HP StorageWorks Disaster Tolerant Solution for mySAP Business Suite on EVA

Executive summary 3

Business needs 3

Solution design and design rules 3

Component review 4
- Continuous Access EVA 5
- Replicating Oracle databases 5

Replication scenarios for Oracle 8i and 9i with Continuous Access EVA 6
- Replicating the entire Oracle database 6
- Replicating Oracle redo log information only 6

Continuous Access disaster management 7

Solution verification Windows 2000 and Windows Server 2003 9
- Microsoft Cluster Service and Oracle failsafe in an SAP environment 10
- Verification workload 10
- Replicating the entire Oracle database 10
- Replicating Oracle redo log information only 11
- Path failure and Continuous Access EVA normalization 11
- Failover and failback operations 12
- Unplanned site failover (disaster) 12
  - Return operations to the home storage system 14

Solution verification HP-UX 15
- MC/ServiceGuard Extension for SAP R/3 15
- Verification workload 15
- Oracle replication scenarios on HP-UX 15
- Path failure and Continuous Access EVA normalization 15
- Failover and failback operations 15

Solution verification Tru64 UNIX ... 16
TruCluster scripts for SAP R/3 ... 16
Verification workload ... 16
Oracle replication scenarios on Tru64 UNIX ... 16
Path failure and Continuous Access EVA normalization ... 16
Failover and failback operations ... 16

Summary ... 17

Solution-specific configuration Windows 2000 and Windows Server 2003 ... 18
A1: Solution verification hardware ... 18
A2: Solution verification software ... 19
A3: Set up zoning for a Continuous Access configuration ... 19
A4: SAN switch port allocation ... 19
A5: EVA storage map and Oracle database layout ... 19
A6: Set up Secure Path ... 20
A7: MSCS SAP cluster setup ... 21
A8: Set up the Oracle standby database ... 21

Solution-specific configuration HP-UX ... 22
B1: Solution verification hardware ... 22
B2: Solution verification software ... 23
B3: Set up zoning for a Continuous Access configuration ... 23
B4: SAN switch port allocation ... 23
B5: EVA storage map and Oracle database layout ... 23
B6: Set up Secure Path and the HP-UX kit for EVA ... 23
B7: MC/ServiceGuard setup ... 24
B8: Set up the Oracle standby database ... 24

Solution-specific configuration Tru64 UNIX ... 25
C1: Solution verification hardware ... 25
C2: Solution verification software ... 25
C3: Set up zoning for a Continuous Access configuration ... 26
C4: SAN switch port allocation ... 26
C5: EVA storage map and Oracle database layout ... 26
C6: TruCluster script setup ... 26
C7: Set up the Oracle standby database ... 26

For more information ... 27
HP Links ... 27
Oracle Links ... 28
Microsoft Links ... 28
SAP Links ... 28

Executive summary

The HP StorageWorks Disaster Tolerant Solution for mySAP Business Suite enhances a high-availability cluster solution on various platforms based on Oracle 8i and 9i with the disaster tolerant capabilities of HP StorageWorks Continuous Access Enterprise Virtual Array (EVA) maintaining application high performance I/O loads with respect to the distance between the two sites.

Replicating the entire SAP database is a robust, high-performing managed solution for SAP customers. The solution measures failover and recovery time in minutes at a remote computing site. This scenario allows existing HP StorageWorks EVA customers a straightforward enhancement of their environment using Continuous Access EVA. Replicating only SAP database redo log information by Continuous Access EVA using the Oracle Standby Database mechanism also provides disaster tolerant functionality up to the latest transactional update. In addition, this scenario requires less bandwidth when the Storage Area Network (SAN) must span a wider distance and allows database changes to be propagated with a time delay at the alternate site to protect the standby database from human error. The trade-off for this scenario is the additional management effort required to maintain the standby database and the Oracle database expertise necessary in case of a site failover.

Business needs

When data security and availability are critical to the success of their businesses, SAP customers require a computing solution that protects their information systems from disasters, such as power outages, earthquakes, fires, floods, or acts of vandalism. The effects of a disaster range from temporary loss of availability to outright physical destruction of a facility and its assets. In the event of such a disaster, the mySAP Business Suite setup must allow customers to shift their information-processing activities to another site as quickly as possible. Therefore, procedures for disaster recovery must be predictable, well-defined, and immune to human error.

Disaster tolerance based on Continuous Access EVA is characterized by a short recovery time and avoidance of data loss. In a disaster tolerant system based on this approach, redundant, active servers and client interconnects are located at geographically separated sites. As SAP applications produce data, this data is replicated by Continuous Access, whose function is to maintain consistent replicas of the data at each site. Should the system at one site suffer a disaster, SAP instances that were running at the now-disabled site can be failed over to a surviving site that has the resources to support them. The process of failing over a mySAP Business Suite application to the alternate node involves making the application's replicated data accessible and starting instances on the destination node to restore application availability.

Solution design and design rules

This document describes a method for configuring a disaster tolerant mySAP Business Suite landscape distributed over distant computer sites by combining Continuous Access EVA with a high-availability cluster solution. In a cluster configuration using Continuous Access, some member systems reside at one site, and the others reside at a different site. A mySAP Business Suite application can run the database server on the primary (local) site and the corresponding central instance or one dialog instance on the alternate (remote) site. All I/Os occur on the storage subsystem on the primary site under non-disaster conditions. Continuous Access has exclusive access to storage at the alternate site, to which it replicates synchronously the I/O performed on the storage of the primary site. If a significant failure occurs at the primary site, data processing can be resumed at the alternate site where the data is intact and consistent.

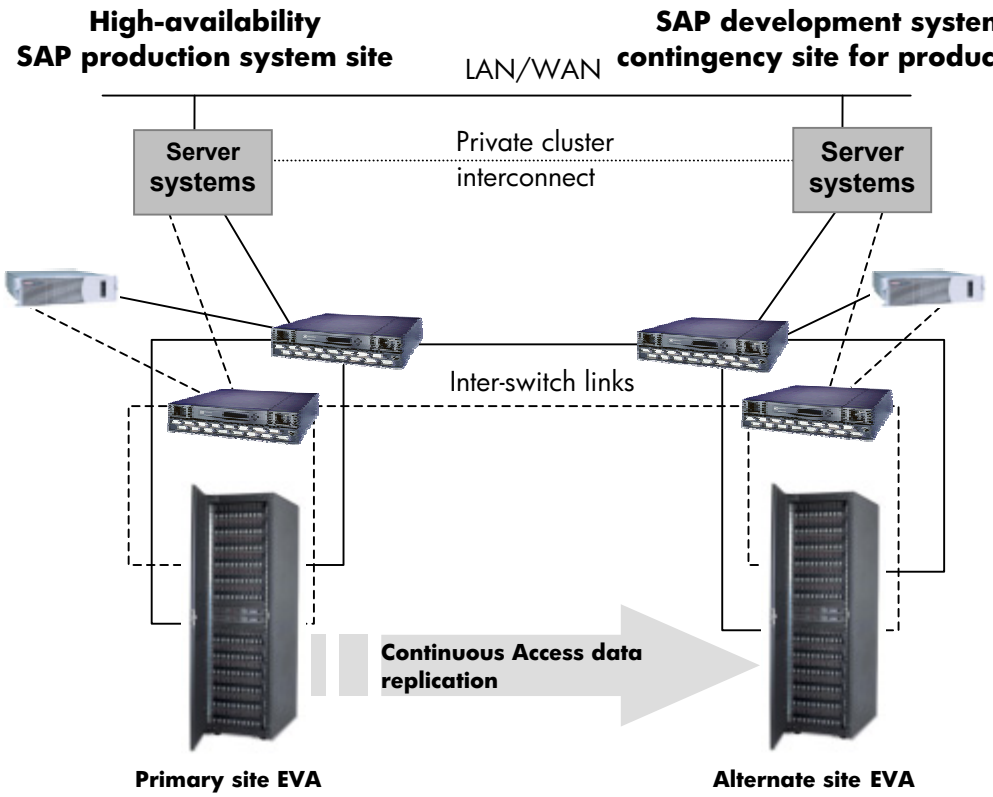

The HP StorageWorks Disaster Tolerant Solution for mySAP Business Suite on EVA takes advantage of the best features of both the Continuous Access and the high-availability technology of the specific-server platform. Cluster members can span distances across a commercial or college campus up to a continental distance depending on the supported platform and the inter-switch link (ISL) type of the SAN. The actual HP StorageWorks Continuous Access EVA data sheet specifies the current certified types of cluster configurations. Data replication hardware ensures correct and consistent mirroring across sites, while the HP management features for various platforms allow you to manage all cluster members, regardless of whether they are at the local or remote site. These capabilities save time during normal system administration and recovery procedures. Although storage failover across sites is a human decision, cluster resources automatically restart mySAP Business Suite applications at the alternate site when the systems are rebooted after a site failover is complete.

