



A CommVault White Paper: Continuous Data Replicator *Continuous Data Protection for Microsoft Exchange Environments*

**CommVault Corporate Headquarters
2 Crescent Place
Oceanport, New Jersey 07757-0900 USA
Telephone: 888.746.3849 or 732.870.4000**

CommVault Continuous Data Replicator:

Continuous Data Protection for Microsoft Exchange Environments

Continuous Data Protection for Microsoft Exchange Environments	1
The Challenge	1
The Solution.....	1
Cost-Effective Replication and Recovery of Exchange.....	2
Primary Use Cases and Key Differentiators	2
Continuous Data Protection for Exchange Servers.....	2
Remote Site Replication for Data Consolidation.....	3
Replication to an Offsite Disaster Recovery Location	3
Flexible Replication Features	4
Consistent Recovery Points	4
Out-of-Band Synchronization.....	4
Bandwidth Throttling.....	5
Data Compression.....	5
Data Encryption	6
Smart Re-Sync	6
Capacity Management	6
CDR Best Practices for Exchange Disaster Recovery	7
Preparing the Production Exchange Server	7
Preparing the DR Exchange Server	7
Setting up the Replication Process.....	8
Recovering the Failed Exchange Server Using the DR Exchange Server.....	8
Recovering the Exchange Production Server	8
Conclusion	9

CommVault Continuous Data Replicator: Continuous Data Protection for Microsoft Exchange Environments

The Challenge

In today's connected world, the reliance on Microsoft Exchange Server based messaging systems has become critical to the daily operations of most businesses. Any loss of Exchange data or server accessibility, regardless of the amount or duration, can trigger a costly and potentially catastrophic effect on communications, customer relations, e-commerce and compliance.

The impact of exponential data growth, mandated compliance and 24 x 7 availability demands require solutions that enable the continuous protection of Microsoft Exchange server. Regardless of the type of disaster you encounter, you need instant access to the latest copy of the data or applications in order minimize data loss and expedite recovery. Depending on your established Recovery Point Objective (RPO) or Recovery Time objective (RTO), your nightly backup may not be sufficient to meet service level agreements.

The Solution

CommVault® Continuous Data Replicator (CDR) is the ideal solution designed to help businesses meet critical recovery point and recovery time objectives. This host-based, multi-platform replication technology helps centralize and simplify the administration and protection of remote office data. It also provides a cost-effective option for replicating data to a disaster recovery site. By continuously capturing and replicating data changes at the byte level, CDR minimizes the amount of data that must be transmitted across the network and reduces impact on performance. One of the many unique features that differentiate CDR from other replication products is the ability to create and maintain multiple, application-consistent snapshots, or point-in-time copies, of your Exchange Server data. This application awareness ensures the integrity and recoverability of replicated Exchange data. The use of snapshots provides administrators with an additional layer of protection against viruses or corruption by allowing them to quickly recover to a point-in-time snapshot prior to the event which caused the data loss.

Key Benefits of CDR:

- Reduces management costs and increases backup reliability by centralizing remote office data
- Provides cost-effective Disaster Recovery for Microsoft Exchange, SQL Server and Oracle
- Reduces potential data loss, increases availability and expedites recovery

Although CDR is not limited to supporting Microsoft Exchange, this white paper focuses on providing use cases, features and detailed information about protecting and recovering Exchange Server data using CDR. For additional information about using CDR with Microsoft SQL and Oracle or for heterogeneous remote office backup consolidation, please reference the CommVault Continuous Data Replicator Overview white paper available from the CommVault corporate website at <http://www.commvault.com>.

Cost-Effective Replication and Recovery of Exchange

Continuous Data Replicator offers a compelling asynchronous alternative to high-cost remote hardware mirroring solutions. It is ideal for customers who are looking to implement a cost-effective disaster recovery solution for key enterprise applications such as Microsoft Exchange Server. CDR maintains updated copies of file system and application data at remote or virtual locations. These copies can be quickly mounted and brought on-line in the event of a site outage or significant data loss.

Key Benefits:

- Quickly recover data and applications in the event of disaster
- Cost effective alternative to hardware mirroring

Primary Use Cases and Key Differentiators

Continuous Data Replicator offers many unique advantages over other, single purpose, replication products currently available on the market. CDR can be deployed as a standalone replication solution to consolidate remote office data and provide continuous data protection from data loss and site failures. Best of all, it can be added to an existing CommVault environment enabling replication and continuous data protection to work seamlessly with the backup and archive modules for simplified management and continuous data protection of local and remote environments.

