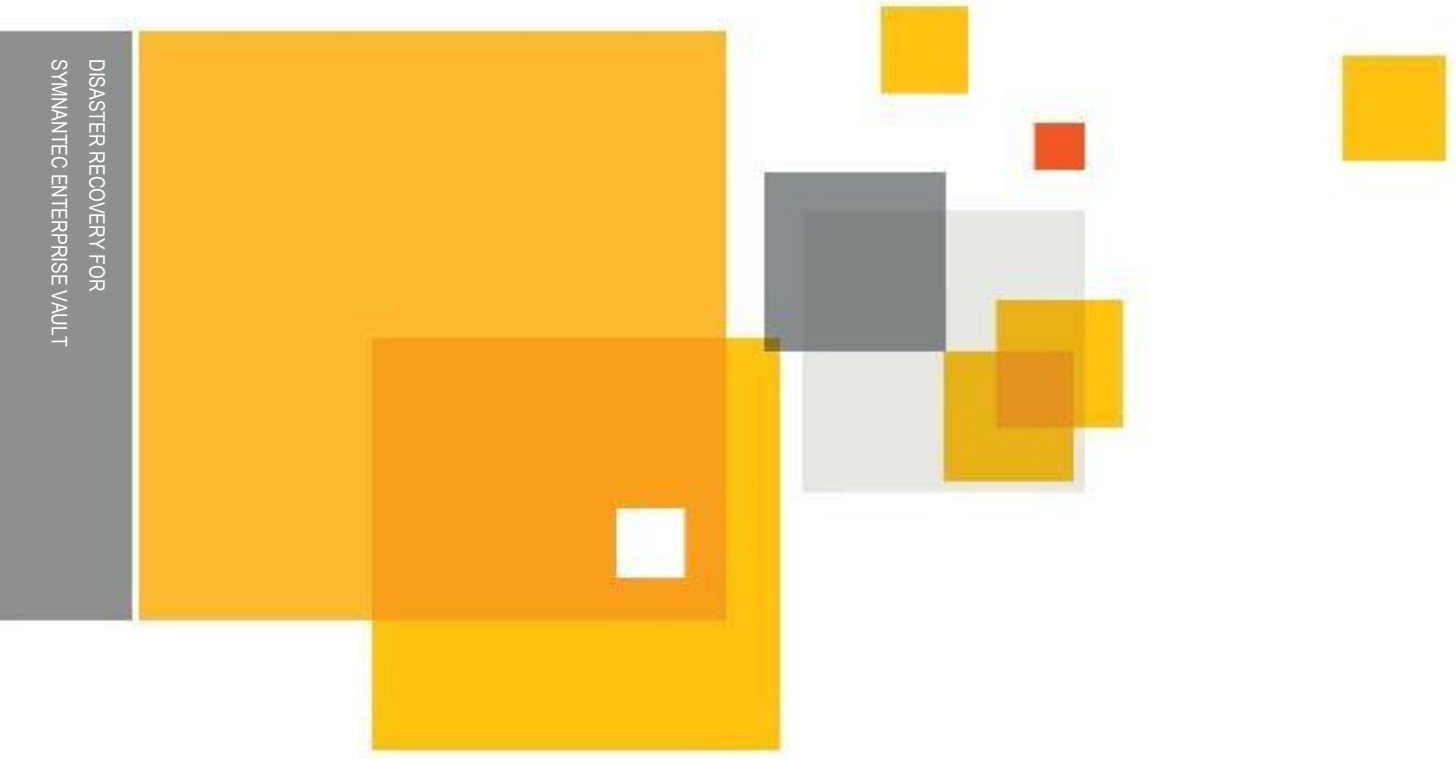

Disaster Recovery for Symantec Enterprise Vault™

Who should read this paper

This Whitepaper is intended to assist customers, partners and service providers prepare for a disaster recovery situation with Enterprise Vault.

If you have any feedback or questions about this document please email them to **iig-tfe@symantec.com** stating the document title.

This document applies to the following version(s) of Enterprise Vault:

10.0.4 and later

Table of Contents

Scope of Document 1
Intended Audience 1
Terminology Used In This Document 1
Executive Summary 2
Disaster Recovery Scenarios 2
    Single Server failure 2
    Enterprise Vault Storage failure 5
    SQL Server failure 7
Datacenter Redundancy 8
    Data Replication 10
    Enterprise Vault Backup Mode 10
    Third Party Storage Replication Software 11
    Standby EV Server 11
    Configuring Enterprise Vault with VCS Global Cluster Option 12
Step by Step Disaster Recovery of Enterprise Vault 13
    Recovery procedure 1: Installing software on the servers 13
    Recovery procedure 2: Restoring Enterprise Vault databases 14
    Recovery procedure 3: Renaming EV servers 15
    Recovery procedure 4: Copy or move the Enterprise Vault data files 16
    Recovery procedure 5: Clearing the directory database entries 16
    Recovery procedure 6: Recreating services and tasks on the Directory service computer 17
    Recovery procedure 7: Recreating services and tasks on Enterprise Vault servers 18
    Recovery procedure 8: Checking the Web Access application settings 19
    Recovery procedure 9: Checking registry entries 20
Steps to recover standalone DA and CA server to a new machine 21
Conclusion 23

Appendices

APPENDIX A – Disaster Recovery Checklist

Document Control

Contributors

Who	Contribution
Dan Strydom	Author

Revision History

Version	Date	Changes
1.0	August 2013	Published

Related Documents

Document Title	Document Location	Version
High Availability Options for Enterprise Vault	www.symantec.com/docs/TECH210099	1.0
Symantec Enterprise Vault Administrator's Guide	www.symantec.com/docs/DOC2200	10.0.4
Enterprise Vault 10 Performance Guide	www.symantec.com/docs/DOC4553	10.0.4
Enterprise Vault Compatibility List	www.symantec.com/docs/TECH38537	10.0.4

Scope of Document

This document aims to provide guidance on disaster recovery for Enterprise Vault. It is recommended that this paper is read in conjunction with the "High Availability Options for Enterprise Vault" whitepaper (referenced documents), as both documents contain information relevant to formulating a business continuity plan.