Component review

As customer applications and 24 x 7 access to data are business critical to SAP customers, requirements for high-availability solutions with no single point of failure increase. The customers' ability to continue application processing and to maintain data access in the event of a catastrophic disaster becomes critical to their business operations. Disaster tolerant solutions provide high levels of availability with rapid data access recovery, no single point of failure, and continued data processing after the loss of one or more components of a configuration.

Continuous Access EVA

Continuous Access EVA is a controller-based data replication software solution for disaster tolerance and data movement that works with HSV-based EVA5000/EVA3000 storage systems and allows replication groups to be mirrored between pairs of storage arrays that can be in separate geographical locations. Each I/O write access is sent to both storage locations, and reads occur only at the local storage location. Continuous Access EVA copies data online and synchronously or asynchronously in real time to remote locations by a local or extended SAN.

Regard Continuous Access EVA on the HSV controller as enhanced pendant to the HP StorageWorks Data Replication Manager MA/EMA (DRM) product on the HSG80 controller to which the existing HP datasafe solutions for SAP are related. The major enhanced features of Continuous Access EVA compared to DRM in a SAP environment are:

- Enhanced capacity and performance per storage subsystem
- Simplified setup and management of the overall disaster tolerant solution
- Bi-directional replication for a multi-instance SAP landscape

The Comparison of HP StorageWorks Continuous Access Enterprise Virtual Array to HP StorageWorks Data Replication Manager Modular Array/Enterprise Modular Array white paper highlights this in more detail.

Continuous Access EVA supports all major operating system platforms and various options to connect the Fibre Channel switches between the primary and the alternate site.  Continuous Access EVA Design Reference Guide list in the SAN solution checklist section actually supported platforms and ISL options that are valid for a mySAP Business Application as well.

In addition to the idea of disaster tolerance, Continuous Access EVA can complete a backup strategy for a distributed SAP landscape by using the EVA snapshot and cloning capabilities. A detailed implementation is beyond the scope of this paper.

Replicating Oracle databases

The Oracle Databases Replication and Solutions white paper highlights various concepts concerning Oracle database replication scenarios and gives a comparison table for different replication scenarios.

Replication scenarios for Oracle 8i and 9i with Continuous Access EVA

In a Continuous Access EVA environment, two major configuration options exist for replicating the Oracle database synchronously to the alternate site with no potential data loss.

The Oracle Storage Compatibility Testing - Remote Mirroring Using Compaq SANworks Data Replication Manager white paper describes a number of considerations for mirroring the entire Oracle database or "shipping" only the database redo log information. For both scenarios Continuous Access EVA provides crash consistency at all times. As mySAP Business Suite applications are based on the underlying database, these suggestions are also valid in a SAP environment.

Replicating the entire Oracle database

In this configuration, all volumes that contain either Oracle data files, online redo log files, or control files are configured equally at both sites and linked to each other by a copy set on the HSV level. In a database environment all copy sets are treated as a single entity if they are in the same Data Replication (DR) group.

- There is a maximum of 128 copy sets per EVA spread over 128 DR groups with a current maximum of eight copy sets per DR group. This means that one single Oracle database must fit on eight LUNs to be treated as a single entity within one DR group and data consistency between the two sites is ensured.
- In this scenario it is not necessary to replicate the Oracle archived redo log files depending on a customer's backup and restore strategy. If they are replicated it is possible to place them in a second DR group to balance the replication load between the two available fabrics. The trade-off is that after a fabric fails, the remaining one must take the full replication load.
- A real advantage of mirroring the entire database is that it is a much simpler solution to manage because it does not require the maintenance of a second database at the standby site.
- A failover to the alternate site in case of a disaster (Continuous Access unplanned failover) is faster when mirroring the entire database because recovery is similar to a standard Oracle instance recovery for the database after a site failover.

Replicating Oracle redo log information only

Here, the Oracle standby database mechanism replicates only Oracle redo log information to the alternate site by Continuous Access to achieve a disaster tolerant state for the SAP Oracle database. Using the Oracle standby database mechanism without Continuous Access EVA is a common approach at SAP customer sites today. These customers accept that the latest transactional updates in the Oracle database might get lost in the event of a disaster at the primary site. The setup of an Oracle standby database is integrated in the SAPDBA utility.

- In this scenario all LUNs that contain control files and online redo log files must be in the same DR group.
- With Continuous Access EVA, currently no server access to the copy sets on the alternate site is allowed under normal operating conditions. Therefore, the archived redo log information must be copied by way of IP to the alternate site and must be applied regularly. Using Oracle functionality is the most effective method to set up this process.
- The archived redo log files should still be replicated by Continuous Access. In the event of a disaster, there is no guarantee that the latest archived redo log has been completely copied by way of IP to the alternate site before the whole site is lost. As a result, it may be that the Continuous Access replicated online redo log files containing the latest transactional updates could not be applied for this reason.

- One advantage of redo log shipping is that transactional updates can be applied to the alternate database with a time delay. If the primary database information is destroyed through human error, the alternate standby database is protected from this kind of error being propagated immediately and a point-in-time recovery is possible.
- Comparing the two replication scenarios, replicating only redo log information requires less bandwidth between the two sites. This is not that important in a campus environment where the customer is more flexible to increase bandwidth at a moderate cost compared to renting additional bandwidth from a telecommunications company.

The ORACLE8i Standby Database technical report covers all details on log shipping as well as design and planning considerations on this topic.

Continuous Access disaster management

For various reasons it might become necessary to use the service that Continuous Access EVA provides and fail over to the alternate site. Table 1 lists possible unplanned failover situations and the recommended actions in a specific situation. If a type of failure requires a site failover, it is important to verify that all components at the alternate site are operational before a failover is initiated. It might be preferable in some situations to fix a single component within an acceptable timeframe and continue processing, rather than performing a complete failover.

Table 1. Unplanned failover situations and recommendations

Type of failure	Recommended action (Error Mode = Failsafe)
Total primary site loss	Manually fail over data processing to target site
Loss of one primary site fabric	Do not fail over
Loss of primary controller pair	Manually fail over data processing to target site
Loss of all intersite links	Decide which side should continue processing
Total target site loss	Manually continue processing at primary site
Loss of target fabric	Manually continue processing at primary site
Loss of target controller pair	Manually continue processing at primary site
Loss of single primary controller	Do not failover
Loss of a single primary switch	Do not failover
Extended power outage at the primary site	Manually fail over data processing to target site
Loss of a host bus adapter	Do not failover
Loss of single disk in redundant storage	Do not failover
Loss of single host of cluster	Do not failover

An essential part of a Continuous Access–based disaster tolerant solution is the mechanism for managing a planned/unplanned failover or failback operation in the event of a disaster or during maintenance operations. Continuous Access EVA functionality can be managed with either the EVA default management interface command view EVA, the command line interface Storage System Scripting Utility (SSSU) for scripting purposes, or the new graphical Continuous Access User Interface (UI).

