

Replication Strategies for SAP Based on the Hitachi Virtual Storage Platform

Solution Guide

By James Stephens

November 2010



Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. Be sure to include the title of this white paper in your email message.

Table of Contents

SAP's Three-tier Strategy	4
Client Tier	4
Application Tier.....	4
Database Tier	5
Extending SAP's Strategy: Virtual Storage Infrastructure Tier.....	5
Tier 4-based Data Replication Technologies.....	6
Local Data Replication Solutions	6
Remote Data Replication Solutions	9
Tier 4-based Data Replication Use Cases.....	11
Data Protection.....	11
System Refresh.....	13
Production Instance Migration.....	14
Combinations of Use Cases.....	14
Conclusion	15

Replication Strategies for SAP Based on the Hitachi Virtual Storage Platform

Solution Guide

The mission-critical nature of SAP implementations implies that the primary mission of all SAP infrastructures is data availability. That means that the data must be available for business access 7x24xforever; however, it also means much more. Access to SAP data is often required for multiple uses simultaneously. One of the challenges that SAP administrators must meet is to satisfy those availability requirements without affecting the performance of the primary business process.

Hitachi replication solutions provide effective and efficient solutions to meet those data availability requirements. They can be used to protect against data loss from failure within the infrastructure, user-introduced corruption and catastrophic failure of a data center. In addition, Hitachi replication solutions enable access to SAP data by multiple processes with no negative effect on the primary application's performance.

Hitachi takes a holistic approach to SAP environments, and the key to that approach is the Hitachi Virtual Storage Platform. The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. Scaling out is required when you add new physical or virtual servers to your environment to meet business demands. Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage.

The Hitachi Virtual Storage Platform is the basis for the advanced functionality that fulfills the complex and demanding requirements of a SAP implementation like data replication. Although data replication can be performed at the server level, this function is more effectively performed within the virtual storage infrastructure.

Multiple replication techniques are available within the Virtual Storage Platform. Each technique has its own unique characteristics. The purpose of this paper is to explain each technique. This knowledge is instrumental in selecting the best replication strategy.

This paper is written for individuals charged with planning and implementing SAP deployments. It assumes that the reader has basic knowledge of IT deployments.

SAP's Three-tier Strategy

SAP applications were initially developed as a mainframe solution. When these applications were ported to the open systems platform, SAP implemented its three-tier architecture.

Note that SAP uses the term *tier* in a very different way than storage companies use the term. All references to tier in this white paper are to SAP's tiered approach to application deployment. SAP's three-tier architecture enables deployment of SAP applications in large centralized global implementations that are accessible from anywhere at any time. This architecture allows SAP applications to be scaled to meet the most complex, dispersed and voluminous requirements.

Each tier performs specific tasks in the application stack. Although all tasks can be performed within a single tier, doing so limits the flexibility and scalability of the application. The tasks that are performed in each tier are assigned to that tier to maximize efficiency, performance and functionality.

The implementation of a virtual storage infrastructure extends the SAP three-tier model. The Hitachi Virtual Storage Platform is the key component of the fourth tier that expands the capabilities of SAP implementations.

Client Tier

The client tier is the front end of the SAP application that is presented to the user. The client tier is made up of client servers or devices that manage end-user facing functions. An SAP deployment might have thousands of clients using a variety of devices, including hand-held scanners, label readers, laptops, thin clients and more. Any device that supports end-user I/O is a potential client.

Simple data entry and display functions are performed on data that is provided by the application tier. All data from the client tier passes through the application tier. The client tier has no direct interaction with the storage infrastructure and has no direct effect on the development of the data storage architecture and infrastructure.

Application Tier

The application tier is made up of application servers that manage data manipulation and calculations. It feeds data to the client tier and receives the data from that tier. All data manipulation occurs within the application tier. The application servers in this tier perform specialized functions to optimize the application tier. For example, a server might be dedicated to performing a specific financial function and all clients that require that specific function are automatically connected to that specific server. The SAP application balances those requests among the application servers. If a specific server fails, the load is spread across the remaining servers. An SAP implementation can have hundreds of application servers and these are candidates for server virtualization.