Intended Audience

This document is aimed at Enterprise Vault customers, partners and service providers. It is assumed that the reader has a thorough understanding of the architecture and operational aspects of Enterprise Vault 10.

Terminology Used In This Document

Term	Description
DR	Disaster Recovery
HA	High Availability
RTO	Recovery Time Objective, refers to the amount of time an application is not available
RPO	Recovery Point Objective, refers to the amount of data lost
USL	Update Service Locations, a built-in Enterprise Vault High Availability feature

Executive Summary

When forming a disaster recovery plan, the two main factors that need to be considered are the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). The RTO refers to the amount of time the application is down, while the RPO refers to the amount of data lost. Organizations should have an RTO (per application) that must be satisfied as well as an RPO that must be met. Most organizations tend to focus on the RTO, or how much downtime is acceptable. However, the amount of data loss an organization can tolerate is just as important—just a few minutes' worth of lost transactions can have a far-reaching negative impact on the business.

Many companies today still rely primarily on tape backup and restoration as the centerpiece of their disaster recovery plan. This usually means at least a day of lost data and a few days of downtime after a disaster. Most organizations have applications that require an RTO that is measured in minutes or seconds, and absolutely no tolerance for data loss.

A comprehensive disaster recovery plan should include data replication or continuous data protection, application clusters, and traditional tape backups. Replication or continuous data protection of the data to a remote site eliminates data loss, while clustering can manage the local and remote failover of applications to minimize downtime. The integration of replication with application clustering, supplemented with tape backups, can satisfy even the most stringent recovery time and recovery point requirements.

Disaster Recovery Scenarios

For every application there are a number of different components that can fail and cause an outage – whether it is shared services such as the network or storage, or application components such as services or the physical server itself. The following sections detail a number of different scenarios and the recommended recovery actions.

Single Server failure

The best way to protect against a single EV server failure is to use one of the high availability options described in the whitepaper "High Availability Options for Enterprise Vault" (referenced documents). Using clustering technologies will protect against server hardware failure, Operating System failure or core server component failure (such as MSMQ, IIS, etc).

Table 1 details the impact of losing a single EV server together with the required action to recover from the outage.

Server Role	Impact	Mitigating Steps	Disaster Recovery
EV Journal Archiving Server	Messages would build up in the Exchange Journal mailboxes. Discovery search/export against indexes held on the EV Journal server will not be available.	Protect against this outage by using clustered EV Journal Archiving Servers. Protect the Exchange Server against this outage by scoping Exchange Journal mailboxes to hold at least 3 days' worth of journal items.	If clustered perform failover. For stand-alone servers perform a building block failover as described in the High Availability paper. For full server recovery perform the server DR procedure outlined in this document.
EV Mailbox Archiving Server	Users will be unable to search and retrieve messages. Vault Cache users enabled to download the full content of archived items will be able to perform an Outlook search, and retrieve items from their cache. No new items will be archived by the Mailbox Archiving task on the failed server until the service is restored.	Protect against this outage by using clustered EV Mailbox Archiving Servers. Enable Vault Cache for end users to cache archived items locally.	If clustered perform failover. For stand-alone servers perform a building block failover as described in the High Availability paper. For full server recovery perform the server DR procedure outlined in this document.
EV File Archiving Server	Users would be unable to retrieve archived placeholders.	Protect against this outage by using clustered EV File Archiving Servers.	If clustered perform failover. For stand-alone servers perform a building block failover as described in the High Availability paper. For full server recovery perform the server DR procedure outlined in this document.

Server Role	Impact	Mitigating Steps	Disaster Recovery
EV SharePoint Archiving Server	Users would be unable to retrieve archived SharePoint items, and SharePoint add-in EV search will be unavailable.	Protect against this outage by using clustered EV SharePoint Archiving Servers.	If clustered perform failover. For stand-alone servers perform a building block failover as described in the High Availability paper. For full server recovery perform the DR procedure outlined in this document.
EV DA / CA Server	DA / CA end user search/export will be unavailable.	For protection against hardware failure deploy DA and CA servers in a VMware or Hyper-V environment, where the Hypervisor is able to fail the application over to a different physical machine. DA and CA can also be deployed with load balancing technology (refer to the High Availability for Enterprise Vault whitepaper).	Perform “Steps to Recover Standalone DA or CA to a new machine” section outlined later in this document.

Table 1 – Failure of a single EV server

Enterprise Vault Storage failure

Table 2 details the impact of losing an Enterprise Vault storage component. Highly available storage should be provided where possible to avoid outages and data loss.

