THALES

Vormetric Data Security Platform

Installation and Configuration Guide

Release 6.1.3

Document Version 2

May 2, 2019

Documentation PN: 50-1000020-08
Document Version 2

Vormetric Transparent Encryption
Installation and Configuration Guide
Release v6.1.3
May 2, 2019
Document Version 2

NOTICES, LICENSES, AND USE RESTRICTIONS

CONTENTS

PREFACE

The *Vormetric VTE Agent Installation and Configuration Guide* describes how to install and configure Vormetric Transparent Encryption (VTE) on hosts. After you install VTE and configure the VTE with the DSM, VTE can protect data on hosts.

DOCUMENTATION VERSION HISTORY

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

Table 1: Documentation Version History

	Date	Changes
6.0.3 v1	03/14/2018	Updated Secure Start chapter for Azure, added sections on CLI upgrade and restricted mode.
6.0.3 v2	05/31/2018	Added information on SecFS support for NetApp snapshots.
6.1.0 v1	8/21/ 2018	Added information for Linux features, enhanced encryption mode, silent install and automatic registration of LDT and Docker.
6.1.1 v1	9/14/18	Added information for Windows features: Exchange DAG encryption, enhanced encryption mode, silent install and automatic registration of LDT.
6.1.2	10/20/18	
6.1.3 v1	03/07/19	Added information on Port Configuration, Scheduled Upgrade, Binary resigning.
6.1.3 v2	5/2/19	Remove obsolete information about Windows scheduled upgrade (on/off not supported).

SCOPE

This document describes the installation and configuration of the VTE Agent on Linux and Windows platforms.

INTENDED AUDIENCE

The *VTE Agent Installation and Configuration Guide* is intended for system administrators who install and configure VTE on host machines.

Assumptions

This document assumes knowledge of network configuration. The system administrator must have root permissions for the systems on which the VTE software is installed.

SERVICE UPDATES AND SUPPORT INFORMATION

The license agreement that you have entered into to acquire the Thales products (“License Agreement") defines software updates and upgrades, support and services, and governs the terms under which they are provided. Any statements made in this guide or collateral documents that conflict with the definitions or terms in the License Agreement, shall be superseded by the definitions and terms of the License Agreement. Any references made to “upgrades” in this guide or collateral documentation can apply either to a software update or upgrade.

For support and troubleshooting issues:

- https://help.thalesesecurity.com/hc/en-us
- Email questions to support@vormetric.com or call 877-267-3247

For Vormetric Sales:

- http://enterprise-encryption.vormetric.com/contact-sales.html
- Email questions to sales@vormetric.com or call 888-267-3732

1

Overview

Introduction

This document describes how to install and configure VTE Agent on host computers requiring data protection. These host computers are called *protected hosts*. VTE is supported in multiple operating system environments and you can deploy it on physical devices as well as virtual environments.

VTE secures data with little impact to application performance. It requires no changes to your existing infrastructure and supports separation of duties between data owners, system administrators and security administrators.

The VTE solution consists of a *Data Security Manager* (DSM) and VTE residing on your hosts, your protected hosts or servers.

- The DSM is the central component of the VTE solution. You can set it up as either a security-hardened physical appliance or a virtual appliance. The DSM stores and manages data encryption keys, data access policies, administrative domains, and administrator profiles.
- The VTE communicates with the DSM and implements the security policies on their GuardPoint systems.
- You can apply VTE to GuardPoints to the servers on-site, in the cloud, or a hybrid of both.

NOTE: Refer to the VDS Compatibility Matrix for a list of VTE versions and supported operating systems.

VTE protects data at rest. VTE can protect data residing on Direct Attached Storage (DAS), Network Attached Storage (NAS) or Storage Area Networks (SAN). This can be a mapped drive or mounted disk, as well as through Universal Naming Convention paths.

The VTE installation and configuration process:

1. Install VTE on the protected host.
2. Add the protected host fully qualified domain name (FQDN) or IP address to the DSM.
3. Register the protected host with the DSM so they can communicate with each other.

How to protect data with VTE

VTE protects data by creating policies that specify file encryption, data access, and auditing on specific directories and drives on your protected hosts. These directories are called *GuardPoints*. Policies specify:

- Whether or not the resting files are encrypted
- Who can access decrypted files and when
- What level of file access auditing is desired by generating fine-grained audit trails

Policies are created by an administrator through the DSM GUI, called the Management Console. You, as an administrator, access the GUI through a Web Browser. After you create the policies and push them to protected hosts, VTE implements those policies.

While a GuardPoint is disabled, VTE cannot enforce security and key selection rules on the files inside an unguarded directory. Therefore, access to data inside an unprotected directory is undetected and ungoverned under the rules of the policy. This may cause corrupted data if data is modified or added to the directory without a protection policy in place.

2 Installing VTE for Windows

This chapter describes how to install and configure VTE on Windows computers.

Both VTE or host installer and the DSM administrator assist in the installation.

This chapter contains the following sections:

Installation Overview

The installation and configuration process consists of three basic steps:

1. Installing VTE on the protected host.
2. Manually adding the protected host fully qualified domain name (FQDN), or IP address, to the DSM. (This is automatically performed using the *Shared Secret Registration* method.)
3. Registering the protected host with the DSM so they can communicate with each other.

Assumptions

- The IP addresses, routing configurations, and DNS addresses allow connectivity between the DSMs and all hosts which contain VTE installations.
- If the protected host is a virtual machine, the VM is deployed and running.

Pre-installation Tasks and Instructions

This section lists tasks you must complete and information you must gather before installing VTE:

Location for VTE installation

You must install VTE on the system drive.

NOTE: Do not install VTE in network shared volumes.

Determine your VTE registration method

You can register the protected hosts with the DSM using either the *Fingerprint method* or the default *Shared Secret method*.

- **Fingerprint method**:

  Requires the DSM Administrator to add the FQDN, or IP address, of each protected host to the DSM before registering VTE.

  During the registration, the DSM generates the certificate and passes it down to VTE along with the fingerprint. The security administrator needs to verify the fingerprint to make sure the certificate is valid.

- **Shared Secret method**:

  Requires the DSM Administrator to create a *shared secret* password—a case-sensitive string of characters—for auto-registering a domain or host group.

  VTE installers use the shared secret to add and register protected hosts to the DSM for a domain or host group. The DSM Security Administrators can optionally

add host names or IP addresses to the DSM. There is no need to verify that the protected host and DSM share valid certificates. You can add multiple protected hosts dynamically with a single shared secret password during VTE installation and registration.

After the DSM Administrator creates a shared secret for the domain or host group in which the new protected host will reside, obtain it and the validity period (one hour, day, week, or month) and register within that period.

Host name resolution

Host name resolution is the method of mapping a host name to an IP address. During this configuration process, enter either the FQDNs, or IP addresses, of your DSM and protected hosts. If you use FQDNs, your protected hosts must be able to resolve their DSM host names, and the DSM must be able to resolve its protected hosts.

NOTE: The exception to this requirement is if you approve of only VTE-initiated communication between the DSM and the protected host. See "One-way communication" on page 8 for more discussion.

A Domain Name Service (DNS) server is the preferred method of host name resolution. Use the following guidelines for host name resolution:

1. If you use DNS, use the FQDNs for the DSM and host for the installation and configuration procedures in this chapter.
2. If you do NOT use a DNS server, complete one of the following tasks on the DSM and the protected hosts:
   - Use the IP addresses of the DSM and protected hosts.
   - Add an entry for the DSM in the `C:\WINDOWS\system32\drivers\etc\hosts` file on the Windows protected hosts.

NOTE: You may have to obtain the DSM server names and IP addresses from the DSM Administrator.

3. Repeat this procedure for each protected host and add an entry for all DSM nodes.

Port configuration

If a protected host must communicate with a DSM through a firewall, open the ports in the firewall as shown in the following figure.

Figure 1: Ports to open between DSM and protected host

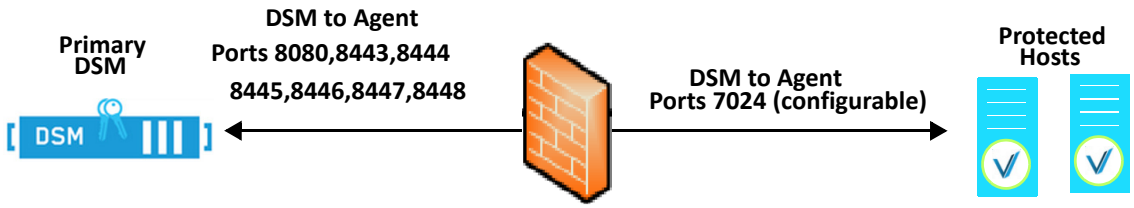

NOTE: See Table 1 to determine which of the above ports must be opened through the firewall.

Table 2 describes the communication direction and purpose of each protected host port you must open. For a complete list of ports required for the DSM, see the *VDS Administration Guide*.

Table 2: Ports to Configure

Port	Protocol	Communication Direction	Purpose
7024	TCP	DSM → Agent	Policy/Configuration Exchange
8080	TCP	Agent → DSM DSM ↔ DSM	Default TCP/IP port for HTTP that is used to exchange certificates between the DSMs in an HA configuration. Also, used once to do the initial certificate exchange between an agent host and DSM.
8443	TCP	Agent → DSM	TCP/IP port through which the agent communicates with the DSM. The agent establishes a secure connection to the DSM, via certificate exchange, using this port.
8444	TCP	Agent → DSM	Agent log messages uploaded to DSM.
8445	TCP	Agent → DSM	Used during initial certificate exchange between agent host and DSM over HTTPS. If not available, then port 8080 will be used instead.
8446	TCP	Agent → DSM	Configuration Exchange using Elliptic Curve Cryptography (Suite B)

8447	TCP	Agent → DSM	Agent uploads log messages to DSM using Elliptic Curve Cryptography (ECC)
8448	TCP	Agent → DSM	Used during initial certificate exchange between agent host and DSM over HTTPS. If not available, then port 8445 or port 8080 will be used instead.

The default port for communication between the DSM and the agent is 7024. If this port is already in use, set the port to a different number by specifying the new port during agent installation.

In the following example, the default communication port between the DSM and agent was changed from 7024 to 8000:

```
Please enter the host name of this machine, or select
from the following list. The name you provide must
precisely match the name used on the "Add Host" page of
the Management Console.

[1] centos-6-0

[2] 192.168.1.160

Enter a number, or type a different host name or IP
address in manually:

What is the name of this machine? [1]: centos-6-0:8000

You entered the host name "centos-6-0:8000"

Is this host name correct? (Y/N) [Y]:
```

Port Usage in One Way Communications Mode

By default, polling from the agent host to the DSM when running in one-way communications mode uses HTTP via port 8080. If the agent is configured to use secure polling, then polling is performed using HTTPS via port 8448 (in suite B mode) or port 8445.

Hardware Association

Vormetric's hardware association feature associates the installation of VTE with the machine's hardware. When enabled, hardware association prohibits cloned or copied versions of VTE from contacting the DSM and acquiring cryptographic keys. Hardware association works on both virtual machines and hardware hosts.

You can enable hardware association during VTE registration process. You can disable hardware association by re-running the registration program.

To verify if hardware association (cloning prevention) is enabled on the protected host, launch the Windows command line and run the following command:

```
C:\Program Files\Vormetric\DataSecurityExpert\agent\vmd\bin>
vmsec.exe hwok
```

To change the status from enable to disable or vice versa:

1. Open the system tray and right-click on the Vormetric icon.
2. Select **Register Host**.
3. Follow the prompts to re-register VTE with the DSM.
4. Select **Enable hardware association** in the wizard.

One-way communication

In some deployments, VTE might not be visible to the DSM through normal network communications. For example, when the host on which VTE is installed:

- is behind NAT
- is behind a firewall
- is not permanently connected to a communication channel to the DSM
- is unable to resolve the host name to an IP address

In these situations, VTE can initiate VTE-only communication to the DSM. This feature is called one-way communication and works by having VTE poll the DSM for any policy messages or changes, then downloading changes as required.

The downside of one-way communication is that the DSM cannot issue any queries to VTE. For example, the DSM Administrator cannot browse host directories or User IDs.

If you want the DSM to be able to contact VTE directly, abandon the Register Host process if you get the following message during the registration process:

Determine the installation method

There are two methods for installing VTE:

- **Typical installation:** This is the most common and recommended type of installation. Use this for installing VTE on one host at a time. See "Windows Typical Install" on page 10.
- **Silent installation:** Create pre-packaged installations by providing information and answers to the installation questions. Use silent installations when installing on a large number of hosts. See "Windows Silent Install" on page 15.

VTE Install Checklist

Use this table to verify prerequisites and collect the information you need for the installation.

Table 3: VTE Install Checklist

Checklist item	Status
Obtain VTE installation file from Vormetric support. The format for VTE file names is: vee-<*agent_type-build-system*>*.exe* Example: `vee-fs-6.1.3.0-132-win64.exe`	
Fully Qualified Domain Name (FQDN) of the DSM	
IP address or FQDN of the host	
Administrator password for the host	
If using Shared Secret Registration, obtain the following from the DSM Security Administrator: 1) Shared secret 2) Domain 3) Host group, if applicable 4) (Optional) A description for the host.	
If using the Fingerprint Registration ask the DSM Administrator to add the host to the DSM and check the **Registration Allowed** check box. After checking the fingerprint, select the **Communication Enabled** check box.	
Addressed "Host name resolution" (page -5) for the protected hosts and DSM?	
Set "Port configuration" on page 6	

Table 3: VTE Install Checklist

Checklist item	Status
Do you want "Hardware Association" on page 7?	
Is "One-way communication" on page 8 required?	
Synchronize host clock to DSM clock.	
Set network subnet mask on the host (unless you are using one-way communication)	
Preferred DNS Server (if using FQDNs):	

Windows Typical Install

Use the *typical install* feature to install VTE on a host by manually answering each installation question. A *silent install* pre-packages installations by providing information with scripts that answer the installation questions. Use silent installation when installing on a large number of hosts (see "Windows silent install command" on page 15).

Typically, you register VTE with the DSM as part of the installation process as described below. However, you may postpone registration if you have a plan to register VTE later. VTE cannot protect data on the host until you register VTE with the DSM.

Install VTE on the protected host

NOTE: If you are upgrading a protected host, stop all activity on the application before upgrading.

1. Log on to the host as a Windows user with administrative privileges.
2. Copy the installation file onto the Windows system.
3. Double-click the installation file.

4. The *InstallShield Wizard for Vormetric File System Agent* window opens. Click **Next**.
5. Verify the version of VTE you are installing and click **Next**.
6. The ***License Agreement*** appears. Accept the ***License Agreement*** and then click **Next**.
7. The ***Destination Folder*** window opens. Click **Next** to accept the default folder.

NOTE: If you have VTE already installed on the system and are upgrading, you can not change the Destination Folder.

8. The *Ready to Install the Program* window opens. Click **Install**. VTE software installs. This may take a few minutes. When the installation is finished, the *Install Shield Wizard Completed* window opens.
9. Click the **Register Vormetric File System Agent now** option.
10. During host registration, a new feature allows you to automatically register and enable LDT during registration. The installation script adds a question to the Installation Wizard to ask if you want to enable the automatic registration.
11. To register now using the *Certificate Fingerprint* method, go to "To register VTE using the Certificate Fingerprint method" on page 13.
    To register now using the *Shared Secret* method, go to "To register VTE using the Shared Secret Registration method" on page 11.
12. If you want to register VTE later, clear the check box for **Register Vormetric File System now**, then click **Finish**.

To register VTE using the Shared Secret Registration method

1. Select **Register Vormetric File System Agent now** and click **Finish** in Step 11 in the previous section.
2. Click **Next**. The *Select components to register* window opens. Verify that you have selected the correct VTE type.

   Register your host by using a domain-specific shared secret created by the DSM Security Administrator. Alternatively, use the fingerprint method that requires you to verify fingerprints displayed by this registration program with those shown in the DSM Dashboard (see "Determine the installation method" on page 8).
3. Click **Next**. You are prompted for the shared secret registration information:

- **Shared secret**: Password for the domain to which the host, or host group, will be added. Contact the DSM Security Administrator for this value.
- **Domain name**: Name of the DSM domain to which the host will be added. Contact the DSM Security Administrator for this value.
- **Host group** (optional): Name of the host group to which the host will be added. Contact the DSM Security Administrator for this value.
- **Host description** (optional): Description of the host to be registered.

Warning! Be sure to enter the case-sensitive shared secret, domain name, and host group correctly. If any of these are entered incorrectly, an error message displays. If you exceed the number of tries defined in the **Maximum Number of Login Tries** setting on the DSM **Password Preferences** page (System > General Preferences > Password), you are locked out of the system for a period define in the **User Lockout Time** setting.

4. Click **Next**. VTE prompts you to enter the host name of the Primary DSM. If necessary, ask the DSM Security Administrator for the **Server name** on the dashboard of the DSM Management Console. Type the name of the DSM and click **Next**.
5. VTE prompts you to enter the protected host name. Enter the protected host name or IP address that the DSM Administrator entered in the DSM. If it is a name, it must be resolvable by the DNS server (see “Host name resolution” on page 5). You can type the name/IP address or select it from the drop-down menu. To enable cloning prevention, select the check box called **Enable Hardware Association**, see “Hardware Association” on page 7. Click **Register.**

NOTE: If you get the message “Only VTE-initiated communication is possible ...”, read “One-way communication” on page 8 and choose accordingly.

6. If successful, the *Register Host* window displays that the File System Component was successfully registered with a Management Console.
7. Verify the installation by checking VTE processes.

a. In the system tray of the protected host, right-click the Vormetric icon.

b. Select **Status**. Review the information in the **Vormetric Status** window to confirm that the correct VTE are installed and registered.

To register after Installation

1. When registering after the installation, change to the following directory:

   `C:\Program Files\Vormetric\DataSecurityExpert\agent\shared\bin\`

2. Type:

   `> register_host.exe`

To register VTE using the Certificate Fingerprint method

1. Select **Register Vormetric File System Agent now** and click **Finish** in Step 11 in the previous section.
2. Click **Next**. The *Select components to register* window opens. Verify that you have selected the correct VTE type.

   Register your host by using a domain-specific shared secret created by the DSM Security Administrator. Alternatively, use the fingerprint method that requires you to verify fingerprints displayed by this registration program with those shown in the DSM Dashboard (see “Determine the installation method” on page 8).
3. Uncheck the **Use Shared Secret Registration** checkbox and click **Next**.
4. VTE prompts you to enter the host name of the Primary DSM. If necessary, ask the DSM Administrator for the **Server name** on the dashboard of the DSM Management Console. Type the name of the DSM and click **Next**.
5. VTE prompts you to enter the protected host name. Enter the protected FQDN host name or IP address, or select it from the drop-down menu, that the DSM Administrator entered in the DSM (if it is a name, it must be resolvable by the DSM server).
6. To enable hardware association (cloning prevention), select the check box called **Enable Hardware Association**, see “Hardware Association” on page 7. Click **Register.**

NOTE: If you get the message “Only agent-initiated communication is possible ...”, read “One-way communication” on page 8 and choose accordingly.

If successful, the *Register Host* window displays the fingerprint of the elliptic curve (EC) Certificate Authority (CA) certificate.

NOTE: If you get the error message *File System component service stopped 'Couldn't resolve host name',* it means the DSM host name could not be resolved by the protected host. See "Host name resolution" on page 5 to fix.

At this stage of the installation, you and the DSM Administrator must exchange information to confirm that VTE host and DSM share valid certificates. This is done to verify that nobody is intercepting and modifying traffic between the DSM and VTE.

7. **VTE Installer**: Send the fingerprint to the DSM Administrator and wait for confirmation.
8. DSM **Security Admin**:
   a. Log on to the DSM **Management Console** and navigate to the domain where the host was added.
   b. Click the **Dashboard** tab.
   c. Match the fingerprint from VTE Installer with the EC CA fingerprint on the Dashboard.
   d. Advise the Host Admin of the results.
9. **VTE Installer**: If the fingerprints match, click **Yes**. A fingerprint for the protected host certificate displays.
10. Click **Ok**. Host has been successfully registered with the DSM. Click **Finish**. Restart the System.
11. Ask the DSM Administrator to select the **Communication Enabled** checkbox on the Management Console for the new protected host.
12. Verify the installation by checking VTE processes.
    a. In the system tray of the protected host, right-click the Vormetric icon.
    b. Select **Status**. Review the information in the **Vormetric Status** window to confirm that the correct VTE are installed and registered.

Windows Silent Install

The silent install is a command-line automated installation that minimizes installer interaction. Use silent install to roll out large numbers of hosts or to reduce your time and interaction as an administrator.

Windows silent install command

VTE silent install runs on the Windows command line with a specific set of options shown in Table 4 and Table 5. The most common format for VTE installation is:

Installation_executable `/s/v"/qn`
REGISTERHOSTOPTS=`\"`***REGISTERHOSTOPTS_Options***`\""`

For example:

```
vee-fs-6.1.1-1-win64-unsigned.exe /s /v"/qn REGISTERHOSTOPTS=\"
SecSrv.demo.com -useip\""
```

Table 4: Command line options for silent install

Windows Attribute	Description	Required?
`/s`	Runs the installation in silent mode.	Yes
`/v`	Use the `/v` option to pass command-line options and values of public properties through to the installer.	Yes
`/qn`	No UI	Yes
`REGISTERHOSTOPTS`	Options for registering VTE with the DSM. The full set of options is specified below in Table 5.	No
`INSTALLDIR`	Specifies an alternate location for installing VTE. NOTE: If syntax is incorrect, silent install will fail. See examples. NOTE: The alternate location must be on the system drive.	No
`REBOOT=ReallySuppress`	By default, the machine reboots after installation. Supplying this option prevents that. However, a reboot is still required to fully install VTE.	No
`-enableldt`	Set to 1 to automatically enable and register LDT during silent install.	No

Table 5: REGISTERHOSTOPTS Options for silent install

Windows Attribute	Description	Required?
The host name of the DSM	FQDN of the DSM. Example: `thl.example.com`	Yes
`-useip`	Use the IP address of the protected host instead of host name. Used when `-agent` is not supplied.	No
`-agent=`***`your.agent.name.com`***	FQDN of protected host.	No
`-onewaycomms`	Set when VTE-initiated-only communication is required. See "One-way communication" on page 8	No
`-usehwsig`	Sets when you want to associate this installation with the machine hardware for cloning prevention. See "Hardware Association" on page 7	No
-log=*filename*	Sets a log file to output results to.	No
-port=*port*	Specifies the port number.	No
`-secret=`***`password`***	Specifies the password for a shared secret registration. See "Determine your VTE registration method" on page 4	No
`-domain=`***`domain_name`***	Specifies domain for the shared secret.	No
`-hostgroup=`***`hostgroup_name`***	Specifies the optional host group for the shared secret.	No
`-description="`***`description`***`"`	Specifies a description for the protected host with a Shared Secret registration, however, it does not overwrite an existing description. It does not work for the Fingerprint method, NOTE: If syntax is incorrect, silent install fails. See examples.	No
`-log=%temp%\vor-agent-reg.log`	Specifies a log file particular to the registration. Useful for documenting why the registration failed (if it does).	No

Windows silent install examples

These examples fall into two categories: installations using the *Shared Secret* registration method and installations using the *Certificate Fingerprint* method (see "Determine your VTE registration method" on page 4). The Fingerprint method requires that you enter the host name, or IP address, in the DSM.

Silent install using the Shared Secret Registration method

The Shared Secret method requires that the DSM Security Administrator create a shared secret password for a domain/host group and share the password with you. Collect the following information from the DSM Security Administrator:

Table 6: Shared Secret Information for VTE Registration

Shared secret password	
Domain assigned to host in the DSM for the shared secret password	
Host group for the shared secret password (if created)	
Description for protected host (optional)	
Shared secret validity period (registration must be completed in this period)	

NOTE: If either the `-description` or `-INSTALLDIR` syntax is incorrect, silent install fails and you must correct the syntax and reinstall. See examples for correct syntax.

Example 1: Default install without description and using 2-way communications. Installs VTE using the Shared Secret method, specifies DSM named `thl.example.com`, host uses IP address to register, supports hardware association (cloning prevention), specifies a secret key for the DSM domain `domain1`. Also automatically registers LDT during install.

```
 vee-fs-6.0.0-28-win64.exe /s /v" /qn
registerhostopts=\"thl.example.com -useip -enableldt -usehwsig -
secret=MaCarena45# -domain=domain1 \""
```

Example 2: Installs VTE using the Shared Secret method, specifies DSM named `thl.example.com`, host uses IP address to register, supports 1-way communications, specifies a shared secret for the DSM domain `domain1` with custom installation directory.

```
 vee-fs-6.0.0-28-win64.exe /s /v" /qn
INSTALLDIR=\"c:\opt\vormetric2\"
registerhostopts=\"thl.example.com -useip -onewaycomms -usehwsig
-secret=MaCarena45#
-domain=domain1 -description=\"\"Silent Install\"\"\""
```

Example 3: Installs VTE using the Shared Secret method, specifies DSM named `thl.example.com`, host uses IP address to register, supports hardware association (cloning prevention), specifies a secret key for the DSM domain `domain1`, installs in default directory, and has a host description “Silent Install”.

```
 vee-fs-6.0.0-16-win64.exe /s /v" /qn
registerhostopts=\"thl.example.com -useip -usehwsig -
secret=MaCarena45# -domain=domain1 -description=\"\"Silent
Install\"\"\""
```

Silent install using the Fingerprint Registration method

The Fingerprint Registration method requires that the DSM Administrator add all hosts on which you will install VTE to the DSM with the Registration Allowed and Communication Enabled features enabled. Once those are added, you may proceed with the silent install.

Example 1: Installs VTE by FQDN host name.

```
 vee-fs-6.0.0-16-win64.exe /s /v" /qn
registerhostopts=\"30181.example.com -onewaycomms \""
```

Example 2: Installs VTE using Fingerprint method, specifies DSM named `thl.example.com`, host uses IP address to register, supports hardware association (cloning prevention), and installs in a custom directory.

```
 vee-fs-6.0.0-16-win64.exe /s /v" /qn
INSTALLDIR=\"c:\abc\vormetric2\"
registerhostopts=\"thl.example.com -useip -usehwsig \""
```

To install in Windows silently

NOTE: Before performing a silent install, disable the firewall because the firewall may block the installation. After installation, enable the firewall.

1. Either gather the appropriate Shared Secret information (for Shared Secret Registration) or have the DSM Security Administrator add the protected host name or IP address to the DSM (Fingerprint registration method).
2. Add the target host to DSM and enable communication for silent mode
3. Log on to the protected host as an administrator and run the silent install command on each host from the Windows command shell.
4. After running the silent install command, the system reboots.
5. Repeat for additional silent installs.

Verify the Windows installation

Verify the installation by checking VTE processes.

1. In the system tray, right-click the Vormetric icon.
2. Select **Status**. Review the information in the Vormetric Status window to confirm the correct VTE are installed and registered.

VTE Scheduled Upgrade

Scheduled upgrade allows you schedule an upgrade of the VTE agent to occur the next time the server on which an agent is installed reboots normally. Scheduled upgrade can minimize VTE service interruptions. Also, scheduled upgrade can reduce coordination issues in organizations where the security roles are separated.

NOTE: The VTE scheduled upgrade feature is compatible with Windows Server 2008 R2 and higher versions.

Schedule VTE Upgrade

To schedule VTE to upgrade the next time the system reboots, type:

```
# voradmin upgrade <VTE setup executable path> -reboot false
```

Example

```
> voradmin upgrade C:\6.0.3.15\vee-fs-6.0.3-15-win64-
  unsigned.exe -reboot false
```

System Response

```
Creating and installing service to upgrade. VTE agent will be
upgraded on next reboot.
```

Warning! If you have scheduled an upgrade on reboot and the system crashes or is not shutdown gracefully, you must restart the system again to upgrade the agent.

Show Scheduled VTE Upgrades

To display all scheduled VTE agent upgrades, type:

```
> voradmin upgrade show
```

System Response

```
Current version:          6.0.3.12
Target upgrade version:   6.0.3.15
Upgrade on reboot:        Enabled
```

Cancel a Scheduled VTE Upgrade

To cancel/cleanup a scheduled VTE agent upgrade, type:

```
> voradmin upgrade cancel
```

System Response

```
VTE agent upgrade canceled successfully.
```

Uninstall VTE from a Windows Host

Caution: After removing VTE software, access to data is no longer controlled. If data was encrypted, it remains encrypted. If decrypted or copied out of the host, the data is visible as clear text.

Before removing VTE from a Windows host, verify the following:

1. All applications protected by hosts are stopped.
2. The DSM Administrator has evaluated the current hosts in the *Guard FS* tab to avoid data loss or compromise.
3. The DSM Administrator removed **System Locked** and **FS Agent Locked** settings for this host (if set).
4. All hosts have been removed.
5. VTE is removed from the host **before** the host is removed from the DSM.
6. Any data you want to use after uninstall is decrypted.

NOTE: If VTE installation fails because a host is in use, determine which applications are using the hosts and stop them. Then run the uninstall again.

To uninstall VTE

NOTE: VTE for Windows must be removed from the host before the host is removed from the DSM **Hosts** pane.

If multiple VTE types are installed on the Windows host, they must be removed or uninstalled separately. Remove all non-VTE items before VTE.

1. Stop any application accessing files in the host.
2. Log on to the host that is running VTE as the system administrator.
3. Run the Windows **Add** or **Remove Programs** utility to remove VTE software.
4. Click **Yes** to reboot the host.

3

Installing VTE for Linux

This chapter describes how to install and configure VTE on Linux systems.

It contains the following sections:

Installation Overview

The installation and configuration process consists of three basic steps:

1. Installing VTE on the host.
2. Adding the host fully qualified domain name (FQDN) or IP address to the DSM. The DSM Administrator can perform this task manually, using the Fingerprint registration method, or automatically, using the Shared Secret registration method.
3. Registering the host with the DSM so they can communicate.

Before you can perform these steps, complete the "Prerequisites" on page 24.

Prerequisites

This section lists the tasks you must complete, and the information you must obtain, before installing VTE:

General setup information

- Vormetric recommends that you install VTE in the default location.
- Do not install VTE on network-mounted volumes such as NFS.
- Make the Installation root directory `/opt` a real directory.
- If `/opt` is a symlink, you **must** use the `-d` option to specify the installation directory, which must be a real directory.
  Example:
  ```
  # ./vee-fs-6.0.3-11-rh7-x86_64.bin -y -d /home/hello/
  ```
- Ensure read/write permission is granted to other users accessing your shared resource.

Network setup

Before you install VTE, the following must be in effect:

- The IP addresses, routing configurations, and DNS addresses allow connectivity of the DSM(s) to all hosts where you install VTE.
- If the host is a virtual machine, the VM is deployed and running.

Determine your VTE registration method

You can register protected hosts with the DSM using either the *Fingerprint method* or the default *Shared Secret method:*

- **Fingerprint method**: Requires the DSM Security Administrator to add the FQDN or IP address of each host to the DSM before registering VTE.

  During the registration, the DSM generates the certificate and passes it down to the VTE with the fingerprint. The security administrator must verify the fingerprint to make sure the certificate is valid.

- **Shared Secret method**: Requires the DSM Security Administrator to create a *shared secret* registration password—a case-sensitive string of characters—for auto-registering a host in a domain or host group.

  VTE installers use the shared secret to add and register hosts to the DSM for a domain or host group. This method can automatically add host names or IP addresses to the DSM. This eliminates the need to verify that the host and DSM share valid certificates. You can add multiple hosts dynamically, during VTE installation and registration, with a single shared secret password.

  If you choose the Shared Secret method, ask the DSM Administrator to create a shared secret for the domain, or host group, in which the new host will reside. Then, obtain the shared secret and the validity period (one hour, day, week, or month) and register within that period.

NOTE: In the DSM Registration Shared Secret page (Hosts > Registration Shared Secret) there is an option: **Require that hosts are first added**. If you select this option, you must manually add the hosts to the DSM first.

Host name resolution

Host name resolution maps host names to an IP address. During this configuration process, enter either the FQDNs, or IP addresses, of your DSMs and hosts. If you use FQDNs, your hosts must be able to resolve their DSM host names, and the DSMs must be able to resolve to their hosts.

NOTE: The exception is if you approve of only VTE-initiated communication between the DSM and the Host. See "Determine the random number generation method" on page 29 for more information.

Setting the host name with DNS

The Domain Name Service (DNS) is the preferred method of host name resolution. If you use DNS, use the DSM and host name FQDNs for the installation and configuration procedures in this chapter.

Setting the host name without DNS

If you do not use a DNS server, perform one of the following tasks on all of the DSMs and the hosts:

- Request that the DSM Security Administrator add an entry in the `/etc/hosts` file on the DSMs for each Host. The administrator must use the DSM Admin CLI, and add an entry to *each* DSM in an HA deployment because entries in the `/etc/hosts` file are not replicated across DSMs.
- Use the IP addresses of the DSMs and hosts.

Port configuration

If a protected host must communicate with a DSM through a firewall, open the ports in the firewall as shown in the following figure.

Figure 2: Ports to open between DSM and protected host

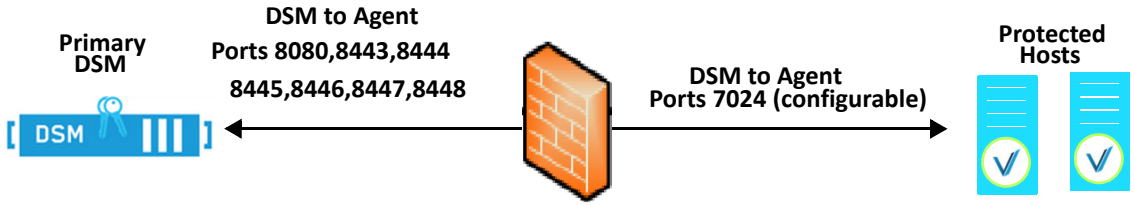

NOTE: See Table 1 to determine which of the above ports must be opened through the firewall.

Table 7 describes the communication direction and purpose of each protected host port you must open. For a complete list of ports required for the DSM, see the *VDS Administration Guide*.

Table 7: Ports to Configure

Port	Protocol	Communication Direction	Purpose
7024	TCP	DSM → Agent	Policy/Configuration Exchange
8080	TCP	Agent → DSM DSM ↔ DSM	Default TCP/IP port for HTTP that is used to exchange certificates between the DSMs in an HA configuration. Also, used once to do the initial certificate exchange between an agent host and DSM.
8443	TCP	Agent → DSM	TCP/IP port through which the agent communicates with the DSM. The agent establishes a secure connection to the DSM, via certificate exchange, using this port.
8444	TCP	Agent → DSM	Agent log messages uploaded to DSM.
8445	TCP	Agent → DSM	Used during initial certificate exchange between agent host and DSM over HTTPS. If not available, then port 8080 will be used instead.
8446	TCP	Agent → DSM	Configuration Exchange using Elliptic Curve Cryptography (Suite B)
8447	TCP	Agent → DSM	Agent uploads log messages to DSM using Elliptic Curve Cryptography (ECC)
8448	TCP	Agent → DSM	Used during initial certificate exchange between agent host and DSM over HTTPS. If not available, then port 8445 or port 8080 will be used instead.

The default port for communication between the DSM and the agent is 7024. If this port is already in use, set the port to a different number by specifying the new port during agent installation.

In the following example, the default communication port between the DSM and agent was changed from 7024 to 8000:

```
Please enter the host name of this machine, or select
from the following list. The name you provide must
precisely match the name used on the "Add Host" page of
the Management Console.

[1] centos-6-0

[2] 192.168.1.160

Enter a number, or type a different host name or IP
address in manually:

What is the name of this machine? [1]: centos-6-0:8000

You entered the host name "centos-6-0:8000"

Is this host name correct? (Y/N) [Y]:
```

Port Usage in One Way Communications Mode

By default, polling from the agent host to the DSM when running in one-way communications mode uses HTTP via port 8080. If the agent is configured to use secure polling, then polling is performed using HTTPS via port 8448 (in suite B mode) or port 8445.

Determine the installation method

There are two methods for installing VTE on Linux platforms:

- **Typical:** Most common and recommended type of installation. Use this method for installing VTE on hosts concurrently. See "Typical Install" on page 31.
- **Silent:** Create pre-packaged installations by providing information and answers to a set of installation questions. Use silent installations when installing on a large number of hosts. See "Silent Install" on page 38.

Determine the random number generation method

On Linux systems, you can choose from two methods for generating the random number used to create a certificate:

- `/dev/urandom`: This method is recommended. It is efficient and provides an acceptable level of security.
- `/dev/random`: This method provides more security but it can lengthen installation times and decrease startup speeds.