Data used by the application tier is not stored locally but routed through the database tier. Because no application server maintains a static copy of SAP data, these servers do not need and do not have any direct interaction with the storage infrastructure for data. Consequently, the structure of the application tier has no effect on the data storage infrastructure architecture or implementation.

Database Tier

The database servers that make up the database tier are responsible for all data management functions. All data that is kept within the data storage infrastructure is managed by these servers. These servers are generally clustered to provide high availability. The database tier is usually the source of any performance problems and scalability constraints in an SAP deployment. If these servers are tasked with functions such as host-based mirroring, they can become I/O constrained and can have a negative effect on the entire deployment. The limitation of the servers in the database tier adds to the importance of developing a fourth data tier.

Extending SAP's Strategy: Virtual Storage Infrastructure Tier

The logical next step in the evolution of SAP tiered strategy is a fourth virtual storage infrastructure tier.

This tier provides the framework that hosts your SAP mission-critical data. It also provides protection, migration, replication and long-term availability of data. Moving these tasks to a fourth tier is necessary for the massive mission-critical SAP deployments that are common today. This strategy also brings a higher level of efficiency and functionality to SAP implementations.




Figure 1

The Hitachi Virtual Storage Platform is the key component of the virtual storage infrastructure tier. The virtual storage infrastructure tier is based on the advanced storage technologies available from Hitachi Data Systems. The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the new Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. The Virtual Storage Platform provides virtual storage that meets the growing demands of server virtualization.

The growth and consolidation of SAP systems consolidates and intensifies the I/O demand on the storage system. Implementing fourth tier functionality also places more workload on the storage system. Storage systems must be able to dynamically add more storage resources to keep up with I/O and functionality demand. The 3D scaling capability of the Virtual Storage Platform meets that requirement.

Scaling up allows you to implement consolidation, improve utilization of resources, and reduce costs. With the Hitachi Virtual Storage Platform, you can increase performance, capacity and connectivity by individually adding cache, processors (VSDs), connections (Fibre Channel ports) and disks (SSD, SAS HDDs, and SATA HDDs) to the base system. All SAP servers that require access to the storage system have managed access to all these resources, which act as one system managed as a common pool of resources.

Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. By scaling out you can support increased resource needs in SAP environments.

Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform's pool of storage resources. Once virtualized, external data can then be migrated, tiered, replicated and managed by the Virtual Storage Platform. In this manner, older data storage systems can gain a longer useful life. You can extend functionality to lower-cost, lower-function storage systems by virtualizing them behind a Virtual Storage Platform.

The switch matrix architecture of the Virtual Storage Platform makes all of this possible. It connects the basic components, front-end directors, back-end directors, global cache modules and virtual storage directors. You can add redundant pairs of directors and cache modules as required without disruption to connected host servers. All these resources are tightly coupled through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

Tier 4-based Data Replication Technologies

The data availability requirements of SAP implementations cannot be efficiently met with a single type of data replication. The advanced functionality of the Hitachi Virtual Storage Platform includes the capability to perform the multiple types of data replication that are required for SAP implementations. The Virtual Storage Platform's ability to replicate SAP data at the storage infrastructure level is an advanced technology that brings great value to SAP deployments. The primary benefit is that it helps ensure the quality of service of the production environment. Functionality that would otherwise be performed within the data base tier is performed within the virtual data tier. Being able to remove this from the data base tier helps businesses face the challenge of not only ensuring that their data is protected from loss but also that the production instance performance is not negatively affected by any secondary operations such as back up or data mining.

Several distinct tier 4 replication techniques are available. The effective and efficient application of these technologies requires a thorough understanding of the attributes of each technique.

Local Data Replication Solutions

While remote replication receives much of the attention from disaster recovery planners, local, in-system copy solutions add a valuable complementary layer to a data resilience program. For example, point-in-time (PiT) in-system replication disk copies permit you to rapidly recover from logical data corruption introduced by human or application errors, a particular aid in satisfying the stringent requirements characteristic of SAP implementations.