Storage object	Impact	Disaster Recovery
Loss of EV Journal Index Storage	DA/CA users will not be able to search the Journal Index Volume. EV Journal task should be stopped until the issue is resolved. In an extended outage it is recommended that an alternative EV server is used to journal messages to prevent items building up in the Exchange journal mailbox.	Restore index volumes from backup. Run Verify operation on index volumes following restore. Index may require synchronize operation to ensure it is up to date. If backup is not available the index volumes can be rebuilt using the vault store data. **Important Note:** Large journal index volumes can take an extensive amount of time to rebuild. This process should not be considered without the assistance of a professional services partner. For more index operations including troubleshooting refer to **www.symantec.com/docs/HOWTO58947**
Loss of EV Mailbox Index Storage	End users will be unable to search archived items. Retrieval of items from Outlook will continue to work. Vault Cache will not synchronize until the index volume is available. Mailbox archiving schedules and manual archiving should be disabled until the issue is resolved.	Restore index volumes from backup. Run Verify operation on index volumes following restore. Index may require synchronize operation to ensure it is up to date. If backup is not available the index volumes can be rebuilt using the vault store data. Although mailbox volumes are typically much smaller than journal volumes, the rebuild process is still likely to take a long time depending on the number and size of volumes. For more index operations including troubleshooting refer to **www.symantec.com/docs/HOWTO58947**
Loss of EV File or SharePoint Index Storage	Users will be unable to search archived items through browser search, Archive Explorer or EV Web search in SharePoint. Retrieval of items will continue to work. Archiving schedules should be disabled until the issue is resolved.	Restore index volumes from backup. Run Verify operation on index volumes following restore. Index may require synchronize operation to ensure it is up to date. If backup is not available the index volumes can be rebuilt using the vault store data – searching will not be available until the rebuild process is completed.

Storage object	Impact	Disaster Recovery
Loss of EV Journal Vault Store Partition Storage	DA/CA users will be able to search the Journal Index volumes up to the point of when the failure occurred. No DA or CA export operation will be possible. No EV Journal archiving operations will be possible, and items will queue in Exchange Journal mailboxes. In an extended outage an alternative EV Journal Archiving server can be used to archive the journal items to an alternative or new vault store.	Restore vault store partitions from backup, using the same folder paths where possible. If restoring to a different path the Vault Store Partition entry table in the EV Directory database will need to be updated. Run EVSVR to validate items in vault store partitions against SQL Vault Store databases. For more information on EVSVR see **www.symantec.com/docs/HOWTO37650**
Loss of EV Mailbox Vault Store Partition Storage	Users will be unable to retrieve or restore archived items. End user search operations will continue to work, and HTML versions of items will still be available through search applications. Users will not be retrieve any archived item from Outlook unless it is stored in their local Vault Cache. Mailbox archiving schedules and manual archiving should be disabled until the issue is resolved.	Restore vault store partitions from backup, using the same folder paths where possible. If restoring to a different path the Vault Store Partition entry table in the EV Directory database will need to be updated. Run EVSVR to validate items in vault store partitions against SQL Vault Store databases. For more information on EVSVR see **www.symantec.com/docs/HOWTO37650**

Storage object	Impact	Disaster Recovery
Loss of EV File or SharePoint Vault Store Partition Storage	Users will be unable to retrieve or restore archived items. End user search operations will continue to work, and HTML versions of items will still be available through search applications. File/SharePoint archiving schedules should be disabled until the issue is resolved.	Restore vault store partitions from backup, using the same folder paths where possible. If restoring to a different path the Vault Store Partition entry table in the EV Directory database will need to be updated. Run EVSVR to validate items in vault store partitions against SQL Vault Store databases. For more information on EVSVR see **www.symantec.com/docs/HOWTO37650**

Table 2 – Failure of Enterprise Vault Storage component

SQL Server failure

SQL Server is absolutely essential to the operation of Enterprise Vault – any disruption to SQL Server will cause a complete outage to Enterprise Vault. For a list of high availability and recovery options for SQL Server refer to **http://technet.microsoft.com/en-us/library/ms190202.aspx**.

Note that at the time of publication Enterprise Vault 10 did not provide support for the AlwaysOn feature - please check the Enterprise Vault Compatibility List (Referenced Documents) for the latest updates.

Table 3 summarizes the impact of losing SQL services and recommended recovery actions.

Database	Impact	Disaster Recovery
Loss of all EV Databases (Directory, Vault Store, Fingerprint, DA)	All EV functionality will cease to work.	Recover from Log shipping or equivalent DR function. Alternatively restore from backup. Use EVSVR to sync and re-create SQL entries for data archived since last successful backup. For more information on EVSVR see **www.symantec.com/docs/HOWTO37650**

Database	Impact	Disaster Recovery
Loss or corruption of a single EV database	If the Directory database is unavailable all EV functionality will be disabled. If the Fingerprint or Vault Store database is corrupted then users associated with that database will be impacted. In the event of a DA or CA database corruption the DA/CA application will not be available.	Recover from Log shipping or equivalent DR function. Alternatively restore from backup. Use EVSVR to sync and re-create SQL entries for data archived since last successful backup. For more information on EVSVR see **www.symantec.com/docs/HOWTO37650**

Table 3 – Loss of SQL database

Datacenter Redundancy

It is not uncommon for Enterprise Vault to be deployed in environments with multiple datacenters. Depending on the network bandwidth and latency between the hosting locations, it may be possible to configure Enterprise Vault in an Active/Active site configuration. In this scenario both sites hosts EV servers actively archiving from the targets within the local datacenter, and replicating archive and index data in both directions so either site can host the other site's servers/services.

A more common scenario however is where Enterprise Vault is configured in an Active/Passive or Primary/Secondary datacenter configuration, where all Enterprise Vault data is replicated to a secondary location in order to provide redundancy in the event of an outage to the primary site.

Figure 1 shows an example of an Active/Passive site configuration for Enterprise Vault Mailbox Archiving.