Figure 1. Command View EVA

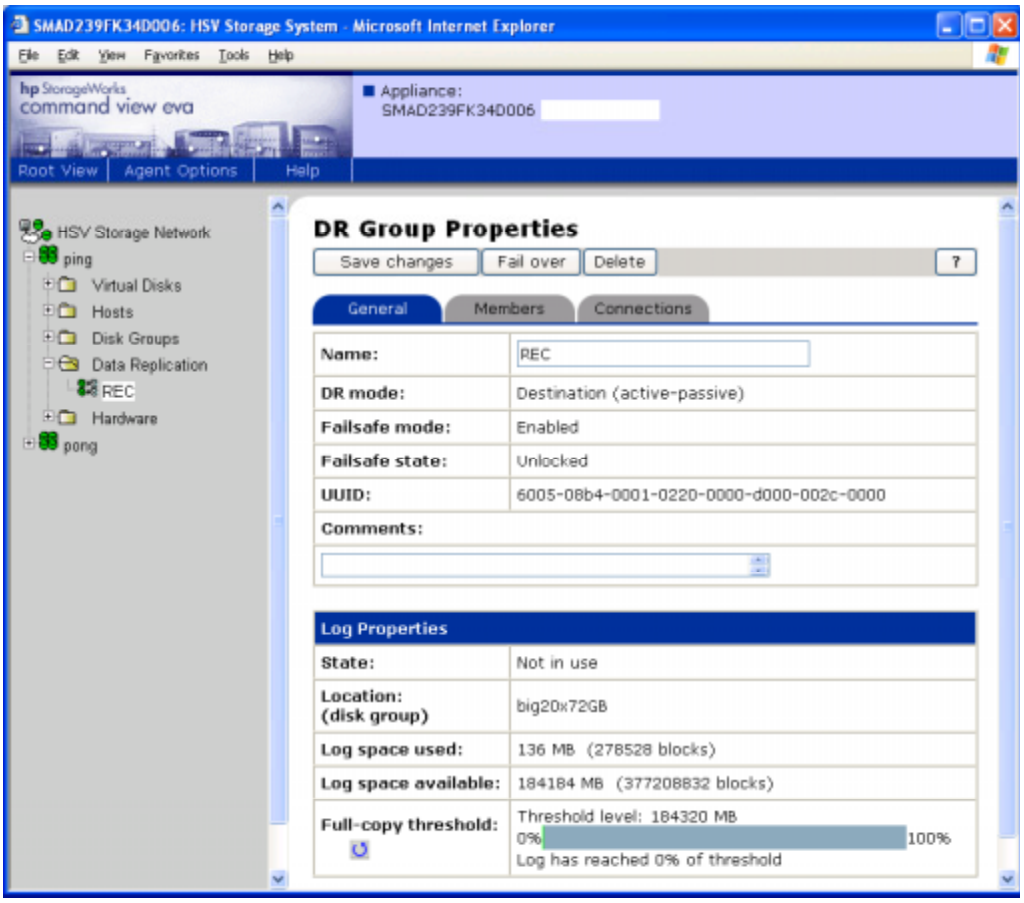

Which interface an SAP customer plans to use depends on the customers' preference and the complexity of the Continuous Access EVA implementation. The advantage of command view EVA is that the customer uses one single interface to manage an EVA for all management activities while the advantage of the Continuous Access UI is that this interface is dedicated to Continuous Access and the probability of a human error in a disaster situation is reduced to a minimum. Furthermore, only UI allows managing multiple DR groups as a single entity. This is interesting for SAP customers replicating more than one SAP database instance or replicating the Oracle archived redo logs in a separate DR group. Use SSSU command line interface for scripting purposes when data structures must be recreated in combination with the powerful CAPTURE command or LUN access (Selective Storage Presentation) must be managed.

Solution verification Windows 2000 and Windows Server 2003

To verify the functionality and performance of Continuous Access EVA for mySAP Business Suite in a Windows 2000 and Windows Server 2003 cluster environment, HP has set up a configuration running an SAP workload with the two different Oracle replication scenarios. The different configurations were tested using direct fiber connections in a 0-km SAN. See the Solution-specific configuration Windows 2000 and Windows Server 2003 section for setup details and software versions. The configuration used for the solution verification is a Microsoft Cluster Service (MSCS) cluster setup with Continuous Access EVA, as shown in Figure 2. Everything within the gray box is a standard MS cluster setup. Under normal operating conditions, the shared storage for the SAP database and the central instance runs at the primary site. One SAP dialog instance has been set up on the cluster member at the alternate site. Virtual disks D and Q, containing the SAP database and the MSCS quorum disk, are replicated through Continuous Access to D' and Q' on the alternate site. Figure 2 displays only two virtual disks to simplify the graph.

Figure 2. Non-disaster situation

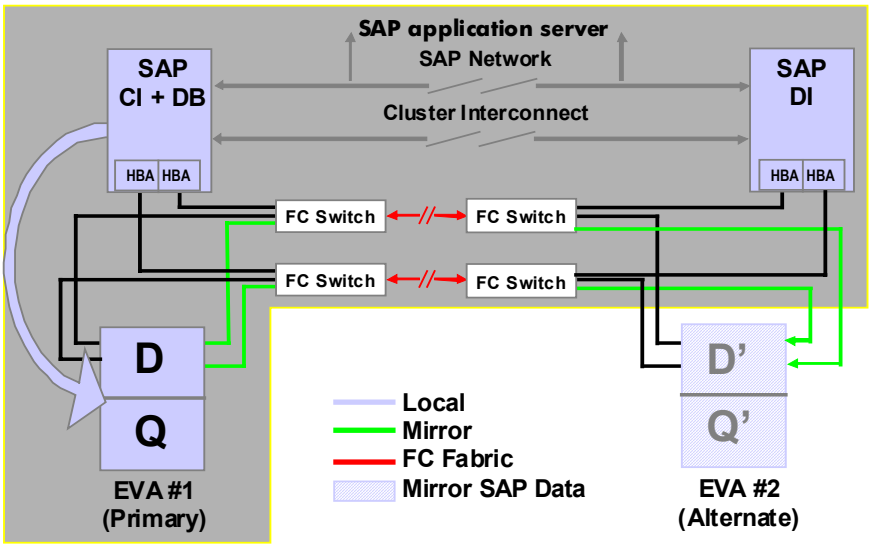

In the event of a node failure on the primary site, the SAP database service starts automatically on the node at the alternate site and accesses the shared storage on the primary site while Continuous Access replication continues. This process could overload the ISLs depending on the SAP load and the distance between the sites because in this situation database access as well as Continuous Access replication might use the same ISL. Various configuration options for this situation include:

- Have an additional cluster member at the primary site in the event of a node having downtime but without being in a disaster situation.
- Have more than one ISL per fabric and use static routes on the Fibre Channel switches or use the 2-Gb Fibre Channel switch products and the licensed *trunking* feature.
- Fail over the entire site in the event of a node failure on the primary site.

Microsoft Cluster Service and Oracle failsafe in an SAP environment

The MSCS provides high availability for services and resources in a two-node advanced server and up to four nodes in a data center configuration. MSCS allows every node in a cluster to be actively running. In case of a failure, the protected SAP database, the central instance, or a dialog instance fails over to a surviving node that assumes the additional workload. The cluster server groups resources, such as network names, IP addresses, or disks, and forms "virtual servers" with which clients communicate. The group or virtual server can run on any physical server at any point in time.

The Oracle Fail Safe product, integrated with MSCS, is responsible for failing over and restarting the SAP database on a surviving node in the solution configuration. The SAP database in an Oracle active-passive configuration with a single instance runs on one of the cluster members.

Verification workload

The verification workload is an SAP ABAP program, started by transaction SE38 in the SAP front-end. This ABAP inserts a specified number of 200-byte records into five Oracle tables containing unique indices into the USER1 table space of an R/3 standard database. This scenario simulates the behavior of a generic R/3 batch job. The tables are deleted and recreated after each run to ensure equal conditions for different runs. The size of the configured SAP Oracle database has no direct impact on the workload.

To verify the solution in terms of functionality and Continuous Access EVA overhead, while not focusing on high-water benchmarking for a specific type of server hardware, the ABAP parameters have been adjusted as follows to ensure that neither the servers nor the network becomes a bottleneck in the verification scenario.

The default workload specifies an insert/update of 1.5 million records by way of three SAP D+W processes. A commit occurs every 1,000 records. In this workload, 5 x 100-MB archived transaction log files are generated and replicated to the alternate site. The ABAP provides wall clock time for the whole run (transaction response time), as well as inserted records per second.

The ABAP program is completed in 41 seconds and provides the 100% write-intensive baseline for a non-Continuous Access scenario. The Continuous Access overhead is calculated using this baseline in the two Continuous Access replication scenarios for Oracle.

Replicating the entire Oracle database

As previously discussed, in this scenario the entire R/3 directory structure and all SAP Oracle database files are replicated by Continuous Access to the alternate site. All remote copy sets are synchronous and in failsafe mode. All database-related copy sets belong to the same DR group.

The ABAP program running the write-intensive workload completes on average after 42 seconds, having generated 5 x 100-MB achieved redo log information replicated to the alternate site. After deleting all remote copy sets to run without Continuous Access, the same workload is completed after 41 seconds, which means that the Continuous Access overhead under this write-intensive workload can be neglected compared to the non-Continuous Access scenario in a zero-latency SAN, as shown in Figure 3. Although not demonstrated in this configuration environment, there is an overhead of Continuous Access EVA that translates into a performance decrease after the distance between the primary and alternate site is increased. The Continuous Access EVA replication performance estimator helps to analyze the effects of distance using Continuous Access EVA.