In-system Full Volume Copies

The most often used tier 4 replication technology is a full copy of the primary data, also called a clone. A copy of the primary data is made and the copy is then available to be used by secondary applications. This replication process is accomplished with negligible effect on SAP application performance. The unique value of working with a clone is that once the clone is created, any operations on the clone have no effect on the primary data. On the Hitachi Virtual Storage Platform, Hitachi ShadowImage® In-System Replication software is used to implement clones. Following are the key characteristics of ShadowImage software:

- **RAID-protected full volume copies** — This capability allows the cloned copy to be of a different RAID type or disk drive geometry from the primary volume. This takes advantage of cost reductions by using the correct storage type for the secondary application's requirements of the cloned copy. In addition, full volume copies can reside in Dynamic Provisioning pools to further enhance the capabilities and cost effectiveness of the cloned volumes.
- **Non-disruptive to business operations and to application performance** — Multiple ShadowImage volume pairs can be created and split at the same time to create a consistent multi-volume point in time copy. This allows a non-disruptive consistent split of multiple volumes without affecting production applications.
- **Flexible to your changing needs** — The flexibility of ShadowImage allows the correct matching of cloned volume capabilities with the cost and performance requirements of the secondary application. Also as future drive and SSD enhancements are delivered, ShadowImage will be able to take advantage of the new technologies with minimal effort.
- **Host independent** — ShadowImage is an internal application to the Virtual Storage Platform, and as such has no restrictions on the number or type of hosts using the volumes.
- **Supports external storage on the Virtual Storage Platform** — The Virtual Storage Platform allows all ShadowImage functionality to be performed using external virtualized capacity, enhancing the cost effectiveness of ShadowImage.

Figure 2 shows the full volume copy replication process.

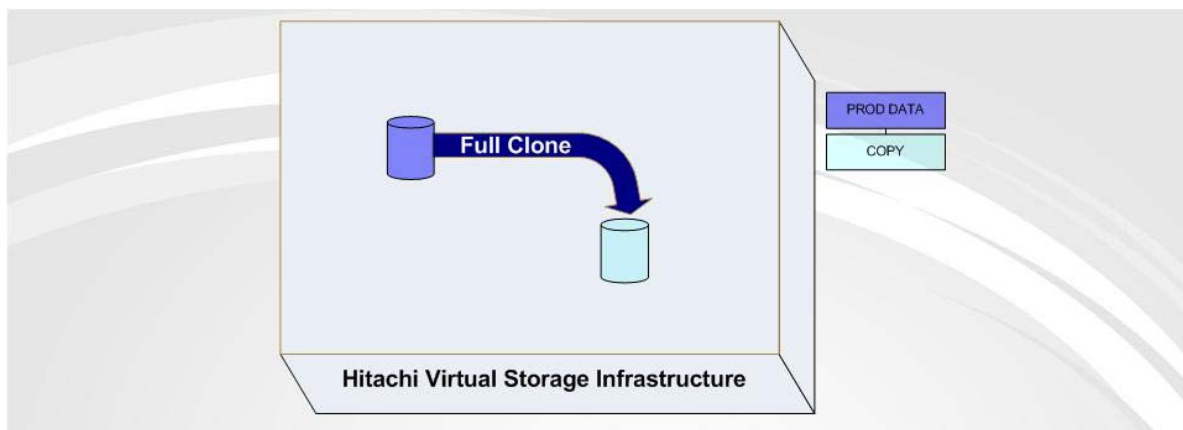


Figure 2

In-system Copy-on-Write Copies (Snapshots)

Snapshots are an effective technique when a point in time virtual copy of the data is required. After a snapshot is created, the Hitachi Virtual Storage Platform tracks any changes to the data and maintains the original data in a pool. The advantage of this technique is that it is an efficient method to maintain point in-time copies. On the Hitachi Virtual Storage Platform, Hitachi Copy-on-Write Snapshot software is used to implement snapshot copies. Following are the key characteristics of Copy-on-Write Snapshot software:

- **Uses less space than full copies or clones** — The Copy-on-Write snapshot preserves the original data when a portion of the primary volume is updated. This is accomplished by creating a pool to house the original data prior to the update. This technique only requires capacity to house changed records. The Copy-on-Write design minimizes storage requirements by saving the original data prior to updating the primary volume with an update. By saving the original data on the first update of a particular block, all succeeding updates to the block are not tracked. When a second point-in-time clone of the same volume is requested, the entire volume is reset for this individual copy while the prior existing copies are maintained, and the process repeats for each copy created. This allows multiple clones to be created with minimal storage utilization.
- **Allows frequent, cost-effective, point-in-time copies** — The pair creation and subsequent split to create a point-in-time copy requires very little overhead to accomplish. Thus, multiple point-in-time copies can be effectively created with minimal overhead.
- **Up to 64 virtual copies of a primary volume** — Copy-On-Write delivers instantaneous creation and restore of just the data you need with versioning of backups for easy restore.
- **Immediate read/write access to virtual copy** — When access to a point-in-time Copy-on-Write clone is requested, the process to present the clone is minimal.
- **Fast restore from any virtual copy** — Copy-On-Write reduces recovery time from data corruption or human error considerably with an immediate restore from a disk-resident point-in-time data snapshot copy.

Figure 3 shows the Copy-on-Write snapshot replication process.

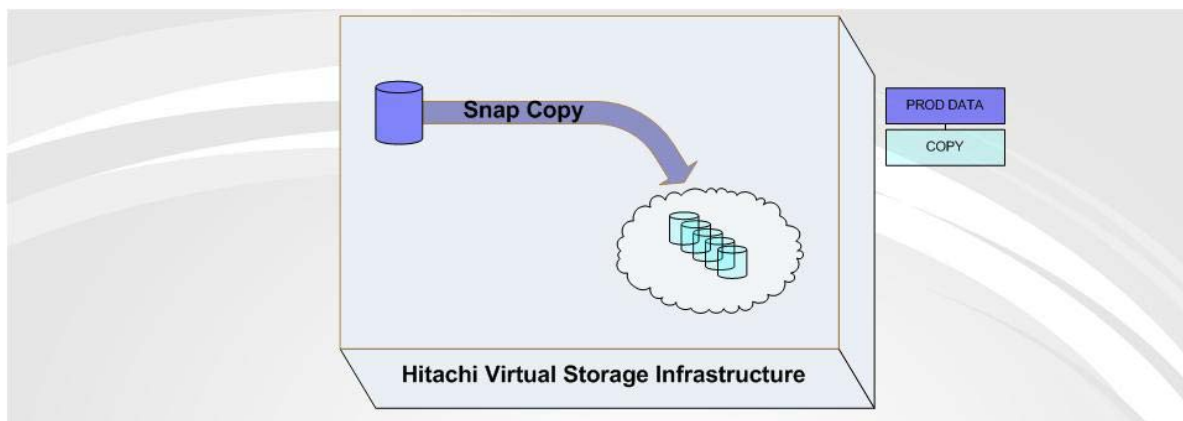


Figure 3

Remote Data Replication Solutions

Replicating data to a remote secondary site represents the most effective insurance policy against catastrophic system failure. Remote data replication offers the fastest recovery time following an outage and the lowest risk of data loss, and it is therefore critical to SAP operations that have aggressive recovery time objectives (RTO) and recovery point objectives (RPO).

Virtual Storage Platform remote replication eliminates the time-consuming, manual and error-prone multistep recovery process required by traditional tape-based backup. It also provides a variety of productivity benefits through secondary access to repurposed data, without affecting regular production workloads.