One-way communication

In some deployments, VTE might not be visible to the DSM through the normal network communications. For example, when the host that contains VTE:

- uses NAT so its IP address will change
- is protected behind a firewall
- is not permanently connected to a communication channel on the DSM
- is unable to resolve the host name to an IP address

In these situations, VTE can initiate VTE-only communication to the DSM. This feature is called one-way communication and works by making VTE poll the DSM for any policy messages or changes. Then it downloads the required changes.

The downside of one-way communication is that the DSM cannot issue any queries to VTE. For example, the DSM Administer cannot browse host directories or User IDs.

VTE Install Checklist

Use the following table to verify prerequisites and collect the information you need for the installation.

Table 8: VTE Install Checklist

Checklist item	Status

Obtain VTE installation image from Vormetric. The format for VTE file names is:
`vee-<agent_type-build-system>.bin`
Example:
`vee-fs-6.0.2-49-rh6-x86_64.bin`	

Fully Qualified Domain Name (FQDN) of the DSM	
IP address or FQDN of the host	
Administrator password for the host	
If using Shared Secret registration, obtain the following from the DSM S**ecurity Administrator**: 1) Shared secret password 2) Domain 3) Host group, if applicable 4) (Optional) Description of the host	
If using the Fingerprint registration, ask the DSM **Administrator** to add the host to the DSM and check the **Registration Allowed** check box. After checking the fingerprint, select the **Communication Enabled** check box.	
Resolved "Host name resolution" on page 25 for the hosts and DSM	
Set "Port configuration" on page 26	
Do you want to "Determine the random number generation method" on page 29	
Is "Determine the random number generation method" on page 29 required	
Synchronize the host clock to the DSM clock.	
Set the network subnet mask on the host (unless you are using one-way communication)	

Table 8: VTE Install Checklist

Checklist item	Status
Preferred DNS Server (if using FQDNs):	

Typical Install

This section describes the typical install and registration process of VTE on a Linux system.

Typically, you register VTE with the DSM as part of the installation process. However, you can postpone registration if you have a specific plan to register VTE later.

The data on the host is not protected until you complete the configuration of the host. Communication to the DSM (and retrieval of any policies and keys) cannot happen until you register VTE on the DSM, and enable communication between VTE and the DSM.

NOTE: Do not install VTE on network-mounted volumes like NFS.

Checks for Supported Kernels

At install or upgrade time, the installer checks and warns the administrator if the currently running kernel is unsupported and if it is unsupported, logs a warning in the `syslog`. The administrator is advised not to proceed with the installation in case the kernel is unsupported, but he is not prevented from installing VTE on an unsupported kernel. In the case where the administrator has chosen the silent install option or is using configuration management tools (for example, Chef or Puppet), the installer will go ahead with the installation even if the running kernel is unsupported. It will log a warning that the kernel is unsupported in the `syslog`.

The check for supported kernels is also done whenever the VTE services start up, just before loading the VTE kernel modules. If the currently running kernel is

unsupported, a warning message is logged in the `syslog` and the startup of VTE continues.

Checking for supported kernels is an advisory feature. The administrator is advised to not use VTE on an unsupported kernel, but if he decides to ignore the warning, VTE installation and startup of services will proceed as if the kernel were supported. In some cases, the VTE kernel module loading might fail or hit some unexpected error. In such cases, the administrator is advised to run VTE on supported kernels as published in the Compatibility Matrix.

Before you begin

Verify that the DSM Security Administrator has added all of the hosts that you need for VTE installation to the DSM with the following functionality enabled:

- Registration Allowed
- Communication Enabled

NOTE: If registration appears to freeze, verify that the DSM and VTE can communicate with each other over the network.

NOTE: If you are installing VTE using the shared secret method, you do not need to add the hosts to the DSM before installation.

Installation

1. Log on to the host where you will install VTE.

NOTE: You must have root access.

2. Copy or mount the installation file to the host system. If necessary, make the file executable with the `chmod` command.
3. Install VTE. A typical installation and registration, in the non-verbose mode, uses the following syntax:

```
#   ./vee-<product>-<version>-<build-system>.bin
```

Example:

```
#   ./vee-fs-6.0.2-49-rh6-x86_64.bin
```

4. The Vormetric License Agreement displays. Enter '**Y**' and **Enter** to accept.

NOTE: VTE is installed on the host, but not yet registered. The Vormetric Encryption Expert File System prompts you to register.

5. Enter the host name of the machine running the Security Server (the host name displays on the Dashboard window of the Management Console)

NOTE: If you are using the Fingerprint method, VTE's host machine must be pre-configured on the DSM with the **Reg. Allowed** option enabled.

6. Enter '**Y**' and **Enter** to accept.
7. During host registration, a new feature allows you to automatically register and enable LDT during registration. The installation script adds a question to the Installation Wizard to ask if you want to enable the automatic registration.
8. Select one of the following registration options:
   - Register VTE using the *Certificate Fingerprint* method. See "To register VTE using the Certificate Fingerprint method:" on page 33.
   - Register VTE now using the *Shared Secret* method. See "To register VTE for Linux using the Shared Secret Registration method:" on page 36.
   - Register VTE later by entering N.

NOTE: Use the command `register_host` at `/opt/vormetric/DataSecurityExpert/agent/vmd/bin/` to register without the installation program.

Register using the certificate fingerprint

To register VTE using the Certificate Fingerprint method:

1. Enter Y when you see the following prompt:

```
Do you want to continue with agent registration? (Y/N) [Y]: Y
Please enter the primary Security Server host name:
```

2. Enter the DSM FQDN and then `Y`. Ask the DSM Administrator to obtain the FQDN from the dashboard of the DSM Management Console.

   Example:

   ```
   thl1490.i.vormetric.com

   You entered the host name thl1490.i.vormetric.com
   Is this host name correct? (Y/N) [Y]: Y
   ```

3. Enter the host name when prompted:

   ```
   Please enter the host name of this machine, or select
   from the following list. If using the "fingerprint"
   registration method, the name you provide must precisely
   match the name used on the "Add Host" page of the
   Management Console.

   [1] host14.i.example.com
   [2] Host-RHEL-14.i.example.com
   [3] 10.3.14.90
   [4] 192.168.122.

   Enter a number, or type a different host name or IP address in
   manually:
   What is the name of this machine? [1]: 1
   ```

4. Enter the host name. This must match the name used on the **Add Host** page of the Management Console. Once completed, the installation prompts you for the registration method:

   ```
   You selected "host14.i.example.com".
   Would you like to register to the Security Server using a
   registration shared secret (S) or using fingerprints (F)? (S/F)
   [S]: F
   ```

5. Enter **F** (fingerprints). At the prompt, select **Y** to enable hardware association (see "Determine the random number generation method" on page 29).

   ```
   It is possible to associate this installation with the hardware of
   this
   machine. If selected, the agent will not contact the DSM (GDE
   Appliance) or use any cryptographic keys if any of this machine's
   hardware is changed. This can be rectified by running this
   registration program again. Do you want to enable this
   functionality? (Y/N) [Y]:
   ```

6. Enter **Y** or press **Enter** to enable. The following prompt displays:

```
Do you want to configure the agent to use /dev/random for the source
of random numbers? Doing so has security benefits but could cause
significant delays during installation and startup. The default
behavior (answering No) is to use /dev/urandom, which has no
associated delay.

Would you like to use /dev/random for random numbers? (Y/N) [N]:
Generating certificate...done.
Signing certificate...done.
```

7. Enter **N**. If everything is working, the install program generates certificate signing requests and lists the fingerprint of the EC (elliptic curve) CA (Certificate Authority) certificate:

```
The following is the fingerprint of the EC CA certificate. Please
verify that it matches the fingerprint shown on the Dashboard page
of the Management Console. If they do not match, it can indicate an
unsuccessful setup or an attack.

A5:6D:4B:DE:1C:ED:F7:E5:8C:C7:F3:21:58:31:F2:27:15:C5:8C:C9

Do the fingerprints match? (Y/N) [N]:
```

NOTE: If you see the error message *File System component service stopped 'Couldn't resolve host name',* it means that the DSM could not resolve the host name. See the *Vormetric* DSM Guide for information.

8. This fingerprint must match the certificate on the DSM dashboard. This verifies that nobody is intercepting and modifying traffic between the DSM and VTE. Verify this match with the DSM Administrator, then enter **Y**. VTE fingerprint for the host displays:

```
The following is the fingerprint for this agent on this host.
Please verify that it matches the fingerprint shown for this host
on the
Edit Host window of the Management Console.

01:FE:F9:37:93:36:F7:74:DD:D5:5D:EA:C8:4A:9B:9C:D0:58:73:8C

Successfully registered the Vormetric Encryption Expert File System
Agent with the primary Vormetric Data Security Server on
thl1490.i.vormetric.com.
```

9. Verify with the DSM Administrator that VTE fingerprint matches with the fingerprint shown for this host on the (Hosts > Hostname) **Edit Host** window of the Management Console. VTE is installed and registered.

10. Verify the installation by checking VTE processes on the host:
    a. Run `vmd -v` to check the version of VTE.
    b. Run `vmsec status` to display VTE processes.
    c. Look at the log files in `/var/log/vormetric/install.fs.log.<date>`, especially `vorvmd_root.log`.

Register with Shared Secret

To register VTE for Linux using the Shared Secret Registration method:

1. Verify that the DSM Administrator created a shared secret for the domain or host group in which the new host resides.
2. Enter Y when you see the following prompt:

```
Do you want to continue with agent registration? (Y/N) [Y]: Y
Please enter the primary Security Server host name:
```

3. Enter the DSM FQDN and then `Y`. Ask the DSM Administrator to obtain the FQDN from the dashboard of the DSM Management Console.

   Example: `thl1490.i.vormetric.com`

```
You entered the host name thl1490.i.vormetric.com
Is this host name correct? (Y/N) [Y]: Y
```

4. Enter the host name when prompted:

```
Please enter the host name of this machine, or select from the
following
list. If using the "fingerprint" registration method, the name you
provide must precisely match the name used on the "Add Host" page of
the Management Console.

[1] host14.i.example.com
[2] Host-RHEL-14.i.example.com
[3] 10.3.14.90
[4] 192.168.122.

Enter a number, or type a different host name or IP address in
manually:
What is the name of this machine? [1]: 1
```

5. Enter the host name. This host name must match the name used on the **Add Host** page of the Management Console (adding the host name is not needed for the shared secret method). The installation prompts you for the registration method:

```
You selected "host1490.i.vormetric.com".
Would you like to register to the Security Server using a
registration shared secret (S) or using fingerprints (F)? (S/F)
[S]: S
```

6. Enter S (Shared Secret). You are prompted for the following information (examples are in italics—use your own system information):

```
What is the registration shared secret?
Please enter the domain name for this host: <assigned-domain-name-
in-DSM>
Please enter the host group name for this host, if any:
Please enter a description for this host: Linux RH-6

Shared secret  : *********
Domain name  : <assigned-domain-name-in--DSM>
Host Group  : (none)
Host description  : Linux RH-6
Are the above values correct? (Y/N) [Y]:y
```

7. If the Shared Secret information is correct, enter **Y**. Enter appropriate information when prompted for enabling hardware association (see "Determine the random number generation method" on page 29).

```
It is possible to associate this installation with the hardware of
this
machine. If selected, the agent will not contact the DSM or use any
cryptographic keys if any of this machine's hardware is changed.
This
can be rectified by running this registration program again.
Do you want to enable this functionality? (Y/N) [Y]:
```

8. Enter Y or press **Enter** to enable. If everything is working properly, the install program generates certificate signing requests and then generates the signed certificates. Unlike the fingerprint method, the fingerprints do not display for verification:

```
Generating certificate signing request for the kernel
component...done.
Signing certificate...done.
Generating EC certificate signing request for the vmd...done.
Signing certificate...done.
Generating EC certificate signing request for the vmd...done.
Signing certificate...done.
Successfully registered the Vormetric Encryption Expert File System
Agent with the primary
Vormetric Data Security Server on thl1490.i.vormetric.com.
Installation success.
[root@host15101 Downloads]#
```

9. Verify the installation by checking the VTE version on the host, type:

```
# vmd -v
```

Silent Install

This section describes how to perform a silent installation of the VTE on a single host. The silent installation automates the installation process by storing the answers to installation and registration questions in a separate file that you create. You can also use the silent installation to install VTE on multiple hosts simultaneously.

Before you begin

The silent install method installs VTE on the host, and registers the host with the DSM you specify in the silent installation file.

For the Fingerprint Registration method, the DSM Administrator must add all hosts on which you will install VTE to the DSM. The following functionality must be enabled:

- Registration Allowed
- Communication Enabled

For the Shared Secret Registration method, you may not need to add hosts to the DSM.

Create the silent installation file

The following table shows the required and optional environment variables to be entered in the silent installation file. You can store this file anywhere on your system.

- Initiate a silent install by using the `-s` option in the install command.
- Use `-i` to install VTE and suppress the license text so that it is non-interactive.

Table 9: Register host options for silent install

Variable	Description	Required?
`AGENT_HOST_NAME`	FQDN of this VTE's host	Yes, if HOST IP is being registered.

Variable	Description	Required?
AGENT_HOST_PORT	The port number for VTE. Ignored for other agents.	No
AGENT_USEIP	Uses IP address instead of host name	No
ENABLE_DOCKER	Set to 1 to automatically enable and register Docker during silent install.	No
ENABLE_LDT	Set to 1 to automatically enable and register LDT during silent install.	No
HOST_DESC	Specifies a host Description on the **Hosts** page of the DSM Management Console. Works only with SHARED_SECRET.	No
HOST_DOMAIN	Specifies domain for the shared secret. Required if using Shared Secret method.	Yes
HOST_GROUP	Specifies the optional host group for the shared secret.	No
ONEWAY_COMMS	Set to '1' when VTE-initiated-only communication is required	No
SERVER_HOSTNAME	FQDN of DSM	Yes
SHARED_SECRET	Specifies the passphrase for a shared secret registration. See "Determine your VTE registration method" on page 25	Yes
STRONG_ENTROPY	Use /dev/random on Linux. Set to '1' if desired	No
USEHWSIG	Associate hardware to keys+certs. Set to '1' if desired.	No (default: false)

Silent Install with Shared Secret Registration method

1. Create a parameter file and store it on your system. Following is an example file containing the FQDN of the DSM and the FQDN of the host on which you will install VTE. In this example, the file is called `unattended.txt`.

   Example:

   ```
   SERVER_HOSTNAME=linux64.example.com
   AGENT_HOST_NAME=RH6.example.com
   SHARED_SECRET=Shallacl12345#
   USEHWSIG=1
   HOST_DESC="Linux RH-6"
   ```

2. Log on as an administrator to the host where you will install VTE.
3. Copy or mount the installation file to the host system.
4. Start the installation. Type:

```
# ./vee-<product-version-build-system>.bin -s -i
<dir>/unattended.txt
```

Example:

```
# ./vee-fs-6.1.0-110-rh6-x86_64.bin.bin -s /tmp/unattended.txt
```

5. Verify the installation by checking VTE processes on the host:

   a. Run `vmd -v` to check the version of VTE matches that just installed.

   b. Run `vmsec status` to display VTE kernel status.

   c. Look at the log files in `/var/log/vormetric`, especially `install.fs.log.<date>`, `vorvmd_root.log`.

Silent Install with Fingerprint Registration method

1. Create a parameter file and store it on your system. Following is an example file containing the FQDN of the DSM and the FQDN of the host on which you will install VTE. In this example, the file is called `unattended.txt`.

   Example:

   ```
   SERVER_HOSTNAME=linux64.example.com
   AGENT_HOST_NAME=RH6.example.com
   ```

2. Log on as an administrator to the host on which you will install VTE.
3. Copy or mount the installation file to the host system.
4. Start the installation. Type:

   ```
   # ./vee-<product-version-build-system>.bin -s
   <dir>/unattended.txt
   ```

   Example:

   ```
   #  ./vee-fs-6.1.0-110-rh6-x86_64.bin.bin -s /tmp/unattended.tx
   ```

   The fingerprint displays:

   ```
   D9:0E:B5:FF:51:F8:8F:2F:C9:F1:B0:74:5C:09:5B:45:BF:DA:01:9E
   ```

5. Verify the installation by checking VTE processes on the host:

   a. To check the version of VTE, type:

   ```
   # vmd -v
   ```

   b. To display VTE processes, type:

   ```
   # vmsec status
   ```

c. Look at the log files in `/var/log/vormetric/install.fs.log.<date>`, especially `vorvmd_root.log`.

Automatically Registering LDT and Docker

The registration process now lets you register LDT and Docker during registration, instead of after.

For Linux: there are two new questions during registration:

```
Do you want this host to have docker support enabled on the server?
(Y/N)  [N]:
```

```
Do you want this host to have LDT support enabled on the server?
(Y/N)  [N]:
```

For Windows, there is an option with the following label:

```
Enable LDT Feature (FS agent only)
```

Automatically Registering LDT and Docker During Silent Installation

If you want Docker or LDT to be automatically registered during a silent install, you can set the register_host script to do that.

1. Change to the following directory:

   `# cd /opt/vormetric/DataSecurityExpert/agent/vmd/bin`

2. In a terminal, edit the `register_host` file.
3. Find the section entitled: `ask_enable_feature()`
4. Set the following:

   `ENABLE_DOCKER = 1`

   `ENABLE_LDT = 1`

NOTE: Docker and LDT are not compatible. You cannot use them both on the same system.

5. Save the file.
6. Run the Silent Install as described in the section: "Silent Install" on page 38.

Tracking and Preventing Local User Creation

VTE allows you to track attempts to change user authentication files. This includes, but is not limited to user creation, modification, and deletion, or to deny users.

All VTE versions enable detection and prevention of user accounts on the local host. You can deploy any 5.x or 6.x DSM for protection of the Linux host.

VTE provides the host settings, `protect` and `audit` for this purpose. The `audit` setting is set to on, by default. It audits access to the system credential files but does not prevent account creation. The `protect` setting both audits and prevents local user account creation. You must manually enable the `protect` setting for tracking and prevention of local user account creation.

You can tag the following files with either `audit` or `protect`:

```
/etc/passwd
/etc/group
/etc/shadow
/etc/gshadow
```

The **Protect** setting supersedes the **Audit** setting if both tags are applied to the same file.

NOTE: You do not have to restart VTE after applying or removing these host settings.

This VTE for Linux feature does not require a matching DSM version. You can use a VTE for Linux installation with a v5.x DSM. However, Vormetric highly recommends that you use the this VTE feature with a v6.x DSM. Although a VTE 6.x Linux installation can use this protection feature with a v5.x DSM, audit messages are absent on the v5.0 DSM.

Linux Package Installation

This section describes how to extract and run Linux packages directly so that the Vormetric VTE installation integrates with the distribution software.

Caution: Do not use package installation for SUSE Linux. Instead, use the typical installation or the silent installation.

To extract and run the RPM file

The Vormetric installation `bin` files contain the native packages. Extract them by running the `bin` file with the `-e` flag.

1. Log on to the host system as root and copy or mount the installation file to the host system.
2. Extract the RPM file. Type:

   ```
   # ./vee-<product-version-build-system>.bin -e
   ```

   Example:

   ```
   # ./vee-fs-6.1.0-110-rh6-x86_64.bin.bin -e

    Contents extracted.

   # ls *rpm

   vee-fs-6.1.0-110-rh6-x86_64.bin.rpm
   ```

3. To start the installation using the `RPM` file, Type:

   ```
   # rpm -ivh vee-fs-6.1.0-110-rh6-x86_64.bin.rpm
   ```

4. Follow the prompts until installation and registration are complete.

Restricted Mode

You can install VTE in restricted mode. This mode prevents users, other than root, from accessing the following directories:

- /var/log/vormetric
- /opt/vormetric/DataSecurityExpert

Accessing Utilities

Restricted Mode prevents non-root users from running the following utilities:

- agenthealth
- agentinfo
- check_host
- register_host utility
- secfsd
- vmd
- vmsec
- voradmin

VTE permissions in restricted mode

The following addresses VTE permissions in restricted mode on systems that also use key agents.

Key Agent or VKM

- On systems where VTE is installed in restricted mode, you cannot install key agent (pkcs11) or VKM.
- On systems where key agent (pkcs11) or VKM are already installed, you cannot install VTE in restricted mode.

Restricted Mode installation

To install in restricted mode, use the -r option.

```
# . /vee-<product><version><build-system>.bin -r
```

Example

```
# . /vee-fs-6.1.0-110-rh6-x86_64.bin -r
```

RPM Installation

If installing from an RPM directly, prior to installation, type:

```
# export VOR_RESTRICTED_INSTALL_MODE=yes
```

Upgrade in Restricted Mode

The upgrade mode is the same as the installation mode.

Caution: If you install or upgrade in restricted mode, you cannot revert to unrestricted mode without uninstalling VTE.

Restrictions

Linux does not allow you to guard the following directories:

```
<secfs install root>/agent/secfs/
/etc
/tmp
/usr
/usr/lib
/usr/lib/pam
/var/log/vormetric
```

Linux does not allow you to guard the following directories and all of their subdirectories:

```
<install root>/agent/secfs/bin
<secfs install root>/agent/vmd
/etc/vormetric
/etc/pam.d
/etc/security
/usr/lib/security
/etc/rc*
```

Uninstalling VTE

This section describes how to uninstall VTE on a Linux host.

Before Removing VTE from a Linux host

Consider the following before removing a VTE from a host machine.

- Stop all applications from running on locations where hosts are applied.
- Before you remove VTE, decrypt any data you want to use after uninstall. Once VTE software is removed, access to data is no longer controlled. If data was encrypted, it will remain encrypted. If decrypted or copied out of the host, the data is visible as clear text.
- The DSM Administrator must evaluate the current hosts in the *Guard FS* tab to avoid data loss or compromise.
- The DSM Administrator must remove **System Locked** and **FS Agent Locked** settings for this host (if set).
- Vormetric recommends that you remove all GuardPoints.
- VTE for Linux must be removed from the host before the host is removed from the DSM.
- Database applications like DB2, and Oracle can lock the user space while they run. If a VTE installation fails because a host is in use, determine which applications are using the hosts and stop them. Then run the uninstall again.
- Commands like `fuser` and `lsof` might not reveal an active host because they detect active usage, not locked states. Although it may appear that a host is inactive, it may be in a locked state. Under this condition, software removal may fail and an error like the following may display:

  ```
  /home: device is busy.
  ```

To remove VTE from a Linux host

1. Stop any application accessing files in the host.
2. Log on to the host as root with system administrator privileges.
3. Change directory to an unguarded location (for example, `/..`)

Caution: Do not change (`cd`) into the `/opt/vormetric` directory or any directory below `/opt/vormetric`. If you are in `/opt/vormetric`, or any directory below `/opt/vormetric`, the package removal utility may fail and return the following message:

```
...
You are not allowed to uninstall from the
/opt/vormetric
directory or any of its sub-directories.
Agent uninstallation was unsuccessful.
```

4. Start the uninstall. Type:

   `# /opt/vormetric/DataSecurityExpert/agent/vmd/bin/uninstall`

   ```
   Would you like to uninstall the vee-fs package? (Y/N)
   [Y]: Y Success!
   ```

Upgrade

This section describes the generic instructions for upgrading VTE. For specific instructions, refer to the VTE release notes.

To upgrade VTE

1. Stop any application accessing files in the host.
2. Log on to the host where you will upgrade VTE. You must have root access.
3. Copy or mount the installation file onto the host system.
4. Start the upgrade. Type:

   `# ./vee-<product-version-build-system>.bin`

Example: `vee-fs-6.1.0-110-rh6-x86_64.bin.bin`

5. Type **Y** and press **Enter** to accept the Vormetric License Agreement. The upgrade proceeds.
6. Follow the prompts. During an upgrade, the following message displays:

   ```
   Upgrade detected: this product will be stopped and restarted.
   ```

```
Do you wish to proceed with the upgrade? (Y/N) [Y]: y
Installation success.
```

7. Type `y` or press **Enter** to complete the upgrade. You will not do the registration steps since VTE is already registered with the DSM.

Scheduled Upgrade

This section describes how to run the Scheduled Upgrade feature on VTE for Linux systems. It contains the following sections:

To support the Scheduled Upgrade feature, some options have been added to the `voradmin` command.

NOTE: You can schedule an upgrade on reboot with VTE for Linux v6.1.2.216+ or 6.1.3 release or higher.

NOTE: Scheduled upgrade on reboot can be scheduled for VTE builds that follow the Linux 6.1.3 GA release build.

NOTE: Scheduled upgrade on reboot is supported on VTE 6.1.3.x and later releases. It is supported on the following distros: RH6, SLES11, and AmazonLinux (`inittab` framework); RHEL7, SLES12, SLES15, Ubuntu16, Ubuntu18 (`systemd` framework). For more information, refer to the VTE for Linux v6.1.3 Release Notes and the Compatibility Matrix for Linux/UNIX.

NOTE: Supported distros are RHEL6 and RHEL7 for the 6.1.2-WF release. VTE already (pre-v6.1.3) imposes a restriction requiring it to be connected to DSM at upgrade time. The VTE v6.1.3 Scheduled Upgrade feature retains this restriction. The restriction just mentioned internally constrains the software to ensure that networking services are up when doing the actual upgrade.

NOTE: Scheduled upgrade on reboot is not supported on HDFS nodes.

Warnings for VTE/Linux

1. Incompatible kernel upgrade on reboot will prevent `secfs` services from starting.

   For example, say you have VTE installed on a system running a kernel that is compatible with VTE. Now, suppose you upgrade the kernel version of the system.

   On the next reboot, the new kernel that is running is not compatible with the installed VTE version. In this case, the `secfs` module fails to load due to an incompatible kernel version.
2. Behavior if the upgrade on reboot fails due to crash or power failure (this is similar to a failure during a normal upgrade).

   a. If a crash/power failure occurs before the upgrade executes, the upgrade will not take place, and the currently installed VTE version continues to run after the reboot. To upgrade, the system needs to be restarted again.

   b. If a crash/power failure occurs during the upgrade execution, the worst case is that VTE might enter an inconsistent state. The only way to restore would be to enter single-user mode and delete all the files/directories related to VTE.

   c. If a crash/power failure occurs after the upgrade executes successfully, then the new version will be running on the next reboot. No user intervention is required in this case.
3. As with prior VTE versions, DSM connectivity is required during upgrade.
4. All databases must be configured to automatically stop before VTE services stop during reboot/shutdown.

Basic Case: Using the Scheduled Upgrade feature

Following is a sample workflow on how to schedule an upgrade on reboot of VTE on Linux.

To check which version of VTE for Linux you currently have installed, do the following:

```
$ vmd -v
Version 6, Service Pack 2
6.2.0.9000
2018-11-15 21:02:06 (IST)
Copyright (c) 2009-2018, Vormetric. All rights reserved.
$./vee-fs-6.2.0.9001-rh7-x86_64.bin -u
```

The above command specifies a higher version of the VTE, and the -u option specifies that you want to schedule an upgrade on reboot on the specified binary.

Upon running the above command, the installation script runs and displays the license agreement.

This launches the License agreement. When prompted, you need to enter “Y” to proceed with the installation, or “N” not to proceed.

```
Do you accept this license agreement? (Y/N) [N] Y

Upgrade on reboot detected: this product will be upgraded on
shutdown/reboot.
Do you wish to proceed with the upgrade? (Y/N) [Y/N) [Y]: Y
Created symlink from /etc/systemd/system/multi-
user.target.wants/secfs-upgrade.service to
/usr/lib/systemd/system/secfs-upgrade.service.
Successfully scheduled upgrade on reboot to build 6.2.0.9001.
```

NOTE: Appropriate logs will be logged in syslog.

```
$ reboot
```

Log in and verify that the upgrade was successful.

NOTE: Appropriate logs will be logged in syslog.

```
$ vmd -v
Version 6, Service Pack 2
6.2.0.9001
2018-11-15 21L96L49 (IST)
Copyright (c) 2009-2018 Vormetric. All rights reserved.
```

NOTE: Appropriate logs will be logged in syslog.

This completes the workflow for the basic upgrade on reboot feature.

```
$ vmd -v
Version 6, Service Pack 2
6.2.0.9001
2018-11-15 21:076:40 (IST)
Copyright (c) 2009-2018, Vormetric. All rights reserved.
```

Another option of the scheduled upgrade command is to use the -y option with the -u option. This lets you avoid having to display and agree to the License Agreement. You could use this command to schedule an upgrade on reboot for a higher or later build, for example, **vee-fs-6.2.0-9002-rh2-x86_64.**

```
$ ./vee-fs-6.2.0-9002-rh2-x86_64 -uy
Linux Kernel 3.10.0-862.el7.x86_64 is supported with the VTE
version 6.2.0.9002 being installed.
Created symlink from /etc/systemd/system/multi-
yuser.target.wants/secfs-upgrade.service to
/usr/lib/systemd/system-secfs-upgrade.service.
Successfully scheduled upgrade on reboot to build 6.2.0.9002.
```

Performing an Upgrade Manually when an upgrade is already scheduled

The following describes how to perform an upgrade manually when an upgrade is already scheduled.

Use the following command to check what build is currently installed on your system:

```
$ vmd -v

Version 6, Service Pack 1
6.1.2.9007
2018-11-20 11:09:40 (IST)
Copyright (c) 2009-2018, Vormetric. All rights reservved.
```

Then schedule an upgrade on reboot of the 6.1.2-9008 build, and give options so you will not have to accept the license agreement (`-uy`).

```
$ ./vee-fs-6.1.2-9008-rh7-=x86_64.bin -uy

Created symlink from /etc/systemd/system/multi-
user.target.wants/secfs-upgrade.service to
/usr/lib/systemd/system/secfs-upgrade.service.

Successfully scheduled upgrade on reboot to build 6.1.2.9008.
```

Now the scheduled upgrade on reboot is supposed to happen on the next reboot cycle.

Consider if the administrator tries to do a manual upgrade before the next scheduled upgrade on reboot.

```
$ ./vee-fs-6.1.2-9008-rh7-x86_64.bin
```

At this point the License agreement displays, which is a couple of screens. At that point, you see:

```
Do you accept this license agreement? (Y/N) [N]: Y
```

At this point, you will get a warning:

```
WARNING: upgrade on reboot is already scheduled for 6.1.2.9008.

Do you want to cancel scheduled upgrade on reboot ? (Y/N) [Y] :
```

At this point, you need to enter “N” for No, or “Y” for yes, to cancel the scheduled upgrade on reboot. If you enter “N”, you will get:

```
Already scheduled upgrade on reboot remains intact.
```

```
Installation failure.
```

What happens here is that the manual installation will fail, because you specified that you wanted the scheduled upgrade on reboot to continue.

Alternatively, if you enter “Y” for yes (you do want to cancel the scheduled upgrade on reboot), you will get the following:

```
Removed symlink /etc/systemd/system/multi-user.target.wants/secfs-
upgrade.service.

WARNING: upgrade on reboot is cancelled for 6.1.2.9008. Proceeding
with manual upgrade.

Upgrade detected: this product will be stopped and restarted.

Do you wish to proceed with the upgrade? (Y/N) [Y]: Y

.............

Upgrade success.
```

At this point, the manual upgrade succeeded.

Voradmin commands available to run the scheduled upgrade feature

This section describes how to use some `voradmin` commands to run the upgrade on reboot feature.

```
$ vmd -v

Version 6, Service Pack 2

6.2.0.9001

2018-11-15 21L96L49 (IST)

Copyright (c) 2009-2018 Vormetric. All rights reserved.
```

To schedule an upgrade on reboot , use the `voradmin upgrade schedule` command.

```
$ voradmin upgrade schedule ./vee-fs-6.2.0-9002-rh7-x86_64.bin -y
```

The license agreement displays. When prompted to accept this license agreement, enter “Y”.

```
Do you accept this license agreement? (Y/N) [N]: Y

Linux Kernel 3.10.0-862.el7.x86_64 is supported with the VTE
version 6.2.0.9002 being installed.

Upgrade on reboot detected: this product will be upgraded on
shutdown/reboot.

Do you wish to proceed with the upgrade? (Y/N) [Y]:
```

```
Created symlink from /etc/syystemd/system/multi-
user.target.wants/secfs-upgrade.service to
/usr/lib/systemd/system/secfs-upgrade.service.

Successfully scheduled upgrade on reboot tobuild 6.1.0.9002.

$
```

To get information about a scheduled upgrade:

```
$ voradmin upgrade show

Upgrade on reboot is currently scheduled.

Current VTE version is 6.2.0.9000, upgrade on reboot scheduled for
VTE 6.2.0.9002.
```

At this point, reboot so that the scheduled upgrade on reboot can run.

```
$ reboot
```

Log in and verify that the upgrade was successful.

```
$ vmd -v

Version 6, Service Pack 2

6.2.0.9002

2018-11-15 21:09:27 (IST)

Copyright (c) 2009-2018, Vormetric. All rights reserved.

$
```

You can use the following command to cancel any existing scheduled upgrade on reboot, at any point in time:

```
$ voradmin upgrade cancel

Removed symlink /etc/systemd/system/multi-user.target.wants/secfs-
upgrade.service.

Successfully cancelled upgrade on reboot

$
```

You can use the `vordadmin upgrade show` to confirm your scheduled upgrade. There can be only one upgrade scheduled at a time.

```
$ voradmin upgrade show

Upgrade on reboot has not been scheduled.
```

To schedule an upgrade on reboot, use command:

```
voradmin upgrade schedule <path to VTE installer> [-y]
```

$

4 Installing VTE on Hadoop

This chapter describes how to protect an HDFS cluster with VTE. It contains the following sections:

Overview

The Hadoop Distributed File System (HDFS) is a file system that supports large files and directory structures distributed across hundreds, or even thousands, of commodity DataNode hosts in a cluster. Previously, VTE could only protect directories and files on the local file system rather than the actual HDFS files and directories. Now, VTE can protect files and directories.

DSM Administrators can:

- Define an encryption policy for HDFS files and directories in HDFS name space
- Selectively encrypt HDFS folders with different keys providing multi-tenancy support.
- Define user-based I/O access control rules for HDFS files in HDFS name space.

At the heart of an HDFS cluster is the *NameNode* that provides the framework to support a traditional hierarchical file and directory organization. The NameNode is a master server that manages the HDFS name space and regulates access to files by clients.

HDFS files are split into one or more data blocks that are distributed across DataNode hosts in a cluster. The NameNode maintains the namespace tree and the mapping of data blocks to DataNodes. To deploy VTE, install VTE on all the NameNode and DataNode hosts in a cluster.

Overview of VTE on HDFS

This section lists the high-level steps for implementing VTE protection on your HDFS. The process requires that the HDFS Administrator and DSM Security Administrator to work in tandem to complete separate tasks.

You can keep the HDFS cluster alive a active if you enable HDFS data replication and activate the nodes individually. Following are the high-level steps:

HDFS Administrator:

1. Compile a list directories specified by `dfs.datanode.data.dir`. If these directories do not already exist in the NameNode local file system, create them.
2. Pass the directory list to the DSM Security Administrator.
3. Ask DSM Security Administrator to:

   a. Add the NameNode to the HDFS Host Group.

   b. Create a GuardPoint for the HDFS Host Group on each of these directories.

DSM Security Administrator:

1. Create an HDFS host group to contain the HDFS nodes.
2. Create a host group GuardPoint on each of the datanode directories obtained from the previous step.
3. Add the NameNode to the HDFS Host Group.

HDFS Administrator:

1. For each DataNode, take the node offline and perform a data transformation (see the *VTE Data Transformation Guide*).
2. Ask the DSM Security Administrator to add the DataNode to the host group. (After the DataNode is added to the host group, it can be brought online.)
3. Repeat this process until all of the nodes have been encrypted and added to the HDFS Host Group.
4. Modify the host group policies to protect specific HDFS files and directories, as needed.

VTE on HDFS implementation assumptions

This chapter assumes the following:

- You have installed, configured and registered VTE on all of the NameNodes and DataNodes in the Hadoop cluster.
- The HDFS Administrator has knowledge and experience with HDFS and Ambari.
- The DSM Administrator has knowledge and experience with VTE
- The two can work and communicate in tandem with each other.

Implementing VTE on HDFS

This section describes how to implement VTE protection on your HDFS NameNode or DataNode. If you enable HDFS data replication, protecting one node at a time allows you to maintain the HDFS cluster. When all the nodes are configured, you can create GuardPoints on specific HDFS files and directories.

These instructions require that the HDFS Administrator and DSM Security Administrator work in tandem to complete separate tasks.

NOTE: The instructions below assume that each NameNode and DataNode exist on their own separate host. If you have a NameNode and DataNode on the same host, see "Implementing VTE on HDFS on a single host" on page 66.

