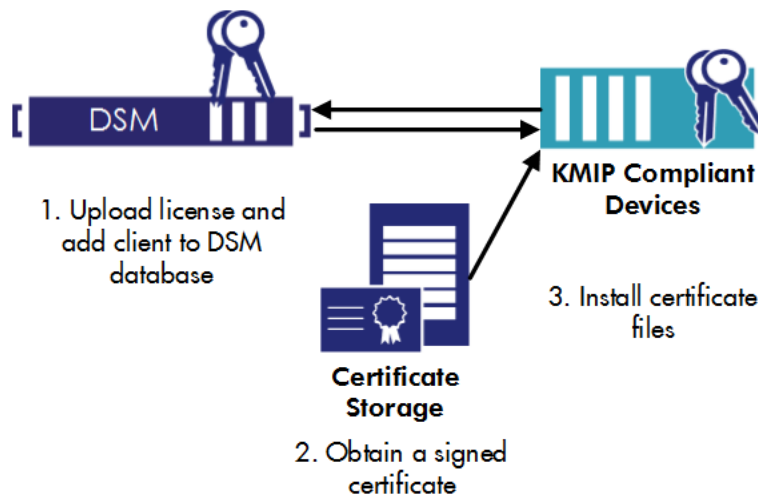
The KMIP client must be registered with the DSM to ensure secure communication.

Enable the DSM for KMIP

To enable KMIP on the DSM, upload the license that enables KMIP, then add the KMIP client name to the DSM using the Management Console. This is the first step in the process of registering the KMIP client.

Upload a KMIP license

If your license is already KMIP enabled, skip to "Add a KMIP client to the DSM" on the facing page.

1. Log on to the Management Console on the initial DSM as an All/System administrator.

2. Select **System > License** in the menu bar. The License window opens.

3. Click **Upload License File**. The Upload License File window opens.

> **Note**
> If you are in a domain, the **Upload License File** button is disabled. Click **Domain > Exit Domain**.

4. In the License File box, enter the full path of the license file or click **Browse** to locate and select the license file.

5. Click **Ok**.

6. Confirm that an **Agent Type** of KMIP is displayed on the License page.

7. Restart the DSM after installing the KMIP license.

> **Note**
> You cannot run a Key Agent and KMIP agent on the DSM at the same time (unless you use an IP for the host name).

Add a KMIP client to the DSM

> **Note**
> If you are enabling KMIP for the first time on the DSM, this is not an upgrade. Any existing domains will not be KMIP-enabled after a valid KMIP license is uploaded. Only domains created after uploading a KMIP license can be KMIP-enabled. Therefore, you must upload a valid KMIP license **before** creating a domain where KMIP clients will be used.

1. Log on to the Management Console as type: Security Administrator or type: All.

2. Switch to the domain that serves the KMIP clients. Click **Domain > Switch Domains**, and then select the domain and click **Switch to Domain**.

3. Select **Hosts > Hosts** in the menu bar. The Hosts window opens.

4. Click **Add**. The Add Host window opens.

5. In the **Host Name** field, enter the name of your KMIP client.

> **Note**
> This is the name to use when you generate the certificate.

6. Leave **Password Creation Method** as **Generate**.

> ⚠️ **WARNING**
>   **Do not select Automatically Assign to a Server.**

7. In the **Description** field, describe this client, (optional).

8. In the **License Type** list, select the option specified in your license. Options are Perpetual, Term, and Hourly.

9. For **Registration Allowed Agents**, select the **KMIP** option.

10. Select the **Communication Enabled** option.

11. Click **Ok**. The Hosts window opens.

12. Click on the hostname of the host you just added. The Edit Host window opens.

> **Note**
> The **Certificate Fingerprint** column should be empty.

13. Under the General tab, the Registration Allowed and Communication Enabled columns will have check marks indicating that they have been enabled.

14. Repeat steps 4 through 13 for each of the KMIP clients.

Establish Trust between DSM and KMIP Client

The KMIP protocol requires that you use a "mutually authenticated TLS connection" between a KMIP client and a KMIP server. In other words, the client has to cryptographically trust that it's talking to the server, and the server has to cryptographically trust that it's talking to the client. This trust is built through the use of certificates.

You must use a trusted third party CA certificate as described in "Import CA Certificate for KMIP" below, or you can create a self-signed certificate as described in "Extract the Signing Certificate from the DSM" on the facing page.

Import CA Certificate for KMIP

Use the KMIP Trusted CA Certificates page to import or update CA certificates for KMIP client certificate validation. All of the uploaded CA certificates display on this page. You may choose to use different certificates for different hosts.

> **Note**
> Once you choose to import a CA to validate KMIP clients, validation is enabled for all KMIP clients, including existing ones. You will have to re-import valid certificates for those KMIP clients. These certificates must have client authentication to be valid.

You can import the following types of certificates:

- Self-signed certificate

- Single CA certificate, provided the certificate issuer exists on the DSM

- Certificate trust chain, if required by the CA, provided the certificate issuer exists on the DSM

> **Note**
> The KMIP client host name **must** match the Common Name (CN) in the CA certificate, otherwise certificate import will fail.

Import a CA certificate

1. Log on to a DSM as a system/all administrator.

2. Navigate to the **System > KMIP Trusted CA Certificates** page.

3. Click **Browse/Choose** File to select a certificate file to import.

4. Click **Import/Update Certificate** to import the file.

5. A warning message displays, click **OK** to continue with the import/update, or **Cancel** to cancel the operation.

    The CA certificate or trust chain is uploaded and the details displayed in the table.

After uploading the CA file for the KMIP validation, upload the related certificate for each of the KMIP clients. If you have not already added a KMIP client, see "Add a KMIP client to the DSM" on the previous page. To upload the client certificate:

1. Log on to the Management Console.

2. Switch to the domain that serves the KMIP clients. Click **Domains > Switch Domains**, select the KMIP domain that you want, and click **Switch Domains**.

3. Navigate to **Hosts > Hosts** and select the KMIP client where the certificate file, to be uploaded, is located.

4. On the Edit Host page, click **Import KMIP Cert**.

5. On the Import KMIP Client Certificate page, click **Browse** to navigate to the location where the client certificate has been saved, select the file and click **Open**.

6. Click **Ok** to upload the file to the KMIP client.

7. Confirm that a fingerprint was entered in the **Certificate Fingerprint** column of the KMIP agent row.

Now that you have imported a CA certificate, validation has been turned on for all KMIP clients, including any existing KMIP clients you may have had. You must re-import valid certificates for each of those clients. The DSM now requires KMIP certificates to contain client authentication. If your client certificates do not contain client authentication, consult your third-party certificate authority for them.

Re-import a valid certificate

1. Navigate to **Hosts > Hosts** and click the KMIP client that needs to re-import a valid certificate.

2. In the **Agent Information** table on the Edit Host page, click **Import KMIP Cert** and import a valid certificate.

3. Confirm that a fingerprint was entered in the **Certificate Fingerprint** column of the KMIP agent row.

Extract the Signing Certificate from the DSM

This section describes how to create and import certificates so that the DSM trusts the client and the client trusts the DSM.

First, acquire the CA certificate used internally by the DSM. This certificate is used to establish trust between the KMIP client and the DSM—that the KMIP client is communicating with the DSM.

The name of the file containing the CA certificate has a special format. *<IP address>_CA.pem*. For example, if the DSM IP address is 1.2.3.4, then the file name must be, `1.2.3.4_CA.pem`. (This example is used throughout the document.)

Acquiring the certificate from Windows involves using a browser to connect to the KMIP proxy server. Acquiring the certificate from Linux involves using the `openssl` command to communicate with the KMIP proxy port.

> **Note**
> You **must** use KMIP port 5696 to extract the correct root certificate from the DSM, in both Windows and Linux.

Extracting the CA Certificate on Windows with a Web Browser

1. Using your web browser, navigate to the Management Console.

2. Each browser manages certificates differently. Refer to the browser documentation for the specific method for managing certificates.

3. Access the certificate manager as the browser documentation instructs.

> **Note**
> The following instructions are somewhat general. Because all browsers manager certificates differently, use the following information as a guide only.

4. Click the **Certification Path** tab.

5. Select the top most certificate – it starts with "CG CA (S) on …"

6. Click **View Certificate**.

7. Click the **Details** tab and click **Copy to File**.

8. Select **Next**, then select "Base-64 encoded X.509" and click **Next** again

9. Supply a file name for the certificate.

10. Click **Next** and **Finish**.

11. Locate and open the file you just saved with your favorite text editor. It should start with
    `"-----BEGIN CERTIFICATE-----"` and end with `"-----END CERTIFICATE-----"`.
    Copy and paste the contents of this file to the file "1.2.3.4_CA.pem"

Extracting the CA Certificate on Linux with OpenSSL

Acquire the same certificate through a different procedure on Linux using the OpenSSL program.

1. Locate a Linux machine with the `openssl` utility installed.

2. Run the following command:

```
# openssl s_client -connect <HostName>:5696 -showcerts
```

Lots of output scrolls past. The second block of base-64 encoded text (between
`"-----BEGIN CERTIFICATE-----"` and `"-----END CERTIFICATE -----"`)
is the certificate of interest.

3. Copy this text, **including** the text for Begin Certificate and End Certificate, into the file `1.2.3.4_CA.pem`.

Testing Your CA Certificate

A quick test of your certificate can be performed using the `openssl` command. Run and execute the following command at the prompt:

```
# openssl s_client -connect kmip-interop.vormetric.com:5696 -showcerts -CAfile 1.2.3.4_CA.pem
```

A lot of output will flow past; look for the very last line:

```
Verify return code: 0 (ok)
```

If you see this, your certificate file is valid.

Creating a Certificate for the KMIP Client

Once the certificate identifying the DSM has been created, you need to create a certificate to identify the KMIP client. You can use openssl on a Linux server to do this, and create a self-signed certificate.

1. Create a 2048-bit RSA key:

```
$ openssl genrsa 2048 > client_private.pem
```

2. Create a self-signed certificate using that key:

```
$ openssl req -new -x509 -key client_private.pem -out client.pem -days 365
```

3. Follow the prompts. When prompted for the "Common Name", enter the host name that you supplied when you added the KMIP client to the DSM.

> **Note**
> This name must be a perfect match.

You now have two files; `client_private.pem` and `client.pem`. The first contains the key and the second contains the certificate.

Register the KMIP client using Import KMIP Cert

The self-signed certificate created for the KMIP client must be uploaded to the DSM:

1. Obtain the signed CSR as described above, and copy it to any location that the DSM can access.

2. Switch to the domain that serves the KMIP clients. Click **Domains > Switch Domains**, select the KMIP domain that you want, and click **Switch Domains**.

3. Navigate to **Hosts > Hosts**, click the KMIP client you want to register with the DSM. The Edit Hosts page opens.

4. Click **Import KMIP Cert** and browse to the location of the signed CSR. In our example;

   `client.pem`

5. Click **Open**.

6. Click **OK** on the Import KMIP Client Certificate page.

   The Import KMIP Client Certificate page closes displaying the Edit Host page.

7. Confirm that a fingerprint was entered in the **Certificate Fingerprint** column of the **KMIP agent** row.

Create a KMIP-enabled domain

In order to create a KMIP domain, KMIP must be explicitly enabled for that domain when it is created. The steps are as follows:

1. Log on to the DSM as an All/System Administrator.

2. Navigate to the **Domain > Manage Domains** page.

3. Click **Add**, the General tab on the Add Domain page displays.

4. Enter details for the domain, the **Name** field is mandatory.

5. Select the **Enable KMIP** option to create a KMIP domain.

6. Click **Apply** to save the domain information.

7. Click the **Assign Admin** tab to assign an administrator.

8. (Optional) Click the **License** tab to allocate licenses or license hours per agent on this domain.

9. Click **Ok**. The Domains window opens with the name and description of the new domain and a check mark is displayed in the KMIP Supported column.

   > **Note**
   > You can migrate a standard domain to a KMIP domain simply by editing the domain and selecting the Enable KMIP option.

Backing up a KMIP-enabled domain

Back up a KMIP-enabled domain exactly as you would a standard domain. The only difference is that a backup containing a KMIP-enabled domain must be restored with the same name. You cannot restore it with a different name, unlike a standard domain.

Managing KMIP CA Certificates

Certificates may sometimes need to be deleted, for example, if they compromised or are corrupted, or you may want to turn off KMIP client certificate validation.

Turn off validation

To turn off KMIP client certificate validation, you need to delete all the CA certificates in the table on the KMIP Trusted CA Certificates page. When the table is empty, validation is turned off. To turn it back on, see "Import CA Certificate for KMIP" on page 303.

Delete a CA certificate

Deleting a CA or Intermediate CA certificate will result in authentication failure for any KMIP clients that use client certificates issued by that CA or Intermediate CA certificate. All the affected KMIP clients will have to re-import valid certificates.

1. Log on to the Management Console and Navigate to **System > KMIP Trusted CA Certificates** page.

2. Select the CA certificate file to delete from the table.

3. Click **Delete**.

4. A warning dialog displays, click **OK** to continue with the deletion or **Cancel** to cancel the operation.

KMIP clients using client certificates issued by the CA being deleted will no longer be able to communicate with the KMIP server and authentication will fail. All the affected KMIP clients will have to re-import valid certificates.

Export a CA Certificate

> **Note**
> You may want to back up a certificate by exporting it, before deleting it.

To export a certificate:

1. Log on to the Management Console and Navigate to **System > KMIP Trusted CA Certificates** page.

2. Select the CA file to export from the table.

3. Click **Export**.

   The certificates are exported in a text format to the default location on your computer.

Managing KMIP Objects

Once a KMIP client has been successfully registered with the DSM, you can manage KMIP objects and control their use.

Viewing KMIP objects

You must be inside the KMIP domain to view KMIP objects. Switch to the KMIP domain and select **Keys > KMIP Objects**. The KMIP Objects page displays.

The top panel of the page toggles between **Show Search** and **Hide Search**.

1. Search using any one of the following fields:
   - **UUID**: the Unique Identifier of the object.
   - **Creation (From)**: This field in conjunction with the **Creation Time (To)** field are use to search for objects created within a specified date range. Set the start date for the range search.
   - **Creation (To)**: Set the end date for the range search.
   - **Type**: Select the type of object to search for from the drop-down list.

      ° **State**: Select the state of the object to search for from the drop-down list.

2. Click **Go** to search based on the selected filters.

   The panel under **Search** contains a table that displays the following information about objects stored on the KMIP server:

   ° **Name**: The name attribute of the object if any was assigned when the object was created.

   ° **Unique Identifier**: The universally unique identifier of that object. Clicking this unique identifier displays the KMIP Object Attributes page with the attributes for that object.

   ° **State**: The cryptographic state of the object

   ° **Object Type**: The type of the managed object

   ° **Creation Time**: The timestamp of when the object was created.

KMIP High Availability

You can configure a DSM high availability (HA) cluster with KMIP. See Chapter 13: "High Availability (HA)" on page 108 for information about how create a high availability cluster.

When the HA replication is done, you need to restart the Ha node from the CLI to bring up the KMIP server. To restart the server:

1. Start a CLI session on an HA node DSM.

2. At the prompt type,

```
0001:dsm$ system
0002:system$ server restart
Do you want to restart the server software ? (y/n):y
Restarting now...
Stopping Security Server...done.
Stopping the data store...done.
Starting Security Server...done.
SUCCESS: The security server software is restarted.
0003:system$
```

KMIP clients can make KMIP key read requests to another node. In order to enable read requests to the other node, you must first make a key retrieval request to the initial DSM once you have configured HA. After that, read requests to the other HA nodes are enabled.

Using Efficient Storage with a KMIP Server

Refer to the chapter 'Using VTE with Efficient Storage Devices" in the VTE Agent Installation and Configuration guide for information on the Efficient Storage device.

To protect data with Efficient Storage:

1. Create a KMIP-enabled domain. See "Create a KMIP-enabled domain" on page 379.

2. Create a key for Efficient Storage:

   ° Algorithm: **AES256**

   ° Encryption Mode: **XTS**

   ° KMIP Accessible: **Selected**

3. Create a standard policy. See Chapter 22: "Policies" on page 207.

4. Set up and initialize the Efficient Storage device. See the chapter 'Using VTE with Efficient Storage Devices" in the VTE Agent Installation and Configuration guide.

5. Create and apply Efficient Storage GuardPoints.

See "Using VTE with Efficient Storage Devices" in the *VTE Agent Installation and Configuration Guide* for more information.

Chapter 26: Security Administrator Preferences & Logs

Viewing Preferences

Although most preferences for viewing the various windows and panels on the DSM Management Console are set by the DSM System Administrator, as a DSM Security Administrator you can still set certain viewing preferences within the domains you are authorized to access. From the **System > General Preference** window, you can set parameters for the following pages:

- Domain Page

- Administrator Page

- Host Page

- Policy Page

- Key/Certificate Page

- Signature Page

- Log Page

You can also set the Management Console Timeout limit for your sessions.

From the **System > Log Preferences** window, you can set the following parameters for:

- Server
  - Logging Settings such as Logging Level (DEBUG, INFO, WARN, ERROR, FATAL)
    - Log Upload DB Retry (secs)
    - Log Buffer Size (messages)
    - Log Buffer Flush Time (secs)
  - Communication Settings
    - Update Host Frequency (secs)
    - Default Host Communication Port

- Agent Logs—the available tabs will depend on the agents for which you have a licenses installed.

> **Note**
> We recommend turning on **Log to File** or **Log to Syslog** instead of **Upload to Server** for INFO and DEBUG levels. For general day-to-day operation, we recommend enabling and setting only ERROR Level (so that only ERROR, WARNING, and FATAL log entries are received). Setting Upload to Server to INFO or DEBUG level for policy evaluation can affect DSM performance.

Viewing Logs

The entries displayed in the Message Log depend on the DSM administrator type (System, Domain, Security, All), the domain in which the administrator is working, and, for Security Administrators, the administrator role (Audit, Key, Policy, Host, Challenge & Response, Client Identity).

Security Administrators can see log entries for the management of Security Administrators by Domain Administrator, GuardPoint application, and policy evaluation.

Log entries are displayed in the Management Console based on the current administrator type and the domain in which the administrator is working. The combined list of this log information is available in the `server.log` file on the DSM.

Chapter 27: DSM Command Line Interface

CLI administrators are system users with login accounts. That is, they are entered in `/etc/passwd` and they have directories under `/home`. CLI administrators perform the tasks to set up and operate the DSM and any tasks that need to be performed from the CLI. CLI administrators exist only in the CLI and they cannot access the UI.

Table 27-1: Differences between CLI administrators and Management Console administrators

CLI Administrators	Management Console (UI) Administrators
CLI administrators are created and administered in the CLI only.	Management Console administrators are created and administered in the Management Console only.
CLI administrators cannot log on to the Management Console.	Management Console administrators cannot log on to the CLI.
CLI administrators are not included in the backup.	Management Console administrators are included in a DSM backup.
The CLI administrator exists only on the appliance or system on which they were created.	A Management Console administrator can open a Web browser session on HA nodes using the same password.

The password requirements for both CLI and DSM administrators are set by the password policy in the Management Console.

The DSM Command Line Interface (CLI) enables you to configure the DSM (represented in the code as a Security Server) network and do other system-level tasks.

Procedures for the DSM are divided between the Management Console and the CLI. This is usually because the procedures require a mix of network, DSM database, or system access, such as for upgrades. The Management Console Web interface (GUI) is used to upload DSM application upgrade images and DSM OS upgrade images, because the DSM CLI does not support file uploading.