Remote data replication increases data availability in these ways:

- Automating procedures to reduce the duration of planned events, such as system maintenance, application testing and development, and data backups
- Allowing non-disruptive backup of current production data with no effect or acceptable effect on the production application, depending on whether asynchronous or synchronous replication is selected
- Speeding failover and data restoration in the event of an outage by replacing slow and labor intensive tape-based restores with continuously available online backups
- Allowing secondary sites to take over primary processing to eliminate scheduled downtime
- Enabling frequent, nondisruptive disaster recovery testing with an online copy of current and accurate production data

Two basic variations of remote data replication are available: synchronous and asynchronous.

Synchronous Replication

Synchronous replication ensures that a remote copy of the data, identical to the primary copy, is maintained at the same time the primary copy is updated. In synchronous replication, an I/O update operation is not considered complete until confirmed at both the primary and remote sites, guaranteeing that the remote copy is always an exact image of the primary. This strategy meets most aggressive RPO requirements but has performance implications as the distance between sites increases. On the Hitachi Virtual Storage Platform, Hitachi TrueCopy® Synchronous software is used to implement remote synchronized copies. Following are the key characteristics of TrueCopy Synchronous software:

- **Replication across any distance with some host performance impact** — The replication circuit distance has a direct effect on the application disk response times. When distances exceed practical performance needs, the asynchronous replication alternative must be considered. Another factor affecting performance is the bandwidth of the replication circuit. Even when distances allow adequate responses for synchronous replication, the capacity of the circuit might be exceeded and performance again is affected.
- **Acknowledgement from secondary site required** — For each write operation, the data is secure at the secondary site before the write is acknowledged to the server.
- **No data loss or data integrity issues** — The synchronous replication process ensures that the data at the secondary site is current with the data at the primary site.

Asynchronous Replication

Asynchronous replication meets the challenge of maintaining a remote copy of SAP data without having a performance impact on the primary data. In asynchronous replication, an I/O update operation is considered complete when it is confirmed at the primary site. The Virtual Storage Platform manages the process of replicating the changes to the secondary site. The advantage of this strategy is that distance between sites does not affect the performance of the primary process.

On the Hitachi Virtual Storage Platform, Hitachi Universal Replicator software is used to implement remote asynchronous copies. Following are the key characteristics of Universal Replicator software:

- **Replication across any distance without significant negative effect on host performance** — Hitachi Universal Replicator has been successfully used in replication configurations that span thousands of miles.
- **No acknowledgement dependencies from secondary site** — The Hitachi Universal Replicator mechanism allows replication to a remote site without the performance impact of waiting for each individual record to be acknowledged. Instead, Hitachi Universal Replicator manages the remote relationship at a controller level. If the communication to the remote unit is disrupted or the capability of the replication circuit is exceeded, Hitachi Universal Replicator retains a replicated write data in local journals and catches up write data when the condition is rectified.
- **Potential for some data lag between remote and primary sites, particularly at longer distances** — The RPO is managed with the configuration of the data communication lines. When I/O activity at the primary site exceeds the capacity of the communication channel, the data is staged and moved to the secondary site in the same order as it was written at the primary site.

Figure 4 shows the remote replication process.