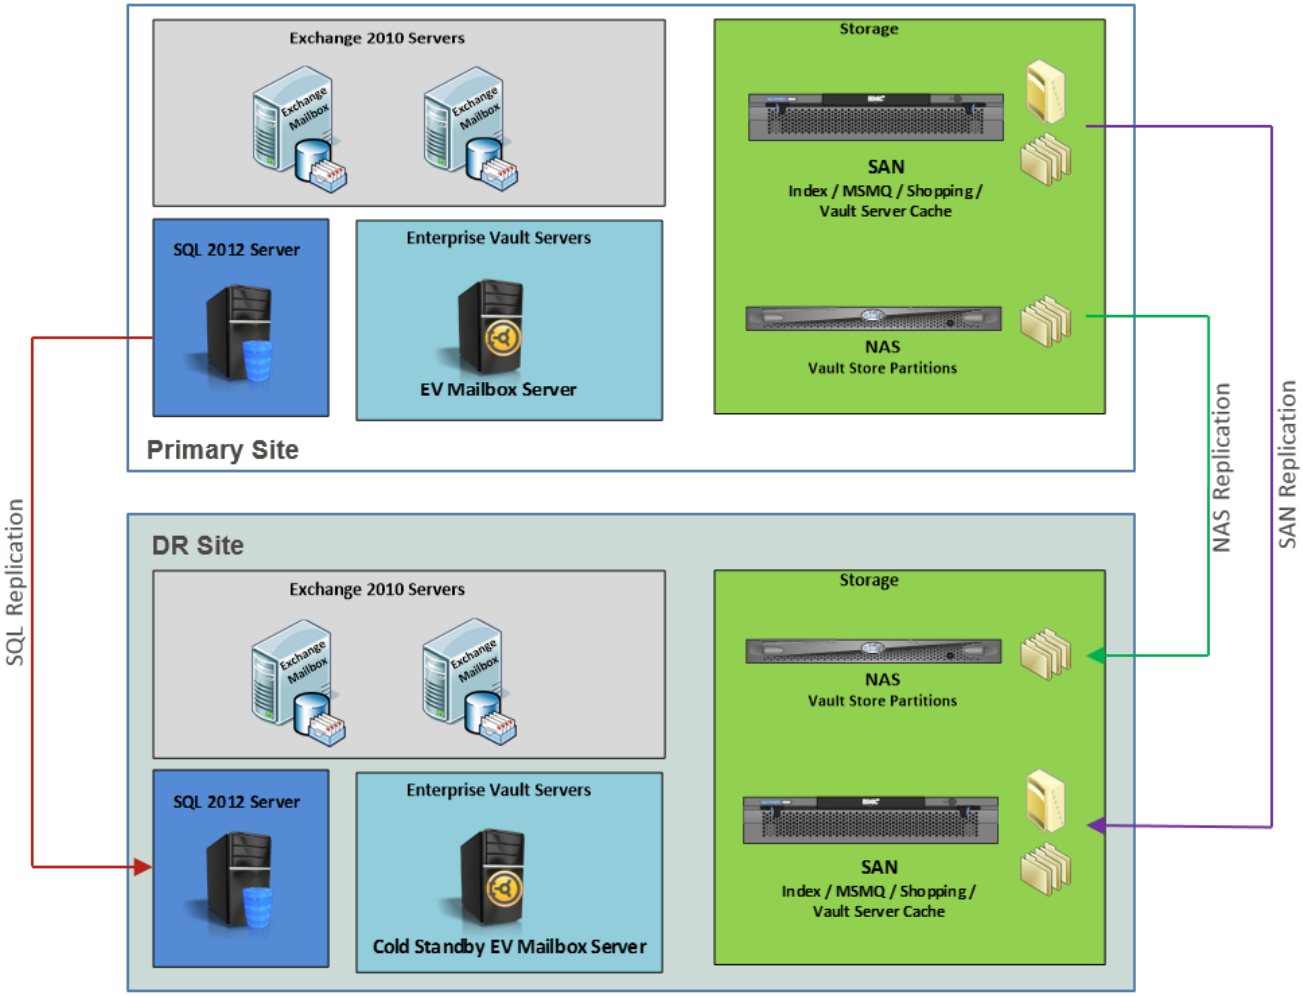

Figure 1 - Example Active/Passive site configuration

In this example the Microsoft Exchange Server mailbox databases are being replicated to secondary site using the Distributed Availability Groups feature of Microsoft Exchange Server. The business requested that Enterprise Vault should have a similar SLA as Exchange, and should therefore be configured in such a way that the application will still be available following a complete outage of the primary datacenter.

Data Replication

In order to achieve this level of redundancy, it is required that the following Enterprise Vault objects are replicated to the DR site:

- All Enterprise Vault databases
- Index Locations
- Vault Store Partitions
- Enterprise Vault Shopping location
- Vault Server Cache
- PST Holding areas
- EMC Centera Staging area (if used)

It is important to note that the various database and storage components that make up the Enterprise Vault solution needs to be in "sync" for a complete recovery to take place in the DR Datacenter. If for example the Vault Store SQL database was replicated at 9.30pm, and the last known good replica of the Vault Store Partition on the NAS took place at 9.45pm, then it is likely that any changes to the SQL database during the 15 minute difference window will be lost (not considering any log replay functionality in this example). In such an event the EVSVR tool can be used to build the missing SQL entries, but ideally this complication should be avoided.

Enterprise Vault Backup Mode

One way to ensure consistency across the databases and storage areas is to make use of the Backup Mode[1] option in Enterprise Vault. Typically EV servers would perform scheduled nightly archiving runs, where data is collected from archive targets outside of normal business hours. In most environments a backup of EV would follow the archive run, in order to clear the Safety Copies and remove pending state items from the archive targets. During the backup Enterprise Vault would go into a read-only mode known as Backup Mode, where no change is allowed in either SQL, Vault Store or Index locations. This is done to ensure the data from the 3 areas are backed up in a consistent state and a full restore of data is possible.

Backup mode can also be used to ensure the environment is replicated consistently to an alternative location, especially useful in cases where different replication technologies are used such as the example in Figure 3 (SQL database replication, SAN replication and NAS snapshot replication).