Figure 3. ABAP completion time in seconds

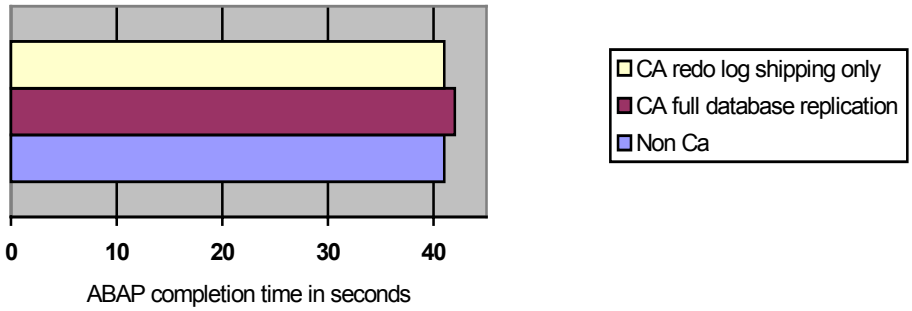

The perfmon utility on the eight CPU database server reported a constant CPU idle time within the range of 30% and an average of 17MB/s throughput with a maximum of 65 MB/s and a peak of 5,500 I/Os per second with an average of 650 disk transfers per second.

Replicating Oracle redo log information only

When replicating only the EVA Vdisks containing Oracle archived redo log, online redo log, and control file information, there is less data to be transferred to the alternate site. As there is hardly no overhead of Continuous Access to be seen with this workload in a zero-latency SAN replicating the entire database, the effect of a performance advantage for redo log shipping can be seen only in a Continuous Access configuration having the intersite links spanning a distance beyond a local or metropolitan area. This is shown in Figure 3. See the considerations in the design trade-off section of the Continuous Access EVA Design Reference Guide.

Path failure and Continuous Access EVA normalization

At both sites, a Continuous Access configuration has no single point of failure in the I/O path from the server to the data on disk. There are at least two paths in two distinct fabrics to ensure that an unplanned site failover (disaster) can only happen if a series of failures occur. To test this functionality in a mySAP Business Suite environment, a path failure was simulated by powering off one Fibre Channel switch while the write-intensive workload was running. The path failure was acknowledged and completed in less than 10 seconds and the running job completed within the range of 51 seconds compared to the 41 seconds in which the write-intensive job is completed without error conditions.

Figure 4. ABAP completion time in seconds

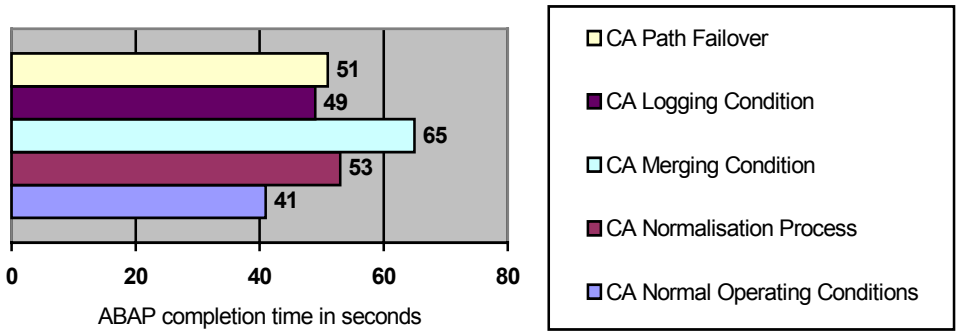

The normalization process in a Continuous Access configuration is not a full copy of all available blocks of a Vdisk between the primary and the alternate site controllers. Only primary site–written blocks containing data are copied to the alternate site. This full copy must happen when a copy set is created or after the two sites are out of synchronization under a disaster condition. In the verification configuration with direct fiber ISLs, the portperfshow utility for the Fibre Channel switches reported up to 80 MB/s during this process. The job completion time of the ABAP program is increased to 53 seconds when a full normalization is in progress during the program run. Therefore, even in a worst-case write-intensive situation, it might be acceptable to continue normal operations in terms of application response time for a SAP customer when a full copy must be initiated. In a 1-Gb fabric environment, the EVA-connected host ports stepped beyond the 100 MB/s border during this scenario.

In a situation when the ISLs are down and failsafe mode must be disabled to continue normal operations at the primary site, the Continuous Access logging process starts automatically. The DR group log space is filled up to 1.7 GB by running the ABAP job, which completes within 56 seconds in a 1-Gb fabric and in 49 seconds within a 2-Gb fabric.

Running the ABAP job while merging the Continuous Access logfile with the alternate site, the job takes 71 seconds to complete in a 1-Gb fabric and 65 seconds in a 2-Gb fabric.

Failover and failback operations

In a Continuous Access EVA environment after a site failover, the data is already available at the alternate site. Failback moves data operations back to the primary storage array after the primary site is online.

Unplanned site failover (disaster)

In the event of a series of failures at the primary site, a total loss of access to the storage on this site might result. This leads to an SAP production system halt situation at the primary site. A human decision must be made to initiate a site failover to the alternate site.

Figure 5. Disaster situation

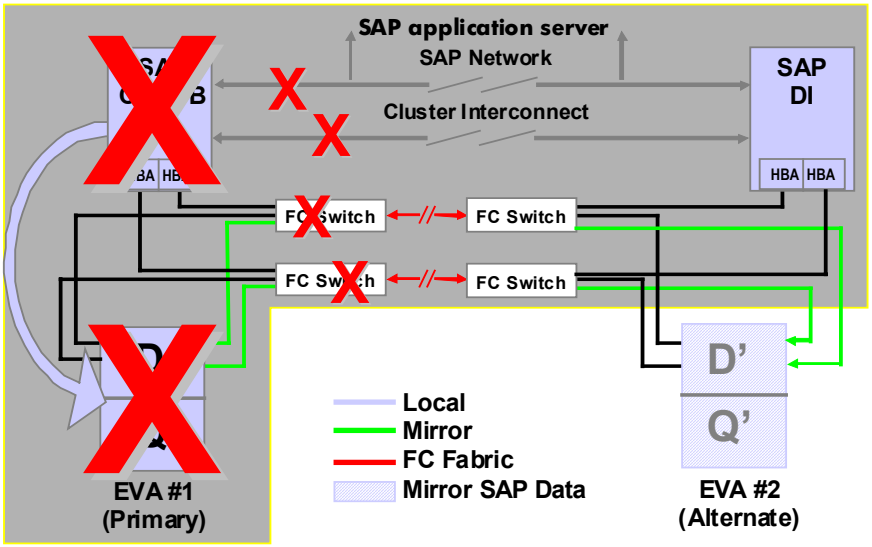

Depending on which Continuous Access management interface is used, the following steps must be taken to complete a site failover for both configurations after the remote HP OpenView Storage Management Appliance has access to alternate site EVA:

1. Press the failover button in either the Continuous Access UI or within Command View EVA or run a SSSU script.
2. Scan the SCSI bus or reboot the systems at the alternate site.
3. Recover the SAP Oracle database :
   – If the entire database has been replicated, the defined MSCS SAP database service will automatically start the Oracle database instance. During database startup, Oracle will automatically perform an instance recovery. The time it takes until the database is available depends on the number of open transactions during the disaster situation.
   – If only redo log information has been replicated, the following steps must be taken:

- Start up and recover the standby database to apply archived redo logs that were not successfully transferred by way of IP while the disaster occurred.
- Shut down the standby database after all archived redo logs have been applied.
- Apply the replicated online redo log information to the standby database using the latest replicated control file from the primary site using the following options:
  SVRMGR> STARTUP MOUNT <SID> ;
  SVRMGR> RECOVER DATABASE ;
  SVRMGR> ALTER DATABASE OPEN ;
- Start the SAP application.