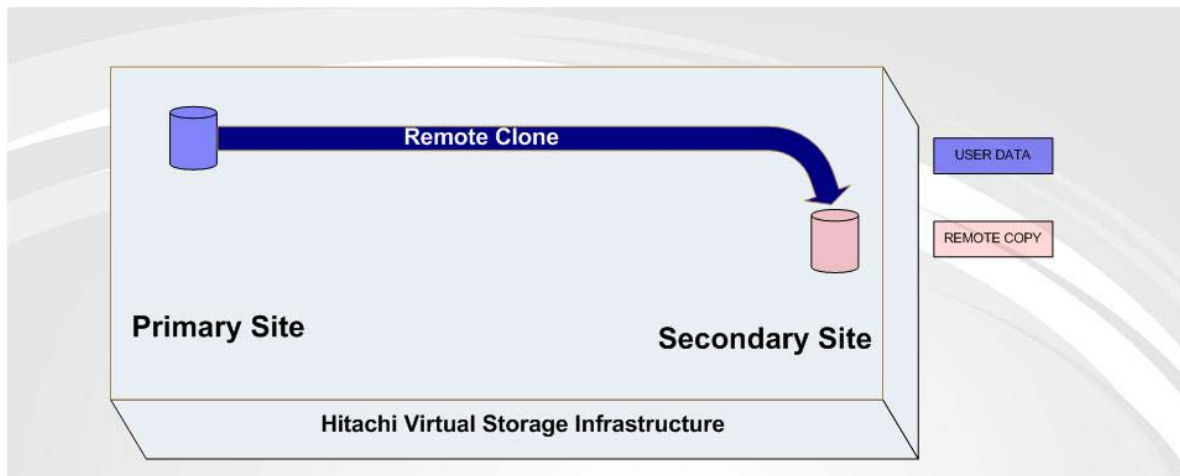


Figure 4

Tier 4-based Data Replication Use Cases

Tier 4 replication can address a variety of business needs. Each specific use case has unique requirements and those requirements determine which replication technology is most appropriate. In SAP environments, the primary concern is the quality of service (QOS) for the production environment. Proper selection and application of replication technology can ensure that the product environment's QOS requirements are met.

Data Protection

SAP data needs to be protected from loss or corruption. When selecting a data protection strategy, the key criteria are RPO and RTO. As with all technology decisions, the cost of the solution is important but is not the most significant concern. It is important to understand how to protect not just against data loss but also loss of access to the data during a recovery operation.

Backup to Tape

Tape backups are the traditional method used to protect against data loss in SAP environments. A daily tape backup process is combined with online transaction logs to meet the RPO. The tape backup is done during system down periods or during periods of low I/O activity. When a recovery is needed, the environment must first be restored from tape and then the transaction logs are applied to recover the transaction activity post backup.

Issues with Tape Backup

The backup-to-tape process can affect the performance of mission-critical SAP operation. In SAP environments, the system downtime window is either small or non-existent. A tape backup can be performed during a production period, but it has the potential to affect the production environment during the backup operation, for two reasons.

The first is that as a backup is performed, the system is online and in hot backup mode, which ensures that a consistent set of data is available for backup. This mode causes additional I/O activity within the application and has a negative effect on the performance. Tape backups can take several hours for SAP databases and this additional I/O activity is required during the entire process.

The second is that during the backup window, two separate processes attempt to access the disk drives simultaneously. The production process is doing the normal business I/O and the backup process is accessing data to transfer it to tape. The advanced technologies of the Hitachi Virtual Storage Platform help to minimize the disk contention issue.

Replication Solutions

The two technologies that potentially address these issues are snap copies and clone copies. The technology that is used to create a snap copy does not require hot backup mode during the backup process. Although the data base must be put in hot backup mode to get a consistent point in time copy the virtual copy can be created almost instantaneously. The production process's performance is not noticeably affected. The remaining issue is not addressed by snap copies. Disk drive contention is still a potential issue with snap copies. Although a point-in-time copy seems to exist, in reality only one physical copy of the data exists. As the tape backup is performed on the virtual copy, only one copy of the data exists and disk drive contention remains as a potential performance issue for the production process.

The application of clone copies is a more suitable solution to resolve these issues. When performance of the production instance is priority, a better solution is to create a full clone of the production data. The clone data is re-synchronized with production data with no negative effect on the production performance. When a point-in-time copy is required, the database is put into hot backup mode for a short period of time while the clone is split. The SAP instance is put into hot backup mode only for that short period of time. The key advantage of clone copies over snap copies is that after that clone is split, the backup-to-tape operation can be performed with no negative effect on the production system.

Online Backup

Although tape backups and redo logs can address the RPO requirements of the business, they cannot address of RTO requirements. Recovering from tape can be a long and tedious operation and access to the production data is lost during this process. This is unacceptable due to the mission-critical nature of SAP deployments. Fortunately, multiple strategies are available to address this issue.