The Management Console cannot be used to restart the DSM and the CLI cannot be used to download files across the net.

A mixture of DSM CLI and Management Console activities is required for some procedures to reduce the potential for software hacks or other misuse.

Overview

CLI administrators are system users with login accounts. That is, they are entered in `/etc/passwd` and they have directories under `/home`. CLI administrators do the tasks to set up and operate the DSM installation and any tasks that need to be done from the CLI. DSM administrators only access the Management Console.

Table 27-2: Differences between CLI administrators and Management Console administrators

DSM CLI Administrators	Management Console Administrators
CLI administrators are created and administered in the CLI only.	Administrators are created and administered in the Management Console only.
CLI administrators cannot log on to the DSM Management Console.	Management Console administrators cannot log on to the CLI.
CLI administrators are not included in the backup.	Included in a DSM backup.
The CLI administrator exists only on the appliance or system on which they were created.	A DSM administrator can open a Web browser session on all of the HA nodes.

The password requirements for both CLI and Management Console administrators are set by the password policy in the Management Console.

DSM CLI Navigation

These are the CLI command categories:

- `network`
- `system`
- `HSM` (Hardware Security Module)
- `maintenance`
- `HA` (High Availability)
- `IPMI` (Intelligent Platform Management Interface)
- `user`

As a DSM CLI administrator, log on to the CLI, then enter a command category by typing the category name at the command line prompt. For example, type `system` to enter the `system` category. While in the category, you can execute the commands for that category.

Enter the entire category name, command, or argument, or enter just enough characters to uniquely identify the category, command, or argument. For example, both of these commands achieve the same result:

```
ip address add 10.3.5.100/16 dev eth1
i a a 10.3.5.100/16 d eth1
```

You can use the <Tab> key to complete a category, command, or argument. Enter enough characters to uniquely identify a category, command, or argument, and then press the <Tab> key. The CLI will complete it for you.

For example:

At the top level, enter `m` and press <Tab>, the CLI expands it to `maintenance`.

Inside the `maintenance` category, you can enter `di`<Tab> and it expands to `diag`. Type `d`<Tab> and it expands to `diskusage`. Note that you must enter `di` because there are other `d` commands in the maintenance category, like `date` and `delver`.

Other supported CLI navigation methods are:

- Enter a question mark (`?`) to display the next command or argument that is expected. Think of it is as a shorthand form of help.

- Enter "up" to return to the top level so that you can enter another category. You can enter another category only from the top level.

- Enter "exit" at any time to end the current CLI session.

Network Category Commands

The network category is used to set, modify or delete IP addresses on the system, and set up DNS servers. DHCP is supported and is enabled by default on a fresh installation. DHCP must be enabled on an upgraded appliance.

The network category supports the following commands:

Table 27-3: Network category commands

ip	Configures the network interface.
dns	Sets one or more DNS servers for the appliance.
host	Configures an IP address to a host name.
ssh	Enables Secure Shell (SSH) port
ping	Pings an IP address, host name, or FQDN.
traceroute	Traces route to IP address or host name.
rping	Sends an ARP (Address Resolution Protocol) request to a neighbor host.
arp	Displays the system ARP cache.
checkport	Checks local and remote TCP port status.
nslookup	Queries DNS to domain name to IP address mapping.

ip

The ip command configures the network interface.

It includes the following elements.

Table 27-4: Network category `ip` command elements

address	Adds, deletes, or initializes the IP address a network interface.
route	Configures network routing.
link	Sets the physical components of the network interface, such as connection speed, mode, set bond mode for bonded NICs, and MTU.
dhcp	Manages Dynamic Host Configuration Protocol (DHCP) settings.

ip address

Use the ip address command to add, initialize (set to default), delete, or show different addresses on the interface, or to assign an IP address to a bonded NIC. The DSM also supports IPv6 addresses. Examples are included below.

Syntax

```
ip address {init|add|delete} ip_address dev {eth0|eth1|bond0} [label {diag|this}]
ip address {show|flush} {eth0|eth1|bond0} [label {diag|this}]
```

The `ip address` command takes the following arguments:

Table 27-5: Network category ip address command

add	Adds an IP address to the specified interface.
delete	Deletes an IP address from an interface.
show	Displays the current addresses on the interfaces.
flush	Removes the IP addresses on the specified interface.

Example 1

The following example assigns an IP address to the bonded NIC interface, bond0:

```
0000: dsm$ network
0001:network$ ip address init 1.2.3.4/16 dev bond0
```

for IPv6:

```
0001:network$ ip address init fa01::3:15:130/64 dev bond0
```

Example 2

The following example changes the current `eth0` IP address:

```
0001:dsm$ network
0002:network$ ip address init 1.2.3.4/16 dev eth0
```

for IPv6

```
0002:network$ ip address init fa01::3:15:130/64 dev eth0
```

Example 3

The following example deletes the IP address for the `eth1` network interface and assigns the IP address to `bond0`:

```
0003:network$ ip address delete 1.2.34/16 dev eth1 label diag
WARNING: Changing network ip address requires server software to be restarted.
Continue? (yes|no)[no]:yes
SUCCESS: delete ip address. Please restart server software to pick up the changes.
0004:network$ ip address show
   Device         Prefix      Broadcast        Label
   eth0    192.168.10.1/16        192.168.255.255 diag
   Show ip address SUCCESS
0005:network$ ip address add 1.2.3.4/16 dev bond0 label diag
WARNING: Changing network ip address requires server software to be restarted.
Continue? (yes|no)[no]:yes
SUCCESS: add ip address. Please restart server software to pick up the changes.
```

To view the IP address changes, use the `show` command:

```
0006:network$ ip address show
   Device Prefix      Broadcast         Label
   eth0 192.168.10.1/16 192.168.255.255    diag
   bond0 1.2.3.4/16        1.2.255.255     diag
   Show ip address SUCCESS
```

ip link

The `ip link` command establishes how the various interfaces connect to the other nodes in the network. The `ip link` command is used to specify the bandwidth of the `eth0` and `eth1` interfaces and sets the Maximum Transmission Unit (MTU). It is also used to set the mode for the bonded NIC interface `bond0`. See the *DSM Installation & Configuration Guide* for more information about bonded NICs. See the following examples for the different modes that can be set for the `bond0` interface.

Syntax

```
ip link set (eth0|eth1|bond0) [mtu {100..1500}] [{up|down}] [mode {0..6}]|[speed
{auto|10mb_half|10mb_full|100mb_half|100mb_full|1000mb_half|1000mb_full}]

ip link show [eth0|eth1|bond0]
```

> **Note**
> When an IPv6 configured DSM Ethernet interface link is brought down using the command,
> `ip link set {eth0|eth1|bond0} down`
> the IPv6 address is lost. You will need to reconfigure the IPv6 address for that Ethernet interface when you bring it back up

The `ip link` command can take the following arguments:

Table 27-6: Network category ip link command arguments

`eth0`	Network interface card 1.
`eth1`	Network interface card 2.
`bond0`	Bonded NIC device type interface.
`mtu`	Sets the Maximum Transmission Unit value. The default MTU is 1500.
`pdelay`	Raises the delay value. Options are 0-2147483600 (ms).
`set`	Enables the parameter settings below for the ip link command.
`show`	Displays information about the IP link connections.
`speed`	Sets the link speed of the interface.
`xmithashpolicy`	Transmits the hash policy.

> **Note**
> Use auto detect to set the data rate of all interfaces and set the MTU value to the default, 1500.

Table 27-7: Bonding driver modes

Mode	Name	Description	Load-balancing	Fault tolerance
0	balance-rr	Round-robin policy. Transmit packets in sequential order from the first available through the last. This is the default mode for the bonded NICs.	Yes	Yes

Table 27-7: Bonding driver modes (continued)

Mode	Name	Description	Load-balancing	Fault tolerance
1	active-backup	Active-backup policy: Only one slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch.	No	Yes
2	balance-xor	XOR policy: Transmit based on the selected transmit hash policy. The default policy is a simple [(source MAC address XOR'd with destination MAC address) modulo slave count].	Yes	Yes
3	broadcast	Broadcast policy: transmits everything on all slave interfaces.	No	Yes
4	802.3ad	IEEE 802.3ad Dynamic link aggregation. Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification.	Yes	Yes
5	balance-tlb	Adaptive transmit load balancing: channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.	Yes	Yes
6	balance-alb	Adaptive load balancing: includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the slaves in the bond such that different peers use different hardware addresses for the server.	Yes	Yes

Example 1

The following example configures the `eth1` interface to operate at 100 Mb/s, in full-duplex mode, and then activates the interface so that it is network accessible:

```
0002:network$ ip link set eth1 speed 100mb_full
ip link speed SUCCESS
```

```
0003:
```

Example 2

The following example sets the `bond0` interface mode to mode 2:

```
0003:network$ ip link set bond0 mode 2
```

xmithashpolicy

Bonding, or link aggregation, is the process of combining several network interfaces (NICs) into a single link. This allows for benefits such as high availability, load balancing, maximum throughput, or a combination of these benefits.

When using the Link Aggregation Control Protocol (LACP) to bond/link network interfaces together, you have to create a Transmit Hash Policy. This is a policy that uses upper layer protocol information, when available, to generate the hash. This allows for traffic to a particular network peer to span multiple child links.

The xmithashpolicy command transmits the hash policy.

Table 27-8: Hash Policy Options

Policy	Definitions
layer2only	Layer 2
layer2+3	Layer 2 and 3
layer3+4	Layer 3 and 4
encap2+3	Encapsulation 2 and 3
encap3+4	Encapsulation 3 and 4

Syntax

```
0001: network$ ip link set bond0 [hash_policy]
```

Example

```
0001: network$ ip link set bond0 encap2+3
```

ip link show

The `ip link show` command displays the physical link settings on the system. Also use it to verify any changes made to the physical link settings:

```
0003:network$ ip link show
Device   State   MTU     Mediatype       Speed
eth0     UP      1500    copper          auto
eth1     UP      1500    copper          auto
Device   State   MTU     Mode
bond0    UP      1500    0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
Bonding Mode: load balancing (round-robin)
MII Status: down
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
SUCCESS: show ip link
0004:network$
```

ip route

Use the `ip route` command to set up IP routes. If the `eth01` and `eth1` interfaces are set on the same subnet, you do not have to include a netmask. If they are on separate subnets, include the netmask for the other subnet.

> **Note**
> Configure a default route connection outside of the subnet.

The `ip route` command uses the following arguments:

Table 27-9: Network category ip route command arguments

add	Adds a static route.
delete	Deletes a static route.
get	Shows information for a specific route.
replace	Changes the table, gateway and/or source of an existing IP route.
show	Displays all the currently configured route

Syntax

```
ip roudd|delete|replace} [ip|default]
table main.table [dev {eth0|eth1|bond0} | via {ip}] src ip
ip route get ip
ip route show
```

Example 1

The following example adds a gateway to the `eth1` interface, which has 1.2.3.4 as it's IP address and then displays the results:

```
0044:network$ ip route add default table main.table dev eth1 via 1.2.6.7
```

> **Note**
> Ignore the separation of routes into unique tables. All routes are considered members of the main routing table, as reflected in the Management Console. Separate routing tables have been deprecated.

```
ip route SUCCESS
0045:network$ ip route get 1.2.3.4
local 1.2.3.4 dev lo src 1.2.3.4
cache <local>
ip route SUCCESS
0046:network$
```

A default route specifies the gateway to which IP packets are sent when the local routing table is unable to resolve a destination. Always configure a default route. The following example configures a default route on the `eth0` interface:

```
ip route add default table main.table dev eth1 via 1.2.6.7
```

The default interface is `eth0`.

Example 2

The following example adds a default gateway to the bond0 interface:

```
0005:network$ ip route add default table main.table dev bond0 via 1.2.6.7
```

Example 3

The `ip route show` command displays the IP routes that have been assigned to the system. Use the `ip route show` command to verify the changes you made to the IP route tables:

```
0020:network$ ip route show
Main routing table
1.2.0.0/16 dev eth1 proto kernel scope link src 1.2.3.4
192.168.0.0/16 dev eth0 proto kernel scope link src 192.168.10.1
ip route show SUCCESS
```

The following example displays the IP routes that have been assigned with the bond0 interface configured:

```
0000:dsm$ network
0001:network$ ip route show
Main routing table
default via 1.2.6.7 dev bond0
1.2.0.0/16 dev bond0 proto kernel scope link src 1.2.3.4
6.2.0.0/16 dev bond0 scope link metric 1004
6.2.0.0/16 dev bond0 scope link metric 1005
192.168.0.0/16 dev eth0 proto kernel scope link src 192.168.10.1
ip route show SUCCESS
0002:network$
```

ip dhcp

Use the `ip dhcp` command to manage DHCP settings. Note that when DHCP addressing is released, all network configuration is removed, you will have to reconfigure the gateway and DNS information. The current DSM DHCP implementation does not support IPv6 addresses.

Syntax

```
ip dhcp {enable|release|renew|show} {eth0|eth1|bond0} version {4|6}
```

Table 27-10: Network category ip dhcp command arguments

`enable`	Enables DHCP IP address leasing for a specified interface.
`release`	Releases DHCP IP address leasing for a specified interface.
`renew`	Renews DHCP IP address leasing for a specified interface.
`show`	Displays DHCP IP address leasing status for all interfaces, there are no additional parameters for this command.

Example 1

The following example enables DHCP on the `bond0` interface for an IPv4 address:

```
0004:network$ ip dhcp enable bond0 version 4
WARNING: Changing network ip address may disconnect your session and will require the
server software to be restarted.
Continue? (yes|no)[no]:yes
DHCP operations may take some time, please wait....
SUCCESS: Please restart server software to pick up the changes.
0005:network$
```

Example 2

The following example releases DHCP IP address leasing for the `eth0` interface for an IPv4 address:

```
0000:dsm$ network

0001:network$ ip dhcp release eth0 version 4

WARNING: Changing network ip address may disconnect your session and will require the
server software to be restarted.

Continue? (yes|no)[no]:yes

DHCP operations may take some time, please wait....

SUCCESS: Please restart server software to pick up the changes.

0002:network$
```

Example 3

The following example renews DHCP IP leasing for the `eth0` interface:

```
0008:network$ ip dhcp renew eth0 version 4

WARNING: Changing network ip address may disconnect your session and will require the
server software to be restarted.

Continue? (yes|no)[no]:yes

DHCP operations may take some time, please wait....

SUCCESS: Please restart server software to pick up the changes.
```

Example 4

The following example displays the DHCP IP leasing status for all interfaces. In this example, the `bond0` interface has been enabled so the 'Active' column and the `eth0` and `eth1` interfaces are not in use.

```
0000:dsm$ network

0001:network$ ip dhcp show

Device Active? DCHP? DHCP Addr DHCPv6? DHCPv6 Addr

------ ------- ----- --------- ------- --------------

eth0

eth1

bond0    Y      Y 1.3.2.4/16

SUCCESS

0004:network$
```

dns

The DNS command sets the DNS domain servers that the DSM will use for HA communication. This is equivalent to editing the `/etc/resolv.conf` file. You can configure just the DNS server name, just the DNS server IP addresses, or both the DNS server name and IP addresses.

Syntax

```
dns [search domainname] [dns1 ip] [dns2 ip] [dns3 ip] [offwithdhcp][onwithdhcp]
[switchhosts][clear][show]
```

The `dns` command includes the following elements:

Table 27-11: Network category DNS elements

`clear`	Removes all of the DNS settings.

Table 27-11: Network category DNS elements (continued)

`dns1`	Specifies settings for domain name server 1.
`dns2`	Specifies settings for domain name server 2.
`dns3`	Specifies settings for domain name server 3.
`offwithdhcp`	Disallows static DNS settings with DHCP.
`onwithdhcp`	Allows static DNS settings with DHCP.
`search`	Defines the DNS domain name to search.
`show`	Shows all of the currently configured Domain Name Servers and the order of DNS resolution.
`switchhosts`	Switch hosts source from DNS server to local `/etc/hosts` file.

DNS Search

Example

The following example sets the domain to `i.vormetric.com` and the `dns1` lookup IP address to `192.168.2.254`.

```
0002:network$ dns search i.vormetric.com dns1 192.168.2.254
DNS SUCCESS
```

DNS Clear

To remove all the DNS settings use the `dns clear` command:

```
0003:network$ dns clear
DNS SUCCESS
```

DNS1 | DNS2 | DNS3

The following example sets the DNS server lookup address for dns2.

```
0003:network$ dns dns2 192.168.110.224
```

DNS OnwithDHCP

The following example allows a static IP address for a DNS server with DHCP.

```
0012:network$ dns onwithdhcp
Enabled static DNS with DHCP.
```

DNS OffwithDHCP

The following example does not allow a static IP address for a DNS server with DHCP.

```
0013:network$ dns offwithdhcp
Disabled static DNS with DHCP. DHCP DNS only.
```

DNS Switchhosts

Switch the order of name resolution to improve the response time for network operations. If DNS is listed first, it searches the DNS server first. If files is listed first, it searches the `/etc/hosts` file.

```
0015:network$ dns switchhosts
Switched hosts sources in nsswitch.conf.
0016:network$ dns show
nameserver 10.3.110.104
nameserver 10.3.110.224
hosts:       files dns
0017:network$ dns switchhosts
Switched hosts sources in nsswitch.conf.
0018:network$ dns show
nameserver 10.3.110.104
nameserver 10.3.110.224
hosts:       dns files
```

host

The `host` DSM CLI command is used to add and remove static IP addresses to and from the `/etc/hosts` file of an appliance-based DSM. By default, only hosts with resolvable host names or FQDNs can be configured in the DSM database. The `host` DSM CLI command allows the DSM to communicate with other DSM and hosts without using DNS.

This feature is provided on appliance-based DSMs only. Administrators on appliance-based DSMs cannot edit system files directly. Administrators on software-only DSMs can edit system files directly and so do not need this feature.

The name of a host in the Management Console and the host's network identity are one and the same.

To name a host with a valid network host name without DNS so that the network host name resolves to a valid IP address, run the `host` command on an appliance-based DSM or edit the `/etc/hosts` file on a software-only DSM.

Check that the network host names and FQDNs resolve successfully on the DSM. You can use the `host` command on an appliance-based DSM to link host names to IP addresses or FQDNs.

Host names cannot contain spaces and IP addresses must be in the standard `xxx.xxx.xxx.xxx` format. You cannot assign multiple host names to an IP address like you can if you were editing `/etc/hosts` directly. Also, if an IP address is already assigned multiple names, the `host show` command will display the first name only and the DSM uses the first entry only.

For example, `/etc/hosts` can contain:

```
1.3.5.7 deptsys deptsys.domain.com
```

but `host show` will display:

```
name=deptsys ip=1.3.5.7
```

Syntax

```
host add name ip
host delete name
host show
```

where, name is the host name of an HA node DSM or agent system, and **ip** is the IP address to use to contact that node or agent system.

The `host` command has the following options:

Table 27-12: Network category host command options

`add`	Inserts a host:IP pair in `/etc/hosts`.
`delete`	Removes a host:IP pair from `/etc/hosts`.
`show`	Shows the `/etc/hosts` file except for blank lines, comment lines, and the `localhost` entry. Displayed entries are not sorted.