[1] For more information refer to the whitepaper Backing up Enterprise Vault available at **www.symantec.com/docs/TECH147148**

It is important to note that if this method is to be relied upon for full recovery the changes to the EV server and data outside of this backup window should be considered. For example if end users are allowed by policy to do manual archiving in Outlook, it's likely that the EV environment will see changes outside of the nightly archive run and therefore there will be an amount of data loss should a full recovery from the DR site be required. Safety Copies can be used to counter this particular issue, as the pending items can be set to revert to the original item after a set amount of time, but it won't be suitable for PST Migrations for example. It is also worth considering whether Vault Store Collections are enabled, as this is set to run between 10am and 4pm by default and will change the file structure and SQL database during the day.

It's highly recommended that a 24hr schedule of the Enterprise Vault environment is put together to show what activities happen at different times of the day and night, so as to work out when would be suitable to backup and replicate the environment to a state where a full recovery with no data loss will be possible.

There are many different replication technologies on the market, and depending on the SLA and Recovery Point Objective assigned to the application it may be that a bigger investment is required in selecting a technology that allows all the data sources to be replicated in sync at all times. It is however known that such solutions are likely to suffer write commit performance issues, as the data needs to be replicated before the transaction is committed.

A more cost effective and less complicated solution would be to set Backup Mode at intervals during the day, where archiving is paused and data can be replicated consistently during that window. This method would also be suitable for environments where there are constant changes to the EV environment, such as Exchange Journal Archiving for example. Depending on the replication technology used it may only be required to pause the archiving tasks a short while during which time the replication snapshot takes place.

Third Party Storage Replication Software

A number of vendors provide software and hardware based replication technologies for Enterprise Vault. It is always recommended to confirm with the vendor whether the solution is certified with the particular version of Enterprise Vault you plan to deploy. The Enterprise Vault Compatibility List (Referenced Documents) and the Partner Solutions for Enterprise Vault[2] web site show more information on certified solutions.

Standby EV Server

As depicted in Figure 1 at least one Enterprise Vault server will be required in the secondary/DR datacenter.

[2] http://www.symantec.com/theme.jsp?themeid=enterprise-vault-extensions

This server can form part of a stretched cluster, but in most cases a site outage will allow sufficient time to perform the failover in USL method described in the High Availability Options for Enterprise Vault paper. The server can be a VMWare or Hyper-V server that is only powered up when required, but it should be as powerful as the primary server if archiving will be resumed while in the failover state, or powerful enough to allow retrieval only if archiving is not required while in the failed over state.

Most importantly the server must be configured in the exact the same way as the primary server, with the exception that the *Enterprise Vault Configuration Wizard* is not run on the standby server. The standby server should have access to the replicated data using the same drive letters as the primary server, at the point when the replicated storage is enabled for write access.

The server must meet all standard build pre-requisites, including anti-virus exclusions. The server must also have the Enterprise Vault license files installed locally and any registry modifications done to the primary server must be replicated on the standby server.

Finally any scripts (for example PowerShell scripts to put the EV servers in backup mode) required during failover must also be modified with the new name of the standby EV server.

Configuring Enterprise Vault with VCS Global Cluster Option

For details on how to configure Enterprise Vault in a VCS Global Cluster with Veritas Volume Replication refer to technote **www.symantec.com/docs/TECH75241**.

Step by Step Disaster Recovery of Enterprise Vault

The information in this section is based on the Recovery section of the Administrator's Guide. It applies to a scenario where you are performing a recovery of a server locally using data-only backups, or when doing a recovery in a DR Datacenter where you have the databases and data (as outlined in the Site Redundancy section of this document) available in an alternative location.

Use the following recovery procedures when you have backed up only Enterprise Vault data, including the registry, and have not backed up the system disks on your Enterprise Vault servers.

Note: The following procedure may not be suitable for all Enterprise Vault environments, and it is always recommended that disaster recovery procedures are carefully planned and tested before used in production environments.

The procedures described in this section require backups of the following Enterprise Vault SQL databases:

- EnterpriseVaultDirectory
- EnterpriseVaultMonitoring
- EnterpriseVaultAudit
- Each FSA Reporting database you have set up, if you use FSA Reporting.
- Fingerprint databases
- Vault store databases

You must also have backups of the following Enterprise Vault data:

- Vault store partitions
- Index locations

You can use these procedures when you need to recover only one Enterprise Vault server, or to recover multiple servers.

To recover each server, you need to know which Enterprise Vault services it was running before the disaster occurred. If you are unsure which Enterprise Vault services were running on each server, run the SQL script ServiceLocations.sql, which is installed in the Enterprise Vault installation folder, for example C:\Program Files (x86)\Enterprise Vault.

Note: Before you can run the script you must first restore your Enterprise Vault Directory database.

Recovery procedure 1: Installing software on the servers

All the data relating to your previous Enterprise Vault installation needs to be recovered onto new servers. For each server that has failed you need to set up a new computer. Ideally, set up each computer with the same name as the original computer that it is replacing.

Note: If this is not possible the recovery steps tell you what to do to accommodate a change in computer name.

Build each new system, starting with the installation of Windows and then all the prerequisites for Enterprise Vault. Refer to the Enterprise Vault documentation if you are not sure which prerequisite software you must install on each computer. When you have set up the correct prerequisite software on each server, install Enterprise Vault on the server.

Note the following:

- Install Enterprise Vault on each new server, into the same folder as on the original server.
- Install the same version of Enterprise Vault as is being used in your current environment.

Do not run the Enterprise Vault Configuration wizard at the end of completing the installation of the Enterprise Vault software.