CommVault Continuous Data Replicator supports three primary use cases:

1. Continuous data protection for datacenter and remote office environments
2. Consolidation of remote Exchange Server data to a central site
3. Replication of Exchange data to an offsite disaster recovery location

Continuous Data Protection for Exchange Servers

CDR protects file systems and application data such as Exchange, by replicating data from a source machine to a destination machine in near real-time mode. Key data at the source or primary site can be designated as a member of a replication set pointed to a target location across the WAN at the central site. Moreover, CDR for Exchange automatically detects all associated Exchange folders that should be replicated such as: Email database files .EDB, Log files .LOG, etc. and includes them automatically in the

replication set. CDR may also be configured for local replication to service data availability scenarios. Once the replication set is created, CDR captures byte level changes as they are written at the source and rolls them up to a log file. Once the log reaches a defined size or time limit, the changes are transferred to the destination or secondary location where the writes are applied. This process reduces performance impact on the application or file server and minimizes bandwidth use.

Remote Site Replication for Data Consolidation

In a many-to-one configuration, CDR performs the replication of critical datasets from multiple remote locations to a central site in order to provide consolidated remote office data protection.

Replication to an Offsite Disaster Recovery Location

In a one-to-many configuration, CDR allows the replication of critical datasets from a central site to one or more physical or virtualized disaster recovery locations.

Figure 1 shows one possible configuration for Disaster Recovery. The mail server and database server at the central site are both replicated over the WAN to two separate locations. As the Exchange data is replicated to the remote datacenters, Recovery Points are created at specified intervals ensuring an available snapshot that can be brought online in the event of a site disaster.




Figure 1. Disaster Recovery Configuration

Flexible Replication Features

CDR provides a number of enterprise-class features to ensure data integrity, maximize the use of available network resources and protect data in transit.

Consistent Recovery Points

CDR enhances data protection of Exchange by providing the ability to create point-in-time, application consistent snapshots of replicated data. These snapshots, called Consistent Recovery Points, do not place any additional load on the Exchange servers or the network. Unlike other competitive point replication products that create snapshot images at the source and then replicate these potentially very large images over the network, CDR creates application consistent snapshots of the replicated data at the destination, after ensuring the consistency of data. Application awareness ensures that a single Consistent Recovery Point includes consistent snapshots of all volumes associated with the Exchange replication set, including .EDB files, Log files, etc..

Figure 2 shows the replication process and consistent recovery point creation. Prior to creating a Consistent Recovery Point for Exchange on replicated data, it is important that the application data on the source be in a consistent state. When creating a Recovery Point, CDR interacts with the Exchange application and signals it to put its data at the source into an application consistent state. When the Exchange data is in a consistent state, a marker is placed into the replication logs which indicates the point where a consistent snapshot can be taken. This marker, when encountered while applying the logs

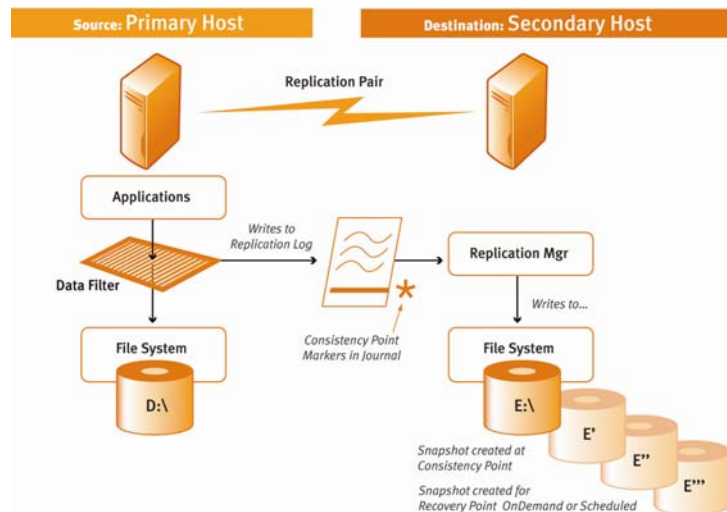


Figure 2. Replication Process and Recovery Point

at the destination, triggers a recovery point snapshot of the replicated data thus ensuring Exchange data consistency at the target. These recovery points, residing on the destination machine, can also serve as data sources for automatic backup operations. This seamlessly integrates replication with your centralized backup strategy.

