THALES

CTE Agent Quick Start Guide

FOR CTE V7.6.0

CTE Agent for AIX Integration Guide

This document covers the following information:

Using CTE with Oracle

This section describes how to install and configure CTE on Oracle RAC ASM, how to use an ASM Filter Driver as well as install and use CTE for AIX with Oracle Automated Storage Management (ASM™) Cluster File System (ACFS™) . It contains the following topics:

CTE on Oracle ACFS Overview

CTE enables data protection of Oracle Automatic Storage Management Cluster File System (Oracle ACFS) on `secvm` volumes as part of the Oracle ASM stack. Oracle ACFS configured with `secvm` block devices is intended for use solely by the Oracle RAC application set to store related Oracle generated data such as:

- Oracle-generated related database files:
    - database datafile
    - control files
    - redo log files

    - archive log files
- Oracle-generated database backup files:
    - hot/cold
    - rman
    - datapump exports
- Oracle-generated database TDE local wallet files

> **Note**
>
> CTE on ACFS only provides encryption. It does not provide access control.

For other files such as manually created shell scripts that require staging in a shared storage device, use other shared storage setups such as Veritas shared storage or share NFS mount.

Oracle RAC	Oracle RAC
Oracle ACFS (File System)	Oracle ADVM (Volume Manager)
Oracle ASM (Storage Manager)	SecVM

On Oracle, ACFS is layered on ASM disks, which in turn are built on `secvm` block devices. `secvm` is a proprietary device driver that supports GuardPoint protection to raw devices. secvm is inserted in between the device driver and the device itself.

Key Managers and SecVM

Server-side administrators must ensure that all `secvm` guards for an Oracle cluster use the same policies for encryption and access control.

Host Groups and Identical Keys and Policies

Thales recommends that you deploy host groups to ensure that identical policies and keys are applied on all nodes of the ACFS cluster. This is faster and less error-prone.

Restrictions and Caveats

- Thales does not support `secfs` layered on ACFS.

- Oracle ACFS encryption in conjunction with `secvm` encryption might impact performance.

Oracle RAC ASM

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

Using CTE with an Oracle RAC ASM

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

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

Caution

If you drop a disk from an ASM diskgroup, then add it back to the diskgroup without cleanly wiping the disk, the ASM diskname will be corrupted. To avoid this problem, clear out the disk before you add it back to diskgroup. Example :

```
dd if=/dev/zero of=/dev/secvm/dev/mapper/asmdg-asmlv002 bs=32k
```

- Important ASM Commands and Concepts
- Determining Best Method for Encrypting Disks
- General Prerequisites
- Specific Prerequisites

Determining Best Method for Encrypting Disks

A diskgroup can contain one or multiple disks. You must determine if the diskgroup contains enough disks and free space for encryption. If the diskgroup contains only one

disk, or multiple disks but not enough free space, then you must use the **Offline** (backup/restore) method for encryption.

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

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

Online Method (No Application / Database Downtime)

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

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

Offline Method (Backup the DB)

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

1. Backup the database.
2. Make an ASM disk available for protection by either removing a disk from an existing diskgroup, or allocating a new disk.
3. Stop the Oracle database.
4. Delete the diskgroup.
5. Apply CTE encryption to the disk.
6. Recreate the diskgroup.

7. Add the protected disk to the diskgroup.
8. Restart the nodes and the failover test.
9. Repeat the previous steps for each disk in the diskgroup.

Important ASM Commands and Concepts

Rebalancing Disks

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

Example ASM Command:

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

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

Mapping Raw Devices

You can map raw devices for this configuration using:

- **EMC PowerPath**

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

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

Checking Rebalance Status

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

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

This command returns information about the:

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

Note

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

Oracle cautions against this issue:

Caution

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

General Prerequisites

Setup

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

Note

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

Altering ASM_DISKSTRING on ASM

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

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

```
SQL> SHOW PARAMETER ASM_DISKSTRING
```

2. To modify the setting, type:

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

Where the path added is the path to SecVM.