Example

The following example adds a **host**:**IP** pair to the `/etc/hosts` file and then displays all the configured **host**:**IP** pairs.

```
0029:network$ host add deptsys 1.3.5.9
SUCCESS: add host
0030:network$ host show
name=vmlinux10 ip=1.3.5.10
name=vmlinux11 ip=1.3.5.11
name=vmlinux12 ip=1.3.5.12
name=vmlinux13 ip=1.3.5.13
name=vmlinux14 ip=1.3.5.14
name=deptsys ip=1.3.5.9
SUCCESS: show host
0031:network$
```

The following example deletes a host from the `/etc/hosts` file:

```
0031:network$ host delete deptsys
SUCCESS: delete host
0032:network$
```

ssh

The `ssh` command enables the secure shell (SSH) port.

Syntax

```
ssh [on|off|show]
```

Table 27-13: Network category ssh command options

`on`	Enables the SSH port.
`off`	Disables the SSH port.
`show`	Shows whether SSH port is enabled or not.

Example

The following example displays the SSH port status:

```
0000:dsm$ network
0001:network$ ssh show
ssh port : on
SUCCESS: ssh port status shown.
0002:network$
```

ping

The `ping` command sends ICMP (Internet Control Message Protocol) echo request packets (ECHO_REQUEST) to a specified network host. The `ping` command uses the ICMP protocol's mandatory echo request datagram to elicit an ICMP echo response (ECHO_RESPONSE) from a host or gateway. The `ping` command sends six packets to the network host and then reports the results.

Syntax

```
ping {ipaddress|FQDN}
```

Table 27-14: Network category ping command options

ipaddress	IP address of the host from which you want a response.
FQDN	Fully qualified domain name of the host from which you want a response.

Example

The following example sends a `ping` request to the host `vmlinux04_RH5`:

```
0022:network$ ping deptsys
PING deptsys (1.3.5.9) 56(84) bytes of data.
64 bytes from deptsys (1.3.5.9): icmp_seq=1 ttl=64 time=3.07 ms
64 bytes from deptsys (1.3.5.9): icmp_seq=2 ttl=64 time=0.477 ms
64 bytes from deptsys (1.3.5.9): icmp_seq=3 ttl=64 time=0.121 ms
64 bytes from deptsys (1.3.5.9): icmp_seq=4 ttl=64 time=0.136 ms
64 bytes from deptsys (1.3.5.9): icmp_seq=5 ttl=64 time=0.131 ms
64 bytes from deptsys (1.3.5.9): icmp_seq=6 ttl=64 time=0.214 ms
--- deptsys ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5003ms
rtt min/avg/max/mdev = 0.121/0.691/3.070/1.071 ms
ping SUCCESS
```

traceroute

The `traceroute` command uses the IP-protocol time field to elicit an ICMP time exceeded (`TIME_EXCEEDED`) response from each gateway along the path to a specified host.

Specify the target IP address or FQDN. The `traceroute` command supports a timeout option.

Syntax

```
traceroute (ipaddress|FQDN) {timeout}
```

Table 27-15: Network category traceroute command options

ipaddress	IP address of the host for which you want the path information.
FQDN	Fully qualified domain name of the host for which you want the path information.
timeout	The time period in seconds after which the request is dropped, range is from 1 to 60 seconds.

Example

The following example sends a `traceroute` command request to an IP address:

```
0028:network$ traceroute 192.168.60.7
```

```
traceroute to 192.168.60.7 (192.168.60.7), 30 hops max, 40 byte packets
 1  10.3.244.3  3000.605 ms !H  3000.571 ms !H  3000.548 ms !H
Traceroute Completed
0029:network$
```

rping

The `rping` command sends Address Resolution Protocol (ARP) requests to a neighbor host, pings the address on the device interface by ARP packets and informs how many users are using a particular IP address.

Syntax

```
rping ipaddress {eth0|eth1}
```

Example

```
0024:network$ rping 1.3.5.9 eth0
ARPING 1.3.5.9 from 1.3.5.7 eth0
Unicast reply from 1.3.5.9 [00:0C:29:36:9E:B3]  2.518ms
Unicast reply from 1.3.5.9 [00:0C:29:36:9E:B3]  0.817ms
Unicast reply from 1.3.5.9 [00:0C:29:36:9E:B3]  0.866ms
Sent 3 probes (1 broadcast(s))
Received 3 response(s)
Arping SUCCESS
0025:network$
```

arp

The `arp` command displays the current Address Resolution Protocol (ARP) cache of the DSM.

Syntax

```
arp
```

Example

The following example displays the current ARP cache:

```
0001:network$ arp
1.3.5.25 dev eth0 lladdr 00:08:a1:59:c1:cc REACHABLE
1.3.5.254 dev eth0  FAILED
1.3.11.14 dev eth0 lladdr 00:17:31:6f:58:16 STALE
link info
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue \    link/loopback
00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 1000\
link/ether 00:0c:29:60:f9:3e brd ff:ff:ff:ff:ff:ff
3: sit0: <NOARP> mtu 1480 qdisc noop \    link/sit 0.0.0.0 brd 0.0.0.0
arp SUCCESS
0002:network$
```

If a connection is STALE, ping it and check again. It should change to REACHABLE. If it does not change, or it changes to FAILED, the connection is no longer available.

checkport

The `checkport` command is used to scan a port on a network-accessible system to verify that a TCP connection can be made to the system using the specified port. It does not guarantee that you can log on, just that a communication channel can be opened on the DSM or on a host. It is typically used to check the status and availability of the ports through which to administrate and run the DSM. These are ports such as 22, 7024, 8443, 8444, and 8445. The `checkport` command returns the transport layer protocol and the service using that port. The transport layer protocol is always TCP. The service is a system service like ssh, vmsvc, and *.

A "Connection refused" message can be returned for various reasons, such as a port is not assigned and/or is not in a LISTEN state.

> **Note**
> If a DSM port refuses a connection, you must troubleshoot the TCP connection.

`checkport` activity is logged in the Management Console, and is displayed when operating outside of a domain. A sample **Logs** window entry is shown below.

```
18713  2010-08-27 13:07:11.944 PDT  I  vmSSA05  CLI0003I: [cliadmin] network checkport
vmlinux101 7024
```

When `checkport` is executed in the Management Console interface, rather than on the command line, the log entry is appended with "`timeout x`", where *x* is either the value you entered on the command line or the default timeout.

Syntax

```
checkport host port [timeout x]
```

where, host is an IP address, FQDN, hostname, or even "localhost". Typically, it is a valid DSM or agent host, as configured in the Management Console.
port is a single TCP port number or a range of port numbers. A port number range is a hyphen/dash-separated list and is entered in the form "startnum-endnum". For example, `8440-8449`.
x is an integer between 1 and 600, inclusive. It is the timeout threshold and is expressed in seconds. The default is `180` seconds.

Example

The following example checks the availability of port (`8445`) used to run the Management Console on a DSM node:

```
0004:network$ checkport vmSSA06 8445
Connection to vmSSA06 8445 port [tcp/*] succeeded!
SUCCESS: invoked checkport(nc) command.
0005:network$
```

The following example checks the availability of port (`7024`) used to download configuration data to an agent host:

```
0005:network$ checkport solaris120 7024
Connection to solaris120 7024 port [tcp/vmsvc] succeeded!
SUCCESS: invoked checkport(nc) command.
0006:network$
```

The following example checks the availability of a range of ports on the local system, a DSM, and includes a 10 second timeout.

```
0081:network$ checkport localhost 8440-8449 timeout 10
nc: connect to localhost port 8440 (tcp) failed: Connection refused
nc: connect to localhost port 8441 (tcp) failed: Connection refused
nc: connect to localhost port 8442 (tcp) failed: Connection refused
nc: connect to localhost port 8446 (tcp) failed: Connection refused
```

```
nc: connect to localhost port 8447 (tcp) failed: Connection refused
nc: connect to localhost port 8448 (tcp) failed: Connection refused
nc: connect to localhost port 8449 (tcp) failed: Connection refused
Connection to localhost 8443 port [tcp/pcsync-https] succeeded!
Connection to localhost 8444 port [tcp/pcsync-http] succeeded!
Connection to localhost 8445 port [tcp/*] succeeded!
SUCCESS: invoked checkport(nc) command.
```

nslookup

The `nslookup` command is used to query the DNS to get hostname to IP address mapping. Specify the FQDN or IP address of the server for which you want the IP address or host name information.

Syntax

```
nslookup HOST_NAME [timeout {1..600} | port {1..65535}]
```

Example

```
0010:network$ nslookup linuxhost.domain.com
Server: <dns server>
Address: <dns server ip address>
Name: linuxhost.domain.com
Address: 1.2.3.4
SUCCESS: invoked nslookup command.
```

System Category Commands

The `system` configuration category enables you to set the appliance host name, enable/disable the console port, create certificates, restart the DSM, and reboots/shuts down the DSM.

> **Note**
> These DSM CLI commands work only in DSM configurations that run on a DSM appliance. Software-only DSM installations do not support the `console`, `reboot`, `setinfo`, and `shutdown` commands. For software-only DSMs, use the regular operating system commands to name and restart the system.

Enter the `system` configuration category by typing:

```
0001:dsm$ system
```

The `system` category supports the following commands:

Table 27-16: System category commands

setinfo	Sets the host name or FQDN of the DSM.
console	Enables or disables the serial console port.
security	Creates the CA signer certificate and the DSM certificate. It also signs the DSM certificate.
mfauth	Enables, disables, or displays the configuration status of multi-factor authentication.
tls1	Enable/Disable TLS 1.0/1.1 support
shutdown	Stops the DSM software and powers off the appliance.

Table 27-16: System category commands (continued)

`reboot`	Reboots the DSM appliance and restarts the software.
`server`	Provides the options to restart, start, and stop the DSM as well as the option to check the status of the DSM.

setinfo

The `setinfo` command enables you to set the host name of the DSM and display appliance-related information such as the hardware UUID, serial number, and uptime.

The assigned name is used to identify the appliance and identify the certificate owner. If you change the host name after generating the CA signer and DSM certificates, you must regenerate the certificates because the host name is used in the certificates to identify the DSM.

Syntax

```
setinfo [show | hostname | sshbanner ]
```

The `setinfo` command can take the following arguments:

Table 27-17: System category setinfo command arguments

`hostname`	Sets the host name for your system. This option takes one argument, the network name to assign the appliance.
`sshbanner`	Defines the `/etc/ssh/ssh-banner` file. Available only on Vormetric-provided physical and virtual DSMs. Edit the banner shown when logging on to the DSM CLI. The default is "Welcome to the Vormetric Data Security Manager."
`show`	Shows the current `setinfo` settings.

Example

The following example sets the DSM host name to `vmSSA001`:

```
0005:system$ setinfo hostname SSA666
SUCCESS: setinfo hostname. If the DSM certificate is already generated, please re-sign
the server certificate to reflect the hostname changes.
0006:system$
```

setinfo show

The `setinfo show` command displays general appliance information. The following example was taken on an appliance-based DSM.

```
0017:system$ setinfo show
hostname = SSA666
UUID = 53D19F64-D663-A017-8922-003048C497D4
serial number = 999X9120411
part number = 30-1010002-01
uptime =  10:36:56 up 15:47,  2 users,
          load average: 0.09, 0.05, 0.01
ssh banner = Welcome to the Vormetric Data Security Manager.
Show setinfo SUCCESS
0018:system$
```

console

The console option displays the state of the serial console. By default the serial console is always on. If you turn off the serial console port the only access to the appliance will be through the network. We strongly recommend that you leave the serial console on.

- `console on`—This command turns the serial console on. It is on by default.

- `console off` —This command turns the serial console off. You cannot use the serial console to log on when the console is off.

Syntax

```
console [on | off | show]
```

Example

```
0013:system$ console on
Be prepared to wait for a few minutes
0014:system$ console show
console on
```

security

The system category security command creates the SSL credentials used to authenticate DSMs and their agents.

Table 27-18: System category security command arguments

`masterkey`	Master key management
`signcert`	Re-signs the DSM certificate.
`gencert`	Re-generates the DSM certificate.
`genca`	Generates the CA signing certificate on the initial DSM.
`suiteb`	Suite B mode configuration (deprecated, use 'mode' instead)
`mode`	Security mode configuration
`cc`	Common Criteria Mode configuration
`boot-passphrase`	Set a passphrase to unlock the DSM at system boot time, to maintain the security of the encrypted filesystem. This is feature is available only on a fresh installation of v6.0.2 or later.
`legacyregistration`	Manually close/open port 8080 for new deployment or backwards compatibility.

masterkey

The `security masterkey` command displays the DSM master key. It displays the master key identifier and the date on which it was created. You can also rotate the masterkey from this menu.

Syntax

```
security masterkey [show | rotate]
```

Table 27-19: System category masterkey command arguments

show	Show master key information.
rotate	Rotate the master key.

Example 1

```
0001:system$ security masterkey show
identifier=4fc24a6b
creation_date=2016-04-08
SUCCESS: showed master key info
```

Example 2

```
0002:system$ security masterkey rotate
WARNING: Rotating the master key will restart DSM software automatically.
Continue? (yes|no)[no]:yes
SUCCESS: rotated master key.
0003:system$
```

> ⚠️ **WARNING**
>
> **Master key rotation from the DSM (`system > security masterkey rotate`) overwrites the existing key on the Luna, thereby rendering your previous backups nonfunctional. Therefore, you MUST make a new backup after rotating the master key.**

signcert

This utility signs the DSM certificate for the HA node. Usually it is used to re-sign expired certificates.

The utility does the following, in the following order:

- Generates a new certificate request from the existing key pair in the keystore
- Gets the certificate request signed by the CA (Certificate Authority, located on the initial DSM)
- Imports the new DSM certificate back to the keystore

DSM and VTE Agent communication is not affected by this change.

Syntax

```
security signcert
```

Example

```
0037:system$ security signcert
WARNING: The server certificate will be resigned, and the security server software will
be restarted automatically!
Continue? (yes|no)[no]:yes
This computer may have multiple IP addresses. All the agents will have to connect to
Security Server using same IP.
Enter the host name of this computer. This will be used by Agents to talk to this
Security Server.
Security Server host name[vmlinux03_RH5]:
Please enter the following information for key and certificate generation. Security
Server Certificate Configuration
```

```
        What is the name of your organizational unit? []:UnitX

        What is the name of your organization? []:Widgets, Inc.

        What is the name of your City or Locality? []:Santa Clara

        What is the name of your State or Province? []:CA

        What is your two-letter country code? [US]:

        What is your email address? []

        What is the validity period of the generated certificate (from 2 to 10 years)? [10]:

        Regenerating the server certificates now...

        Deleting existing key with alias cgss_server_app

        Renaming new key with alias cgss_server_app_new to cgss_server_app

        Generating certificate signing request

        Signing certificates

        Deleting old signer certificate from keystore

        Importing new signer certificates into keystore

        Importing new server certificates into keystore

        DB20000I The SQL command completed successfully.

        Server certificate has been re-signed by the Certificate Authority successfully.

        Starting the Security Server.
```

After you enter this data, the utility creates certificates, completes the installation process, and then starts the DSM. You are then returned to the DSM CLI prompt.

gencert

The DSM CLI `security gencert` command generates the DSM certificate for the current DSM.

Regenerate the DSM certificate when:

- The DSM key has been compromised

- The DSM certificate has expired

- When the host name of the DSM changes

- One of the certificate fields (such as organization, city, and so on) of the certificate has changed

- When the DSM is restored on another appliance with a different host name

The utility does the following, in the following order:

- Checks for an existing DSM certificate

- Generates a new key pair in the keystore

- Swaps the master key encryption to use the new key pair

- Deletes the old key pair in the keystore

- Generates a new certificate request based on the new key

- Gets the certificate request signed by the CA (Certificate Authority, located on the initial DSM)

- Imports the new DSM certificate back to the keystore

DSM and VTE Agent communication isand the validity period not affected by this change. The information that you provide is displayed when the signer-certificate is viewed. You are prompted to specify:

- Your organizational unit, which is frequently a department or group name

- Organization name, which is frequently the company name

- City or locality in which the organization is located

- State or province in which the organization is located

- The country in which the organization is located

- The validity period for the certificate

After you enter this data, the utility creates certificates, completes the installation process, and then starts the DSM. After which, you are returned to the DSM CLI prompt.

Syntax

```
security gencert
```

Example

```
0036:system$ security gencert

WARNING: All Peer node and agent certificates will need to be re-signed after CA and
server certificates are regenerated!

The security server software will be restarted automatically.

Continue? (yes|no)[no]:yes

This node may have multiple IP addresses. All of the agents will have to connect to the
Security Server using the same IP.

Enter the host name of this node. This will be used by Agents to talk to this Security
Server.

This Security Server host name[dsm15104.i.thales.com]:

Please enter the following information for key and certificate generation.

What is the name of your organizational unit? []:Really Fine Stuff

What is the name of your organization? []:Widgets, Inc.

What is the name of your City or Locality? []:Santa Clara

What is the name of your State or Province? []:CA

What is your two-letter country code? [US]:

What is your email address? []:

What is the validity period of the generated certificate (from 2 to 10 years)? [10]:

Regenerating the server certificates now...

SUCCESS: The security certificates are re-generated and the Security Server software is
restarted.

0037:system$
```

genca

The `security genca` command regenerates the Certificate Authority (CA) on the initial DSM.

The administrator should run this utility in one of the following situations:

- Setting up a new DSM

- When the signer key is compromised

- When the signer certificate expires

- Any of the fields of the signer certificate has changed

- Restoring a backup configuration to a different DSM (Recommended)

The command does the following, in the following order:

- Generates a new signer certificate

- Deletes the old signer certificate from the keystore

- Imports the new signer certificate into the keystore

- Generates a new certificate request from the existing DSM certificate

- Signs the DSM certificate with new CA

- Imports the new DSM certificate into the keystore

- Restarts the DSM

Do the following operations after running the `security genca` command.

- If HA nodes DSMs are configured, there is now a certificate mismatch and the HA certificates must be re-signed. Establish a DSM CLI connection to each HA node and run the `signcert` command.

Every agent registered with the DSM must be re-registered. No agent > DSM communication will occur until the following steps are completed:

1. Disable the agent's registration from the Management Console to remove the agent's certificates.

2. Re-enable that agent's registration on the DSM.

3. Run register_host on that agent.

The information that you provide is displayed when the signer-certificate is viewed. You are prompted to specify:

- Your organizational unit, which is frequently a department or group name

- Organization name, which is frequently the company name

- City or locality in which the organization is located

- State or province in which the organization is located

- The country in which the organization is located

After you enter this data, the utility creates certificates, completes the installation process, and then starts the DSM. You are then returned to the DSM CLI prompt.

Syntax

```
security genca
```

Example

```
0001:system$ security genca
WARNING: All Agents and Peer node certificates will need to be re-signed after CA and
server certificate regenerated, and the Security Server software will be restarted
automatically!
Continue? (yes|no)[no]:yes
This computer may have multiple IP addresses. All the agents will have to connect to
Security Server using same IP.
Enter the host name of this computer. This will be used by Agents to talk to this
Security Server.
Security Server host name[vmSSA05]:
Please enter the following information for key and certificate generation. Security
Server Certificate Configuration
```

```
What is the name of your organizational unit? []:Widgets

What is the name of your organization? []:Excelsior

What is the name of your City or Locality? []:S.C.

What is the name of your State or Province? []:CA

What is your two-letter country code? [US]:

What is your email address? []:

What is the validity period of the generated certificate (from 2 to 10 years)? [10]:

Regenerating the CA and server certificates now...

SUCCESS: The CA and security certificates are re-generated and the Security Server
software is restarted.

0002:system$
```

suiteb

The `suiteb` command has been deprecated. Use mode instead.

mode

The `mode` command is used to activate or deactivate suiteb , compatible, or RSA mode.