Figure 6. Failover situation

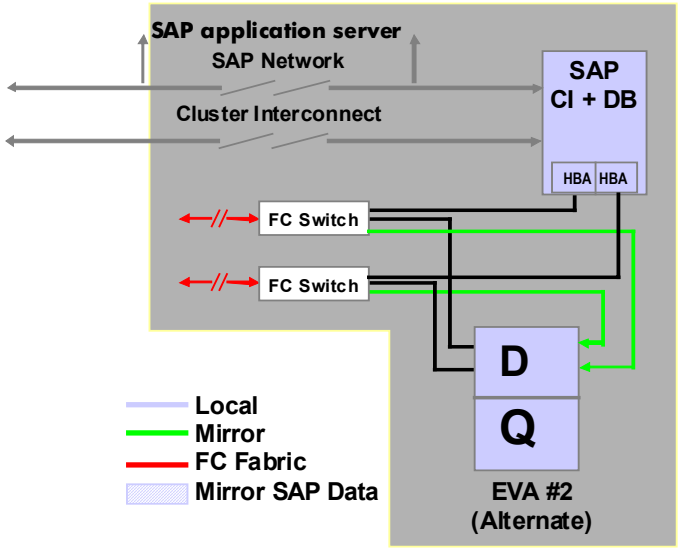

After the site failover has been completed, the SAP database and the central instance are running on the alternate site, as shown in Figure 6. When a node on the primary site becomes available again, it can join the cluster. Although the Continuous Access failover happens immediately, it takes time to scan a bus or reboot a clustered system. The total failover time in the verification scenario is less than 9 minutes for replicating the entire SAP database and within the range of 13 minutes for the standby database scenario, depending on the amount of redo log information that must be applied. In the verification scenario, one 100-MB archived redo log was applied within 25 seconds.

Return operations to the home storage system

After an unplanned site failure has occurred, a Continuous Access configuration still has no single point of failure (Disk, HBA, Cable, HSV), but is no longer in a disaster tolerant state. To achieve this status again, the necessary actions depend on the customer's disaster plan and strategy, the type of disaster that had occurred, and the replication scenario that the customer is using.

To fail back the verification configuration to the original primary site, the following steps are taken after an unplanned site failover when no hardware replacement at the primary site is required:

In the scenario, normalize and fully replicate the entire database.

In the scenario for replicating redo log information only, establish a standby database on the node at the primary site on the primary site storage. The Oracle Storage Compatibility Testing - Remote Mirroring White Paper suggests various ways to achieve this:

- Reverse role by database copy and by restoring backup
- Reverse role by recovery
- Direct fallback by DB copy
- Direct fallback by restoring backup

Fail back the copy sets for this scenario as previously described. The performance of the SAP service will hardly be affected, as only the LUN containing redo log information must be normalized.

Solution verification HP-UX

To verify the functionality and performance of Continuous Access EVA for mySAP Business Suite in an HP-UX environment, HP has set up a configuration running an SAP workload with the two different Oracle replication scenarios. The different configurations were tested using direct fiber connections in a 0-km SAN. See the Solution-specific configuration HP-UX section for setup details and software versions.

The configuration used for the solution verification is a two-node MC/ServiceGuard cluster setup at the primary site and a single server at the alternate site with Continuous Access EVA. Under normal operating conditions, the shared storage for the SAP database and the central instance runs at the primary site. All LUNs containing the SAP database and profiles are replicated through Continuous Access to the alternate site. After a failover in case of a disaster, the single server at the alternate site previously prepared runs both the database and the central instance. A MetroCluster configuration with Continuous Access EVA and MC/ServiceGuard nodes placed at the primary and alternate site as shown in Figure 2 is expected to be certified at the end of this year.

MC/ServiceGuard Extension for SAP R/3

The MC/ServiceGuard Extension for SAP R/3 enhances the MC/ServiceGuard failover capabilities to SAP R/3 environments. The health of each SAP R/3 node is continuously monitored with automatic response to failures or threshold violations. MC/ServiceGuard protects the SAP R/3 central instance and database by defining them in MC/ServiceGuard packages. A detailed description of the functionality and the management of the SGeSAP provide the Managing MC/ServiceGuard Extension for SAP R/3 documentation.

Verification workload

The verification workload for the HP-UX setup is the same SAP ABAP program used for the OS platform previously discussed. Because of the platform independency on ABAP level, this is possible without any changes or porting issues.

Oracle replication scenarios on HP-UX

Replicating the entire Oracle database or replicating Oracle redo log information only slightly differs under HP-UX from other platforms in terms of performance. The results shown in Figure 3 relate to HP-UX as well.

As there is a minimum overhead of Continuous Access to be seen with this workload in a zero-latency SAN when replicating the entire database, the effect of a performance advantage for redo log shipping described earlier can be seen only in a Continuous Access configuration having the intersite links spanning a distance beyond a local or metropolitan area. See the considerations in the design trade-off section of the Continuous Access EVA Design Reference Guide.

Path failure and Continuous Access EVA normalization

The verification tests for fabric, path, and ISL failures previously described have been successfully executed under HP-UX as well. The results shown in Figure 4 reflect the HP-UX test results as well, with the exception that it took slightly longer to get a path failure acknowledged under HP-UX.

Failover and failback operations

Failover and failback operations on the SAP instance in the Continuous Access EVA environment with HP-UX have been successfully verified by using the same mechanisms and steps described in the Unplanned site failover (disaster) section.

Solution verification Tru64 UNIX

To verify the functionality and performance of Continuous Access EVA for mySAP Business Suite in a Tru64 UNIX environment, HP has set up a configuration running an SAP workload with the two different Oracle replication scenarios. The different configurations were tested using direct fiber connections in a 0-km SAN. See the Solution-specific configuration Tru64 UNIX section for setup details and software versions.

The configuration used for the solution verification is a TruCluster cluster setup with Continuous Access EVA, as shown in Figure 2. Everything within the gray box is a standard TruCluster cluster setup with the TruCluster scripts for SAP R/3 installed. Under normal operating conditions, the shared storage for the SAP database and the central instance runs at the primary site. Virtual disks D and Q, containing the SAP database and the TruCluster Quorum device, are replicated through Continuous Access to D' and Q' on the alternate site. Figure 2 displays only two virtual disks to simplify the graph.

TruCluster scripts for SAP R/3

The TruCluster scripts for SAP/R3 enhance the Tru64 UNIX cluster capabilities like Single System Image (SSI) management, the cluster-wide file system (CFS), and Cluster Application Availability (CAA) services to SAP R/3 environments. The health of each SAP R/3 node is continuously monitored with automatic response to failures or threshold violations. A detailed description of the functionality and the installation provide the TruCluster scripts for SAP/R3 documentation.

Verification workload

The verification workload for the Tru64 UNIX setup is the same SAP ABAP program used for the OS platform previously discussed. Because of the platform independency on ABAP level, this is possible without any changes or porting issues.

Oracle replication scenarios on Tru64 UNIX

Replicating the entire Oracle database or replicating Oracle redo log information only slightly differs in the Tru64 UNIX test setup from other platforms in terms of performance. The results shown in Figure 3 relate to Tru64 UNIX as well.

As there is a minimum overhead of Continuous Access to be seen with this workload in a zero-latency SAN when replicating the entire database, the effect of a performance advantage for redo log shipping previously described can be seen only in a Continuous Access configuration having the intersite links spanning a distance beyond a local or metropolitan area. See the considerations in the design trade-off section of the Continuous Access EVA Design Reference Guide.

Path failure and Continuous Access EVA normalization

The verification tests for fabric, path, and ISL failures previously described have been successfully executed under Tru64 UNIX as well. The results shown in Figure 4 reflect the Tru64 UNIX test results.

Failover and failback operations

Failover and failback operations on the SAP instance in the Continuous Access EVA environment with Tru64 UNIX have been successfully verified by using the same mechanisms and steps described in the Unplanned site failover (disaster) section.

Summary

The HP StorageWorks Disaster Tolerant Solution for mySAP Business Suite on EVA (Oracle 8i and 9i) enhances the high-availability features of various OS platforms with the disaster tolerant capabilities of Continuous Access, maintaining high-application performance with respect to the distance between the two sites. The Continuous Access overhead in the described SAP-specific workload can be neglected in a zero-latency SAN, but must be taken into account as the distance between the sites increases.

Replicating the entire SAP database is a robust, managed solution for SAP customers. The solution measures failover/recovery time in minutes at a remote-computing site. A medium-sized SAP configuration placed on six Vdisks can fail over to a recovery site in less than 9 minutes. This scenario allows existing HP StorageWorks EVA customers a straightforward enhancement of their environment using Continuous Access.