Specific Prerequisites

Establishing a Starting Point

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

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

The Importance of Device Mapping

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

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

Important Note about Raw Devices on AIX

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

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

Note

- Attempting to apply a GuardPoint on a character device that does not have a corresponding block device may result in a GuardPoint that never encrypts data. The status of the GuardPoint never shows as guarded.
- The WebUI does not support browsing for the character devices. You would need to manually paste the name into the WebUI.

Once guarded, CTE creates both a character and block mode version of the guarded device. Oracle ASM can use either device.

Creating an Oracle ASM Disk Group for Guarding

Creating a new ASM Disk Group

Note

This document is for Oracle 19c on AIX 7.

1. List the available cluster shared disks, type:

```
# xiv_devlist
```

Response

```
IBM storage devices
-------------------------------------------------------------------
-------------------------------------------------------------
Device       Size (GB)  Paths  Vol
Name                                Vol ID  Storage ID  Storage Ty
pe  Hyper-Scale Mobility
-------------------------------------------------------------------
-------------------------------------------------------------
/dev/hdisk2  51.6       2/2    sjaix81lpar069_sjaix81lpar079_vol_0
01  1053    7825664     XIV            Idle
-------------------------------------------------------------------
-------------------------------------------------------------
/dev/hdisk3  154.9      2/2    sjaix81lpar069_sjaix81lpar079_vol_0
02  1054    7825664     XIV            Idle
-------------------------------------------------------------------
-------------------------------------------------------------
/dev/hdisk4  154.9      2/2    sjaix81lpar069_sjaix81lpar079_vol_0
03  1055    7825664     XIV            Idle
-------------------------------------------------------------------
-------------------------------------------------------------
/dev/hdisk5  154.9      2/2    sjaix81lpar069_sjaix81lpar079_vol_0
04  1056    7825664     XIV            Idle
-------------------------------------------------------------------
-------------------------------------------------------------
/dev/hdisk6  51.6       2/2    sjaix81lpar069_sjaix81lpar079_vol_0
05  1075    7825664     XIV            Idle
-------------------------------------------------------------------
-------------------------------------------------------------
```

2. Ensure that the disks are available, type:

```
# /usr/sbin/lsdev -Cc disk
```

Response

```
hdisk0 Available          Virtual SCSI Disk Drive
hdisk1 Available          Virtual SCSI Disk Drive
hdisk2 Available C5-T1-01 MPIO 2810 XIV Disk
```

```
hdisk3 Available C5-T1-01 MPIO 2810 XIV Disk
hdisk4 Available C5-T1-01 MPIO 2810 XIV Disk
hdisk5 Available C5-T1-01 MPIO 2810 XIV Disk
hdisk6 Available C5-T1-01 MPIO 2810 XIV Disk
```

3. To identify the device names for the physical disks that you want to use, type the following on any node:

```
# /usr/sbin/lspv | grep -i none
```

Response

```
hdisk1          00fa0087e313f5c9                    None
hdisk2          none                                None
hdisk3          none                                None
hdisk4          none                                None
hdisk6          none                                None
```

4. Select available candidate disks for a new ASM disk group. On the Oracle system, type:

```
SYS@+ASM2> COLUMN path format a20

SYS@+ASM2> SELECT name, header_status, path FROM V$ASM_DISK;
```

Response

```
NAME                           HEADER_STATUS PATH
------------------------------ ------------ --------------------
                               FORMER       /dev/rhdisk1
                               FORMER       /dev/rhdisk2
                               CANDIDATE    /dev/rhdisk3
TDE1_0000                      MEMBER       /dev/rhdisk4
GRID_0000                      MEMBER       /dev/rhdisk6
```

5. Prepare the Targeted Disk for CTE and ASM Diskgroup creation:

```
# chown grid:dba /dev/rhdisk3
# chmod 660 /dev/rhdisk3
# dd if=/dev/rhdisk3 of=zzz bs=4k count=1
```