Syntax

```
security mode [set [suiteb | compatible | rsa]] | [show]
```

The `mode` command can take the following arguments:

Table 27-20: Security mode command arguments

`set`	This command activates or deactivates suiteb, compatible, or RSA mode `[suiteb	compatible	rsa]`
`show`	Show security mode configuration

The DSM is in compatible mode by default. The suiteb or RSA modes must be enabled to take effect.

Compatibility mode uses both RSA and ECC certificates. The DSM uses the ECC certificate to communicate with other DSMs, newly installed or upgraded agents. It uses the RSA certificate to communicate with older agents. When the agent negotiates a transaction with the DSM, the handshake determines which certificate to use.

Suiteb uses only ECC certificates. Older agents that do not support ECC must be upgraded or they will fail to communicate with the DSM.

RSA mode uses only RSA certificates. Any agents that registered when RSA mode was enabled must re-register because the ECC port is now closed. In order to communicate with the DSM, agents must re-register with that DSM.

Examples

The following example activates suiteb mode:

```
0008:system$ security mode set suiteb
```

System Response:

```
Important!: Ensure DSM ports 8446-8448 are not blocked by corporate firewall.

In addition, all other DSMs in this cluster must be manually set to the suiteb mode.

The Security Server will be restarted.

Continue? (yes|no)[no]:
```

The following example shows whether suiteb is configured:

```
0003:system$ security mode show
Current mode is: suiteb
SUCCESS: showed security mode status
```

The following example activates RSA mode:

```
0004:system$ security mode set rsa
Important!: Ensure DSM ports 8443-8445 are not blocked by corporate firewall.
In addition, all other DSMs in this cluster must also be in rsa mode.
The Security Server will be restarted.
Continue? (yes|no)[no]:
```

Type `yes` to enable RSA mode.

cc

The `cc` command is used to enable or disable Common Criteria mode.

Syntax

```
security cc [on | off | show]
```

The `cc` command takes the following arguments:

Table 27-21: Security cc command arguments

on	Enable Common Criteria mode
off	Disable Common Criteria mode
show	Show console port status

Example

The following example enables Common Criteria mode:

```
0001:system$ security cc on
cc (Common Criteria) mode is node specific configuration, and need to be configured in
each cluster node individually. Turning on cc (Common Criteria) will improve the
security level but limit some functionality, and server will restart automatically,
continue? (yes|no)[no]:
```

The following example shows whether Common Criteria is enabled:

```
0002:system$ security cc show
SUCCESS: Common Criteria mode is off
0003:system$
```

boot-passphrase

The `boot-passphrase` command sets and manages a passphrase required at DSM system boot time to unlock the system. Refer to the *DSM Installation & Configuration Guide* for more information about this feature.

Syntax

```
security boot-passphrase [ set | clear | show ]
security boot-passphrase recovery [ show [<filename>] | delete <filename> ]
```

Table 27-22: Security boot-passphrase command arguments

set	Set a boot-passphrase to unlock the DSM filesystem at system boot up
clear	Clears the boot passphrase
show	Shows whether a boot-passphrase has been set or not
recovery show	Displays the passphrase recovery file and displays the contents—the encrypted passphrase, the public key used to encrypt the passphrase, and a sample command to recover the passphrase using the associated private key
recovery delete	Deletes the passphrase recovery file

Example 1

The following example shows how to set a boot passphrase. After setting the passphrase, the DSM reboots and the SSH console connection is lost. You need to have IPMI Java console access, or if using a virtual appliance, you can connect to the DSM via the console available from the virtualization application in use. Refer to the DSM Installation & Configuration Guide for details about setting a boot passphrase.

```
0000:dsm$ system

0001:system$ security boot-passphrase set

An RSA public key with minimum length of 2048 bits is required for boot passphrase
recovery. Please enter one now, ending with an empty line:

-----BEGIN PUBLIC KEY-----

MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAwYIf0Z04nzne9j78BY7m

Q9kMTgh8YErtklECnVVhxExob/UvAWOvSBcGDVgixpeMCywWVh8OgTIbj751PVfa

TI8/C+gP4Rd6cdtO7fGzsYsAZxN9OCssRQlCJfCe6y6fNep3dDOh1noTFyFNTqOy

c3WW0gAlJ9ILPwn6uxVRgtXPgLnFfP9zNieyWmHTLw6He8BZAAYkWbESMgnA5BoJ

mcxdpv/i/8ZODTMMo/6Ji4oYpQPa8i9Ex7qTZinl5hxjIjC8eIcUOMNdAhvslNzs

T6FZPJ2BEYBU6TAQpxDPLwPAQIEw1x/NzcYUUfgaP1pZIAdhWFJUZkx4FqmEA5od

MwIDAQAB

-----END PUBLIC KEY-----

Enter new boot passphrase:

Enter new boot passphrase again:

WARNING: After setting the new boot passphrase, the system will be rebooted
automatically and the new passphrase must be entered on the console. If you do not have
direct or IPMI access to the console, then choose 'no' to cancel. DSM will not boot up
until a correct boot passphrase is entered.

Continue? (yes|no)[no]: yes

NOTE: run this command on every server node in the cluster to keep them at a uniform
security level.

SUCCESS: custom boot passphrase has been set.

DSM server is rebooting...
```

Example 2

The following example shows whether a boot passphrase has been set or not:

```
0008:system$ security boot-passphrase show

Prompt-On-Boot mode - the system disk is encrypted and there is a custom boot passphrase
set.

SUCCESS
```

```
0009:system$
```

Example 3

The following example shows the recovery file and the contents of that file

```
0000:dsm$ system

0001:system$ security boot-passphrase recovery show

SUCCESS

The following passphrase recovery files are available:

0. 201710031407


Type the number of a file to view the contents, or 'q' to quit: 0

Encrypted passphrase (base64 encoded):

fqWOGbKe4x6R3vmWtBMFvoAauaEpOnQ9OGLmFW9eZhFbv+w1+u0LPgIGYx9e5ATH8GOvuJvht7UzBodMA07DHNMp
yMnOEsy6Nz+ouWsMWhHen5JFNMXKWM9TYQ9/yr

W/2McKIYuBqgeaOefzL2jr8vyyFudq6TGgTjRJe1edLDCqTJbcK100o036U0vyn

Lpes6Zp1ud5usWngn2J2X6PrlAugHp4nMMDIRLQBgzX95x7Fb7VLebcb/eIGn39xh/f6azXhHpjahwjirzfpZl03

00VFYT0P9o5xg==

Public key used for encryption:

-----BEGIN PUBLIC KEY-----

MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAwYIf0Z04nzne9j78BY7Q9kMTgh8YErtklECnVVhxExob
/UvAWOvSBcGDVgixpeMCywWVh8OgTIbj751PVf

TI8/C+gP4Rd6cdtO7fGzsYsAZxN9OCssRQlCJfCe6y6fNep3dDOh1noTFyFNTqO

c3WW0gAlJ9ILPwn6uxVRgtXPgLnFfP9zNieyWmHTLw6He8BZAAYkWbESMgnA5Bo

mcxdpv/i/8ZODTMMo/6Ji4oYpQPa8i9Ex7qTZinl5hxjIjC8eIcUOMNdAhvslNz

T6FZPJ2BEYBU6TAQpxDPLwPAQIEw1x/NzcYUUfgaP1pZIAdhWFJUZkx4FqmEA5o

MwIDAQAB

-----END PUBLIC KEY-----

Example command for decrypting the passphrase, given the matching private key:

    base64 -d <file-containing-the-ciphertext-above> | openssl rsautl -inkey <private-
key> -decrypt

SUCCESS
```

Example 4

The following example clears the boot passphrase. When you clear the passphrase the recovery file is not deleted, you can delete it later, see next example:

```
0004:system$ security boot-passphrase clear

Enter current boot passphrase:

WARNING: After clearing the custom boot passphrase, the system will be rebooted
automatically!

Continue? (yes|no)[no]: ^C

0005:system$ security boot-passphrase clear

Enter current boot passphrase:

WARNING: After clearing the custom boot passphrase, the system will be rebooted
automatically!

Continue? (yes|no)[no]: yes

NOTE: run this command on every server node in the cluster to keep them at a uniform
security level.

This operation will take some time, please wait....

SUCCESS: custom boot passphrase has been cleared, unattended boot enabled.
```

```
DSM server is rebooting
0006:system$
```

Example 5

The following example deletes the passphrase recovery file:

```
0003:system$ security boot-passphrase recovery delete 201710031407
SUCCESS: removed passphrase recovery file '201710031407'
0004:system$
```

legacyregistration

Port 8080 is no longer used for registration, but you can manually close/open this legacy port for new deployment.

Syntax

```
# security legacyregistration [ on | off | show ]
```

Example

```
# security legacyregistration show
```

mfauth

The `mfauth` command enables or disables the multi-factor authentication of DSM Management Console administrators. When enabled, the Management Console login screen displays the usual **Login** and **Password** boxes.

- DSM Management Console administrators with multi-factor authentication must enter the value displayed on their RSA SecurID device into the **Password** box.

- DSM Management Console administrators without multi-factor authentication should enter the DSM administrator password in the Password box.

The `mfauth` command includes a `clean` option to remove all configured administrator:device bindings. To remove the administrator:device bindings of individual Management Console administrators, open the Edit Administrator window and delete the value in the **RSA User Name** box.

The `mfauth` command displays the enabled/disabled status of multi-factor authentication. The current `mfauth` status is also indicated in the **System > General Preferences > System** tab. The **Multi-Factor Authentication Status** checkbox is a display indicator only and cannot be set in the Management Console.

Do not enable multi-factor authentication until after the RSA Authentication Agent is configured on the DSM.

> **Note**
> You must enable `mfauth` on all HA nodes. If not, it results in an HA node database mismatch.

Syntax

```
mfauth on | off | clean | show
```

Example

The following example shows the current multi-factor configuration status of the DSM. The DSM is configured for multi-factor authentication. It's just not enabled.

```
0001:system$ mfauth show
Administrator multiple factor authentication : off
RSA secret file configured : on
SUCCESS: administrator multiple factor authentication status showed.
```

```
0002:system$
```

The following example enables multi-factor authentication for DSM administrators.

```
0010:system$ mfauth on

WARNING: After enabling the administrator multiple factor authentication, the security
server software will start to validate the extra one-time password!

Continue? (yes|no)[no]:yes

SUCCESS: administrator multiple factor authentication enabled.

0011:system$
```

The following example deletes the RSA node secret file from the DSM, effectively breaking all communication between the DSM and the RSA Authentication Manager, which, in turn, disables multi-factor authentication of DSM administrators. If you remove the node secret using `mfauth clean`, you must also delete the node secret on the RSA Security Console by running `Clear Node Secret`. A new node secret will be automatically generated the next time any Vormetric administrator logs into the Management Console. To break all DSM administrator device bindings on the DSM:

```
0001:system$ mfauth clean

WARNING: Cleaning RSA secret file will break the communication between the security
server and RSA server!

Continue? (yes|no)[no]:yes

SUCCESS: RSA secret file is removed.

0002:system$
```

tls1

The `tls1` command enables or disables support for TLS protocols v1.0 and 1.1. This command turns on or turns off support for both v1.0 and v1.1 of the TLS protocol, you cannot select one or the other.

Syntax

```
tls1 [on | off | show]
```

The `tls1` command takes the following arguments:

Table 27-23: System tls1 command arguments

on	Enable TLS 1.0/1.1 support
off	Disable TLS v1.0/1.1 support
show	Show status of TLS v1.0/1.1 support (enabled or not)

Example

The following example turns on TLS v1.0/1.1:

```
0002:system$ tls1 on

WARNING: After enabling TLS 1.0/1.1, the security server software will restart!

Continue? (yes|no)[no]:yes

SUCCESS: TLS 1.0/1.1 enabled and server restarted.

This change only affected this node. Run the same tls1 command on all other nodes in the
cluster.

0003:system$
```

The following example shows whether TLS v1.0/1.1 is enabled or not:

```
0004:system$ tls1 show
```

```
TLS 1.0/1.1 is disabled
SUCCESS: TLS 1.0/1.1 status shown
0005:system$
```

shutdown

The shutdown command stops the DSM software, brings down the appliance operating system, and then powers off the appliance. Configuration changes are automatically saved. Afterwards, the appliance can be safely turned off.

Syntax

```
shutdown
```

Example

The following example shuts the system down:

```
0038:system$ shutdown
Do you want to shutdown the system ? (y/n):y
Shutting down now...
Shutdown SUCCESS
0039:system$
```

The last message displayed on the appliance LCD before it powers down is:

```
Power off or reboot in approx 15 secs
```

You can remove the power cords and power modules after the appliance powers down. Reapply power by reattaching the power cords and pressing the power switch. The power-interrupt alarm may sound. If it does, press the red reset button on the back of the appliance.

reboot

The reboot command reboots the DSM appliance.

Syntax

```
reboot
```

Example

The following example reboots the system immediately:

```
0001:system$ reboot
Reboot the system y/n?
Rebooting now...
Reboot SUCCESS
system$
Broadcast message from root (Sun Feb 9 02:44:20 2014):
The system is going down for reboot NOW!
```

server

Previous CLI commands limited the DSM CLI administrator to do a restart of the DSM. The commands have been enhanced so that DSM CLI administrators can start and stop the DSM based on the need for maintenance intervals, test cycles and so on. Available server commands are listed below:

Table 27-24: System category server commands

`restart`	Restarts the DSM software. Shuts down the DSM software and then restarts it.
`start`	Starts the DSM software.
`stop`	Stops the DSM software.
`status`	Displays the DSM software running status.

restart

The `restart` command stops and then starts the DSM software. It does not reboot the appliance. The `reboot` command restarts the DSM operating system and, in the process of coming up, starts the DSM software. Use the `reboot` command only if `restart` does not correct a problem.

Syntax

```
restart
```

Example

The following example restarts the DSM:

```
0033:system$ server restart
Do you want to restart the server software ? (y/n):y
Restarting now...
Stopping Security Server...done.
Stopping the data store...done.
Starting Security Server...done.
SUCCESS: The security server software is restarted.
0034:system$
```

status

The status command displays the current running status of the DSM software.

Syntax

```
status
```

Example

The following example shows the status display.

```
0038:system$ server status
Security Server is running.
Security Server uptime:  2 days, 09:27:27
SUCCESS: The security server software status is shown.
```

HSM Category Commands

> **Note**
> HSM commands apply only to HSM-enabled DSM appliances. Do not execute HSM commands if the DSM Management Console System > About window, or the DSM CLI `hsm$ show` command, displays "HSM is disabled".

The HSM category includes commands to display configuration and status information about the HSM hardware device that is inside the DSM appliance or system.

Table 27-25: HSM category commands

`show`	Displays the HSM ready status and if the HSM password is manually or automatically entered during DSM startup.
`status`	Shows HSM card status.
`statistics`	Display the HSM hardware statistics.
`slots`	Shows all the HSM slots.
`replaceacs`	Replaces the Administrator Card Set
`test`	Test the HSM hardware
`remoteadmin`	Enables or disables remote administration support, and displays remote administration status.
`connect`	Adds or removes nShield Connect network HSM, and displays status
`secworldupdate`	Updates Security World data on the this DSM node
`luna`	Add/Delete/Show Luna SA network HSM.

show

The show command displays the configuration status of the HSM.

Syntax

```
show
```

Example

```
0037:hsm$ show
```

If an HSM device is not present in the appliance or system, the following message is displayed:

```
0003:hsm$ show
HSM is disabled
SUCCESS show hsm
0004:hsm$
```

If an HSM device is present in the appliance or system, a message like the following is displayed. The `unattendedstart` parameter is set to `on`, indicating that the appliance is configured to automatically provide the HSM password to DSM operations that require the password. The `unattendedstart` parameter is not user-configurable.

```
0001:hsm$ show
HSM is enabled
unattendedstart: on
SUCCESS show hsm
0002: hsm
```

status

The status command displays the configuration status of the HSM, as well as other important information. It displays the HSM model, manufacture date, serial number, board revision, firmware version, security mode, and so on.

Syntax

```
status
```

Example: V6100 appliance output

```
HA auto recovery:  enabled
              HA recovery mode:  activeEnhanced
 Maximum auto recovery retry:  500
 Auto recovery poll interval:  60 seconds
                   HA logging:  disabled
           Only Show HA Slots:  yes
          HA Group Label:  g30606
         HA Group Number:  11394399181044
        HA Group Slot ID:  4
        Synchronization: enabled
           Group Members:  1394399181044
               Needs sync:  no
           Standby Members:  <none>
Slot #    Member S/N                      Member Label    Status
======    ==========                      ============    ======
------  1394399181044                     dsmpartition0    alive
```

statistics

The statistics command displays the HSM hardware statistics.

