IBM® Guardium Data Encryption

Installation & Configuration Guide

Release 4.0.0.2

IBM Guardium Data Encryption 4.0.0.2 is the same product as Vormetric Data Security (VDS) Release 6.4.0. VDS Release 6 consists of Data Security Manager and Vormetric Agents.

THALES

Vormetric Data Security Platform

Guardium Data Encryption

Release 6

Version 6.4.0

Installation and Configuration Guide

Documentation v1
December 23, 2019

Vormetric Data Security Platform
Installation and Configuration Guide
6.4.0

Table of Contents

PREFACE

The *Installation and Configuration Guide* describes how to install and configure the hardware and virtual appliances, including the V6100 DSM with Hardware Security Module (HSM). This document is intended for system administrators who install the GDE and connect it to a network.

DOCUMENTATION VERSION HISTORY

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

Documentation Changes

Document Version	Date	Changes
6.1.0 v1	09/07/2018	GA release of V6000 and Virtual appliances can now be HMS-enabled by connecting them to an nShield Connect appliance. DSM is now available in the Azure marketplace.
6.1.0 v2	08/06/2018	Removed extraneous text.
6.2.0 v1	4/11/2019	GA release; HA is now active/active, new CLI commands, new API calls expanded KMIP features. This guide contains new troubleshooting and IPMI information. Added rules for hostnames.
6.2.0 v2	3/26/19	Added rules for hostnames.
6.3.0 v1	08/09/19	Supports Efficient Storage with VTE 6.2.0, Excluding files from encryption, fixed security vulnerabilities.
6.3.0 v2	9/16/19	Made clearer the concept that with 6.2.x and 6.3.x user is creating a new cluster, they cannot upgrade a 6.1.x (or lower) cluster.
6.4.0 v1	12/19/19	DSM now compatible with the Luna HSM and Smart cards, users can create GuardPoints for Cloud Object Storage devices, System admins can prevent domain admins from deleting other admins, LDAP limits raised.

ASSUMPTIONS

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

For more information about what's new in this release, refer to the *DSM 6.4.0 Release Notes*. Refer to the *DSM Administrators Guide* for how to administer your DSM Appliance and to the various agent guides for information about Vormetric Data Security Agents.

DSM Hypervisor Support

The virtual DSM can be installed on the following hypervisors:

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

Getting Help

SALES AND SUPPORT

For support and troubleshooting issues:

- **Support portal:** https://supportportal.thalesgroup.com/csm
- **Email address:** technical.support.DIS@thalesgroup.com

For Thales Sales:

- http://enterprise-encryption.vormetric.com/contact-sales.html
- sales@thalesesecurity.com
- (408) 433-6000

1

Installing & Configuring GDE

This chapter describes how to install the IBM Guardium Data Encryption (GDE) virtual appliance as a standalone server or HA node. It contains the following sections:

Overview

The the IBM GDE virtual appliance helps you protect structured and unstructured data and meet compliance requirements. It provides centralized encryption key and policy management to simplify data security management.

In conjunction with the GDE appliance, VTE/VAE/VTS/VPTD agents enable data-at-rest encryption and the collection of security intelligence logs without re-engineering applications or infrastructure.

Installation Process

The installation process includes the following steps:

1. Download the *.tar (Linux) or *.zip (Windows) file from the IBM Passport website and extract the files. The file contains the following files:
   - GDE_PreReq_Download_Credential.bin
   - GDE_4_0_0_1.bin
   - GDE_README.txt
2. Run the GDE_PreReq_Download_Credential.bin file to generate the credentials required to log on to the Thales eSecurity website and download the GDE Appliance OVA file.

NOTE: The GDE_PreReq_Download_Credential.bin must be run in a Linux GUI. Console (CLI) mode is not supported.

3. Run the enabler: GDE_4_0_0_1.bin, to extract a license for the GDE appliance. You will need to upload this license once the GDE appliance is deployed and configured.
4. Deploy the OVA file on your ESXi server.
5. Configure the appliance.

Extract the GDE appliance license

1. Run the enabler file (GDE_4_0_0_1.bin) downloaded from the IBM Passport website. The file must be run on a RedHat or CentOS 6/7 system. To run the file type the following at the prompt:
   ```
   ./GDE_4_0_0_1.bin
   ```
2. Select a language to display the instructions and the EULA, by entering a number that corresponds to that language.
3. Accept the default location to install the license, or follow the on-screen instructions to save it to another location. Press ENTER to continue.
4. Accept the license agreement.
5. The license will be saved on your system in the default location or to the one you specified.

Make sure you can access the system on which you have saved the license from the GDE appliance. You will need to upload this license file to start using the appliance.

Installing the GDE Appliance

This section describes the steps to build a GDE appliance.

System Requirements

- VMware ESXi v5.5 or later with v9 hardware or later
- VMware vSphere Client
- GDE virtual appliance OVA file

These instructions assume the IP address, routing configuration, and DNS addresses for the GDE, appliance allow connectivity to all hosts where the Vormetric Agents are installed.

Hardware Requirements

The hardware hosting the virtual machine must meet the following requirements:

Table 1: Virtual machine hardware requirements

	Number of Agents			
	1 to 10	**11 to 50**	**51 to 250**	**Over 250**
Number of CPUs	2	4	4	6
RAM (in GB)	4	8	12	16
HD (in GB)[a]	250	250	250	above 250

[a] The disk size change was introduced in v5.3.1, however you can still use "thin" provision to minimize storage utilization.

Installation Plan

1. Assemble configuration information using the checklist, see "GDE Appliance Installation Checklist".
2. Complete the pre-configuration tasks described here "Installation and Pre-Configuration tasks", if applicable.
3. Deploy the GDE appliance as described here, "Deploying the GDE Appliance".
4. Setup initial and basic configurations as described here, "Configure the appliance".
5. Verify Web access as described here, "Verify web access".

GDE Appliance Installation Checklist

Table 2: Installation Checklist

REQUIREMENT	VALUE
Software requirements	
Virtual Machine file from Support.	
Hardware requirements for Virtual Machine	
1 virtual socket, 4 cores per socket	
4GB memory	
2 virtual NIC cards	
250GB virtual disk	
Network Information	
eth0—dhcp by default.	IP address netmask default gateway (optional)
eth1—this comes configured with a default IP address 192.168.10.1. We recommend that you retain this configuration in the event that you need a recovery option to access the appliance.	IP address netmask default gateway (optional)
bond0—this interface is used when the eth0 and eth1 interfaces are aggregated into a single logical interface for load balancing./fault tolerance. If configured, the bond0interface supersedes the eth0 and eth1 interfaces, and must be used to access the GDE appliance.	IP address netmask default gateway (optional)
GDE appliance Initial HA node Hostname: FQDN (lowercase only)	
GDE appliance HA node Hostname: FQDN (lowercase only)	
Domain Name Server (DNS) addresses - up to 3 plus optional DNS search domains.	
NTP server FQDN or IP address (if applicable)	
Certificate Information	
GDE appliance Hostname: FQDN (must be an exact match to the hostname)	
Name of your organizational unit	

Name of your organization	
Name of your city or locality. Must be fully spelled out, no abbreviations.	
Name of your state or province. Must be fully spelled out, no abbreviations, e.g., California *not* CA	
Two-letter country code	

Installation and Pre-Configuration tasks

Specify host name resolution method

You can map a host name to an IP address using a Domain Name Server (DNS). DNS is the preferred method of host name resolution. DNS names are case sensitive, make sure host names are correctly entered while configuring DNS and registering hosts. A valid hostname must:

- Be an FQDN
- Be in all lowercase
- Match exactly with:
  - Name set in the CLI (system$ **`set hostname`**)
  - Hostname used when running (`system$ security genca)`
  - Hostname used when running (`system$ security gencert)`

You can also modify the `/etc/hosts` file on the GDE Appliance or identify a host using only the IP address.

- If you use DNS to resolve host names, use the FQDN for the host names.
  - Both forward and reverse address resolution is required for nodes in a cluster.
  - FQDN name must be lowercase
- If you do NOT use a DNS server to resolve host names, do the following on all of the GDE Appliances and the protected hosts:
  - Modify the *host* file on the GDE Appliance: To use names like serverx.domain.com, enter the host names and matching IP addresses in the `/etc/hosts` file on the GDE Appliance using the `host` command under the `network` menu. For example:

```
0011:network$ host add <hostname> 192.168.1.1
SUCCESS: add host
0012:network$ host show
name=localhost1.localdomain1 ip=::1
name=<host name>.<domain name>.com ip=192.168.10.8
name=<host name> ip=192.168.1.1
```

```
SUCCESS: show host
```

You must do one of the following on *each* GDE Appliance, since entries in the host file are not replicated across GDE Appliances.

- Modify the *host* file on the protected hosts: Enter the GDE Appliance host names and matching IP addresses in the `/etc/hosts` file on the protected host.

NOTE: You must do this on EACH protected host making sure to add an entry for all GDE Appliance nodes (if using HA).

- Use IP addresses: You may use IP addresses or the FQDN to identify the host simultaneously. In other words, they don't all have to use an IP address or FQDN.

Port configuration

If a GDE appliance must communicate with a device behind a firewall, you must open various ports in the firewall.

The following table lists the communication direction and purpose of each port you must open.

Table 3: Ports to configure

Port	Protocol	Communication Direction	Purpose
	ICMP	All ICMP	Used to run Ping
22	TCP	Management Console → GDE GDE → SSHD Server	CLI SSH Access Auto-backup via SCP
443	TCP	Browser →GDE Appliance GDE Appliance ↔ GDE Appliance Agent → GDE Appliance	Redirects to either port 8445 or 8448 depending on the security mode. (8445 is used in compatible & RSA modes; 8448 is used in Suite B mode, for secure communication between GDE Appliances in an HA cluster and for LDT registration.)
5432	TCP	GDE HA node 1 ↔ GDE HA node 2	Uses SNMP to get HA node response time.
7025	TCP/UDP	SNMP Manager → GDE GDE ↔ GDE	SNMP queries from an external manager and to get HA node response times.

Table 3: Ports to configure

8080	TCP	Agent → GDE	Port 8080 is no longer used for registration, but you can manually close/open this legacy port for new deployment, for backward compatibility if you use previous versions of the agent and need to register to 8080. Default is on, (open). **Syntax** `# security legacyregistration [ on | off | show ]`
8443	TCP	Agent → GDE	Fallback RSA TCP/IP port through which the agent communicates with the GDE, in case 8446 is blocked. The agent establishes a secure connection to the GDE appliance, via certificate exchange, using this port.
8444	TCP	Agent → GDE	Fallback RSA port via which the Agent log messages are uploaded to GDE, in case 8447 is blocked.
8445	TCP	Browser → GDE GDE ↔ GDE (fallback)	Management Console, VMSSC, and fallback for RSA HA communication in case port 8448 is blocked.
8446	TCP	Agent → GDE	Configuration Exchange using Elliptic Curve Cryptography (Suite B).
8447	TCP	Agent → GDE	Agent uploads log messages to GDE using Elliptic Curve Cryptography (Suite B).
8448	TCP	Browser → GDE GDE ↔ GDE	GUI management during enhanced security using Elliptic Curve Cryptography (Suite B). Also for secure communication between GDE appliances in HA clusters.
9004	TCP	GDE Appliance ↔ network HSM	GDE Appliance communication with nShield Connect and its associated RFS
9005	TCP	GDE Appliance ↔ remote admin	Used by Remote Administration Service process to accept connections from the Remote Administration Client.