Configure the HDFS NameNodes

The first step to implementing VTE on HDFS is for the **HDFS Administrator** to compile a list of the DataNode HDFS local file system directories, and create them on the NameNode local file systems. After this, the DSM **Security Administrator** must add the NameNodes to an HDFS Host Group:

1. Compile a list of directories specified by `dfs.datanode.data.dir`. Obtain this from `hdfs-site.xml` or using Ambari go to:
   HDFS > **Configs** > **Settings** > **DataNode** > **DataNode directories**

Figure 3: DataNode directories on Ambari

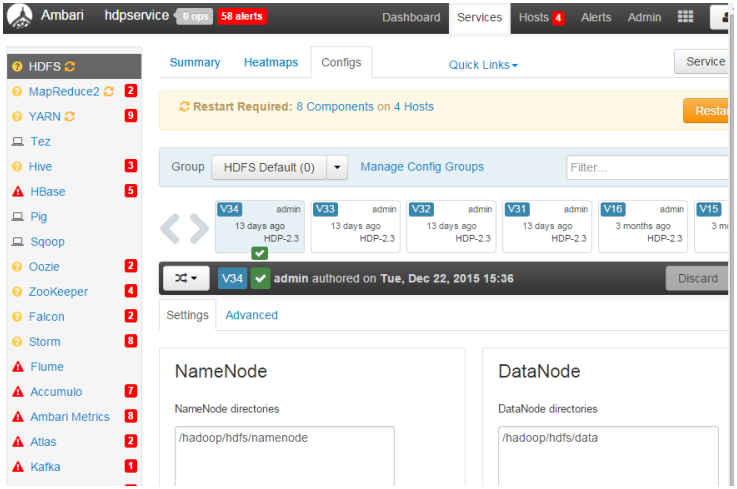

2. If these directories do not already exist in the NameNode local file system, create them on each NameNode in your Hadoop cluster
3. Pass the following information to the DSM Security Administrator:
   a. The directory list and instructions to create a GuardPoint for the HDFS Host Group on each of these directories.
   b. Instructions to add the NameNodes IP addresses or host names to the HDFS Host Group.

Create an encryption zone in HDFS name space for AWS EMR

HDFS requires the following manual steps if you want to have an encryption zone in the HDFS name space for AWS EMR, (Elastic MapReduce).

1. Add the following properties to `hdfs-site.xml`.

NOTE: The default .sec folder name is `/opt/vormetric/DataSecurityExpert/agent/secfs/.sec`.

```
<property>
   <name>dfs.vte.ioctl.lib</name>
   <value>vorhdfs</value>
</property>
```

```
<property>
   <name>dfs.vte.rename.check</name>
   <value>true</value>
</property>
<property>
   <name>dfs.vte.rename.check</name>
   <value>true</value>
</property>
<property>
   <name>dfs.vte.ioctl.device</name>
   <value><.sec folder name, up to the VTE installation
location></value>
         </property>
```

2. Save the file.
3. Restart HDFS NameNode and DataNode services.

Using the Original information from HDFS

Update or add the following properties to `hdfs-site.xml` if you want VTE to use the original user information from HDFS.

1. Add the following properties to `hdfs-site.xml`.

```
<property>
   <name>dfs.block.access.token.enable</name>
   <value>true</value>
      </property>
      <property>
   <name>dfs.client.read.shortcircuit</name>
   <value>false</value>
    </property>
<property>
   <name>dfs.vte.user.push</name>
   <value>true</value>
    </property>
```

2. Save the file.
3. Restart HDFS NameNode and DataNode services.

Create a HDFS Host Group and Host Group GuardPoint

After configuring the NameNodes, the next steps in activating VTE on HDFS is for the DSM **Security Administrator**.

Create an HDFS Host Group to contain the HDFS nodes:

1. In the DSM Management Console, click **Hosts** > **Host Groups** > **Add**.
2. Enter a **Host Group Name** for the Hadoop cluster.
3. Select **HDFS Cluster** for the Cluster type.
4. (Optional) Enter a description and click **Ok**.

Figure 4: Add HDFS Host Group

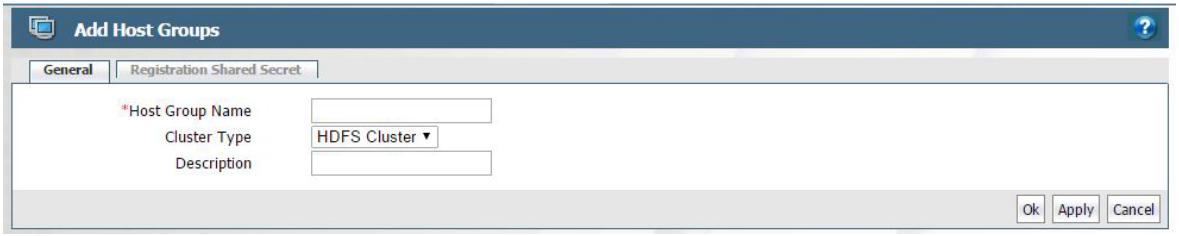

5. In the **Edit Host Group** page, click the HDFS tab.
6. For a Hadoop authentication configured as **Simple** mode, enter the NameNode URL information in the URL format: `hdfs://<host>:<port>`.

   By default the port number is 8020. Check the HDFS configuration to verify this. HDFS HA cluster requires the URLs for both active and standby.

Figure 5: Editing the HDFS Cluster Host Group in Hadoop *Simple* authentication mode

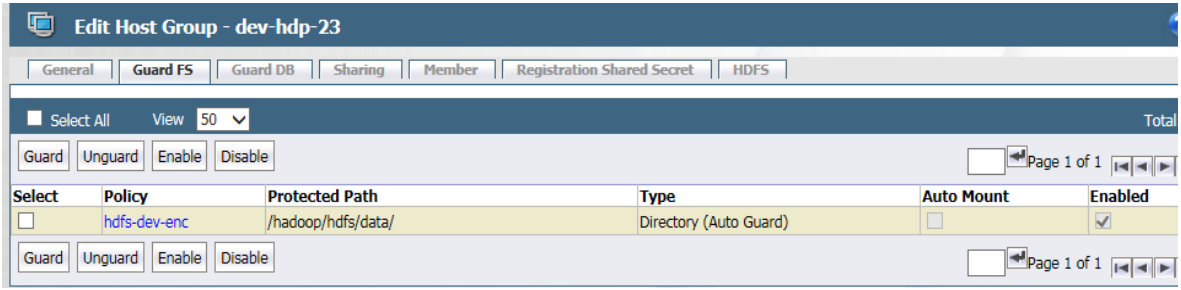

7. For Hadoop authentication configured as **Kerberos**, enter the NameNode URL information in the URL format: host name (not IP address).
8. Check the **Requires Kerberos Authentication** option and enter the following Kerberos information used for authentication:
   - **Kerberos Principal**: Unique identity to which Kerberos can assign tickets. Format is: `primary/instance@REALM`.
   - **Kerberos Realm**: Typically your domain name.
   - **KDC Host**: Hostname of your domain controller.
   - **Keytab File**: File containing pairs of Kerberos principals and encrypted keys.

Figure 6: Editing the HDFS Cluster Host Group in Hadoop *Kerberos* authentication mode

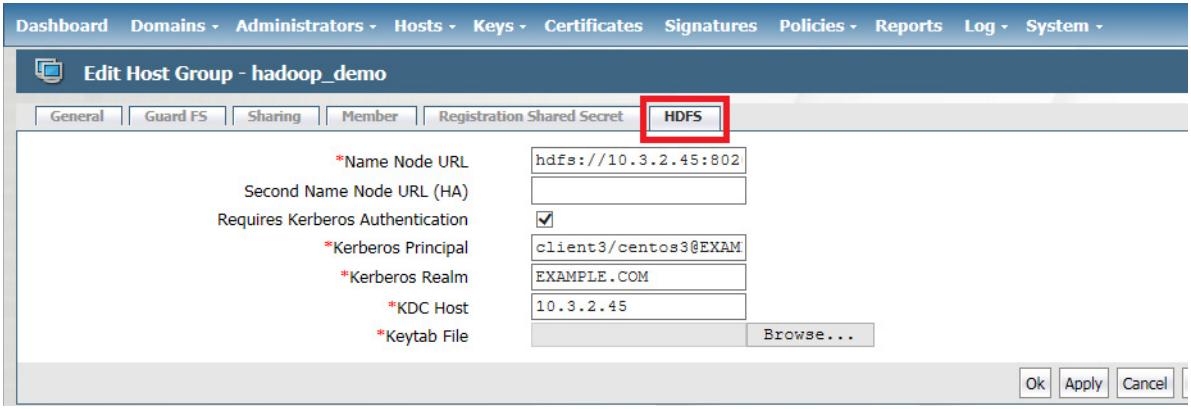

9. You can create any policy you choose, but the example encryption policy, `hdfs-dev-enc`, uses the following rules:
   - For the user set `hdfs-user`, the action is *all_ops*, the effect is *Audit*, *Apply Key*, *Permit*

- For other users the action is *READ*, the effect is *Permit*.
- For the resource set `hdfs-dev-data-1`, the key is `hdfs-dev-key-1`.
- For the resource set `hdfs-dev-data-2`, the key is `hdfs-dev-key-2`.

Figure 7: Example HDFS Host Group policy

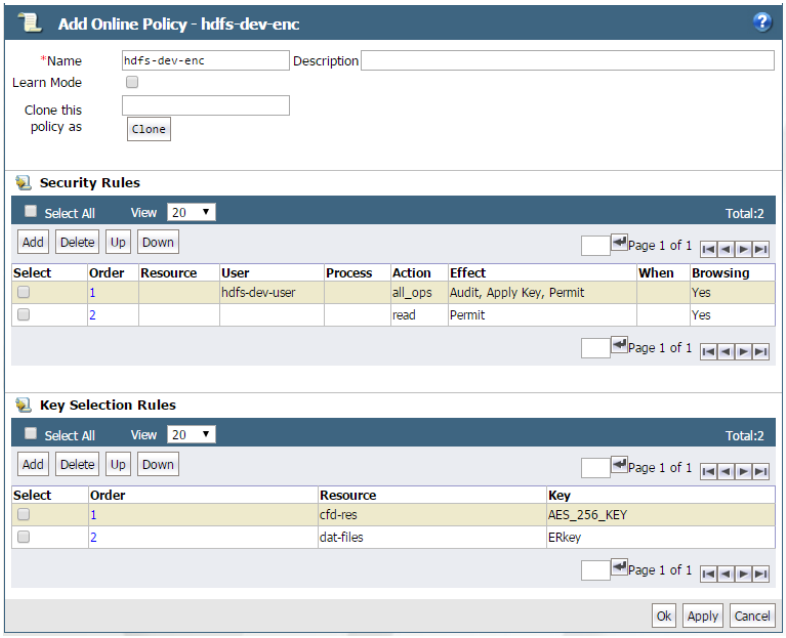

10. If you haven't already done so, add the NameNode obtained from the HDFS Admin to the HDFS Host Group.

NOTE: Vormetric highly recommends Auto Guard for HDFS. You can use manual guards, but this might result in data corruption if some nodes in a running cluster are guarded, while other are not.

Take a DataNode offline and perform data transformation

The next step in activating VTE on HDFS is to switch a DataNode to offline and transform (encrypt) its sensitive data. Once the data is transformed, the HDFS Admin can add the DataNode to the HDFS Host Group. Then they can switch the DataNode back to online. Most of these procedures are completed by the **HDFS Administrator** although one is done by the DSM **Security Administrator**.

1. **HDFS Administrator**: Switch a DataNode to offline.
2. **HDFS Administrator**: Encrypt the files in the directories specified by `dfs.datanode.data.dir` (see "Configure the HDFS NameNodes" on page 59). Read the VTE *Data Transformation Guide* for instructions on how to encrypt files and directories.
3. DSM **Security Administrator**: Create encryption keys and a data transformation policy to transform the data.

   Figure 8: "Example data transformation policy" transforms the Resource Set `hdfs-dev-data-1` from `clear_key` to `hdfs-dev-key-1`. It also transforms the Resource Set `hdfs-dev-data-2` from `clear_key` to `hdfs-dev-key-2`. Resource Set `hdfs-dev-data-1` consists of `/tmp/data1` and the Resource Set `hdfs-dev-data-2` consists of `/tmp/data2`.

Figure 8: Example data transformation policy

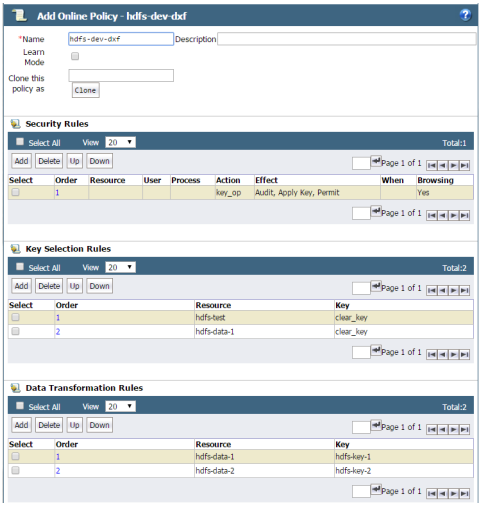

4. **HDFS Admin**: After encrypting the data in those directories, ask the DSM Security Administrator to add the DataNode to the HDFS Host Group.
5. DSM **Security Admin**: Add the DataNode host to the HDFS Host Group.
6. **HDFS Admin**: After the DataNode is added to the HDFS Host Group, bring the DataNode online.
7. Repeat this procedure for all the DataNodes in your HDFS cluster.

Implementing VTE on HDFS on a single host

It is possible, though not recommended, that an HDFS NameNode and DataNode exist as separate processes on the same host. If this is your deployment, use the following VTE deployment guidelines:

1. Configure the HDFS NameNodes (see "Configure the HDFS NameNodes" on page 59):

NOTE: The directories specified by `dfs.datanode.data.dir` already exist on the local file system so you do not have to create them.

2. Pass the following information to the DSM Security Administrator:
   - The `dfs.datanode.data.dir` directory list and instructions to create a GuardPoint for the HDFS Host Group on each of these directories.
   - Instructions to add the NameNodes IP addresses, or host names, to the HDFS Host Group.
3. Create an HDFS Host Group and Host Group GuardPoint (see "Create a HDFS Host Group and Host Group GuardPoint" on page 62):
   a. DSM Security Administrator must create an HDFS Host Group to contain the HDFS nodes.
   b. The DSM Security Administrator must create a GuardPoint for the Host Group on each of the directories specified by `dfs.datanode.data.dir`
4. Take the DataNode offline and perform a data transformation.
5. Add the NameNode/DataNode host to the Host Group.

Adding a New DataNode to a VTE-protected HDFS

Use the following procedure to add a new DataNode to a VTE-protected HDFS. If not followed, HDFS encrypted files could be exposed in cleartext.

NOTE: If you already have VTE installed on the cluster nodes **before** Ambari installs the Hadoop software, see "VTE installed on the cluster nodes before Ambari installs Hadoop" on page 67.

1. Install the HDFS client on the host. This option is available in Ambari when adding a new DataNode to the cluster.
2. Add the new node to DSM database and make sure that the host settings of the new node is the same as existing nodes in the cluster. See the *VTE Installation for Hadoop* chapter in the *VTE Installation and Configuration Guide.*
3. Install VTE on the new node, register to DSM, and run `config-hadoop.sh` to prepare the libraries. See the *Configuring Hadoop to use VTE* section in the *VTE Installation and Configuration Guide.*
4. Make sure that the data directories (specified in `dfs.datanode.data.dir` property) exist on the new node. They must have the same permission and ownership as the other existing nodes in the cluster. If necessary, create them.
5. Add the host to DSM HDFS Host Group that is guarding the cluster. This is important: do not rely on the DataNode to create the data directories as the data replication can occur before GuardPoints are in effect.
6. Add the DataNode service to the new node. Again this option is available through Ambari.
7. If using Kerberos, check that the keytab files are created correctly.
8. Start the DataNode service on the new node.
9. Execute some `hdfs dfs` shell commands to ensure that encryption/decryption of data works correctly.

VTE installed on the cluster nodes before Ambari installs Hadoop

If VTE is already installed on the cluster nodes **before** Ambari installs the Hadoop software, Ambari can mistakenly pick up the `.sec` directory in configuration steps to store the HDFS data. Make sure the following properties do not contain the `.sec` directory:

- DataNode data directory
- NameNode data directory
- Secondary NameNode checkpoint directory

- Zookeeper directory
- yarn.nodemanager.local-dirs
- yarn.nodemanager.log-dirs
- yarn.timeline-service.leveldb-timeline-store.path
- yarn.timeline-service.leveldb-state-store.path

NOTE: This list is not exhaustive. Depending on the Hadoop ecosystem packages installed, there can be others.

HDFS Upgrade with VTE

To upgrade Hadoop from version 2.6.0 to version 2.7.0 and higher, configure VTE to integrate with the new HDFS instance.

Upgrading one node at a time

Once VTE is installed and configured on the node:

1. Upgrade Hadoop.
2. Make sure that HDFS services are shut down on the node.
3. Type:

   ```
   /opt/vormetric/DataSecurityExpert/agent/secfs/hadoop/bin/config-hadoop.sh -i -y
   ```

4. Start HDFS services on the node.

Upgrade VTE with LDT in an HDFS Cluster

If you are using LDT with HDFS cluster, follow these steps when upgrading VTE, in order to maintain your LDT GuardPoints.

1. Suspend rekey on all data nodes.
2. Shutdown your namenodes/datanodes.
3. Upgrade VTE in the namenode first.

NOTE: Always upgrade namenodes before datanodes.

4. After VTE upgrade succeeds, type:
   ```
   # config-hadoop.sh -i -y
   ```
5. On the Ambari admin console, start the namenode.
6. Verify that the Vormetric java process successfully launched in the namenode. (You should not see an error message.) Type:
   ```
   # ps -ef | grep java | grep vormetric
   ```
7. Check the DSM status. The DSM should show LDT rekeyed status.
8. Check the namenode status. It should display the GuardPoint status and match the state before upgrade. Type:
   ```
   # secfsd -status guard

   GuardPoint        Policy          Type   ConfigState  Status
   Reason----------        ------          ----   -----------  -
   -----   ------
   /hadoop/hdfs/data  LDT_HDFS_Sanity  local  guarded     guarded
     N/A
   ```
9. Repeat the above steps for all of the datanodes in the HDFS cluster.

Rolling Upgrades

Hortonworks Data Platform has introduced rolling upgrades to automate the Hadoop upgrade process (`http://bit.ly/2pQrFo3`). The upgrade process is controlled by the Upgrade Pack (`http://bit.ly/2rkutvF`) that is predefined and certified by Hortonworks.

To integrate VTE with the upgrade, you need to temporarily change the Ambari scripts before performing the rolling upgrades and then restore the scripts after the upgrades.

1. On Ambari server machine, type:
   ```
   # cd /var/lib/ambari-server/resources/common-
   services/HDFS/2.1.0.2.0/package/scripts
   ```

2. Copy the `utils.py` file, type:

   # cp utils.py `utils.py.org`

3. Using a text editor, add the following commands to `utils.py`:

   if action == "start":

   if name == "namenode" or name == "datanode":

   ```
   Execute(format("/opt/vormetric/DataSecurityExpert/agent/secfs/hado
   op/bin/c onfig-hadoop.sh -i -h {hadoop_bin}/../ -j <java home> -p
   hdp -d", not_if=service_is_up, user=params.root_user)# For Redhat
   6.x, uncomment the following command#
   Execute(format("/etc/init.d/secfs secfsd restart"),
   not_if=service_is_up, user=params.root_user)# For Redhat 7.x,
   uncomment the following command#
   Execute(format("/etc/vormetric/secfs restart"),
   not_if=service_is_up, user=params.root_user)
   ```

   before

   Execute(daemon_cmd, not_if=service_is_up,
   environment=hadoop_env_exports

   The Java home of your HDFS instance should be used to replace <java home>:

   ```
   if action == "start":
     if name == "namenode" or name == "datanode":
   Execute(format("/opt/vormetric/DataSecurityExpert/agent/secfs/hado
   op/bin/config-hadoop.sh -i -h {hadoop_bin}/../ -j <java home> -p
   hdp -d"), not_if=service_is_up, user=params.root_user)

   # For Redhat 6.x, uncomment the following command

   # Execute(format("/etc/init.d/secfs secfsd restart"),
   not_if=service_is_up, user=params.root_user)

   # For Redhat 7.x, uncomment the following command

   # Execute(format("/etc/vormetric/secfs restart"),
   not_if=service_is_up, user=params.root_user)

   Execute(daemon_cmd,
           not_if=service_is_up,
           environment=hadoop_env_exports)
   ```

4. Type:

   # ambari-server restart

5. Perform rolling upgrades.

6. During the upgrade process, many of the intermediate service status checks can fail. Skip over them by clicking on **Proceed to Upgrade**.
7. Click **Finalize** to complete the upgrade. If the active NameNode fails to activate due to the incompatible HDFS layout version, manually start the NameNode with '`-upgrade`' option to correct the layout version file.

   ```
   # /var/lib/ambari-server/ambari-sudo.sh su hdfs -l -s
   /bin/bash -c 'ulimit -c unlimited ; /usr/hdp/current/hadoop-
   client/sbin/hadoop-daemon.sh --config /usr/hdp/current/hadoop-
   client/conf start namenode -upgrade'
   ```

8. If there are excessive under-replicated blocks, run the following command to isolate them and manually start the replication:

   ```
   # su - <$hdfs_user>
   # hdfs fsck / | grep 'Under replicated' | awk -F':' '{print
   $1}' >> /tmp/under_replicated_files
   # for hdfsfile in `cat /tmp/under_replicated_files`; do echo
   "Fixing $hdfsfile  :" ;  hadoop fs -setrep 3 $hdfsfile; done
   ```

9. Restart the HDFS services. Wait for the replication to complete and the NameNodes to exit safe mode.
10. When Hbase is restarted after upgrades, it tries to rename from: `/apps/hbase/data/.tmp/data/hbase/namespace` to: `/apps/hbase/data/data/hbase/namespace`, which may cause key conflict if the GuardPoint is set incorrectly (for example, `/apps/hbase/data/data` is guarded, but not `/apps/hbase/data/.tmp`). This results in Hbase shutting down.

    Before re-starting Hbase, make sure the GuardPoint policies on the Hbase files are set correctly to cover all Hbase-related files. A broader GuardPoint (`/apps/hbase/data` instead of just `/apps/hbase/data/data` and other folders) could fix this issue.
11. Check cluster upgrade by verifying the `hadoop version`.
12. Run a few mapreduce jobs and Hbase commands to make sure that the entire Hadoop stack is working properly.
13. Rename `utils.py.org` to `utils.py`

Configure the Hadoop Cluster for VTE

Configure the Hadoop cluster to use VTE before installing and configuring VTE on the nodes. Use Ambari to perform this configuration.

Create a Vormetric Configuration Group

The Vormetric Configuration Group will eventually contain all the hosts in your Hadoop cluster. At first you will create an empty group and later populate it with the hosts on which you will install and configure agents.

1. On Ambari, go to **HDFS** > **Configs** > **Manage Config Groups**.
2. Add a new configuration group *Vormetric*.
3. Make group *Vormetric* the current group.

Figure 9: VTE in a Hadoop environment

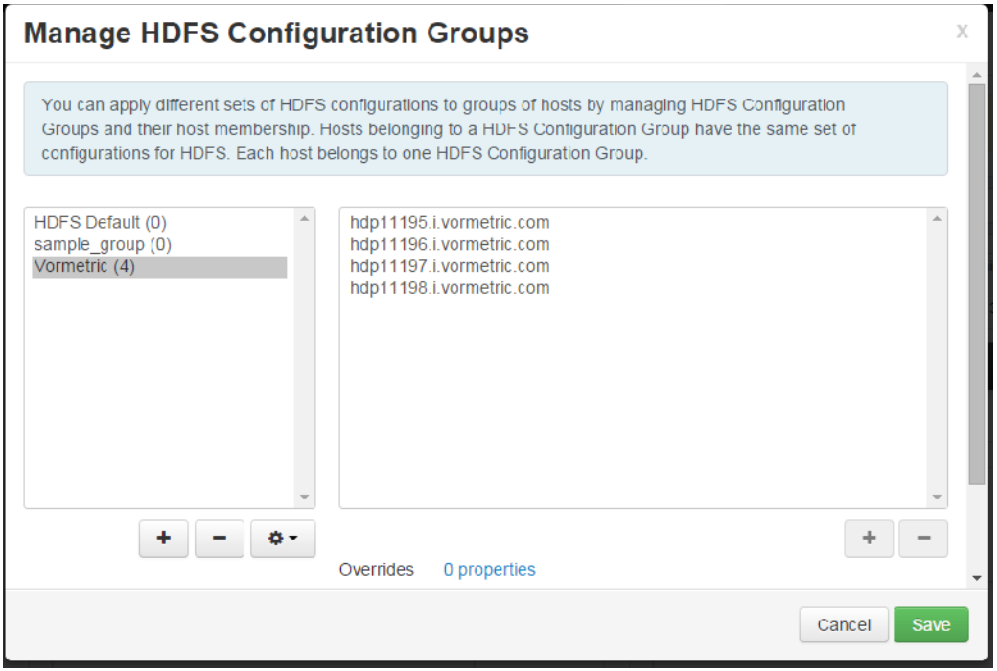

Update the Hadoop-env template with VTE settings

1. Go to **HDFS** > **Configs** > **Advanced** > **Advanced hadoop-env** > **hadoop-env template**
2. Copy and paste the original hadoop-env templates into the Vormetric template and add the following **two** export lines to specify that the VTE Java agent is instrumented into NameNode and DataNode.

```
export HADOOP_NAMENODE_OPTS="-
javaagent:/etc/vormetric/hadoop/jar/vormetric-hdfs-
agent.jar=voragent ${HADOOP_NAMENODE_OPTS}"

export HADOOP_DATANODE_OPTS="-
javaagent:/etc/vormetric/hadoop/jar/vormetric-hdfs-
agent.jar=voragent ${HADOOP_DATANODE_OPTS}"
```

Figure 10: Hadoop environment template and line specification

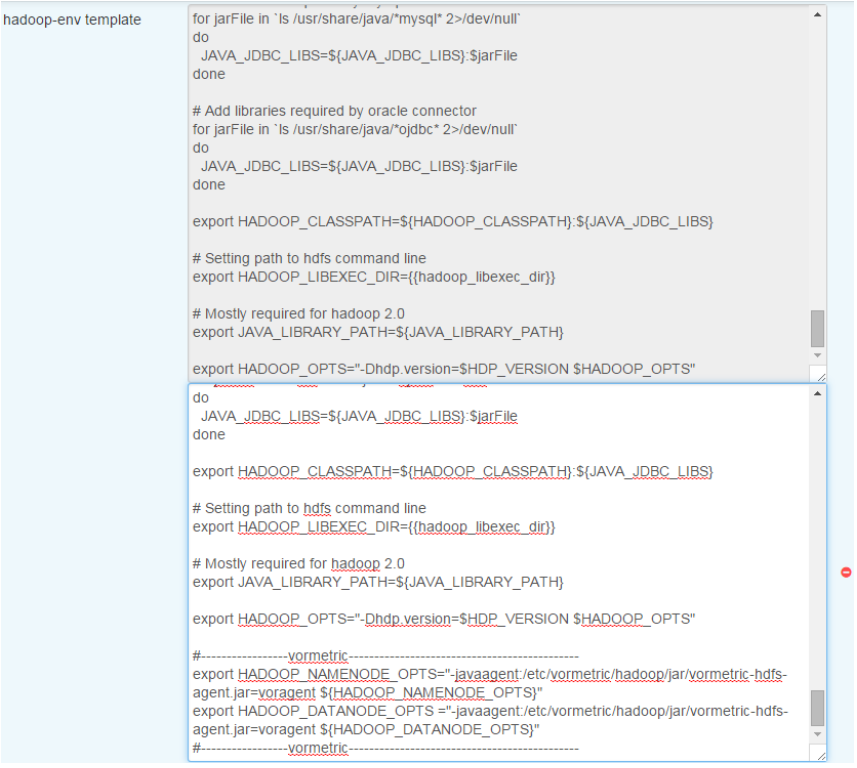

Modify the HDFS IOCTL

1. Go to **HDFS** > **Configs** > **Advanced** > **Custom hdfs-site**.
2. In the **dfs.vte.ioctl.lib** field, type: `vorhdfs`
3. In the **dfs.vte.ioctl.device** field, type:

   `# /opt/vormetric/DataSecurityExpert/agent/secfs/.sec`

Change the HDFS file rename check

1. Go to **HDFS** > **Configs** > **Advanced > Custom hdfs-site**.
2. Set **dfs.vte.rename.check** to **true**

User information push

You only need this configuration if you want to use the original user information with HDFS operation for IO access check.

NOTE: It has a performance cost.

1. Go to **HDFS** > **Configs** > **Advanced** > **Advanced hdfs-site** and uncheck **HDFS Short-circuit read**.
2. Set **dfs.block.access.token.enable** to **true**.
3. Go to **HDFS** > **Configs** > **Advanced** > **Custom hdfs-site** and set **dfs.vte.user.push** to **true**.

Create Kerberos principal for VTE

If Hadoop is configured in the secure mode with Kerberos enabled (`hadoop.security.authentication=Kerberos`), you need to create the Kerberos principal for VTE. Make the principal renewable with the `maxlife` property (Maximum ticket life) larger than 1 hour and smaller than 1 day, and Maximum ticket life smaller than Maximum renewable life (which by default is 7 days).

Configure the *keytab* file and principal with HDFS using these steps:

1. Go to **HDFS->Configs->Custom hdfs-site**
2. Set **dfs.vte.keytab.file**=*<VTE keytab file>.*
3. Set **dfs.vte.kerberos.principal**=<*VTE principal name*>.

VTE calls `kinit` to initialize the Kerberos ticket and renew the ticket once per hour. You can execute the following steps from the command line to verify that the Kerberos principal was created and configured for VTE correctly:

```
kinit -r 1440m -k -t <VTE keytab file> <VTE principle>

kinit -R
```

Uninstalling VTE for the Hadoop cluster

This section explains how to remove VTE and restore the environment back to the non-VTE cluster environment. Vormetric recommends uninstalling the agent from

HDFS nodes one by one, starting from the DataNode. Uninstall the agent from the NameNode last.

Follow the steps below:

1. Shut down one DataNode.
2. Perform the normal agent uninstall procedure
3. Go to Ambari, remove the DataNode host from the Vormetric Configuration Group.
4. Start the DataNode.
5. Delete the Vormetric Configuration Group from Ambari when agent is uninstalled from all DataNodes.
6. Repeat these steps for each DataNode.
7. Repeat these steps for each NameNode.

VTE Installation and Configuration

After configuring the Hadoop cluster for VTE:

1. Install and register VTE on the HDFS nodes.
   - You can do this to all the nodes at once, but the HDFS is unavailable during VTE installation and configuration.
   - You can also do this one node at a time. If you install and register VTE notes one at a time, you must start from NameNode, then DataNode, and always keep NameNode service up once NameNodes are configured.
2. In either case, add the FQDN of the node to the Vormetric Configuration Group, then proceed with agent installation and configuration. See “Installing and configuring VTE on an HDFS node” on page 77.
   - Modify the Host Group. See “Modifying host settings for HDFS hosts on the DSM” on page 77.
   - Configure VTE by running `config-hadoop.sh` on the HDFS node. See “Configuring Hadoop to use VTE” on page 79.
   - Review the `SecFS` configuration variables that support the HDFS name cache. “HDFS name cache” on page 81.

Installing and configuring VTE on an HDFS node

1. Using Ambari, add the FQDN of the node to the Vormetric Configuration Group. See "Create a Vormetric Configuration Group" on page 72.
2. Install, configure, and register VTE as described in "Installing VTE for Linux" on page 23.
3. Modify the host settings for each node. See "Modifying host settings for HDFS hosts on the DSM" on page 77.

Modifying host settings for HDFS hosts on the DSM

The Hadoop service can start as root and then downgrade to an unprivileged user. If the unprivileged user is not authenticated by password, VTE flags the user as fake. VTE cannot allow a user to access a resource protected by a user rule when the user is faked, even if the user matches the permit rule. Because of this, modify DSM host setting as follows:

- On Ambari, go to **HDFS** > **Configs** > **Advanced** > **Advanced core-site** and find out if **hadoop.security.authentication** mode is set to **simple** (no authentication) or **Kerberos**.

Simple Modification

To use **simple**, ask the DSM Administrator to add the following lines to the Host Group in the DSM Management Console:

NOTE: `/usr/jdk64/jdk1.8.0_40/bin/java` is the Java executable used to launch the HDFS services. Change the Java `jdk` path to reflect your end-user environment.

```
|authenticator+arg=+class=org.apache.hadoop.hdfs.server.namenod
e.NameNode|/usr/jdk64/jdk1.8.0_40/bin/java
|authenticator+arg=+class=org.apache.hadoop.hdfs.server.datanod
e.DataNode|/usr/jdk64/jdk1.8.0_40/bin/java
```

The entire host settings will look like this:

```
|authenticator|/usr/sbin/sshd
|authenticator|/usr/sbin/in.rlogind
|authenticator|/bin/login/
|authenticator|/usr/bin/gdm-binary
```

```
|authenticator|/usr/bin/kdm
|authenticator|/usr/sbin/vsftpd

|authenticator+arg=+class=org.apache.hadoop.hdfs.server.namenode
.NameNode|/usr/jdk64/jdk1.8.0_40/bin/java
|authenticator+arg=+class=org.apache.hadoop.hdfs.server.datanode
.DataNode|/usr/jdk64/jdk1.8.0_40/bin/java
```

Using Kerberos

To use **Kerberos**, ask the DSM Administrator to add the following two lines to the **Host Settings** in the DSM Management Console:

NOTE: `/usr/jdk64/jdk1.8.0_40/bin/java` and `/usr/lib/bigtop-utils/jsvc` are the Java executables used to launch the HDFS services. Change the versions accordingly to fit your environment.

```
|authenticator+arg+class=org.apache.hadoop.hdfs.server.namenode.
NameNode|/usr/jdk64/jdk1.8.0_40/bin/java
|authenticator+arg+class=org.apache.hadoop.hdfs.server.datanode.
SecureDataNodeStarter|/usr/lib/bigtop-utils/jsvc
```

The entire Host Group will look like this:

```
|authenticator|/usr/sbin/sshd
|authenticator|/usr/sbin/in.rlogind
|authenticator|/bin/login/
|authenticator|/usr/bin/gdm-binary
|authenticator|/usr/bin/kdm
|authenticator|/usr/sbin/vsftpd

|authenicator+arg=+class=org.apache.hadoop.hdfs.server.namenode.Na
meNode|/usr/jdk64/jdk1.8.0_40/bin/java
|authenicator+arg=+class=org.apache.hadoop.hdfs.server.datanode
.SecureDataNodeStarter|/usr/lib/bigtop-utils/jsvc
```

Modifying Host Group for HDFS NameNodes HA on DSM

To enable high availability (HA) for your HDFS NameNodes, ask the DSM Administrator to add the following lines to the Host Group in the DSM Management Console

NOTE: `/usr/jdk64/jdk1.8.0_40/bin/java` is the Java executable used to launch the HDFS services. Change the Java `jdk` path to reflect your end-user environment.

```
|authenticator+arg+class=org.apache.hadoop.hdfs.qjournal.server.Jo
urnalNode|/usr/jdk64/jdk1.8.0_40/bin/java

|authenticator+arg+class=org.apache.hadoop.yarn.server.application
historyservice.ApplicationHistoryServer|/usr/jdk64/jdk1.8.0_40/bin
/java

|trust+arg=+class=org.apache.hadoop.hdfs.tools.DFSZKFailoverContro
ller|/usr/jdk64/jdk1.8.0_40/bin/java
```

The entire Host Group for HA (in this example, with Kerberos) will look like this:

```
|authenticator|/usr/sbin/sshd
|authenticator|/usr/sbin/in.rlogind
|authenticator|/bin/login/
|authenticator|/usr/bin/gdm-binary
|authenticator|/usr/bin/kdm
|authenticator|/usr/sbin/vsftpd

|authenticator+arg+class=org.apache.hadoop.hdfs.server.namenode.
NameNode|/usr/jdk64/jdk1.8.0_40/bin/java
|authenicator+arg=+class=org.apache.hadoop.hdfs.server.datanode
.SecureDataNodeStarter|/usr/lib/bigtop-utils/jsvc

|trust+arg+class=org.apache.hadoop.hdfs.qjournal.server.JournalNod
e|/usr/jdk64/jdk1.8.0_40/bin/java

|trust+arg=+class=org.apache.hadoop.yarn.server.applicationhistory
service.ApplicationHistoryServer|/usr/jdk64/jdk1.8.0_40/bin/java

|trust+arg=+class=org.apache.hadoop.hdfs.tools.DFSZKFailoverContro
ller|/usr/jdk64/jdk1.8.0_40/bin/java
```

Configuring Hadoop to use VTE

1. On the HDFS node, type:

   # /opt/vormetric/DataSecurityExpert/agent/secfs/hadoop/bin/

   config?hadoop.sh –i.

