



Best practices for fully automated disaster recovery of Microsoft SQL Server 2008 using HP Continuous Access EVA with Cluster Extension EVA



Table of contents

Overview.....	2
Key findings.....	3
Solution configuration.....	4
Server configuration	7
Disk array and SAN configuration	7
Array-based replication with HP Continuous Access	10
Failover cluster configuration.....	11
Disaster recovery with Cluster Extension EVA.....	11
Application load simulation.....	15
Testing	15
Disaster Scenario 1: All nodes in data center fail.....	16
Disaster Scenario 2: Primary EVA storage array becomes unavailable	18
Disaster Scenario 3: Entire data center failure	22
Disaster Scenario 4: Failover during resynchronization	23
Best Practices.....	25
Conclusion	26
For more information	27

Overview

As database availability becomes critical in every segment of the market, administrators are looking for recommendations and best practices to configure their database, disk array, server, software for remote replication, and the mechanisms available for disaster recovery automation. When properly configured and deployed, these infrastructure elements are capable of providing timely recovery from a variety of failures, including the loss of an entire data center site.

HP offers a range of solutions that protect your valuable information against the threat of downtime, making your operations resilient regardless of external or internal events. HP StorageWorks Continuous Access EVA Software solutions provide disaster tolerance to keep your corporate information accessible and available through adverse events. Any amount of system downtime can mean lost productivity, lost revenue, lost customers, and lost opportunities. HP provides proven technologies, strategies, and services to reduce your exposure and vulnerability.

HP StorageWorks Cluster Extension EVA (HP CLX EVA) is a proven, full-site, disaster-tolerance software solution designed for the mid-range storage or tiered enterprise storage markets. HP CLX EVA integrates with cluster software to provide automated failover of storage disks between data centers, allowing applications to fail over between data centers that are separated by metropolitan distances with minimal downtime.

HP CLX EVA brings HP Continuous Access EVA functionality to Microsoft Failover Clustering for Windows environments. In doing so, HP CLX EVA allows a failover cluster to take advantage of HP Continuous Access EVA by allowing cluster nodes and shared storage to span across two data centers, eliminating storage as a single point of failure. HP CLX EVA provides high availability of storage by automatically making the decision of whether or not it is appropriate for the cluster software to fail over the application from one cluster node to another, and it also monitors the health of the storage. HP CLX EVA provides helpful services, such as swapping the personalities of the replicated disk pair to facilitate read/write access to the disaster recovery group LUNs.

This paper provides a set of test-proven best practices for configuring, deploying, and operating Microsoft SQL Server 2008 for high availability.

The solution consists of the following components:

- Microsoft SQL Server 2008 (SP1)
- Windows Server 2008 Failover Clustering
- HP StorageWorks Enterprise Virtual Array (EVA) 4400 Disk Array
- HP StorageWorks Continuous Access EVA
- HP StorageWorks Cluster Extension EVA

Specifically, this paper addresses the following fundamental deployment tasks:

- Using existing recommended best practices for configuring a SQL Server 2008 database and an EVA4400 for the following benefits:
 - Optimal array-based replication performance
 - Optimal database performance, using existing best practices, set-recommended recovery point objectives (RPO) and recovery time objectives (RTO)
 - Proper sizing of the inter-site link (ISL)
 - Accounting for the behavior of the bandwidth and latency of the ISL influence replication
- Using existing recommended best practices for configuring the following major environment elements:
 - Servers
 - Storage device
 - Storage replication
 - Network
 - Storage area network (SAN)
 - Database
- Configuring the SQL Server volumes (LUNs) with HP Continuous Access data replication groups for optimal availability and consistency.
- Deploying and configuring the HP StorageWorks Cluster Extension EVA as a cluster resource for SQL Server 2008. This configuration uses parameters that will ensure the fully automated recovery of an SQL Server 2008 in disaster scenarios. These disaster scenarios disrupt the cluster and cause the database resources to failover from a primary data center to a secondary data center at a remote site.

Key findings


- Ensuring consistency between the shared database volumes of a SQL Server instance, which requires adding each volume used in the database (log and data files) to data replication groups.
- Ensuring fully automatic recovery between data centers requires disabling the failsafe data replication property. Also, the HP CLX EVA resource parameter *Use non-current data Ok* must be set to **YES**.
- By default, the data replication property *Suspend on links down* is enabled, preventing automatic failback after recovering from a data center failure, which occurs when the replication link between two EVAs is lost.

- If a failover between data centers occurs during a resynchronization, the HP CLX EVA parameter *RESYNCWAIT* can be enabled to ensure that failover to the remote data center is successful.
- Microsoft Distributed Transaction Coordinator (MSDTC) is required by SQL Server 2008 in Failover Cluster. MSDTC should have its own EVA data replication group—if it is configured in a separate application group with a configured HP CLX EVA resource—for the MSDTC to have the capability to fail over between all cluster nodes.

Solution configuration

The solution outlined in this paper incorporates Microsoft SQL Server 2008 with SP1 in a high-availability configuration, using HP StorageWorks Cluster Extension EVA integrated with Windows Server 2008 Failover Clustering and HP Continuous Access EVA. The configured solution addresses a common situation in which there are two sites in a disaster recovery plan. During testing, Node and File Share Majority is used as the quorum mechanism for the cluster. Ideally, Node and File Share Majority should have a third location where the file share witness is located. As [Figure 1](#) shows, the configuration uses HP Continuous Access EVA to replicate the database data and log files between the two sites at the disk array level.

Figure 1. Solution configuration



[Table 1](#) lists the complete solution components used in testing.