Access the Command Line Interface (CLI)

The CLI commands are used to configure the appliance. The commands are grouped into the following categories or *submenus*. Entering ? on the CLI command line lists those categories:

```
0000:dsm$ ?
network      Networking configuration
system       System configuration
hsm          HSM configuration
maintenance  System maintenance utilities
```

```
ha            HA configuration
ipmi          IPMI configuration
user          User configuration
exit          Exit
```

To enter a submenu, enter a name or just the first few letters of the name. To display the commands for that submenu, enter a ?. For example, the submenu `maintenance` is used to provide maintenance utilities:

```
0001:dsm$ main
0038:maintenance$ ?
showver       Show the installed VTS version
ntpdate       Set ntp services
date          Set system date
time          Set system time
gmttimezone   Set system time zone
diag          OS diagnostics
up            Return to previous menu
exit          Exit
```

Every command has usage and example input. Type the command without a value:

```
0039:maintenance$ ntpdate

usage: ntpdate {sync | add SERVER_ADDRESS | delete SERVER_ADDRESS | on |
off | show }
0040:maintenance$ date

month=Mar day=17 year=2015

Show system date SUCCESS

  0041:maintenance$ time

hour=11 min=11 sec=36 zone=PDT

Show system time SUCCESS

  0042:maintenance$ gmttimezone

usage: gmttimezone {list|show|set ZONE_NAME}

  0043:maintenance$ diag

usage: diag [log [ list | view LOG_FILE_NAME] | vmstat | diskusage |
hardware | osversion | uptime ]

  0044:maintenance$
```

You must enter the submenu to execute the submenu commands. For example, the reboot command is in the system submenu, so you would enter system, then enter reboot. To return to the main level when finished, enter up.

A complete description of the CLI commands can be found in the *Administrators Guide*.

Deploying the GDE Appliance

This section describes how to deploy the OVA file to create the appliance. The GDE appliance uses static IP addresses and cannot be assigned an address by DHCP.

1. Open the VMware vSphere Client.
2. Click **File > Deploy OVF template**.
3. Click **Browse** and locate the OVA file. Select the file and click **Next**. The *OVF Template Details* page appears.
4. Click **Next**. The *Name and Location* page opens.
5. Type in a name for the Virtual Appliance and then click **Next**. The *Storage* page opens.
6. Select a destination for the Virtual Appliance and then click **Next**. The *Disk Format* page opens.

   Select the type of provisioning based on the storage characteristics for your system. The options are:

   - **Thick Provisioned Lazy Zeroed**: creates the VM and allocates all the blocks for the VM but doesn’t zero them.
   - **Thick Provisioned Eager Zeroed**: creates the VM, allocates and zeros all the blocks.
   - **Thin Provision**: creates the VM with just the header information, but it does not allocate or zero blocks.

   In the following example, we use **Thick Provisioned Lazy Zeroed**.
7. Select **Thick Provisioned Lazy Zeroed** and click **Next**. The *Ready to Complete* window opens.
8. Click **Finish** to deploy the Virtual Appliance. This takes a few minutes.
9. At the message **Completed Successfully**, click **Close**. The main screen of the vSphere Client appears.
10. In the left pane, select the Virtual Appliance you just created and then click the power on icon in the tool bar. It takes about a half hour to provision the VM and build the appliance.
11. To watch the output as the installation progresses, click the Console tab and click inside the console window. When the installation is finished, continue to the next section.

Configure the appliance

This section describes how to configure network settings, NTP, Time Zone and Date/Time, and the hostname. It describes how to configure a bonded NIC device type should you choose to use this feature. It also describes how to generate a certificate authority (CA), add console administrators, and verify Web access.

If you are setting up the GDE appliance in a high availability (HA) deployment, configure each node as a standalone appliance, which is the same procedure as described here. Then add each node to the initial node and join the HA network.

Configure network settings

1. Access the GDE appliance CLI and log in with the default login and password:

```
Login: cliadmin
Password: cliadmin123
```

2. The Thales EULA is displayed, type ‘y’ to accept and press **Enter**.
3. When prompted, type in a new password and press **Enter**. Reconfirm your password.

Warning! Do not lose this password.

4. Navigate to the *network commands* menu. Type:

```
0000:dsm$ network
```

5. Add an IP address for the GDE appliance. Type:

NOTE: We recommend that you retain the default eth1 IP address configuration in the event that you need a recovery option to access the GDE appliance.

```
0001:network$ ip address init <GDE IP address>/<subnet mask (e.g. 16 or
24)> dev eth0/eth1
```

Example: `ip address init 192.168.10.2/16 dev eth1`

IPv6 Example: `ip address init fa01::3:15:130/64 dev eth1`

NOTE: If you are connected via eth0 and you choose to configure eth0 with a new IP address, you will be disconnected at this step. Reconnect on the new IP address.

6. (Optional) You may choose to configure the eth0 interface instead of retaining the default IP address 192.168.10.1, if for example, you want the GDE appliance to communicate with agents on a different subnet, or access the Management Console from a different subnet. To configure an IP address for eth0, type:

   ```
   0001:network$ ip address init <eth0 IP address>/<subnet mask (e.g., 16
   or 24)> dev eth0
   ```

   Example: `ip address init 192.168.10.3/16 dev eth0`

   IPv 6 Example: `ip address init fa01::3:15:130/64 dev eth0`

   The following warning is displayed:

   ```
   WARNING: Changing the network ip address requires server software to be
   restarted.
   Continue? (yes|no) [no]:
   ```

   Type ‘yes’ to continue with the IP address configuration.

7. Add the IP address for the default gateway. Type:

   ```
   0001:network$ ip route add default table main.table dev [eth0 or eth1]
   via <IP address for the default gateway>
   ```

   Example: `ip route add default table main.table dev eth0 via 192.168.1.5`

   IPv 6 Example: `ip route add default table main.table dev eth0 via fa01::3:15:120`

8. Verify interface settings. Type:

   ```
   ip address show
   ```

9. Verify route settings. Type:

   ```
   ip route show
   ```

10. If you are using DNS, set the initial DNS server for the GDE appliance. Type:

    ```
    dns dns1 <ip address for dns server 1>
    ```

11. If you have a second or third DNS server, set them for the GDE appliance. Type:

    ```
    dns dns2 <ip address for dns server 2>
    ```

12. If you want to set the search domain, type:

```
dns search <search_domain>
```

13. Show the DNS settings. Type:

```
dns show
```

14. Return to the main menu. Type:

```
up
```

Configure a bonded NIC device

This section describes how to aggregate the two NICs on the GDE appliance into a single logical interface to provide load balancing and/or fault tolerance. The bonded NIC device is called `bond0`.

On the virtual appliance, you must configure at least two NICs and define them as `eth0` and `eth1` in order to enable the bond0 device type. Any additional physical/virtual NICs are ignored. For virtual appliances where only one network connector is configured for a virtual machine, the `bond0` interface cannot be enabled—the network interface itself can be up but, no IP address can be assigned to it.

The NIC bonding setting is system specific. If it is to be used for all nodes in a cluster, it must be enabled on all nodes individually.

1. Access the GDE appliance CLI and login with your login credentials. If this is the first time you are logging in, then you will be required to accept the license agreement and change the default password, see “Configure network settings”.
2. Navigate to the network commands menu;

```
0000:dsm$ network
```

```
0001:network$
```

3. Enable the bonded NIC;

```
0001:network$ ip address init <ip_address>/<subnet_mask> dev bond0
```

Example: `ip address init 1.2.3.4/16 dev bond0`

In the event that a bonded NIC is being configured after the initial configuration, or after the GDE appliance has been upgraded, if you want to reuse an IP address that was originally assigned to `eth0` or `eth1`, then you must delete that address from `eth0` or `eth1` first, and then reassign it to the `bond0` interface.

4. Add a default gateway for the `bond0` device;

```
0001: ip route add default table main.table dev bond0 via
<gateway_ip_address>
```

Example: `ip route add default table main.table dev bond0 via 1.2.7.8`

If a `bond0` interface is configured after setting up the `eth0` and/or `eth1` interfaces, and it is configured with an IP address that is on the same subnet as a default gateway, that gateway configuration continues to apply. However, if you configure `bond0` with an IP address on a different subnet, you will have to reconfigure the default gateway.

5. You can change the bonding driver mode based on your requirements. There are seven modes available from 0-6. See “Bonding driver modes” for more information. Note however, that only the default options are available with each of the modes and these options cannot be changed.

   When the mode option is specified the speed option cannot be specified (i.e. the options mode and speed are mutually exclusive). In other words, bond0 does not take the speed option and both eth0 and eth1 don't take the mode option. However, the MTU and up/down options can still be used for the bond0 device.

   To set or change the mode type:

   ```
   0002:network$ ip link set bond0 mode <mode>
   ```

   Example: `ip link set bond0 mode 2`

   To see what mode is currently in use type:

   ```
   0002: network$ ip link show bond0
   ```

6. To disable or break up a bonded NIC type, you can use either the delete or flush command. Delete will only delete a specific IP address (multiple can be assigned) and flush will clear all assigned IP addresses.

   ```
   0003:network$ ip address delete <ip_address>/<subnet_mask> dev bond0

   0003:network$ ip address flush bond0
   ```

   Routes that are associated with this bonded NIC device will also be deleted.

Bonding driver modes

The modes specify the bonding policies. The following modes are supported (see Table 4 below), but none of the options for the modes are configurable and take the default values for those modes, except for the `miimon` setting. The `miimon` setting specifies the MII link monitoring frequency in milliseconds, which determines how often the link state of each slave is inspected for link failures. The `miimon` setting has a value of 100 instead of the default value of 0.

Table 4: Bonding driver modes

Mode	Name	Description	Load-balancing	Fault tolerance
0	balance-rr	Round-robin policy. Transmit packets in sequential order from the first available through the last. This is the default mode for the bonded NICs.	Yes	Yes

Mode	Name	Description	Load-balancing	Fault tolerance
1	active-backup	Active-backup policy: Only one slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch.	No	Yes
2	balance-xor	XOR policy: Transmit based on the selected transmit hash policy. The default policy is a simple [(source MAC address XOR'd with destination MAC address) modulo slave count].	Yes	Yes
3	broadcast	Broadcast policy: transmits everything on all slave interfaces.	No	Yes
4	802.3ad	IEEE 802.3ad Dynamic link aggregation. Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification.	Yes	Yes
5	balance-tlb	Adaptive transmit load balancing: channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.	Yes	Yes
6	balance-alb	Adaptive load balancing: includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the slaves in the bond such that different peers use different hardware addresses for the server.	Yes	Yes

Configure NTP, time zone, date, time

You must have the correct time set on your GDE appliance(s) as this will affect system functions such as agent registration, log timestamps, high availability cluster synchronization, and certificate exchange. Although configuring an NTP server is not mandatory, it is strongly recommended.

1. Navigate to the *maintenance commands* menu. Type:

   ```
   0000:dsm$ maintenance
   ```

2. Show the current ntpdate settings. Type:

   ```
   0001:maintenance$ ntpdate show
   ```

3. Add a new ntpdate server. Type:

```
0002:maintenance$ ntpdate add <IP address/Hostname for the ntpdate server>
```

Repeat this step for each ntpdate server.