   HDFS prompts you for the following information:

Hadoop product name: (i.e. *hdp)*

Hadoop product version: (i.e. *2.6.0.2.2.0.0-2041)*

Path to JAVA_HOME used by Hadoop:c(i.e. */usr/jdk64/jdk1.8.0_40)*

NOTE: Alternatively, you can use the automated installation option:
`/opt/vormetric/DataSecurityExpert/agent/secfs/hadoop/bin/config-hadoop.sh -i -p hdp -v 2.6.0.2.2.0.0-2041 -j /usr/jdk64/jdk1.8.0_40 2.`

Verify the configuration the using ?s option:

 # /opt/vormetric/DataSecurityExpert/agent/secfs/hadoop/bin/config-hadoop.sh -s

System Response:

```
Vormetric-Hadoop Configuration Status
PRODUCT_NAME=hdp- PRODUCT_VERSION=3.0.0.0-2557
HADOOP_HOME=/usr/hdp/current/hadoop-client/sbin/../-
HADOOP_VERSION=2.7.1 HADOOP_PRODUCT_VERSION=2.7.1.2.3.0.0-2557-
HADOOP_VERSION_MAJOR=2.7
LIBVORHDFS_SO=/usr/hdp/current/hadoopclient/sbin/..//lib/native/libvorhdfs.so LIBHDFS_SO=/etc/vormetric/hadoop/lib/libhdfs.so
VORMETRIC_HADOOP=/etc/vormetric/hadoop-
#----------vormetric-------------------------------------------------
export HADOOP_NAMENODE_OPTS="-
javaagent:/etc/vormetric/hadoop/jar/vormetric-hdfs-
agent.jar=voragent ${HADOOP_NAMENODE_OPTS}"v6 . . . . . 62 export
HADOOP_DATANODE_OPTS="-
javaagent:/etc/vormetric/hadoop/jar/vormetric-hdfs-
agent.jar=voragent ${HADOOP_DATANODE_OPTS}"
#--------------vormetric-----------------------------------------
/etc/vormetric/hadoop/lib/libhdfs.so ...ok
/usr/hdp/current/hadoop-
client/sbin/..//lib/native/libvorhdfs.so ...ok
/etc/vormetric/hadoop/gen-vor-hadoop-env.sh ...ok
/etc/vormetric/hadoop/vor-hadoop.env ...ok
Looks ok.
```

Vormetric Transparent Encryption Agent 6.0.3 Installation and Configuration Guide v6 . . . . .62 export HADOOP_DATANODE_OPTS="-javaagent:/etc/vormetric/hadoop/jar/vormetric-hdfs-agent.jar=voragent ${HADOOP_DATANODE_OPTS}"

```
#--------------vormetric ---------------------------------------
/etc/vormetric/hadoop/lib/libhdfs.so ...ok /usr/hdp/current/hadoop
```

```
client/sbin/..//lib/native/libvorhdfs.so ...ok
/etc/vormetric/hadoop/gen-vor-hadoop-env.sh ...ok
/etc/vormetric/hadoop/vor-hadoop.env ...ok

Looks ok
```

Verify secfsd is running with Hadoop environment

Using a text editor, view the `/etc/init/secfsd?upstart.conf` file. The file should contain `env` entries, type:

```
# cat /etc/init/secfsd?upstart.conf
```

HDFS name cache

Obtaining the HDFS file name from the NameNode is network intensive, so the map from HDFS block file name to HDFS file name is cached in a hash table. The following `secfs` configuration variables are used to support the hash cache. They are provided in case you need to tune the memory management of the name cache for better performance.

hdfs_cache_entry: Default is 1,024.000, which could cover up to 125TB HDFS data because the default HDFS block size is up to 128MB (128MB * 1024000 = 125TB).

hdfs_cache_bucket: Default is 10240.

hdfs_cache_timeout: Default to 30 minutes.

hdfs_cache_interval: Default wake up interval for a worker thread to update the cache entry whose timeout has expired is 10 seconds.

On Linux, you can configure each secfs configuration variable in the following file:

```
# /opt/vormetric/DataSecurityExpert/agent/secfs/.sec/
conf/<variable name>
```

For example, you can configure the variable `hdfs_cache_entry`

```
../.sec/conf/hdfs_cache_entry
```

NOTE: Because HDFS rename is a metadata operation inside the HDFS NameNode and does not call into the local file system, the hash cache might contain expired data. The HDFS NameNode is coded to prevent renamed data from crossing a key boundary to prevent data corruption. However, other access checks based on the

HDFS file name may give incorrect information if the name is expired. Understand this risk before using this feature.

Enabling VTE on HDFS

To enable VTE on your HDFS:

1. Restart VTE agent on the node:
   - For Redhat 6.x, type:
     ```
     # /etc/init.d/secfs restart
     ```
   - For Redhat 7.x, type:
     ```
     # /etc/vormetric/secfs restart
     ```
2. Restart the Hadoop Services in the cluster.

You can now create GuardPoints to protect your entire HDFS.

Deleting Metadata in HDFS when Migrating Out of LDT

In an HDFS deployment, if you migrate from an LDT to a non-LDT environment, the administrator must delete the LDT mdstore file.

In the following example:
`/hadoop/hdfs` is the mount point
`/hadoop/hdfs/data` is the guardpoint

To manage the migration:

1. In the DSM console, click **Host > Host Groups.**
2. Click <*host group name*>. The Edit Host Group - <*host group name*> window opens.
3. Click **Guard FS**.
4. Select the appropriate HDFS directory with an LDT GuardPoint, and click **Unguard**.
5. Using the Ambari admin console, shutdown all NameNode/DataNode one by one. Ensure that no HDFS guardpoints are busy.
6. Ensure that no guardpoints are configured on any HDFS node in the cluster, type:
   ```
   # secfsd -status guard
   ```

System Response

```
No guardpoints configured
```

7. On the node running secfs, type:

   # voradmin ldt attr delete <guard path>

   Example

   # voradmin ldt attr delete /hadoop/hdfs/data

   System Response

   ```
   LDT metadata has been removed from all files in guardpoint
   /hadoop/hdfs/data
   ```

8. On the system, verify that the metadata store has been removed from the secfs mount points, type:

   # voradmin ldt rmstore <mount_point>

   Example

   # voradmin ldt rmstore /hadoop/hdfs

   System Response

   ```
   Enter YES if /hadoop/hdfs does not include any guardpoints
   associated with an LDT policy ->YES
   MDS file /hadoop/hdfs/::vorm:mds:: has been removed.
   ```

9. Verify that the metadata store has been removed from the secfs mount points, type:

   # ls -altr <mount_point>

   Example

   # ls -altr /hadoop/hdfs

   You should **not** see `/hadoop/hdfs/::vorm:mds::` listed.

10. Repeat the above steps for each node in the HDFS cluster.

5 Using VTE with Oracle

This chapter describes how to install and configure VTE on Oracle RAC ASM, Linux/It contains the following sections:

Oracle RAC ASM and ASMLib

This section describes how to install and configure VTE on an Oracle RAC ASM and ASMLib.

Using VTE with an Oracle RAC ASM

You can apply VTE when the Oracle DB is active or inactive. If you choose to use it while the Oracle DB is active, it eliminates any downtime. You can apply VTE during low volume traffic time frames. If you choose to use this option, then use the **rebalance** function of ASM. This allows you to:

1. Migrate data off of a disk so that it can be dropped/removed from a **Diskgroup.**
2. Apply VTE protection.
3. Add the disk back into the diskgroup.

ASMLib

ASMLib is an optional support library for the Automatic Storage Management feature of the Oracle Database. If the customer is using ASMLib, then management is performed through an Oracle ASM command line. Using this can be simpler than the setup required for standard ASM. The commands and details of the procedure differ as well.

Important ASM Commands and Concepts

Rebalancing Disks

When you drop/remove a disk from the diskgroup, it is important to apply the proper value for the power setting for rebalance and to use the **`WAIT`** command.

Example ASM Command:

```
SQL> ALTER DISKGROUP <DiskGroupName> DROP DISK <diskName>
REBALANCE POWER 8 WAIT;
```

- The **`rebalance`** command moves the data off of the disk that you are removing from the diskgroup, distributing the data across the remaining DISKS.
- The **`power`** setting is a number from 1 to 11. It determines how much processing power is dedicated to the **rebalance**, versus normal operations. Unless the encrypting occurs during heavy traffic volume, the minimum value you should use is 6. Otherwise, consult the customer's DBA for the proper setting. An appropriate value to start with is 8.

Mapping Raw Devices

You can map raw devices for this configuration using:

- **Multipath I/O**

  This is typically evident when the path for the mapped devices is set to: `/dev/mapper/<device-name>`.

- **Raw devices**

  Some customers use raw devices to map a name like raw3 to a specific device name. You usually find this mapping in a file called: `/etc/sysconfig/rawdevices`.

NOTE: It is important to understand how the device names are used and if they are the same across all of the RAC nodes.

- **EMC PowerPath**

  If using EMC PowerPath then the device names are similar to the following: `/dev/emcpowerXXXX`.

When browsing the DSM through the local host, you cannot find Power Path devices. You must manually input the paths. The guarded disk names are prepended with: `/dev/secvm`.

Checking Rebalance Status

The **Wait** command is very important when ASM performs a rebalance. When you specify **`wait`**, the command prompt does not display until all of the data is rebalanced and migrated off of the disk. If you do not specify wait, the command prompt returns immediately, and you must issue the following ASM command to check the status of the rebalance:

```
SQL> select * from v$asm_operation;
```

This command returns information about the:

- State
- Current power level
- Current amount rebalanced
- Estimated work until completion
- Rate
- Estimated minutes
- Any error codes

NOTE: It is highly recommended that you always specify the **WAIT** command when performing a **Drop Disk** with Rebalance. If it is not specified, ASM may prematurely release the disk, thereby allowing VTE to place a GuardPoint on the disk before the rebalance completes. This action may corrupt the data.

Oracle cautions against this issue:

Caution: The `ALTER DISKGROUP...DROP DISK` statement returns before the drop and rebalance operations complete. Do not reuse, remove, or disconnect the dropped disk until the `HEADER_STATUS` column in the `V$ASM_DISK` view for this disk changes to `FORMER`. You can query the `V$ASM_OPERATION` view to determine the amount of time remaining for the drop/rebalance operation to complete. For

more information, refer to the *Oracle Database SQL Language Reference* and the *Oracle Database Reference*.

Determining Best Method for Encrypting Disks

A diskgroup can contain one or multiple disks. You must determine if the diskgroup contains enough disks and free space for encryption. If the diskgroup contains only one disk, or multiple disks but not enough free space, then you must use the **Offline** (backup/restore) method for encryption.

If the diskgroup contains more than one, you can use the **Online** (rebalancing) method. During rebalancing, additional disks allow for migrating data from the original disk so that it can be encrypted, added back into the diskgroup, and then migrated back to the source disk. Therefore, if the customer does not want to permanently add extra disks, they can add disks temporarily, just for rebalancing.

In general, once you have completed the initial setup for the operating system with which you are working, for both ASM or ASMLib, the high-level process is the same for applying VTE protection to raw devices and using them.

Online Method (No Application / Database Downtime)

Typically, when using the online method, follow these steps:

1. Make an ASM disk available for protection by either removing a disk from an existing diskgroup, or allocating a new disk.
2. Apply VTE encryption to the disk.
3. Add each protected disk to the diskgroup.
4. Restart the nodes and the failover test.
5. Repeat the previous steps for each disk in the diskgroup.

Offline Method (Backup the DB)

Typically, when using the offline method, follow these steps:

1. Backup the database.
2. Make an ASM disk available for protection by either removing a disk from an existing diskgroup, or allocating a new disk.
3. Stop the Oracle database.

4. Delete the diskgroup.
5. Apply VTE encryption to the disk.
6. Recreate the diskgroup.
7. Add the protected disk to the diskgroup.
8. Restart the nodes and the failover test.
9. Repeat the previous steps for each disk in the diskgroup.

General Prerequisites

Follow these guidelines for best results.

Setup

- Ensure that you have a current backup of the database
- Install and register VTE agents on **all** RAC node Hosts
- Create a **Host Group** and add all RAC node hosts as members
- Create an encryption key for the Oracle RAC Database / Application
- Create an Oracle policy using the proper encryption key

NOTE: If the raw device mappings for the disk(s) are **not** identical across all nodes in the RAC, then you cannot use a Host Group for managing the GuardPoint within the DSM. You **must** apply the GuardPoint to each Host individually. This is typically not optimal, as a Host Group is the most effective and consistent way to manage GuardPoints for Oracle RAC environments.

Modify the UDEV Rules

NOTE: For RHEL 5 operating systems.

Update the UDEV rules to ensure that they recognize the VTE protected raw devices:

1. Edit the file `/etc/udev/rules.d/90-dm.rules`.
2. Find and comment out the following line by adding a hash tag (**#**) to the beginning of the line.

```
# KERNEL=="dm-[0-9]*", ACTION=="add", OPTIONS+="ignore_device"
```

3. Restart the UDEV rules by executing the following command:

   /sbin/start_udev

NOTE: The VTE guide states that a restart is usually required once the UDEV rules have changed. However, that would defeat the purpose of an online method for conversion to VTE. Using the above `start_udev` command removes the need to restart the Host. If this command does not work, then a system restart may be required.

Altering ASM_DISKSTRING on ASM

ASM uses the `asm_diskstring` setting to identify the path where ASM will attempt to locate available disks to use. If you are using device names when adding the disk, you must modify the string to include the path to SecVM.

1. To retrieve the `ASM_DISKSTRING` setting, type:

   ```
   SQL> SHOW PARAMETER ASM_DISKSTRING
   ```

2. To modify the setting, type:

   ```
   SQL> ALTER SYSTEM SET ASM_DISKSTRING='/dev/mapper/*',
   '/dev/secvm/dev/mapper/*';
   ```

   Where the path added is the path to SecVM.

ASMlib manages the binding, not ASM. ASMlib creates ASMlib devices on the SecVM devices and presents it to ASM. ASM automatically recognizes the new device. This creates the need to alter diskstrings for ASM. In addition, Oracle ASM sees a new device created using ASMlib and Raw, by default.

Example

Use Oracle ASMlib to bind the device:

```
# oracleasm createdisk <devicename> /dev/secvm/dev/<blockdev>
```

Using raw command to bind the device:

```
# raw /dev/raw/rawN /dev/secvm/dev/<blockdev>
```

Specific Prerequisites

Establishing a Starting Point

In many production environments, you may find that it has been a very long time since the RAC nodes have had the services restarted or have been completely rebooted. This can result in a lack of understanding of the actual state of the RAC cluster and its ability to survive a reboot on its own, prior to installing VTE.

Restarts can uncover issues in the RAC environment that are unrelated to VTE. To avoid issues after a VTE installation, Thales recommends that you restart each RAC node **AFTER** VTE is installed and **PRIOR** to establishing any GuardPoints. This may not be feasible in a single node configuration. However, by doing so, VTE is installed but inactive, and you can ensure that the platform is in a workable state prior to getting started.

The Importance of Device Mapping

It is important to use device naming and mapping in a multi-node RAC configuration. Verify the device names to ensure that the disks are mapped to the same disks on each RAC node before applying any GuardPoints. Thales **recommends** that RAC nodes use the same device names across all nodes. If they do not match, then problems can occur.

If the RAC nodes use the same device names, use a Host Group to create GuardPoints. If they do not match, do not use a Host Group to create GuardPoints. Set them up independently on each Host.

Important Note about Raw Devices on AIX & UNIX

In general, raw devices are created as either character or block mode devices. Any I/O performed on character devices is non-buffered, while I/O on block devices is buffered and performed in defined block sizes (that is, 4K bytes).

While the Oracle documentation for using ASM with raw devices indicates that you can use either character or block devices**, VTE REQUIRES a block device for guarding.**

NOTE: Attempting to apply a GuardPoint on a character device that ***does not*** have a corresponding block device may result in a GuardPoint that never encrypts data. The status of the GuardPoint never shows as guarded.

NOTE: WebUI does not support browsing for the character devices. You would need to manually paste the name into the WebUI.

Oracle RAC ASMLib Multi-Disk Online Method

The online method describes how to remove, protect and add disks to a diskgroup.

Assumptions

Using the Online Method assumes that there is enough free space in the diskgroup so that you can drop/remove, protect and then add the disk back into the diskgroup.

NOTE: During the initial investigation, you may want to ensure that you have the correct raw device name for each disk that you plan on protecting. Before making any changes to the ASM configuration, obtain the definitive device names for **each** disk by running the following from the command prompt:
`# oracleasm querydisk –p <diskName>`

To add the disk to the diskgroup using the online method and make it ready for use:

1. Open a terminal session on both RAC Nodes.
2. On **RAC Node 1**, perform the following:

   On the ASM, type the following to remove a disk for the diskgroup. **`SQL> ALTER DISKGROUP <diskGroupName> DROP DISK <diskName> REBALANCE POWER 11 WAIT;`**

   c. Delete the disk from ASM, type:

   `# oracleasm deletedisk <diskName>`

   d. Verify that the disk is deleted from the ASM and therefore, it is not listed, type:

   `# oracleasm listdisks`

NOTE: If you are planning to apply a GuardPoint to a raw device that is currently in an ASM diskgroup, you must remove and delete the disk from the diskgroup before you apply the GuardPoint. ASMLib will not see the guarded disk if you skip this step. When deleting the disk, make sure that the deletion completes before continuing.

About Oracle RAC ASM Raw Devices

When Not Using ASMLib

Before starting the VTE implementation, investigate how the customer is using raw devices for their ASM configuration.

Devices using Raw Bindings

Typically, a device that uses a raw binding looks like the following to ASM:

```
/dev/raw/raw1
```

If the device is mapped this way, you must locate where the mapping is performed. Typically, you can find this in the following configuration file:

```
/etc/sysconfig/rawdevices
```

The underlying binding could be to either a **standard device** name or a **multipath** I/O device name. Either way, you must find where the bind commands are run so that you can modify them for SecVM.

NOTE: If raw bindings are in use, then typically no changes are needed for the `asm_diskstring`. Because the binding to the actual device is created through the `bind` command, locate where the binding occurs and change the binding to SecVM.

Multipath I/O Devices

Devices using multipath I/O are typically found with the name:

```
/dev/mapper/mpath1
```

Generally, when using multipath I/O, you create SecVM on the multipath device name.

NOTE: If you use multipath I/O devices in the ASM configuration to add its disk, you must modify the `asm_diskstring` parameter to include the `/dev/secvm/dev/*` path.

Standard Devices

In many cases the ASM configuration may be using plain device names, like the following:

```
/dev/sda1
```

NOTE: If you use standard device names in the ASM configuration to add a disk, you must modify the `ASM_DISKSTRING` parameter to include the `/dev/secvm/dev/*` path.

Consistent Naming of Devices across RAC Nodes

As previously stated, if the raw device mappings for the disk(s) are **NOT** identical across all nodes in the RAC, then you **CANNOT** use a Host Group and you **MUST** apply the GuardPoints to each Host individually. This is typically NOT optimal, as a Host Group is the most effective way to manage an Oracle RAC environment.

Oracle RAC ASM Multi-Disk Online Method

Performing encryption with the online rebalancing method requires sufficient free space to allow the drop of the largest ASM disk.

Checking for Space

In the Oracle system, use the following commands to check for available disk space:

1. Check total free space in the disk group:

```
SQL> SELECT name, free_mb, total_mb, free_mb/total_mb*100 as
percentage    FROM v$asm_diskgroup;
```

System Response:

```
NAME                            FREE_MB   TOTAL_MB PERCENTAGE
------------------------------ ---------- ---------- ----------
DATA                                 7       2109 .331910858
```

2. Check individual ASM disk size and usage:

```
SQL> select a.name DiskGroup, b.disk_number Disk#, b.name
DiskName, b.total_mb, b.free_mb, b.path, b.header_status FROM
v$asm_disk b, v$asm_diskgroup a where a.group_number (+)
=b.group_number order by b.group_number, b.disk_number, b.name
```

System Response:

```
DISKGROUP  DISK# DISKNAME   TOTAL_MB    FREE_MB PATH  HEADER_STATU
---------- ---------- ------------- ---------- ---------- ---------
DATA    0 DATA_0000           1874       1273
/dev/oracleasm/disks/DATA3     MEMBER

DATA    1 DATA_0001           1992        608
/dev/oracleasm/disks/DATA4     MEMBER

DATA    3 DATA_0003            117          0
/dev/oracleasm/disks/DATA2     MEMBER

        0 DATA_ENC_0000        109         28
/dev/oracleasm/disks/DATA1_ENC MEMBER
```

Adding a Disk to the Diskgroup

Using the Online Method assumes that there is enough free space in the diskgroup so that you can drop/remove a disk, protect it with VTE, and then add it back into the diskgroup.

To add the disk to the diskgroup:

1. Open a terminal session on both RAC Nodes.
2. On **RAC Node 1**, on the ASM, remove the disk from the disk group, type.

   SQL> **ALTER DISKGROUP <*diskGroupName*> DROP DISK <*diskName*> REBALANCE POWER 11 WAIT;**

3. On the DSM, in the Host Group, apply a GuardPoint to the Raw Device: **<*rawDevice1Name*>**
4. From **RAC Node 1**, display the status of the guarded disks, type:

```
# secfsd -status guard
```

5. On both **RAC Node 1 and 2** type:

   ```
   # chown oracle:oinstall /dev/secvm/<rawDevice1Name>
   # chmod 660 /dev/secvm/<rawDevice1Name>
   ```

6. From **RAC Node, on the ASM,** add the protected disk to the disk group:

   ```
   SQL> ALTER DISKGROUP <diskGroupName> ADD DISK
   /dev/secvm/<rawDevice1Name> NAME <disk1Name>;
   ```

   The disk is now added to the diskgroup and ready for use.

7. The system is now ready for a reboot and failover test. Go to the section "Surviving the Reboot and Failover Testing" on page 98.

Troubleshooting

Occasionally, settings do not persist when the system is rebooted. To ensure they do persist, edit the `/etc/rc.local` file and add the following lines:

```
Echo "Changing Permission for secvm devices"
chown oracle:oinstall /dev/secvm/dev/<rawDevice1Name>
chmod 660 /dev/secvm/dev/<rawDevice1Name>
```

Oracle RAC ASM Multi-Disk Offline Method (Backup/Restore)

Using the Offline Method assumes that there is not enough free space in the diskgroup.

1. Open a terminal session on both RAC Nodes.

   On RAC Node 1, on the ASM, type the following to remove the disk group. **SQL> DROP DISKGROUP <diskGroupName> FORCE INCLUDING CONTENTS;**

NOTE: Make sure that the disk is removed before guarding the raw devices.

2. On the DSM, in the Host Group, apply GuardPoints to the three raw devices:

   <rawDeviceName1>

   <rawDeviceName2>

<*rawDeviceName3*>

3. On **RAC Node 1**, perform the following:

   a. Display the status of the guarded disks, type:

   ```
   # secfsd -status guard
   ```

4. On both **RAC Node 1** and **2**, type:

   ```
   # chown oracle:oinstall /dev/secvm/<rawDeviceName1>
   # chmod 660 /dev/secvm/<rawDeviceName1>
   # chown oracle:oinstall /dev/secvm/<rawDeviceName2>
   # chmod 660 /dev/secvm/<rawDeviceName2>
   # chown oracle:oinstall /dev/secvm/<rawDeviceName3>
   # chmod 660 /dev/secvm/<rawDeviceName3>
   ```

5. From **RAC Node 1**, **on the ASM**, add the protected disk to the disk group, type:

   ```
   SQL> ALTER DISKGROUP <diskGroupName> ADD DISK
   /dev/secvm/<rawDeviceName1> NAME <diskName1>;
   SQL> ALTER DISKGROUP <diskGroupName> ADD DISK
   /dev/secvm/<rawDeviceName2> NAME <diskName2>;
   SQL> ALTER DISKGROUP <diskGroupName> ADD DISK
   /dev/secvm/<rawDeviceName3> NAME <diskName3>;
   ```

   The disks are now added to the diskgroup and ready for use.

6. On **RAC Node 1**, restore the database.
7. The system is now ready for a reboot and failover test. Go to the section "Surviving the Reboot and Failover Testing" on page 98.

Troubleshooting

Occasionally, settings do not persist when the system is rebooted. To ensure they do persist, edit the `/etc/rc.local` file and add the following lines:

```
Echo "Changing Permission for secvm devices"
# chown oracle:oinstall /dev/secvm/dev/<rawDeviceName1>
# chmod 660 /dev/secvm/dev/<rawDeviceName1>
# chown oracle:oinstall /dev/secvm/dev/<rawDeviceName2>
# chmod 660 /dev/secvm/dev/<rawDeviceName2>
# chown oracle:oinstall /dev/secvm/dev/<rawDeviceName3>
```

```
# chmod 660 /dev/secvm/dev/<rawDeviceName3>
```

Surviving the Reboot and Failover Testing

Preparing for Failover Testing with ASMLib

When using ASMLib with the `createdisk` command, there is no requirement to make additional changes in `rc.local` or other areas for mapping device names or to use `chmod` or `chown` for SecVM. This is because it is managed for you by the `createdisk` function and you can verify this by running the following command:

```
# ls -l /dev/oracleasm/disks
```

VTE Load Order and Startup Scripts

The last change is to ensure that VTE starts before ASM starts in the startup scripts.

Failover Testing

Confirm that everything is functional:

- Ensure that the GuardPoints are all operational.
- Ensure that you receive valid results when you query the database.
- Verify that the load order ensures that VTE starts before ASM .

Once verified, you can start the failover testing for each RAC Node.

1. Reboot the RAC Node 1 and monitor the startup.
2. Once the restart is clean, reboot RAC Node 2 and monitor the startup.

Issues with Device Mapper and Invalid Guard Path

If VTE is unable to apply a GuardPoint on a raw device, the logs may generate an error similar to the following:

```
[SecFS, 0] EVENT: Failed to guard /dev/mapper/devicename (reason:
Invalid Guard Path flags 0x2 gypped 0x4 status 0x11) – Will retry
later
```

If you receive this error, use the `setup` command to check the status of the disks, type:

```
#setup info <deviceName>
```

Before attempting to establish a GuardPoint, look closely at the open count value and ensure that it is 0 on **all** nodes.

Basic Troubleshooting Techniques

Following are some of the most common configuration issues that prevent the Oracle ASM configuration from working properly.

If you encountering errors similar to:

- ORA-15075: disk(s) are not visible cluster-wide
- ORA-15032: not all alterations performed

This could be the result of improper settings for the I/O layer, meaning that your disks are not properly configured, etc.

Perform the following tasks to verify that the settings are correct:

1. On the DSM **WebUI**, in the Host Group that was created for the RAC cluster, verify that the host group for this configuration does **NOT** have the Cluster Group option set (it is only for GPFS).
2. Ensure that the GuardPoints for the block devices are set at the Host Group level. This ensures that each node receives identical GuardPoints.
3. Verify that the GuardPoints are active on all nodes. When the GuardPoints are set, go to each node and verify that they are set and guarded, using the WebUI or the **`secfsd -status guard`** command. If they do not guard correctly:
   - The udev rules are not set correctly, see "Modify the UDEV Rules" on page 89
   - The device names are not the same across all nodes
4. From ASM, make sure that the `asm_diskstring` parameter is modified to include the VTE devices and that the proper pathname is used, see "Altering ASM_DISKSTRING on ASM" on page 90.

Verifying Database Encryption

Option 1

The best way to verify the state of the data, without impacting anything in the existing environment, is to use the Oracle `kfed` command. You can run this command against the native path of the existing GuardPoints and make sure it returns with valid header information. If it returns valid information with the GuardPoint in place, then this confirms that the data is properly encrypted. If it returns with invalid header information, then that indicates that the data is either clear, or not in the expected encrypted state. The syntax for running this command would look similar to the following but will vary based on your environment.

```
# /app/oracle/grid/product/11.2.0/grid/bin/kfed read
/dev/rdisk/<diskName>
```

If the location is properly encrypted, following is an example of the viewable output:

```
# /app/oracle/grid/product/11.2.0/grid/bin/kfed read
/dev/rdisk/<diskName>
```

System Response:

```
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          242 ; 0x001: 0xf2
kfbh.type:                          124 ; 0x002: *** Unknown Enum
***
kfbh.datfmt:                         66 ; 0x003: 0x42
kfbh.block.blk:              1088904227 ; 0x004: blk=1088904227
kfbh.block.obj:              1558192170 ; 0x008: file=8234
kfbh.check:                  3321251423 ; 0x00c: 0xc5f6465f
kfbh.fcn.base:                932956641 ; 0x010: 0x379bc9e1
kfbh.fcn.wrap:               3040493590 ; 0x014: 0xb53a4016
kfbh.spare1:                 3806015223 ; 0x018: 0xe2db2ef7
kfbh.spare2:                 3794962182 ; 0x01c: 0xe2328706
60000000000D8000 01F27C42 40E75C23 5CE0202A C5F6465F
[..|B@.\#\. *..F_]
60000000000D8010 379BC9E1 B53A4016 E2DB2EF7 E2328706
[7....:@......2..]
```

```
600000000000D8020 CA2F30AD 522B4D21 99292639 004EBB34
[./0.R+M!.)&9.N.4]

600000000000D8030 A3896BE8 BD839D23 2204E19E 946C575C
[..k....#"....lW\]

600000000000D8040 4CE2218F 35E1B101 AF751A70 780E6D6E
[L.!.5....u.px.mn]

600000000000D8050 5E7E6A38 C600ED5F 929047C4 DF372A8E
[^~j8..._..G..7*.]

600000000000D8060 E103152D BA87CC03 11A7D963 9D72FCE1

[...-.......c.r..]

600000000000D8070 1EC6B48B 03EE869F 61D651F9 E7614957
[........a.Q..aIW]

600000000000D8080 810E0353 9C461F49 69569733 501D19EF
[...S.F.IiV.3P...]

600000000000D8090 B268002B 4F9457B6 BDB04AC5 D3D07446
[.h.+O.W...J...tF]

600000000000D80A0 FD9EE5E0 9B46CB66 30D10B22 F63AB77E
[.....F.f0.."..:.~]

600000000000D80B0 6FF79075 4BBD1FAD 8F226188 7774300D
[o..uK...."a.wt0.]

600000000000D80C0 A809B6FB E1F1C80B B5760E68 9747D97D
[.........v.h.G.}]

KFED-00322: Invalid content encountered during block traversal:
[kfbtTraverseBlock][Invalid OSM block type][][124]
```

Option 2

The second option to verify the state of the data is to use the `dd` command. This requires you to specify some blocks and write it out to a file. After this completes, read the file using the strings command. You can do this while the device is in use. In the example below some sectors are skipped and it only dumps 10000 counts.

For example:

```
# dd if=/dev/mapper/asm_data2p1 of=/tmp/dd2.out skip=1047
count=10000
```

Option 3

The third option to verify the state of the data without impacting the existing environment is to use the strings command.

NOTE: The strings command cannot read a busy or large device.

You can run this command against the native path (`/dev/<deviceName>`) of the existing GuardPoints (`/dev/secvm/dev/<deviceName>`). The new path to SecVM would be similar to `/dev/secvm/dev/<deviceName>`. By executing the strings command against the native path **`strings /dev/devicename | more`**, this does not go through the SecVM device and therefore is not be decrypted. If it is encrypted the output should contain illegible text.

6 Configuring Support for SAP HANA

This chapter describes SAP HANA, which provides automatic host-failover support. It contains the following sections:

Overview

SAP HANA provides automatic host-failover support. VTE works with HANA fibre storage systems to enable and disable GuardPoints when a protected host starts, stops, or fails over to standby host.

SAP HANA supports non-shared storage where each HANA node has its own separate storage volumes. VTE provides customized scripts to support startups, shutdowns, and fail overs.

HANA attaches logical unit number (LUNs) or logical volume management (LVMs) using a Fibre Storage Connector (fcClient) providers. Vormetric provides hooks that are called by the HANA fcClient providers that manage guarding or unguarding of storage locations.

NOTE: Vormetric recommends using host groups to manage configuring in a clustered host environment.

Customizing VTE for SAP HANA

1. Go to the installation directory:

   ```
   # cd /opt/vormetric/DataSecurityExpert/agent/secfs/saphana
   ```

2. If required, edit the appropriate VTE `fcClient` refined script:

Script file	Use
`fcClientRefinedVTE.py`	fcClient provider for LUN
`fcClientLVMRefinedVTE.py`	fcClientLVM provider for LVMs

3. Copy the appropriate script file to a shared location that is accessible to all nodes. In a HANA cluster environment, all nodes require access to the VTE scripts.
4. Edit the storage section of the `global.ini` file to indicate the corresponding VTE `fcClient` as the High Availability (HA) provider and to point to the location of the VTE script.

LUN Example:

```
[storage]
   ha_provider = fcClientRefinedVTE
   ha_provider_path = /hana/shared/myFcClient
```

If necessary, enable debug tracing:

```
   [trace]
   ha_provider = debug
   ha_fcclient = debug
   ha_fcclientrefinedvte = debug
```

LVM Example:

```
[storage]
ha_provider = fcClientLVMRefinedVTE
ha_provider_path = /hana/shared/myFcClient

If necessary, enable debug tracing

[trace]
ha_provider = debug
ha_fcclientlvm = debug
ha_fcclientlvmrefinedvte = debug
```

5. Use the same VTE agent with all hosts, including standby hosts.
6. Ensure that `/etc/sudoers` includes the following:

   ```
   <sid>adm ALL=NOPASSWD: /usr/bin/secfsd
   ```

7. Enable the guard paths at the mount-point level.
   For example, individual guards were placed on `/hana/data/HAN/mnt00001`, `/hana/data/HAN/mnt00002`, and so forth.

   a. Use a similar naming practice for log partitions, such as `/hana/log/HAN/mnt00001`, and so forth.

   b. Place the guard at the mount-point level. The guarded paths must match the corresponding data and log mount paths.

8. From the DSM, configure GuardPoints as type `manual`. You must enable and disable the guardpoints immediately after the device is attached, or just prior to detachment.

   The reason for manual GuardPoints is that it invokes guarding and unguarding from within the HANA during the startup, shutdown, or failover process. The process resembles that of mounting and unmounting guarded auto-mount points.

9. Configure all GuardPoints so that they are available in the standby host, so that any data and log partitions that fail over from any host can be guarded on the standby.

   Vormetric recommends that you configure all GuardPoints on all hosts, because a failed-over active host can then become the new standby, and will require all available GuardPoints.

Example:

Following is an example of the data and log volumes for the host that are mounted.

```
/dev/mapper/VG_HAN_DATA_1-LV_HAN_DATA_1    793971096  3059712
750579916   1% /hana/data/HAN/mnt00001

   /hana/data/HAN/mnt00001                    793971096  3059712
750579916   1% /hana/data/HAN/mnt00001

   /dev/mapper/VG_HAN_LOG_1-LV_HAN_LOG_1      496233160  2461764
468564152   1% /hana/log/HAN/mnt00001

   /hana/log/HAN/mnt00001                     496233160  2461764
468564152   1% /hana/log/HAN/mnt00001
```

Note that partition `mnt00002` is also configured, although not currently mounted by HANA. The `secfsd` status output should show the GuardPoint configuration as follows:

```
  GuardPoint              Policy    Type    ConfigState  Status
Reason
  ----------              ------    ----    -----------  ------
------
  /hana/data/HAN/mnt00001  my-pol   manual  guarded    guarded
N/A
  /hana/log/HAN/mnt00001   my-pol   manual  guarded    guarded
N/A
   /hana/data/HAN/mnt00002  my-pol    manual  unguarded    not
guarded  Inactive
   /hana/log/HAN/mnt00002   my-pol    manual  unguarded   not
guarded   Inactive
```

For more information, see the *SAP HANA Fiber Channel Storage Connector Admin Guide*
http://www.sap.com/documents/2016/06/84ea994f-767c-0010-82c7-eda71af511fa.html

Using SAP HANA with LDT

SAP HANA is compatible with LDT. You must add additional VTE commands to the HANA administrator entry.