Table 1. Solution components

Number on diagram	Function	Model	Specification
1, 2, 3, 4	SQL Server cluster node	HP Proliant BL465c G1	CPU: 2 Dual-Core MD Opteron, 3000 MHz Memory: 32 GB HBA: 2 LPe1105 MPIO DSM 3.01.00 SQL Server version: SQL Server 2008 Enterprise SP1 x64 Operating System: Windows Server 2008 SP1 x64
5, 6	Command View EVA management server. Callout 6 also provides file share witness used for quorum. This is normally located in a third data center.	DL380?	Command View Version: 9.00.00 CPU: Intel Xeon 3.60Hz Memory: 1 GB HBA: 2 LP1050 Operating System: Windows Server 2003 Enterprise SP2
13	Database transaction load generator	HP Proliant BL465c G1	CPU: 2 Dual-Core MD Opteron, 3000 MHz Memory: 32 GB HBA: 2 LPe1105 Operating System: Windows Server 2008 SP1 x64
7, 8	EVA Storage Array	HP EVA 4400	Firmware version: 09501200
9,10,11, 12	Fiber Chanel Switch	HP 8.24c SAN Switch Pwr pk+ BladeSystem c-Class	

The ideal HP CLX EVA environment consists of at least five servers configured as follows:

- Two servers at each data center site
- One non clustered host used for the file share witness arbitrator at a third site
- Separate, redundant communications links for cluster heartbeats, client access, and HP Continuous Access EVA.

Note: All communications interfaces must be installed in pairs to serve as failover components to prevent single points of failure.

Multipathing software, such as HP MPIO Full Featured DSM for EVA4x00/6x00/8x00 family of Disk Arrays for Windows operating systems, is used for host-to-storage connections. For more information, see [Other Multipath Options for HP Arrays](#).

For network communications links and components between the dispersed data centers, the best practice is to have redundant communications and physical components routed differently to prevent the *backhoe* issue that is created when all links between data centers are cut together. By using redundant and separately routed communications and components, you can eliminate single points of failure, which prevents application service startup.

To ensure reliable data backup between the local (primary) and remote (secondary) storage systems, at least two HP Continuous Access EVA links must be available in the SAN when the storage systems are connected. Doing so provides redundancy and protection against single points of failure. Although TPC/IP communications links can cover considerable distances, each network segment must be extended to the dispersed data center to maintain a heartbeat among all servers.

Two cluster nodes are located at each data center. In the event of a single node failure, the clustering software gives priority to the second local node when failing over SQL Server resources. This response avoids the more time-consuming operation of failing over to a node that is located in the remote data center.

When failing between data centers, the HP CLX EVA must reconfigure the storage systems to change the mirroring direction. At the remote site, two systems should be available in case the primary failover system experiences a hardware or power failure.

Server configuration

HP BladeSystem servers are used in the test environment. They offer correctly sized Intel or Advanced Micro Devices (AMD) processors, sufficient memory slots to service the SQL Server requirement, and multiple benefits that are based on integrated infrastructure and power-intelligent design.

Disk array and SAN configuration

Databases traditionally require a high-performing disk array to accommodate the potentially demanding I/O throughputs that they can generate. For this reason, the EVA is used in testing because it is a mid-range system that uses virtualization technology to increase

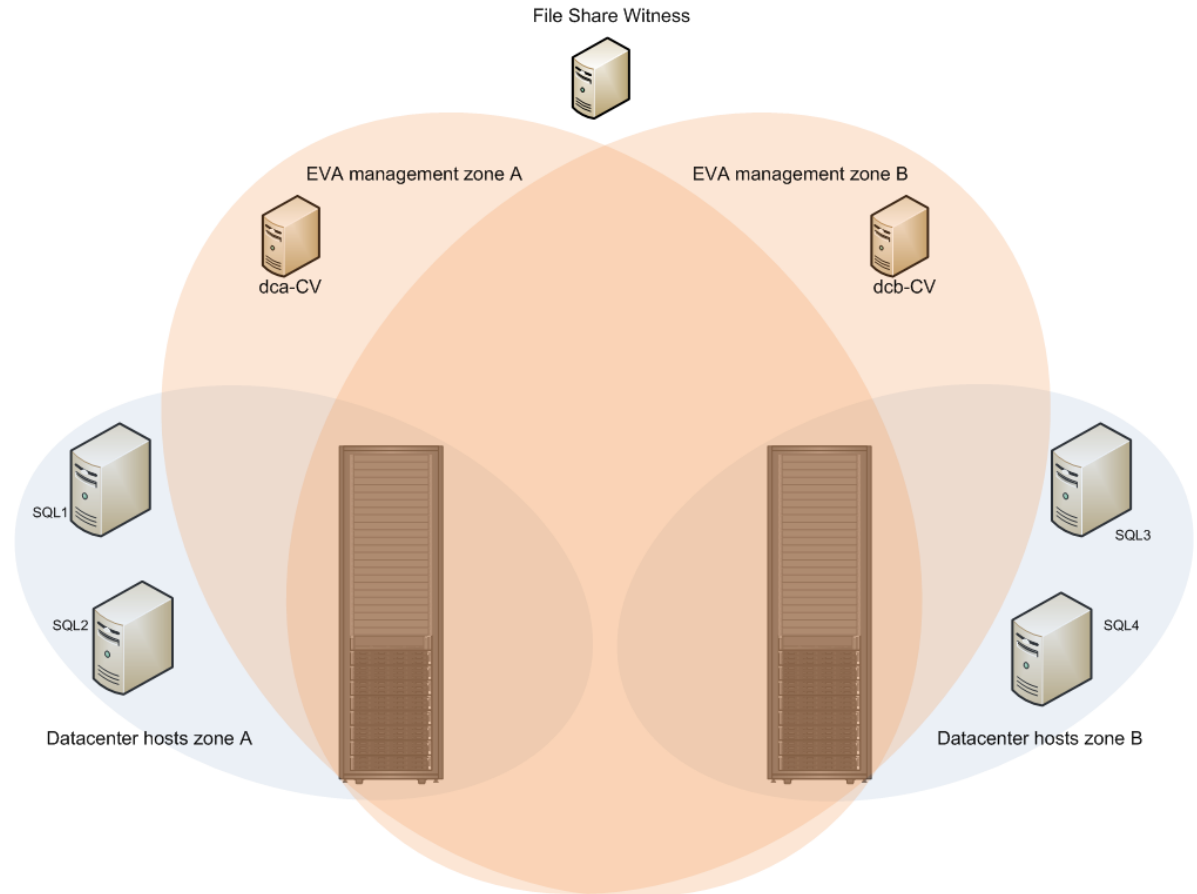
performance with a given number of disks (that is, as the number of disk spindles increase, performance will increase).

