Technical Report

SAP with Oracle on UNIX and FCP and NetApp Storage

SAP Competence Center, NetApp
February 2010 | TR-3533

BEST PRACTICES

This document provides customers and partners with the best practices for deploying NetApp® storage in support of SAP® Business Suite solutions running in a UNIX® and FCP environment using an Oracle® Database.

TABLE OF CONTENTS

1 INTRODUCTION

This document addresses the challenges of designing storage solutions to support SAP Business Suite products using an Oracle Database. The primary focus is on the common storage infrastructure design, deployment, operation, and management challenges faced by business and IT leaders utilizing the latest generation of SAP solutions. Recommendations are generic and are specific neither to any given SAP application nor to the size and scope of the SAP implementation. This guide assumes a basic understanding of the technology and operation of NetApp and SAP products and was developed based on the interaction of technical staff from NetApp, SAP, Oracle, and our customers.

BUSINESS CHALLENGES FACING THE SAP CUSTOMER

Corporations deploying SAP software today are under pressure to reduce cost, minimize risk, and control change by accelerating deployments and increasing the availability of their SAP landscapes. Changing market conditions, restructuring activities, and mergers and acquisitions often result in the creation of new SAP landscapes based on the SAP NetWeaver® platform. Deployment of these business solutions usually exceeds a single production instance of SAP. Business process owners and project managers must coordinate with IT management to optimize the scheduling and availability of systems to support rapid prototyping and development, frequent parallel testing or troubleshooting, and appropriate levels of end user training. The ability to access these systems as project schedules dictate with current data sets and without affecting production operations often determines whether SAP projects are delivered on time and within budget.

TECHNOLOGY CHALLENGES OF AN EXPANDING SAP LANDSCAPE

A typical SAP production landscape today consists of several different SAP systems. Just as important as the successful operation and management of these production instances are the many nonproduction instances used to support them.

SAP recommends customers to maintain separate development and test instances for each production instance. In practice, standard SAP three-system (development, quality assurance, and production) landscapes often expand to include separate instances such as sandbox and user training systems. It is also common to have multiple development instances as well as more than one system used for quality assurance, testing, or perhaps a final staging system prior to releasing applications into production. Compound this with the many different SAP applications, such as ERP, CRM, BI, SCM, SRM, and Enterprise Portal, and the number of systems to support can become very large.

Adding to the challenge of maintaining these SAP systems is the fact that each of these instances has different performance and availability requirements. These requirements vary depending on the phase of project and whether the project is focused on an existing SAP implementation or a new one. Projects rely on frequent refreshes of the nonproduction instances so that testing and training can occur with the most current data.

As more test and training systems are needed to accelerate test cycles by allowing parallel independent operation, the demand on the IT infrastructure increases. If the infrastructure that is supporting SAP systems and related applications is inflexible, expensive, and difficult to operate or manage, the ability of business owners to deploy new and improve existing business processes might be restricted.

As SAP landscapes have expanded, the technology also has changed. SAP has evolved to take advantage of the latest technology trends. Along with traditional ABAP-based systems, SAP also has Java instances that need to be managed. Database technologies such as Oracle Real Application Clusters have introduced additional complexity into the database layer. Virtualization or cloud technologies have become more predominant as corporations look to leverage efficient computing methods to maximize their investment and reduce data center expenses. Without a storage infrastructure that can adapt to the needs of the changing technology, IT organizations would be unable to meet the business needs of the company.

NETAPP SOLUTIONS FOR SAP

NetApp minimizes or eliminates many of the IT barriers associated with deploying new or improved business processes and applications. The combination of SAP solutions based on the NetWeaver platform and a

simplified and flexible NetApp storage infrastructure allows business owners and IT departments to work more efficiently and effectively toward the goal of improving enterprise business processes.

Storage consolidation with NetApp assures the high availability and performance of SAP data and applications so that stringent service-level agreements (SLAs) are met. In addition, NetApp helps to reduce the administration and management costs associated with deploying these new business applications and processes.