4. Activate the ntpdate server connection. Type:

```
0003:maintenance$ ntpdate on
```

5. Show the current timezone settings. Type:

```
0004:maintenance$ gmttimezone show
```

6. Set the country and city where the GDE appliance resides. Type:

```
0005:maintenance$ gmttimezone set <country/city>
```

7. Set the date. (If you used `ntpdate synch`, this step is not necessary.) Type:

```
0006:maintenance$ date <mm/dd/yyyy>
```

8. Set the time. (If you used `ntpdate synch`, this step is not necessary.) Type:

```
0007:maintenance$ time <hh:mm:ss>
```

Where *hh* is 00 to 23.

9. Verify your settings. Type:

```
0008:maintenance$ time
0009:maintenance$ date
```

10. Return to the main menu. Type:

```
0010:maintenance$ up
```

Configure the hostname

1. Navigate to the *system* menu. Type:

```
0001:dsm$ system
```

2. Show the current setting. Type:

```
0002:system$ setinfo show
```

The default host name in the output is *your. name.here*.

3. Set the hostname. You must enter the fully qualified domain name for the GDE appliance. Type:

```
0003:system$ setinfo hostname <FQHN>
```

Example:

```
0003:system$ setinfo hostname securityserver.company.com
```

Generate the Certificate Authority

1. Generate a new certificate authority for the GDE appliance. Type:

   ```
   0004:system$ security genca
   ```

2. A warning is displayed, informing you that all agents and peer node certificates will need to be re-signed after the CA and server certificate have been regenerated, and the GDE appliance server software will be restarted. Type 'yes' to continue, the default is 'no'.
3. Enter the FQDN of this appliance, the name displayed in 'This Security Server host name [FQDN of the GDE appliance]', should be correct if you entered the host name information in the previous sections correctly. Press **Enter** to accept the name.
4. Next, enter the information required to generate the certificate. Answer the prompts:

   a. What is the name of your organizational unit? []:

   b. What is the name of your organization? []:

   c. What is the name of your City or Locality? []:

   d. What is the name of your State or Province? []:

   e. What is your two-letter country code? [US]:
5. Once the certificate is signed, return to the main menu. Type:

   ```
   up
   ```

Add CLI administrators (optional)

With separation of duties for good security practices, CLI administrators can only log into the CLI and administer the GDE appliance. Management Console administrators can only log on to the Management Console to administer the GDE appliance.

1. Navigate to the *user commands* menu. Type:

   ```
   0001:dsm$ user
   ```

2. Add an administrator. Type:

   ```
   0002:user$ add <administrator name>
   ```

3. When prompted, enter a password. The password criteria are:
   - Does not have repeating characters
   - Uses at least 1 upper and 1 lower case character
   - Uses at least 1 special character
4. Return to the main menu. Type:

   ```
   0003:user$ up
   ```

Verify web access

The Management Console is a Web-based GUI used for day-to-day security and administration tasks. Open a browser and confirm access over HTTPS to either the GDE appliance hostname (if configured in DNS) or the IP address defined in "Configure network settings" on page 16.

Example URL:

```
https://securityserver.vormetric.com
```

If the URL doesn't work because, for example, port 443 is blocked by a firewall, specify port 8448.

Example:

```
https://securityserver.vormetric.com:8448
```

or specify port 8445. Example

```
https://securityserver.vormetric.com:8445
```

If the link still does not work, make sure all the necessary ports are open, see Table 3, "Ports to configure," on page 12.

The first time you connect to the appliance via a web browser, a self-signed certificate is used by default. Your browser will display a warning about the SSL certificate, follow the instructions on your browser to continue with the default self-signed certificate. You can configure the GDE appliance to use third party signed certificates after you have logged in for the first time. Refer to the *GDE Administrators Guide*, chapter 6 for procedures to do this.

The default user name and password to log on to the GDE appliance the for first time are; `admin` and `admin123`. You will be prompted to reset the password. The password criteria are:

- Does not have repeating characters
- Uses at least 1 upper and 1 lower case character
- Uses at least 1 special character

Upload a license file

The first time you log on to the GDE appliance, the dashboard displays "License file not found," and all you will see are the *Dashboard* and *System* tabs. You need to click **System**, select **License**, and then **Upload the license file**.

Upload the license file that you extracted from the enabler package.

After uploading your license file, all the other tabs for which you have licenses are displayed.

Upgrading GDE Appliance Software

2

This chapter contains instructions for migrating data from GDE 2.0 release 2h to GDE 4.0.0.0.

It contains the following sections:

Upgrade Paths

The following table describes the GDE upgrade path based on your current version:

Figure 1: Upgrade Path for GDE

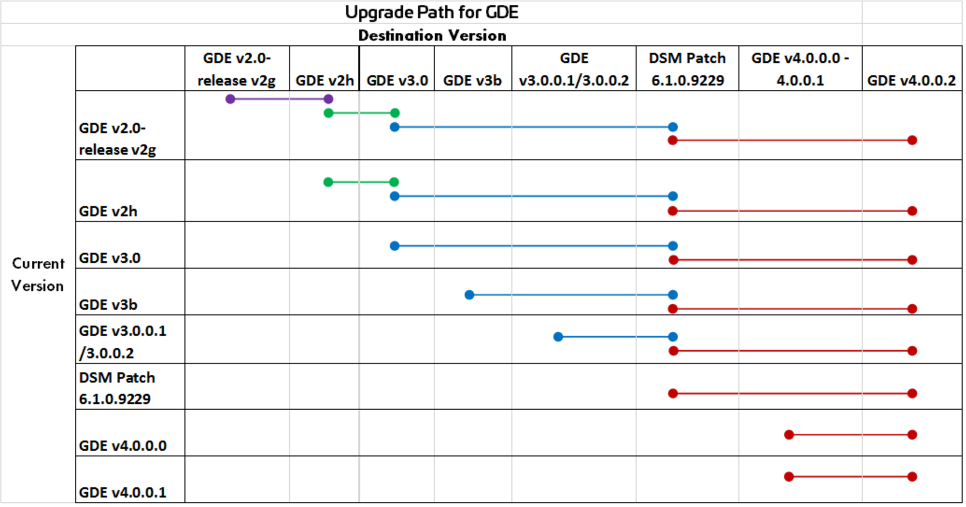

Upgrade Path for GDE

Destination Version

Current Version		GDE v2.0-release v2g	GDE v2h	GDE v3.0	GDE v3b	GDE v3.0.0.1/3.0.0.2	DSM Patch 6.1.0.9229	GDE v4.0.0.0 - 4.0.0.1	GDE v4.0.0.2
	GDE v2.0-release v2g	●	●						
			●	●					
				●	─	─	●		
							●	─	●
	GDE v2h		●	●					
				●	─	─	●		
							●	─	●
	GDE v3.0			●	─	─	●		
							●	─	●
	GDE v3b				●	─	●		
							●	─	●
	GDE v3.0.0.1 /3.0.0.2					●	●		
							●	─	●
	DSM Patch 6.1.0.9229						●	─	●
	GDE v4.0.0.0							●	●
	GDE v4.0.0.1							●	●

- If you are on a GDE version prior to GDE 2.0 release 2h, you must upgrade to release 2h.
- The change from GDE r2h to GDE r3.0 involves an operating system change. Therefore, you must back up your current system, build a v3.0 system, and then restore to your r3.0 system.

- The change from DSM patch 6.1.0.9229 to GDE 4.0.0.2 involves a database migration. That migration is built into 6.1.0.9229 which is why all users must upgrade to 6.1.0.9229 first, and then upgrade from 6.1.0.9229 to GDE 4.0.0.2.

Migrating to a GDE 3.0 Appliance

As illustrated in Figure 1, you must first upgrade your GDE appliance to GDE 3.0. To complete this task, you must:

1. “Backup the current configuration”
2. “Install and configure a GDE appliance”
3. “Restore backup to new GDE appliance”

NOTE: If you are upgrading from a version higher than GDE 3.0, you can upgrade directly to the DSM patch 6.1.0.9229. Prior to upgrading, follow the steps below to backup your current configuration, in case the upgrade fails to upgrade your system properly.

Backup the current configuration

A backup is a snapshot of a GDE appliance configuration. When a backup is restored, the GDE appliance Management Console will contain and display the same information captured at the time the backup was originally made.

Create or import a wrapper key

GDE appliance backup files are encrypted with a wrapper key to keep them secure. This wrapper key must be created, or imported from a previous create operation, before creating a backup. The same wrapper key used to encrypt a backup is also required to restore that backup. For additional security, wrapper keys can be broken up into key shares—pieces of a wrapper key. These key shares can then be divided amongst two or more custodians, such that each custodian must contribute their key share in order to assemble a complete wrapper key. This is also referred to as split key knowledge or M of N configuration.

For example, you can break up the wrapper key amongst a total of five custodians and set the minimum number of required custodians at two. When the wrapper key is required, at least two of the custodians must contribute their key share in order to assemble a complete wrapper key. The wrapper key must be created by a System administrator or an AllAdmin.

1. Log on to the Management Console as a System administrator or an AllAdmin.
2. Select **System > Wrapper Keys** from the menu bar.

3. In the *Wrapper Keys* window, select **Operation > Create**, then click **Apply** to create the wrapper key.
4. Select **System > Backup and Restore > Manual Backup and Restore** from the menu bar.

   A confirmation message also displays on this tab, stating that the wrapper key exists. You can now proceed with creating a backup.
5. Click **Backup** tab and select **Ok**.

NOTE: Some Browsers will automatically save and download the file. Some will display a Save as dialog.

6. Click **Save** in the File Download dialog box, if your browser displays one.
7. Save the file to a secure location that you are sure will still be accessible if the server fails.

   By default, the file name will be in the format: backup_config_<gde *server name*>_yyyy_mm_dd_hhmm.tar (.zip for Windows). Where <*gde_server name*> is the FQDN of the GDE appliance that is being backed up.
8. Return to the **System > Wrapper Keys** menu option and select **Operation > Export** to export key shares.
9. Set a number for both the **Minimum Custodians Needed** and the **Total Number of Custodians**.

   This setting splits the wrapper key value among multiple custodians. If only a single administrator is to control the wrapper key, enter a value of 1 in both fields.
10. Select the GDE appliance administrators who will serve as custodians for the wrapper key shares.

    Administrators of type *System Administrator* and *All* are listed. You can select any of these administrators, with the exception of the default initial log-on administrator *admin*, as a custodian.
11. Click **Apply** on the bottom right hand corner.

    If you have selected more than one custodian, each of them is given a share of the wrapper key. The wrapper key share is displayed on their **Dashboard** page, beneath the fingerprint for the CA, when they log into the Management Console. The generated wrapper key, or key shares, are exported and are visible on the **Dashboard**, beneath the fingerprint for the CA. The **Wrapper Key Share** displayed on the *Dashboard* is a toggle. Click **Show** to display the wrapper key share value. Each administrator must see a unique wrapper key share displayed on the dashboard beneath the fingerprint for the CA.
12. On the Dashboard, click **Wrapper Key Share** string to hide the value and display ‘Show’.
13. Ensure the administrator(s) or wrapper key custodian(s) securely store a copy of this key or key share. This is required, as part of their role in a GDE appliance restore operation.

NOTE: Do NOT lose the wrapper key used to create the backup. You cannot restore the backup without the wrapper key that was used to create it.

14. Create a backup of the GDE appliance configuration after the wrapper key has been created.

Create a backup

1. Log on to the Management Console as an administrator of type *System Administrator* or *All*.
2. Select the **System > Backup and Restore** menu option. The *Manual Backup and Restore* page opens.
3. Click the **Backup** tab and then click **Ok**.
4. Click **Save** in the **File Download** dialog box. Save the file to a secure location that you are sure will still be accessible if the server fails. By default, the file name will be in the format:

   ```
   backup_config_<gde server name>_yyyy_mm_dd_hhmm.tar (.zip)
   ```

   Where <*gde server name*> is the FQDN of the GDE appliance that is being backed up.
5. Save the backup to a secure location. Access to the backup should be limited to only a few employees and should be audited.