The SAN incorporates multiple switches operating at 4 Gb/s, of which four (Brocade BladeSystem 4/24) reside in each HP BladeSystem blade enclosure.

The design solution includes a dual fabric implementation. This redundant design means that the solution continues to operate, regardless of the failure or configuration problems of any SAN component. Multipath Input Output (MPIO) software is configured with the HP-provided device-specific module (DSM) on the participating servers to manage redundancy and load balancing at the HBA level.

The SAN was zoned by using the configuration recommended by HP Cluster Extension EVA. Each cluster node is zoned to observe the local EVA, but not the remote EVA. Each management server recognizes both EVAs. [Figure 2](#) shows the zoning configuration.

Figure 2. Zoning configuration



- The storage and disk layout design takes into consideration both disaster tolerance and performance.

- The database data files are separate from the transaction log files on independent disk groups and are, therefore, on separate spindles.
- The disk group layout design provides redundancy, which allows for database reconstruction if either of the disk groups fails.
- The disk group layout provides isolation of disk I/O patterns because the database data files are accessed in a random pattern, and the transaction log data access is sequential. These characteristics are typical of OLTP workloads.

[Table 2](#) and [Table 3](#) list the disk details and LUN layout. [Figure 3](#) shows a pictorial overview of the storage configuration.

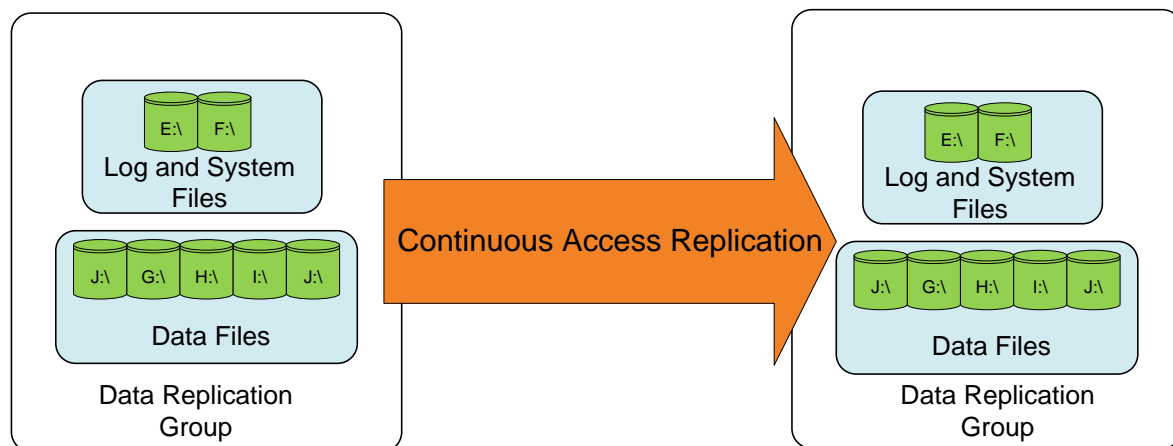
Table 2. Disk array configuration

Disk group	Number of disks	Capacity	Contents
Transaction logs	8	1104 GB	Transaction logs
Data files	16	3908 GB	Data files

Table 3. LUN (vdisk) configuration

Disk drive letter or mount point	Disk group	LUN size (GB)
E:\SQL_SYS	Master database, temp_db, and so on	20 GB
F:\SQL_LOG	Transaction logs	512 GB
G:\SQL_DATA1	Data files	512 GB
H:\SQL_DATA2	Data files	512 GB
I:\SQL_DATA3	Data files	512 GB
J:\SQL_DATA4	Data files	512 GB

Figure 3. Storage replication configuration for SQL Server disk resources



Calculating the required number of disk spindles

Because the primary focus of this paper is disaster recovery of SQL Server by using HP Continuous Access EVA with HP CLX EVA, disk performance is not the focus for the test environment. However, there are different approaches to plan and calculate the number of disk spindles required by a workload. The approach you use depends on what information is available about the type of workload supported by the disk spindles.

For more information about estimating disk group size for SQL Server and the HP transaction Processing Storage Planning Calculator for SQL Server, see the [SQL Server 2008 remote replication and high availability using HP Continuous Access EVA and ProLiant Servers](#) white paper.

Note: Even though performance is not heavily stressed for this project, we test a database environment with 1,000 users.

Array-based replication with HP Continuous Access

To guarantee consistency among all disk resources used in a clustered instance of SQL Server, each EVA virtual disk used for SQL Server is configured to reside together within a single EVA data replication group, as shown in [Figure 3](#).

Because HP CLX EVA supports only synchronous replication, we use the replication mode in the test environment.

With synchronous replication, all write requests from the server are first transferred to the remote storage system. Each I/O is mirrored in the cache area of the remote array, and an acknowledgement is sent to the local storage system. The write request is then acknowledged to the server.