Response

```
1+0 records in.
1+0 records out.
```

6. In CipherTrust Manager, create a key with the following characteristics:
   - Encryption mode: CBC
   - Algorithm: AES
   - Size: 128 or 256
7. Create a CipherTrust Transparent Encryption policy for Oracle on AIX.
   a. Create a CBC key with CBC-AES128 or CBC-AES256.
   b. Create a Security Rule:
      - **Action**: all_ops
      - **Effect**: Audit, Permit, Apply Key
   c. Create a Key Selection Rule:
      - **Key: cte_cbc_aes256_key**
   d. Guard your targeted RAC raw devices so that you can use the secvm disk to create a guarded Oracle RAC ASM or ASMLib disk group.

```
/dev/rhdisk3
```

      - **Type = Raw or Block Device (Auto or Manual Guard)**

   Once you guard your target, CipherTrust Transparent Encryption creates the following:

```
/dev/secvm/dev/rhdisk3
```

8. Install the same version of CipherTrust Transparent Encryption on all nodes in the cluster. To check the version, type:

```
# vmd -v
```

Response

```
Version 7, Service Pack 3
7.3.0.35
2022-09-23 02:08:46 (PDT)
Copyright (c) 2009-2022, Thales Inc.  All rights reserved.
```

9. Guard targeted disk on all cluster nodes.

   a. Check the guard status of the disk on all cluster nodes, type:

```
#  secfsd -status guard
```

Response

```
GuardPoint    Policy                    Type        ConfigState
Status   Reason
----------    ------                    ----        -----------
------   ------
/dev/rhdisk3  encrypt_cbc_aes256_all  rawdevice  guarded
guarded  N/A
```

   b. List the devices, type:

```
# ls -l /dev/secvm/dev/rhdisk3

crw-rw----    1 grid     dba          43,  1 Sep 28 15:21 /dev/
secvm/dev/rhdisk3
```

10. Add the following client settings for both RAC nodes that are set in the $GRID_HOME & $ORACLE_HOME variables.

```
$ echo $GRID_HOME
```

Response

```
/u01/app/19.0.0/grid

$ echo $ORACLE_HOME
```

Response

```
/u01/app/oracle/product/19.0.0/db_1
```

For each node in the cluster, in the client settings, type:

```
|authenticator_euid|/u01/app/19.0.0/grid/bin/grid
|authenticator_euid|/u01/app/19.0.0/grid/bin/orarootagent.bin
|authenticator_euid|/u01/app/oracle/product/19.0.0/db_1/oracle
|authenticator_euid|/u01/app/oracle/product/19.0.0/db_1/bin/oracle
```

> **Note**
>
> This step is optional because it does not effect Oracle behavior. However, without these setting, CipherTrust Transparent Encryption can generate authentication error messages in the CTE log in `/var/log/vormetric`. These errors do not interfere with Oracle functions.

11. Launch the GUI for the GRID `(./asmca)` to create the new CipherTrust Transparent Encryption guarded ASM disk group:
12. Update the discovery path to the following in order for both the baseline and guarded disks to be found:

```
/dev/rhdisk*,/dev/secvm/dev/rhdisk*
```