Replicating only SAP database redo log information by Continuous Access using the Oracle standby database mechanism also provides disaster tolerant functionality up to the latest transactional update. In addition, this scenario requires less bandwidth when spanning longer distances and allows database changes to be propagated with a time delay at the alternate site to protect the standby database from human error. The trade-off for this scenario is the additional management effort required to maintain the standby database and the Oracle database expertise necessary in the event of a site failover, creating a longer failover time.

An important consideration in a SAP customer's disaster tolerant plan is the necessary time it takes to be in a disaster tolerant state again following a disaster and subsequent failover. Resynchronizing of EVA virtual disks in a zero-latency SAN for a SAP database is in the range of up to 100 MB/s with direct fiber ISLs.

Solution-specific configuration Windows 2000 and Windows Server 2003

A1: Solution verification hardware

The solution is not limited to PL8500 servers or a specific amount of memory or the number of switches. Every HA/F500 enhanced DT configuration meets the requirements of HP StorageWorks Disaster Tolerant Solution for mySAP Business Suite on EVA when a stretched cluster configuration is planned.

	Primary site	**#**	**Alternate site**	**#**
Server	PL8500	1	PL8500	1
	CPU	8	CPU	8
	4-GB memory		4-GB memory	
	KGPSA-CB	2	KGPSA-CB	2
Storage	EVA 2C2D	1	EVA 2C2D	1
	10 K RPM FC disk drives	28	10 K RPM FC disk drives	28
Fibre Channel infrastructure	SAN Switch/16	1	SAN Switch/16	1
	SAN Switch 2/8-EL	1	SAN Switch 2/8-EL	1
SAN management		1		1
Client (SAP)	ML370			
Network	Servers and clients are connected by 10/100 NICs.			

Database I/O performance on the EVA is directly related to the available number of spindles per disk group as outlined in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite. Enhancing the number of spindles in the solution verification configuration provides a higher level of I/O performance in the various test scenarios.

A2: Solution verification software

Software	Version	
Windows	W2K AS SP3	Windows Server 2003
HP StorageWorks Secure Path	4.0	4.0B
Array Controller Version VCS	3.000	3.000
Fabric OS	2.6h/3.0.2k1	2.6h/3.0.2k1
SAP R/3	4.6D SR2	4.7 SR1
Oracle	8.1.7.4.1	9.2.0.2.1

A3: Set up zoning for a Continuous Access configuration

The Continuous Access EVA Des Reference Guide explains when and how zoning on the Fibre Channel switches must be set up in a Continuous Access environment for the HSVs, the Fibre Channel adapters (FCAs), and the management zones for Storage Managements Appliances at the primary and alternate site.

The Continuous Acces lease Notes specify a current restriction that must be followed when a server has more than one FC per fabric:

"All members of a DR group must be presented to the same Fibre Channel adapter (FCA) on hosts with more than one FCA per fabric (for example, multiple FCA pairs, multiple dual-channel FCAs, or a combination of single and dual-channel FCAs). This restriction is required to keep the DR group members using the same host FCA to EVA path, so in the event of a path or controller failure, the members will collectively fail over to the other path."

A4: SAN switch port allocation

	Port	0	1	3	5	7
Primary	TopFabric	PRI TOP1 (HSV)	PRIBOT2 (HSV)	PL85-PRI-HBA1	PRI-SMA-TOP	ISL
	BottomFabric	PRI BOT1 (HSV)	PRITOP2 (HSV)	PL85-PRI-HBA2	PRI-SMA-BOT	ISL
Target	TopFabric	TARTOP1 (HSV)	TARBOT2 (HSV)	PL85TAR-HBA1	TAR-SMA-TOP	ISL
	BottomFabric	TARBOT1 (HSV)	TARTOP2 (HSV)	PL85TAR-HBA2	TAR-SMA-BOT	ISL

A5: EVA storage map and Oracle database layout

Planning and configuring EVA disk groups and Vdisks for this disaster tolerant solution follow the guidelines and best practices referenced in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite and the general EV est Practices. The following graph shows the properties of one of the six Vdisks grouped for th C SAP SID in a Windows 2000 folder.

The Continuous Access Release es specify two restrictions that must be followed when using multi-member DR groups with Microsoft Windows clusters:

- When presenting Vdisks to cluster nodes, all the members of a group must be presented to the same set of FCAs. The group cannot be split across multiple sets of FCAs.
- When making LUN assignments, each shared Vdisk must have the same LUN number on every host.

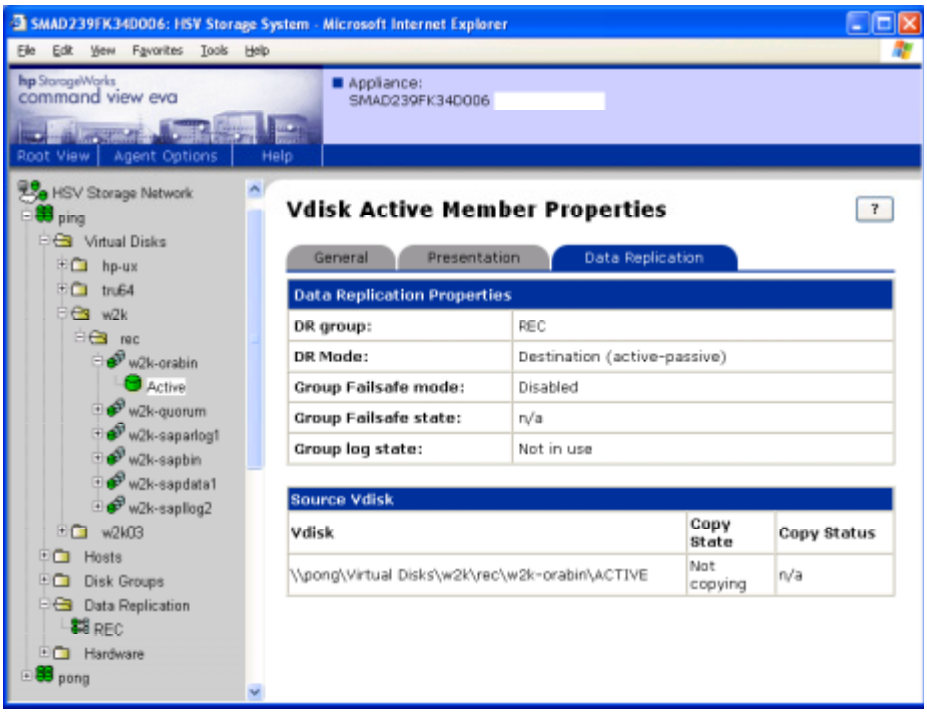

A6: Set up Secure Path

Because only one fabric represents a single point of failure, the maintenance of two separate fabrics is a prerequisite for Continuous Access. To maintain two distinct paths to the storage, Secure Path software provides the functionality for switching between paths in case one path has a problem with completing application I/Os. The following graph shows the Secure Path Manager utility for the solution verification scenario.

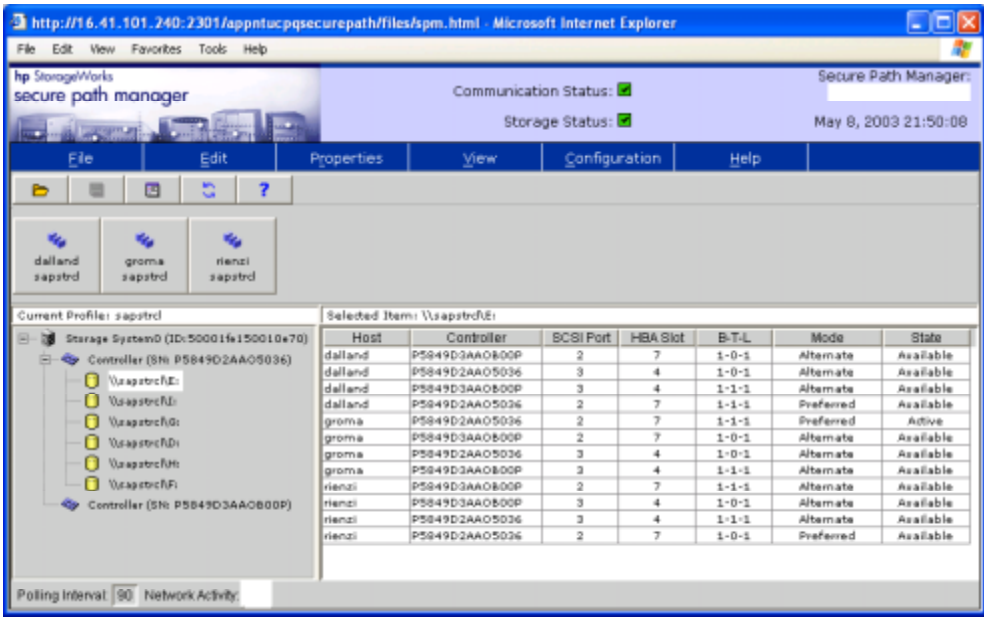

A7: MSCS SAP cluster setup

After the installation of Oracle and Oracle FailSafe, the configuration of SAP in a MSCS environment is configured as shown in the following image.