Recovery procedure 2: Restoring Enterprise Vault databases

Perform the steps in this section if the SQL server is unavailable and the databases require restoring.

Restore the following Enterprise Vault SQL databases:

- EnterpriseVaultDirectory
- EnterpriseVaultMonitoring
- EnterpriseVaultAudit
- Each FSA Reporting database you have set up, if you use FSA Reporting.
- Fingerprint databases
- Vault store databases

If you have restored EnterpriseVaultMonitoring or the FSA Reporting databases to a SQL server other than the one that previously hosted them, you must update the Directory database.

To update the monitoring settings in the Directory database run the following SQL script on the SQL server that hosts the Directory Database:

```
USE EnterpriseVaultDirectory
UPDATE MonitoringSettings
SET SQLServer = 'SQL_server_name'
where SQL_server_name is the name of the new SQL server.
```

To update the FSA reporting settings in the Directory database:

1. On the SQL server that hosts the Directory database, run the following SQL script to determine which SQL server hosted each FSA Reporting database:

```
USE EnterpriseVaultDirectory
Select SQLServer,DatabaseName From FSAReportingDatabase
2 Run the following SQL script:
USE EnterpriseVaultDirectory
UPDATE FSAReportingDatabase
SET SQLServer = 'SQL_server_name'
WHERE DatabaseName = 'FSA_reporting_database_name'
```

Where:

- SQL_server_name is the name of the new SQL server.
- FSA_reporting_database_name is the name of the FSA Reporting database that you restored.

Recovery procedure 3: Renaming EV servers

Ideally, you should set up each server with the same name as the original server that it is replacing. However, if this is not the case, you must perform the following extra procedure.

Warning: If you are running Enterprise Vault in a clustered environment, do not perform this operation unless Symantec Support advises you to do so.

To set up a server with a different name than the old server, repeat the following steps for each server that you are recovering:

1. Run SQL Query Analyzer and connect to the server that is running the Enterprise Vault Directory service. Enter and run the following SQL command:

   ```
   USE EnterpriseVaultDirectory
   UPDATE ComputerEntry
   SET ComputerNameAlternate = 'Name of new server'
   WHERE ComputerNameAlternate = 'Name of old server'
   ```

2. Check that the DNS alias you set up for the old server points to the name of the new server. If you are unsure what the DNS alias is, run the following SQL query against the EnterpriseVaultDirectory database.

   ```
   USE EnterpriseVaultDirectory
   SELECT ComputerName FROM ComputerEntry
   ```

3. If you are recovering the system that provided the vault site alias (usually the first server that was added to the site), then you need to update the vault site alias to point to the new server. To do this, perform the following steps in the order listed:

   - Run SQL Query Analyzer and connect to the server running the Enterprise Vault Directory service.
   - Enter and run the following SQL command:

     ```
     USE EnterpriseVaultDirectory
     SELECT SiteEntryId
     FROM SiteEntry
     ```

   The value returned contain the vault site alias at the end of a long string of numbers. For example, if the command returns the following then the vault site alias is sitealias:

   ```
   10354B15D38FE5B41BAAC212490EBA5351d10000sitealias
   ```

   In DNS, change the DNS alias entry so that it points at the new server.

Recovery procedure 4: Copy or move the Enterprise Vault data files

You now need to restore the backups of the Enterprise Vault data files to their locations on the Enterprise Vault servers.

Depending on the original Enterprise Vault components that existed on the servers you are recovering you must restore only the following data files:

- If you are restoring a server that used to run a Storage service, or a server that is configured in a cluster, you need to restore onto this server the saveset files for any vault stores managed by the original Storage service.
- If you are restoring a server that used to run an Indexing service, or a server that is configured in a cluster, you need to restore onto this server the indexing data files managed by the original Indexing service.
- If you are restoring a server that used to run a Shopping service, or a server that is configured in a cluster, you need to restore onto this server the shopping files managed by the original Shopping service.

The Enterprise Vault data should be restored to the locations where they existed on the original servers. For example, if you are recovering the server running the Indexing service and the indexing data was originally stored in the following location: `I:\Indexing`

Then this indexing data should be restored to the same location on the new server. To reorganize and move any SQL database devices on the disks, you can perform the procedures as listed in the following Microsoft Knowledge Base article: **http://support.microsoft.com/?kbid=181602**

This must be correct before you start any of the Enterprise Vault services, otherwise some cleanup operations may occur, resulting in information loss.

Recovery procedure 5: Clearing the directory database entries

You can clear the directory database entries for all of the Enterprise Vault servers in your environment, or for selected servers. The SQL query that is provided in this section clears the entries in the database for all the Enterprise Vault servers. If you have multiple Enterprise Vault servers in your environment, you may want to recover only some of the servers. The following technical note provides alternative SQL scripts that let you specify the servers for which you want to clear directory entries:

http://www.symantec.com/docs/TECH127004

To clear the directory database entries for all of the Enterprise Vault servers:

1. Run SQL Query Analyzer and connect to the server running the Enterprise Vault Directory service.
2. Enter and run the following SQL command:

```
USE EnterpriseVaultDirectory
UPDATE StorageServiceEntry
```

```
SET StorageArchive = '', StorageRestore = '',
StorageReplayIndex = '', StorageSpool = ''
UPDATE RetrievalTask
SET RetrievalSpoolQueue = ''
UPDATE ArchivingRetrievalTask
SET MessageQueue = ''
UPDATE RetrievalTask
SET MessageQueue = ''
UPDATE JournalTask
SET MessageQueue = ''
UPDATE PublicFolderTask
SET MessageQueue = ''
```

Recovery procedure 6: Recreating services and tasks on the first Directory service computer

If you are recovering all the Enterprise Vault servers that run a Directory service, you must use this procedure when you recover the first of these servers. When you recover the subsequent servers including other servers that run a Directory service, use the procedure described in Recovery procedure 7.