Failover cluster configuration

Because a quorum disk requires read/write access from all nodes in the cluster, it cannot be applied in a HP CLX EVA cluster. This is because only nodes in one data center have such disk access. The method for implementing cluster quorum in a HP CLX EVA cluster is Node Majority or Node and File Share Majority. For this paper, Node and File Share Majority is used. Ideally, the file share witness should reside on a host outside of the cluster that is situated in a third location.

SQL Server 2008 cluster installation requires that Microsoft Distributed Transaction Coordinator (MSDTC) is configured. Because MSDTC uses a shared disk resource that spans two data centers, the MSDTC disk resource must reside in its own data replication group (separate from the SQL Server data replication group). It also must be configured with HP Cluster Extension EVA so that it can fail over between all the cluster nodes.

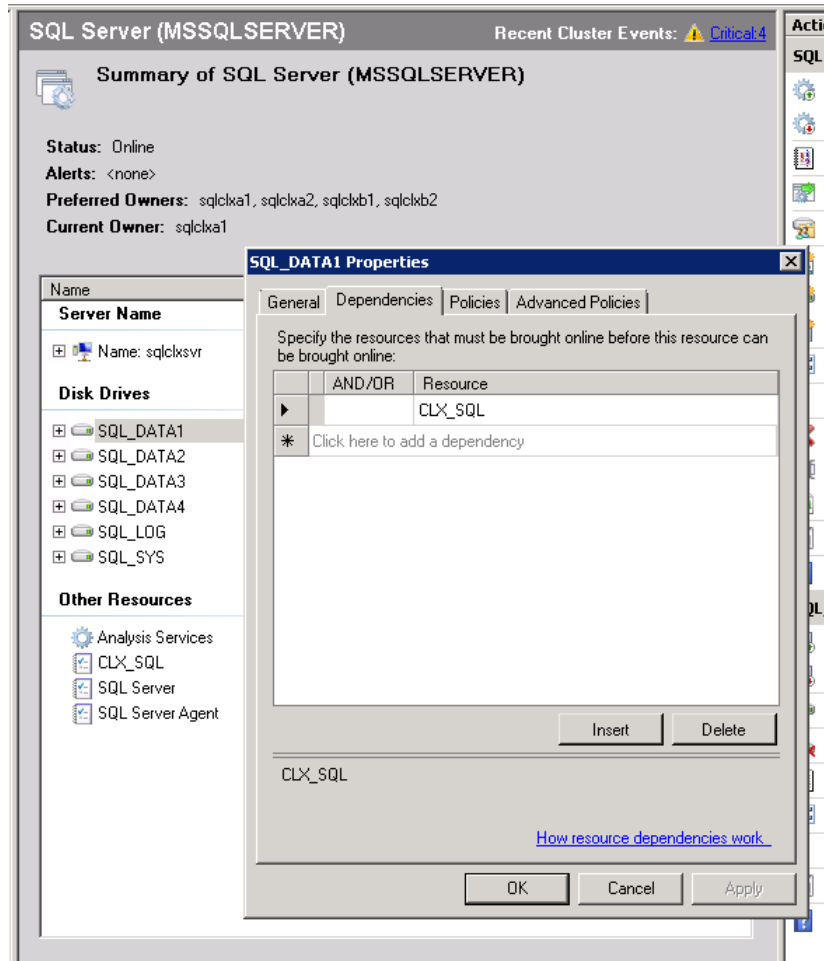
Because HP CLX EVA is part of the SQL Server startup process, resource dependencies must be set before storage system virtual disks (vdisks) are accessed. This allows the HP CLX EVA resource associated with the SQL Server instance to come online in the cluster prior to the disk resources.

Disaster recovery with Cluster Extension EVA

HP Cluster Extension EVA Software requires Microsoft Failover Cluster to automatically fail over among systems on a local site and between sites. Because HP CLX EVA is part of the SQL Server startup process, resource dependencies must be set before storage system virtual disks (vdisks) are accessed. This allows the HP CLX EVA resource associated with the SQL Server instance to come online in the cluster prior to the disk resources. This is done by setting the disk resource dependencies properties so that they all depend on the common HP CLX EVA resource to come online first.

Figure 4 shows that resource dependency is configured using the Failover Cluster manager.

Figure 4. Configuring a disk resource to depend on HP CLX EVA



[Figure 5](#) shows the dependency relationship between one of the SQL data disk resources with the HP CLX EVA resource.