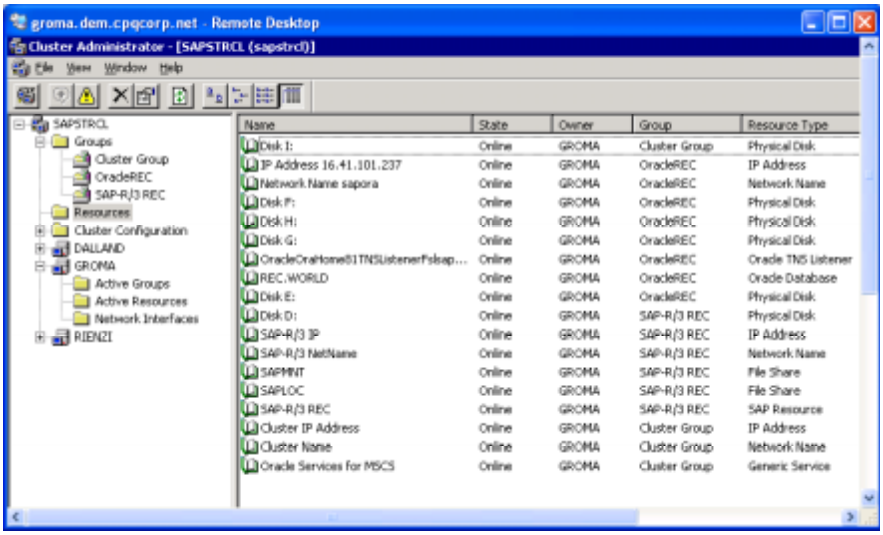

The *R\3 Installation on Windows: Oracle Database* guide in the SAP Library provides a step-by-step description for Oracle and MSCS integration within a mySAP.com environment.

A8: Set up the Oracle standby database

To create a standby database for Oracle on a alternate site:

1. Install the standby server with the same characteristics and the same SAP SID like the production server.
2. Restore primary server datafiles.
3. Restore standby controlfiles after created on the primary database with SVRMGR> ALTER DATABASE CREATE STANDBY CONTROLFILES AS <filename> ;
4. Modify init.ora files (if applicable).
5. Set up primary database tnsnames.ora file and test connection (if applicable).
6. Set up listener on the standby side (if applicable).
7. Mount the standby database using the standby controlfile
   SVRMGR> STARTUP NOMOUNT ;
   SVRMGR> ALTER DATABASE MOUNT STANDBY DATABASE ;
   SVRMGR> RECOVER STANDBY DATABASE ;

The detailed steps for creating the initial standby database can be found in the *Major Preparation* section of the ORACLE8i Standby Database - Technical Report.

Solution-specific configuration HP-UX

B1: Solution verification hardware

The solution is not limited to the verification hardware configuration. An actual overview of the supported server classes, EVA configurations, and Fibre Channel infrastructure can be found in the Continuous Access EVA Design Reference Guide.

The solution verification configuration is based on local clustering. A HP-UX configuration for metropolitan-wide distances based on Continuous Access EVA is expected to be certified at the end of this year.

	Primary site	**#**	**Alternate site**	**#**
Server	rp5430	1	rp5430	1
	750-MHz CPU	2	750-MHz CPU	2
	2-GB memory		2-GB memory	
	A6795A	2	A6795A	2
Storage	EVA 2C2D	1	EVA 2C2D	1
	10 K RPM FC disk drives	28	10K RPM FC disk drives	28
FC infrastructure	SAN Switch/16	1	SAN Switch/16	1
	SAN Switch 2/8-EL	1	SAN Switch 2/8-EL	1
SAN management		1		1
Client (SAP)	ML370			
Network	Servers and clients are connected by 10/100 NICs.			

Database I/O performance on the EVA is directly related to the available number of spindles per disk group as outlined in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite. Enhancing the number of spindles in the solution verification configuration provides a higher level of I/O performance in the various test scenarios.

B2: Solution verification software

The solution is not limited to the verification software configuration. An actual overview of the supported versions can be found in the Continuous Access EVA QuickSpecs.

Software	Version	
HP-UX	11.11i	
MC/ServiceGuard	11.14	
HP StorageWorks Secure Path	3.0B SP1	
Array Controller Version VCS	3.0	
Fabric OS	2.6h/3.0.2k1	
SAP R/3	4.6D SR2	
Oracle	8.1.7.4.1	9.2.0.3.0

B3: Set up zoning for a Continuous Access configuration

See the A3: Set up zoning for a Continuous Access configuration section.

B4: SAN switch port allocation

The following image shows the port allocation of the first eight ports of the primary and target site for the two fabrics.

	Port	0	1	4	5	7
Primary	TopFabric	PRI TOP1 (HSV)	PRIBOT2 (HSV)	HPUX1-PRI-HBA1	PRI-SMA-TOP	ISL
	BottomFabric	PRI BOT1 (HSV)	PRITOP2 (HSV)	HPUX1-PRI-HBA2	PRI-SMA-BOT	ISL
Target	TopFabric	TARTOP1 (HSV)	TARBOT2 (HSV)	HPUX2TAR-HBA1	TAR-SMA-TOP	ISL
	BottomFabric	TARBOT1 (HSV)	TARTOP2 (HSV)	HPUX2TAR-HBA2	TAR-SMA-BOT	ISL

The cable length of the ISLs was 15 m.

B5: EVA storage map and Oracle database layout

Planning and configuring EVA disk groups and Vdisks for this disaster tolerant solution follow the guidelines and best practices referenced in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite and the general EVA best Practices.

B6: Set up Secure Path and the HP-UX kit for EVA

To maintain two distinct paths to the storage, Secure Path software provides the functionality for switching between paths in case one path has a problem with completing application I/Os. The following graph shows the Secure Path Manager utility output for a c20t0d7 sample device having four I/O paths with two adapter td0 and td1.

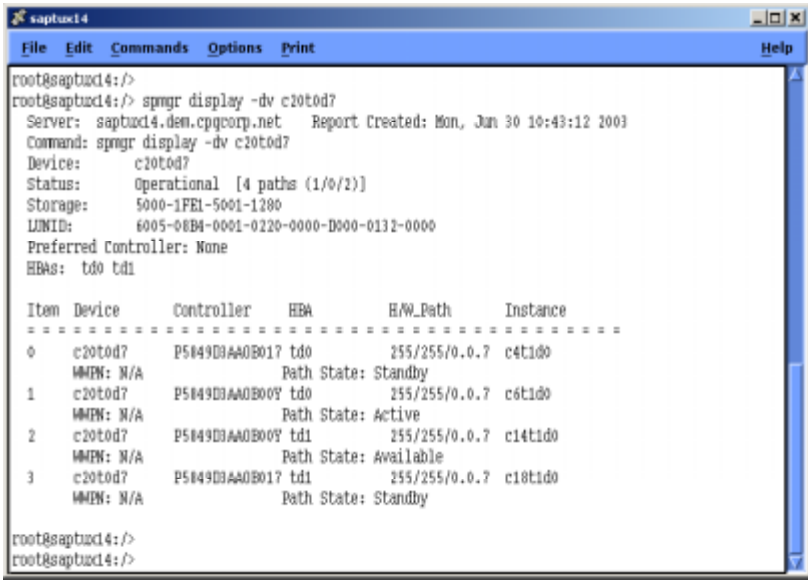

The HP-UX Kit for the EVA installs the supported FCA driver and the SSSU. The HP-UX kit EVA installation and configuration guide and the HP-UX kit for EVA Release Notes specify which HP-UX minimum patch revisions must be installed.

B7: MC/ServiceGuard setup

A detailed step-by-step description of the setup and the management of ServiceGuard in a SAP environment provides the Managing MC/ServiceGuard Extension for SAP R/3 documentation.