Syntax

```
statistics
```

Example 1: V6100 appliance

```
+ServerGlobals:
   -Uptime              7114
   -CmdCount            6082
   -CmdBytes            255096
   -CmdMarshalErrors    0
   -ReplyCount          6087
   -ReplyBytes          245564
   -ReplyMarshalErrors  0
   -ClientCount         9
   -MaxClients          9
   -DeviceFails         0
   -DeviceRestarts      0
+Connections:
   +#1:
      -Uptime              7114
      -CmdCount            0
      -CmdBytes            0
      -CmdMarshalErrors    0
```

```
          -ReplyCount              0
          -ReplyBytes              0
          -ReplyMarshalErrors      0
          -DevOutstanding          1
          -QOutstanding            0
          -LongOutstanding         0
          -RemoteIPAddr            (local)
          -ClientNumber            1
          -ClientProcessID         0
          -ClientProcessName
          -ObjectCountTotal        0
          +PerModule:
             +#1:
                -ObjectCount            0
       +#8:
          -Uptime                  7081
          -CmdCount                3
          -CmdBytes                132
          -CmdMarshalErrors        0
          -ReplyCount              3
          -ReplyBytes              492
          -ReplyMarshalErrors      0
          -DevOutstanding          0
          -QOutstanding            0
          -LongOutstanding         0
          -RemoteIPAddr            (local)
          -ClientNumber            8
          -ClientProcessID         2566
          -ClientProcessName       /large/opt/nfast/sbin/raserv
          -ObjectCountTotal        0
          +PerModule:
             +#1:
                -ObjectCount            0
       +#20:
          -Uptime                  7057
          -CmdCount                51
          -CmdBytes                8436
           -CmdMarshalErrors        0
          -ReplyCount              51
          -ReplyBytes              3140
          -ReplyMarshalErrors      0
          -DevOutstanding          0
          -QOutstanding            0
          -LongOutstanding         0
          -RemoteIPAddr            (local)
```

```
        -ClientNumber           20
        -ClientProcessID        4001
        -ClientProcessName      /large/opt/nfast/bin/ppmk
        -ObjectCountTotal       1
        +PerModule:
            +#1:
                -ObjectCount            1
    +#19:
        -Uptime                 7057
        -CmdCount               10
        -CmdBytes               1556
        -CmdMarshalErrors       0
        -ReplyCount             10
        -ReplyBytes             460
        -ReplyMarshalErrors     0
        -DevOutstanding         0
        -QOutstanding           0
        -LongOutstanding        0
        -RemoteIPAddr           (local)
        -ClientNumber           19
        -ClientProcessID        4001
        -ClientProcessName      /large/opt/nfast/bin/ppmk
        -ObjectCountTotal       1
        +PerModule:
            +#1:
                -ObjectCount            1
    +#21:
        -Uptime                 7055
        -CmdCount               57
        -CmdBytes               8612
        -CmdMarshalErrors       0
        -ReplyCount             57
        -ReplyBytes             4204
        -ReplyMarshalErrors     0
        -DevOutstanding         0
        -QOutstanding           0
        -LongOutstanding        0
        -RemoteIPAddr           (local)
        -ClientNumber           21
        -ClientProcessID        4001
        -ClientProcessName      /large/opt/nfast/bin/ppmk
        -ObjectCountTotal       1
        +PerModule:
            +#1:
                -ObjectCount            1
```

```
+#22:
    -Uptime                 7055
    -CmdCount               306
    -CmdBytes               7372
    -CmdMarshalErrors       0
    -ReplyCount             306
    -ReplyBytes             17780
    -ReplyMarshalErrors     0
    -DevOutstanding         0
    -QOutstanding           0
    -LongOutstanding        0
    -RemoteIPAddr           (local)
    -ClientNumber           22
    -ClientProcessID        4001
    -ClientProcessName      /large/opt/nfast/bin/ppmk
    -ObjectCountTotal       1
    +PerModule:
        +#1:
            -ObjectCount            1
+#128:
    -Uptime                 90
    -CmdCount               6
    -CmdBytes               164
    -CmdMarshalErrors       0
    -ReplyCount             6
    -ReplyBytes             304
    -ReplyMarshalErrors     0
    -DevOutstanding         0
    -QOutstanding           0
    -LongOutstanding        0
    -RemoteIPAddr           (local)
    -ClientNumber           128
    -ClientProcessID        4001
    -ClientProcessName      /large/opt/nfast/bin/ppmk
    -ObjectCountTotal       1
    +PerModule:
        +#1:
            -ObjectCount            1
+#129:
    -Uptime                 90
    -CmdCount               18
    -CmdBytes               452
    -CmdMarshalErrors       0
    -ReplyCount             18
    -ReplyBytes             1156
```

```
                -ReplyMarshalErrors    0

                -DevOutstanding        0

                -QOutstanding          0

                -LongOutstanding       0

                -RemoteIPAddr          (local)

                -ClientNumber          129

                -ClientProcessID       4001

                -ClientProcessName     /large/opt/nfast/bin/ppmk

                -ObjectCountTotal      1

                +PerModule:

                    +#1:

                        -ObjectCount          1

            +#133:

                -Uptime                0

                -CmdCount              55

                -CmdBytes              2140

                -CmdMarshalErrors      0

                -ReplyCount            54

                -ReplyBytes            4896

                -ReplyMarshalErrors    0

                -DevOutstanding        0

                -QOutstanding          1

                -LongOutstanding       0

                -RemoteIPAddr          (local)

                -ClientNumber          133

                -ClientProcessID       0

                -ClientProcessName

                -ObjectCountTotal      0

                +PerModule:

                    +#1:

                        -ObjectCount          0

        +PerModule:

            +#1:

                +ModuleObjStats:

                    -ObjectCount           22

                    -ObjectsCreated        199

                    -ObjectsDestroyed      177

                +ModuleEnvStats:

                    -MemTotal              62398464

                    -MemAllocKernel        798720

                    -MemAllocUser          0

                    -CurrentTempC          56.00

                    -MaxTempC              56.50

                    -MinTempC              55.00

                +ModuleJobStats:
```

```
        -CmdCount                6269
        -ReplyCount              6267
        -CmdBytes                248948
        -ReplyBytes              278404
        -HostWriteCount          6245
        -HostWriteErrors         0
        -HostReadCount           12272
        -HostReadErrors          0
        -HostReadEmpty           0
        -HostReadDeferred        6096
        -HostReadTerminated      0
        -PFNIssued               2905
        -PFNRejected             0
        -PFNCompleted            2904
        -ANIssued                7
        -CPULoadPercent          4
      +ModulePCIStats:
        -HostIRQs                12426
        -HostReadCount           6179
        -HostReadDeferred        6098
        -HostReadReconnect       6097
        -HostReadErrors          0
        -HostReadPushedDMA       21
        -HostReadPushedPIO       6155
        -HostWriteCount          6247
        -HostWriteErrors         0
        -HostDebugIRQs           0
        -HostUnhandledIRQs       0
        -HostKernReadCount       0
        -HostKernReadDeferred 0
        -HostKernReadReconnect 0
        -HostKernReadErrors      0
        -HostKernReadPushedDMA 0
        -HostKernReadPushedPIO 0
        -HostKernWriteCount      0
        -HostKernWriteErrors     0
    +RemoteServers:
    +PerDevice:
      +#1:
        +ModuleDriverStats:
        -DriverIRQs              12574
        -DriverReadIRQs          6253
        -DriverWriteIRQs         6328
        -DriverWriteFails        0
        -DriverWriteBlocks       6328
```

```
        -DriverWriteBytes      308552

        -DriverReadFails       0

        -DriverReadBlocks      4

        -DriverReadBytes       32768

        -DriverEnsureFail      0

        -DriverEnsure          6254

    +ModuleServerStats:

        -JobsOutstanding       1

        -LongJobsOutstanding   0

        -CmdCount              6351

        -ReplyCount            6350

    SUCCESS: show HSM statistics
```

slots

The `slots` command displays the current configuration of the HSM slots. There are two slots with corresponding tokens.

The examples below show the output for the Vormetric V5800 appliance and the Vormetric V6100 appliance.

Syntax

```
    slots
```

Example 1: V5800 appliance

In this example, the first token is labeled "dse" and the last is "AdminToken".

```
    0016:hsm$ slots
    ProtectToolkit C Status Utility $Revision: 1.1.1.2 $
    Copyright (c) Safenet, Inc. 2009
    ShowAllSlots:2 slots, 2 with tokens
    Slot ID 0
        Description      : ProtectServer Gold:44743
        Manufacturer     : SafeNet Inc.
        Hardware Version : 66.00
        Firmware Version : 3.00
    Token for Slot ID 0
        Label            : dse
        Manufacturer     : SafeNet Inc.
        Model            : KSE:PL220
        Serial Number    : 92404255:44743
        Hardware Version : 66.00
        Firmware Version : 3.00
    Slot ID 1
        Description      : ProtectServer Gold
        Manufacturer     : SafeNet Inc.
        Hardware Version : 66.00
        Firmware Version : 3.00
    Token for Slot ID 1
```

```
      Label           : AdminToken (92404255)

      Manufacturer    : SafeNet Inc.

      Model           : KSE:PL220

      Serial Number   : 92404255

      Hardware Version : 66.00

      Firmware Version : 3.00

  SUCCESS: show slots information
```

Example 2: V6100 appliance

```
  0002:hsm$ slots
  PKCS#11 library CK_INFO
        interface version 2.01
                   flags 0
          manufacturerID "nCipher Corp. Ltd              "
       libraryDescription "nCipher PKCS#11 2.2.1          "
    implementation version 2.02
  slots[0] CK_SLOT_INFO
         slotDescription "
          "
          manufacturerID "nCipher Corp. Ltd              "
                   flags 5
                   flags & CKF_TOKEN_PRESENT
                   flags & CKF_HW_SLOT
          hardware version 0.00
          firmware version 0.00
  slots[0] CK_TOKEN_INFO
                   label "loadshared accelerator        "
          manufacturerID "nCipher Corp. Ltd              "
                   model "                  "
          serialNumber "                  "
                   flags 201
                   flags & CKF_RNG
                   flags & CKF_DUAL_CRYPTO_OPERATIONS
       ulMaxSessionCount CK_EFFECTIVELY_INFINITE
     ulMaxRwSessionCount CK_EFFECTIVELY_INFINITE
             ulMaxPinLen 18446744073709551615
             ulMinPinLen 0
      ulTotalPublicMemory CK_UNAVAILABLE_INFORMATION
       ulFreePublicMemory CK_UNAVAILABLE_INFORMATION
     ulTotalPrivateMemory CK_UNAVAILABLE_INFORMATION
      ulFreePrivateMemory CK_UNAVAILABLE_INFORMATION
          hardware version 0.00
          firmware version 0.00
                   utcTime "                  "
  slots[1] CK_SLOT_INFO
```

```
              slotDescription "1
               "
            manufacturerID "nCipher Corp. Ltd                "
                      flags 1
                      flags & CKF_TOKEN_PRESENT
           hardware version 0.00
           firmware version 0.00
    slots[1] CK_TOKEN_INFO
                      label "1                              "
            manufacturerID "nCipher Corp. Ltd                "
                      model "                  "
               serialNumber "c0b3cb0fa3acb1e8"
                      flags 20D
                      flags & CKF_RNG
                      flags & CKF_LOGIN_REQUIRED
                      flags & CKF_USER_PIN_INITIALIZED
                      flags & CKF_DUAL_CRYPTO_OPERATIONS
          ulMaxSessionCount CK_EFFECTIVELY_INFINITE
        ulMaxRwSessionCount CK_EFFECTIVELY_INFINITE
                 ulMaxPinLen 18446744073709551615
                 ulMinPinLen 0
         ulTotalPublicMemory CK_UNAVAILABLE_INFORMATION
          ulFreePublicMemory CK_UNAVAILABLE_INFORMATION
        ulTotalPrivateMemory CK_UNAVAILABLE_INFORMATION
         ulFreePrivateMemory CK_UNAVAILABLE_INFORMATION
           hardware version 0.00
           firmware version 0.00
                    utcTime "                  "
    SUCCESS: show slots information
```

replaceacs

The `replaceacs` command is used to replace Administrative Card Set (ACS). This command was introduced with the release of the V6100 appliance.

The ACS needs to be replaced in case any cards have been compromised, or if cards have been lost, or if a card or cards have been damaged. This command should be used with care, as any data that has been backed up using the card set that is to be replaced, cannot be restored with the replacement card set. While creating the new card set you will be prompted to optionally erase your old card set after you create the replacement card set, you should keep the old ACS if you want to restore a backup made with the old ACS.

Syntax

```
replaceacs
```

Example

```
0000:dsm$ hsm
0001:hsm$ replaceacs
```

```
Before you start to replace an ACS, you must ensure that you have enough blank cards to
create a complete new ACS.

If you start the procedure without enough cards, you will have to cancel the procedure
part way through.

Have a quorum of cards from the current ACS ready. Please wait...

Insert ACS to authorise ACS replacement:
 Module 1: 0 cards of 1 read
 Module 1 slot 0: empty
 Module 1: completed.
Writing new ACS:
 Module 1: 0 cards of 2 written
 Module 1 slot 0: Admin Card #2
 Module 1 slot 0: empty
 Module 1 slot 0: unknown card
 Module 1 slot 0:- no passphrase specified - overwriting card
 Module 1: 1 card of 2 written
 Module 1 slot 0: remove already-written card #1
 Module 1 slot 0: empty
 Module 1 slot 0: unknown card
 Module 1 slot 0:- no passphrase specified - overwriting card
Card writing complete.
admin cardset replaced; you should erase the old admin cards
Examining/erasing card(s):
 Module 1 slot 0: Admin Card #2
 Module 1 slot 0: empty
 Module 1 slot 0: unknown card
 Module 1 slot 0:- confirmed - erasing card
 Module 1 slot 0: unknown card: Card erased
 Module 1 slot 0: blank card
 Module 1 slot 0: empty
 Module 1 slot 0: unknown card
 Module 1 slot 0:- confirmed - erasing card
 Module 1 slot 0: unknown card: Card erased
 Module 1 slot 0: blank card
Card erasure complete.
SUCCESS: ACS replaced.
```

test

The `test` command is used to test the nShield Connect network HSM hardware. This command was introduced with
the release of the V6100 appliance.

Syntax

```
test
```

Example

```
0003:hsm$ test
```

```
Running a universal statistical test on random numbers returned by the module:

Initialising 6-bit Universal Statistical Test

Processing...

46K

6-bit test results: Xu is 5.221571 = 0.003865 from mean (0.582749 s.d.'s)

Checking modulo exponentiations performed on the module against the test data located in
the /opt/nfast/testdata directory:

File testdata-misc:

16 completed OK

File testdata-s1024-n1:

1 completed OK

File testdata-s1024-n100:

100 completed OK

File testdata-s2048-n1:

1 completed OK

File testdata-s2048-n100:

100 completed OK

File testdata-s3072-n1:

1 completed OK

File testdata-s3072-n9:

9 completed OK

File testdata-s512-n1000:

1000 completed OK

File testdata-s512-n300:

300 completed OK

File testdata-vlong:

105 completed OK

File testdata-vshort:

122 completed OK

Performing hardware speed-testing by using modular exponentiation:

 Hardware module #1 speed index 552 recommended minimum queue 19

Found 1 module; using 19 jobs

 1,          90 36, 90 overall
 2,         622 234.4, 311 overall
 3,        1151 352.24, 383.667 overall
 4,        1674 420.544, 418.5 overall
 5,        2225 472.726, 445 overall
 6,        2760 497.636, 460 overall
 7,        3305 516.582, 472.143 overall
 8,        3854 529.549, 481.75 overall
 9,        4392 532.929, 488 overall
 10,        4932 535.758, 493.2 overall
 11,        5464 534.255, 496.727 overall
 12,        6005 536.953, 500.417 overall
 13,        6529 531.772, 502.231 overall
```

```
14,      7046 525.863, 503.286 overall
15,      7599 536.718, 506.6 overall
16,      8124 532.031, 507.75 overall
17,      8678 540.818, 510.471 overall
18,      9208 536.491, 511.556 overall
19,      9722 527.495, 511.684 overall
20,      10238 522.897, 511.9 overall
SUCCESS: All HSM hardware tests passed.
0004:hsm$
```

remoteadmin

The `remoteadmin` command enables remote administration of the DSM V6100 appliance and eliminates the need for administrators to be physically present in the lab to administer the DSM. Smart cards can be presented remotely from a laptop or PC. This feature must first be enabled on the CLI before it can be used to administer the DSM appliance. Refer to the *DSM Installation and Configuration Guide* for details about this feature.

Syntax

```
remoteadmin [ show | on | off ]
```

Example

Display remote administration status:

```
0001:hsm$ show
HSM is disabled
SUCCESS show hsm
0002:hsm$
```

To enable the feature:

```
0002:hsm$ remoteadmin on
HSM remote administration is enabled.
SUCCESS: remoteadmin command ran successfully.
```

connect

The `connect` command is used to enable or disable a network HSM for a DSM appliance that does not have a built-in HSM—DSM V6000, and the virtual appliance.

Syntax

```
connect [ add nShield Connect IP address RFS IP address | delete | show ]
```

where,

nShield Connect IP address is the IP address of the nShield Connect appliance

RFS IP address is the IP address of the computer that has the RFS installed

The connect command is used as follows:

Table 27-26: HSM Category connect command

add	Add a nShield Connect Network HSM

Table 27-26: HSM Category connect command (continued)

delete	Delete a nShield Connect Network HSM
show	Show currently configured nShield Connect Network HSMs

Example

```
0001:hsm$ connect add 1.2.3.16 1.2.3.4
```

This DSM is being connected to a nShield Connect for the first time(i.e. it is being converted into HSM enabled). A new DSM master key in the HSM will replace the existing master key in the Java keystore. Once that is done, this DSM cannot be converted back to non-HSM enabled without all the data being destroyed with 'config load default' to reset it back to factory configuration.

An administrator card from the ACS of the Security World the nShield Connect belongs to is required if the Security World is FIPS 140-2 level 3 compliant. If you don't have the administrator card currently, you need to abort now.

Do you want to continue? (yes|no)[no]: yes

Please remove the administrator card from the reader.

Stopping the Security Server

Stopping the data store [OK]

Self test in progress: passed

Starting Security Server

Security Server started in compatible mode

nShield Connect HSM with IP address 1.2.3.16 is added successfully.

SUCCESS: connect command ran successfully

```
0002:hsm$
```

secworldupdate

The `secworldupdate` command is used to synchronize the DSM with the nShield Connect when the Security World on the configured nShield Connect appliance has been upgraded. A Security World update may be triggered for various reasons, for example the ACS has been replaced. If the DSM is in an HA cluster, the command must be run all cluster nodes.

Syntax

```
secworldupdate
```

Example

```
0001:hsm$ secworldupdate
SUCCESS: Security World data on this DSM node updated
0002:hsm$
```

Luna

The `Luna` command is used to connect and register a DSM to a Password-authenticated Luna or a PED-authenticated Luna.

Syntax

```
luna [ add [HSMHOST] | delete [HSMHOST]| show | copass ]
```

where,

HSMHOST is the IP address, or host name, of the Luna HSM.

Table 27-27: Luna commands

add	Add a Luna HSM to a DSM or DSM HA cluster
delete	Delete a Luna HSM.
show	Show currently configured Luna HSMs
copass	Update the DSM with the Luna partition password //PED-authenticated Luna only

Add

Example

```
0015:hsm$ luna add 192.168.59.214
WARNING: Converting to an HSM configured setup will delete current data!
Continue? (yes|no) [no]:yes
192.168.59.214: UZHnfG5tTURxZ8etW0VQHitywmiN5H8NgObKdF20j/M
The fingerprint above should match the RSA output of the
'sysconf fingerprint ssh' lunash command on the Luna SA
Enter HSM administrator username at 192.168.59.214: lunaadmin
Enter password for lunaadmin@192.168.59.214:
Enter HSM partition ID number: 1396098836034
Enter partition crypto officer password:  //If using a Luna PED model
Enter hostname or IP address of this DSM: 192.168.15.105
WARNING: All Peer node and agent certificates will need to be re-signed after CA and
server certificates are regenerated! The security server software will be restarted
automatically.
Continue? (yes|no) [no]:yes
This node may have multiple IP addresses. All of the agents will have to connect to the
Security Server using the same IP.
Enter the host name of this node. This will be used by Agents to talk to this Security
Server.
This Security Server host name[dsm51015.i.vormetric.com]:
Please enter the following information for key and certificate generation.
What is the name of your organizational unit? []:Tech Pubs
What is the name of your organization? []:Thales
What is the name of your City or Locality? []:San Jose
What is the name of your State or Province? []:CA
What is your two-letter country code? [US]:
What is your email address? []:hpotter@bulleit.com
What is the validity period of the generated certificate (from 2 to 10 years)? [10]:
Regenerating the CA and server certificates now...
Please wait ......\
```

Delete

Example

```
0015:hsm$ luna delete 192.168.59.214
SUCCESS: deleted HSM host
```

> **Note**
> You cannot delete the last Luna from the DSM. This causes the DSM to stop functioning properly. You must do a config reset to remove the last Luna.

Show

Example

```
0016:hsm$ luna show
        HA auto recovery:  enabled
        HA recovery mode:  activeEnhanced
        Maximum auto recovery retry:  500
        Auto recovery poll interval:  60 seconds
        HA logging:  disabled
        Only Show HA Slots:  yes


        HA Group Label:  g30606
        HA Group Number:  11394399181044
        HA Group Slot ID:  4
        Synchronization: enabled
        Group Members:  1394399181044
        Needs sync:  no
        Standby Members:  <none>
        Slot #    Member S/N                      Member Label     Status
        ======    ==========                      ============     ======
        ------   1394399181044                    dsmpartition0     alive
        1394399181044@10.10.59.214
        SUCCESS: Showed HSM configuration
```

Copass

The copass command updates the PED-authenticated Luna crypto officer's partition password. In the Password-authenticated Luna, a DSM can only register to an empty, unconfigured partition. It registers to the partition, creates a random password for the partition, and stores it on the DSM. Only that one DSM (or DSM cluster) can access that partition. In the PED-authenticated Luna, multiple DSM clusters can register to the same partition. Therefore, the password is now user-controlled because all subsequent DSMs to register to that partition will need that password.