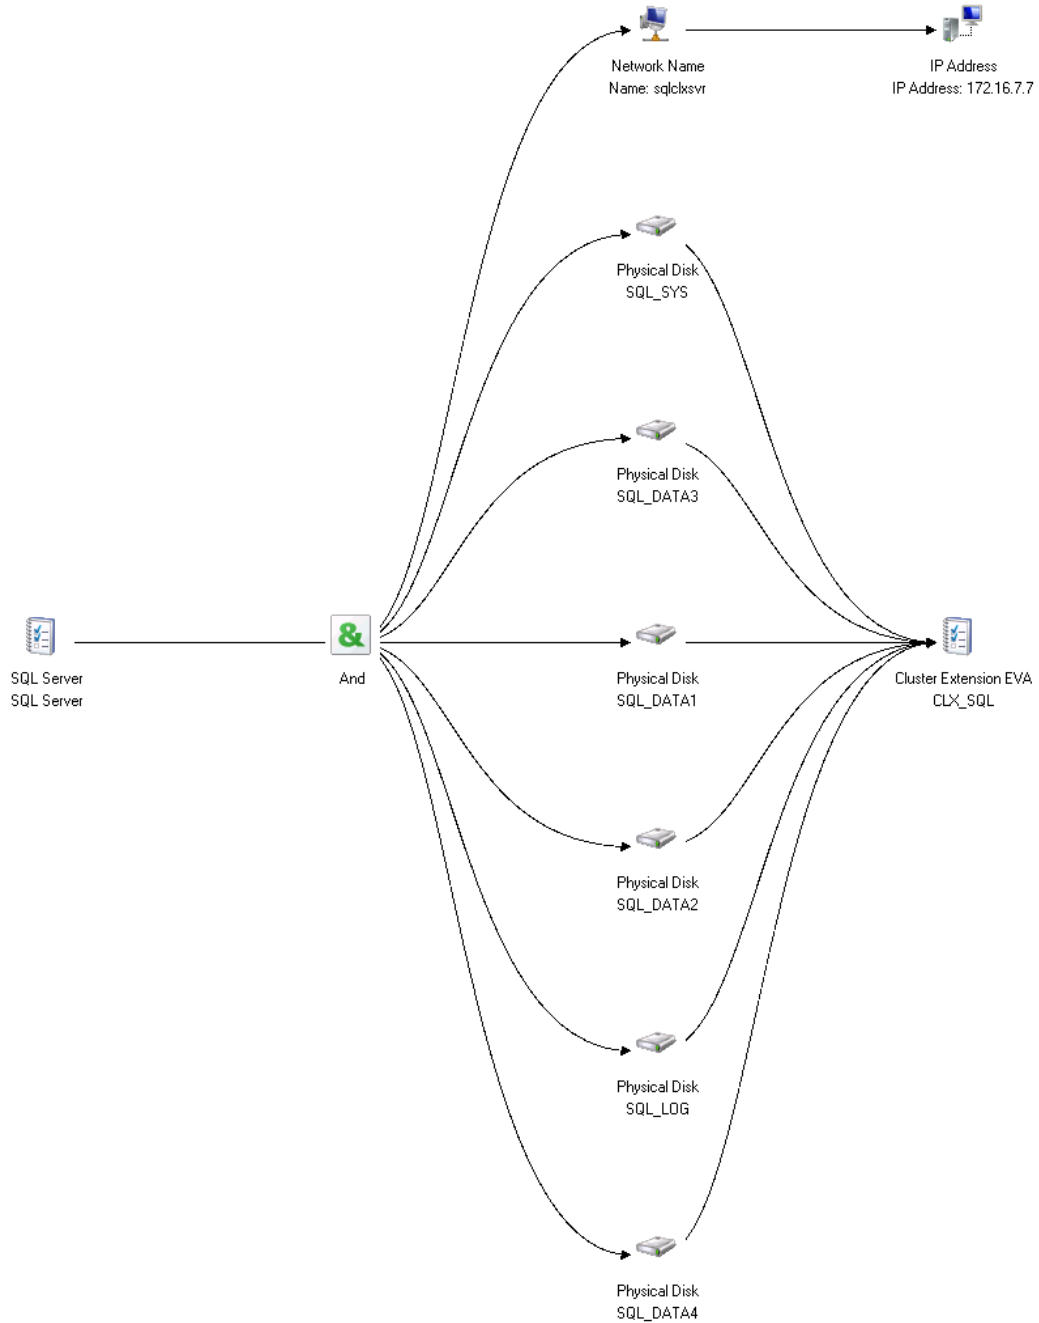
Figure 5. Physical disk resource SQL_DATA1 depends on the HP CLX EVA resource CLX_SQL



Figure 6 shows the entire dependency structure for SQL Server.

Note: The network resources do not depend on the HP CLX EVA resource.

Figure 6. All physical disk resources that SQL Server is dependent on the HP CLSX EVA resource **CLX_SQL**



Application load simulation

All failover testing is performed while SQL transactions are in progress. After the SQL Server 2008 instance fails over to the second data center, transactions are resumed to verify database availability. The OLTP workload is generated by using BenchCraft, which simulates a brokerage firm database.

Testing

The purpose of conducting these tests is to find a single HP CLX EVA configuration that allows fully automated disaster recovery and also to provide users with considerations when selecting a HP Continuous Access EVA and HP CLX EVA configuration. For testing, HP testers conduct failover scenarios that cause a clustered instance of SQL Server 2008 to failover from one data center to the next. Once a successful failover is complete and SQL Server is back online in the recovery data center, failback procedures are documented with the goal of restoring the clustered SQL Server instance back to the primary data center to its state prior to failover. The following failover scenarios are chosen for this testing


- All nodes in a data center fail
- Primary EVA storage array becomes unavailable
- Entire data center fails
- Failover occurs during resynchronization

The disaster scenarios occur during OLTP transactions that are generated with BenchCraft test suite running on a non-clustered host. After Microsoft Failover Cluster reports that the SQL Server is back online in the secondary data center, transaction activity is manually resumed by using BenchCraft to verify that the SQL Server and the database are available. Once SQL Server availability is verified, steps are taken to manually fail back the SQL Server instance to the primary data center to its state prior to the disaster scenario.

Disaster Scenario 1: All nodes in data center fail

Figure 7 shows the two nodes in the primary data center failing while active SQL Server transactions are being replicated from Data Center A to Data Center B

Figure 7. Disaster Scenario 1: All nodes in Data Center A fail



For Disaster Scenario 1, all cluster nodes in the production site (Data Center A) fail while SQL Server is performing transactions. The node failure triggers the SQL Server cluster resource to fail over to Data Center B. For testing purposes, the cluster node blades in Data Center A are ejected from the blade enclosure to simulate an abrupt failure. However, the same type of failure can occur in different scenarios, such as disabling the product

environment network connectivity from the rest of the cluster (including the fileshare witness server), or blade chassis failure. In this type of failure scenario, the SAN environment remains intact.