- Using a text editor, edit `/etc/sudoers` and add entries for `/usr/bin/voradmin` and `/usr/bin/vmsec`:

  Example:

  ```
  # hanadm ALL=NOPASSWD:
  /usr/bin/secfsd,/usr/bin/voradmin,/usr/bin/vmsec,/sbin/multipa
  th,/sbin/multipathd,/etc/init.d/multipathd,/usr/bin/sg_persist
  ,/bin/mount,/bin/umount,/bin/kill,/usr/bin/lsof,/sbin/vgchange
  ,/sbin/vgscan
  ```

If you are using an ext3 file system, you must mount it with extended attributes.

- Using a text editor, edit the storage section of the `global.ini` file, type:

  ```
  partition_*_data__mountOptions = -o user_xattr

  partition_*_log__mountOptions = -o user_xattr
  ```

Setting Memory Allocation

There is a limitation in memory allocations for SAP HANA with asynchronous direct I/O. When you use VTE in conjunction with applications like SAP HANA that can process large numbers of direct I/O writes through the Linux AIO interface, VTE can allocate more memory than is desirable.

To limit the amount of memory that VTE consumes for AIO buffers, use the following configuration to limit the amount of memory VTE consumed for AIO buffers:

```
# max_aio_memory_limit <MB>
```

The MB value specifies how much memory to allocate to temporary DIO buffers.

NOTE: If you do not specify a value, the default is 0, which has no memory bounding effect.

Set the option by echoing a value into the `opt/vormetric/DataSecurityExpert/agent/secfs/.sec/conf/` configuration file. For example:

```
echo 1024 > /opt/vormetric/DataSecurityExpert/agent/
secfs/.sec/conf/max_aio_memory_limit
```

This limits the memory consumed by AIO buffers to 1GB.

NOTE: You **must restart** VTE after changing any values in the configuration directory to make the changes effective.

7

Using VTE with Microsoft SQL

This chapter discusses using VTE with Microsoft SQL AlwaysOn and SQL File Tables. It contains the following sections:

Using VTE with SQL

You must stop the SQL service before guarding the SQL DB. When this occurs, the SQL Server replication may become unsynchronized. When restarting, it may take a brief period of time for the SQL Server replication to resynchronize with the other node. The SQL Server issues a warning against any attempted failovers during that brief period.

NOTE: Minimizing the duration for which the SQL Server service is stopped is beneficial for reducing the resynchronization period.

Using VTE with SQL FileTables

SQL FileTables allows you to store files and documents in special tables in the SQL Server called FileTables, but access them from Windows applications as if they were stored in the file system, without making any changes to your client applications. For some of the use cases, you can use FileTables with VTE.

Considerations

Be aware of the following considerations:

Installation

- DO NOT install VTE agent on the SQL Server
- Install VTE agent on the remote system where SQL server does not reside

SQL FileTables

- If multiple servers access the SQL FileTable:
  - Install VTE agent on all of the servers
  - Protect all of the FileTable folders with the same VTE policy

Caution: Accessing the FileTable without VTE may corrupt the data.

- When you create a new FileTable, alter, or drop FileTables, this may require applying a new GuardPoint
- Every FileTable has a separate FileTable Folder so you must apply separate GuardPoints for each FileTable
- You must apply a unique GuardPoint to each VNN path.

  For example, if you configure two FileTables on an SQL Server, then the remote SQL administrator system must apply one GuardPoint to each configured VNN name.
- Guarding on a VNN name is similar to guarding a network path with VTE.
- If you want to access the FileTables from multiple remote systems, you must install VTE agent on those systems and apply the GuardPoints.

Caution: LDT is **not** supported with SQL FileTables. Only use offline Data Transformation to transform the initial SQL data.

Advantages

- System administrator cannot see the data locally on the SQL server because no VTE Agent is installed on the SQL server.
- The data transferring between servers is also encrypted.

Supported Use Cases

VTE supports the following FileTables use cases:

VTE Data Transformation of existing files in FileTables

Configuration guidelines:

1. Install VTE agent on the remote server.
2. Create a new FileTable, or Identify the FileTable folder for the existing FileTable.
3. Create an offline Data Transformation policy and apply to the GuardPoint on the FileTable folder.
4. Run the Dataxform utility to transform the data.

Protect files in SQL FileTables with VTE

Configuration guidelines:

1. Install VTE agent on the remote server.
2. Create a new FileTable, or Identify the FileTable folder for the existing FileTable.
3. Create a production policy and apply the GuardPoint on the FileTable folder.
4. Once the GuardPoint is active, you can use the file table to load and access files.

Protect files with SQL AlwaysOn Availability Groups with VTE

When the database that contains the FILESTREAM, or FileTable data, belongs to an AlwaysOn availability group, the FILESTREAM and FileTable functions accept or return virtual network names (VNNs) instead of computer names.

Configuration guidelines:

1. Install VTE agent on the remote server.
2. Create a new FileTable, or Identify the virtual network names (VNNs) for the existing FileTable.

3. Create a production policy and apply the GuardPoint to the VNN name
4. Once the GuardPoint is active, you can use the FileTable to load and access files.
5. When you enable FILESTREAM on an instance of SQL Server, it creates an instance-level share to provide access to the FILESTREAM data. Access this share by using the computer name in the following format:

   `\\<computer_name>\<filestream_share_name>`
6. In an AlwaysOn availability group, the computer name is virtualized by using a Virtual Network Name, (VNN). When the computer is the primary replica in an availability group, and databases in the availability group contain FILESTREAM data, then SQL creates a VNN-scoped share to provide access to the FILESTREAM data. Applications that use the file system APIs have to use the VNN-scoped share, which has a path in the following format:

   `\\<VNN>\<filestream_share_name>`

Install VTE on remote systems and guard the SQL Server VNN names

In this use case, VTE is installed on the SQL administrator system (a separate system from where the SQL Server resides) and a GuardPoint is applied to the VNN name.

Unsupported Use Cases

VTE does not support the following use cases:

1. Install VTE agent on the SQL Server and locally apply the GuardPoint on the SQL Server storage.
2. Access FileTables with Transact-SQL.
3. Access FileTables with File I/O APIs on the SQL server. Perform all file I/O on the remote system running the VTE agent.

Installing VTE on Microsoft SQL AlwaysOn

This section describes how to implement VTE with Microsoft SQL AlwaysOn in a variety of configurations for primary and secondary replica servers.

Configurations & Information

This section will document the four configurations that you may deploy for Microsoft SQL AlwaysOn.

Assumption

This guide assumes that you possess a basic understanding of Microsoft SQL database.

Additional Information

You may want to keep the primary server decrypted to serve all users, and use the secondary database for running reports or backups.

- If the database is encrypted, then the Volume Shadow copy-related backups will snapshot and backup encrypted protected data.
- Administrators with the `apply_key` permission can run a query and pull down reports from the secondary database server without affecting the performance of the primary database server.
- The secondary server could be in a remote Data Recovery location. You may want to secure it with encryption.
- LDT is supported with SQL AlwaysOn. See "Using LDT with SQL AlwaysOn" on page 117 for more information.

Methods for Initial Encryption

There are multiple methods for performing the initial encryption of the databases. Decide on which of the following methods best fits your environment. For more information on transforming data, see the "*VTE Data Transformation Guide*."

- Data Transformation – Encrypt data in place
- Backup and Restore to a GuardPoint
- Copy and paste the data into a GuardPoint

Configuration 1

- Databases on primary server and secondary replica servers require encryption

- Database name and location of secondary replica server are the same as the primary server

To perform the procedure:

1. Perform a full backup of the primary database.
2. Change the primary database to offline mode.
3. Confirm the creation of a data transformation and/or operational policy.
4. Guard the folder containing the primary database files with that policy:
   a. If using 'Encrypt data in place' as the selected method of encryption, execute the data transformation and then apply the operational policy.
   b. If using the 'Copy/Restore ' method of encryption, apply the operational policy on an empty folder/device.
5. On the secondary server, create a new folder to store the replicated database.

NOTE: The folder name and the path must be the same as the primary server.

6. Guard the folder with the operational policy.
7. Perform step 4 above for additional secondary server(s).
8. Put the primary database back into online mode.
9. Setup SQL AlwaysOn High Availability group to perform FULL Data Synchronization.

   This copies the primary database and replicates it to secondary replica servers.
10. Verify that the databases in the secondary server are in “Synchronized” mode.

Configuration 2

- Database on the primary server does not require encryption, but the secondary replica database requires it
- Database names and locations for the secondary replica servers are the same as the primary server

To perform the procedure:

1. Perform a full backup of the primary database.
2. Confirm the creation of a data transformation and/or operational policy.
3. On the secondary server, create a new folder to store the replicated database.

NOTE: The folder name and the path must be the same as the primary server.

4. Guard the folder with the operational policy.
5. Perform step 3 & 4 above for additional secondary server(s).
6. Setup SQL AlwaysOn High Availability group to perform **FULL Data Synchronization**.

   This copies the primary database and replicates it to secondary replica servers.
7. Verify that the databases in the secondary server are in “Synchronized” mode.

Configuration 3

- Databases on the primary and secondary replica servers require encryption
- Database name is the same, but the location of the secondary replica server is in a different location from that of the primary server

To perform the procedure:

1. Perform a full backup of the primary database.
2. Change the primary database to offline mode.
3. Confirm the creation of a data transformation and/or operational policy.Guard the folder containing the primary database files with that policy:

   a. If using 'Encrypt data in place' as the selected method of encryption, execute the data transformation and then apply the operational policy.

   b. If using the 'Copy/Restore ' method of encryption, apply the operational policy on an empty folder/device.
4. On the secondary server, create a new folder to store the replicated database.

NOTE: The folder name and the path must be the same as the primary server.

5. Guard the folder with the encryption policy.
6. From secondary server, perform the restore to the primary database.

   a. Select the options **Restore with norecovery** and **Relocate all files to folder**.

   b. Specify the path of the new folder from step 5.
7. Repeat steps 4 & 5 above for any additional secondary server(s).

8. Setup SQL AlwaysOn High Availability group to perform **JOIN ONLY Data Synchronization.**

   This joins the secondary database to the SQL Always High Availability Group. It also establishes replication of new data and logs from the primary to the secondary replicated server.
9. Verify that the databases in the secondary server are in **Synchronized** mode.

Configuration 4

- Database on the primary server does not require encryption, but the secondary replica database requires encryption
- Database name is the same, but the location on the secondary replica server is in a different location than that of the primary server

To perform the procedure:

1. Perform a full backup of the primary database.
2. Confirm the creation of a data transformation and/or operational policy.On the secondary server, create new folder to store the replicated database.
3. Guard the folder with the operational policy.
4. From secondary server, perform restore the primary database:

   a. Select the options **Restore with norecovery** and **Relocate all files to folder**.

   b. Specify the path of the new folder from step.
5. Setup SQL AlwaysOn High Availability group to perform **JOIN ONLY Data Synchronization.**

   Joins the secondary database to the SQL Always HA Group. It also establishes replication of new data and logs from the primary to the secondary replicated server.
6. Verify that the databases in the secondary server are in **Synchronized** mode.

Data Transformation (Encryption in place)

For more information on transforming and encrypting data-in-place, see the "*VTE Data Transformation Guide*."

Copy/Restore

For more information on transforming data using the copy and replace method, see the "*VTE Data Transformation Guide*."

SQL Server Policy Tuning

In this section, you created and defined a process set for SQL Server that grants certain executables –in this case `sqlservr.exe`– unrestricted access to the database files. The need may arise to allow other executables, and/or users, access to the files.

You can grant this access by:

- Adding to the existing process set
- Creating a new one

The best option depends on the access requirements. The key decision is whether or not to select the **Apply Key** effect along with **Permit** or not. Omitting **Apply Key** on a security rule that still contains **Permit** allows the specified user or process to access to the data, but does not apply the encryption key, so therefore only shows them the data in its encrypted, cypher-text format. This is useful for anti-virus or backup software that may need to scan or copy the file, but does not necessarily need to see the contents.

Using LDT with SQL AlwaysOn

To guard a directory with an LDT (Live Data Transformation) policy, you must temporarily close all files in that directory. In a SQL Server AlwaysOn environment, this may entail temporarily stopping the SQL Server service on the node that is being guarded. Once the directory is guarded, then you can start the SQL Server service immediately.

It is important to remember the SQL Server AlwaysOn replication standard operating procedures.

- If one SQL Server service is taken offline for any reason, then once it is brought back on line, it takes the SQL Server a moment to re-synchronize the database nodes.
- The longer that secondary service was down, and the more inserts/updates and deletes that occurred on the still active node during that downtime, then the longer the synchronization period takes.
- During that synchronization period, any attempted fail over results in the SQL Server warning that data loss may occur if the fail over continues. However, once the SQL Server has completed re-synchronizing that secondary node, then any fail over is safe and does not result in loss of data.

8 Concise Logging

This section describes Concise Logging and selective filtering.

This chapter contains the following sections:

Overview of Concise Logging

Thales's standard operational logging sends audit messages for each file system operation. An audit message is sent each time a file is opened, read, updated, or written. Thales's standard logging can generate high volumes of log data. Most of these messages might not be useful or required by security administrators to monitor file system activity on the system.

Concise Logging allows you to focus on relevant audit messages and important actionable messages, such as errors and warnings. It can eliminate the repetitive and less important audit messages generated by read and write activity on a file, reading and writing directory attributes, and other file system activity.

Concise Logging eliminates the following types of messages:

- Audit messages for each and every block read by the user or application. It sends only one audit message for each read/write activity.
- Audit messages that read the attributes, read the basic information of file-set attributes, and other event-based messages.
- Audit messages for directory open, read directory attributes, and directory close.

Using Concise Logging

You can enable and disable the Concise Logging option from the DSM. You can configure Concise Logging for the following:

- All registered hosts in all domains; see "Do not use Learn mode with Concise Logging." on page 120
- A host that has registered with the DSM; see "Configuring Concise Logging for a registered host" on page 121

Considerations

- Concise Logging changes the set of log messages that are sent to Security Information and Event Management (SIEM) software systems. If this results in loss of data required for customer reports, then disable Concise Logging.
- Concise Logging is only supported by VTE `secfs`.
- Enable and disable Concise Logging on the host. VTE applies it to all GuardPoints and for all users on the host for which it is selected. There is no finer-grained control, such as per-GuardPoint, user, or message type.
- When you enable this setting at the DSM level, it applies to all hosts in all domains, that are added to the DSM, but does not apply to any existing hosts. Hosts added after this setting is enabled inherit this setting. The default global setting is off.
- Do not use Learn mode with Concise Logging.

Configuring global Concise Logging

You can enable or disable Concise Logging at any time. The DSM controls the function. Any change in the Concise Logging is reflected on any newly registered hosts and their domains.

To configure global Concise Logging:

1. Login to the DSM with System Admin privileges.
2. Click **System > Log Preferences**. Your system may contain multiple log tabs.
3. Click on a **Log** tab.
4. In the Duplicate Message Suppression Settings field, click **Enable Concise Logging**.
5. Click **Apply**.
6. Repeat steps for any other logs, as appropriate.

The host sends the following message after the administrator has enabled Concise Logging for an individual host:

```
DAO00821: Administrator "voradmin" updated Security Server configuration
"Concise Logging Enabled" from "true" to "false".
```

Configuring Concise Logging for a registered host

You can enable Concise Logging for a host after you have registered it with the DSM. Hosts that are added to the DSM after enabling Concise Logging inherit the global settings from the DSM. This setting can be customized at any time.

To enable Concise Logging on the DSM for a registered host:

1. Log into your host with DSM security admin privileges.
2. Select the host that you would like to customize.
7. Select a **Log** tab.
8. In the Duplicate Message Suppression Settings, click **Enable Concise Logging**.
9. Click **Apply**.

After you enable or disable Concise Logging, VTE generates a log message to record that event:

```
"[CGA] [INFO] [CGA3201I] [11/11/2016 10:57:18] Concise logging
enable
"[CGA] [INFO] [CGA3202I] [11/11/2016 10:57:27] Concise logging
disabled
```

9

Container Security

This chapter describes securing data in container environments, Docker or RedHat OpenShift, using VTE. It contains the following sections:

NOTE: Docker and OpenShift are for Linux only.

Installing Docker Automatically

A new option allows for automatically registering the host for Docker during registration. This allows for easily enabling Docker for a large number of hosts.

During the installation phase, the agent indicates to the DSM that it wants DOCKER. Only registration is performed at this time. After registration, the DSM will validate the Docker host to ensure that the agent and the DSM use compatible versions, and that LDT is not enabled for the host.

NOTE: LDT and Docker are not compatible.

During the installation phase, the agent indicates to the DSM that it wants DOCKER. Only registration is performed at this time. After registrat

Container Security Overview

VTE provides data security for container environments. You can set up data protection policies for container images. In addition to data encryption, VTE also provides container-level access control and audit logging. DSM Administrators can create GuardPoints in container images through the DSM Management Console. Users can use either of the following container options:

- **Docker Container**: CLI tool for creating containers.
- **Red Hat OpenShift with Customized Docker Container**: GUI tool for creating containers and persistent storage that mounts like NFS.

Container Terminology

- **Containers**

  The basic unit of OpenShift/Docker Application are called containers. A container runs one or more processes inside of a portable Linux environment. Containers are started from an Image and are usually isolated from other containers on the same machine.

- **Image**

  A layered Linux file system that contains application code, dependencies, and any supporting operating system libraries. An image is identified by a name that can be local to the current cluster or point to a remote Docker registry.

- **Pods (OpenShift only)**

  A POD is a set of one or more containers that reside on a host/node and share a unique IP address and volume, (persistent storage). OpenShift leverages the Kubernetes concept of a pod.

- **Project and Users (OpenShift only)**

  A namespace that provides a mechanism to scope resources in a cluster. Users interact with OpenShift. It grants permission to access applications.

Docker Containers with VTE

Docker allows for containerizing an environment for application deployment. Docker is an infrastructure built on top of Linux containers and various namespace

components. Typically, an application deployment is bundled with all of the dependent packages in a single image called a Docker image. Docker images are ready to run applications in containers. A user can instantiate any number of Docker containers using from one or more Docker images. This makes it easy for users to move applications around on their servers. The containerization removes the pain of setting up an environment for applications and provides isolation for applications.

A Docker container can run any application, for example, a Postgres database server. Docker containers are used widely to run micro-services, which are stateless in nature. But a micro-service, when it is running, might generate a log trail inside the Docker container. This log trail might contain sensitive information. This warrants encrypting directories inside a Docker container. If a Docker container runs a database-server type of service, it will need data protection including encryption of data, granular access control and the ability to audit the log that accesses the data.

VTE: Virtual Machine versus Docker

VTE installs on the Docker container host and orchestrates security from the Docker host. It does not leave a footprint inside the Docker image or containers. You setup the GuardPoints inside individual containers at runtime. The diagram below depicts the difference between VTE deployment on a virtual machine and a Docker container host.

Figure 11: VTE in Docker container environment versus virtual machine environment

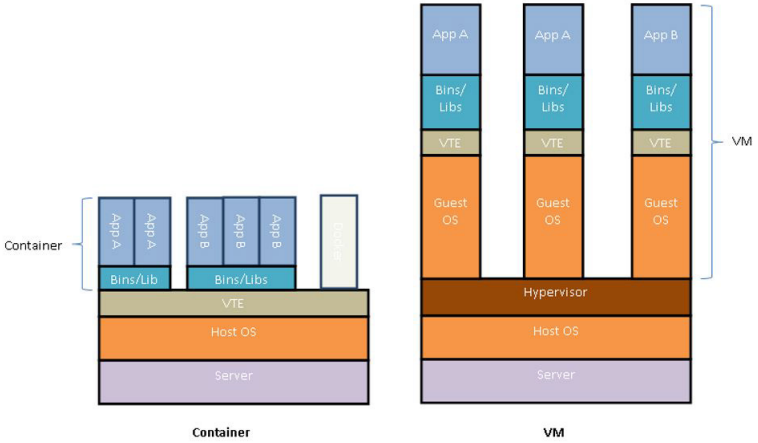

Using the VTE Agent

In order to use the VTE Agent to protect Docker images and containers, you must obtain a VTE Agent license for Docker. Contact Vormetric Support for information on obtaining a license.

There is no change to the VTE installation, however VTE agent must be installed on the Docker host system.

Set the Docker Storage Driver

On RHEL 7.x operating systems, the Docker engine default storage driver has changed from `devicemapper` to `overlayfs` for v1.13.1 and above. Currently, VTE only supports `devicemapper`. If your Docker engine uses `overlayfs` or any other storage driver, then you must change it to `devicemapper` before using VTE to protect your data.

Before you change your storage driver, verify your Docker version:

1. Login to your Docker agent CLI.
2. Check your Docker version, type:

   ```
   # docker version
   ```

   System Response:

```
Client:
 Version:      17.04.0-ce
 API version:  1.28
 Go version:   go1.7.5
 Git commit:   4845c56
 Built:        Mon Apr  3 18:01:50 2017
 OS/Arch:      linux/amd64

Server:
 Version:      17.04.0-ce
 API version:  1.28 (minimum version 1.12)
 Go version:   go1.7.5
 Git commit:   4845c56
 Built:        Mon Apr  3 18:01:50 2017
 OS/Arch:      linux/amd64
 Experimental: false
```

3. Check to what option the storage driver is set, type:

   `# docker Info`

   System Response:

```
Containers: 2
 Running: 1
 Paused: 0
 Stopped: 1
Images: 2
Server Version: 17.04.0-ce
Storage Driver: overlay
```

Change the Storage Driver

To change the storage driver:

1. Login to your Docker agent CLI.
2. Stop the Docker daemon, type:

   `# systemctl stop docker.service`

3. Open another CLI session of agent so you can monitor Docker.
4. In `/etc/docker/` directory, create a file called `daemon.json` if the file does not yet exist.
5. Add a storage driver to the file, type:

```
{
  "storage-driver": "devicemapper"
}
```

6. Save the file.
7. Restart the Docker daemon, type:

   # systemctl start docker.service

8. Verify that the storage driver is set to device mapper, type:

   # docker Info

   System Response:

```
Containers: 2
 Running: 1
 Paused: 0
 Stopped: 1
Images: 2
Server Version: 17.04.0-ce
Storage Driver: devicemapper
```

   Verify that the docker containers are active, type:**# docker ps**

   System Response:

```
CONTAINER ID | IMAGE | COMMAND | CREATED | STATUS | PORTS | NAMES
bb953565e866  abc/rh:v1   "bash" 4 weeks ago Up 4 weeks jdoe
```

Administering the Docker Host

To protect data inside Docker images or containers, you need to create GuardPoints in the DSM, inside a Docker image or container, to which a VTE Agent policy is applied.

This section describes the administrative tasks related to Docker hosts; creating policies, creating GuardPoints, configuring audit logs, and generating reports.

Creating Policies

Policy creation for Docker hosts is largely the same as VTE's existing policy creation procedure with a few differences described below.

The basic procedure to create a policy is as follows:

1. Log on to your DSM as an administrator of type Security, Domain and Security, or All.
2. Navigate to **Policies > Manage Policies**.
3. Click **Add** to open the *Add Online Policy* page
4. From the **Policy Type** drop down list, select **Standard**.

5. Type in a name for the policy in the **Name** field.
6. Add a description for the policy (optional).

Refer to “Configuring Policies” in the *DSM Administrators Guide* for details about creating a policy. Once you have created a policy, you must add rules to the policy to encrypt data and control access to files and directories.

Adding Security Rules

This section describes how to create security rules in the context of Docker images and containers.

Create Resource Set

A ***Resource Set*** specifies the hosts, files, and directories to which a user or process will be permitted or denied access.

1. After creating the policy in the previous section, click **Add** in the **Security Rules** panel to open the *Add Security Rule* page.
2. Click **Select** next to the **Resource** field to open the *Select Resource Set* page. The page displays resource sets if any currently exist.
3. Select an existing resource set that meets your requirements, otherwise click **Add** to open the *Add Resource Set* page.
4. Enter a name for the resource you want to create.
5. Click **Add** again to open the *Add Resource* page to add a resource to the resource set you want to create.
6. Click **Select** next to the **Host** field to select a Docker host from which to choose resources. Select the radio button next to the host and click **Select**.
7. Since this a Docker host, another field displays: **Docker Image**. Click **Browse** to open the **Remote Docker Browser** to select a Docker image or container, from which to select a resource.
8. Click **Browse** next to the **Directory** field to open the **Remote File Browser**.
9. Browse the directories and files on the image or resource that you want to add to the resource set. (Select **Directory Only** or **Directory and File** to browse only directories or files and directories.)
10. Select the resources to add and click **Ok**.
11. Click **Ok** to add the resource set. The *Select Resource Set* page opens

12. Select the resource you just created and click **Select Resource Set**.
13. Check the **Exclude** box, to the right of the **Resource** field.

    This excludes the resources in the resource set and includes all other host resources. Uncheck the box to include just the resources in the resource set.

Create User Set

A ***User Set*** specifies users that are permitted or denied access to files and directories in a GuardPoint.

1. After creating the policy in the previous section, click **Add** in the **Security Rules** panel to open the *Add Security Rule* page.
2. Click **Select** next to the **User** field to open the *Select User Set* page. The page displays user sets if any currently exist.
3. Select an existing user set if it meets your requirements, otherwise click **Add** to open the *Add User Set* page.
4. Enter a name for the user set and click **Add** to open the *Add User* page.
5. Enter information for the **uname**, **uid**, **gname**, **gid**, and **osDomains** fields. Refer to the online help for more details.
6. 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.
7. 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 displays: **Docker Image/Container**.
8. Click **Browse** to open the **Remote Docker Browser** to select a Docker image or container from which to select users.
9. From the **Remote Docker Browser**, expand the file icon to view the Docker image and containers from which to select users to add to the User Set.
10. Once you've made your selections, click **Ok**, a tabulated list of available users is displayed.
11. Select the appropriate users. Click **Ok** to return to the *Add User Set* page.
12. Select users. Click **Ok** to return to the *Select User Set* page.
13. Select the newly created user set and click **Ok**.

Create Process Set

A ***Process Set*** specifies the executables that are permitted or denied access to GuardPoint data.

1. After creating a policy in the previous section, click **Add** in the **Security Rules** panel to open the *Add Security Rule* page.
2. Click **Select** next to the **Process** field to open the *Select Process Set* page. The page displays process sets if any currently exist.
3. Select an existing process set if it meets your requirements, otherwise click **Add** to open the *Add Process Set* page.
4. Enter a name for the process set you are about to create and click **Add** to open the *Add Process* page.
5. Click **Select** next to the **Signature Set** field, a list of existing signature sets is displayed.
6. Select an existing signature set if it meets your requirements, otherwise click **Add** to open the *Add Signature Set* page.
7. If you selected **Add** then you need to provide a **Name** for the signature set.
8. Enter the name and click **Ok**, the *Signature Reference* page opens.
9. Click the name of your signature set to edit the signature set, the *Edit Signature Set* page opens.
10. Click the **Source** tab, click **Select** next to the Host field, the **Select a host to continue** page opens, select the Docker host and click **Select** to go back to the **Source** tab.
11. Click **Browse** next to the **Docker Image/Container** field, select a Docker image or container from the **Remote Docker Browser**.
12. Select a binary from the Docker image and sign it similar to a binary on a VTE host.
13. Click **Back**, the *Signature Reference* page opens.
14. Select the signature set you just created and click **Select Signature Reference**, the *Add Process* page opens.
15. Select a host. Once you select a Docker host, the **Docker Image/Container** field displays.
16. Select a Docker image or container.
17. 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.

18. Choose the appropriate (newly created) process set and click **Ok**. This returns you to *Select Process Set* page.
19. Select the process set and click **Select Process Set**. This returns you to the Add Security Rule page.
20. Click **Ok** to add the process set to the security rule.

Enable Docker through Host Settings

You must enable Docker in the Agent.

1. At the DSM management console, click **Hosts > Hosts**.
2. Click **Add Host**, or click on an **existing** Host name to edit the host.
3. In the General tab, select **Docker Enabled**.
4. Click **Apply**.
5. Click **Host Settings** tab.
6. Add the Docker Daemon to the host settings:

   For Docker v1.12 and above, type:

   ```
   |authenticator|/usr/bin/dockerd
   ```

   For Docker v1.12 and below, type:

   ```
   |authenticator|/usr/bin/docker
   ```
7. Click **Apply**.

VTE Docker GuardPoints

Docker typically provides two types of containers:

- **Transient**: Run micro services which are stateless applications.
- **Long running**: Host stateful applications similar to a database application.

The VTE Docker security feature protects data in both type of containers. You can use VTE to protect either type of containers. You select the container type while configuring the GuardPoint.

VTE provides two types of GuardPoints:

- Image-based
- Container-based

Image-based GuardPoints

You can set up a data protection policy on a Docker image. After an image-based GuardPoint is created, all of the instances running from the protected Docker image inherit the policy and its settings. Any change to the policy is reflected across Docker containers that are started from protected Docker images. You an also refer to Image-based protection as templatized protection. The GuardPoints set up on a Docker image serves as a template for protection of all the Docker containers created as VTE protected Docker images.

Users can browse a Docker image to select the path for protection and configure security rules using information from a Docker image. This process is described in the DSM Guide.

Container-based GuardPoints

You can set up a GuardPoint for a specific Docker container. You can browse a Docker image to select the path for protection, and to configure security rules using information from a Docker image.

GuardPoints for Docker Containers

Before creating GuardPoints on Docker images and containers, the following must be taken into consideration:

- You must add the Docker engine process to the Host Settings.
- When applying GuardPoint policies to Docker containers, users must ensure that the root user has at least 'permit' effect on the GuardPoint. Otherwise, the GuardPoint is inaccessible to all users, even for users with ‘apply_key’, and ‘permit’ effects.

Creating GuardPoints

1. Log on to your DSM as an administrator of type Security, Domain and Security, or All.
2. Navigate to **Hosts**.
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.

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.
10. Click **Ok**, the *Edit Host* page opens with the newly created GuardPoint listed in the table.

NOTE: Automount is not supported in a Docker environment.

Viewing GuardPoints

You can view GuardPoints from the Management Console GUI and from the Docker host using the VTE Agent `secfsd` utility. To view GuardPoints using the `secfsd` utility:

1. Log on to your Docker Host as root.
2. At the prompt, type:

```
# secfsd -status guard -tree
```

The output is displayed in a tabular format. The table displays the following information:

- GuardPoint location on the image or container
- Name of the policy applied to the GuardPoint
- Type of directory being guarded
- Container ID
- GuardPoint configuration status; whether or not the GuardPoint has been enabled
- GuardPoint status; whether or not the GuardPoint is currently guarded or not
- Reason for the GuardPoint not being guarded

To view information for each GuardPoint;

1. Log on to your Docker Host as root.
2. At the prompt, type the following;

```
# secfsd -status guard -v
```

The output is displayed for each GuardPoint configured on the host. The following information is displayed for each GuardPoint;

- Name of the policy applied to the GuardPoint
- Directory to which the GuardPoint is applied
- Type of directory being guarded
- GuardPoint configuration status; whether or not the GuardPoint has been enabled
- GuardPoint status; whether or not the GuardPoint is currently guarded or not
- Reason for the GuardPoint not being guarded
- Space usage on the GuardPoint location
- Container ID

Data Security for Docker Images and Containers

The VTE Agent supports data security for directories within Docker images and containers. If a new GuardPoint is added to a directory withing an image or container, and that GuardPoint contains data, that data must be transformed before VTE can apply an encryption policy. Therefore, before creating a GuardPoint, determine which of the following conditions is applicable to your situation.

Setting up an image based GuardPoint

If you are setting up an image based GuardPoint:

1. Create a container from that image using the *docker run* command. For example;

   ```
   # docker run Ubuntu
   ```

   VTE creates a container <*container_name*>.

2. Export the container to a TAR file in a directory using the *docker export* command. The following example creates a directory and exports the TAR file:

   ```
   # mkdir -p /tmp/export/GP1

   # docker export <container_name> >
   /tmp/export/GP1/<container_name>.tar
   ```

3. Extract TAR file using the command *tar -xvf*. The following example creates a directory and extracts the TAR file to that directory;

   ```
   # mkdir -p /tmp/extract/GP1
   ```

```
# tar -xfv /tmp/export/GP1/<container_name>.tar -C
/tmp/extract/GP1/
```

4. Guard extracted folder with a data transform policy, for example guard a folder with sensitive data under */tmp/extract*.
5. Transform files using VTE.
6. Unguard the transformed folder.
7. Create a TAR file from the extracted files using the command `tar -czf`. The following example creates a directory, a TAR file and places it in that directory;

   ```
   # mkdir -p /tmp/import/GP1

   # cd /tmp/extract/GP1

   # tar -czf  /tmp/import/GP1/<container_name>.tar *
   ```

8. Import the TAR file back to the image using the Docker command *docker import*. For example,

   ```
   # cat /tmp/import/GP1/<container_name>.tar | docker import -
   <image>

   sha256:d303c8b357084291f02f5e033660f469705f8e2c1744bf14516783
   b58b0f3bd7
   ```

   The image *<imageName>* is created in this example, and its ID is; sha256:d303c8b357084291f02f5e033660f469705f8e2c1744bf14516783b58b0f3bd7
9. You can now guard directories within this image or container with a production policy.

NOTE: This procedure only works with an image's local file system and cannot transform a Docker data volume or NFS mount.

NOTE: Data transformation only occurs on the directory that is guarded, and not the entire Docker image.

Setting up a container-based GuardPoint

If you are setting up a container based GuardPoint, stop the Docker container before setting up GuardPoints:

1. Export the container to a TAR file in a directory using the *docker export* command. The following example creates a directory and exports the TAR file;

   ```
   # mkdir -p /tmp/export/GP1

   # docker export <container_name> >
   /tmp/export/GP1/<container_name>.tar
   ```

2. Extract TAR file using the command *tar -xvf*. The following example creates a directory and extracts the TAR file to that directory;

   ```
   # mkdir -p /tmp/extract/GP1

   # tar -xfv  /tmp/export/GP1/<container_name>.tar -C
   /tmp/extract/GP1/
   ```

3. Guard extracted folder with a security policy.
4. Transform files using the VTE.
5. Unguard the transformed folder.
6. Create a TAR file from the extracted files using the command *tar -czf*. The following example creates a directory and creates a TAR file and places it in that directory:

   ```
   # mkdir -p /tmp/import/GP1

   # cd /tmp/extract/GP1

   # tar -czf  /tmp/import/GP1/<container_name>.tar *
   ```

7. Import the TAR file back to the image using the Docker command *docker import*. For example,

   ```
   # cat /tmp/import/GP1/<container_name>.tar | docker import -
   <image>
   ```

   ```
   sha256:d303c8b357084291f02f5e033660f469705f8e2c1744bf14516783b58b0
   f3bd7
   ```

   The image, *<image>* is created in this example, and its ID is;
   sha256:d303c8b357084291f02f5e033660f469705f8e2c1744bf14516783b58b0f3bd7

8. You can now guard this image or container with a production policy.

NOTE: This procedure only works with an image's local file system and cannot transform a Docker data volume or NFS mount.

NOTE: Data transformation only occurs on the directory that is guarded, and not the entire Docker image.

Setting up a GuardPoint for an exported Docker volume

If you are setting up a GuardPoint for a Docker volume exported from a Docker host:

1. Guard a folder on a host with a data transform policy.
2. Transform files using the Vormetric executable dataxform.
3. Unguard the folder on the host.
4. You can now guard this image or container with a production policy.

This approach works with Docker data volume and NFS mount, and can not transform the image/container local file system.

Configuring Audit Logging

Configure Log settings for the VTE Agent (FS Agent Log) at the System level on the DSM. These settings are inherited by all the domains on the DSM. However, you can fine tune log settings for a specific host, and those settings will override the system settings. With the introduction of Docker support, you can now configure log settings for Docker images and containers. Docker logs evaluate GuardPoint policies.

Configure Docker Log Settings

1. Log on to your DSM and switch to a domain. Alternately, log on to a DSM as a local domain administrator of type Security 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.
4. Click **Docker Log**.
5. Enter the following information in the **Configure Docker Log Setting** panel:
   a. **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 created, 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.

b. **Policy Evaluation Level**: Select a log message level. For more information about log levels, refer to the *DSM Administrators Guide*.

c. **Policy Evaluation Duplicated**: Select to suppress or allow duplicate messages. The default is SUPPRESS.

6. Click **Ok**. VDS saves the Policy Evaluation settings in a tabular format under the **Configure Docker Log Setting** panel.

Searching for Docker Log Messages

Docker log messages display on the Logs page. To search for Docker specific log messages:

1. Click **Logs > Logs**.
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 server, 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 the 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 display in the table under the **Search** panel.

Generating Reports

The following DSM reports have been updated to include Docker information:

System Level Reports

To view system level reports, log on to a DSM as an administrator of type System or All. The system level reports that contain Docker information are:

- **System License Usage Summary**: Includes information about the total number of Docker licenses in use for the entire DSM.
- **License Usage by Domain**: Includes information about the total number of hosts with the Docker license enabled.