Out-of-Band Synchronization

Replication involves creating an initial, one-time, baseline replica or “mirror” on the target. Once the baseline is established, changes on the source are then transferred on a

near real time basis. The initial size of a replication data set is typically much larger than the amount of data that changes on a daily basis. Therefore, available network bandwidth, sufficient for ongoing replication, may not be sufficient to create the initial replica in a reasonable time frame.

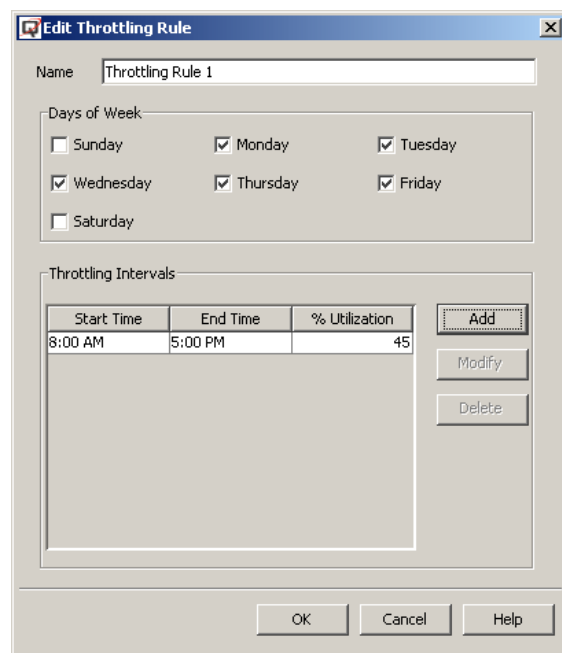
CDR addresses the issue of creating large replica sets by leveraging a built-in utility to capture the state of the initial replica while the administrator uses an alternate, out of band, option to quickly transfer the large initial dataset to the target host. Alternate options might include using a backup or disk clone sent via overnight shipping to the target location. While the data is in transit, CDR captures and logs ongoing changes to the source machine. Once the initial replica copy is complete, CDR then transfers the changes to the remote replica and the normal replication process continues from that point forward.

Bandwidth Throttling

Network throttling can be leveraged to control the bandwidth that is available to transfer replication data. CDR allows administrators to explicitly define how much of the available network's bandwidth should be allocated to replication activity. This allocation can be automatically adjusted on an administrator-defined schedule to increase network productivity.

It is important to note that the rate of change must be such that the changes that occur in a day can be transferred within that same day otherwise, the target will not be able to catch up with the changes on the source.

Also, if bandwidth is reduced to a percentage where the transfer of replication logs is delayed, additional storage space on the source will be consumed. This may exceed storage space and thus cause replication to stop. CommVault Technical Services can assist with the configuration and implementation of your CDR environment.



Data Compression

When network bandwidth is at a premium, administrators can enable compression of the replication stream before transferring it to the destination. At the destination, the data is decompressed as it is written to the target volume. The bandwidth savings is highly dependent upon the type of data within the replication stream.

Data Encryption

When replicating critical business data between remote offices, a dedicated network is not always available or cost effective. In order to protect the privacy of the data when replicating over open networks, administrators can encrypt the replication stream at the source. Upon reaching its destination, the replication stream is decrypted as it is written to the replica volume.

Smart Re-Sync

The reliability of any remote replication solution is dependent on the underlying network. Network interruptions do and will happen. The ability to recover from network outages is critical to the viability of the overall disaster recovery solution. CDR is unique in its ability to quickly and reliably recover from network outages ensuring that a replicated data set is brought back into sync in an efficient manner.

When a network outage occurs, the replication system simply continues to capture changes within the replication log area of the source machine. When connectivity is restored the accumulated logs are sent to the destination. If the duration of the outage is so long that the changes on the source exceed available log space, CDR automatically re-establishes the replica integrity without the need to perform a full re-sync of the destination replica. This Smart Sync of the replication set is fast and consumes minimum resources on the systems, regardless of the size of the dataset. This is an important consideration when working with the large datasets.

Capacity Management

When used in conjunction with Data Archiver (DA), CDR can leverage rule-based, capacity management capabilities to easily move large or infrequently accessed files to secondary storage. Data Archiver leaves behind a small stub file of the migrated data. CDR can replicate that stub file without needing to recall the underlying data. After the stub is replicated, it can be accessed normally. Data Archiver can be deployed at the replication source to reduce the size of the replication data set. This significantly reduces the time it takes to create the initial replica and also reduces the storage capacity required at both the source and the destination. This is particularly useful in remote office fan-in environments.