During the failover process, the HP CLX EVA cluster resource for the SQL Server disk resources fails over the HP Continuous Access EVA data replication direction for the data replication group. The result is that the data replication group is now in source mode in Data Center B, and in destination mode in Data Center A. Once SQL Server is back online, transactions can resume, with all data now replicating from Data Center B to Data Center A. For this type of failover scenario, the user-defined parameters for the HP CLX EVA resource have no effect on failover behavior or takeover action.


Failback process

Because there is no disruption to the SAN in this failover scenario, the SQL Server application can manually move back to the primary node by using Windows Failover Cluster Manager and HP CLX EVA, once communication to the physical nodes in Data Center A is restored and rejoin the failover cluster. Manual failback is performed by using Windows Failover Cluster Manager, which takes the SQL Server instance offline and brings it back online in Data Center A. HP CLX EVA is called by the cluster during the failover operation and reverts the replication direction of the disk resources to replicate from Data Center A to Data Center B, thus reverting the source/destination relationship of the data replication group to the state prior to the initial failover.

Disaster Scenario 2: Primary EVA storage array becomes unavailable

Figure 8 shows the EVA in the primary datacenter failing while active SQL Server transactions are being replicated from Data Center A to Data Center B.

Figure 8. Disaster Scenario 2: Primary EVA storage array becomes unavailable



For Disaster Scenario 2, the cluster nodes in the production site remain online, but the shared storage resource being used by SQL Server becomes unavailable in the data center, which triggers a cluster failover event. Once the cluster node running SQL Server realizes the disk resources are unavailable, Microsoft failover cluster attempts to move the SQL Server

instance to the second node in the list of preferred cluster nodes. If the second node in the list is the same data center as the node previously running data center, the HP CLX EVA resource detects that the local array is no longer available and fails, which triggers Microsoft Failover Cluster to attempt to move the SQL Server instance to the next node. Once a node in Data Center B is reached to bring SQL Server online, HP CLX EVA places the data replication group in *source* mode in Data Center B. HP CLX EVA can only achieve this if the parameter *Use non-current data OK* is set to **YES** in the Cluster Extension EVA resource property.

The *Parameter Use non-current data OK* specifies whether or not the EVA Cluster Extension can fail over to the destination site in cases where the data at the destination site might not be current. The following might occur if an HP Continuous Access EVA link goes down and the resource is moved to the remote data center:

- Setting this value to **YES** means that regardless of whether or not the data is current in the remote data center, the resource comes online.
- Setting the value to **NO** means that if the data is current in the remote data center, the resource comes online even though *Use non-current data OK* is set to **NO**. If the data is not current in the remote data center, the resource will not come online.¹

Once SQL Server successfully fails over and is online in Data Center B, transactions can resume. At this point, the data replication group goes into logging state, because there is no replication communication between the two arrays. All subsequent I/Os are copied to the data replication group write history log until the log becomes full, or until the replication link between the arrays is restored with resynchronization before the log becomes full. If the write history log becomes full, the data replication group is marked to do a full copy. If the link is restored before the data replication group is marked for full copy and resynchronizes, a delta (merge) resynchronization occurs rather than the bit-for-bit copy of the replicated volumes that occurs during a full copy.

Failback process

For this failover scenario, the EVA in Data Center A still functions as though the data replication group is in *source* mode on its side of the replication link, until the connectivity to the SAN is restored. The data replication group moves into a *source/source* state, in which each array functions as though it owns the primary volumes. When this happens, the two following scenarios can occur, which affects the failback operation, depending on the setting of the HP Command View EVA for the data replication group.

Scenario 1 - Suspend on links down is Enabled (default value) prior to failover

The *Suspend on links down* property is found under the *Connections* tab for the data replication group used by the clustered instance of SQL Server (see [Figure 9](#)).

¹Description from the HP StorageWorks EVA Cluster Extension Software administrator guide.

Figure 9. Data replication group connection behavior modification with HP Command View EVA

DR Group Properties

Save changes		?
General Log Members Connections		
Destination system: LE_DCB		
Connection Attributes		
Write mode:	Requested: Synchronous	Actual: Synchronous
Replication I/O:	Resumed	
Destination host access:	Read only	
Suspend on links down:	Disabled	
Suspend on full copy:	Disabled	
Log state:	Not in use	
Connection state:	Good	
UUID:	6005-08b4-0008-0ca2-0001-1000-0216-0000	

By default, the data replication group property *Suspend on links down* is set to **Enabled**. If this property is enabled and the data replication link between the two EVAs is restored, replication remains suspended, and the data replication group remains in source mode in both data centers, which prevents Cluster Extension EVA from being able to fail back the SQL Server instance back to Data Center A. This suspension functionality can be of use, because resynchronization has an impact on I/O response time, as described in the [QL Server 2008 remote replication and high availability using HP Continuous Access EVA and ProLiant Servers](#) white paper. Therefore resynchronization can be manually initiated at a more convenient time, such as during a period of low transaction activity.

Additionally, when *-Suspend on links down* is set to **Enabled**, the data replication group needs to be manually resynchronized by resuming replication I/O before it can fail back. This is achieved by setting the *Replication I/O* parameter to **Resumed**, which is set by using HP Command View EVA under the Connections property of the data replication group that is used by the clustered SQL Server instance (see [Figure 9](#)).