Domain Level Reports

To view domain level reports, log on to a DSM as an administrator of type Domain, Domain and Security, or All. Administrators of type Domain and Security and type Security must have AUDIT role privileges to access the reports. The domain level reports that contain Docker information are:

- **License Usage by Domain Summary**: Includes information about the total number of hosts with the Docker licenses enabled.
- **Host with GuardPoint Status**: Includes identification information about Docker Images and Docker containers that have GuardPoints. The columns are **Docker Image ID** and **Docker Container ID**.

You can download and save all reports locally in CSV format by clicking **Download**.

RedHat OpenShift Containers with VTE

Red Hat OpenShift Container Platform (OCP) provides an immutable, container-based platform to deploy and run applications and micro services. It is Red Hat's on-premise private PaaS product. It is built around a core of application containers powered by Docker Container Packages and Kubernetes Container Cluster Management, on a foundation of Red Hat Enterprise Linux.

Using the VTE Agent

In order to use the VTE Agent to protect OCP images and containers, you must obtain a VTE Agent license for Docker. The Docker license covers both Docker and OCP. Contact Vormetric Support for information on obtaining a license.

There is no change to the VTE installation, however, VTE agent must be installed on the OCP host system.

VTE: Virtual Machine versus OCP

Similarly to Docker, VTE installs on the OCP container host and orchestrates security from the OCP host. It does not leave a footprint inside the OCP image or containers. You setup the GuardPoints inside the individual containers at runtime.

Set the OpenShift Storage Driver

Although needed for Docker, this is not needed for OCP.

Administering the OpenShift Host

To protect data inside OCP containers, you need to create GuardPoints in the DSM, inside an OCP image or container, to which a VTE Agent policy is applied.

The administrative tasks related to OCP hosts are the same as for Docker hosts. See “Administering the Docker Host” for more information.

Enable OpenShift through Host Settings

If you have obtained a Docker license and enabled it through the DSM Host Settings, then OCP is also enabled.

VTE OCP GuardPoints

OCP typically provides two types of containers:

- **Transient**: Run micro services which are stateless applications
- **Long running**: Host stateful applications similar to a database application.

The VTE OCP security feature protects data in both type of containers. You can use VTE to protect either type of containers. You select the container type while configuring the GuardPoint.

Types of GuardPoint

VTE provides two types of GuardPoints:

- Image-based
- Container (POD)-based

Image-based GuardPoints

You can set up a data protection policy on an OCP image. After an image-based GuardPoint is created, all of the instances, running from the protected OCP image, inherit the policy and its settings. Any change to the policy is reflected across OCP containers that are started from protected OCP images. You an also refer to Image-based protection as templatized protection. The GuardPoints set up on an OCP image serves as a template for protection of all of the OCP containers created as VTE protected OCP images.

Users can browse an OCP image to select the path for protection and configure security rules using information from an OCP image. This process is described in the DSM Guide. It is the same procedure as for a Docker image.

Container/POD-based GuardPoints

You can set up a GuardPoint for a specific OCP container in the same manner that you do for a Docker container. You can browse an OCP image to select the path for protection, and to configure security rules using information from an OCP image.

Creating GuardPoints

Create GuardPoints in the same manner as for Docker.

Viewing GuardPoints

View GuardPoints in the same manner as for Docker.

Data Security for OpenShift Images and Containers

The VTE Agent supports data security for directories within OCP images and PODS/containers. If a new GuardPoint is added to a directory withing an image or container, and that GuardPoint contains data, that data must be transformed before VTE can apply an encryption policy. Therefore, before creating a GuardPoint, determine which of the following conditions is applicable to your situation.

Setting up an Image-based GuardPoint

Set up a GuardPoint in the same manner as for Docker.

Setting up a POD-based GuardPoint

Set up a POD GuardPoint in the same manner as for Docker GuardPoints. Select the POD as you would a container.

Setting up a GuardPoint for an exported OCP volume

Set up a GuardPoint in the same manner as for exported Docker volume.

Configuring Audit Logging

Configure Audit logging in the same manner as for Docker.

Generating Reports

Generate Reports in the same manner as for Docker.

Creating an OCP Project in CLI with API Commands

Creating an OCP Project with a Template

You can create an OCP project in the CLI as well as in the GUI.

NOTE: RedHat OpenShift is OpenSource technology. Therefore, commands, references and documentation are subject to change. Thales is providing CLI commands that are current at this time. Thales cannot guarantee the integrity of these commands. If commands no longer work, consult the OpenShift developer documentation located at:
https://docs.openshift.com/enterprise/3.2/cli_reference/basic_cli_operations.html#cli-reference-basic-cli-operations

1. Login to the OCP server with a user name and password, type:

   ```
   # oc login <ocp-server> -u <username> -p <password>
   ```

2. Create a new project, type:

   ```
   # oc new-project <project-name>
   ```

3. Add an instant application template to your project and deploy it, type:

   ```
   # oc process openshift <instant-app-template-path> | oc create
   -f -
   ```

   Example

   ```
   # oc process openshift//django-psql-example | oc create -f -
   ```

Deploying an OCP Project

Run the following commands after deployment completes.

1. Get pod details, type:

   ```
   # oc get -o name pods
   ```

2. Parse JSON output and get container details, type:

   ```
   # oc get -o json <pod-name>
   ```

3. Create directory to be guarded in containers

   ```
   # oc exec <pod-name> -c <container-name> -- <command>
   ```

   ```
   # oc exec postgresql-1-bag25 -c postgresql -- mkdir
   /var/tmp/gp
   ```

4. Guard path inside all containers and inside all pods, type:

   - Container-based guarding

```
# vmssc host addgp -d <guard-path> -p <policy> -c
<container-id> -i <image-id> <hostname>
```

- Image-based guarding

```
# vmssc host addgp -d <guard-path> -p <policy> -i <image-id>
<hostname>
```

NOTE: You can also guard paths in the DSM UI.

5. Unguard all guarded paths, type:
   - Container-based guarding

```
# vmssc host delgp -d <guard-path> -p <policy> -c <container-
id> -i <image-id> <hostname>
```

   - Image-based guarding

```
# vmssc host delgp -d <guard-path> -p <policy> -i <image-id>
<hostname>
```

NOTE: You can also guard paths in the DSM UI.

6. Delete project, type:

```
# oc delete project <project-name>
```

Available OpenShift commands

Commands	Function
clusterresourcequota	Create cluster resource quota resource.
configmap	Create a configmap from a local file, directory or literal value
deployment	Create a deployment with the specified name.
deploymentconfig	Create deployment config with default options that uses a given image.
identity	Manually create an identity (only needed if automatic creation is disabled).
imagestream	Create a new empty image stream.
namespace	Create a namespace with the specified name

Commands	Function
policybinding	Create a policy binding that references the policy in the targetted namespace.
quota	Create a quota with the specified name.
route	Expose containers externally via secured routes
secret	Create a secret using specified subcommand
service	Create a service using specified subcommand.
serviceaccount	Create a service account with the specified name
user	Manually create a user (only needed if automatic creation is disabled).
useridentitymapping	Manually map an identity to a user.

Available OPC Options

Options	Parameter	Function
-f	--filename	Filename or URL to file to read a template.
-l	-labels	Label that you can set in all resources for this template.
-o	-o, --output='json'	Output format. It is either: describe|json|yaml|name|template|templatefile.
-o	--output-version	Output the formatted object with the given version (default api-version).
-o	--parameters=false	Do not process but only print available parameters.
-o	--raw=false	If true, output the processed template instead of the template's objects. Implied by -o describe.
-t	--template	Template string or path to template file to use when -o=template or -o=templatefile. The template format is golang templates. [http://golang.org/pkg/text/template/#pkg-overview]
-v	-v, --value=[]	Specify a key-value pair (ex: -v FOO=BAR) to set/override a parameter value in the template.

Container secfsd Utilities

Move to the Container-appropriate name space for the instance.

Use the following commands for more information.

```
# secfsd -[command] [option]
```

Commands	Function
-status guard [-v/-tree]	list all GuardPoints
-status keys	show current encryption key state
-status auth	list authentication settings
-status lockstat	show status of system and agent lock

Commands	Function
-status logger	list logging details
-status policy	list configured policies
-status pslist	list protected process
-status devmap	list guarded devices
-guard path [*containerID*]	manually guard path
-unguard path [*containerID*]	manually unguard path
-version	show version of kernel module secfs2
cmd -c debug.<level>.[on|off]>	set debug logging on/off
-debug <on|off>	enable verbose logging
-help	Displays this help message

10 NetApp Snapshot Directory

This chapter describes SecFS support for NetApp .snapshot directory over NFS. It contains the following sections:

Overview

The NetApp snapshot directory contains ONTAP snapshot data entries for a specific live volume. Each snapshot is a read-only volume that is automatically mounted over NFS.

A snapshot copy is a read-only image of a traditional, or FlexVol volume, or an aggregate, that captures the state of the file system at a specific point in time.

Data ONTAP maintains a configurable snapshot copy schedule that creates and deletes snapshot copies automatically for each volume.

Accessing snapshots

By default, every volume contains a directory named .snapshot through which users can access previous versions of files. Users can gain access to snapshot copies depending on the file-sharing protocol used, NFS or CIFS. You can also prevent access to snapshot copies.

Snapshot files carry the same read permissions as the original file. A user who has permission to read a file in the volume, can also read that file in a snapshot copy. A user without read permission to the volume cannot read that file in a snapshot copy.

NOTE: Snapshot copies do not have write permissions.

Snapshot directories only display at the mount point, although they actually exists in every directory in the tree. This means that the .snapshot directory is accessible by name in each directory, but is only seen in the output of the `ls` command at the mount point. The snapshots are stamped with the date and time.

Enabling Snapshots

The NetApp storage administrator, or the OnTap device, must configure this feature. No configuration is required through VTE. VTE guards the client directory mounting the OnTap data volume over NFS.

NOTE: NetApp documentation is located here: https://nt-ap.com/2vEnEeJ

Dataxform Considerations

You cannot transform snapshot directory entries with Dataxform with a new key, because the snapshots are read only. You must keep previous keys and alter the running security policy accordingly to maintain access to the older snapshot entries alongside any new snapshots taken with the new key.

Also, any snapshots that get created during the data transform process (this may take a long time) have to be discarded/deleted as it may contain a mix of data blocks encrypted with both old and new keys.

Best Practices

Maintaining keys for access to older snapshots can be tedious and cumbersome. Therefore, the simplest and safest practice is to delete all old snapshots once the data is transformed with a new key.

This allows for all new snapshots to be readable with the new key while old keys can be discarded, unless used in other security policies.

11

Secure Start

This chapter describes encrypting an Microsoft Active Directory (AD) with the Secure Start feature. It contains the following sections:

Secure Start Overview

Secure Start offers data protection for applications which start earlier in the boot sequence than VMD (Vormetric Daemon).For example, the Microsoft Active Directory system service starts very early in the boot sequence. To determine if another application qualifies, contact Thales technical support.

NOTE: Secure Start is included with VTE v6.0.2. You do not have to purchase it separately.

NOTE: Secure Start is supported on Windows Server 2008 R2 and later versions. It is not supported on Linux.

In VTE v6.0.1 and below, VTE cannot encrypt the AD system service because it boots earlier in the boot sequence than the VMD agent service. After it boots, VMD makes a secure connection with the DSM and retrieves the encryption keys from the DSM. Therefore, since AD boots before that, the AD system service is not encrypted.

VTE offers Secure Start through a new type of GuardPoint. A Secure Start GuardPoint starts before the AD services, and can, therefore, encrypt those services.

There are three methods for encrypting the directory:

- "Encrypt by Moving the AD Service into a Guarded Directory" on page 153
- "Encrypt Data in Place with Offline Transformation" on page 155
- "Encrypt with an LDT Transformation Policy" on page 156

Prerequisites

Prior to using Secure Start to guard your AD database:

1. Backup your AD database:
   a. Navigate to **Administrative Tools**.
   b. Click **Windows Server Backup.**
   c. Click **Action > Backup Once**.
   d. Follow the instructions in the Backup Wizard to create a backup of the server in a local drive.

NOTE: When the backup operation completes, it saves the server backup in <*backup drive*>:\WindowsImageBackup\<*BackupComputerName*>.

2. Perform a system state backup.
3. Obtain the Microsoft DSRM (Data Services Restore Mode) password.
4. Ensure that your AD database is not in `c:\Windows\NTDS`.

Warning! Do not put your AD database in `c:\Windows` or `c:\Program files`. Secure Start cannot encrypt or decrypt any files in those folders.

Encrypt by Moving the AD Service into a Guarded Directory

You can move the AD service into a directory protected by a standard or LDT production policy. This method does not require the initial data transformation step. When you move the AD service into this directory, VTE immediately encrypts the data with either policy.

NOTE: This step occurs when the system is in DSRM mode, so users have no access to the AD service.

Create the AD GuardPath directory

Create the directory in which the AD service will reside.

1. Log in to the Active Directory Server in DSRM mode using the DSRM password. User ID is Administrator.
2. Create a folder to which you will move the AD database.

Apply Secure Start GuardPoints to a Directory

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 GuardPoints:

1. In the DSM, click **Hosts > Hosts > *<hostName>***
2. In the General host information section, select the option: **Secure Start GuardPoint**.
3. Click **Guard FS.**
4. Select the directory and click **Guard**.
5. In the Policy field, select an **LDT** or **Standard Production** policy.
6. Set Type to **Directory (Auto Guard)**.
7. Click **Browse** and navigate to the folder that you just created for the AD database.
8. Select the option: **Secure Start**.
9. Click **OK**.

10. Select the GuardPoint and click **Secure Start On.**

Verify the Secure Start GuardPoint with CLI

After the DSM pushes the policy to the Active Directory Server, verify the GuardPoints.

To verify the GuardPoints, type:

```
> voradmin ss verify <GuardPoint_path>
```

System Response

```
Successfully completed the command verify
Success from kernel -Successfully verified the secure start GP
```

Move the AD Database into the Secure Start GuardPoint

Move your AD database from the default location (`c:\windows\NTDS`) to this newly created protected folder.

To move the AD database:

1. In DSRM mode, login using the DSRM password. User ID is Administrator.
2. Start NTDSUTIL utility, type:

   ```
   > activate instance ntds
   ```

3. Type:

   ```
   > files
   ```

4. Type:

   ```
   > move db to \<GuardPoint>
   ```

5. Type:

   ```
   > move logs to \<GuardPoint>
   ```

6. Exit NTDSUTIL utility.
7. Reboot the system into normal mode. The Active Directory Services automatically starts after rebooting.

NOTE: This step occurs when the system is in DSRM mode, so users have no access to the AD service.

Encrypt Data in Place with Offline Transformation

Encrypting the AD database with a standard (production), or offline policy is very similar to encrypting other data with a standard (production), or offline policy. For information on standard policies and encrypting data offline or with a production policy, see the *VTE Data Transformation* guide.

The advantage to encrypting data in place is that it saves space. When you copy/move a directory into a guarded directory, you will need twice as much space to store the data because you leave a copy of the data in the original folder, as a precaution, until the original directory has been successfully moved and encrypted. Once the data is transformed, then you can delete the directory that contains the decrypted/clear data.

Using this method, you perform an Initial Data Transformation using the `dataxform` command line utility. During this transformation, access to the GuardPoint data is blocked. After initial transformation, you remove the initial policy, and then apply a production policy, so users can access the data.

NOTE: This step occurs when the system is in DSRM mode, so users have no access to the AD service.

NOTE: If your AD service is installed in the default directory, `C:\Windows\NTDS`, you must move it to another directory before you can encrypt it. See "Encrypt by Moving the AD Service into a Guarded Directory" on page 153 for more information.

To encrypt the data:

1. In DSRM mode, login using the DSRM password. User ID is Administrator.
2. Create and apply a `dataxform` policy to the GuardPoint directory.

3. Run the `dataxform` command.
4. Remove the `dataxform` policy on the GuardPoint and replace it with a production policy.
5. Reboot out of DSRM mode.

Encrypt with an LDT Transformation Policy

Encrypting the AD database with an LDT policy uses the same steps as encrypting with a standard production policy. The only difference is that you select an LDT policy instead of a standard one. See "Encrypt by Moving the AD Service into a Guarded Directory" on page 153 for more information.

NOTE: If your AD service is installed in the default directory, `C:\Windows\NTDS`, you must move it to another directory before you can encrypt it.

Configure the Time Out Failure

During the initial access to a Secure Start GuardPoint, the VTE agent sets a timer. The default duration is 120 seconds, but you can configure the duration. Minimum duration is one second, maximum duration is 300 seconds.

Data inside the GuardPoint is accessible without DSM connectivity until the timeout is reached. VMD service activates and makes a secure connection to the DSM. After the VMD makes a secure connection, the agent verifies that it is connected to correct DSM. If the VMD fails to connect to the DSM, the timeout is reached, and if AD is installed, the agent shuts down the system for data security purposes.

NOTE: In DSRM mode, when the timeout occurs, VTE removes the keys from memory. However, VTE does not shut down the system.

In normal mode, VTE shuts down the AD server. For any other application, or if AD is not installed, Secure Start does not shut down the server. However, the data

inside the GuardPoint becomes inaccessible until DSM connectivity is restored, or you issue a challenge/response, or password. After the timer has expired, VTE denies any further access to the Secure Start GuardPoint.

1. To configure the timeout duration in seconds, type:

   `C:\> voradmin ss settimeout <timeout>`

   Example

   `C:\> voradmin ss settimeout 220`

   System Response

   ```
   Successfully completed the command settimeout
   Successfully set the Secure Start timeout value to 220 seconds
   ```

   To verify the timeout duration, type: **`C:\> voradmin ss gettimeout`**

   System Response

   ```
   Successfully set the Secure Start timeout value to 220 seconds
   ```

Recover a Server after it loses connection to the DSM

User must unlock the GuardPoint by entering the Challenge/Response or password to restore the DSM connectivity. Once the GuardPoint is unlocked, you can start the AD services manually.

NOTE: The challenge response pop up dialog does not display in DSRM mode when the host loses DSM connectivity.

To activate the challenge/response:

1. Navigate to:

   `C:\Program Files\Vormetric\DataSecurityExpert\agent\shared\bin`

2. Double-click `etray.exe` to start it manually.

If the AD server does not connect to the DSM, then the AD system automatically shuts down. The Administrator must enter DSRM mode and restore the DSM connectivity to recover the server.

Prerequisites

Before rebooting your active directory servers:

- Ensure that DSM connectivity is strong. If it is not strong, restore the DSM connectivity.

NOTE: When trying to fix a DSM connectivity issue, you can log in to DSRM mode. In DSRM mode, there is no requirement to increase the timeout, because in DSRM mode, the AD system does not shut down after timeout expires.

DSRM Mode

The first method for recovering a server relies on manual DSM connection troubleshooting:

1. Boot into DSRM mode.
2. Attempt to resolve why the server is not connecting to the DSM.
3. Fix that DSM connectivity issue.
4. Reboot into normal mode.

Other Use Cases

Using Secure Start GuardPoints, you can also secure an SQL Server on Microsoft Azure in certain scenarios. SQL system services in Azure also boot earlier in the boot sequence than the VMD (Vormetric Daemon) agent service.

NOTE: To determine if another application qualifies, contact Thales technical support.

Boot a Windows Server in Azure

To move and guard the AD database, you must boot the AD server into DSRM mode.

To boot a Windows Server 2012/2016 Domain Controller into DSRM remotely in Azure:

NOTE: The Windows Server 2012/2016 domain controller must be running and accessible through Windows Remote Desktop.

1. Establish a Remote Desktop session on the domain controller.
2. Open an command prompt as Administrator and type:
   ```
   > bcdedit /set safeboot dsrepair
   ```
3. Reboot the domain controller. The Remote Desktop session disconnects.
4. Wait a few minutes, then establish a new Remote Desktop session. The domain controller will be running in DSRM.
5. To reboot into normal mode, open an command prompt as Administrator and type:
   ```
   > bcdedit /deletevalue safeboot
   ```
6. Reboot the domain controller.

Best Practices for Encrypting and Protecting the AD Service

Thales recommends the following best practices when using Secure Start with an AD service.

Access Control with Secure Start

User can setup a restricted access control policy with encryption to prevent the unauthorized access of AD database files. The restricted policy with Secure Start:

1. Prevents a rogue user from logging into the system, and moving or copying the AD database files to another directory and tampering with it.
2. Denies permissions, after you setup and guard files, so that no one can move a file from the guarded directory. Plus it restricts any other unwanted/unnecessary process or users from tampering with AD files.

3. Provides permission for an authorized user who needs access to AD services and files.

Creating a Minimal Policy required for AD with Access Control

When creating a normal, strict policy for access control, you must allow access to the following processes and directories for Active Directory.

- **Processes**

```
secfsd.exe (C:\Program
Files\Vormetric\DataSecurityExpert\agent\secfs\ sec\bin\)
lsass.exe (C:\Windows\System32\)
vds.exe (C:\Windows\System32\)
vssvc.exe (C:\Windows\System32\)
wbengine.exe (C:\Windows\System32\)
ntoskrnl.exe (C:\Windows\System32\)
```

- **Users**

```
NT AUTHORITY\SYSTEM
```

To create a minimal policy:

1. Create a User Set named: **AD_Minimum_User_Set**.
2. Set with the following parameters:

ID	uname	osDomains
1	SYSTEM	NT AUTHORITY

3. Create a Process Set named: **AD_Process_Set**.
4. Set with the following parameters:

ID	Directory	Base Name
1	C:\Program Files\Vormetric\DataSecurityExpert\agent\secfs\sec\bin	secfsd.exe
3	c:\Windows\System32\	ntoskrnl.exe
4	c:\Windows\System32\	vds.exe
5	c:\Windows\System32\	vssvc.exe

ID	Directory	Base Name
6	c:\Windows\System32\	wbengine.exe
7	c:\Windows\System32\	lsass.exe

5. Create a Security rule.
6. Set with the following parameters:

Order	User	Process	Action	Effect	Browsing
1	AD_Minimum_User_Set	AD_Process_Set	all_ops	Audit, Permit, Apply key	Yes
2				Audit, Deny	Yes

Creating a Restricted policy in DSRM mode

Create the following policy for the initial transformation of an AD database in DSRM mode. The policy allows access to the local administrator.

In DSRM mode, you use the `NTDSUTIL` utility to perform maintenance for an Active Directory.

To create a restricted policy:

1. Create a User Set named: **AD_Minimum_User_Set**.
2. Set with the following parameters:

ID	uname	osDomains
1	SYSTEM	NT AUTHORITY
2	Administrator	localhost

3. Create a Process Set named: **AD_Process_Set**.
4. Set with the following parameters:

ID	Directory	Base Name
1	C:\Program Files\Vormetric\DataSecurityExpert\agent\secfs\sec\bin	secfsd.exe

ID	Directory	Base Name
2	c:\Windows\System32\	ntdsutil.exe
3	c:\Windows\System32\	ntoskrnl.exe
4	c:\Windows\System32\	vds.exe
5	c:\Windows\System32\	vssvc.exe
6	c:\Windows\System32\	wbengine.exe
7	c:\Windows\System32\	lsass.exe

5. Create a Security rule.
6. Set with the following parameters:

Order	User	Process	Action	Effect	Browsing
1	AD_Minimum_User_Set	AD_Process_Set	all_ops	Audit, Permit, Apply key	Yes
2				Audit, Deny	Yes

Guard Directories

The best practice for guarding a directory with a Secure Start GuardPoint is to:

1. Create a directory.
2. Guard that directory with a standard production or LDT policy. Follow the steps in "Apply Secure Start GuardPoints to a Directory" on page 153.
3. Move the AD service into that directory.

Perform Subsequent System State Backups

After you move an AD service into a guarded directory, or out of a guarded directory:

1. Perform another system state backup.
2. Save this subsequent backup to a different location.

12 Enhanced Encryption Mode

This chapter describes the enhanced AES-CBC-CS1 encryption mode for keys. It contains the following sections:

The AES-CBC-CS1 encryption is superior to the existing AES-CBC mode because it uses a unique and unpredictable (random) IV (initialization vector) generated for each individual file. The per-file IV object is generated only at file creation time. It is stored as file metadata.

NOTE: AES-CBC-CS1 encryption does not require any additional license.

Table 10: Features Comparison

	AES-CBC	AES-CBC-CS1
Security Improvements		
Unique IV per-file	No	Yes
IV predicatibility	Yes	No
File System Support		
Local FS (Linux)	EXT3/EXT4/XFS/VXFS	EXT3/EXT4/XFS/VXFS
Remote FS (Linux)	NFS3/NFS4/CIFS	NFS3/NFS4
Local FS (Windows)	NTFS/ReFS	NFS3/NFS4/CIFS
Remote FS (Windows)	CIFS	CIFS (if the backend storage for the CIFS share is Windows-based storage).
Block Device Support (secvm)	Fully supported	No. When a policy contains a key with CBC-CS1 encryption mode, the guarding fails on the DSM, and an error message displays.

Compatibility

VTE is backward compatible with and fully supports the existing AES-CBC mode, both for new and existing datasets, after the Agent is upgraded to VTE

Caution: AES-CBC-CS1 encryption is **only supported** with VTE 6.1.0 and later versions. A pre-6.1.0 VTE host is incapable of supporting AES-CBC-CS1. On these earlier versions, attempting to guard using a policy containing an AES-CBC-CS1 key will fail.

- LDT and Offline dataxform support AES-CBC-CS1 encryption on VTE Linux and Windows environments.
- VTE hosts supporting AES-CBC-CS1 encryption can be added to host groups

Table 11: Data Transformation

	AES-CBC	AES-CBC-CS1
Offline data transformation	Supported	Supported
Live transformation	Supported	Supported

Difference between AES-CBC and AES-CBC-CS1

The two encryption modes are completely different from a file format standpoint.

- AES-CBC-CS1 encryption only applies to file system directories; AES-CBC encryption applies to both files and block devices.

NOTE: If you attempt to use an AES-CBC-CS1 key to guard a block device or partition, the guarding fails with an error reported on the DSM, similar to: "Raw or Block Device (Manual and Auto Guard) Guardpoints are incompatible with Policy “policy-xxx" that contains a key that uses the CBC-CS1 encryption mode.”

IMPORTANT: While AES-CBC-CS1 encryption is supported on both Linux and Windows environments, the file formats are incompatible. An encrypted file

created with a specific AES-CBC-CS1 key on Windows cannot be read on Linux, even if that specific key were to be used and vice versa.

- AES-CBC-CS1 uses cipher-text stealing to encrypt the last partial block of a file whose size is not aligned with 16 bytes.
- Each file encrypted with an AES-CBC-CS1 key is associated with a unique and random base IV.
- AES-CBC-CS1 implements a secure algorithm to tweak the IV used for each segment (512 bytes) of a file.

Disk Space

Files encrypted with AES-CBC-CS1 keys consume additional disk space in contrast to files encrypted with AES-CBC keys. This is because AES-CBC-CS1 encryption requires file IVs to be created and persistently stored -- in contrast to AES-CBC encryption which does not consume any additional disk storage.

Therefore, Administrators need to plan and provision additional disk capacity prior to deploying AES-CBC-CS1 encryption.

Table 12: File size changes

	AES-CBC	AES-CBC-CS1
Local FS (Linux)	No change to file size. No extended attribute allocation	Internal use of extended attribute per file. Possible file size increase depending on LDT policy
Remote FS (Linux)	No change to file size. No extended attribute allocation	Extra 4KB allocation in the form of an embedded header per file. With VTE guarding enabled, file size expansion is hidden.
Local FS (Windows)	No change to file size. No ADS allocation.	Extra 4KB allocation (at minimum) in the form of an embedded header per file. With VTE guarding enabled, file size expansion is hidden.

Encryption Migration

You can use both LDT or offline dataxform to:

- Transform data encrypted by AES-CBC to AES-CBC-CS1 and vice versa
- Transform AES-CBC-CS1 encrypted data to clear contents

File Systems Compatibility

On Linux, Windows, , you can use AES-CBC-CS1 keys to guard currently supported file systems.

Local and Remote File Systems

AES-CBC-CS1 encrypted files on Linux remote file systems like NFS and CIFS increase the file size compared to encrypted files on Linux local file systems which retain the original file size.

AES CBC CS1 encrypted files on Linux local file systems, in conjunction with LDT policies, can result in additional space consumption. Unlike the current AES CBC encryption where encrypted files on all file systems, both remote or local, have the same file format, AES CBC CS1 encrypted file formats differ based on whether or not they were created on local or remote file systems.

AES-CBC-CS1 files on Linux remote file systems such as NFS and CIFS embed the IV in a 4K-byte header within the file. When these files are guarded, VTE masks the file header -- to applications and system utilities. The expanded file is only apparent when VTE guarding is disabled.

NOTE: The remote file system must have enough extra space to store the extra 4K bytes of the embedded header. You can run the following script in the GuardPoint to identify how much space to reserve before data transformation:

```
x=$(find . -type f | wc -l); y=$(echo "$x * 4 /1024" | bc);
echo ${y}MB
```

File System Requirements

Unlike with AES-CBC encryption, files encrypted with AES-CBC-CS1 on remote file systems cannot be copied over to local file systems in the absence of VTE guarding.

Similarly, AES-CBC-CS1 encrypted files on local file systems cannot be copied over to remote file systems in the absence of VTE guarding.

The fundamental reason for this incompatibility is the usage of extended attributes on local file systems to store the IV, in contrast to its storage as a part of the file metadata on remote file systems. This is why files cannot be transferred across these file system boundaries in the absence of VTE guarding.

Table 13: Physical file system requirements

	AES-CBC	AES-CBC-CS1
EXT3/EXT4 on RHEL6	None	EXT3/EXT4 must be mounted with *user_xattr* mount option
EXT3/EXT4 on other Linux distros	None	No limitation
XFS/VXFS on Linux	None	No limitation
NFS v3,v4/CIFS on Linux	No limitation	No limitation

Samba Share

The remote Samba share server does not support ADS so you cannot use the CBC-CS1 key type on these GuardPoints.

Storing Metadata

- AES-CBC-CS1 encrypted files on Linux store the base IV in either the extended attributes or in the file metadata. On local FS (EXT3/EXT4/XFS/VXFS), it saves it as an extended attribute associated with the file. It saves the base IV of a file on remote FS (NFS and CIFS) in the embedded header of the file.
- AES-CBC-CS1 encrypted files on Windows stores the IV in a Windows ADS (Alternate Data Streams) associated with the file. The size required for saving the CS1 key depends on the allocation size of the file system. If the allocation size is set to 4k, then the new IV will require 4K of extra space on the disk. User can run `fsutil fsinfo` tool to find out the allocation size of the file system.
  - On Windows, CS1 key is supported on following file systems:
    - **NTFS**: Disks formatted with NTFS file system on all platforms
    - **REFS**: File system on Windows 2012 R2 and later

- **CIFS**: If the backend storage for the CIFS share is Windows-based storage.

Some network storage servers do not support multiple ADS associated with a file.
To get the value of the base IV, type:

```
# voradmin secfs iv get <file-name>
```

NOTE: The base IV of a file is protected. It cannot be set/modified/removed by commands and applications. However, if a GuardPoint is unguarded, the files in the GuardPoint are no longer protected. An adversary can then corrupt the content of the files, as well as the IVs.

- AES-CBC-CS1 encryption for VTE Linux local file systems like XFS, EXT and VxFS file systems store the file IV as a part of the file metadata. The underlying file system requires that you mount it with the extended attribute mount option.
- AES-CBC-CS1 depends on the physical file system's support for extended attributes in a manner similar to the LDT feature.

Missing IV file

If you are using the CBC-CS1 key, then each file has a unique IV associated with the file. This IV is stored in ADS on Windows.

If the IV for a file is missing, or VTE is unable to read the IV, then VTE denies access to the file. This access denied message may trigger an application to display an error message. This message may vary from application to application.

Table 14: IV Storage

	AES-CBC	AES-CBC-CS1
Local FS (Linux)	No change	Internal extended attribute for each file
Remote FS (Linux)	No change	4KB embedded header for each file
Local FS (Windows)	No change	Alternate Data Streams

HDFS

The AES-CBC-CS1 key is compatible with current Hadoop File System support.

Backups

Backups and other data protection utilities should take into account the extended attributes present in each AES-CBC-CS1 encrypted file on a Linux local file system to ensure that they are safely backed up. An AES-CBC-CS1 encrypted file whose IV is corrupted, renders the files to be corrupted and therefore unreadable. Hence all data protection software must preserve the file's extended attributes.

VTE Linux can inspect a file's IV using the following command:

```
# voradmin secfs iv get file
```

On Linux, the backup utility specified in the guarding policy should automatically backup/restore extended attributes as well. For example, you must use the options to preserve extended attributes when running `cp` or `rsync.normal`.

Due to the different file formats, the backup/restore across the local and remote file systems are not allowed. If you want to backup a GuardPoint from a local directory, you must restore it to a local directory. If a GuardPoint is backed up on a remote file system, you must restore it to a remote system.

Table 15: Backup and Restore

	AES-CBC	AES-CBC-CS1
General backup utility requirement (all platforms)	Backup utility defined in guarding policy with clear view	Backup utility defined in guarding policy with clear view
Special requirement for backup/restore local fs on Linux	No	Backup utility must be run with user extended attribute enabled
Special requirement for backup/restore remote fs on Linux	No	No
On Linux, backup local fs and restore to remote fs	Allowed	Not allowed, the restored files cannot be accessed with I/O error
On Linux, backup remote fs and restore to local fs	Allowed	Not allowed, the restored files cannot be accessed with I/O error

FileTable Support (Windows Only)

The CBC-CS1 key does not support FileTables. This is because FileTables do not support alternate data streams. The CS1 key requires the ability to write the per-file IV into an alternate data stream on each file.

Container Compatibility

The CBC-CS1 key is compatible with current Docker and OpenShift support.

Using the new Encryption mode

Deploy the new encryption mode (AES-CBC-CS1) by using the new symmetric agent key type created in DSM v5.3:

1. In the DSM, click **Keys > Agent Keys > Keys.**
2. Click **Add**.
3. In the Encryption Mode dropdown, select **CBC-CS1**.
4. In the Algorithm dropdown, select **AES128** or **AES256** to create an AES-CBC-CS1 key.
5. Add the key to your policy.

Exceptions and Caveats

On RHEL6, EXT3/EXT4 are not mounted with user extended attributes enabled, by default. If a GuardPoint is on EXT3/EXT4, then remount EXT3/EXT4 with `user_xattr` as an option. Otherwise, guarding fails with the error "Extended attribute not enabled for GuardPoint."

Guarding existing files without data transformation

You must convert an existing file with clear text through offline data transformation or LDT. If you do not transform the file, then after you guard using an AES-CBC key, the file displays garbled characters.

If you use an AES-CBC-CS1 key, access to the file is blocked with an I/O error.

Best Practices

The following are the recommended practices for deploying host groups with AES-CBC CS1 keys:

- In a host group, do not deploy policies associated with AES-CBC and AES-CBC CS1 keys unless all hosts are intended to run VTE 6.1.0 or later versions.
- If VTE 6.1.0 and older VTE versions are intended to be a part of the same host group, Thales recommends that you use policies without AES-CBC CS1 keys.

13 Exchange DAG

This chapter describes encrypting email databases using Microsoft Exchange database availability group (DAG). It contains the following sections:

Exchange DAG Overview

A DAG is a high-availability (HA) and data-recovery feature of the Microsoft Exchange Server. A DAG, which can consist of up to 16 Exchange mailbox servers, automates recovery at the database-level after a database, server or network failure. You can now use VTE to encrypt Exchange DAG mailboxes.

NOTE: Microsoft Exchange DAG is for Windows only. It is not compatible with Linux.

Figure 12: High Level Overview of Exchange DAG functions

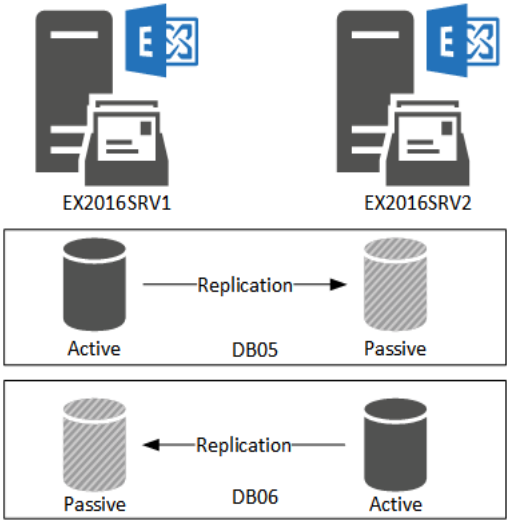

You can encrypt the Exchange databases with a standard (offline) policy or an LDT policy. In an offline policy, users cannot access the database during initial data encryption. LDT is Live Data Transformation. It encrypts the data while users and applications are accessing applications within it. LDT is used for Initial data transformation as well as transparent encryption/decryption.