> **Note**
> If you change the partition password directly on the Luna, then you **must** update the partition password in the DSM using the `Luna copass` command.

Example

```
0007:hsm$ luna copass
```

```
WARNING: This command updates this DSM with the partition crypto officer password.
The Security Server will be restarted.
Continue? (yes|no) [no]: yes
Enter partition crypto officer password:
Success: DSM updated with crypto officer password. Security service will now restart
```

Maintenance Category Commands

> **Note**
>
> Most of these DSM CLI commands work in a DSM configuration that runs on a DSM appliance.
>
> Software-only DSM installations do not support the `config load`, `date`, `diag`, `gmttimezone`, `ntpdate`, and `time` commands. If you have a software-only installation, use the regular operating system commands to configure these operating system components.

The maintenance category is used to restore the DSM to factory defaults, upgrade the current DSM installation, and, on appliance-based DSMs, set operating system attributes, such as date, time, and time zone.

Enter the maintenance category by typing:

```
0009:dsm$ maintenance
0010:maintenance$
```

The maintenance category consists of the following commands:

Table 27-28: CLI maintenance category commands

`config`	This command restores the appliance image and configuration to the same state and version at which the appliance was shipped from the factory. On a software-only DSM, this command removes just the data that was added after the software was installed, leaving the current DSM version intact but without policies, keys, hosts, and so on. On both appliance and software-only DSM deployments, this command also configures automatic DSM backup and remote archival.
`showver`	Displays the DSM versions that are on the system and indicates the version that is currently running.
`delver`	Deletes a DSM image from the DSM or the system hosting the software-only DSM.
`ntpdate`	Configures one or more Network Time Protocol (NTP) servers with which to synchronize the system clock. DSM Appliance only.
`date`	Sets the system date. DSM appliance only.
`time`	Sets the system time. DSM appliance only.
`gmttimezone`	Sets the system time zone. DSM appliance only.
`diag`	Displays DSM and system logs, available system disk space, system OS version, and system uptime. DSM appliance only.
`repair`	System configuration recovery

config

This section describes using the DSM CLI maintenance category `config` command to delete the DSM configuration or restore the DSM to its original factory configuration.

The `config reset` and `config load default` commands restore the DSM installation to an unconfigured state.

The `config reset` command removes all the configuration data that was added after the current DSM software was installed. The command preserves the currently installed DSM software, but removes all configuration data except the network configuration and the original Security World configuration on V6100 appliances and the DSMs configured with nShield Connect HSM. When the command is run on a V6100 appliance, the option to recover the original Security World is available as the first choice. Alternatively, you could choose to destroy the old security world and create a new one using new cards.

> **Note**
> The config load default command causes a reboot of the DSM. When the command is issued, wait for the system to fully reboot and restart. Do not cycle power. Critical system files are installed on the reboot following a config load default so it is important to wait until it has completed

The `config load default` command is an extreme form of `config reset`. The `config load default` command deletes everything from the appliance and restores the same DSM installation with which the appliance was shipped. The partitions that contain DSM installations are deleted from the appliance, so there is no hope of retrieving any data once this command is executed. The `config load default` command is available on appliance-based DSMs only.

The `config load default` command produces the same result as the Kill switch:

- Execute this DSM CLI command when there is a serial console connection to the appliance, or when there is a terminal window from which to SSH onto the appliance.

- Press the Kill Switch when you do not have console or terminal access, but you do have access to the physical appliance.

The `config reset` and `config load default` commands discard the CA signer certificate. The CA signer certificates must be restored from a backup or regenerated later.

Syntax

The following is the DSM CLI `config` command syntax:

```
config load default
config reset
```

The `config` command supports the following arguments:

Table 27-29: Maintenance category config commands

load	Removes the current DSM installation and restores the manufacturer default DSM installation. It removes all policies, hosts, keys, and so on from the DSM.
reset	Preserves the currently installed DSM software but removes all configuration data, except the network configuration and, in the case of the V6100 appliance, the original Security World.

The following example deletes the partitions that contain DSM installations. Everything that has been added since the appliance was first started is deleted. The appliance reboots and loads the original DSM installation.

> **Note**
> The original DSM installation will not contain any patches or upgrades.

Example 1

```
0011:maintenance$ config load default

Loading manufacture default will wipe out all the configuration data and set the machine
configuration to the manufacture default. System will reboot automatically.

Continue? (yes|no)[no]:yes
```

```
config load SUCCESS.
0012:maintenance$
```

Example 2

The following example resets the current DSM installation back to its initial unconfigured state. This command returns database and configuration files to their original, fresh installation state without changing or reinstalling the current DSM version.

```
0003:maintenance$ config reset

Reset configuration will wipe out all the configuration data and set the configuration
data to the manufacture default. System will reboot automatically.

Continue? (yes|no)[no]:yes

config reset SUCCESS. You can reboot the Security Server now or it will reboot
automatically in 60 seconds.
```

Retrieving Security World

To retrieve the original Security World after running the config reset command, do the following:

To reset the current DSM installation to its initial unconfigured state—network configuration remains intact—and retrieve the original security world, do the following;

1. Log on to the DSM CLI console using the CLI Administrator credentials and at the prompt type `maintenance` to enter the maintenance category of commands, then type `config reset`;

```
0000:dsm$ maintenance

0001:maintenance$ config reset

Reset configuration will wipe out all the configuration data and set the configuration
data to the manufacture default. System will reboot automatically.

Continue? (yes|no)[no]:yes

config reset SUCCESS. You can reboot the Security Server now or it will reboot
automatically in 60 seconds.

0002:maintenance$
```

2. You will need to generate the DSM certificate authority (CA) again. This will require a quorum from the original ACS used to create the Security World. Wait until the system has rebooted and the `vormetric$` prompt is displayed, then run the `security genca` command.

   Type `up` at the prompt to return to the main menu and then type `system` to access the System category sub-menu;

```
0002:maintenance$ up

0003:dsm$ system
```

At the prompt, type `security genca` to generate the CA. A warning message is displayed informing you that all agent and peer node certificates will need to be resigned after the new certificate authority is created and that the DSM software will be restarted, type 'yes' to generate the certificate;

```
0004:system$ security genca

WARNING: All Agents and Peer node certificates will need to be re-signed after CA and
server certificate regenerated, and the security server software will be restarted
automatically!

Continue? (yes|no)[no]:yes
```

3. The following message is displayed. Read it, enter the required information to generate the CA, and ensure the DSM host name is correct, press enter:

```
This computer may have multiple IP addresses. All the agents will have to connect to
Security Server using same IP.

Enter the host name of this computer. This will be used by Agents to talk to this
Security Server.

This Security Server host name[mycompany.com]:

Please enter the following information for key and certificate generation.

What is the name of your organizational unit? []:

What is the name of your organization? []:

What is the name of your City or Locality? []:

What is the name of your State or Province? []:

What is your two-letter country code? [US]:

What is your email address? []:

What is the validity period of the generated certificate (from 2 to 10 years)? [10]:

Regenerating the CA and server certificates now...
```

4. A message is displayed informing you that a Security World exists and provides you the option to reuse it. This is the Security World in place when the `config reset` command was run. Type `yes` to reuse the existing Security World. Keep the quorum of cards from the ACS used to create this original Security World available as these will be required for the next few steps. Follow the instructions on the screen;

```
There is an existing Security World.

Would you like to reuse it? (yes|no)[no]: yes

Please provide a quorum of cards from the ACS of the existing Security World.

15:47:02 WARNING: Module #1: preemptively erasing module to see its slots!

Indoctrinating Module:
 Module 1: 0 cards of 1 read
 Module 1 slot 0: empty
 Module 1 slot 2: empty
 Module 1 slot 3: empty
Card reading complete.

security world loaded on 1 module; hknso = 3546197a6456c5e3bfb28d7facd063072b7a8f52

Do NOT remove the smart card from the reader yet.

Creating CA keys and signer certificates...

done.

Generating server private key...

done.

You may now remove the smart card from the reader.

Creating and signing the server certificates...

done.

CA and Server certificates have been generated successfully.

JBoss vault keystore password have been completed successfully.

Self test in progress: passed

SUCCESS: The CA and security certificates are re-generated and the Security Server
software is restarted.
```

The DSM appliance is now ready to use with the original Security World. For procedures to restore a backup of your previous configuration, see Chapter 12: "Backing Up and Restoring" on page 94.

To reset the current DSM installation to its initial unconfigured state—network configuration remains intact—and create a new Security World:

1. Log on to the DSM CLI console using the CLI Administrator credentials and at the prompt type `maintenance` to enter the maintenance category of commands, then type `config reset`;

```
0000:dsm$ maintenance

0001:maintenance$ config reset

Reset configuration will wipe out all the configuration data and set the configuration
data to the manufacture default. System will reboot automatically.

Continue? (yes|no)[no]:yes

config reset SUCCESS. You can reboot the Security Server now or it will reboot
automatically in 60 seconds.

0002:maintenance$
```

2. You will need to generate the DSM certificate authority (CA) again. This will require a quorum from the original ACS used to create the Security World. Wait until the system has rebooted and the `vormetric$` prompt is displayed, then run the `security genca` command.

Type `up` at the prompt to return to the main menu and then type `system` to access the System category sub-menu;

```
0002:maintenance$ up

0003:dsm$ system
```

At the prompt, type `security genca` to generate the CA. A warning message is displayed informing you that all agent and peer node certificates will need to be resigned after the new certificate authority is created and that the DSM software will be restarted, type 'yes' to generate the certificate;

```
0004:system$ security genca

WARNING: All Agents and Peer node certificates will need to be re-signed after CA and
server certificate regenerated, and the security server software will be restarted
automatically!

Continue? (yes|no)[no]:yes
```

3. The following message is displayed. Read it, enter the required information to generate the CA, and ensure the DSM host name is correct, press enter:

```
This computer may have multiple IP addresses. All the agents will have to connect to
Security Server using same IP.

Enter the host name of this computer. This will be used by Agents to talk to this
Security Server.

This Security Server host name[mycompany.com]:

Please enter the following information for key and certificate generation.

What is the name of your organizational unit? []:

What is the name of your organization? []:

What is the name of your City or Locality? []:

What is the name of your State or Province? []:

What is your two-letter country code? [US]:

What is your email address? []:

What is the validity period of the generated certificate (from 2 to 10 years)? [10]:

Regenerating the CA and server certificates now...
```

4. A message is displayed informing you that a Security World exists and provides you the option to reuse it. This is the Security World in place at the when the `config reset` command was run. Type no to destroy the old Security World and create a new Security World. Have the new set of cards available for this step. For information about the ACS and best practices, refer to the Administrative Card Set section in chapter 1 of the *DSM Installation and Configuration Guide*. Follow the instructions on the screen;

```
There is an existing Security World.

Would you like to reuse it? (yes|no)[no]: no

Enter the total number of cards (N) you would like to use in your Administrator Card Set
(ACS).

Note: The system can handle at most 64 cards but you should not enter more than the
number of available cards in your possession currently.

This value must be at least 2 and no higher than 64: 2

Enter the number of cards (K) required to authorize an action. This number K is known as
the quorum.

Note: The value for K must be less than N. Creating card sets in which K is equal to N
is not allowed because an error on one card would render the whole card set unusable.

This value must be at least 1 and less than 2: 1

17:10:29 WARNING: Module #1: preemptively erasing module to see its slots!

Create Security World:
 Module 1: 0 cards of 2 written

 Module 1 slot 0: empty

 Module 1 slot 2: empty

 Module 1 slot 3: empty

 Module 1 slot 4: empty

 Module 1 slot 5: empty

 .

 .

 .

Module 1 slot 2:- no passphrase specified - overwriting card

 Module 1: 1 card of 2 written

 Module 1 slot 2: remove already-written card #1

 Module 1 slot 2: empty

 Module 1 slot 2: unknown card

 Module 1 slot 2:- no passphrase specified - overwriting card

Card writing complete.

security world generated on module #0; hknso = 781eb7d9ae3abad631bfc4c7487279eadbcef4a8

Do NOT remove the smart card from the reader yet.

Creating CA keys and signer certificates...

done.

Generating server private key...

done.

You may now remove the smart card from the reader.

Creating and signing the server certificates...

done.

CA and Server certificates have been generated successfully.

JBoss vault keystore password have been completed successfully.

Self test in progress: passed
```

```
SUCCESS: The CA and security certificates are re-generated and the Security Server
software is restarted.
```

The DSM appliance is now ready to use with the new Security World. For procedures to restore a backup of your previous configuration, see Chapter 12: "Backing Up and Restoring" on page 94.

> ⚠️ **CAUTION**
> **Restoring a backup of the previous configuration will restore the old Security World and the new one just created will be destroyed.**

showver

The `showver` command displays the DSM software images that have been uploaded and that are available for use. The DSM comes from the factory with one image pre-installed. The software-only DSM installation utility installs one image. Up to two images can be installed and configured at one time. Software patches are not displayed by this or any other command.

Syntax

```
showver
```

Example

Enter the `showver` command without any arguments to display the current image. For example:

```
0001:maintenance$ showver
ver_count=1
cur_ver=6.0
show version SUCCESS
0002:maintenance$
```

delver

The `delver` command deletes the inactive DSM image from the system. Up to two images can be installed on the appliance or system at one time. As image updates become available, you will cycle through the installed images, usually deleting the older of the two images. You cannot delete an active image, nor can you simply stop the DSM.

Syntax

```
delver
```

You are prompted to continue.

Example

The following example deletes the image from the system.

```
0010:maintenance$ delver
You are deleting the alternative software version.
Continue? (yes|no)[no]:yes
Delete version SUCCESS
0011:maintenance$
```

ntpdate

The `ntpdate` command:

- Configures one to four Network Time Protocol (NTP) servers for the current DSM

- Enables and disables NTP on the appliance

- Forces immediate clock synchronization with an NTP server

- Shows the current NTP configuration status

When NTP is configured and enabled, at one hour intervals the DSM CLI daemon synchronizes the system clock of the DSM with the first available NTP server. If, within one second, the DSM cannot connect with the NTP server, the DSM CLI daemon tries the next NTP server in the list. The NTP server can reside in any time zone.

Syntax

```
ntpdate sync | add SERVER_ADDRESS | delete SERVER_ADDRESS | on | off | show
```

The `ntpdate` command takes the following arguments:

Table 27-30: ntpdate command arguments

`sync`	forces clock synchronization with the first available NTP server.
`add SERVER_ADDRESS`	adds the named NTP server to the list of servers to contact for time synchronization. At least one server must be configured before you can enable (turn on) time synchronization. You may configure up to four NTP servers.
`delete SERVER_ADDRESS`	removes the named NTP server from the list of servers to contact for time synchronization. Time synchronization is disabled (turned off) when the last NTP server is removed from the list.
`on`	enables NTP time synchronization. At least one NTP server must be configured before you can enable synchronization.
`off`	disables time synchronization and leaves the current NTP server list intact. You can re-enable synchronization without having to reconfigure the NTP servers.
`show`	Displays the NTP server configuration and state. The `ntpdate show` command does not sort the output. It displays all the configured NTP servers in the same order that they were added to the DSM.

`sync`, `add`, `delete`, `on`, `off`, and `show` are literals that are entered as shown or in abbreviated form.

Examples

The following examples:

- displays the default NTP configuration environment

- adds four NTP servers

- enables NTP synchronization

- displays a fully-configured NTP environment

- synchronizes the appliance clock with the first available NTP server clock

- swaps the last two NTP servers in the list to change access order

```
0001:maintenance$ ntpdate show
Total ntpdate server number  : 0
ntpdate is off

ntpdate SUCCESS

0007:maintenance$ ntpdate add 172.16.78.110

ntpdate SUCCESS

0008:maintenance$ ntpdate add search.domain.com
```

```
ntpdate SUCCESS
0009:maintenance$ ntpdate add 172.30.45.115
ntpdate SUCCESS
0010:maintenance$ ntpdate add 172.20.244.75
ntpdate SUCCESS
0011:maintenance$ ntpdate on
ntpdate SUCCESS
0012:maintenance$ ntpdate show
Total ntpdate server number  : 4
ntpdate server [1] : 172.16.78.100
ntpdate server [2] : search.domain.com
ntpdate server [3] : 172.30.45.115
ntpdate server [4] : 172.20.244.75
ntpdate is on
ntpdate SUCCESS
0013:maintenance$ ntpdate sync
ntpdate SUCCESS
0014:maintenance$ ntpdate delete 172.16.78.100
ntpdate SUCCESS
0015:maintenance$ ntpdate add 172.16.78.100
ntpdate SUCCESS
0016:maintenance$ ntpdate show
Total ntpdate server number  : 4
ntpdate server [1] : 172.30.78.100
ntpdate server [2] : search.domain.com
ntpdate server [3] : 172.20.244.75
ntpdate server [4] : 172.16.78.100
ntpdate is on
ntpdate SUCCESS
0017:maintenance$
```

date

The `date` command in the `maintenance` category is used to set or to display the date on the system. The `date` command without any arguments displays the current system date. If a parameter is included with the `date` command it resets the system date to the specified date.

Syntax

The syntax for the `date` command is:

```
date MM/DD/YYYY
date
```

Example

To set the date on the system to December 20th, 2014, enter the following:

```
0001:maintenance$ date 12/20/2014
```

The following example displays the system date:

```
0004:maintenance$ date
month=Dec day=20 year=2014
Show system date SUCCESS
0005:maintenance$
```

time

The `time` command sets or to displays the time on the system using a 24-hour clock. When no parameters accompany the `time` command, it displays the current system time. If a parameter is included with the time command, it resets the system time to the specified value.

Syntax

The syntax for the `time` command is:

```
time HH:MM:SS
time
```

Example

To set the time on the system enter the following:

```
0001:maintenance$ time 02:23:00
```

This sets the system to 2:23 AM.

The following example uses the time command to display the system time:

```
0003:maintenance$ time
hour=18 min=22 sec=38 zone=PDT
Show system time SUCCESS
```

gmttimezone

The `gmttimezone` command in the maintenance category is used to set the system time zone. If a parameter is included with the `gmttimezone` command, it sets the time to the zone specified. To see a list of supported time zones, enter `gmttimezone list`.

Syntax

The syntax for the `gmttimezone` command is:

```
gmttimezone {list|show|set zonename}
gmttimezone list
gmttimezone show
```

Example

To list and set the `gmttimezone` on the system enter the following:

```
0025:maintenance$ gmttimezone list
...
(GMT-07:00) America/Phoenix (Mountain Standard Time)
(GMT-07:00) America/Shiprock (Mountain Standard Time)
(GMT-07:00) America/Yellowknife (Mountain Standard Time)
(GMT-08:00) America/Dawson (Pacific Standard Time)
(GMT-08:00) America/Los_Angeles (Pacific Standard Time)
(GMT-08:00) America/Tijuana (Pacific Standard Time)
```

```
       (GMT-08:00) America/Vancouver (Pacific Standard Time)

       (GMT-08:00) America/Whitehorse (Pacific Standard Time)

       (GMT-08:00) Pacific/Pitcairn

       (GMT-09:00) America/Anchorage

       ...

       0026:maintenance$ gmttimezone show

       Timezone is set to : US/Pacific

       Show timezone SUCCESS

       0030:maintenance$ gmttimezone set America/Tijuana

       Set timezone SUCCESS

       0031:maintenance$ gmttimezone show

       Timezone is set to : America/Tijuana

       Show timezone SUCCESS
```

diag

> ⚠ **CAUTION**
>
> **When one hard drive fails on the DSM, the system will continue to operate using the second HDD.**
>
> **The DSM is a secure device and "hot-swapping" the HDD is not permitted. Replacing the failed HDD requires a reboot of the system and reconfiguration to restore the RAID capability again.**

The `diag` command in the `maintenance` category displays OS system information and related log files. This command is available in appliance-based installations only.