13. Select targeted `rhdisk3` disk with the guarded path of `secvm`.

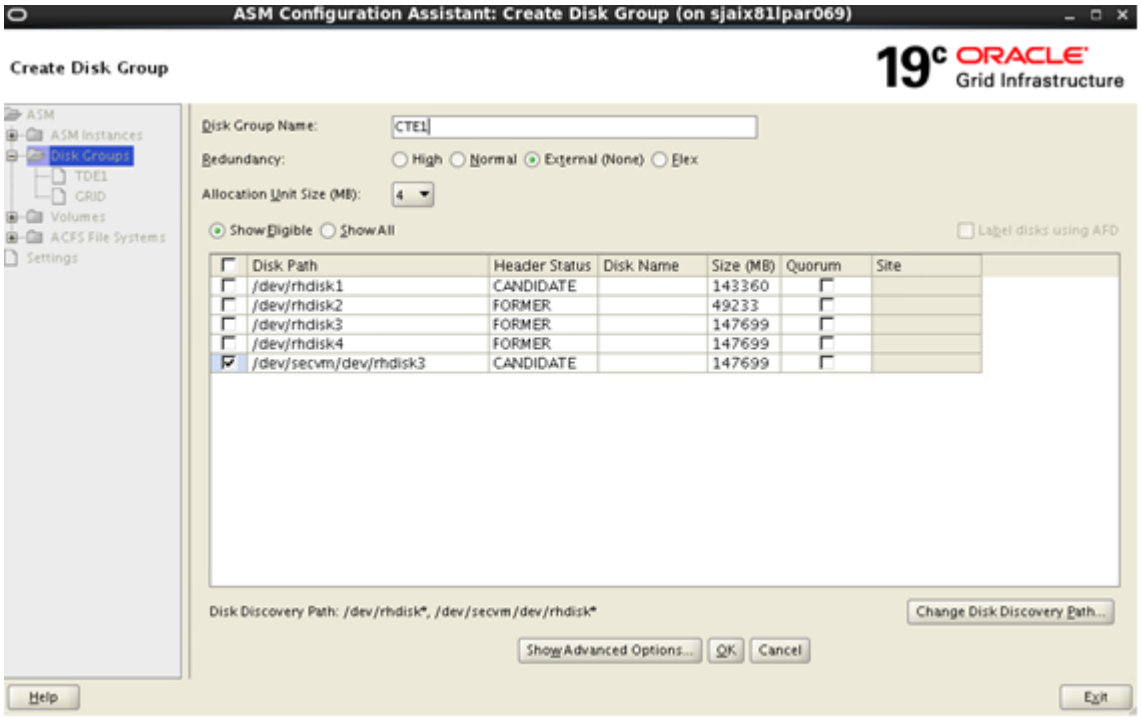

14. The end result should show your new CTE guarded ASM disk group called CTE1.

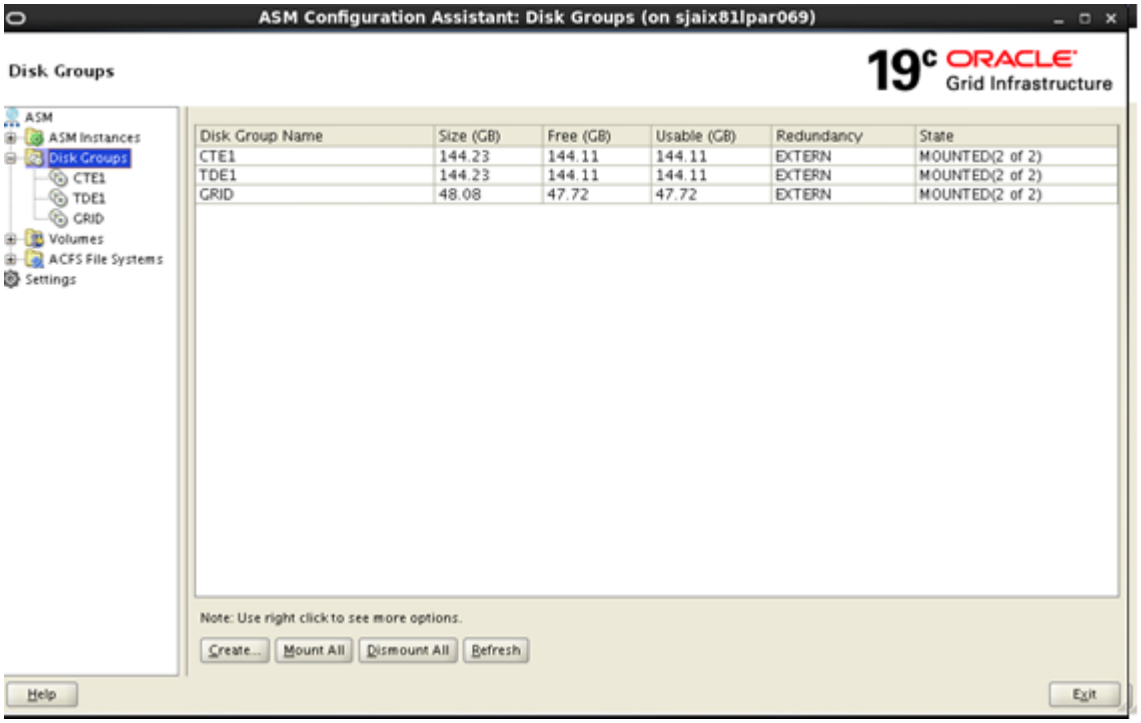

15. When creating your RAC database, choose DB files to reside in the CTE guarded ASM disk group that you just created:

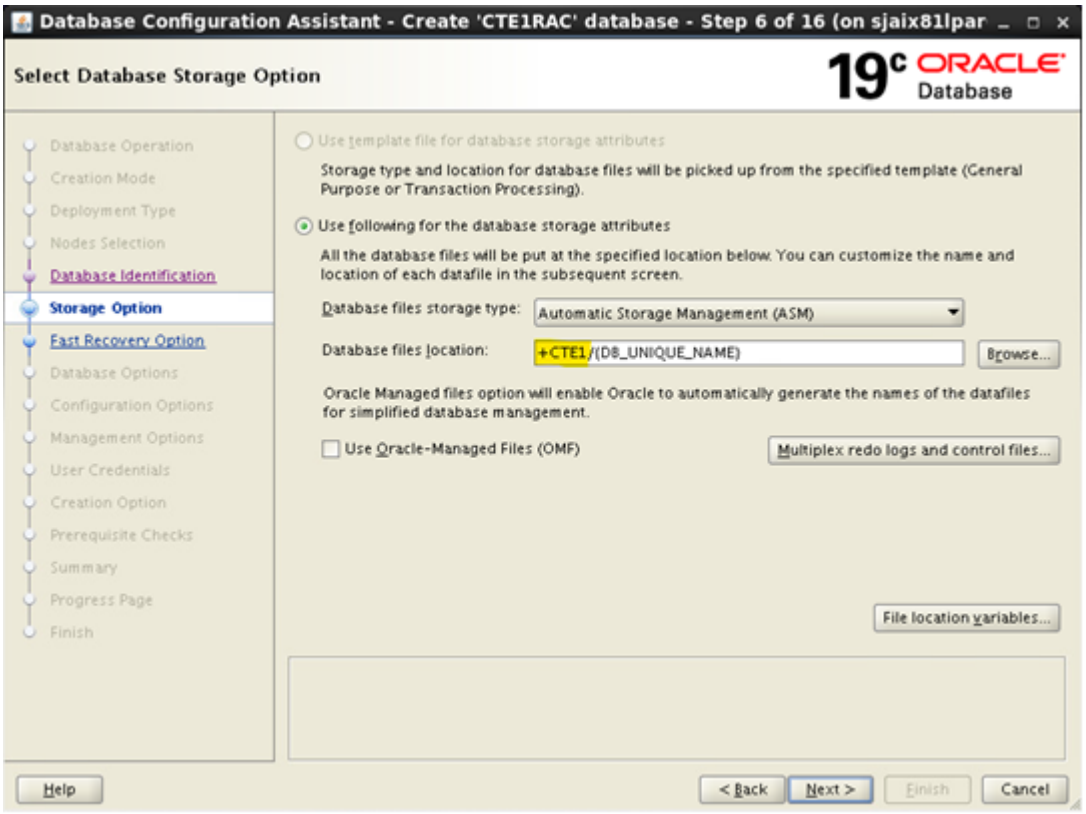

You can now use the secvm disk that you created to create a guarded Oracle RAC ASM or ASMLib disk group.

About Oracle RAC ASM Raw Devices

Standard Devices

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

```
/dev/hdisk1
```

Note

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

Consistent Naming of Devices across RAC Nodes

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

Oracle RAC ASM Multi-Disk Online Method

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

Checking for Space

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

1. Check total free space in the disk group:

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

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

2. Check individual ASM disk size and usage:

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

DISKGROUP  DISK#  DISKNAME        TOTAL_MB    FREE_MB
PATH                           HEADER_STATUS
---------  -----  --------       ---------    -------
```

```
------------------------------  -------------
DATA         0     DATA_0000            1874        1273   /dev/
oracleasm/disks/DATA3      MEMBER
DATA         1     DATA_0001            1992         608   /dev/
oracleasm/disks/DATA4      MEMBER
DATA         3     DATA_0003             117           0   /dev/
oracleasm/disks/DATA2      MEMBER
             0     DATA_ENC_0000         109          28   /dev/
oracleasm/disks/DATA1_ENC  MEMBER
```

Adding a Disk to the Diskgroup

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

To add the disk to the diskgroup:

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

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

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

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

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