A straightforward solution is to use the clone as a source for recovery in conjunction with the redo logs. For recovery points within the backup cycle, this strategy eliminates the requirement to recover from tape. For failures that require recover points beyond the routine backup cycle, a recovery from tape is still required. An additional shortcoming is that a failure during the relatively short cycle when the clone is being re-synchronized recovery from tape is required. Optionally, multiple clones can be maintained with a round robin approach to online backups. This is an effective but resource-intensive strategy.

To overcome these challenges, snap copies can provide a cost effective and efficient solution for online backups. The purpose of the virtual copies is only for recovery in the case of failures and no secondary processes are required to cause disk conflicts with production. Also multiple virtual copies can be maintained to represent multiple points in time for recovery that minimize the number of redo logs that need to be applied.

In most implementations, tape backup is still required to meet business requirements and those are performed using clones but the virtual copies provide an excellent solution to meet RPOs and RTOs in a cost-effective manner.

Business Continuation

Because SAP deployments are mission-critical, it's important to create a business continuation or disaster recovery strategy. Tape transfers combined with transmission of redo logs can provide a solution but issues of non-structured data and the manpower-intensive nature of rebuilds caused by changes in the production environment limit the effectiveness of this solution. Because SAP disaster recovery strategies also need to be tested and validated and must have a failback strategy, replicating data in tier 3 is an inadequate solution. Instead of server-based or tape-based replication, using the asynchronous remote replication capabilities of the Virtual Storage Platform is ideal.

A second strategy for Business Continuation is to use Synchronous Replication to maintain a copy of data that is updated concurrently with the production data. This type of solution falls more under the umbrella of High Availability solutions rather than DR solutions. Using Synchronous Replication can be an effective solution to reach the most aggressive RPO and RTO requirements. The limiting factor is potential performance impact on the primary data base. All writes must be transmitted to the secondary site and acknowledgement must be received before the write considered complete. The physical limitation of time required to transmit the data and receive acknowledgement and the impact on performance are limiting factors for this solution.

System Refresh

The ability to make data from the production process available to other processes is important in SAP deployments. The challenge is to meet these availability requirements in a timely manner without affecting the production environment.

Reloading tapes or using database data movement technology are solutions, but each has major issues. Reloading data from tapes is a time-consuming and manpower-intensive process. Using the processing power of the database server to move the data can have negative effect on the production process.

Just as in data protection, the ability to create copies of data at the storage infrastructure level provides an excellent solution to meet the business needs of SAP deployments. For system refreshes where the source of the information is production and performance is a priority, a full clone is the strategy of choice. If performance is not a priority, a snapshot can be a cost effective solution. This is often true in non-production environments where multiple copies of data are needed for functions such as testing. The following sections describe sample landscape refresh scenarios.

QA

All SAP deployments require a QA environment that is based on the production data. The QA environment requires a refresh from the production environment on a somewhat regular basis. Using the replication technologies of the storage system is the fastest and most effective method of copying all of the production data into the QA environment. In most cases, pent-up demand exists for more frequent refreshes of QA that is beyond the reach of tape-based replication.

The fact that secondary operations in the QA environment could potentially affect performance in the production environment dictates that clone technology is the correct choice. The synchronization process is performed without affecting the production environment. After that process is complete, the production environment is put into a hot backup mode and the clone is split. This short process does not affect the performance of the production environment. The clone can be used for any purpose without affecting the production environment.

Test, Training and Development

Many environments require multiple copies of SAP data that can be used for concurrent testing or training. In these cases, where performance is not a concern, snapshot technology can be used to create these virtual copies. The strategy is to have a data set that is generally a clone of the production data. Then, multiple virtual copies are created from this complete copy. This can be a cost-effective method to meet the various requirements of test and training environments.

Data Mining

Frequently, a multitude of other processes, for example data mining, require a copy of production landscape data. The I/O requirements of each of the secondary processes generally require that a separate clone is available and dedicated to just that operation. Business needs determine the number of copies needed.