Table 27-31: Maintenance category diag command arguments

`diskusage`	Displays system disk space usage.
`hardware`	View RAID and motherboard status.
`log`	Lists and displays system messages and logs.
`osversion`	Displays the system kernel version.
`tlsmon`	Monitor TLS connections and generate audit logs.
`uptime`	Displays how long the system has been running since the last reboot, the current number of administrators logged into the system, and CPU load usage.
`vmstat`	Displays CPU and memory usage.

diskusage

The `diskusage` argument to the `diag` command displays information about the system disk, such as partitions, amount of used and available disk space, percentage of free space, and partition names.

Syntax

```
diag diskusage
```

Example

```
0017:maintenance$ diag diskusage
Filesystem    1M-blocks   Used Available Use% Mounted on
/dev/sda6     9389     4403     4510  50% /
/dev/sda9   254458     939   240594   1% /partitions/large
/dev/sda1      935      22      866   3% /grub
tmpfs         1963       0     1963   0% /dev/shm
/dev/sda2     7511    2307     4823  33% /partitions/std/2
/dev/sda8     7513     155     6977   3% /tmp
SUCCESS: Show disk usage
0018:maintenance$
```

hardware

The `hardware` argument to the `diag` command displays information about the RAID and motherboard status.

Syntax

```
diag hardware
```

Example

```
0004:maintenance$ diag hardware
--------------------------------------------------------
Motherboard Information
--------------------------------------------------------


--------------------------------------------------------
RAID Status Information
--------------------------------------------------------
Controller information
--------------------------------------------------------
  Controller type                   : SAS2008
  BIOS version                      : 7.17.00.00
  Firmware version                  : 9.00.00.00
  Channel description               : 1 Serial Attached SCSI
  Initiator ID                      : 0
  Maximum physical devices          : 255
  Concurrent commands supported     : 1952
  Slot                              : Unknown
  Segment                           : 0
  Bus                               : 6
  Device                            : 0
  Function                          : 0
  RAID Support                      : Yes
--------------------------------------------------------
IR Volume information
--------------------------------------------------------
IR volume 1
```

```
        Volume ID                                 : 323
```

log

The `log` argument to the `diag` command is used to list and view system files on the DSM.

Syntax

```
diag log list
diag log view <file>
diag log kmipdebug [show | on | off]
```

The `diag log` command supports three additional arguments, `list`, `view` and `kmipdebug`.

The `list` argument displays the system files that are available for viewing. It takes no additional input. The `view` argument takes the name of the log file to display. The `view` argument calls the "`more`" system command to display the file. Some of the "`more`" command display options are supported. The `kmipdebug` argument turns on or turns off logging of KMIP server SSL traffic, and indicates whether KMIP debugging is turned on or off, the default is off.

> **Note**
> As a best practice, we recommend that you turn off KMIP logging once you have accomplished your task as the log files fill up rapidly and could impact performance.

Active logs are log files that being currently written to and updated by DSM processes. Inactive logs are logs that have been filled to capacity and then closed. The name of the closed log file is the original name usually appended with the date. For example, the name of the active DSM log is `cgss.log`. When it reaches the configured capacity, it is made inactive and renamed to `cgss.log.`*YYYY-MM-DD*.

The DSM log files that you can view are described below:

- The `server.log` file contains details about agent backup and restore requests, connection status, Management Console interaction, Java exceptions, JBoss start and stop processes, and more. This file contains diverse information and it should be the first file you check for problems that are related to DSM operation.

- The `cgss.log` file contains a record of the events that make up the BEK generation process for an agent requesting to make a backup, as well as the names of uploaded audit files. This file does not contain events that pertain to restore operations. Check this file if the agent fails to back up a database, even though agent/DSM authentication is correctly configured, and the policy for this agent permits the backup operation.

- The `messages` file is generated by syslog. It contains kernel entries for enabling/disabling the log service, memory usage, CPU usage, system calls, device initialization, and so on. It also contains log entries that would otherwise be displayed in the Message Log but, for some reason, cannot be uploaded to the DSM. The `messages` file follows the standard naming convention to cycle large files. For example, the active file is `messages`. The cycled files are `messages.1`, `messages.2`, `messages.3`, and so on.

Example 1

To list the files that are available for viewing:

```
0011:maintenance$ diag log list
messages
messages.1
messages.2
messages.3
messages.4
cgss.log
cgss.log.2014-01-08
```

```
server.log

server.log.2014-01-15

SUCCESS: list log file

0012:maintenance$
```

Example 2

To display a log file, execute the `diag log view` command followed by the name of the file to view. For example:

```
0018:maintenance$ diag log view cgss.log

...

2014-01-19 19:09:22,025 INFO  [com.vormetric.server.sdk.user.UserManager] Entering
createUser

2014-01-19 19:09:22,025 INFO  [com.vormetric.server.sdk.user.UserManager] User is
authorized. Generating password...

2014-01-19 19:09:22,025 INFO  [com.vormetric.server.sdk.user.UserManager] Verifying
whether the password meets the PasswordPolicy conditions...

2014-01-19 19:09:22,027 INFO  [com.vormetric.server.sdk.user.UserManager] Verifying
complete: Password meets all the PasswordPolicy conditions

2014-01-19 19:09:22,032 INFO  [com.vormetric.server.sdk.user.UserManager] Password
generation complete

...

0019:maintenance$
```

Example 3

To turn on KMIP debugging run the `kmipdebug on` command followed by the `kmipdebug show` command to show whether it was successfully turned on. For example:

```
0010:maintenance$ diag log kmipdebug on

SUCCESS: KMIP debug logging turned on.

0011:maintenance$ diag log kmipdebug show

SUCCESS: KMIP debug mode is on
```

osversion

The `osversion` command displays the operating system version and kernel that is running on the DSM. The `osversion` command is equivalent to the Linux "`uname -a`" command.

Syntax

```
osversion
```

Example

To display the appliance operating system version and kernel:

```
0014:maintenance$ diag osversion

Linux SSA666 2.6.18-128.el5PAE #1 SMP Wed Jan 21 11:19:46 EST 2009 i686 i686 i386
GNU/Linux

SUCCESS: Show version
```

uptime

The `uptime` command displays the amount of time the operating system has been running since the last bootup. It also displays the system load and the number of DSM CLI administrators that are currently running CLI sessions on the DSM. Administrators that are configured in the DSM Management Console GUI are not included in the count

because DSM CLI administrators are actual system users and Management Console administrators exist only in the DSM database. The uptime command is equivalent to the Linux "`uptime`" command.

Syntax

```
uptime
```

Example

To display the amount of time that the appliance has been running, system load, and the number of current DSM CLI sessions:

```
0019:maintenance$ diag uptime
 17:02:20 up 3 days, 22:02,  2 users,  load average: 0.14, 0.06, 0.01
SUCCESS: Show uptime
0020:maintenance$
```

vmstat

The `vmstat` argument to the `diag` command displays information about the system disk, such as partitions, amount of used and available disk space, percentage of free space, and partition names.

Syntax

```
diag vmstat
```

Example

```
0001:maintenance$ diag vmstat
procs -----memory- --swap---io---system-----cpu-----
 r  b   swpd   free   buff  cache   si   so    bi    bo    in   cs us sy id wa st
 1  0      0 4984848 203208 3793404    0    0     0    12    10    2  0  0 99 00
 SUCCESS: Show vmstat result
0002:maintenance$
```

repair

vacuum

The vacuum feature compacts and defragments the postgres database.

Syntax

```
repair [vacuum | prunewal | purgepolicyhistory ]
```

Example

```
0003:maintenance$ repair vacuum
running vacuum will temporarily lock all tables.
Continue? (yes|no)[no]: yes
SUCCESS: vacuum db done
```

> **Note**
> Running the `repair vacuum` function is disruptive. Thales recommends that you perform this function when server traffic is light.

prunewal

If you have a large number of postgres transaction logs, they may potentially take up a large amount of storage space and impact performance. The prunewal command clears (or prunes) the processed transaction logs to improve performance.

Syntax

```
repair [vacuum | prunewal | purgepolicyhistory ]
```

Example

```
0002:maintenance$ repair prunewal

running prunewal...

Continue? (yes|no)[no]: yes

SUCCESS: prunewal done
```

purge policy history

Each time any modification to a policy results in a new policy version, the old version is saved. These older versions may potentially take up a large amount of storage space and impact performance. Purging the policy history deletes the old policy versions, thereby freeing up storage space and improving performance.

Syntax

```
repair [vacuum | prunewal | purgepolicyhistory ]
```

Example

```
0013:maintenance$ repair purgepolicyhistory

purge policy history will delete all of the archived history policy data and
modification logs.

Continue? (yes|no)[no]: yes

SUCCESS: purged all the policy history logs and data
```

High Availability Category Commands

High Availability (HA) is the configuration of multiple DSMs in an HA cluster. Each DSM HA node runs in parallel in an active-active format. When one DSM HA node fails, the other HA nodes in the HA cluster continue to run. When the downed DSM HA node is up and running again, it synchronizes with the other DSMA nodes.

After configuring and managing any DSM HA node, it synchronizes with the other nodes within seconds so that they remain mirror copies of each other.

HA configuration and maintenance is done in both the DSM CLI and the UI Management Console.

Table 27-32: HA category commands

add	Add an HA node to the HA cluster, without configuring it for replication.
cleanup	Cleanup HA configuration data on a local node, after it has been removed from an HA cluster.
join	Join this node to an HA cluster and configure it for replication.
remove (reassign host)	Remove a node from an HA cluster. If a host is assigned to the node, you must reassign hosts to another node before removing the node from the cluster. You cannot remove the node from the cluster while hosts are assigned to it.

Table 27-32: HA category commands (continued)

show	Lists all the DSM HA nodes in the HA cluster.

add

Add a node to the HA cluster without configuring it for replication. You must add the node before it can join the cluster. You can add the node from either the UI or CLI of node 1 in the cluster.

You must add the new node to the HA cluster, from node 1.

Syntax

The syntax for the command is:

```
add <haNodeName>
```

where, hostname is the host name or FQDN of an HA node.

Example

1. Log on to HA node 1.

2. Switch to the HA menu, type:

   ```
   0000:dsm$ ha
   ```

3. Add the HA node to the cluster, type:

   ```
   0003:ha$ add <haNodeName>
   ```

Example:

```
0003:ha$ add HAnode2.i.vormetric.com
```

System Response:

```
This command adds node 'HAnode2.i.vormetric.com' to the HA cluster (with
'HAnode1.i.vormetric.com'), without configuring it for replication.

To configure it for replication, login to the CLI of '192.168.15.96' and execute the
'join' command under 'ha'

Continue? (yes|no) [no]: yes

SUCCESS: Node added to HA cluster.

0004:ha$
```

cleanup

The `cleanup` command cleans up HA configuration data on the current HA node and restarts it.

> **Note**
> You must remove this node from the HA cluster before you can run this command. Perform a cleanup before you regenerate the node certificates and re-register the HA node with the HA cluster.

Syntax

```
cleanup
```

1. Log on to an HA node in the cluster, type:

   ```
   0002:ha$ cleanup
   ```

   System Response:

   ```
   WARNING: This command cleans up HA configuration data of, and restarts, this server.
   ```

```
        This may take several minutes.
        Continue? (yes|no)[no]:
```

2. Type **yes** to continue.

```
        SUCCESS: cleanup
```

join

Joins the current node to the HA cluster. If you are joining an HA cluster after an upgrade, and the node previously had a host assigned to it, after it successfully joins the cluster, the DSM asks if you want the host restored to the node.

Syntax

The syntax for the command is:

```
        join [longwait]
```

Example

To join an HA node to the HA cluster:

1. Log on to the node that you want to join to the cluster.

2. Switch to the HA menu, type:

```
0000:dsm$ ha
```

3. Join the node to the cluster. Type:

```
0000:dsm$ join
```

> **Note**
> Sometimes, when DSM nodes are spread far apart geographically, or are in a cloud environment, the Join function takes so long that the ssh session times out and terminates automatically before the Join can finish. If the Join function fails, type: **join longwait** to make the join command proceed in a 'longwait mode' (as opposed to the 'normalwait' mode). The difference between the two modes is the duration that it waits for the node replication status state to be set to 'ready.' Longwait waits for approximately twice as long as the normalwait mode.

System Response:

```
        WARNING: This server node is about to join an HA cluster.
        Please make sure the HA cluster is running and has this server node in its HA node list.
        This may take several minutes.
        Continue? (yes|no)[no]:yes
        HA Initial Server host name:HAnode1.i.vormetric.com
        Initial Security Server system administrator name: tesadmin
        Initial Security Server system administrator password:
        This node may have multiple IP addresses. All the agents will have to connect to
        Security Server using same IP.
        Enter the host name of this computer. This will be used by Agents to talk to this
        Security Server.
        This Security Server host name[HAnode2.i.vormetric.com]:
        Please enter the following information for key and certificate generation.
        What is the name of your organizational unit? []:TP
        What is the name of your organization? []:Thales
        What is the name of your City or Locality? []:SJ
```

```
What is the name of your State or Province? []:CA
What is your two-letter country code? [US]:US
What is your email address? []:
What is the validity period of the generated certificate (from 2 to 10 years)? [10]:
Restore original host assignment back to this node (yes|no)[yes]:
1 hosts now assigned to dsm15099.i.vormetric.com
```

remove

The remove command removes an HA node from an HA cluster. You must issue the remove command for the node from the initial node in the cluster. You cannot remove a node when you are logged into it.

If hosts are assigned to the node that you want to remove from the HA cluster, you must reassign the hosts to another node before removing the node from the cluster. You cannot remove a node from the cluster while hosts are assigned to it.

You can choose to move hosts to specific HA nodes by naming the host, or you can have the DSM evenly distribute the hosts to balance the load in the HA cluster.

To move the hosts to a specific HA node, type:

```
0001:ha$ remove <HA_Node1> reassignhost <HA_node2>
```

Example

```
0001:ha$ remove dsm15099.i.vormetric.com reassignhost dsm15100.i.vormetric.com
```

To have the DSM move the hosts to HA nodes and evenly distribute the load, type:

```
0001:ha$ remove <HA_Node1> reassignhost rr
```

Example

```
0001:ha$ remove dsm15099.i.vormetric.com reassignhost rr
```

System Response

```
WARNING: This command removes the given server node from the HA cluster. After deletion,
running the "ha cleanup" command from its CLI will be required.
This may take several minutes.
Continue? (yes|no)[no]:yes
Restore original host assignment back to this node (yes|no)[yes]:
1 hosts now assigned to <HA_Node1>
```

show

The `show` command displays all of the HA DSM nodes in the HA cluster. All of the HA nodes are listed, regardless if they are registered or not. You can execute the `show` command on any HA node in the cluster. The node that you are on is listed first in the list.

Syntax

```
show
```

Example

```
0001:ha$ show
```

System Response:

```
List of Nodes in the Cluster:
```

```
HAnode1.i.vormetric.com
HAnode2.i.vormetric.com
HAnode3.i.vormetric.com
HAnode4.i.vormetric.com
```

IPMI Category Commands

The Intelligent Platform Management Interface (IPMI) provides remote access to the V6000 and V6100 hardware appliances. It allows administrators to remotely monitor appliance health (example: temperature, power consumption, physical drive status, chassis intrusion), perform cold boots (power-off and power-on), and access the DSM CLI. IPMI is not supported on DSM hardware appliances earlier than V6000/V6100.

Table 27-33: IPMI category commands

`ip`	Set, delete, or show ip address of machine using IPMI.
`mask`	Set, delete, or show subnet mask of machine using IPMI.
`gateway`	Set, delete, or show gateway of machine
`disable`	Disable IPMI network
`user`	Add, show, or delete user. Also change password and privilege level.
`clearint`	Clear chassis intrusion. If chassis cover is open, a chassis intrusion toggle is set that shows up in the IMPI GUI sensor reading, CLI maintenance diagnostic, and DSM log. This command resets the toggle.
`reset`	Reset IPMI configuration to factory default.
`selftest`	Triggers a test called the IPMI Baseboard Management Controller (BMC) self-test. Use this if you suspect the IPMI controller is not working. The BMC chip runs even when the rest of the system is down.
`version`	Show IPMI version.
`psinfo`	Show power supply information.
`dhcp`	Configure DHCP for the IPMI NIC
`port`	Configure IPMI ports.

ip

The `ip` command configures the IPMI network IP address.

The `ip` command includes the following elements:

Table 27-34: IPMI category ip command elements

`set`	Sets the IPMI IP address.
`delete`	Deletes the IP address.
`show`	Shows the IP address.

ip set

Set the IPMI IP address. Setting the IP address automatically sets the mask to 16-bit mask.

Syntax

```
ip set ip_address
```

Example

The following example sets the IPMI network interface IP address.

- 0001:dsm$ ipmi
  0002:ipmi$ ip set 10.3.99.77
  IP=10.3.99.77
  SUCCESS: ip set

ip delete

Delete the IPMI IP address. This sets the IP to 0.0.0.0.

Syntax

```
ip delete
```

Example

The following example deletes the IPMI network interface IP address.

```
0002:ipmi$ ip delete
IP=0.0.0.0
SUCCESS: ip delete
```

ip show

Show the IPMI IP address.

Syntax

```
ip show
```

Example

The following example shows the IPMI network interface IP address.

```
0002:ipmi$ ip show
IP=0.0.0.0
MAC=00:25:90:F7:12:52
SUCCESS: ip show
```

mask

The `mask` command sets, deletes or shows the subnet mask for the IP address. The `mask` command includes the following elements:

Table 27-35: IPMI category `mask` command elements

set	Sets the subnet mask for the IP address.
delete	Sets the IP subnet mask 0.0.0.0.
show	Shows the subnet mask for the IP address.

mask set

Set the subnet mask for the IP address.

Syntax

```
mask set subnet_mask
```

Example

```
0001:dsm$ ipmi
0002:ipmi$ mask set 255.255.0.0
Subnet Mask=255.255.0.0
SUCCESS: subnet mask set
```

mask delete

Set the subnet mask for the IP address to 0.0.0.0.

Syntax

```
mask delete
```

Example

```
0002:ipmi$ mask delete
Subnet Mask=0.0.0.0
SYUCCESS: subnet mask delete
```

mask show

Show the subnet mask for the IPMI IP address.

Syntax

```
mask show
```

Example

```
0002:ipmi$ mask show
Subnet Mask=255.255.0.0
SUCCESS: subnet mask show
```

gateway

The `gateway` command sets the IPMI gateway. The `gateway` command includes the following elements:

Table 27-36: IPMI category gateway command elements

set	Sets the IPMI gateway.
delete	Deletes the IPMI gateway.
show	Shows the IPMI gateway.

gateway set

Use the `gateway set` command to set the IPMI IPMI gateway.