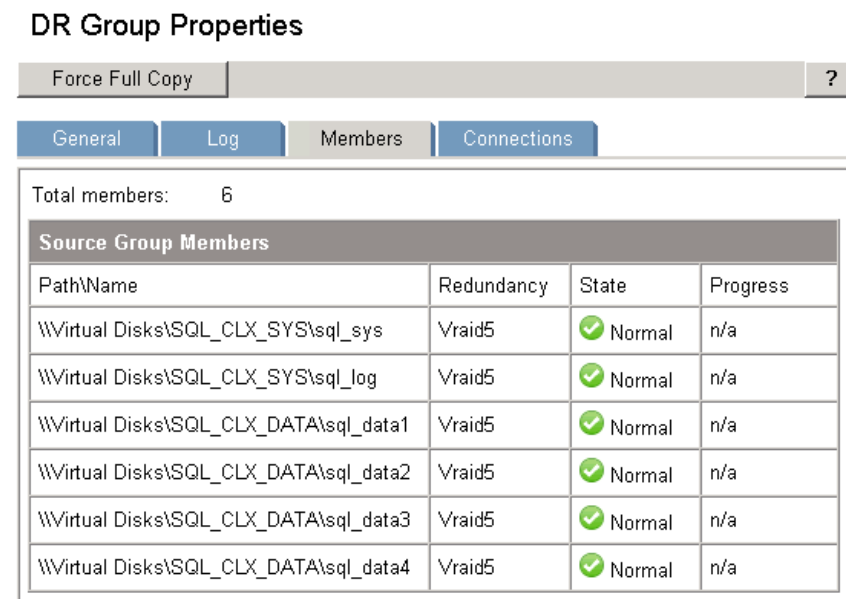
Once replication resumes, the data replication group changes from *source/source* mode to *source/destination* mode, with source being in Data Center B replicating to the destination volumes in Data Center A. A full-copy or merge resynchronization takes place, and the synchronization status can be monitored by using HP Command View EVA or Storage System Scripting Utility (SSSU).

For maximum SQL Server instance availability, resynchronization should complete before attempting to move the database instance back to Data Center A. Otherwise, the SQL Server

instance remains offline until resynchronization completes. After resynchronization completes, with all data replication members (vdisks) now in the Normal state (see [Figure 10](#)), manual failback is performed by using Windows Failover Cluster Manager. Doing so takes the SQL Server instance offline in Data Center B and brings it back online in Data Center A. HP CLX EVA is called by the cluster during the failover operation and reverts the replication direction of the disk resources to replicate from Data Center A to Data Center B, thus enabling disk resources and SQL Server to move back to Data Center A.

[Figure 10](#) shows how a data replication group member synchronization state can be monitored by using HP CommandView EVA.

Figure 10. Data replication group member synchronization state monitored with HP CommandView EVA




Scenario 2- Suspend on links down is disabled prior to failover

If *Suspend on links down* is disabled prior to failover, data replication group resynchronization will take place as soon as the physical link between the two EVAs is reestablished. As in Disaster Scenario 1, it is best to let resynchronization complete before attempting to move the SQL Server instance back to Data Center A. Otherwise the SQL Server instance will be offline until resynchronization completes. After resynchronization has completed, with all data replication members now in the Normal state, manual failback can be performed using Windows Failover Cluster Manager, which takes the SQL Server instance offline in Data Center B and brings it back online in Data Center A. HP CLX EVA is called during the failover operation and reverts the replication direction of the disk resources to replicate from Data Center A to Data Center B.

Disaster Scenario 3: Entire data center failure

Figure 11 shows an entire data center failure in the primary site while active SQL Server transactions are being replicated.

Figure 11. Disaster Scenario 3: Entire data center failure



Disaster Scenario 3 is an event that is basically the combination of Disaster Scenarios 2 and 3, simulating a full data center outage, including all nodes, servers and SAN infrastructure. In this case, the cluster will immediately try to bring the SQL Server instance online in Data Center B. Automatic recovery of the clustered SQL Server instance once again relies on the Cluster Extension EVA parameter *Use non-current data OK* in order for Microsoft Failover Cluster to be able to successfully failover the SQL Server instance and disk resources to Data Center B.


Failback process

For failback to be successful, not only to the cluster nodes in Data Center A need to rejoin the cluster as in Disaster Scenario 1, but planned failback of SQL Server, along with its shared disk resources, once again rely on the data replication group containing the shared disk resources to be synchronized prior to failback as in failover Disaster Scenario 2.

Disaster Scenario 4: Failover during resynchronization

Figure 12 shows both nodes in the primary data center failing while the SQL Server replicated disk resources are resynchronizing from data center A to data center B.

Figure 12. Disaster Scenario 4: Failover during resynchronization

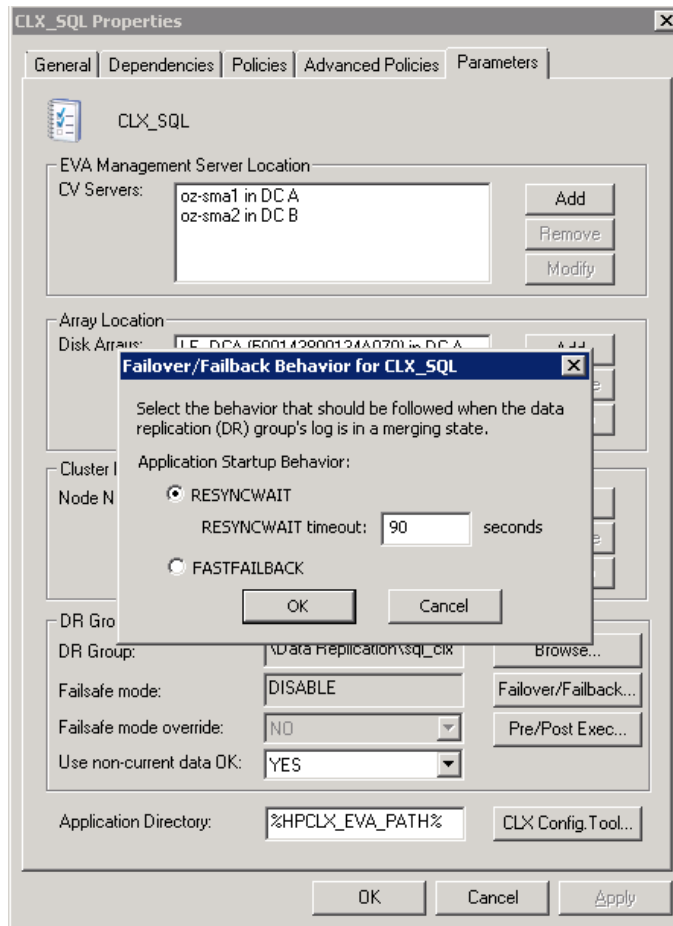


In this scenario, failover behavior of SQL Server is much different from the previous three test scenarios, mainly because failover is attempted during resynchronization of the SQL Server