Install and configure a GDE appliance

In order to ensure the continuity of your GDE deployment, you must configure the GDE 3.0 appliance with the same hostname as the GDE appliance that is being migrated. You can assign a new IP address to the appliance, however you must ensure that the hostname resolution method in use is correspondingly updated.

NOTE: If you are using a third party SSL certificate, then a change in the GDE appliance hostname will cause a cause a conflict when you restore the backup.

You will have to upload a new third party certificate with the new GDE appliance hostname.

After you configure the new GDE 3.0 appliance with the same hostname and IP address, you must take the old appliance off of the network, otherwise any registered agents will try and communicate with both the old and the new GDE appliance and cause conflicts in your system.

If you configure a GDE 3.0 appliance and give it a new hostname, then when you reach your destination version of the GDE software, VTE agents that were registered with the earlier GDE appliance backup will have to re-register with the new GDE appliance. Refer to the *VTE Agent Installation and Configuration Guide* for detailed procedures to re-register agents.

For procedures to install and configure a GDE appliance, see “Installing the GDE Appliance”.

Restore backup to new GDE appliance

The GDE appliance backup is restored via the Management Console.

1. Locate the backup that is to be restored.
2. Log on to the Management Console as an administrator of type *System Administrator* or *All*.

NOTE: If you already have the Wrapper Key imported, skip to Step 7.

3. Import wrapper keys. Select **System > Wrapper Keys** from the menu bar.
4. Select **Import** from the **Operation** dropdown menu. Click **Add**.
5. If key shares have created from the wrapper key, paste a Key Share value from one previously stored with a custodian into the **Key Share** text field and click **Ok**.

   Repeat steps 5 and 6 for each administrator selected as a key custodian if you have chosen to have more than one custodian for the wrapper key. A key share must be imported for at least as many as were specified by the Minimum Number of Custodians value when the wrapper key was exported.
6. Click **Apply** to finish importing the wrapper key.
7. Restore the backup file. Select **System > Backup and Restore** from the menu bar.
8. Select the **Restore** tab.
9. Click **Browse**. Locate and select the backup file to restore.
10. Click **Ok**. The restored file uploads and the GDE appliance disconnects from the Management Console. The restore operation takes up to 30 minutes to complete.

    If the browser has not refreshed automatically after the restore operation, you must manually refresh the browser to log back on to the Management Console.

    If you were using a third party SSL certificate, this certificate will now also be restored as part of this operation. See “Verify web access” on page 23 for more details.
11. Log back on to the Management Console as an administrator of type *System* or *All*. Verify that the configuration is restored correctly

Your GDE appliance is now migrated to version 3.0.

Upload a license

As part of the process of configuring a new GDE 3.0 appliance, you will have already uploaded the GDE 3.0 license. However, once you restore a backup of the earlier GDE version, you will need to upload the license once again. Follow the steps to upload the license as described here, “Upload a license file”.

Upgrade to the DSM patch 6.1.0.9229

After upgrading to GDE 3.0 or higher, you must upgrade to the DSM patch 6.1.0.9229:

1. Click **System > Software upgrade**.
2. Delete the idle version if it exists.
3. Select the 6.1.0.9229 upgrade tar file.
4. Click **Upgrade.**

Upgrade to GDE 4.0.0.2

After upgrading to 6.1.0.9229, you can migrate to GDE 4.0.0.2.

1. Click **System > Software upgrade**.
2. Delete the idle version if it exists.
3. Select the GDE 4.0.0.2 upgrade tar file.
4. Click **Upgrade.**

HA for V6x00 and Virtual Appliances

3

This chapter describes how to set up High Availability (HA) for V6x00 hardware and virtual appliances. Refer to the High Availability chapter in the *Administrators Guide* for details about managing an HA deployment.

This appendix contains the following sections:

HA Overview

To configure High Availability (HA) for GDE Appliancess, you need to be a System/All administrator and have GDE Appliances CLI privileges. A GDE Appliances HA configuration consists of two or more GDE Appliances HA nodes.

As of GDE Appliances v6.2.0, HA is now configured as Active-Active. This means that there is no longer a primary or a failover node. All nodes are peers. When one node fails, the other nodes continue operating normally. When the failed node is working properly again, it synchronizes with the other HA nodes. that if you are migrating from 6.1.x or an earlier version, then you can't upgrade or migrate. You have to create a new cluster.

NOTE: If you are migrating from an HA cluster that is GDE Appliances v6.1.x or an earlier version, then you cannot upgrade or migrate your cluster to 6.2.x or 6.3.x. You must create a new cluster. See "Migrating from DSM v6.1.0.9229 to DSM 6.4.0" on page 111 for more information.

Supported HA Deployments

You must have at least two GDE Appliance HA nodes installed on the same network to create an HA cluster. The maximum number of nodes allowed in an HA cluster is eight.

To ensure reliable operation, the appliances in an HA cluster must run the same version of the GDE Appliances software and have the same hardware configurations.

Configuring HA for Virtual Appliances

This section describes how to configure an HA cluster for a virtual appliance in an HA cluster.

Prerequisites

Refer to “Configure network settings”, for how to configure a virtual appliance.

Before you set up your HA cluster, do the following:

1. Specify a hostname resolution method, see “Specify host name resolution method”.

   You can map a host name to an IP address using a Domain Name Server (DNS). DNS is the preferred method of host name resolution.

   You can modify the `hosts` file on the HA node:

   a. Log in to the CLI menu.

   b. Type: **`network`**

   c. Type: **`host add <hostname> <IP address>`**

   You can also identify a host using only the IP address.

   - If you use DNS to resolve host names, use the FQDN for the host names.
     - Both forward and reverse address resolution is required for nodes in a cluster.
     - FQDN name can be lower or uppercase, however, the GDE Appliances converts and displays all node names as lowercase.

2. Open all required ports. To see the ports to configure, see Table 4, “Ports to Configure,” on page 67.

NOTE: For upgrades and fresh installations of GDE Appliances 6.2.0, if you are using HA, you must open port 5432 in your firewall to allow communication between GDE Appliances HA nodes. For Azure and AWS platforms, you will need to add this port to your security groups. You can now close port 50000 as it is no longer used.

3. Perform a ‘ping’ operation on all of the GDE Appliancess to ensure that network communication is working between the GDE Appliances HA nodes.

Network Latency

If the network latency between the HA nodes exceeds 100ms, you may experience delays in HA replication, especially if you have many policies, or you have large policies that contain many resource sets, user sets, etc.

Another factor in network latency is the Policy Version History setting (System > General Preferences > System > Maximum Number of Saved Policy History). Each time changes are made to a policy, a new version of that policy is created. This setting determines how many previous versions of the policy will be kept. The more versions that are kept, the longer the delay because it increases the time required to replicate policy data to the cluster nodes. See Table 3, “General Preferences System tab attribute values and use,” on page 36 for details about this setting. We recommend changing this value to 0 or 5 from the default of 10 if you experience network latency.

Adding Nodes to an HA Cluster

Only add one HA host to the cluster at a time. Adding multiple hosts at the same time does not work.

1. Install and configure your GDE Appliances as described in previous chapters of this guide.

NOTE: The license must be installed on the GDE Appliance designated as HA node 1 before you can configure the other HA nodes.

2. On HA node 1 (the Initial Server), log on to the Management Console as an administrator of type System, or All.
3. Click **High Availability** in the menu bar. The *High Availability Servers* window opens.
4. Click **Add**. The **Add High Availability Server** window opens.
5. In the **Server Name** field, enter the host name or FQDN of a GDE Appliance node.
6. Click **Ok**. The GDE Appliance node is listed in the High Availability Servers window. It is designated as ‘Not Configured’.

Figure 2: Node added but not configured/joined to the cluster

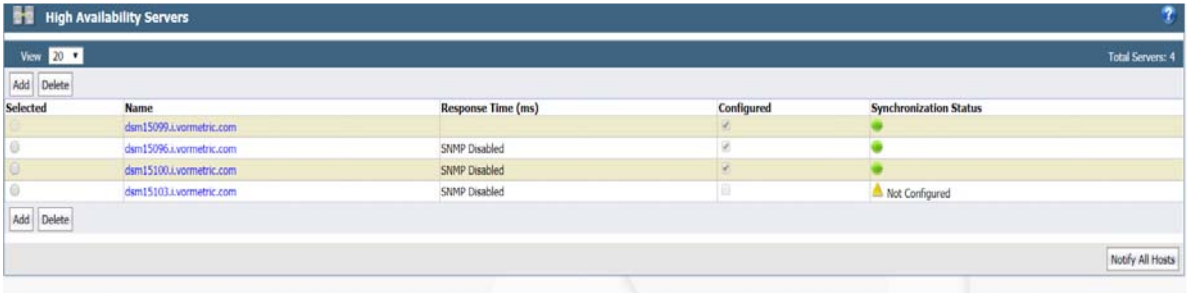

Selected	Name	Response Time (ms)	Configured	Synchronization Status
○	dsm15099.i.vormetric.com		✓	●
○	dsm15096.i.vormetric.com	SNMP Disabled	✓	●
○	dsm15100.i.vormetric.com	SNMP Disabled	✓	●
○	dsm15103.i.vormetric.com	SNMP Disabled	□	⚠ Not Configured

NOTE: You can also add nodes in the CLI. See the High Availability Category section in the CLI chapter in the GDE Appliances Admin guide.

Join a Node to an HA Cluster

Joins the current node to the HA cluster. If you are joining an HA cluster after an upgrade, and the node previously had a host assigned to it, after it successfully joins the cluster, the GDE Appliances asks if you want the host restored to the node. Only join one HA host to the cluster at a time. Joining multiple hosts at the same time does not work.

To join an HA node to the HA cluster:

1. In the CLI, log on to HA Node 2 on the Cluster.
2. Switch to the HA menu, type:

   ```
   0000:dsm$ ha
   ```

3. Join the node to the cluster. Type:

   ```
   0000:dsm$ join
   ```

   System Response:

   ```
   WARNING: This server node is about to join an HA cluster.
   Please make sure the HA cluster is running and has this server node in its HA
   node list.
   This may take several minutes.
   ```

NOTE: Sometimes, when GDE Appliances nodes are spread far apart geographically, or are in a cloud environment, the Join function takes so long that the ssh session times out and terminates automatically before the Join can finish. If the Join function fails, type: **`join longwait`** to make the join command proceed in a 'longwait mode' (as opposed to the 'normalwait' mode). The difference between the two modes is the duration that it waits for the node replication status state to be set to 'ready.' Longwait waits for approximately twice as long as the normalwait mode.

4. Follow the prompts:

   a. Type **yes** to continue.

   b. For **HA Initial Server host name**, type the hostname of HA Node 1.

   c. For **Initial Security Server system administrator name**, type the UI admin name for the **Current** node.

   d. For **Initial Security Server system administrator password**, type the UI admin password for the **Current** node.