Syntax

```
gateway set ip_address
```

Example

```
0001:dsm$ ipmi
0002:ipmi$ gateway set 10.10.79.254
Gateway=10.10.79.254
```

gateway delete

Use the `gateway delete` command to delete the IPMI gateway.

Syntax

```
gateway delete
```

Example

```
0002:ipmi$ gateway delete
Gateway=0.0.0.0
SUCCESS: gateway delete
```

gateway show

Use the `gateway show` command to show the IPMI gateway.

Syntax

```
gateway show
```

Example

```
0003:ipmi$ gateway show
Gateway=0.0.0.0
SUCCESS: gateway show
```

disable

The `disable` command disables the IP, and mask. It sets them network addresses to 0.0.0.0.

Syntax

```
disable
```

Example

```
0001:dsm$ ipmi
0002:ipmi$ disable
Do you want disable IPMI network? Access to IPMI will not be available afterwards. (yes|no)
[no]:yes
SUCCESS: Disable IPMI network
```

user

The `user` command adds, shows, or deletes users. Also changes user password and privilege level. The `user` command includes the following elements:

Table 27-37: IPMI category user command elements

add	Add an IPMI user.
delete	Delete an IPMI user.
show	Show the IPMI users.

Table 27-37: IPMI category user command elements (continued)

password	Change the IPMI user password.
level	Set the IPMI user privilege level.

user add

Add an IPMI user. You can modify the password for this user later, using the `user password` command.

Syntax

```
user add userID User_name privilege_level
```

Where:

- *userID* - IPMI user ID
- *User_name* - IPMI user name up to 15 characters
- *privilege_level* - Administrator, Operator or User

Example

```
0001:dsm$ ipmi
0002:ipmi$ user add 3 user1 3
Enter new password:
Enter password again:
SUCCESS: ipmi user modified
```

user password

Change or add IPMI user's password. The password constraints are as follows; 8 characters minimum, 20 characters maximum. At least one capital letter, one number, and one special character are also required.

Syntax

```
user password userID
```

Example

```
0002:ipmi$ user password 3
Enter new password:
Enter password again:
SUCCESS: user password modified
```

user delete

Delete user.

Syntax

```
user delete userID
```

Example

```
0002:ipmi$ user show
0003:ipmi$ user delete <userid>
```

user show

Use `user show` to show the IPMI users.

Syntax

```
user show
```

Example

```
0003:ipmi$ user
    User ID | User Name | Privilege Level | Enable
    ------- | --------- | --------------- | ------
    2       | ADMIN     | Administrator   | Yes
    SUCCESS: user show
```

user level

Change IPMI user's privilege. Don't assign administrative privileges to all users. Assign operator or user privilege instead.

Syntax

```
user level userID privilege_level
```

Example

```
user level 3 3
```

clearint

Clear chassis intrusion. When someone opens the chassis cover of the DSM, chassis intrusion will get flagged. When the cover is put back, run `clearint` to clear it.

Syntax

```
clearint
```

Example

```
clearint
```

reset

The reset command resets the IPMI configuration or resets the IPMI BMC controller. The reset command includes the following elements:

Table 27-38: IPMI category reset command elements

factorydefault	Reset IPMI configuration to factory default.
bmc	Reset IPMI BMC controller.

factorydefault

Wipe out IPMI configuration and return to IPMI factory default: ip=192.168.10.10, mask=255.255.0.0, gateway=0.0.0.0. ADMIN will be the only user left with default password ADMIN. All others users are deleted. Also disables DHCP and sets the IPMI to a dedicated non-share port.

Syntax

```
reset factorydefault
```

Example

```
reset factorydefault
```

reset bmc

Reset IPMI controller. When you want to do a reboot on the IPMI BMC controller chip because the IPMI is failing to respond, or a self-test failed, use this command to reset IPMI chip.

Syntax

```
reset bmc
```

Example

```
reset bmc
```

selftest

Test that the BMC chip is working.

Syntax

```
selftest
```

Example

```
selftest
Selftest: Passed.
SUCCESS: ipmi selftest
```

version

Show IPMI version.

Syntax

```
version
```

Example

```
version
Firmware Version: 03.40
SUCCESS: ipmi show version
```

psinfo

Show power supply information. If only one power module is plugged in, the output will display only one module.

Syntax

```
psinfo
```

Example

```
0001:ipmi$ psinfo

[SlaveAddress = 78h] [Module 1]

Item                            |                Value
----                            |                -----
Status                          |       [STATUS OK](00h)
AC Input Voltage                |              116.5 V
AC Input Current                |               0.34 A
DC 12V Output Voltage           |              12.00 V
```

```
DC 12V Output Current         |               2.25 A
Temperature 1                 |               27C/81F
Temperature 2                 |               34C/93F
Fan 1                         |               5472 RPM
Fan 2                         |               0 RPM
DC 12V Output Power           |               26 W
AC Input Power                |               36 W
PMBus Revision                |               0x8B22
PWS Serial Number             |          P406PCE24AT1144
PWS Module Number             |             PWS-406P-1R
PWS Revision                  |               REV1.1

[SlaveAddress = 7Ah] [Module 2]

Item                          |               Value
----                          |               -----
Status                        |          [STATUS OK](00h)
AC Input Voltage              |               117.0 V
AC Input Current              |               0.37 A
DC 12V Output Voltage         |               12.00 V
DC 12V Output Current         |               2.87 A
Temperature 1                 |               35C/95F
Temperature 2                 |               40C/104F
Fan 1                         |               6560 RPM
Fan 2                         |               0 RPM
DC 12V Output Power           |               36 W
AC Input Power                |               43 W
PMBus Revision                |               0x8D22
PWS Serial Number             |          P406PCE24AT1150
PWS Module Number             |             PWS-406P-1R
PWS Revision                  |               REV1.1

SUCCESS: Show power supply information
```

dhcp

Enable or disable Dynamic Host Configuration Protocol (DHCP). Forces the IPMI IP address to be static. This is turned off by default.

Syntax

```
dhcp [enable, disable]
```

Example

```
dhcp [enable, disable]
```

port

Enable, disable, or check the status of the IPMI port. The default ports are https=443, keyboard/video/mouse (kvm)=5900, vmedia=623, web=80. IPMI users can change the port numbers but the service will still get enabled or disabled with the same command.

Table 27-39: IPMI category port command elements

enable	Enable IPMI port.
disable	Disable IPMI port.
status	Show IPMI port status.

enable

Enable IPMI port.

Syntax

```
port enable [https | kvm | vmedia |web]
```

Example

```
0001:dsm$ ipmi
0002:ipmi$ port enable https
```

disable

Disable IPMI port.

Syntax

```
port disable [https | kvm | vmedia |web]
```

Example

```
0002:ipmi$ port disable https
```

status

Check the status or IPMI ports.

Syntax

```
port status [https | kvm | vmedia |web]
```

Example

```
0003:ipmi$ port status https
IPMI https web port is enabled
```

User Category Commands

The `user` category enables you to add, modify, delete, and display DSM CLI administrators. When setting up a new appliance-based or software-only installation, access the appliance or system through the DSM CLI and do basic appliance configuration, like IP address and host name. After the appliance or system is setup, you can run the Management Console to configure policies, keys, and GuardPoints.

CLI administrators are system users, and are not related to the administrators configured and displayed in the Management Console. DSM CLI administrators configure the appliance network, configure High Availability, and do general appliance administrative tasks. A DSM CLI administrator cannot log into the Management Console and a Management Console administrator cannot log into the DSM CLI.

The `user` category supports the following commands:

Table 27-40: CLI user category commands

add	Adds a new CLI administrator.
delete	Deletes a CLI administrator.
modify	Changes a CLI administrator password.
show	Lists all configured CLI administrators.

add

The `add` command enables you to add new DSM CLI administrators.

CLI administrators are system administrators and are not related to the administrators configured and displayed in the Management Console or to the HSM administrator. Administrators created in the Management Console are placed in the DSM database, and follow the strong password requirements set the Management Console Password window. Administrators created via the CLI are also placed in the `/etc/passwd` file on the system, and use a modified set of requirements. If the administrator already exists in `/etc/hosts` as a regular system administrator, the CLI administrator will not add the administrator to the DSM database.

The default CLI user, `cliadmin` password is `cliadmin123`. The password that you enter can be include a-z, A-Z, 0-9, and the special characters (@#$%^&*(){}[]).
Other characters, such as spaces and periods, are not supported.

Password complexity is defined in the Password tab in the General Preferences window. The password may be from 8 to 31 characters long. The **Password Complexity** group of parameters set requirements for uppercase/lowercase characters, special characters, and integers in the password.

Syntax

```
new <name>
```

Example

```
0016:user$ new hpotter
    Enter new password  : hpotter!0565
    Enter password again: hpotter!0565
    Add user SUCCESS
0017:user$ show
    user[1]: name=cliadmin
    user[2]: name=hpotter
    total: 2
    Show user SUCCESS
0018:user$
```

This is an example only. The password is not actually displayed when it is entered.

If the administrator already exists as a regular system administrator, an error like the following will be returned.

```
0001:user$ show
    user[1]: name=cliadmin
    total: 1
    Show user SUCCESS
0002:user$ add hpotter
    Enter new password  :
    Enter password again:
    ERROR: User name hpotter already exists
0003:user$
```

Though the user `hpotter` does not exist on the DSM, user `hpotter` cannot be added to the database because it already exists in `/etc/passwd`. If you want to add the user to the DSM, delete that user from `/etc/passwd` and then run the DSM CLI command again.

delete

The `delete` command removes a DSM CLI administrator from the system.

Syntax

```
delete name
```

Example

The following example deletes an administrator named hpotter:

```
0010:user$ delete hpotter
    Delete user SUCCESS
0011:user$
```

modify

The `modify` command is used to change a DSM CLI administrator password.

You must know the current password of the administrator to execute this command. The requirements for an acceptable password are set in the Management Console Password window.

The password that you enter can include a-z, A-Z, 0-9, and the special characters, (!@#$%^&*(){}[])
Other characters such as spaces and periods are not supported.

Password complexity is defined in the Password tab in the General Preferences window. The password may be from 8 to 31 characters long. The **Password Complexity** group of parameters set requirements for uppercase/lowercase characters, special characters, and integers in the password.

Syntax

```
modify name passwd
```

Example

The following example changes the password of the DSM CLI administrator `hpotter`.

```
0022:user$ modify hpotter passwd
    Enter old password  : hpotter!0565
    Enter new password  : 0957#hpotter
    Enter password again: 0957#hpotter
    Modify user SUCCESS
0023:user$
```

This is an example only. The password is not displayed when it is entered.

show

The `show` command displays configured DSM CLI administrators.

Syntax

```
show
```

Example

The following example displays information about all currently configured DSM CLI administrators on the DSM.

```
0017:user$ show
    user[1]: name=cliadmin
```

```
user[2]: name=hpotter

total: 2

Show user SUCCESS
```

Chapter 28: Suite B Cryptography

Suite B is a group of publicly available cryptographic algorithms approved by the United States National Security Agency (NSA). The algorithms consist of:

- Advanced Encryption Standard (AES)

- Elliptical Curve Digital Signature Algorithm (ECDSA)

- Elliptic Curve Diffie–Hellman (ECDH)

- Secure Hash Algorithm 2 (SHA-256 and SHA-384)

These algorithms enhance security by adding up to 384-bit encryption to the communication between a Web browser and the DSM, the DSM and Agent, and DSM to DSM in HA environments.

Hardware and Software Requirements

Before you can use Suite B for browser/Agent/DSM communication you need:

- DSM (v5.2.1 or greater)

- VTE Agent (v5.2.x or greater) installed on a host machine

- Microsoft Internet Explorer 11 or greater

- Ports 8446, 8447 and 8448 must be open.

For dedicated Suite B communication you need an agent version of 5.2.x or greater.

Suite B Modes and Ports

The Suite B mode is set to 'compatible' by default, so as to be backward compatible with older versions of the agents. In compatibility mode the `genca` command generates both RSA and ECC certificates. The DSM uses the ECC certificate by default to communicate with other DSMs, newly installed or upgraded agents, and uses the RSA certificate to communicate with older agents. When the agent negotiates a transaction with the DSM, the handshake determines which certificate is to be used. If Suite B mode is enforced, only ECC certificates are used and older agents that do not support ECC will need to be upgraded or will fail to communicate with the DSM.

Operating Modes

With respect to Suite B encryption, the DSM operates in one of two modes.

- **Compatibility Mode** - Version 4.4.1 through 5.1.x Agent software can operate with version 5.2.x DSM software using up to 256-bit encryption. If the Agent and DSM are operating with 5.2.x software, they are compatible and use up to 384-bit encryption.

- **Suite B Mode** - Version 5.2.x DSM and Agent software are compatible and operate at up to 384-bit encryption. A DSM with version 5.2.x software in this mode is not compatible with agent software older than 5.2.x.

Communication Ports

If the DSM is in 'compatible' mode, the browser can connect to the DSM on port 8445 and use the RSA certificate for communication, or, the browser can connect on port 8448 and use the ECC certificate for communication.

If Suite B mode is enforced, only port 8448 is supported. Enabling Suite B mode switches the port redirect from 443 >8445 to 443 >8448. If you don't specify a port in the browser URL, it automatically uses port 443 by default and the DSM will forward it to port 8445 or 8448 depending on the mode.

> **Note**
> In case of IPv6 addresses, there is no port forwarding and the ports will have to be manually specified.

In Suite B mode, any attempts to connect to port 8445 will automatically redirect to port 8448 and use the ECC certificate.

Not all versions of browsers support ECC, currently the only browser that supports the DSM ECC is Microsoft Internet Explorer 11.

Figure 28-1: Suite B modes and ports
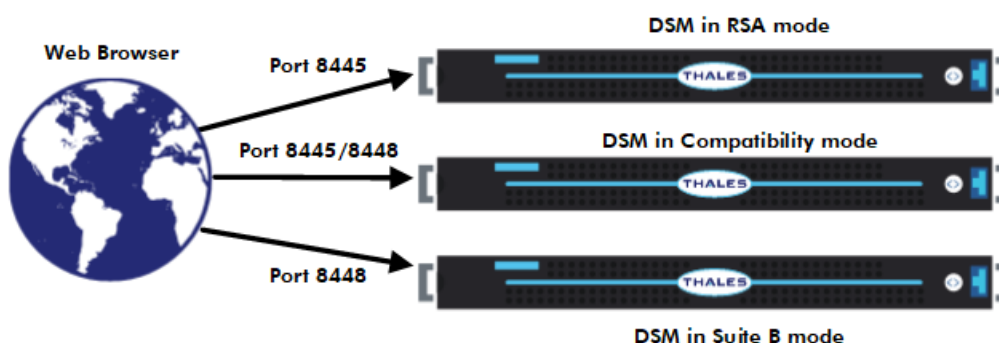
Suite B DSM and Agent Compatibility

The current DSM and Agent software was designed for backward compatibility. All older agent software beginning with 5.2.3 will work with the current DSM software in compatibility mode, but not in Suite B only mode. The following table shows more detail.

Table 28-1: DSM and Agent compatibility

DSM Software	Agent Compatibility	
Version/Mode	Agent 5.1.x	Agent 5.2.x
DSM (5.1.x)/Normal mode	Yes (up to 256-bit encryption)	Yes (up to 256-bit encryption)
DSM (5.2.x)/ Compatibility mode	Yes (up to 256-bit encryption)	Yes (up to 384-bit encryption)
DSM (5.2.x)/ Suite B mode	No, not compatible	Yes (up to 384-bit encryption)

> **Note**
> If KMIP is enabled for key management then the DSM must be in 'compatible' mode as KMIP does not support ECC.

Suite B Mode Status

By default, all 5.2.x and newer DSMs will operate in Compatibility Mode. The dashboard shows the following fingerprints when Suite B is active.

RSA CA fingerprint: The RSA server fingerprint represents the certificate used to establish secure communications between the DSM and the browser. The RSA fingerprint represents up to 256-bit encryption and should be compared with fingerprints generated during DSM and Agent registration. Only the RSA fingerprint is displayed on the dashboard when both DSM and Agent software are earlier than 5.2.1.

EC CA fingerprint: The Elliptical Curve (EC) server fingerprint represents the certificate used to establish secure communications between the DSM and browser. The EC fingerprint represents up to 384-bit encryption and should be compared with fingerprints generated during DSM and Agent registration. In compatibility mode, the EC and RSA fingerprints are both displayed when the DSM is at version 5.2.1 or newer. In Suite B mode, only the EC fingerprint is displayed when the DSM is at version 5.2.1 or newer.

In addition to displaying status, you can also activate or deactivate Suite B mode from the DSM CLI utility.

Activate Suite B mode

1. Open a CLI session on the DSM.

2. Type 'system' to access the System command menu and press enter.

   ```
   0001:vormetric$ system
   ```

3. At the prompt, type the following:

   ```
   security suiteb set suiteb
   0002:system$ security suiteb set suiteb
   WARNING: After setting to suiteb mode, the security server software will be restarted
   automatically!
   Continue? (yes|no)[no]:yes
   ```

Deactivate Suite B mode

1. Open a CLI session on the DSM.

2. Type System to access the System command menu and press enter.

3. At the prompt, type the following:

```
0002:system$ security suiteb set compatible
```

Show Suite B mode

1. Open a CLI session on the DSM.

2. Type System and press enter.

3. At the prompt, type the following:

   ```
   0003:system$ security suiteb show
   Current mode is: suiteb
   SUCCESS: showed suite b mode
   0003:system$
   ```

Chapter 29: Other Data Center Administrative Roles

Personnel doing the initial DSM setup and configuration using the CLI can also be thought of as administrators. They may include any of the following job titles and are system users with login accounts. Note that different companies will have different titles for the roles and responsibilities (and may combine two or more) for these personnel.

Although they may temporarily be granted access to the DSM through the CLI, they will not have access to the DSM

System Administrators

A system administrator is responsible for the upkeep, configuration, and reliable operation of computer systems in a data center. Some system administrators have access to the data on the machines that they administer; however, they don't need this access and it is a security liability.

Data Center Administrators

Data center administrators set up, run, and maintain data centers, and are responsible for the day-to-day operation and interoperability of the sometimes large and complex data centers.

Database Administrators

Database administrators (DBA) maintain a database system, and are responsible for the integrity of the data and the efficiency and performance of the system.

Network Administrators

Network administrators maintain network infrastructure such as hubs, bridges, switches, and routers, and diagnose problems with these or with the behavior of network-attached computers.

Security Administrators

Security administrators are specialists in computer and network security, including the administration of security devices such as firewalls, as well as consulting on general security measures. Security administrators may be part of a company's loss prevention team or as members of a separate group in a company's data center. A specific security administrator (or security administrators) in an organization may or may not have DSM administrator privileges depending on the organization's separation of duties policies.

Web Administrators

Web administrators maintain web server services (such as Apache or IIS) that allow for internal or external access to web sites. Tasks include managing multiple sites, administering security, and configuring components and software. Web administrator responsibilities may also include software change management.

Storage Administrators

Storage (SAN) Administrators create, provision, add, or remove storage to/from computer systems. Storage can be attached local to the system or from a Storage Area Network (SAN) or Network Attached Storage (NAS). Storage administrators also create file systems from newly added storage.

Computer Operators and Lab Technicians

Computer operators and lab technicians do routine maintenance and upkeep, such as changing backup tapes or replacing failed drives in a RAID. Such tasks usually require physical presence in the room with the computer; and while less skilled than system administrator tasks, require a similar level of trust, since the operator has access to possibly sensitive data.