The Enterprise Vault Configuration wizard is able to detect missing services and tasks provided that the server name is identical to that in the original installation, or you have correctly followed Recovery procedure 3.

To recreate services and tasks on the first Directory service computer:

1. On the Windows Start menu, click All Programs > Enterprise Vault > Enterprise Vault Configuration.
2. Select Yes to create a new Directory service, and then click Next.
3. Enter the details of the Vault Service account, and then click Next.
4. The Enterprise Vault Configuration wizard does the following:
   - Converts the login for the Enterprise Vault Admin service so that it runs under the Vault Service account
   - Adds the Vault Service account to the local Administrators group on the computer
   - Grants the user rights Log on as a service and Debug programs to the Vault Service account
   - Creates and starts the Enterprise Vault Directory service.
5. When prompted for the name of the SQL Server that will host the directory database, enter the name of SQL Server used to host the directory database for the original configuration of Enterprise Vault, and then click Next.
6. The Enterprise Vault Configuration wizard checks that the SQL Server exists and can connect to it. As long as you have recovered the Directory service database, the Enterprise Vault Configuration wizard now recreates the services and tasks installed on the Directory service computer.
7. To recreate the Enterprise Vault services on the Directory service computer enter the password of the Vault Service account.

8. When the repair has finished, a success message is displayed.
9. If the Enterprise Vault Configuration wizard does not display a message, do not continue to run the wizard. Close the wizard and do the following:
    - Check that all previous steps have been successful, repeat any missed steps and then run the Enterprise Vault Configuration wizard again.
    - Create a String registry value called UseLanManNameForSCM under the following registry key:
        - On a 64-bit installation of Windows:

        `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\KVS\Enterprise Vault\Admin`

        - On a 32-bit installation of Windows:

        `HKEY_LOCAL_MACHINE\SOFTWARE\KVS\Enterprise Vault\Admin`

        Give UseLanManNameForSCM a value of 1. Run the Enterprise Vault Configuration wizard again. If you are sure you have followed all steps correctly and setting the registry key does also not help, contact your Enterprise Vault Support Representative for further assistance.

Recovery procedure 7: Recreating services and tasks on Enterprise Vault servers

If you are recovering all the Enterprise Vault servers that run a Directory service, for the first one you must use Recovery procedure 6.

When you recover the subsequent servers including other servers that run a Directory service, use the procedure described in this section. The Enterprise Vault Configuration wizard is able to detect missing services and tasks provided that the server name is identical to that in the original installation, or you have correctly followed Recovery procedure 3.

To recreate services on other Enterprise Vault servers:

1. Make sure the server running the Directory service is available on the network and the Directory service is started.
2. Make sure the Admin service is started on the local computer.
3. Run the Enterprise Vault Configuration wizard on the server by clicking Start > Programs > Enterprise Vault > Enterprise Vault Configuration.
4. When asked whether you want to create a directory or use an existing one, select No, use existing remote Vault Directory and enter the name of the server running the Directory service.
5. Enter the password of the Vault Service account. This is necessary to recreate the Enterprise Vault services on the computer.
6. The Enterprise Vault Configuration wizard recreates the Enterprise Vault services and tasks that used to run on the server and displays a message to indicate success.

7. If the Enterprise Vault Configuration wizard does not display a success message, do not continue to run the wizard. Close the wizard and then do the following:
   - Check that all previous steps have been successful, repeat any missed steps, and then run the Enterprise Vault Configuration wizard again.
   - Create a String registry value called UseLanManNameForSCM under the following registry key:
     - On a 64-bit installation of Windows:

     `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\KVS\Enterprise Vault\Admin`
     - On a 32-bit installation of Windows:

     `HKEY_LOCAL_MACHINE\SOFTWARE\KVS\Enterprise Vault\Admin`
   - Give UseLanManNameForSCM a value of 1. Run the Enterprise Vault Configuration wizard again.

   If you are sure you have followed all steps correctly and setting the registry key does also not help, contact your Enterprise Vault Support Representative for further assistance.

8. Start all the Enterprise Vault services. The message queues should automatically be recreated on the new server. If the Storage service is configured to start multiple processes, it may stop during message queue creation. This is because of a conflict between the processes creating the queues. To fix the problem, restart the Storage service.

   If the Indexing service finds any inconsistency in the index metadata, it automatically synchronizes the metadata. You may see the following events:
   - Event 41395 Index Volume metadata upgrade required
   - Event 41372 Index Volume metadata synchronization started

   During the synchronization the Indexing service logs progress events every 10 minutes. At the end of the synchronization, one of the following events is logged:
   - Event 41373 Index Volume metadata synchronization completed
   - Event 41377 Index Volume metadata synchronization completed

The index synchronization may take some time. For example, an Enterprise Vault recommended specification server takes approximately 10 minutes to process 5,000 index volumes. If any other index housekeeping is required there will be other progress messages every few minutes.

Recovery procedure 8: Checking the Web Access application settings

You must now ensure that the port and protocol settings for the Web Access application are correct.

To check the Web Access application settings:

1. Open the Administration Console.
2. Expand the Enterprise Vault and Directory containers.
3. Right-click the Site entry, and then select Properties.

4. View the General page. Check that the port and protocol set for accessing the WebAccess application virtual directory, /EnterpriseVault, match the settings on the Default Web Site in IIS.
5. Run the OWAUser script

Recovery procedure 9: Checking registry entries

Check that the Enterprise Vault registry entries are all set correctly on the newly-recovered servers.