   System Response:

   ```
   This node may have multiple IP addresses. All the agents will
   have to connect to the Security Server using the same IP.

   Enter the host name of this node. This will be used by Agents
   to talk to this Security Server.

   This Security Server host name[dsm15100.i.vormetric.com]:

   Please enter the following information for key and certificate
   generation.
   ```

5. The HA cluster will issue the certificate using the information you provide in the following steps:

   a. What is the name of your organizational unit? []:

   b. What is the name of your organization? []:

   c. What is the name of your City or Locality? []:

   d. What is the name of your State or Province? []:

   e. What is your two-letter country code? [US]:

   f. What is your email address? []:

   System Response:

   ```
   WARNING: The following information you entered will be used to
   join this server to the HA cluster, please make sure the
   information is correct

   Initial Security Server host name:HaNode1.i.vormetric.com

   Initial Security Server system administrator name:voradmin

   Initial Security Server system administrator password:xxxxxxxx

   This Security Server host
   name[dsm15100.i.vormetric.com]:HaNode2.i.vormetric.com

   The name of your organizational unit: TP

   The name of your organization: Thales
   ```

```
The name of your City or Locality: SJ

The name of your State or Province: CA

Your two-letter country code[US]: US

What is your email address: groot@thales.com

Restore original host assignment back to this node
(yes|no)[yes]:

Continue? (yes|no)[no]: yes
```

6. Type **yes** to continue.
7. The installation utility creates certificates, completes the installation process, and then starts the HA node. This may take a few minutes.

   The CA certificate fingerprint is displayed.
8. On HA node 1 on the Management Console, click the **Dashboard** tab.
9. Match the fingerprint from the output on HA node 2 with the **RSA CA fingerprint** on the HA node 1 **Dashboard**.

   Sample output:

```
Initial_Server=HaNode1.i.vormetric.com
CAs_Fingerprint=8F:104:BE:78:0E:BB:28:4F:64:4D:54:5A:B1

Ensure the fingerprint listed above matches the one on the

Security Server web console dashboard.

Self test in progress: passed

Starting data store

Starting Security Server

Security Server started in compatible mode

SUCCESS: joined to the HA cluster. The server is started.
Please verify the fingerprint.

0009:ha$
```

10. In the GUI, click the **High Availability** tab. In the row for the HA node 2, the **Synchronization status** should contain a green circle and the Configured column should contain a check.

Configuring High Availability for Network HSM-enabled Nodes

GDE Appliance appliances, which do not have a built-in HSM, can be configured to use a network HSM via an nShield Connect HSM.

When configuring high availability (HA) for network HSM-enabled GDE Appliance, Thales recommends the following:

- Configure at least two nShield Connect appliances in the Security World for fault tolerance. This means that in the event that one of the appliances is not reachable, the Security World is still available. Refer to the nShield Connect user documentation for a description of procedures to configure an nShield Connect HSM.

NOTE: Client licenses are required for each nShield Connect appliance that is configured for the HA Cluster—the number of client licenses required per Connect appliance is equal to the number of GDE Appliances HA nodes connected to the nShield appliance.

You can configure a network HSM-enabled GDE Appliance HA cluster in one of two ways:

Configure HA with standalone nodes

The first method is to configure all of the GDE Appliances as standalone nodes and enable network HSMs for each of them in the same Security World. You can now create a network HSM-enabled GDE Appliance cluster in the same manner as for any other GDE Appliance cluster.

The high-level steps to configure a network HSM-enabled GDE Appliance HA cluster following this method are:

1. Configure two nShield Connect appliances and the associated RFS.
2. Configure the HA cluster.
3. Add the GDE Appliances individually to the nShield Connect Security World to make each GDE Appliance network HSM-enabled. This means you must run the `connect` add command on each GDE Appliance to add them to that Security World.

   Refer to the nShield user documentation for a description of how to configure and deploy the nShield Connect device and the associated RFS.
4. Add both nShield Connect appliances to each of the GDE Appliances as follows:

   a. Add the nShield Connect appliance to the GDE Appliance. Open a CLI session on the GDE Appliance appliance that is a client of the nShield Connect appliance.

NOTE: If the nShield Connect Security World is FIPS 140-2 level 3 compliant, only one card from the associated ACS is required for this step. The card is only required for the first Connect device to be added to the GDE Appliance, it is not required for any subsequent nShield Connect appliances that are added.

   b. Navigate to the HSM category of commands, type the following at the prompt:

```
0000:dsm$ hsm
```

```
0001:hsm$
```

c. Use the `connect` add command to add the nShield Connect to the GDE Appliance. Type the following command at the prompt,

```
0001:hsm$ connect add <nShield_Connect_IP_Address> <RFS_IP_Address>
```

- **nShield Connect IP Address**: IP address of the nShield Connect appliance
- **RFS IP Address**: IP address of the computer that has the RFS installed

Example:

```
0002: hsm$ connect add 192.168.3.18 192.168.3.4
```

A warning displays, informing you that once this GDE Appliance is converted to a network HSM-enabled appliance, it cannot be rolled back.

d. Type **yes** to continue. The GDE Appliance is restarted if the operation is successful.

e. Follow the prompts to add the nShield Connect appliance to the GDE Appliance.

f. To view the nShield Connect that has been added, type:

```
0002: hsm$ connect show
```

g. If there are more nShield appliances in the same Security World, you can add them now using the `connect` add command.

Configure an HA cluster with HSM-enabled nodes

The second way to create a network HSM-enabled HA cluster is to configure a standalone network HSM-enabled GDE Appliance, as HA node 1. Then join other network HSM-enabled GDE Appliance nodes to the HA cluster.

NOTE: Once a GDE Appliance is network HSM-enabled, it must be connected to at least one nShield Connect appliance. If you remove an nShield appliance from a Security World, you must make sure that any GDE Appliances that were connected to it, are now connected to another nShield appliance belonging to that same Security World.
In this case, if more than one nShield appliance is available in the Security World, a GDE Appliance Administrator could choose to use any of the available nShield appliances after the GDE Appliance has joined the cluster.

See "Configuring High Availability for Network HSM-enabled Nodes" on page 36.

See the *VDS Administrators Guide* for instructions on other HA functions such as:

- Assigning Hosts to the GDE Appliances in an HA cluster
- Pushing Configuration Changes to Hosts

- Reassigning Hosts to Another GDE Appliances in the HA Cluster
- Displaying the High Availability Configuration
- Recovering from incomplete node synchronizations

Adding a Host to a new HA node

If an HA node fails, when the HA node is running again, it will synchronize with the other HA nodes. However, if a host requires a connection to an HA node in the cluster and cannot wait for that node to restart, you can manually move the host to another node. To move the hosts:

1. On the GDE Appliance, click **High Availability**.
2. Click on the Name of the node to which you want to move the Agent(s).
3. Click **Host Assignment**.
4. Click **Add**. The Details page opens and displays all of the Agents connected to the HA cluster.

Figure 3: Host Assignments for HA Server

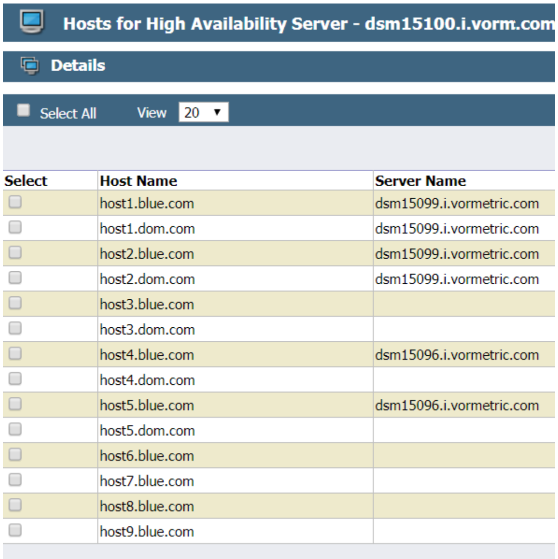

Hosts for High Availability Server - dsm15100.i.vorm.com

Details

□ Select All View 20

Select	Host Name	Server Name
□	host1.blue.com	dsm15099.i.vormetric.com
□	host1.dom.com	dsm15099.i.vormetric.com
□	host2.blue.com	dsm15099.i.vormetric.com
□	host2.dom.com	dsm15099.i.vormetric.com
□	host3.blue.com	
□	host3.dom.com	
□	host4.blue.com	dsm15096.i.vormetric.com
□	host4.dom.com	
□	host5.blue.com	dsm15096.i.vormetric.com
□	host5.dom.com	
□	host6.blue.com	
□	host7.blue.com	
□	host8.blue.com	
□	host9.blue.com	

5. Select and click **OK** for nodes that you want to move to your current node. The GDE AppliancesGDE Appliance moves the selected Agent host from the previous HA node to the current HA node.

NOTE: You can assign both unassigned hosts and hosts currently assigned to other nodes. GDE Applianceswill move those nodes from the previous node to the current node.

6. Once the original HA node is up and running, you can reassign the VTE Agent(s) back to the original node, if desired.

Upgrading an HA Cluster

When upgrading the nodes in an HA cluster, you must break the cluster by removing a node from the cluster, running the HA cleanup function and then upgrading that node independently.

If you are migrating from an HA cluster that is DSM v6.1.x or an earlier version, then you cannot upgrade or migrate your cluster to 6.2.x or 6.3.x. You must create a new cluster.After upgrading all of the HA nodes:

- Add them back into the HA cluster
- Join them to the HA cluster
- Reassign VTE Agents to their nodes

NOTE: Do not reassign hosts in the HA cluster. The host assignment is preserved. When a node is removed from the cluster and then joined back into it, the host(s) are reassigned to the same node.

Prerequisite

- Backup your current GDE Appliances configuration, as described above, “Backup current DSM configuration” on page 111.

NOTE: If synchronization is in progress anywhere in the HA cluster, wait until it completes before upgrading each of the nodes in the cluster.

Remove Nodes from the HA cluster

Breaking up the HA cluster involves removing the nodes from the HA cluster. On HA node 1, the initial node, log in to the CLI menu as CLI Admin.

1. Switch to the HA menu, type:

   ```
   0001:dsm$ ha
   ```

2. Remove the node from the HA menu, type:

   ```
   0002:ha$ remove <FQDN/HA_node_IP>
   ```

   Example

   ```
   0002:dsm$ remove HAnode3.i.vormetric.com
   ```

   System Response:

   ```
   WARNING: This command is going to remove one server node from
   the HA cluster.
   This may take several minutes.
   Continue? (yes|no)[no]:yes
   SUCCESS: Removed server node HAnode3.i.vormetric.com from the
   HA cluster.
   ```

3. Repeat the previous step for all of the nodes in the HA cluster.

4. After removing the nodes, log on to one of the other nodes, (not the initial one) as CLI Admin and switch to the HA menu.

   ```
   0001:dsm$ ha
   ```

5. Cleanup the HA configuration data on the node, type:

   ```
   0002:ha$ cleanup
   ```

   System Response:

   ```
   WARNING: This command cleans up HA configuration data of, and
   restarts, this server.
   This may take several minutes.
   Continue? (yes|no)[no]:
   ```

6. Type **yes** to continue.

   ```
   SUCCESS: cleanup
   ```

7. Repeat the cleanup process for every node that you removed from the HA cluster.

Upgrade the Initial HA node

In the GUI, upgrade the initial node to the latest version of the software.

1. Select **System > Software Upgrade**. The Upgrade Software window opens.

2. If two software images are present, click **Delete Idle Version** to delete the one which is not in use.

3. Click **Browse/Choose File** and select the upgrade file that was provided to you.
4. Click **Open**, and then click **Upgrade** to start the upgrade. Follow the directions on the screen.
5. Refresh your browser to view the login screen after the upgrade completes.
6. For HSMs like a V6100, repeat these steps to upgrade each of the other nodes in the HA cluster.