The following image shows a SAM view of an SAP package in the verification configuration.

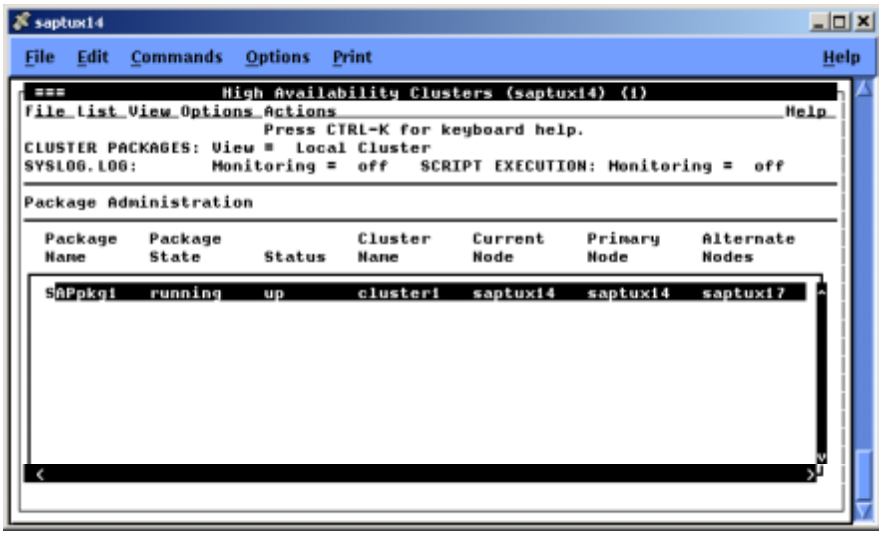

B8: Set up the Oracle standby database

The logical steps for creating a standby database for Oracle on an alternate site under HP-UX follow the same rules as outlined under section A8: Set up the Oracle standby database.

Solution-specific configuration Tru64 UNIX

C1: Solution verification hardware

The solution is not limited to the verification hardware configuration. An actual overview of the supported server types, EVA configurations, and Fibre Channel infrastructure can be found in the Continuous Access EVA Design Reference Guide.

	Primary site	#	Alternate site	#
Server	ES40	1	ES40	1
	600-MHz CPU	2	600-MHz CPU	2
	2-GB memory		2-GB memory	
	KGPSA-CA	2	KGPSA-CA	2
Storage	EVA 2C2D	1	EVA 2C2D	1
	10 K RPM FC disk drives	28	10K RPM FC disk drives	28
FC infrastructure	SAN Switch/16	1	SAN Switch/16	1
	SAN Switch 2/8-EL	1	SAN Switch 2/8-EL	1
SAN management		1		1
Client (SAP)	ML370			
Network	Servers and clients are connected by 10/100 NICs.			

Database I/O performance on the EVA is directly related to the available number of spindles per disk group as outlined in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite. Enhancing the number of spindles in the solution verification configuration provides a higher level of I/O performance in the various test scenarios.

C2: Solution verification software

The solution is not limited to the verification software configuration. An actual overview of the supported versions can be found in the Continuous Access EVA QuickSpecs.

Software	Version	
Tru64 UNIX and TruCluster	5.1B	
TruCluster Scripts for SAP	V011	
Array Controller Version VCS	3.0	
Fabric OS	2.6h/3.0.2k1	
SAP R/3	4.6D SR2	
Oracle	8.1.7.4.1	9.2.0.3.0

C3: Set up zoning for a Continuous Access configuration

See the A3: Set up zoning for a Continuous Access configuration section.

C4: SAN switch port allocation

The following image shows the port allocation of the first eight ports of the primary and target site for the two fabrics.

	Port	0	1	2	5	7
Primary	TopFabric	PRI TOP1 (HSV)	PRIBOT2 (HSV)	TRUX1-PRI-HBA1	PRI-SMA-TOP	ISL
	BottomFabric	PRI BOT1 (HSV)	PRITOP2 (HSV)	TRUX1-PRI-HBA22	PRI-SMA-BOT	ISL
Target	TopFabric	TARTOP1 (HSV)	TARBOT2 (HSV)	TRUX2TAR-HBA1	TAR-SMA-TOP	ISL
	BottomFabric	TARBOT1 (HSV)	TARTOP2 (HSV)	TRUX2TAR-HBA2	TAR-SMA-BOT	ISL

The cable length of the ISLs was 15 m.

C5: EVA storage map and Oracle database layout

Planning and configuring EVA disk groups and Vdisks for this disaster tolerant solution follow the guidelines and best practices referenced in the HP StorageWorks Enterprise Virtual Array Configuration Guide for mySAP Business Suite and the general EVA Practices.

C6: TruCluster script setup

A detailed step-by-step description of the setup and the installation of the TruCluster scripts in a SAP environment provide the TruCluster scripts for SAP/R3 documentation.

The following image shows the implementation of the SAP database (T5O_DB) and the Central Instance (T5O_DVEBMGS00) services in the verification configuration.

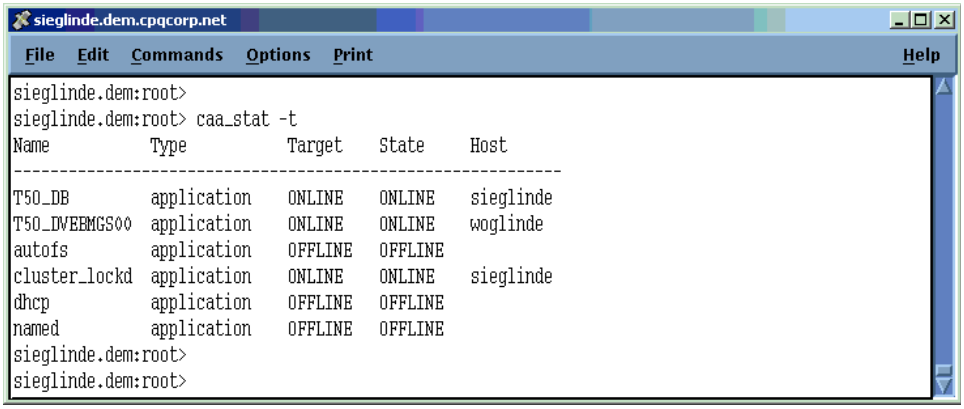

C7: Set up the Oracle standby database

The logical steps for creating a standby database for Oracle on an alternate site under Tru64 UNIX UNIX follow the same rules as outlined under section A8: Set up the Oracle standby database.

For more information

HP Links

HP Network Storage Solutions
www.hp.com/go/storage

HP StorageWorks Continuous Access EVA specifications
http://h18006.www1.hp.com/products/storage/software/conaccesseva/specifications.html

Continuous Access EVA
http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html

HP StorageWorks Continuous Access EVA design reference guide
http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&docIndexId=3124&locale=en_US&prodTypeId=12169&prodSeriesId=316118

HP StorageWorks Enterprise Virtual Array configuration guide for mySAP Business Suite
http://h18006.www1.hp.com/storage/solutionwhitepapers.html

Oracle Databases Replication and Solutions
ftp://ftp.compaq.com/pub/solutions/customsystems/en-orasolrep-wp-02.pdf

Continuous Access EVA Performance Estimator
http://h18006.www1.hp.com/products/storage/software/conaccesseva/relatedinfo.html

Disaster Tolerance-The Technology of Business Continuity
ftp://ftp.compaq.com/pub/products/sanworks/techdoc/drm/12D3-0500A-WWEN.pdf

HP StorageWorks EVA Technical Documentation
http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&docIndexId=3124&locale=en_US&prodTypeId=12169&prodSeriesId=321347

Managing MC/ServiceGuard Extension for SAP R/3
http://docs.hp.com/hpux/pdf/B7885-90013.pdf

TruCluster scripts for SAP/R3
http://saphpcc.bbn.hp.com/Global/ha/hatru/hatru.htm

HP SAN Design Reference Guide
http://h18000.www1.hp.com/products/storageworks/san/documentation.html

SAN product support tables
http://h18000.www1.hp.com/products/storageworks/san/documentation.html

Oracle Links

ORACLE 8i Standby Database
http://technet.oracle.com/deploy/availability/pdf/stby8i_twp.pdf

Microsoft Links

www.microsoft.com/

http://support.microsoft.com/default.aspx?scid=kb;EN-US;q305547

SAP Links

SAP Documentation Library
http://help.sap.com/

5982-1402EN, 09/2003