CDR Best Practices for Exchange Disaster Recovery

Any one who deals with the administration of Exchange knows that recovering a failed Exchange server can be a difficult and time consuming task depending on the mode of failure. In particular, using an alternate server to replace a production environment requires a number of operational steps which can be made even more difficult when under pressure to restore the environment quickly. In order to make the recovery of an Exchange server to a DR Exchange server as painless as possible, we recommend preparation ahead of time.

- ◆ Make sure that the production Exchange server is operational and up to date with the latest service packs and/or required patches.
- ◆ Make sure that the target DR server is running as a member of the domain and that the CommServe has already been installed.

Preparing the Production Exchange Server

The following steps will help prepare the production Exchange server for CDR.

1. Move storage groups off the C: drive. This is usually the case in production environments but the location should be verified.
2. Place the system path on the same volume as the transaction logs. This will simplify the configuration of the replication pairs since the *.chk* file and the logs will be located on the same volume.
3. Install the CDR agent on the server along with the File System agent so that the consistent recovery points can be backed up.

Preparing the DR Exchange Server

The following steps will help prepare the DR Exchange server for CDR.

1. Install the CDR and File System agents on the DR Exchange Server.
2. Create Volumes and Folders for storing Exchange data and logs equivalent to those on the production Exchange server. These will be used as replication targets for the live data as well as operational data and logs locations when the DR Exchange server is being used.
3. Install all the components that were on the production Exchange server including service packs and patches and select the Microsoft Management Tools as you will need to access the System Manager throughout the set up process.
4. Start Exchange services to verify the location of Transaction logs, Stores, public folders and system paths.

Setting up the Replication Process

- ◆ Use Add App Button to set up replication pairs. This enables the discovery of all log file locations and the creation of all necessary Replication Pairs. This ensures that the database as well as all Exchange log destinations are replicated and can be successfully used for Copy back operations.
- ◆ Leverage Out-of-Band Synchronization to create initial replica volume if required.
- ◆ Schedule consistent recovery points to ensure point in time recovery of Exchange.
- ◆ Set up bandwidth throttling to optimize bandwidth use.

Recovering the Failed Exchange Server Using the DR Exchange Server

Exchange dependencies with Active Directory require some procedural steps to be performed in order to recover the Exchange application on the disaster recovery Exchange Server.

- ◆ Make sure the production server is shut down.
- ◆ Reset the production Exchange Server computer entry on the Domain Controller.
- ◆ Remove the DR Exchange server from the Domain.
- ◆ Rename the DR Exchange server to the Production Exchange Server Name.
- ◆ Rejoin the DR Exchange server to the Domain as the production Exchange server.
- ◆ Use the live replica of the data or roll back to a consistent recovery point and perform the Copy back function.
- ◆ Restart Exchange services and mount all mail box and public stores.
- ◆ Change the IP address in the DNS entry of the production Exchange server to match the IP address of the disaster recovery Exchange server. This will enable Outlook clients to connect to Exchange by redirecting them to the DR Exchange server.

Recovering the Exchange Production Server

Simply perform the same steps as you did initially to set up a DR server whereby, the original production server assumes the role of the DR server.

Additionally the out-of-band synchronization capability provides an effective way to perform a fail-back operation of an Exchange Server in a remote DR configuration. If a primary Exchange Server fails and the secondary Exchange Server is used for production, the out-of-band synchronization feature can be used to quickly identify and bring back the primary Exchange Server data that was modified on the secondary Exchange Server during the outage. This allows the administrator to quickly resume a normal operating environment when recovering from a failure at the primary site.

Conclusion

CommVault Continuous Data Replicator is a multi-platform replication solution ideal to help centralize and simplify the administration of remote office data, deliver cost-effective disaster recovery and provide continuous data protection for datacenter and remote office environments. It delivers enterprise-class features such as compression, encryption, advanced link recovery, and the ability to synchronize initial replica sets out-of-band. Unlike competitive offerings of point products, CDR can be used as a stand-alone replication solution or as an integrated component of the CommVault Singular Information Management suite. This integration provides the industry's only policy-based GUI from which to easily manage replication along with backup, recovery and archive. It also provides the granularity to track remote backup sets to the source without time consuming manual processes and allows administrators to browse backup recovery points and perform recovery operations as if the backup was created locally. CommVault CDR is ideal for implementing a cost effective, disaster recovery solution for Exchange Server. It enables you to maintain up-to-date copies of your file systems and application data at remote or virtual locations providing quick access and fast recovery in the event of a disaster.