Use Case tested and supported by Exchange DAG with VTE

Use VTE to encrypt Exchange DAGs in the following scenarios:

- Initial Data Transformation of Exchange Databases using:
  - Live Data Transformation
  - Standard Data Transformation

NOTE: To find out more about LDT and Data Transform, refer the Guides for the software, available from Technical Support.

- Transparent Encryption/Decryption of the Exchange Database on DAG nodes
- Rotate the key using the Live Data Transformation policy

The following Exchange DAG operations were tested using the following use cases:

- Failover/Failback of databases from one node to another node and making both databases active on each node.
- Add new Databases to the existing nodes

Recommendations

To use Microsoft DAG Exchange:

1. Disable all antivirus software as suggested by Microsoft.

   https://technet.microsoft.com/en-us/library/bb332342(v=exchg.160).aspx
2. Disable Windows defender:

   a. Type:

   ```
   > gpedit.msc
   ```

   b. Go to **Computer Configuration > Administrative Templates > Windows Components > Windows Defender**.

   c. Double-click and then, in the context menu, click **Turn Off Windows Defender.**

   For more information see , https://www.windowscentral.com/how-permanently-disable-windows-defender-windows-10.
3. Only guard the Mailbox Database. Do not guard at a higher or lower directory.

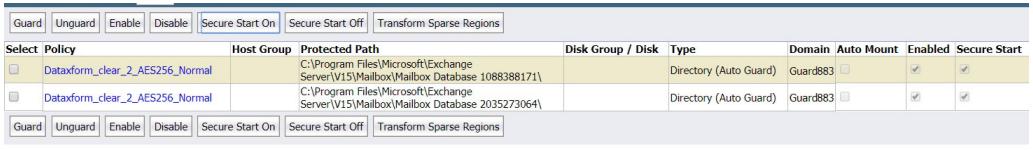

4. Guard one host at a time. Do not apply GuardPoints to hostgroups.
5. Make sure that communication port 7024 is not blocked so that VTE Agent can register with the host.

Requirements

The following is required for encrypting an Exchange DAG database with VTE :

- Windows Exchange Server 2016 Cumulative Patch 5
- Windows Server 2016 for Exchange Server nodes
- Windows 2016 Server as a file share witness
- Two nodes in the Exchange DAG configuration
- Both nodes on the same subnet

NOTE: Thales has not tested an environment with nodes on different subnets, but we do not anticipate it causing an issues.

- Use the same VTE Policy and keys on both nodes.
- You must enable Secure Start on the GuardPoints.
- All users must be offline when applying the initial GuardPoint, (even for LDT).

NOTE: Alternatively, you can choose to make all of the DBs active on one node and perform data transformation on the other node.

- Prepare your environment for suspending databases:

  a. Suspend the Exchange Database to stop replication and access to data.

  b. Disable replication between the nodes.

NOTE: No file access can happen within the target directory

Figure 13: Proper Configuration for Exchange DAG

- Create Keys and Policies for the system

You must use the same Keys and Policies on each node. For more information on Keys and Polices, see the *LDT User Guide* and the *VDS DSM Administrators* guide.

- Install VTE agent on the systems
- Register the agent with the DSM and make sure LDT is enabled, (if they are planning to use LDT. For more information on LDT, see the *LDT Users Guide*.)

Preparing your Exchange DAG Environment for Encrypting/Decrypting with VTE

The following describes how to prepare your environment for encrypting with Exchange.

To use Microsoft DAG Exchange with LDT to encrypt data:

1. In the **Exchange Admin Center**, make Exchange node 1 the primary node.

   Make node 1 the active node and move all of the databases to that node.

Warning! Make sure that all of the Exchange services in node 2 are down and not accessing the Exchange databases. All Exchange Services must be stopped. All databases must be suspended.

2. Make all of the databases active on node 1.
3. Suspend all databases on node 2.

NOTE: Wait for 2-3 minutes for the database to finish with replication so the database will be suspended.

Encrypting with LDT in the Exchange DAG environment

You can use LDT for initial data transformation as well as transparent encryption/decryption.

Warning! You must guard the databases with the same Key/Policy on both nodes.

1. Make sure that you have prepared your environment properly. See "Preparing your Exchange DAG Environment for Encrypting/Decrypting with VTE" on page 177
2. Make sure that the GuardPoint is active on the host. Live Data Transformation starts on the server as soon as the GuardPoint is established.
3. Guard the Mailbox Database directory and apply the Live Data Transformation policy to the directory on node 2.
4. Select the GuardPoint and click **Secure Start On**.

Warning! Only guard the Mailbox Database. Do not guard at a higher or lower directory.

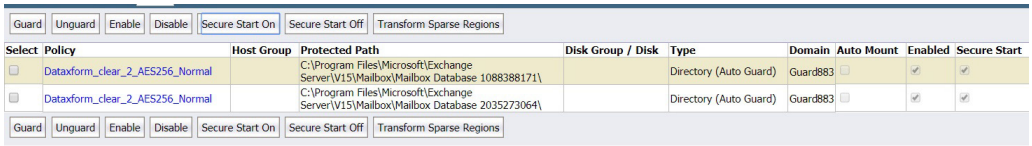

Warning! Guard one host at a time. Do not apply GuardPoints to both nodes using hostgroups.

5. In the **Exchange Admin Center**, go to the Exchange Database tab and **Resume** all Passive database copy on node 2.

NOTE: It may take a few minutes for the Exchange Service to resync. Monitor the Exchange logs on the system and make sure that replication is working. Make sure that database replication finishes and databases are in a healthy state before proceeding.

6. Wait for 10-15 minutes more for the server to move to the healthy state. If not, wait for some more time for the Content Index state to change to Healthy.
7. In the **Exchange Admin Center**, move all of the databases from node 1 to node 2.

   Now the databases on node 1 are mounted as passive. All databases on node 2 are mounted as active.
8. Repeat these steps on node 1.

Decrypting with LDT in the Exchange DAG environment

To use an LDT policy:

1. Make sure that the LDT state is set to REKEYED before unguarding.

2. Make sure that all of the files inside the GuardPoint are at the same version of the key.

   a. Run the LDT report to find the version:

   ```
   > voradmin ldt report <GuardPoint path> [<logfile>]
   ```

   b. Run the Key map report to find the version:

   ```
   > voradmin ldt key [report|map] <key_name, version>
    <GuardPoint path>
   ```

To use Microsoft DAG Exchange with LDT to decrypt data:

1. Make sure that all of the Exchange services in node 2 are down and not accessing the Exchange databases.

NOTE: Suspension can take 2-3 Minutes.

2. In the **Exchange Admin Center**, make Exchange node 1 the primary node.

   This means that node 1 is mounted as the active node and node 2 is mounted as the passive node.

3. Make all of the databases active on Exchange node 1.
4. Go to the Exchange Database tab and suspend all databases on node 2.
5. Unguard the database folders that you previously guarded on node 2.
6. Delete all of the metadata on all of the database folders on node 2, type:

   ```
   > voradmin ldt attr delete [<file name path> | <guard path>]
   ```

7. Guard with an LDT policy set for Encryption to Clear on node 2.

NOTE: You must clone the current version of the encryption key to create the policy.

8. Go to the Exchange Database tab and resume all databases on node 2.

NOTE: After a few minutes, the databases should become healthy automatically. If not wait for the LDT process to decrypt the data. Make sure that all of the data is transformed back to clear and that the LDT state is set to **REKEYED**.

9. Move the database from node 1 to node 2.
10. Repeat Step 2. - Step 8. for node 1.

After both nodes are rekeyed and transformed from encryption to clear, you can unguard them.

1. In the **Exchange Admin Center**, make Exchange node 1 the primary node.

   This means that node 1 is mounted as the active node and node 2 is mounted as the passive node.
2. Make all of the databases active on Exchange node 1.
3. Go to the Exchange Database tab and suspend all databases on node 2.
4. Unguard the database folders that you previously guarded on node 2.

Warning! Always ensure that you are unguarding a passive node.

5. Repeat Step 1. - Step 4. for Node 1.

Encrypting with a Standard VTE policy in the Exchange DAG environment

Refer to the following sections to prepare for encrypting in the Exchange environment.

To use Microsoft DAG Exchange with VTE:

1. In the **Exchange Admin Center**, make Exchange node 1 the primary node.

   This means that node 1 is mounted as the active node and node 2 is mounted as the passive node.
2. Make all of the databases active on Exchange node 1.
3. Go to the Exchange Database tab and suspend all database on node 2.

   Make sure that all of the exchange database services in node 2 are down and not accessing the Exchange databases.

NOTE: Suspension can take 2-3 Minutes.

4. Guard the mailbox database directory with a data transform policy.
5. Run dataxform.exe, type:

   ```
   dataxform --rekey --print_stat --gp <directory>
   ```

6. After the data transformation is finished, unguard the directory and guard with a Production policy. Apply a standard producton data transformation policy on only the node 2 database folders.

NOTE: Use the same Key/Policy on both nodes.

   See “Remove the dataxform policy and apply production policy” in the *VTE Agent DataXform* guide for more information.

7. In the **Exchange Admin Center**, go to the Exchange Database tab and resume all databases on node 2.

   After a few minutes, all nodes should become Healthy.

8. In the **Exchange Admin Center**, try to move a database from node 1 to node 2. If the data move is successful: this means that node 2 is mounted as the active node and node 1 is mounted as the passive node.
9. Repeat the previous steps for node 1.

Unsupported Use cases

The following scenarios are not supported:

- Using different encryption keys on Exchange DAG nodes; both nodes must use the same encryption key
- Adding a new node to the Exchange DAG Environment
- Using the new CBC-CS1 key (not tested for this release)
- Encryption of Exchange Binaries
- Using nodes in a different subnet, data center or site, (Thales is not testing this scenario, but we do not believe it will cause any issues)

14 VTE for Windows Utilities

This chapter describes utilities you can run on Windows. For information on Linux utilities, see Chapter "VTE for Linux Utilities" on page 189.

Vormetric provides a variety of utilities that administrators can use to help manage VTE. This chapter describes the following utilities:

vmsec utility

The `vmsec` utility allows you to manage the security aspect of the VTE agents on the host. The Windows `vmsec` utility is located in `C:\Program Files\Vormetric\DataSecurityExpert\agent\vmd\bin\vmsec.exe`

Syntax

Table 16: `vmsec` Syntax [options]

`check_install`	Verifies that the kernel component is running. This command checks VTE services and reports if any of the services are not running.
`challenge`	Initiates challenge-response on the host. This command displays a File System Agent password challenge string and enter the response string when the DSM is not network accessible.
`status`	Displays kernel configuration.
`vmdconfig`	Displays the vmd configuration.
`check_hwenc`	Determines whether this system supports hardware crypto.
`hwok`	Reports status of hardware signature.
`passwd [-p passwd]`	Enters the host password when the DSM is not network accessible. User can unlock the GuardPoints with this password.

`version`	Displays VTE version.

Agent Health

The `agenthealth.ps` utility validates:

- Super-user privilege
- VTE Agent installation
- VTE registration to DSM Server
- VTE processes/modules that are running
- Available disk resources
- Current GuardPoints

  Tests if the agent can reach the GuardPoints
- VTE log directory resource status

  This directory contains pending VTE log files for upload. This utility reports the size and number of pending files for upload. These text files are logs that contain vmd/SecFS information. They are regenerated whenever secfs restarts. If the number of files is unexpectedly large, this can indicate a problem.

The Agent health check script

The Agent health check script (`agenthealth.ps1`) is located in `C:\Program Files\Vormetric\DataSecurityExpert\agent\shared\bin\`

To run the `Agenthealth` check script:

1. Run the power shell command to enable self-signing for the system.

   Before running agent-health script make sure power shell command has enough privileges to execute the Powershell script. Some Windows operating systems have default execution policy set as **`restricted.`**

   ```
   Use the Powershell command "Set- Execution-policy Remote Signed" to
   change the execution policy if needed.
   ```
2. Open the `Powershell` prompt as administrator.
3. Type:

   `.\agenthealth.ps1`

Example

```
.\agenthealth.ps1
```

System Response:

```
Log file is at
C:\ProgramData\Vormetric\DataSecurityExpert\agent\log\agent_health
.log
Checking super user privilege.............................. OK
Vormetric Agent installation............................... OK
Vormetric policy directory................................. OK
Registration to server..................................... OK
Kernel drivers are loaded.................................. OK
VMD is running............................................. OK
SECFSD is running.......................................... OK
rhat26130.qa1.com is resolvable............................ OK
rhat26130.qa1.com port 8446 is reachable................... OK
rhat26130.qa1.com port 8447 is reachable................... OK
Can communicate to at least one server..................... OK
VMD is listening on port 7024.............................. OK
Time of last update from server2016-12-01         14:39:49.038
Checking available disk space.............................. OK
Checking logging space ........................... OK
        Log directory is "/var/log/vormetric"
        File system for log data is "/", 32G free (17% full)
        Log directory contains 2 of maximum 200 files (1% full)
       Log directory contains 1 of maximum 100 Mbytes used (1% full)
 Testing access to C:\GP2................................... OK
```

agentinfo utility

The `agentinfo` utility collects system logs, VTE agent logs, VTE agent trace information, and system information for diagnostic purposes. All this information is saved in the destination path and compressed into a zip file. The `agentinfo` utility

is available as an `agentinfo.js` Java command and as an `agentinfo.ps1` PowerShell command.

agentinfo utility (Java version)

The `agentinfo.js` utility is a JavaScript file. You can open it in a text editor to see specific functions.

The `agentinfo.js` support collection scripts reside in the following path on systems where the VTE agent is installed.

To run the `agentinfo` script on Windows, navigate to one of the following folders:

```
C:\program files\vormetric\DataSecurityExpert\agent\vmd\bin

or

C:\program files\vormetric\DataSecurityExpert\agent\shared\bin
```

then run the following script:

```
agentinfo.js
```

agentinfo utility (PowerShell version)

The PowerShell version of `agentinfo` supports several parameters.

PowerShell version agentinfo parameters

`Directory` - Specify the directory where all the collection information is saved. By default, this information is saved in the current directory.

`ZipFile` - Specify the name of the compressed file, where all the collected information will be archived. By default, this information is saved in the current directory.

`Logfile` - Specify the name of the files where verbose logs will be saved. By default, this information is saved in the current directory.

Examples for using agentinfo utility (PowerShell version)

To save all the collection information in "`c:\AgentLogs`" folder, run the following command:

```
.\agentinfo.ps1 -Directory 'C:\AgentLogs'
```

To save all the collected information in "`c:\AgentLogs`" folder and verbose logs in "`c:\temp\AgentInfo.log`", run the following command:

```
.\agentinfo.ps1 -Directory 'C:\AgentLogs' -LogFile
'c:\temp\AgentInfo.log'
```

To save all the collected information in "`c:\AgentLogs`" folder, verbose logs in "`c:\temp\AgentInfo.log`", and create the "AgentInfo.zip" archive file, run the following command:

```
.\agentinfo.ps1 -Directory 'c:\AgentLogs' -LogFile
'c:\temp\AgentInfo.log' -ZipFile 'C:\temp\AgentInfo.zip'
```

NOTE: PowerShell 5.1 or later is required. Use the `$PSVersionTable.PSVersion` command to confirm which PowerShell version you are using.

15 VTE for Linux Utilities

This chapter describes VTE for Linux utilities. The Windows utilities are described in "VTE for Windows Utilities" on page 183.

Thales provides a variety of utilities that an administrator can use to help manage VTE. These utilities reside in storage until summoned by the administrator.

The following utilities are described in this chapter:

secfsd utility

The `secfsd` utility displays the following attributes of VTE:

- GuardPoints defined in the *Guard FS* tab
- Authentication parameters defined in the *Host Settings* tab
- Lock status set by enabling FS Agent Locked and System Locked
- Web destination and SSL certificate for uploading log entries
- Policies applied in the **Guard FS** tab
- Status of required processes (`secfsd` and `vmd`)
- Version of `secfs`

The `secfsd` utility is also used to mount GuardPoints for `Directory (Manual Guard)`. Normally, VTE automatically mounts the `secfs` file system when you apply a GuardPoint to a

directory. On Linux/, the `secfsd` utility is located in `<install_dir>/secfs/.sec/bin` and a symbolic link to this file is placed in `/usr/bin/secfsd`.

secfsd syntax

Table 17: `secfsd` Syntax

Command	Description
`-help`	display `secfsd` options
Status Options	
`-status guard [-v | -tree]`	list all GuardPoints
`-status keys`	show current encryption key state
`-status auth`	list authentication settings
`-status lockstat`	show VTE lock status
`-status logger`	list logging details
`-status policy`	list configured policies
`-status pslist`	list protected processes
`-status devmap`	list guarded devices
Manual GuardPoint options	
`-guard path [container ID]`	manually guard path
`-unguard path [container ID]`	manually unguard path
Version option	
`-version`	list version of kernel module `secfs2`

Examples

Updating status file

To create or update the `/var/log/vormetric/statusfile` file, type:

```
# secfsd -status
```

VTE does not remove the file after a configuration change. It updates when you run any of the `secfsd -status` commands.

Display GuardPoint-related information

To display the GuardPoint paths, applied policies, policy type, and guard status, type:

```
# secfsd -status guard
```

System Response

```
# secfsd -status guard
GuardPoint     Policy           Type          ConfigState   Status    Reason
----------     ----------      -----         -----------   ------     -----
/opt/apl/lib   allow AllOps_fs  local          guarded       guarded     N/A
/dev/sdb       watchaccess_rd   rawdevice      guarded       guarded     N/A
/dev/sdc       watchaccess_rd   manualrawdevice guarded      guarded     N/A
/dev/sdd       watchaccess_rd   manualrawdevice  unguarded    not guarded
Inactive
/opt/apl/tmp   MSSQL00123        manual         unguarded     not guarded
Inactive
```

GuardPoint	Full path of the GuardPoint.
Policy	Name of the policy applied to the GuardPoint.
Type	Can be local, automount, manual, raw device, or manual raw device. Configured in the Guard FS tab.
ConfigState	Guard status of the GuardPoint, as recognized by the DSM. It can be guarded or unguarded.
Status	Current guard status, as recognized by VTE. State can vary.

NOTE: Config State and Status can vary. As an example, if you apply a GuardPoint and someone is currently working in the GuardPoint, the policy cannot be applied at that time. In this case, the ConfigState is guarded and the Status is not guarded.

Display GuardPoint-related information in a different format

To display the same information in a different format, include the `-v` argument, type:

```
# secfsd -status guard -v
```

System Response:

```
GuardPoint: 1
        Policy:          allowAllOps_fs
        Directory:       /opt/apps/apps1/tmp
        Type:            local
        ConfigState:     guarded
        Status:          guarded
        Reason:          N/A
GuardPoint: 2
        Policy:          allowAllRootUsers_fs
        Directory:       /opt/apps/apps1/lib
        Type:            local
        ConfigState:     guarded
        Status:          guarded
        Reason:          N/A
GuardPoint: 3
        Policy:          allowAllOps-winusers1_fs
        Directory:       /opt/apps/apps1/doc
        Type:            local
        ConfigState:     guarded
        Status:          guarded
        Reason:          N/A
```

Display GuardPoints in a tree view

Use the `secfsd status guard -tree` to list GuardPoints in a tree view, type:

```
# secfsd -status guard -tree
```

Display host settings

Use the `auth` argument to display the SHA2 hash signature for each VTE host setting, type:

```
# secfsd -status auth
```

System Response:

```
/bin/su
3E765375897E04C39AB17D4C755F50A35195535B6747DBA28DF9BD4AA672DFF9

|authenticator|/usr/sbin/sshd
98FC599D459EDEA52A60AB394B394803B5DAB96B53148DC608732DDA6777FA1A

/usr/sbin/in.rlogind
5C9A0EDD8BF54AE513F039476D21B3032507CF957AA0CB28C368EB8AB6E684FB

/bin/login
0D2EE0B995A30AE382B4B1CA5104715FC8902F457D283BDABAAD857B09259956

/usr/bin/gdm-binary
363780522E3CCF9ABF559F059E437743F9F97BBBB0EE85769007A464AD696BD1

/usr/bin/kdm
BAD41BBCDD2787C7A33B5144F12ACF7ABC8AAA15DA9FDC09ECF9353BFCE614B5
```

The host setting and hash value are also displayed in `/var/log/vormetric/statusfile` Display Key Status

To display the status of VTE keys, type:

```
# secfsd -status keys
```

System Response:

```
Encryption keys are available
```

Display Lock Status

To display the status of VTE locks, type:

```
# secfsd -status lockstat
```

System Response:

```
FS Agent Lock: false
System Lock: false
```

The value is **true** if the lock is applied. The value is **false** if the lock is not applied. **System Lock** corresponds to **System Locked** in the *Host* window. **FS Agent Lock** corresponds to **FS Agent Locked** in the *Host* window.

NOTE: Before you upgrade, remove VTE software, or change operating system files, the status of FS Agent Lock and System Lock must be false.

Display VTE Log Status

To display the status of VTE log service, type:

```
# secfsd -status logger
```

System Response:

```
Upload URL:
https://vmSSA06:8444/upload/logupload,https://vmSSA07:8444/upload/l
ogupload,https://vmSSA05:8444/upload/logupload

Logger Certificate directory:
/opt/vormetric/DataSecurityExpert/agent/vmd/pem
```

This command sequence returns the URL to which the log service sends log data. It also returns the directory that contains the VTE certificate. VTE uses the certificate to authenticate VTE when it uploads the log data to the DSM.

Display Applied Policies

To display the policies that are applied to VTE, type:

```
# secfsd -status policy
```

System Response:

```
Policy: allowAllOps_fs

Type: regular

Policy: allowAllRootUsers_fs

Type: regular

Policy: allowAllOps-winusers1_fs

Type: regular
```

Display VTE processes

To display VTE processes, type:

```
# secfsd -status pslist
```

System Response:

```
Protected pid list:     739   731
```

Display Detail about VTE processes

The example displays the process PID numbers for the `vmd` and `secfsd` processes. The `ps` commands show the processes for those PIDs.

```
# ps -fp <process #>
```

Example

```
# ps -fp 739
```

System Response:

```
      UID        PID     PPID   C    STIME    TTY  TIME CMD
    root  7012404        1   0 11:04:56      -  0:00
/opt/vormetric/DataSecurityExpert/agent/vmd/bin/vmd
```

Display VTE Version Information

To display VTE version information, type:

```
# secfsd -version
```

System Response:

```
version: vee-fs-6.1.0-110-rh6-x86_64.bin
```

Manually Enable a GuardPoint

To manually enable a GuardPoint on a Linux host:

1. Click **Hosts > Hosts > *<hostName>*** > **Guard FS.**
2. Click **Guard**.
3. In the Policy field, select a policy.
4. Set Type to **Directory (Manual Guard)**.
5. Click **Browse** and enter the GuardPoint path.

6. Click **OK**.
7. Log onto the system hosting VTE as the root user.
8. Verify the change, type:

```
# secfsd -status guard
```

System Response:

```
GuardPoint    Policy         Type   ConfigState Status     Reason
----------    ------         ----   ----------- ------     ------
/opt/apps/etc allowAllOps_fs manual unguarded  not guarded Inactive
```

Verifying a GuardPath

Verify that a GuardPath is guarded, type:

```
# secfsd -guard <path>
```

For example:

```
# secfsd -guard /opt/apps/etc
```

System Response:

```
secfsd: Path is Guarded
```

secfsd and raw devices

VTE for Linux only creates block devices.

To display them, type:

```
# ls -l /dev/secvm/dev
```

System Response:

```
brw-------    1 root     system       38,  1 Jan 29 16:37 hdisk1
brw-------    1 root     system       38,  2 Jan 29 16:37 hdisk2
crw-------    1 root     system       38,  3 Jan 29 16:37 rhdisk1
crw-------    1 root     system       38,  4 Jan 29 16:37 rhdisk2
```

vmsec utility

The vmsec utility allows you to manage security aspects of VTE on the host. On Linux-hosts**,** the `vmsec` utility is located in:

```
/opt/vormetric/DataSecurityExpert/agent/vmd/bin/vmsec
```

vmsec syntax

Table 18: vmsec Syntax [options]

`checkinstall`	Show vmd kernel status
`challenge`	Enter the dynamic host password
`vmdconfig`	Display the vmd configuration
`check_hwenc`	Display kernel configuration
`hwok`	Report status of hardware signature
`passwd [-p <password>]`	Enter the static host password
`version`	Display VTE version

Examples

Display VTE Challenge String

To display a VTE password challenge string and enter the response string when the DSM is not network accessible, type:

```
# vmsec challenge
```

System Response:

```
Contact the help desk at 1-800-555-1212 for response generation.
Your host name is "Host120" Your challenge is: HPTQ-ZYLK
Response -> IHFY-W7WG-PDAO-QKKQ
```

The contact information is configured in the DSM Management Console (Domains > Manage Domains) *Add Domain* window. Contact the DSM administrator and give them the challenge string. The DSM administrator will give you the response string. Enter the response string in the **Response** field and press **Enter**. You have 15 minutes to enter the response string.

The Windows equivalent of this command is right-click the tray icon and select **Challenge...-> Response**.... The *VTE Challenge/Response* window opens. Note the Challenge string displayed. If no string is displayed, the host password is static. If a challenge string displays, contact a DSM administrator for the response string.

Display VTE Status

This utility shows you if VTE is configured and running. If it is not running, you might need to start it manually.

To display VTE status, type:

```
# vmsec checkinstall
```

System Response:

```
The kernel component is installed and running.
```

Entering a Password

To enter VTE static host password, type:

```
# vmsec passwd
```

System Response:

```
Please enter password:
OK passwd
```

To enter VTE static host password on the command line so you can specify it in a batch script, type:

```
# vmsec passwd -p myPass123
```

System Response:

```
OK passwd
```

Display Kernel Status

To display the kernel status, type:

```
# vmsec status
```

System Response:

```
FILE_FORMAT=2
FILE_GENERATED=10/27/2017 18:54:10
SA_QOS_STATUS=0
SA_HOST_CPU_UTIL=0
GP_1_Policy=27
GP_1_Dir=/gp
GP_1_lock=1
GP_1_type=1
GP_1_gtype=manual
```

```
GP_1_opt=gtype=2,policy=27,lock=1,type=1,dir=/gp/
GP_1_config_state=unguarded
GP_1_status=not guarded
GP_1_statuschk_tm=0-00-00 00:00:00
GP_1_config_op_retry_cnt=0
GP_1_config_op_attempt_tm=0-00-00 00:00:00
GP_1_flags=0
GP_1_reason=Inactive
GP_1_usage=free
TOTAL_GP=1
KEYS_AVAILABLE=TRUE
sdk_version=6.1.0.73
sdk_builddate=2017-10-17 15:16:46 (PDT)
coreguard_locked=false
system_locked=false
logger_upload_url=https://thl602-
2114.qa.com:8447/upload/logupload,https://thl602-
2116.qa.com:8447/upload/logupload
logger_cert_dir=/opt/vormetric/DataSecurityExpert/agent/vmd/pem
hostname_for_logging=vmd
QOS_PAUSED=false
vmd_STRONG_ENTROPY=false
vmd_URL=https://thl602-2114.qa.com:8446
vmd_SRV_URLS=https://thl602-2114.qa.com:8446, https://thl602-
2116.qa.com:8446
vmd_PRIMARY_URL=https://thl602-2114.qa.com:8446
vmd_SUPPORTS_F8P=TRUE
vmd_SUPPORTS_CR256=TRUE
vmd_RANDHP=TRUE
learn_mode=false
concise_logging=false
vmd_listening_port=7024
vmd_initialization_time=2017-10-25 12:07:14.514
vmd_last_server_update_time=2017-10-25 12:12:04.747
policy_name_27=aes256
policy_version_27=0
policy_keyvers_27=0
policy_type_27=ONLINE
policies=27
logger_suppression_VMD=SUPPRESS
logger_intervaltime_VMD=600
logger_repeat_max_VMD=5
logger_suppression_POL=SUPPRESS
logger_intervaltime_POL=600
logger_repeat_max_POL=5
CONFIG_SA_1=27
```

```
TOTAL_CONFIG_SA=1
SA_1_NAME=27
SA_1_ALIAS=aes256
SA_1_TYPE=0
SA_1_REF=1
SA_1_HIP_REG_TIME=0
SA_1_FLAGS=1
TOTAL_SA=1
TOTAL_AUTH=0
AUTHBIN_1=|authenticator|/usr/sbin/sshd
B92A3D7EEF67B82230F7F76097D65159FCF5722A4154A249EFDC22C20F1B572C
AUTHBIN_2=|authenticator|/bin/login
4F210D1B83ACD79B006BCF7DB247ED002A45FC892C42720390BFA6AE21AEA8DC
TOTAL_AUTHBIN=2
```

Display VTE Build Information

For Linux/, type:

```
# vmsec version
version 5, Service Pack 2
2018-12-17 20:41:51 ()
```

```
Copyright (c) 2009-2018, Vormetric.  All rights reserved.
```

System Response (Linux):

```
Version 6

6.0.2.49

2017-10-17 15:15:23 (PDT)

Copyright (c) 2009-2018, Thales. All rights reserved.
```

Display Contents of Conf files

To display the contents of the `agent.conf` and `.agent.conf.defaults` files, type:

```
# vmsec vmdconfig
```

System Response:

```
appender_syslogdest_Syslog_Appender_0=127.0.0.1

VMSDK_AGENT_CONFIG_FILE=/opt/vormetric/DataSecurityExpert/agent/vmd
/etc/agent.conf

appender_layout_Syslog_Appender_0=Syslog_Layout

VMSDK_AGENT_VERSION=5.2.6.0

VMSDK_AGENT_BUILD_ID=28

PREV_URLS=https://srv.my.vormetric.com:8443

syslog_appender_myhost name=dev.my.vormetric.com

VMD_PORT=7024

...

...

appenders=Upload_Appender, File_Appender, Syslog_Appender_0

layouts=Upload_Layout, File_Layout, Syslog_Layout, Simple

CONNECT_TIMEOUT=180000

URL=https://srv.my.vormetric.com:8443
```

```
STRONG_ENTROPY=false
```

Binary Resigning

Any executable that is part of either a Host Setting or Signature set, and resides in a GuardPoint that uses an LDT policy, will use different signatures for an LDT key rotation. The result is that the Host Settings binaries will no longer be authenticated, or the Signature Set policy rules will no longer trigger for those binaries. To prevent these issues, the Security Administrator must manually resign each affected binary after each key rotation.

As of VTE release 6.1.2, binaries are now signed with a signature that does not change with a key rotation. The Security Administrator will need to do one manual resigning. After that, there is no longer a need to resign after each key rotation.

If upgrading or installing a new machine using the same signature sets that you used previously, do the following:

1. Install release 6.1.2 of the VTE agent (which contains the ability to generate unencrypted signatures of binaries inside GuardPoints). The previous signatures will be used until the next key rotation.
2. Before the next key rotation, the Security Administrator must resign the binaries.
3. Do not remove the old signatures on the DSM until all agents have been upgraded to VTE release 6.1.2 (which has the ability to generate unencrypted signatures on binaries inside GuardPoints). Refer to the DSM Installation and Configuration Guide for information on how to do a manual resign.
4. After all agents have been upgraded, then you can remove the old signatures.

If you are installing the VTE agents for the first time, there are no special steps, if no signatures have been defined. The agent will sign using the new method.

NOTE: In previous releases, if the binary was in a GuardPoint protected directory, but was the same as an unguarded binary, the Administrator could restrict to only the guarded binary. With this change, the unguarded binary is now unrestricted. This means that if a user uses the unguarded binary and its SHA matches the guarded binary, it will now match as if it was the guarded binary.

Enable Automatic Signing for Host Settings

A new feature of VTE blocks automatic re-signing of the host settings. Some users may have established procedures for updating system software. The user created these procedures based on their assumption that restarting the `vmd` will generate new signatures when signed software is updated. This is no longer true. To restore this behavior for updating system software, you must disable this new feature.

Disabling on Linux

1. Change to the directory where the `agent.conf` file resides. For example, type:
   `# cd /opt/vormetric/DataSecurityExpert/agent/vmd/etc/`
2. Edit the `agent.conf` file.
3. Change or add the following line:
   `RE_SIGN_HOST_SETTINGS=TRUE`
4. Save your changes and exit the file.
5. Restart the vmd to set the changes, type:**`# /etc/vormetric/secfs restart`**
6. Type the following to verify that the Host settings is set to true:
   `# vmsec vmdconfig`

Warning! Enabling the automatic regeneration of signatures exposes a potential security vulnerability for agents. When enabled, host setting binaries are resigned when it receives a push from the DSM. If an attacker were to replace a binary with a Trojan, and then force a push from the DSM by, for example, restarting the agent, VTE could generate a signature for the malicious binary and pass it to the kernel.

Restricting access overrides from unauthorized identities

In some setups, system administrators can use the host settings > |authenticator| feature with `su` to change identities and gain access to restricted data. Now, you can instruct VTE to not trust any authentication attempt performed by certain

identities by assigning restricted users to a user shell that VTE can block from authenticating other processes.

Any executable path that is marked with a |path_no_trust| host setting marks the process, and all child processes, as not trusted. Non-trusted processes are treated as "User Not Authenticated" to prevent access on user-based policies.

VTE prevents overrides from other host settings authenticators, using the |path_no_trust| status. If a user runs the su command from a non-trusted shell, that new shell is still marked as |path_no_trust|, even if |authenticator|/usr/bin/su is specified in the host-settings. The |path_no_trust| feature overrides any and all authenticators under host settings.To restrict access overrides:

1. At the management console, click **Hosts > Hosts**.
2. Click on an **existing** Host name to edit the host.
3. Click **Host Settings** tab.
4. Add the following to the host settings:

   |path_no_trust|<*path of the binary*>

   Example

   |path_no_trust|/bin/ksh

   The above example indicates that no process under the kshell executable will be authenticated.
5. Click **OK**.

Using Advanced Encryption Set New Instructions (AES-NI)

Determine AES-NI Hardware Support

To verify AES-NI hardware support, type:

```
# vmsec check_hwenc
```

Unlike the -c algo command above, VTE does not have to be running for this command to execute. This command displays one of the following messages to stdout:

"AES-NI hardware encryption is supported on this system."

"AES-NI hardware encryption is not supported on this system. Will default to software encryption."

vmd utility

The `vmd` utility displays VTE software version information.

The `vmd` utility is located in `/opt/vormetric/DataSecurityExpert/agent/vmd/bin` and a symbolic link to this file is placed in `/usr/bin/vmd`.

Syntax

```
vmd [OPTIONS...]
```

`-h`	show utility syntax
`-v`	display VTE version
`-f`	runs `vmd` in the foreground

Display the Installed Version

To display the installed VTE version, type:

```
# vmd -v
```

System Response:

```
Version 6
6.0.2.495.2.7.8
2018-12-17 20:41:51 ()
Copyright (c) 2009-2018, Vormetric.  All rights reserved.
```

agenthealth utility

The `agenthealth` utility validates:

- Super-user privilege
- VTE Agent installation
- VTE registration to DSM Server
- VTE processes/modules that are running
- Available disk resources

- Current GuardPoints

  Tests if the agent can reach the GuardPoints

- VTE log directory resource status

  This directory contains pending VTE log files for upload. This utility reports the size and number of pending files for upload. These text files are logs that contain vmd/SecFS information. They are regenerated whenever secfs restarts. If the number of files is unexpectedly large, this can indicate a problem.

The Agent health check script

To run the `Agenthealth` check script, type:

`/opt/vormetric/DataSecurityExpert/agent/vmd/bin/agenthealth`

System Response:

```
Checking for super-user privilege ..............  OK
Vormetric Agent installation ...................  OK
Vormetric policy directory .....................  OK
Registration to server .........................  OK
Kernel modules are loaded ......................  OK
VMD is running .................................  OK
SECFSD is running ..............................  OK
dsm602-33-101.qa.com is resolvable .............  OK
dsm602-33-101.qa.com port 8446 is reachable ....  OK
dsm602-33-101.qa.com port 8447 is reachable ....  OK
dsm602-63-183.qa.com is resolvable .............  OK
dsm602-63-183.qa.com port 8446 is reachable ....  OK
dsm602-63-183.qa.com port 8447 is reachable ....  OK
Can communicate to at least one server..........  OK
VMD is listening on port 7024...................  OK
Time of last update from server.................  2018-02-13
20:25:37.446
Checking available disk space...................  OK
Checking logging space .........................  OK
    Log directory is "/var/log/vormetric"
```

```
    File system for log data is "/", 32G free (17% full)
    Log directory contains 2 of maximum 200 files (1% full)
    Log directory contains 1 of maximum 100 Mbytes used (1% full)
Testing access to /ofx-fs1 ..................... OK
Testing access to /gp1 ........................... Access denied as
per policy
```

agentinfo utility (Java version)

The `agentinfo` utility collects system and VTE configuration data. The `agentinfo` utility is used to take a configuration snapshot of the system that you will send to Thales Customer Support to debug an issue**.** (This section describes the Java version.)