```
secfsd -status guard
```

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

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

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

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

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

8. The system is now ready for a reboot and failover test. For details, see Surviving the Reboot and Failover Testing.

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

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

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

```
SQL> DROP DISKGROUP <diskGroupName> FORCE INCLUDING CONTENTS;
```

> **Note**
>
> Make sure that the disk is removed before guarding the raw devices.

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

```
chown oracle:oinstall /dev/<rawDevice1Name>
chmod 660 /dev/<rawDevice1Name>
chown oracle:oinstall /dev/<rawDevice2Name>
chmod 660 /dev/<rawDevice2Name>
```

```
chown oracle:oinstall /dev/<rawDevice3Name>
chmod 660 /dev/<rawDevice3Name>
```

4. On the CipherTrust Manager, in the Host Group, apply GuardPoints to the three raw devices:

```
<rawDeviceName1>
<rawDeviceName2>
<rawDeviceName3>
```

5. On **RAC Node 1**, to display the status of the guarded disks, type:

```
secfsd -status guard
```

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

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

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

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

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

8. On **RAC Node 1**, restore the database.

9. The system is now ready for a reboot and failover test. Go to the Surviving the Reboot and Failover Testing.

Surviving the Reboot and Failover Testing

Failover Testing

Confirm that everything is functional:

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

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

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

Basic Troubleshooting Techniques

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

If you encountering errors similar to:

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

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

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

1. On the CipherTrust Manager, in the Host Group that was created for the RAC cluster, verify that the host group for this configuration does **NOT** have the Cluster Group option set (this option is only for GPFS, which is not supported with CTE).
2. Ensure that the GuardPoints for the block devices are set at the Host Group level. This ensures that each node receives identical GuardPoints.

3. Verify that the GuardPoints are active on all nodes. When the GuardPoints are set, go to each node and verify that they are set and guarded, using the WebUI or the `secfsd -status guard` command. If they do not guard correctly, make sure the device names are the same across all nodes.

4. From ASM, make sure that the `asm_diskstring` parameter is modified to include the CTE devices and that the proper pathname is used, see Altering ASM_DISKSTRING on ASM.

Verifying Database Encryption

Option 1

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

```
/app/oracle/grid/product/19.0.0/grid/bin/kfed read /dev/<diskName>
```

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

```
/app/oracle/grid/product/19.0.0/grid/bin/kfed read /dev/<diskName>
```

System Response:

```
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          242 ; 0x001: 0xf2
kfbh.type:                          124 ; 0x002: *** Unknown Enum ***
kfbh.datfmt:                         66 ; 0x003: 0x42
kfbh.block.blk:              1088904227 ; 0x004: blk=1088904227
kfbh.block.obj:              1558192170 ; 0x008: file=8234
kfbh.check:                  3321251423 ; 0x00c: 0xc5f6465f
kfbh.fcn.base:                932956641 ; 0x010: 0x379bc9e1
kfbh.fcn.wrap:               3040493590 ; 0x014: 0xb53a4016
```