disk resources. During testing, a full copy resynchronization takes place at the time of node failure in the primary data center. However, failover behavior during a merge resynchronization exhibits similar characteristics. To ensure that automatic recovery of SQL Server occurs in this type of scenario, the HP Cluster Extension EVA resource property for Failover/Failback behavior is set to RESYNCWAIT.

[Figure 13](#) shows the location of the HP Cluster Extension resource property page under the Parameters tab.

Figure 13. The Failover/Failback settings in HP Cluster Extension EVA on the Microsoft Failover Cluster resource properties page



If a single-node failure occurs during an EVA data replication group resynchronization, Cluster Extension EVA allows the cluster to move SQL Server and all of its resources to the second node located in the same data center as the first node. However, in Disaster Scenario 4, both nodes in Data Center A fail, which causes Microsoft Failover Cluster to move its SQL Server resources over to Data Center B. Because the Failover/Failback property of the HP CLX EVA resource is set to *RESYNCWAIT*, once the HP Cluster Extension EVA resource attempts to fail over the data replication group that contains the SQL Server

disk resources, it detects that the data replication group is in a resynchronization state. HP EVA waits until the resynchronization completes before successfully coming online. The frequency in which HP CLX EVA checks resynchronization status can be set in the Failover/Failback->RESYNCWAIT timeout field (in seconds). By default, HP CLX EVA checks resynchronization status every 90 seconds until it completes.

In testing, we found that the failover operation is timed out by the cluster if the resynchronization process takes longer than the pending time-out value for the cluster resource. The pending time-out period depends on the amount of data being resynchronized.

Failback process

Once SQL Server has successfully failed over to Data Center B, data is replicated from Data Center B to Data Center A. All data is now synchronized between both data centers. The recovery process is identical to that of failover Disaster Scenario 1, in that once the cluster nodes in Data Center A have recovered and rejoined the cluster, SQL Server can be manually moved back to Data Center A by using Microsoft Failover Cluster Manager. Before SQL Server comes back online, HP Cluster Extension EVA is called to swap the data replication group relationship so that all of the SQL disk resources replicate from Data Center A to Data Center B.

Best Practices

When configuring SQL Server 2008 in a Microsoft Failover Cluster with HP Cluster extension EVA, the following best practices should be followed:

- Implement existing best practices for performance and availability when configuring disk resources for SQL Server, such as using separate disk groups for your SQL Server log and data files. For more information, see the [Remote replication and high availability best practices for SQL Server 2008 using HP Continuous Access EVA and ProLiant Servers](#) white paper.
- Implement existing best practices for performance in an HP Continuous Access EVA environment by properly setting the CHECK POINT variable in SQL Server. For more information, see the [Remote replication and high availability best practices for SQL Server 2008 using HP Continuous Access EVA and ProLiant Servers](#) white paper.

When configuring your EVA data replication, place all SQL Server disk resources into data replication groups.

- Make sure the Data Replication group property *Destination Host Access* is set to *read only*.
- To control resynchronization after a site failover occurs when the physical path between EVA arrays has been lost, the *Suspend on full-copy* Command View EVA property should be set to Enabled. Manual resynchronization can be performed at a more convenient time by using the *Replication I/O* property via CommandView.

- In order to implement fully automated failover during a site failure, the HP Cluster Extension EVA property *Use non-current data OK* property must be set to **YES**.
- To automate the failover process of SQL Server during a site failure when the data replication group is resynchronizing, the HP Cluster Extension EVA resource property RESYNCWAIT must be enabled. Doing so allows resynchronization to complete before enabling the replicated disk-pair in the failover site.

Conclusion

By seamlessly integrating with Microsoft Failover Cluster, HP Cluster Extension EVA is the best product available to take advantage of the automated disaster recovery functionality of replicated storage provided by HP Continuous Access EVA. For SQL Server to fully utilize the replication and automated recovery features of HP CLX EVA and HP Continuous Access EVA, implement the established best practices from past findings, along with the best practice findings from the tests performed for this paper.

For more information

- Disaster-proof solutions from HP
<http://h71028.www7.hp.com/enterprise/cache/523434-0-0-0-121.html>
- Microsoft Storage Solutions from HP
www.hp.com/storage/microsoft
- HP Customer Focused Testing solutions
<http://h71028.www7.hp.com/enterprise/cache/599679-0-0-225-121.html>
- Remote replication and high availability best practices for SQL Server 2008 using HP Continuous Access EVA and ProLiant Servers
<http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA2-2513ENW.pdf>
- HP Cluster Extension EVA documentation and information
<http://h18000.www1.hp.com/products/storage/software/ceeva/index.html>
- HP Continuous Access EVA documentation and information
<http://h18006.www1.hp.com/products/storage/software/conaccesseva/>

Technology for better business outcomes

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Oracle is a registered trademark of Oracle® Corporation and/or its affiliates. Microsoft® and Microsoft Windows® are U.S. registered trademarks of Microsoft Corporation. Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license. Intel® Xeon® is a trademark of Intel Corporation in the U.S. and other countries.

4AA2-8179ENW, August 2009



Get connected

www.hp.com/go/getconnected

Current HP drivers, support & security alerts
delivered directly to your desktop