The `agentinfo` utility is a JavaScript file. You can open it in a text editor to see specific functions.

The `agentinfo` utility displays status information on the screen and outputs the results to a compressed tar file. The compressed tar file name format is `ai.<os_name_ver>.qa.com.tar.gz` and it is located in the current working directory.

To create an agentinfo file, type:

```
# /opt/vormetric/DataSecurityExpert/agent/vmd/bin/agentinfo
```

NOTE: The Java version is supported on VTE/Linux and VTE/AIX.

check_host utility

If a VTE software installation fails during the certificate generation and exchange stage, use the `check_host` utility to list the network addresses for the host. The utility checks network interfaces, `/etc/hosts`, DNS, and so on, to compare, test, and evaluate possible addresses for the host, and weights them based upon their network efficiency. FQDNs are the most preferred and stand-alone IP addresses are the least preferred. Some applications, such as silent-mode installation, use

`check_host` to determine the best host address to submit to the DSM during registration.

Run the `check_host` utility on a system that is hosting VTE to display available network host names, FQDNs, and IP numbers for the host.

Type:

```
# /opt/vormetric/DataSecurityExpert/agent/vmd/bin/check_host
```

check_host Syntax

```
check_host [-h | -i | -a] [-s name] [-l name:port]
```

-h	Print the best host name for this machine
-i	Print the best IP address
-a	Print all the host names and IP addresses
-s	The name of the server (optional hint)
-r	The name of the server for name resolution checks
-l	The name and port of the server for listening checks

register_host utility

Use the `register_host` utility to create certificate requests, exchange certificates between the DSM and the host, and to register VTE on the DSM. After the host is registered, you can configure VTE, apply GuardPoints, or perform database backups. Run the Linux register_host utility in text mode on a terminal window.

Caution: 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 might also have to extend the default TCP timeout.

fs_freeze and xfs_freeze (Linux Only)

Users can freeze, snapshot, and unfreeze a file system with an SecFS GuardPoint using `fs_freeze|xfs_freeze` for both XFS and EXT3/4.

SecFS supports freezing with `fs_freeze|xfs_freeze` or by any other program issuing the same type of requests. Freezing SecFS results in freezing the underlying file system, as well as the primary file system.

Restrictions

There are restrictions for using `fs_freeze|xfs_freeze` support with VTE.

Platform Restrictions

The following platform restrictions occur with VTE and `fs_freeze|xfs_freeze`:

- VTE supports the `fs_freeze|xfs_freeze` utility for freezing SECFS GuardPoints on all Linux distributions for kernels >=3.0 for Redhat, SLES, and Ubuntu platforms on EXT3/EXT4/XFS file systems. (Earlier Kernels do not contain the proper freeze_super VFS code).

NOTE: The existing behavior of using `fs_freeze|xfs_freeze` in the underlying XFS file system works for XFS only. We do not support `fs_freeze|xfs_freeze` on GuardPoints located on VxFS file systems.

`fs_freeze|xfs_freeze` is supported with Redhat 6 (2.6 kernel) only with XFS file system.

Target Restrictions

The expected target of the `fs_freeze|xfs_freeze` command is the path of the GuardPoint.

For example, if `/dev/sdb` is mounted as ext4 on `/data` and VTE contains the GuardPoint: `/data/protected`, then the target of `fsfreeze` must be `/data/protected`, not `/data`.

Valid: `# fsfreeze -f /data/protected`

Not valid: `# fsfreeze -f /data`

File System Restrictions

The following file system restrictions occur with VTE and `fs_freeze|xfs_freeze`:

- If multiple GuardPoints exist on the same file system, you only need to freeze one

  For example, if `/dev/sdb` is mounted as ext4 on `/data` and the VTE GuardPoints are `/data/protected1` and `/data/protected2`, then issuing:

  ```
  # fsfreeze -f /data/protected1
  ```

  freezes `/data/protected1`, `/data/protected2` and the underlying ext4 file system.

Caution: Do not unguard a GuardPoint, or restart the VTE agent, while the file system is frozen. The only action permitted on a frozen file system is taking a snapshot or backing up.

- If you try to freeze `/data/protected2` after freezing `/data/protected1`, it returns as busy
- If you are not permitted to freeze one GuardPoint, then you cannot freeze any GuardPoints

LDT Restrictions

- You cannot freeze a file system while it is undergoing an LDT rekey operation. If it detects a rekey, the freeze returns as busy
- You cannot start an LDT rekey on a frozen file system

Offline Data Transformation Restrictions

Do **NOT** use `fs_freeze|xfs_freeze` while an offline transform policy is in effect

16 Support for Systemd

Vormetric Transparent Encryption (VTE) for Linux supports `systemd`. Specifically, Vormetric VTE services are integrated with the systemd framework. A few minor modifications are required.

This chapter contains the following sections:

Overview

Systemd is a system and service manager for Linux designed to replace Linux init which was inherited from the older UNIX System v-init. It consists of a collection of daemons, libraries, and utilities to provide central management and configuration for Linux. While not all Linux systems have implemented systemd, many have.

Specifically, some versions of Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES) and Ubuntu Server LTS use systemd as the default init system and are supported by VTE (see Table 19 on page 212 for a complete list of supported operating systems). In addition to providing enhanced features and performance, systemd is designed to be backwards compatible with SysV and Linux Standard Base init scripts.

- The systemd hosts that are installing on VTE contain some required procedures to fully integrate with VTE.The start, stop, restart, and VTE status have different command locations.
- For each application that uses GuardPoints:
  - Modify the application's system unit configuration file to manage software dependencies.
  - Add your applications to the file `secfs-fs-barrier.service` and run `systemctl daemon-reload`.

Table 19: Systemd Supported Operating Systems

systemd supported Operating Systems
REHL 7.x, SLES 12, Ubuntu 16.04

VTE Administration

This section describes the commands for start, stop, restart, and status of VTE.

Starting VTE

To start VTE, type:

```
# /etc/vormetric/secfs start
```

Stopping VTE

To stop VTE, type:

```
# /etc/vormetric/secfs stop
```

Restart VTE

To restart the VTE, type:

```
# /etc/vormetric/secfs restart
```

Check the Status of VTE

When the VTE services reside on a normal system, their corresponding states are as follows:

- `secfs-init`: active (exited) state
- `secfsd`: active (running) state
- `vmd`: active (running) state
- `secfs-fs`: active (exited) state

To check the status of VTE, type:

```
# /etc/vormetric/secfs status
```

System Response

```
secfs-init service: active (exited) since Tue 2017-09-26 09:04:21
PDT; 1 day 4h ago

secfsd service : active (running) since Tue 2017-09-26 09:04:21
PDT; 1 day 4h ago

vmd service : active (running) since Tue 2017-09-26 09:04:44 PDT; 1
day 4h ago

secfs-fs service : active (exited) since Tue 2017-09-26 09:04:45
PDT; 1 day 4h ago
```

Systemd Software Dependency File

Systemd manages software dependencies by specifying those dependencies in a unit configuration file. An example of a software dependency is a program, `mongodb`, that requires VTE to be running before `mongodb` is started.

Applications that use GuardPoints have a software dependency. You must add them to the `secfs-fs-barrier.service` file. Typically, applications that use GuardPoints run as services and run as long as the operating system is active. Examples include postgres, httpd, mongodb, mysqld and mariadb.

Before you can safely reboot your protected host, or use the files in the GuardPoint, you must perform the steps below.

Setting up your Systemd software dependency files

Prerequisite

VTE and the protected files must be on the same host.

For RedHat and SUSE Linux, the unit configuration file is installed by the application's installer in the directory `/usr/lib/systemd/system/`. On Ubuntu 16.04 its located at `/lib/systemd/system/mariadb.service`.

It has the format `<application_name>.service`.

1. Compile a list of your applications that use GuardPoints.
2. Make sure that none of those applications are running.
3. For each application, open the unit configuration file for that application, type:

RH and Suse:

```
# vi /usr/lib/systemd/system/mariadb.service
```

Ubuntu:

```
# vi /lib/systemd/system/mariadb.service
```

4. Edit the file and add `secfs-fs-barrier.service` to the `Requires=` clause, type:

   `Requires=secfs-fs-barrier.service`
5. Open the file `secfs-fs-barrier.service` for editing.
6. Add the name of the unit configuration file to the end in the "`Before=`" clause.
7. Type: `systemctl daemon-reload`
8. Start application.

NOTE: You must follow the above procedure for each application that uses Vormetric GuardPoints.

Example for modifying the unit configuration file

This example uses mariadb as the dependent application.

```
# cat /usr/lib/systemd/system/mariadb.service

# It's not recommended to modify this file in-place, because it
will be overwritten
# during package upgrades.  If you want to customize, the best way
is to create a file
# "/etc/systemd/system/mariadb.service", containing
#       .include /lib/systemd/system/mariadb.service
#       ...make your changes here...
# or create a file
"/etc/systemd/system/mariadb.service.d/foo.conf", which doesn't
# need to include ".include" call and which will be parsed
# after the file mariadb.service itself is parsed.
#
# For more info about custom unit files, see systemd.unit(5) or
#
http://fedoraproject.org/wiki/Systemd#How_do_I_customize_a_unit_file.2F_add_a_custom_unit_file.3F

# For example, if you want to increase mariadb's open-files-limit
to 10000,
```

```
# you need to increase systemd's LimitNOFILE setting, so create a
file named
# "/etc/systemd/system/mariadb.service.d/limits.conf" containing:
#       [Service]
#       LimitNOFILE=10000

# Note: /usr/lib/... is recommended in the .include line though
/lib/... still works.
# Don't forget to reload systemd daemon after you change unit
configuration:
# root> systemctl --system daemon-reload

[Unit]
Description=MariaDB database server
After=syslog.target
After=network.target
Requires=secfs-fs-barrier.service

[Service]
Type=simple
User=mysql
Group=mysql

ExecStartPre=/usr/libexec/mariadb-prepare-db-dir %n
# Note: we set --basedir to prevent probes that might trigger
SELinux alarms,
# per bug #547485
ExecStart=/usr/bin/mysqld_safe --basedir=/usr
ExecStartPost=/usr/libexec/mariadb-wait-ready $MAINPID

# Give a reasonable amount of time for the server to start up/shut
down
TimeoutSec=300

# Place temp files in a secure directory, not /tmp
PrivateTmp=true

[Install]
WantedBy=multi-user.target
```

17 Ubuntu Upstart Service Support

Vormetric Transparent Encryption (VTE) supports the Ubuntu upstart service framework. VTE services ensure that dependent Upstart application services start only after the VTE process starts. In addition, applications started by traditional scripts can also be synchronized to start after VTE processes have started.

These changes are specific to Ubuntu 14.04, or earlier versions. They do not apply to Ubuntu 16.04, Red Hat and SLES. This chapter contains the following sections:

Administering Vormetric Services

This section describes the startup and shutdown of the VTE services.

Starting the VTE services

Run the following commands to start the VTE services in the order shown:

```
# start secfs-init

# start secfs-fs
```

Stopping the VTE services

Run the following commands to stop the VTE services in the order shown:

```
# stop secfs-fs

# stop secfs-init
```

You cannot stop `secfs-fs` if guarded directories associated with mysql or other third-party applications are in use. Stop the application before stopping the VTE process. The Upstart framework on Ubuntu does not display error messages. Check for error messages in the following location:

`/var/log/upstart/secfs-fs.log`

Querying VTE status

Use the following commands to obtain the VTE status:

```
# status secfs-fs
```

System Response:

```
secfs-fs start/running, process 1501
```

```
# status secfs-init
```

System Response:

```
secfs-init start/running
```

Vormetric Upstart service management logs

You can find upstart services management log in the following locations:

`/var/log/upstart/secfs-init.log`

`/var/log/upstart/secfs-fs.log`

Upgrading VTE

There are no restrictions to upgrade VTE.

Administering Third-party Services

VTE ensures that GuardPoints are available before mysql and all other dependent services are guarded.

Guarding mysql folders (mysql already installed)

- Ensure that the latest patch is installed
- Confirm that secfs-init and secfs-fs Vormetric services are running
- Create and guard mysql folders
- Install and start mysql

Adding new Upstart dependencies

VTE supports the following Upstart services on the following applications:

- mysql
- mongodb
- apache2
- postgres

If your application is not listed above, then modify the `secfs-fs-barrier.conf` as follows:

At the end of the `start on` section located at:

`/etc/init/secfs-fs-barrier.conf`

Add the following:

`or starting <your_service_name>`

For example, for service foo:

```
start on starting mysql or starting mongodb or starting
apache2 or starting postgres or starting foo
```

Configuring rc/sysvinit services

You can synchronize applications started by traditional `sysvinit` scripts to start after VTE starts.

Enabling the barrier for rc services

Run the following command after guarding data to insure all rc (sysvinit) based services start after the VTE starts, type:

```
# update-rc.d secfs-fs-barrier defaults 00 99
```

Disabling the barrier for rc services

Run the following command after unguarding data to allow independent starting of rc (sysvinit)-based services, type:

```
# update-rc.d -f secfs-fs-barrier remove
```

Troubleshooting and Best Practices

A

Windows Systems

VTE will not register with the DSM

- If there is a firewall between the DSM and VTE, configure `vmd.exe` as a firewall exception on VTE for Windows. Otherwise, the DSM is unable to browse VTE.
- If using a Windows XP or Windows 2003 system Firewall, select **Control Panel > Windows Firewall > Exceptions > Add Program...** and browse for `vmd.exe`. The default location is

```
\Program Files\Vormetric\DataSecurityExpert\agent\vmd\bin\vmd.exe
```

- If using a Windows 7 system Firewall, select **Control Panel > System and Security > Windows Firewall > Allowed Programs...**click **Change settings**, click **Allow another program** and browse for `vmd.exe`. The default location is

```
\Program Files\Vormetric\DataSecurityExpert\agent\vmd\bin\vmd.exe
```

Veritas VxFS Environment

On a VxFS file system that contains a large number of files guarded by VTE, when those files are read by a backup process or security scan, VTE does not fully close the file, which eventually leads to kernel errors similar to the following example:

```
Apr 5 18:33:57 drpbvwebprd11 kernel: [7547.903965] vxfs: msgcnt 1 mesg 014: V-
2-14: vx_iget - inode table overflow
```

Workaround — Increase the vxfs max inode count

Increase the `vxfs max inode` count by tuning the `vxfs_ninode` to a large enough value such that the kernel's reuse code is triggered and starts purging unused entries.

For the following platforms:

- RHEL 6.x: in `/etc/modprobe.d/dist.conf`
- SLES 11: in `/etc/modprobe.conf.local`
- SLES 12: in `/etc/modprobe.d/99-local.conf`

- RHEL 7.x: Create a `.conf` file in `/etc/modprobe.d/`

In each of the above listed files, add the following line and reboot:

```
options vxfs vxfs_ninode=<higher_number>
```

This line does not take effect until the VxFS module is reloaded or you do a system reboot.

GLOSSARY

B

access control
The ability of Vormetric Transparent Encryption (VTE) to control access to data on protected hosts. Access can be limited by user, process (executable), action (for example read, write, rename, and so on), and time period. Access limitations can be applied to files, directories, or entire disks.

admin administrator
The default DSM administrator created when you install the DSM. Admin has DSM System Administrator privileges and cannot be deleted.

Administrative Domain
(domains). A protected host or group of protected hosts on which an DSM administrator can perform security tasks such as setting policies. Only DSM administrators assigned to a domain can perform security tasks on the protected hosts in that domain. The type of VTE tasks that can be performed depends on the type of administrator. See also **"local domain"**.

administrator
See **"DSM Administrator and types"**.

Agent utilities
A set of utilities installed with the VTE agents and run on protected hosts. These utilities provide a variety of useful functions such as gathering protected host and agent configuration data, registering agents on the DSM, and encrypting data on the protected host.

All Administrator, **Administrator of type All**
The DSM Administrator with the privileges of all three administrator types: *System*, *Domain* and *Security*.

appliance
The DSM server. Often referred to as a *DSM hardware appliance*, which is a hardened DSM server provided by Vormetric, or as a *DSM virtual appliance*, which is the software version of the DSM to be deployed by the customers as a virtual machine.

asymmetric key cryptography
See *public key cryptographic algorithm.*

asymmetric key pair
A public key and its corresponding private key used with a public key algorithm. Also called a key pair.

authentication
A process that establishes the origin of information, or determines the legitimacy of an entity's identity.

authorization
Access privileges granted to an entity that convey an “official” sanction to perform a security function or activity.

block devices
Devices that move data in and out by buffering in the form of blocks for each input/output operation.

catch-all rule
The last policy rule that applies to any GuardPoint access attempt that did not fit any of the other rules in the policy.

certification authority or CA
A trusted third party that issues digital certificates that allow a person, computer, or organization to exchange information over the Internet using the public key infrastructure. A digital certificate provides identifying information, cannot be forged, and can be verified because it was issued by an official trusted agency. The certificate contains the name of the certificate holder, a serial number, expiration dates, a copy of the certificate holder's public key (used for encrypting messages and digital signatures) and the digital signature of the certificate-issuing authority (CA) so that a recipient can verify that the certificate is real. This allows others to rely upon signatures or assertions made by the private key that corresponds to the public key that is certified. The CA must be trusted by both the owner of the certificate and the party relying upon the certificate.

challenge-response
When a protected host is disconnected from the DSM, the GuardPoint data is not accessible to users. Challenge-response is a password-based procedure that allows users to gain access to their GuardPoint data during disconnection. Users run a utility, `vmsec challenge`, a seemingly random string (the challenge) is displayed. The user calls this in to their DSM Security administrator. The administrator returns a counter-string (the response) that the host user must enter to decrypt guarded data.

Character device
See ***“raw device.”***

ciphertext
Data in its encrypted form. Ciphertext is the result of encryption performed on plaintext using an algorithm, called a cipher.

cleartext or plaintext
Data in its unencrypted form.

cryptographic algorithm
A computational procedure that takes variable inputs, including a cryptographic key, and produces ciphertext output. Also called a cipher. Examples of cryptographic algorithms include AES, ARIA, and DES.

cryptographic key
See "**encryption key**."

cryptographic signature
See "**signing files**."

Database Encryption Key (DEK)
A key generated by Microsoft SQL when TDE is enabled.

Data Security Manager (DSM)
Sometimes called the *Security Server* or *appliance*. A Vormetric server that acts as the central repository and manager of encryption keys and security policies. Receives instructions and configuration from administrators through a GUI-based interface called the *Management Console*. Passes and receives information to and from VTE Agents. Available as a complete hardened hardware system (*DSM Appliance*) or as software solution installed on a UNIX box (*software-only DSM*).

dataxform
A utility to encrypt data in a directory. Short for "data transform."

DB2
A relational model database server developed by IBM.

Decryption
The process of changing ciphertext into plaintext using a cryptographic algorithm and key.

Digital signature
A cryptographic transformation of data that provides the services of origin authentication, data integrity, and signer non-repudiation.

domains
See *administrative domains*.

Domain Administrator
The second-level DSM administrator created by a DSM *System Administrator*. The DSM *Domain Administrator* creates and assigns DSM *Security Administrators* to domains and assigns them their security "**roles**". See "**DSM Administrator and types**".

Domain and Security Administrator
A hybrid DSM administrator who is has the privileges of a DSM Domain Administrator and Security Administrator.

DSM
See *"**Data Security Manager (DSM)."***

DSM Administrator and types
Specialized system security administrators who can access the Vormetric DSM Management Console. There are five types of DSM administrators:

DSM System Administrator - Creates/removes other DSM administrators of any type, changes their passwords, creates/removes, domains, assigns a Domain Administrator to each domain. Cannot do any security procedures in any domain.

Domain Administrator - Adds/removes DSM Security Administrators to domains, and assign roles to each one. Cannot remove domains and cannot do any of the domain security roles.

Security Administrator - Performs the data protection work specified by their roles. Different roles enable them to create policies, configure hosts, audit data usage patterns, apply GuardPoints, and so on.

Domain and Security Administrator - Can do the tasks of DSM Domain and Security Administrators.

All - Can do the tasks of all three of the DSM administrative types

DSM Automation Utilities
Also called VMSSC. A set of command line utilities that is downloaded and installed separately on the protected host or any networked machine. These utilities can be used by advanced users to automate DSM processes that would normally be done with the Management Console. See the *DSM Automation Reference* for complete details.

DSM CLI
A command line interface executed on the DSM to configure the DSM network and perform other system-level tasks. See the *DSM Command Line Interface* documentation

DSM CLI Administrator
A user who can access the DSM CLI. DSM CLI Administrators are actual system users with real UNIX login accounts. They perform tasks to setup and operate the DSM installation. They do not have access to the Management Console.

DSM database
A database associated with the DMS containing the names of protected hosts, policies, GuardPoints, settings, and so on.

DSM System Administrator
The highest level of DSM administrator. This administrator creates/removes other DSM administrators of any type, creates/removes domains, and assigns a Domain Administrator to each domain. The DSM System Administrator cannot perform any security procedures in any domain or system. This administrator is not related to computer or network system administrators.

EKM
See "**Extensible Key Management (EKM)**."

Encryption
The process of changing plaintext into ciphertext using a cryptographic algorithm and key.

encryption agent
See *Vormetric Transparent Encryption agent*.

encryption key
A piece of information used in conjunction with a cryptographic algorithm that transforms plaintext into ciphertext, or vice versa during decryption. Can also be used to encrypt digital signatures or encryption keys themselves. An entity with knowledge of the key can reproduce or reverse the operation, while an entity without knowledge of the key cannot. Any VDS policy that encrypts GuardPoint data requires an encryption key.

Extensible Key Management (EKM)
An API library specification provided by Microsoft that defines a software framework that allows hardware security module (HSM) providers to integrate their product with the Microsoft SQL Server.

failover DSM
A secondary DSM that assumes the policy and key management load when a protected host cannot connect to the primary DSM or when a protected host is specifically assigned to the failover DSM. A failover DSM is almost identical to the primary DSM, having the same keys, policies, protected hosts, and so on.

FF1
See "Format Preserving Encryption (FPE)".

FF3
See "Format Preserving Encryption (FPE)".

file signing
See *signing files*.

File Key Encryption Key (FKEK)
The key used to encrypt the file encryption key that is used to encrypt on-disk data, also known as a wrapper key.

FKEK
See "File Key Encryption Key (FKEK)"

File System Agent
A Vormetric software agent that resides on a host machine and allows administrators to control encryption of, and access to, the files, directories and executables on that host system. For example, administrators can restrict access to specific files and directories to specific users at specific times using specific executables. Files and directories can be fully encrypted, while the file metadata (for example, the file names) remain in cleartext. Also called the "**VTE Agent**".

Format Preserving Encryption (FPE)
An encryption algorithm that preserves both the formatting and length of the data being encrypted. Examples of such algorithms used by Vormetric include FF1 and FF3, both of which are approved by NIST. Vormetric's **FPE tokenization format** uses the FF3 algorithm.

FQDN
Fully qualified domain name. A domain name that specifies its exact location in the tree hierarchy of the Domain Name Server (DNS). For example: `example.vormetric.com`.

GPFS
General Parallel File System is a high-performance shared-disk clustered file system developed by IBM.

GuardPoint
A location in the file system hierarchy, usually a directory, where everything underneath has a Vormetric data protection policy applied to it. The File System Agent intercepts any attempt to access anything in the GuardPoint and uses policies obtained from the DSM to grant or deny the access attempt. Usually, depending on the policies, data copied into a GuardPoint is encrypted, and only authorized users can decrypt and use that GuardPoint data.

Hardware Security Module or HSM
A tamper-resistant hardware device that stores keys and provides stringent access control. It also provides a random number generator to generate keys. The DSM Appliance can come with an embedded Hardware Security Module.

host locks
Two Management Console options, **FS Agent Locked** and **System Locked,** that are used to protect the File System Agent and certain system files. File System Agent protection includes preventing some changes to the File System Agent installation directory and preventing the unauthorized termination of File System Agent processes.

host password
This is not a regular login or user password. This is the password entered by a host system user to unlock a GuardPoint when there is no DSM connection. This password decrypts cached keys when the DSM is not accessible. The host must also be configured with **Cached on Host** keys. See "**challenge-response**".

initial test policy
A first data security policy applied to a GuardPoint that is used to gather directory access information so DSM Security Administrators can create a permanent operational policy. The initial test policy encrypts all data written into the GuardPoint; decrypts GuardPoint data for any user who access it; audits and creates log messages for every GuardPoint access; reduces log message "noise" so you can analyze the messages that are important to you for tuning this policy; is run in the "**Learn Mode**" which does not actually deny user access, but allows you to record GuardPoint accesses.

After enough data is collected, the DSM Security Administrator can modify the initial test policy into an operational policy.

Key Agent
A Vormetric agent that provides an API library supporting a subset of the PKCS#11 standard for key management and cryptographic operations. It is required for the following products: Vormetric Key Management (VKM), Vormetric Tokenization, Vormetric Application Encryption (VAE), Vormetric Cloud Encryption Gateway (VCEG). Sometimes called the *VAE Agent*.

key group
A key group is a collection of asymmetric keys that are applied as a single unit to a policy.

key management
The management of cryptographic keys and other related security objects (for example, passwords) during their entire life cycle, including their generation, storage, establishment, entry and output, and destruction.

key template
A template that lets 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.

key shares
When data is backed up or exported from VTE (for example, symmetric keys or DSM database backups), they can be encrypted in a wrapper key needed to restore the exported data on the new machine. Wrapper keys can be split and distributed to multiple individuals. Each split piece of the wrapper key is called a *key share*. Decrypting the data requires that some specified number of the individuals that received key shares contribute their key share to decrypt the data.

key wrapping
A class of symmetric encryption algorithms designed to encapsulate (encrypt) cryptographic key material. The key wrap algorithms are intended for applications such as protecting keys while in untrusted storage or transmitting keys over untrusted communications networks. Wrapper keys can be broken up into *key shares*, which are pieces of a wrapper key. Key shares are divided amongst two or more *custodians* such that each custodian must contribute their key share in order to assemble a complete wrapper key.

Key Vault
A Vormetric product that provides passive key vaulting. It securely stores symmetric and asymmetric encryption keys from any application and tracks key expiration dates.

KMIP
Key Management Interoperability Protocol. A protocol for communication between enterprise key management systems and encryption systems. A KMIP-enabled device or client software can communicate with the DSM to manage encrypted keys.

Learn Mode
A DSM operational mode in which all actions that would have been denied are instead permitted. This permits a policy to be tested without actually denying access to resources. In the Learn Mode, all GuardPoint access attempts that would have been denied are instead permitted. These GuardPoint accesses are logged to assist in tuning and troubleshooting policies.

Live Data Transformation (LDT)
A separately licensed feature of Vormetric Transparent Encryption (VTE) that allows you to transform (encrypt or decrypt) or rekey GuardPoint data without blocking use or application access to that data.

local domain
A DSM domain in which DSM administration is restricted to Domain Administrators or Security Administrators assigned to that domain. To access a local domain in the Management Console, a DSM administrator must specify their local domain upon login.

Management Console
The graphical user interface (GUI) to the DSM.

Master encryption key (MEK)
The encryption key for Oracle Database used to encrypt secondary data encryption keys used for column encryption and tablespace encryption. Master encryption keys are part of the Oracle Advanced Security Transparent Data Encryption (TDE) two-tier key architecture.

MEK
See *Master encryption key.*

Microsoft SQL Server
A relational database server, developed by Microsoft.

Microsoft SQL Transparent Data Encryption (MS-SQL TDE)
Microsoft SQL Server native encryption for columns and tables.

multi-factor authentication
An authentication algorithm that requires at least two of the three following authentication factors:
1) something the user knows (for example, password); 2) something the user has (example: RSA SecurID); and 3) something the user is (example: fingerprint). VTE implements an optional form of multi-factor authentication for Management Console users by requiring DSM administrators to enter the token code displayed on an RSA SecurID, along with the administrator name each time the administrator logs on to the Management Console.

multitenancy
A VTE feature that enables the creation of multiple local domains within a single DSM. A local domain is a DSM domain in which DSM administration is restricted to Domain Administrators or Security Administrators assigned to that domain. This allows Cloud Service Providers to provide their customers with VTE administrative domains over which the customer has total control of data security. No other administrators, including CSP administrators, have access to VTE security in a local domain.

offline policy
Policies for Database Backup Agents. *Online policies* are for the File System Agent.

one-way communication
A VTE feature for an environment where the DSM cannot establish a connection to the agent, but the agent can establish a connection to the DSM. For example, the protected host is behind a NAT so protected host ports are not directly visible from the DSM, or the protected host is behind a firewall that prohibits incoming connections, or the protected host does not have a fixed IP address as in the cloud. When an agent is registered with one-way communication, changes made for that protected host on the DSM are not pushed to the protected host, rather as the protected host polls the DSM it will retrieve the change.

online policies
Policies for the File System Agent. *Offline policies* are for Database Backup Agents.

policy
A set of security access and encryption rules that specify who can access which files with what executable during what times, and whether or not those files are encrypted. Policies are created by DSM Security Administrators, stored in the DSM, and implemented on protected hosts by a File system Agent. See “**rule (for policies)**”.

policy tuning
The process of creating a simple Learn Mode policy that allows any protected host user to access a GuardPoint; to examine who accesses the GuardPoint, what executables they use, and what actions they require; and to modify the policy such that it allows the right people, using the right executable, performing the right action to do their job, and prevent anyone else from inappropriate access.

process set
A list of processes that can be used by the users in a user set associated with a policy rule.

protected host
A host on which a VTE Agent is installed to protect that host's data.

public key cryptographic algorithm, public key infrastructure
A cryptographic system requiring two keys, one to lock or encrypt the plaintext, 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. Also called asymmetric key cryptography.

raw device
A type of block device that performs input/output operations without caching or buffering. This results in more direct access.

register host
The process of enabling communication between a protected host and the DSM. Registration happens during agent installation. Before registration can happen, the host must be added to the DSM database.

rekeying
The process of changing the encryption keys used to encrypt data. Changing keys enhances data security and is a requirement to maintain compliance with some data security guidelines and regulations. Also called *key rotation*.

roles
A set of Management Console permissions assigned to DSM Security Administrators by DSM Domain Administrators. There are five roles: *Audit* (can generate and view logging data for file accesses), *key* (can create, edit, and delete keys), *Policy* (can create, edit, and delete policies), *Host* (can configure, modify, and delete protected hosts and protected host groups), and *Challenge & Response* (can generate a temporary password to give to a protected host user to decrypt cached encryption keys when connection to the DSM is broken).

RSA SecurID
A hardware authentication token that is assigned to a computer user and that generates an authentication code at fixed intervals (usually 60 seconds). In addition to entering a static password, Management Console administrators can be required to input an 8-digit number that is provided by an external electronic device or software.

rule (for policies)
Every time a user or application tries to access a GuardPoint file, the access attempt passes through each rule of the policy until it finds a rule where all the criteria are met. When a rule matches, the *effect* associated with that rule is enforced. A rule consists of five access criteria and an effect. The criteria are Resource (the file/directories accessed), User (the user or groups attempting access), Process (the executable used to access the data), When (the time range when access is attempted) and Action (the type of action attempted on the data, for example read. write, rename and so on). *Effect* can be permit or deny access, decrypt data access, and audit access attempt. See *policy*.

secfs
1) The File System Agent initialization script. 2) An acronym for Vormetric Secure File System agent. It generally refers to the kernel module that handles policies (locks, protected host settings, logging preferences) and keys, and enforces data security protection.

secvm
A proprietary device driver that supports GuardPoint protection to raw devices. `secvm` is inserted in between the device driver and the device itself.

Security Administrator
The third-level DSM administrator who does most of data protection work like creating policies, configuring protected hosts, auditing data usage patterns, applying GuardPoints and other duties. The privileges of each Security Administrator is specified by the roles assigned to them by the Domain Administrator. See *roles*. See "**DSM Administrator and types**".

Security Server
See "**DSM**".

separation of duties
A method of increasing data security by creating customized DSM administrator roles for individual DSM administrators such that no one administrator has complete access to all encryption keys in all domains of all files.

signing files
File signing is a method that VTE uses to check the integrity of executables and applications before they are allowed to access GuardPoint data. If file signing is initiated in the Management Console, the File System Agent calculates the cryptographic signatures of the executables that are eligible to access GuardPoint data. A tampered executable, such as a Trojan application, malicious code, or

rogue process, with a missing or mismatched signature, is denied access. Also called *cryptographic signatures*.

Suite B mode
A set of publicly available cryptographic algorithms approved by the United States National Security Agency (NSA). These algorithms enhance security by adding up to 384-bit encryption to the communication between the Web browser and the DSM, the DSM and Agent, and between DSMs in HA environments.

Symmetric-key algorithm
Cryptographic algorithms that use the same cryptographic keys for both encryption of plaintext and decryption of ciphertext. The keys, in practice, represent a shared secret between two or more parties that can be used to maintain a private information link. This requirement that both parties have access to the secret key is one of the main drawbacks of symmetric key encryption, in comparison to public-key encryption.

System Administrator (DSM)
See "**DSM Administrator and types**".

Transparent Data Encryption (TDE)
A technology used by both Microsoft and Oracle to encrypt database content. TDE offers encryption at a column, table, and tablespace level. TDE solves the problem of protecting data at rest, encrypting databases both on the hard drive and consequently on backup media.

user set
A named list of users on which a policy rule applies.

VAE Agent
See "**Key Agent**".

vmd
Acronym for Vormetric Daemon, vmd is a process that supports communication between the DSM and kernel module.

VMSSC or Vormetric Security Server Command Line Interface
See *DSM Automation Utilities*.

Vormetric Application Encryption (VAE)
A product that enables data encryption at the application level as opposed to the file level as is done with VTE.
Where VTE encrypts a file or directory, VAE can encrypt a column in a database or a field in an application. VAE is essentially an API library for key management and cryptographic operations based on PKCS#11. See the *Vormetric Application Encryption Installation and API Reference Guide*.

Vormetric Cloud Encryption Gateway (VCEG)
Vormetric product that safeguards files in cloud storage environments, including Amazon Simple Storage Service (Amazon S3) and Box. The cloud security gateway solution encrypts sensitive data before it is saved to the cloud storage environment, then decrypts data for approved users when it is removed from the cloud.

Vormetric Data Security Platform or VDS Platform
The technology platform upon which all other Vormetric products—Vormetric Transparent Encryption (VTE), Vormetric Application Encryption (VAE), Vormetric Key Management (VKM), Vormetric Cloud Encryption Gateway (VCEG), Vormetric Tokenization Server (VTS), Vormetric Key Management (VKM), and Vormetric Protection for Teradata Database—are based.

Vormetric Encryption Expert or VEE
Earlier name of the Vormetric Transparent Encryption (VTE) product. It may sometimes appear in the product GUI or installation scripts.

Vormetric Key Management (VKM)
Vormetric product that provides a standards-based platform for storing and managing encryption keys and certificates from disparate sources across the enterprise. This includes Vormetric encryption keys, 3rd-party software keys, KMIP device keys and so on.

Vormetric Protection for Teradata Database
Vormetric product that secures sensitive data in the Teradata environment.

Vormetric Security Intelligence
Vormetric product that provides support for Security Information and Event Management (SIEM) products such as ArcSight, Splunk and QRadar. Provides solutions that monitor real-time events and analyze long-term data to find anomalous usage patterns, qualify possible threats to reduce false positives, and alert organizations when needed. Documented in the VDS Platform Security Intelligence User Guide.

Vormetric Tokenization Server (VTS)
Vormetric product that replaces sensitive data in your database (up to 512 bytes) with unique identification symbols called tokens. Tokens retain the format of the original data while protecting it from theft or compromise.

Vormetric Transparent Encryption or VTE
Vormetric product that protects data-at-rest. Secures any database, file, or volume without changing the applications, infrastructure or user experience.

Vormetric Vault
A virtual vault to store 3rd-party encryption keys, certificates and other security objects.

VTE Agent
Vormetric agents that are installed on protected hosts to implement data protection. See "**File System Agent**".

wrapper keys
See "**key wrapping**".

WSDL
Web Services Description Language.