The main registry entries are under the following key:

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\KVS\Enterprise Vault`

Additionally, you may have set registry entries under `HKEY_CURRENT_USER` when logged in as the Vault Site account. If so, restore these entries on each server too, under the following key:

`HKEY_CURRENT_USER\Software\Wow6432Node\KVS\Enterprise Vault`

Finally confirm that any customized INI files and Mailbox Welcome messages are copied over to the new server.

Steps to recover a standalone DA and CA server to a new server

It is very important to use the same existing Compliance Accelerator (CA) or Discovery Accelerator (DA) Configuration database after moving the CA or DA installation to a new server. The Configuration database maintains key settings used to fine tune the operations of CA or DA searches, exports, etc. The configuration settings are based on Customer ID. If a connection is made to an existing database instead of following the resolution listed below, the customer ID is changed and the configuration settings will not be preserved. After the changes below have been made the Customer database will automatically be available. If failing over to DR site, the Exports folder can optionally be replicated (if required). Note that the Pre-Fetch cache should not be replicated, the folder should exist but the content should not be replicated.

The following high-level steps are required to fail DA and CA over to a new server (suitable for both local server replacement and DR recovery)

1. Install prerequisite software on the new server as described in the Installing and Configuring Guide (See related documents below).
2. Stop the existing Enterprise Vault Accelerator Manager Service.Backup the CA or DA Configuration and Customer databases.
3. Install but DO NOT configure CA or DA on the new server.
4. If started, stop the Enterprise Vault Accelerator Manager Service.
5. For versions prior to CA or DA 8.x, copy the license into the Accelerator installation folder; by default: C:\Program Files\Enterprise Vault Business Accelerator for 32-bit systems, or C:\Program Files (x86)\Enterprise Vault Business Accelerator for 64-bit systems.
6. Copy the following four files from the original server to the installation folder of the new server (by default: C:\Program Files\Enterprise Vault Business Accelerator for 32-bit systems, or C:\Program Files (x86)\Enterprise Vault Business Accelerator for 64-bit systems):
    - AcceleratorManager.exe.config
    - AcceleratorManagerConsole.exe.config
    - AcceleratorService.exe.config
    - ADSynchroniser.exe.config

    Note: if moving from a 32-bit to 64-bit operating system, these 4 files must be edited to reflect the new path to the *Reports* and *Reports\Share* folders as well as the new path to the Enterprise Vault files noted within these files. For example, if moving from the default installation path on a Microsoft Windows 2003 (32-bit) server to the default installation path on a Microsoft Windows Server 2008 R2 X64 server, the following lines would need to be edited as follows for the Reports and Reports\Share folders:

From:

```
<add key="ReportSourceLocation" value="C:\Program Files\Enterprise Vault
Business Accelerator\Reports" />
 <add key="ReportShare" value="C:\Program Files\Enterprise Vault Business
Accelerator\Report\Share" />
```

To:

```
<add key="ReportSourceLocation" value="C:\Program Files (x86)\Enterprise
Vault Business Accelerator\Reports" />
 <add key="ReportShare" value="C:\Program Files (X86)\Enterprise Vault
Business Accelerator\Report\Share" />
```

There are other lines within each of these files that must be updated with the new path for EV or CA/DA. If the default installation paths were selected for both Enterprise Vault and CA or DA, then edit these files to replace all instances of Program Files\Enterprise Vault with Program Files (x86)\Enterprise Vault.

7. Modify the following three (3) tables in the CA or DA Configuration database. Replace the original NetBIOS server name with the new NetBIOS server name. If the NetBIOS name of the server is not changing, this step can be skipped.

Table Name	Column	Value
tblCustomer	IIS	NewNetBiosName

Table Name	Column	Value
tblGroup	Name	NewNetBiosName

Table Name	Column	Value
tblServer	Name	NewNetBiosName

Open SQL Query Analyzer and run the following queries separately against the Configuration database. Replace **bold** items with the correct server name values:

```
UPDATE tblCustomer SET IIS = 'NewNetBiosName' WHERE IIS = 'OldNetBiosName'

UPDATE tblGroup SET Name = 'NewNetBiosName' WHERE Name = 'OldNetBiosName'

UPDATE tblServer SET Name = 'NewNetBiosName' WHERE Name = 'OldNetBiosName'
```

8. Start the Enterprise Vault Accelerator Manager Service.
9. The CA/DA IIS virtual directories will need to be manually created on the new server:
    - Open the EVBAAdmin web page
    - Right click on the Customers (one at a time)
    - Select Check Virtual Directory.

Conclusion

Symantec's mission is to provide the essential tools to help its customers protect the security and availability of their information. As both the volume and importance of an organization's archive grow over time, Enterprise Vault customers can find peace of mind in the knowledge that their archiving platform excels in the scalability and high availability they need to fulfill their business requirements.

Appendix A: Disaster Recovery Checklist

Item	Check
The Disaster Recovery plan is accessible and up to date	
The Disaster Recovery plan has been tested, and no major changes has been made to the environment since the last test took place	
Backups are done “in-sync”, and a full recovery of SQL, Vault Store and Index is possible with no data loss	
EV Standby Server has a valid SLF license file	
EV Standby Server is configured with access to the storage locations using the same drive letters	
EV Standby Server is configured with Anti-virus Exclusions, even for drive letters that will only be mapped once failover is initiated	
EV Standby Server contains all custom registry modifications applied to primary server	
EV Standby Server has script files (EVPM, Backup, etc) with the correct DR server name	

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help consumers and organizations secure and manage their information-driven world. Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored. Headquartered in Mountain View, Calif., Symantec has operations in 40 countries. More information is available at **www.symantec.com**.

For specific country offices and contact numbers, please visit our website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com