Production Instance Migration

A key need that many SAP implementations eventually face is the requirement to move the location of the production instance to a new or different site. This must be accomplished without loss of availability of the production system during the move.

A key part of that migration must Asynchronous replication is used to create a remote copy of the original site's production data at the new location. The copy process is done with no affect on the performance of the production environment and produces no additional load on the SAP data base server. After the data is synchronized, the transition operation can be completed with minimal interruption to the production environment.

Combinations of Use Cases

These technologies can be combined in a variety of ways, as shown in Figure 5. Note that for the primary data (labeled A), all replication operations are performed without the affecting performance of the primary application.

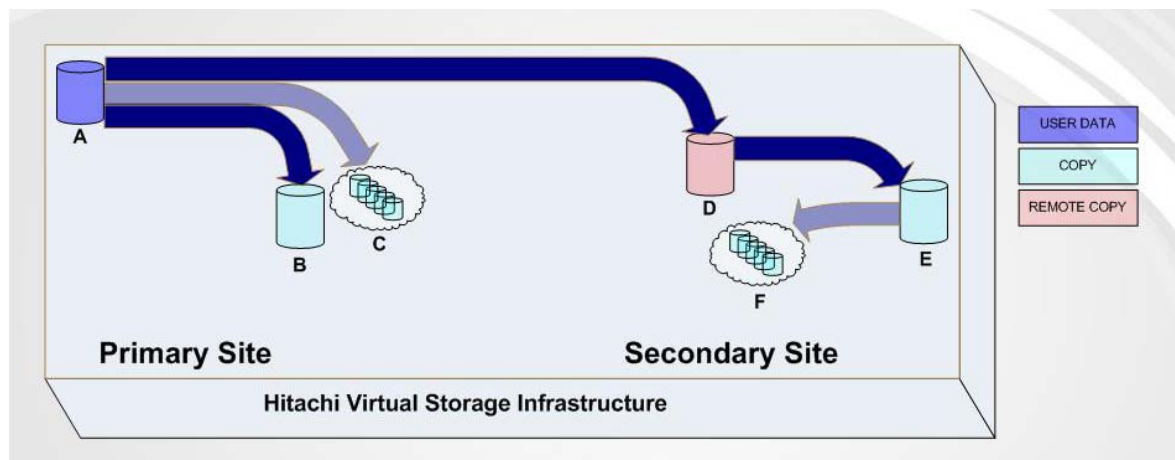


Figure 5

Figure 5 shows the following replication operations:

- **B: Primary data center clone** — A clone of the production data in the primary data center can be used as a source for backup to tape operations or as a QA testing environment. Any operations that are performed on the clone have no performance affect on the production data.
- **C: Primary data center snapshots** — A pool of snapshots of the primary data can be used to create point-in-time copies of the production data for rapid recovery from data corruptions. Because no secondary processes are associated with these volumes, application performance is not affected.

- **D: Remote replication** — A remote copy of the primary data is most often associated with a disaster recovery plan. When asynchronous replication is used to maintain this copy, the performance of the production application is not affected.
- **E: Secondary data center clone** — Clones at the secondary site can be used as a source for backup to tape, testing or data mining operations.
- **F: Clone-based snapshots** — In cases where multiple processes require read and write access to data, snapshots of a clone can be used. When these snapshots are based on a clone, the performance of the production application is not affected.

Conclusion

The data availability demands of SAP implementations are multiple and challenging. Meeting these demands is a key requirement of any SAP infrastructure. You can use the Hitachi Virtual Storage Platform and the associated replication tools to create efficient and cost-effective solutions to meet those demands. For more information about implementing the SAP replication strategies described in this paper, see the Hitachi Data Systems SAP Solutions [web site](#).

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services [web site](#).

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources [web site](#). Click the **Product Demos** tab for a list of available recorded demonstrations.

For more information about Hitachi products, visit the Hitachi Data Systems [web site](#) or contact your sales representative or channel partner.

 **Hitachi Data Systems Corporation**

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-058-00 November 2010

Corporate Headquarters

750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com