NOTE: If you plan to enable nShield Connect integration on a GDE appliance HA cluster see, “Configuring High Availability for network HSM-enabled DSM” on page 52.

Optimize the Upgrading of Nodes in the HA Cluster

The initial node is the only node that needs to be upgraded. This ensures that all of the content of the node: policies, admins, domains, keys, reports, logs, etc. will be saved. However, an upgrade is unnecessary for the other nodes in the cluster, because all of the content of the initial node will be copied over from the initial node when all of the nodes in the HA cluster synchronize. Therefore, to optimize upgrading the virtual GDE Appliance for the additional nodes in the HA cluster:

1. Upgrade the initial node to GDE v4.0.0.2.
2. On each additional node, perform a fresh installation of GDE v4.0.0.2. In other words, use the OVA file for the installation, not the upgrade tar file.
3. Once the software upgrade/installation on each of the nodes is complete, add the nodes to the cluster and join the HA cluster. See “Adding Nodes to an HA Cluster” on page 33 and “Join a Node to an HA Cluster” on page 34 for more information.

Deleting a Node from a Cluster

Deleting a node from a cluster with no host assigned to it is straightforward. When deleting nodes with hosts, you must reassign the hosts to other nodes.

Deleting a Node from a Cluster with no Hosts assigned

Log into the Management Console on the Initial Server as a System/All Administrator.

1. Select **High Availability** in the menu bar.
2. Check the box next to the node to be deleted.
3. Click **Delete** to remove the node from the cluster.
4. Log in to the CLI of the HA node that you are removing from the cluster.

5. Switch to the HA menu, type:

```
0001:dsm$ ha
```

6. To disable the communication between the nodes and stop synchronizations to the node, in the HA menu, type:

```
0001:ha$ cleanup
```

Deleting a Node from a Cluster with Hosts assigned

When deleting a node from an HA cluster, you must reassign the hosts in the CLI or in the UI. You can:

- Re-assign hosts from their designated HA node to another HA node.
- Use the `rr` (round robin) option and have the GDE Appliances evenly distribute the hosts to balance the load in the HA cluster.

Moving a Host to a different Node with the CLI

To move the hosts to a specific HA node, type:

```
0001:ha$ remove <node1> reassignhost <node2>
```

Example

```
0001:ha$ remove dsm15099.i.vormetric.com reassignhost
dsm15100.i.vormetric.com
```

To make the GDE Appliances move the hosts to HA nodes and evenly distribute the load, type:

```
0001:ha$ remove <node1> reassignhost rr
```

Example

```
0001:ha$ remove dsm15099.i.vormetric.com reassignhost rr
```

System Response

```
WARNING: This command removes the given server node from the HA
cluster. After deletion, running the "ha cleanup" command from
its CLI will be required.

This may take several minutes.

Continue? (yes|no)[no]:yes

SUCCESS: Removed server node from the HA cluster.
```

Warning! Remember to ALWAYS run `cleanup` on the node that was removed from the HA cluster.

Moving a Host to a different Node with the UI

When you delete a node in the UI that has hosts assigned to it, a dialog opens providing options for host reassignment. You can cancel the delete and manually reassign the hosts yourself, or the GDE Appliances can perform the reassignment.

1. On the GDE Appliance, click **High Availability**.
2. Select the node to which you want to move the Agent(s).
3. Click **Delete**.

   The Delete HA node dialog opens if the node has hosts attached. The options are:

   - Manually assign the hosts to other nodes before deleting <*node*>. Click **Cancel**.

     See “Adding a Host to a new HA node” on page 39.

   - Leave hosts unassigned after deletion of <*node*>. Click **Delete**.

     See “Deleting a Node from a Cluster with Hosts assigned” on page 43.

   - Let GDE Appliances assign hosts to available nodes in the cluster before deleting <*node*>. Click **Delete**.

     The GDE Appliances evenly distributes the hosts to balance the load in the HA cluster.

   - Assign hosts to a specific node before deleting <*node*>. Click **Delete**.

4. When you select this last option, the “Node to assign to” menu opens. Select a node from the dropdown menu.
5. Click Delete.

NOTE: After deletion, make sure that you log on to the deleted node through the CLI menu and run **HA > Cleanup**.

4

Luna SA HSM

This chapter describes how to set up a Luna SA HSM (hardware security module) with a GDE Appliance The purpose of an HSM is to protect sensitive data from being stolen by providing a highly secure operation structure. HSMs are fully contained and complete solutions for cryptographic processing, key generation, and key storage.

An HSM manages cryptographic keys used to lock and unlock access to digitized information over their lifecycle. This includes generation, distribution, rotation, storage, termination, and archival functions. An HSM also engages in cryptographic processing, which produces the dual benefits of isolation and offloading cryptographic processing from application servers.

This chapter contains the following sections:

Add a Luna SA HSM to an HA Cluster

When you add a Luna SA HSM to your GDE Appliance cluster, you must follow the proper order of operations to be successful:

1. "Configuring an HA Cluster with a Luna SA HSMs"
2. "Backup your Configuration"

3. "Break Apart the Cluster"
4. "Add a Luna to the Initial Node of the HA Cluster"
5. "Add Nodes to a Luna-enabled HA Cluster"
6. "Add Nodes to a Luna-enabled HA Cluster"
7. "Join a Node to an HA Cluster"

NOTE: Add one node at a time. **DO NOT** add multiple nodes simultaneously. This applies to the initial node and all subsequent nodes.

Configuring an HA Cluster with a Luna SA HSMs

You must create partitions on the Luna. Each cluster registers to one partition. You can have multiple clusters using the same Luna, or they can use multiple Lunas for increased fault tolerance. When you add the Luna to the DSM, you will have to enter the partition number to which the GDE Applianceswill register.

Configure for Redundancy and Load Balancing

For redundancy and load balancing, you can configure multiple Lunas as an HA group. Then, you can add GDE Appliances to partitions on different Lunas. If one Luna fails, the other Luna, and GDE Appliances added to it, continue to run, thereby reducing downtime and failures.

The following illustration shows the recommended configuration for a GDE Appliance cluster with a Luna HA group:

Figure 1: Recommended Configuration

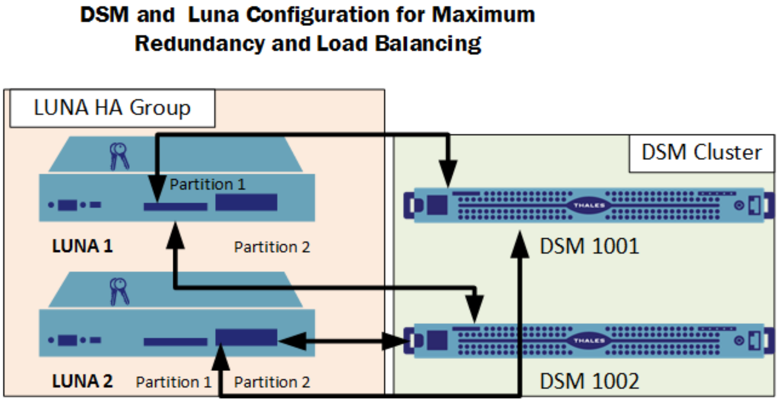

If you are using multiple Lunas with your GDE Appliance cluster, the proper order of operations is:

1. Node 1 (Initial node): Add Luna #1
2. Node 1: Add Luna #2
3. Node 2: Join GDE Appliance Cluster

NOTE: Specify one of the partitions to which Node 1 is registered.

4. Node 2: Add Luna #2

If you are using multiple GDE Appliance clusters, you will make additional partitions on the Lunas and repeat the above steps with the second cluster.

Creating a Partition on the Luna

To create a partition on a Luna:

1. SSH to the Luna. Type:

```
[root@dsm15096 ~]#ssh admin@<LunaHost>
```

Example

```
[root@dsm15096 ~]#ssh admin@192.168.59.214
```

2. Enter the Luna admin password.
3. To view the existing partition list, at the Luna prompt, type:

```
lunash:> partition list
```

Table 1: Initial Partition List

Partition	Name	Objects	Total	Used	Free
1394399181013	Luna1_Par100	0	409782	0	409782
1394399181014	Luna1_Par101	1	409782	200	409582
1394399181015	Luna1_Par102	2	409782	400	409382

4. To create a partition for a GDE Appliance, type:

```
lunash:> partition create -partition <PartitionName>
```

Example

```
lunash:> partition create -partition dsm51005
```

5. Type the following to finish creating the partition:

```
lunash:> proceed
```

If the HSM Administrator is not yet logged in, the following error message displays:

```
Error:  'partition create' failed. (1010000 : LUNA_RET_USER_NOT_LOGGED_IN)
Error:  The HSM Administrator is not logged in.
        The HSM Administrator must be logged in to execute this command.
```

6. At the prompt, type the following:

```
lunash:> hsm login
```

7. Type the HSM Administrators' password.
8. Repeat step 4 - 5 to create the partition.
9. To verify that your partition was created, type:

```
lunash:> partition list
```

Table 2: Initial Partition List

Partition	Name	Objects	Total	Used	Free
1394399181013	Luna1_Par100	0	409782	0	409782
1394399181014	Luna1_Par101	1	409782	200	409582
1394399181015	Luna1_Par102	2	409782	400	409382
1394399181016	DSM51005	0	409782	0	409782

Backup your Configuration

Prior to attaching the Luna to the initial GDE Appliance, you **MUST** perform a complete backup of your system.

You will not have to restore the configurations for any other GDE Appliance in your cluster. When that GDE Appliance synchronizes with the initial node, the initial node overwrites everything so that the cluster nodes are peers.

For more information, see the chapter entitled “Backing Up and Restoring” in the GDE Appliance *Admin* Guide.

Break Apart the Cluster

Before you can add a Luna to a GDE Appliance cluster, you must first break apart the Cluster so that the nodes are not connected and function independently.

To break apart the cluster, see “Upgrading an HA Cluster” on page 40.

Add a Luna to the Initial Node of the HA Cluster

When you add the Luna to a GDE Appliance cluster, you must add it to the initial node of the cluster. After the Luna is setup with the initial node, then you can add/join the remaining nodes to your GDE Appliance cluster.

To add the Luna to a the initial node of your GDE Appliance:

1. Log on to the CLI menu of the GDE Appliance.
2. Change to the HSM menu. At the prompt, type:

   ```
   0001:GDE Appliance$ hsm
   ```

3. In the HSM menu, add the Luna, type:

   ```
   0002:hsm$ luna add <hostName/IP_addr>
   ```

Example

```
0002:hsm$ luna add 192.168.59.214
```

The following warning displays:

```
WARNING: Configuring HSM will restart Security Server service
Continue? (yes|no)[no]:yes
```

Warning! If you have not yet backed up your GDE Appliance, do it now!

4. To proceed, type: **yes**

   The Luna hostname/IP address displays with a fingerprint:

Example

```
192.168.59.214: UZHnfG5tTURxZ8etW0VQHitywmiN5H8NgObKdF20j/M
The fingerprint above should match the RSA output of the
'sysconf fingerprint ssh' lunash command on the Luna SA
```

5. At the prompt, enter the HSM administrator username.
6. At the prompt, enter the HSM administrator password.
7. At the prompt, enter the HSM partition ID number for the partition you created, for registering the GDE Appliance.

NOTE: You can only have one GDE Appliance cluster per partition.