```
kfbh.spare1:                 3806015223 ; 0x018: 0xe2db2ef7
kfbh.spare2:                 3794962182 ; 0x01c: 0xe2328706
60000000000D8000 01F27C42 40E75C23 5CE0202A C5F6465F
[..|B@.\#\. *..F_]
60000000000D8010 379BC9E1 B53A4016 E2DB2EF7 E2328706  [7....:@......
2..]
60000000000D8020 CA2F30AD 522B4D21 99292639 004EBB34  [./0.R+M!.)&9.N.
4]
60000000000D8030 A3896BE8 BD839D23 2204E19E 946C575C
[..k....#"....lW\]
60000000000D8040 4CE2218F 35E1B101 AF751A70 780E6D6E  [L.!.
5....u.px.mn]
60000000000D8050 5E7E6A38 C600ED5F 929047C4 DF372A8E  [^~j8..._..G..
7*.]
60000000000D8060 E103152D BA87CC03 11A7D963 9D72FCE1
[...-.......c.r..]
60000000000D8070 1EC6B48B 03EE869F 61D651F9 E7614957
[........a.Q..aIW]
60000000000D8080 810E0353 9C461F49 69569733 501D19EF  [...S.F.IiV.
3P...]
60000000000D8090 B268002B 4F9457B6 BDB04AC5 D3D07446  [.h.
+O.W...J...tF]
60000000000D80A0 FD9EE5E0 9B46CB66 30D10B22 F63AB77E
[.....F.f0.."
.:.~]
60000000000D80B0 6FF79075 4BBD1FAD 8F226188 7774300D
[o..uK...."a.wt0.]
60000000000D80C0 A809B6FB E1F1C80B B5760E68 9747D97D
[.........v.h.G.}]
KFED-00322: Invalid content encountered during block traversal: [kfbtT
raverseBlock][Invalid OSM block type][][124]
```

Option 2

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

For example:

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

Option 3

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

> **Note**
>
> The strings command cannot read a busy or large device.

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

Using CTE with PowerTech Antivirus

When Powertech Antivirus is configured to use `on-access scanning,` CipherTrust Transparent Encryption Data Transformation may find files within the Data Transformation GuardPoint that are busy. These files are not processed by Data Transformation and are left unencrypted.

> **Note**
>
> This problem only occurs if Powertech is configured to use `on-access scanning.` This problem **does not** occur when using `on-demand scanning.`

Issue

When Powertech uses `on-access scanning,` it opens a file and keeps it open. Before the Data Transformation program encrypts a file, it checks to see if the file is in use by another process. If the file is in use, Data Transformation does not encrypt the file but rather appends the path to the file in the `dataxform_status_error-<guardpoint>` file

located in `/var/log/vormetric/`. This prevents any retry of Data Transformation from succeeding.

Resolving PowerTech/Data Transformation "in-use" Errors

For the following scenario, Powertech is configured with `on_access scanning` and both Powertech and the Thales CTE agent have been installed at default locations. The example assumes the Data Transformation GuardPoint is `/test-dataxform`.

Running Data Transformation with Powertech Antivirus code deactivated

1. Type the following to prevent the Powerteck package from starting on system boot.

   ```
   /opt/sgav/avsvcctl disable
   ```

2. Reboot the system.
3. After the system restarts, verify that Powertech is not running, type:

   ```
   /opt/sgav/avsvcctl status
   ```

   Response

   ```
   The following output indicates Powertech is not running:
   Subsystem Group PID Status
   avsvc powertech inoperative
   Device driver is not loaded
   ```

4. Perform a Data Transformation conversion on the Data Transformation GuardPoints.
   Example

   ```
   dataxform --rekey --gp /test-dataxform
   ```

5. Enable Powertech to start after system boot, type:

```
/opt/sgav/avsvcctl enable
```

6. Start Powertech, type:

```
/opt/sgav/avsvcctl start
```

Response

```
There is a delay before Powertech is completely loaded and
active. The
Powertech status may be found by the following command.
```

7. Verify the status, type:

```
/opt/sgav/avsvcctl status
```

Response

```
The Powertech avsvc subsystem should be `active' and the device
driver loaded:
Subsystem Group PID Status
avsvc powertech 5964286 active
Device driver is loaded
The avsvc subsystem is configured to run at boot
```

Recovering from Data Transformation/Powertech Antivirus code errors

A Data Transformation rekey performed while Powertech is active can result in failed file conversion due to "busy" faults. The following section describes how to recover from this type of error.

Note

There are many reasons that a data transformation may fail. This topic only considers the failure due to the interaction of Powertech with CTE. For other error recovery procedures, refer to the Data Transformation guide

The following section illustrates the process using an example data transformation and describes the steps needed to recover. The Data Transformation GuardPoint `/test-dataxform`.

1. Run Data Transformation.