8. Enter the hostname or IP address of this GDE Appliance. A warning displays:

```
WARNING: All Peer node and agent certificates will need to be re-signed after
CA and server certificates are regenerated!
The security server software will be restarted automatically.
```

9. To continue, type: **yes**.
10. Enter the host name of the initial node. If the name is already correct, hit **Enter**.
    - This Security Server host name [GDE Appliance08648.i.thales.com]:
11. Enter the following information for key and certificate generation.
    - What is the name of your organizational unit? []:
    - What is the name of your organization? []:
    - What is the name of your City or Locality? []:
    - What is the name of your State or Province? []:
    - What is your two-letter country code? [US]:
    - What is your email address? []:
12. If you are adding a second Luna to the initial GDE Appliance node, repeat the steps in this section to add it now before adding GDE Appliance nodes to the cluster.

Verifying the Luna status

Confirm that the Luna connected properly.

- Change back to the HSM menu and, type:

  ```
  0001:hsm$ luna show
  ```

```
HA auto recovery:  enabled
              HA recovery mode: activeEnhanced
   Maximum auto recovery retry: 500
   Auto recovery poll interval: 60 seconds
                    HA logging: disabled
            Only Show HA Slots: yes
HA Group Label: g30606
        HA Group Number: 11394399181044
       HA Group Slot ID: 4
       Synchronization: enabled
          Group Members: 1394399181044
             Needs sync: no
        Standby Members: <none>


Slot #    Member S/N                      Member Label    Status
======    ==========                      ============    ======
------  1394399181044                     dsmpartition0     alive

1394399181044@192.168.88.8044
SUCCESS: Showed HSM configuration
```

Add Nodes to a Luna-enabled HA Cluster

Once the Luna is set up with the initial node in your cluster, you can add subsequent nodes to it using the `HA add` and `HA join` commands.

To add subsequent nodes:

1. Log in to the CLI menu of your initial node.
2. Switch to the HA menu, type:

```
0000:GDE Appliance$ ha
```

3. Add the HA node to the cluster, type:

```
0003:ha$ add <haNodeName>
```

Example:

```
0003:ha$ add GDE Appliance95459.i.thales.com
```

System Response:

```
This command adds node '5459.i.thales.com' to the HA cluster
(with 'GDE Appliance08648.i.thales.com'), without configuring it
for replication.

To configure it for replication, login to the CLI of 'GDE
Appliance95459.i.thales.com' and execute the 'join' command under
'ha'

Continue? (yes|no) [no]: yes

SUCCESS: Node added to HA cluster.

0004:ha$
```

Join a Node to an HA Cluster

When executed, the join command first checks to see if there is a Luna connected to the initial node. If it finds one, then it implicitly performs a Luna Add function while joining the HA cluster.

NOTE: Add one node at a time. **DO NOT** add multiple nodes simultaneously. This applies to the initial node and all subsequent nodes.

1. Logon to node B.
2. Switch to the HA menu, type:

```
0000:GDE Appliance$ ha
```

3. Join the node to the cluster. Type:

```
    0000:GDE Appliance$ join [longwait]

WARNING: This server node is about to join an HA cluster.

Please make sure the HA cluster is running and has this server node in its HA
node list.

This may take several minutes. After join, please make sure this server node
is in the same security mode as other nodes.
```

4. To continue, type: **yes**.

5. Enter the name of the HA Initial Server host name to which you are joining this system.

   ```
   Initial_Server=dsm95459.i.thales.com
   CAs_Fingerprint=2F:C3:56:00:22:6D:8C:71:4A:3B:D8:39:09:62:23:18:A0:FF:77:6D
   ```

6. At the prompt, enter the administrator name for the Initial Security Server system, (this is the node that you just connected to the Luna).

7. At the prompt, enter the administrator password for the Initial Security Server system:

8. To continue, type: **yes.**

9. Enter the IP address or host name of the Luna.

   ```
   192.168.59.214: UZHnfG5tTURxZ8etW0VQHitywmiN5H8NgObKdF20j/M

   The fingerprint above should match the RSA output of the

   'sysconf fingerprint ssh' lunash command on the Luna SA
   ```

10. Enter HSM/Luna administrator username.

11. Enter password for HSM/Luna administrator.

12. Enter the HSM/Luna partition ID number to which you are registering the GDE Appliance.

NOTE: Enter the same partition number that you used for the initial node.

13. Enter hostname or IP address of **this** GDE Appliance.

    ```
    This node may have multiple IP addresses. All of the agents will have to
    connect to the Security Server using the same IP.

    Enter the host name of this node. This will be used by Agents to talk to this
    Security Server.
    ```

14. Enter the host name of this node. If the name is already correct, hit **Enter**.

    - This Security Server host name [<*hostName*>]:

15. Enter the following information for key and certificate generation.

    - What is the name of your organizational unit? []:
    - What is the name of your organization? []:
    - What is the name of your City or Locality? []:
    - What is the name of your State or Province? []:
    - What is your two-letter country code? [US]:
    - What is your email address? []:

16. If you are using a second Luna for redundancy, use `HSM > luna` add to add it to the DSM.

Monitoring the Luna

To monitor the Luna:

- Check the status, type:

  ```
  0011:hsm$ status
  ```

HSM Slots

The slot refers to the slot on the Luna SA in which the HSM card is installed. You can monitor the slot as well as the entire Luna SA.

To monitor the HSM slot:

- Type:

  ```
  0011:hsm$ slots
  ```

Upgrading a GDE Appliance attached to a Luna

When you upgrade the initial GDE Appliance that is attached to the Luna HSM, you do not have to detach the GDE Appliance from the Luna. They can stay connected. The upgrade does not affect it at all.

Troubleshooting

Partitions

On the Luna

If you want to change the partition that a GDE Appliance Cluster is registered to, then you need to delete the original partition. At the Luna type:

```
lunash:> partition delete -partition <partitionName>
```

Registering Again

On the GDE Appliance

If you need to register the Cluster to the same Luna SA again, delete the Luna appliance from the GDE Appliance's local config. Type:

```
0001:hsm$ luna delete <hostName/IP address>
```

On the Luna

When you successfully register to a Luna, the Luna adds the GDE Appliance hostname/IP address to its client list. Sometimes a client (GDE Appliance) does not successfully register to the partition, but the Luna shows it as a valid client in the client list. In this situation, delete the client from the client list and register the GDE Appliance again. At the Luna type:

```
lunash:>client delete -client <clientHostName/IP address>
```

Consult the Luna documentation for more information.

Troubleshooting

A

This section describes some troubleshooting procedures for your appliance.

It contains the following sections:

Loss of Connection

If you have created GuardPoints and for some reason the appliance cannot be reached, the GuardPoints will continue to function with no issues. However, if the system is rebooted, the agent cannot access its configuration from the appliance and the GuardPoints cannot use the encryption key to encrypt or decrypt data, unless you are using a cached-on-host key. Challenge and response and manual passwords are good way to provide business continuity in these situations.

Is the Management Console accessible?

1. Try to open a web browser with the correct address to the appliance (example: https://192.168.10.11:8445 or 8448 for Suite B mode).
2. Check if the appliance is a trusted site in your web browser's Security Options.

Check whether Agent communication ports are open from the UI

1. Use the Network Diagnostic checkport tool in the Management Console (or CLI) to check those ports.
2. Refer to the DSM Installation and Configuration guide for information about ports that need to be configured.

Reset DSM Appliance and Remove All Data

The `config reset` command removes all configuration data added after the current DSM software is installed. This command is available on both appliance-based and software-only DSM installations.

The command preserves the currently installed DSM software, but removes all data except network configuration, and in the case of the V6100 appliance, the original security world created using the ACS is also preserved. When you reconfigure the V6100 appliance, you can then recover that security world with the ACS quorum used to create it. Alternatively, you can choose to destroy the old security world and create a new one with new cards.

Caution: If you choose to create a new Security World, we strongly recommend that you use a **new** set of cards (ACS) to create the new Security World. If you reuse the original ACS to create the new Security World, the cards will be overwritten. Any backups created with that original ACS will be **unrecoverable**.

Reset Original Security World with Original ACS Quorum

To reset the current DSM installation to its initial unconfigured state—network configuration remains intact—and retrieve the original security world, do the following;

1. Log on to the DSM CLI console using the CLI Administrator credentials.
2. To enter the maintenance menu, type:

   ```
   $ maintenance
   ```

3. Reset the configuration, type:

   ```
   0001:maintenance$ config reset
   ```

   System Response

   ```
   Reset configuration will wipe out all the configuration data and
   set the configuration data to the manufacture default. System
   will reboot automatically.

   Continue? (yes|no)[no]:yes

   config reset SUCCESS. You can reboot the Security Server now or
   it will reboot automatically in 60 seconds.

   0002:maintenance$
   ```

You will need to generate the DSM certificate authority (CA) again. This will require a quorum from the original ACS used to create the Security World. Wait until the system has rebooted and the `vormetric$` prompt is displayed, then run the `security genca` command.

1. Type `up` at the prompt to return to the main menu.
2. Type `system` to access the System category sub-menu.
3. To generate the CA, type:

   ```
   0004:system$ security genca
   ```

   System Response

   ```
   WARNING: All Agents and Peer node certificates will need to be
   re-signed after CA and server certificate regenerated, and the
   security server software will be restarted automatically!

   Continue? (yes|no)[no]:yes
   ```

4. The following message displays. Read it, enter the required information to generate the CA, and ensure the DSM host name is correct, press enter:

   ```
   This node may have multiple IP addresses. All the agents will
   have to connect to Security Server using same IP.

   Enter the host name of this node. This will be used by Agents to
   talk to this Security Server.

   This Security Server host name[mycompany.com]:

   Please enter the following information for key and certificate
   generation.

   What is the name of your organizational unit? []:

   What is the name of your organization? []:

   What is the name of your City or Locality? []:

   What is the name of your State or Province? []:

   What is your two-letter country code? [US]:

   What is your email address? []:

   Regenerating the CA and server certificates now...
   ```

5. A message displays informing you that a Security World exists and provides you with the option to reuse it. This is the Security World in place when the `config reset` command was run. Type yes to reuse the existing Security World. Keep the quorum of cards from the ACS used to create this original Security World available as these will be required for the next few steps. Follow the instructions on the screen:

   ```
   There is an existing Security World.

   Would you like to reuse it? (yes|no)[no]: yes

   Please provide a quorum of cards from the ACS of the existing
   Security World.
   ```

```
15:47:02 WARNING: Module #1: preemptively erasing module to see
its slots!


Indoctrinating Module:

 Module 1: 0 cards of 1 read

 Module 1 slot 0: empty

 Module 1 slot 2: empty

 Module 1 slot 3: empty

Card reading complete.

security world loaded on 1 module; hknso =
3546197a6456c5e3bfb28d7facd063072b7a8f52

Do NOT remove the smart card from the reader yet.

Creating CA keys and signer certificates...

done.

Generating server private key...

done.

You may now remove the smart card from the reader.

Creating and signing the server certificates...

done.

CA and Server certificates have been generated successfully.

JBoss vault keystore password have been completed successfully.

Self test in progress: passed

SUCCESS: The CA and security certificates are re-generated and
the Security Server software is restarted.
```

The DSM appliance is now ready to use with the original Security World. For procedures to restore a backup of your previous configuration, refer to the *DSM Administrators Guide.*

Create New Security World with New ACS

To reset the current DSM installation to its initial unconfigured state—network configuration remains intact—and create a new Security World, do the following;

1. Log on to the DSM CLI console using the CLI Administrator credentials.
2. Enter the maintenance category of commands

   ```
   0001:dsm $ maintenance
   ```

3. Reset the configuration, type:

   ```
   0001:maintenance$ config reset
   ```

   ```
   Reset configuration will wipe out all the configuration data and
   set the configuration data to the manufacture default. System
   will reboot automatically.

   Continue? (yes|no)[no]:yes

   Config reset SUCCESS. You can reboot the Security Server now or
   it will reboot automatically in 60 seconds.
   ```

4. You will need to generate the DSM certificate authority (CA) again. This will require a quorum from the original ACS used to create the Security World.
5. Wait until the system has rebooted. Type `up` at the prompt to return to the main menu.
6. Type `system` to access the System category sub-menu.
7. To generate the CA, type:

   ```
   0004:system$ security genca
   ```

   System Response

   ```
   WARNING: All Agents and Peer node certificates will need to be
   re-signed after CA and server certificate regenerated, and the
   security server software will be restarted automatically!

   Continue? (yes|no)[no]:yes
   ```

8. Read the message and enter the required information to generate the CA. Ensure the DSM host name is correct, press **Enter**:
9. A message displays informing you that a Security World exists and provides you the option to reuse it. This is the Security World in place at the when the `config reset` command was run. Type **no** to destroy the old Security World and create a new Security World. Have the new set of cards available for this step. For information about the ACS and best practices, refer to the Administrative Card Set section in chapter 1 of the D*SM Installation and Configuration Guide*. Follow the instructions on the screen.