```
dataxform --rekey --gp /test-dataxform
```

Response

```
Checking if /test-dataxform is a guardpoint with a rekey policy a
pplied
/test-dataxform is a guardpoint with a rekey policy applied
About to perform the requested data transform operation
- Be sure to back up your data
- Please do not attempt to terminate the application
Do you wish to continue (y/n)?y
Scan found 19 files (14 MB) in 1 directories for guardpoint /test-
dataxform
File /test-dataxform/xab is busy
File /test-dataxform/xaa is busy
Transformed 17 files (13 MB) of 19 files (14 MB) for guardpoint /t
est-dataxform
Data transform got errors on some files
File /test-dataxform/xab is busy
File /test-dataxform/xaa is busy
Number of files in error due to file being busy: 2
The data transform operation took 0 hours, 0 minutes and 3 seconds
The data transform program ran from Thu Feb 16 10:58:12 2023
until Thu Feb 16
10:58:15 2023
Data transform for guardpoint /test-dataxform finished but 2
```

```
files were not
processed due to errors
```

2. Inspect the error file (`dataxform_status_error-_test-dataxform`) for `/test-dataxform`. The file is located in `/var/log/vormetric`. The contents for the file contain the error messages generated by Data Transformation.

   Example

```
Skipped, file is busy : /test-dataxform/xab
Skipped, file is busy : /test-dataxform/xaa
```

> **Note**
>
> Other Data Transformation errors may be present besides those caused by the Powertech antivirus code.

3. Create a "todo" file. Using Data Transformation, extract the names of the files that need to be retried for conversion.

```
dataxform --recovery --file_list my-output --gp /test-dataxform
```

   This generates two files in the local directory

```
my-output_done
my-output_todo
```

   The `my-output_todo` file contains the list of files to retry applying Data Transformation:

```
/test-dataxform/xaa
/test-dataxform/xab
```

> **Note**
>
> The `my-output_todo` file may need to be edited to remove entries not caused by the `Skipped, file is busy` error.

4. Turn off Powertech on boot, type:

```
/opt/sgav/avsvcctl disable
```

5. Reboot the system.

6. Verify the status.

```
/opt/sgav/avsvcctl status
```

Response

```
Subsystem Group PID Status
avsvc powertech inoperative
Device driver is not loaded
```

7. Re-issue Data Transformation to transform files listed in the `my-output_todo` file.

```
dataxform --rekey_list --gp /test-dataxform --file_list ./my-
output_todo
```

Response

```
Checking if /test-dataxform is a GuardPoint with a rekey policy
applied
/test-dataxform is a GuardPoint with a rekey policy applied
Previous status information does not relate to a --rekey_file
operation.
Number of files previously in error due to file being busy: 2
About to perform the requested data transform operation
- Be sure to back up your data
- Please do not attempt to terminate the application
Do you wish to continue (y/n)?y
Starting data transform of /test-dataxform for files listed in ./
my-output_todo
The data transform operation took 0 hours, 0 minutes and 2 seconds
The data transform program ran from Thu Feb 16 12:33:47 2023 until
```

```
Thu Feb 16
12:33:49 2023
```

8. Re-enable Powertech on boot.

```
/opt/sgav/avsvcctl enable
```

9. Start Powertech.

```
/opt/sgav/avsvcctl start
```

There is a delay before Powertech is completely loaded and active. The Powertech status may be found by the following command:

```
/opt/sgav/avsvcctl status
```

The Powertech avsvc subsystem should be `active' and the device driver loaded:

Response

```
Subsystem Group PID Status
avsvc powertech 5964286 active
Device driver is loaded
The avsvc subsystem is configured to run at boot
```

Support Contacts

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

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

Customer Support Portal

The Customer Support Portal, at Thales Customer Support, is where you can find solutions for most common problems. The Customer Support Portal is a comprehensive, fully searchable database of support resources, including software and firmware downloads, release notes listing known problems and workarounds, a knowledge base, FAQs, product documentation, technical notes, and more. You can also use the portal to create and manage support cases.

Tip

You require an account to access the Customer Support Portal. To create a new account, go to the portal and click on the REGISTER link.

Telephone Support

If you have an urgent problem, or cannot access the Customer Support Portal, you can contact Thales Customer Support by telephone at +1 410-931-7520. Additional local telephone support numbers are listed on the support portal.

Email Support

You can also contact technical support by email at technical.support@Thales.com.