   ```
   There is an existing Security World.

   Would you like to reuse it? (yes|no)[no]: no

   Enter the total number of cards (N) you would like to use in your
   Administrator Card Set (ACS).

   Note: The system can handle at most 64 cards but you should not
   enter more than the number of available cards in your possession
   currently.

   This value must be at least 2 and no higher than 64: 2

   Enter the number of cards (K) required to authorize an action.
   This number K is known as the quorum.
   ```

```
Note: The value for K must be less than N. Creating card sets in
which K is equal to N is not allowed because an error on one card
would render the whole card set unusable.

This value must be at least 1 and less than 2: 1

17:10:29 WARNING: Module #1: preemptively erasing module to see
its slots!

Create Security World:

 Module 1: 0 cards of 2 written

 Module 1 slot 0: empty

 Module 1 slot 2: empty

 Module 1 slot 3: empty

 Module 1 slot 4: empty

 Module 1 slot 5: empty

.
.
.
Module 1 slot 2:- no passphrase specified - overwriting card

 Module 1: 1 card of 2 written

 Module 1 slot 2: remove already-written card #1

 Module 1 slot 2: empty

 Module 1 slot 2: unknown card

 Module 1 slot 2:- no passphrase specified - overwriting card

Card writing complete.

security world generated on module #0; hknso =
781eb7d9ae3abad631bfc4c7487279eadbcef4a8

Do NOT remove the smart card from the reader yet.

Creating CA keys and signer certificates...done.


Generating server private key...

done.

You may now remove the smart card from the reader.

Creating and signing the server certificates...

done.

CA and Server certificates have been generated successfully.
```

```
JBoss vault keystore password have been completed successfully.

Self test in progress: passed

SUCCESS: The CA and security certificates are re-generated and
the Security Server software is restarted.

00053:system$
```

The DSM appliance is now ready to use with the new Security World. For procedures to restore a backup of your previous configuration, refer to the *DSM Administrators Guide*.

Caution: Restoring a backup of the previous configuration will restore the old Security World and the new one just created will be destroyed.

Chassis Issues

If the Chassis intrusion message indicates that the unit's lid is open, try the following solutions:

In the CLI, try to clear the chassis intrusion:

1. Log on to the CLI menu for the DSM.
2. Change to the IPMI menu, type:

   ```
   $ IPMI
   ```

3. Use the Clear IPMI chassis intrusion command, type:

   ```
   $ clearint
   ```

4. If you are logged into the system with root access and you are not in the CLI menu, type:

   ```
   # /usr/sbin/sdt -r "chassis intrusion”
   ```

Alternatively, from the IPMI GUI:

1. Navigate to **server health > sensor reading**.
2. Change the category to **Physical Security**.

Indicator Definitions

Table 3: HSM Status Led Indicators

LED Status	Description
Off	Power off
On, blinks occasionally	Operational mode
Flashes two short pulses, followed by a short pause	Initialization mode
Flashes two long pulses, followed by a pause	Maintenance mode
Flashes SOS, the Morse code distress code (three short pulses, three long pulses, three short pulses). After flashing SOS, the status LED flashes an error code in Morse code	Error mode: For example, the Morse codes for temperature error is as followed. 3 Fast Flashes 3 Slow Flashes 3 Fast Flashes 1 Medium Flash

Boot-Up Messages

If the following message is observed during boot-up, it means the following.

```
i8042: No controller found
```

This message is not an error and can be safely ignored. It is normal during boot up. This occurs because the DSM does not have a PS/2 keyboard attached. This is not a hardware failure so **DOES NOT** require an RMA.

If the following message is observed on boot-up, it means the following:

```
“Some configured disks have been removed from your system, or are no longer
accessible. Please check your cables and also ensure that all disks are
present.

Press any key to continue, or C to load the configuration utility.”
```

Contact Technical Support if you see this message and arrange a RMA.

Boot-up Issues

If the DSM has boot-up issues, capture the boot-up logs through the IMPI port and provide them to the Thales technical support team.

Once the IPMI has been configured on the DSM, there is a feature for recording the boot-up sequence for troubleshooting.

1. Make sure that the iKVM port is enabled.

   You can enable iKVM either through the CLI command or through the IPMI GUI.

2. Open the remote control and you will see the iKVM console.
3. There is a **Record** option to record the virtual screen. The recording length is limited but you can select low video resolution, like 800x600, to increase recording time length. This is very useful for troubleshooting boot failures.

Figure 2: IPMI Control

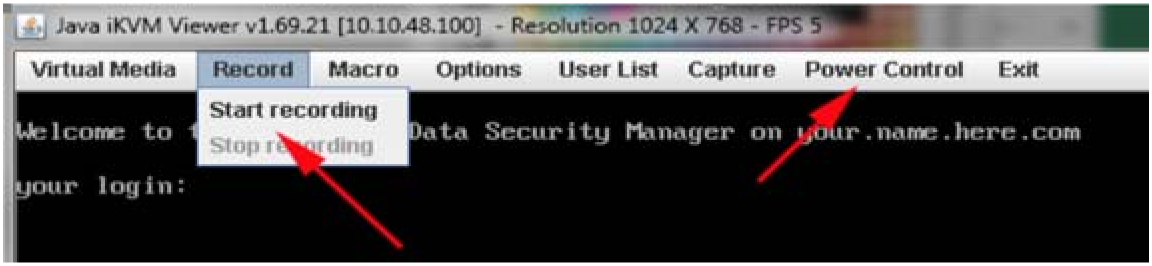

```
Welcome to t                  Data Security Manager on your.name.here.com

your login:
```

B

Ports

This section describes all of the ports that you must configure for your appliance.

Ports to Configure

The following table lists the communication direction and purpose of each port you must open.

Table 4: Ports to Configure

Port	Protocol	Communication Direction	Purpose
	ICMP	All ICMP	Used for Ping
22	TCP	Management Console → DSM	CLI SSH Access
161	TCP/UDP	SNMP Manager → DSM	SNMP queries from an external manager
443	TCP	Browser → DSM DSM↔ DSM Agent → DSM	Redirects to either port 8445 or 8448 depending on the security mode. (8445 is used in compatible & RSA modes; 8448 is used in Suite B mode, for secure communication between DSMs in an HA cluster and for LDT registration.)
5432	TCP	DSM (HA node 1) ↔ DSM(HA node n)	HA information exchange.
5696	TCP	KMIP client → DSM	Allows communication between the KMIP client and DSMs
7025	TCP/UDP	DSM↔ DSM	Uses SNMP to get HA node response time.

8080	TCP	Agent → DSM
		DSM↔ DSM
			Port 8080 is no longer used for registration, but you can manually close/open this legacy port for new deployment, for backward compatibility if you use previous versions of the agent and need to register to 8080. Default is on (open).
			Syntax
			`0001:system$ security legacyregistration`
			`[ on | off | show ]`

Table 4: Ports to Configure

Port	Protocol	Communication Direction	Purpose
8443	TCP	Agent → DSM	RSA TCP/IP port through which the agent communicates with the DSM, in case 8446 is blocked. The agent establishes a secure connection to the DSM, through certificate exchange, using this port.
8444	TCP	Agent → DSM	RSA port via which the Agent log messages are uploaded to DSM, in case 8447 is blocked.
8445	TCP	Browser → DSM DSM↔ DSM(fall back)	Management Console, VMSSC, and fall back for HA communication in case port 8448 is dropped.
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) and RSA
8448	TCP	Browser → DSM DSM ↔ DSM Agent → DSM	GUI Management during enhanced security using Elliptic Curve Cryptography (Suite B). Also for secure communication between DSMs in an HA cluster. Also used for communication between host with LDT host and DSM during Agent registration.
8449	TCP	Smart Card → DSM	Smart card used with RSA mode
8450	TCP	Smart Card → DSM	Smart card used with ECC/compatible mode
9004	TCP	DSM↔ network HSM	DSM communication with nShield Connect and its associated RFS
9005	TCP	DSM↔ remote admin	Used by Remote Administration Service process to accept connections from the Remote Administration Client.

IPMI Ports

The following table lists all of the IPMI ports that you can configure.

Table 5: IPMI Ports

Port	Protocol	Communication Direction	Purpose
80	TCP	Browser → IPMI	This port is disabled by default, for security reasons, and should not be used.

Table 5: IPMI Ports

443	TCP	Browser → IPMI	This port is enabled by default. It is used for the IPMI GUI. If you change the port through which you access IPMI through your browser (“Change the port through which you access IPMI” on page 169).
5900	TCP	Browser → DSM	This port is disabled by default. It is used for remote KVM (Keyboard Video Mouse) management. It should be disabled when not in use for security reasons.
623	UDP	Browser → DSM	This port is disabled by default. Enable only if you want to attach virtual media. It should be disabled when not in use for security reasons.

C

Bonding Driver Modes

This section describes all of the supported Bonding Driver modes.

Bonding Driver Modes to Configure

The following Bonding Driver modes are supported:

Table 6: Bonding Driver Modes

Mode	Name	Description	Load-balancing	Fault tolerance
0	balance-rr	Round-robin policy. Transmit packets in sequential order from the first available through the last. This is the default mode for the bonded NICs.	Yes	Yes
1	active-backup	Active-backup policy: Only one slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch.	No	Yes
2	balance-xor	XOR policy: Transmit based on the selected transmit hash policy. The default policy is a simple [(source MAC address XOR'd with destination MAC address) modulo slave count].	Yes	Yes
3	broadcast	Broadcast policy: transmits everything on all slave interfaces.	No	Yes
4	802.3ad	IEEE 802.3ad Dynamic link aggregation. Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification.	Yes	Yes

Table 6: Bonding Driver Modes

Mode	Name	Description	Load-balancing	Fault tolerance
5	balance-tlb	Adaptive transmit load balancing: channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed relative to the speed) on each slave. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.	Yes	Yes
6	balance-alb	Adaptive load balancing: includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the slaves in the bond such that different peers use different hardware addresses for the server.	Yes	Yes

GLOSSARY

G

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 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.

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.

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 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.

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 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.

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.