2 STORAGE PROVISIONING AND MANAGEMENT

2.1 STORAGE VIRTUALIZATION

In today's rapidly changing business climate, enterprises demand cost-effective, flexible data management solutions that can handle the unpredictable and explosive growth of storage in heterogeneous environments. To enable global data management, provide business continuity, satisfy regulatory and compliance standards, and improve resource utilization, a flexible and scalable storage network solution is required. The solution must also minimize complexity and reduce total cost of ownership (TCO).

NetApp offers highly available, scalable, and cost-effective storage consolidation solutions that incorporate the NetApp unified storage platform and the feature-rich functionality of data and resource management software to deliver storage that improves enterprise productivity, performance, and profitability, while providing investment protection and enhanced asset utilization. NetApp enterprise-class storage solutions are proven interoperable across all platforms. NetApp fabric-attached storage (FAS) systems integrate easily into a complex enterprise and simultaneously support NAS, SAN (Fibre Channel), Fibre Channel over Ethernet (FCoE), and IP SAN (iSCSI) protocols.

NetApp FlexVol® technology delivers true storage virtualization solutions that can lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of the enterprise. FlexVol technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks (aggregate).

Based on the number of physical disks within the aggregate, a certain amount of capacity and throughput performance is available. The capacity and throughput performance are available as a resource pool for all FlexVol volumes that are created on top of the aggregate. The FlexVol volumes are created within a few seconds, and the data is always striped over all available spindles in the aggregate, eliminating issuessuch as “hot spindles” or “hot RAID groups” based on badly distributed tables, table-spaces, or data files. The size of the FlexVol volumes can also be easily increased or decreased.

This flexibility means that operations can be simplified, utilization and efficiency can be increased, and changes can be applied more quickly and seamlessly. NetApp storage solutions enable customers to add storage when and where they need it, without disruption and at the lowest incremental cost.

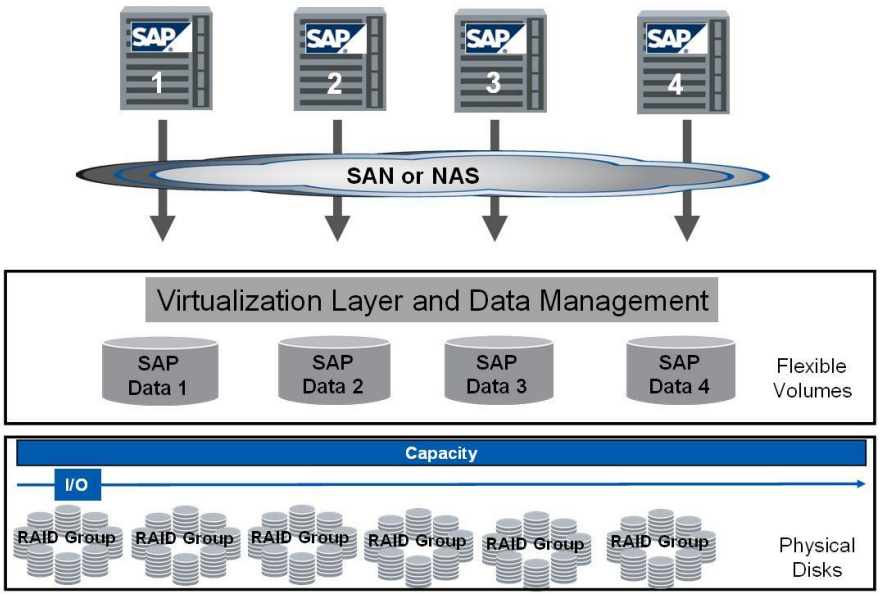

Figure 1) FlexVol technology.

Most data management functionality such as:

- Snapshot™, and SnapRestore® for backup and restore
- SnapVault® or SnapMirror® for replication as disk-to-disk backup or disaster recovery
- FlexClone® for SAP system copies
- FlexShare® for prioritization of FlexVol volumes
- and others

is done on FlexVol volume level. Physical disks and RAID groups are configured once with the basic setup and will then not be part of daily storage management any more.

WORKLOAD PRIORITIZATION

FlexShare is a Data ONTAP® software feature that provides workload prioritization for a storage system. FlexShare gives administrators the ability to leverage existing infrastructure and increase processing utilization without sacrificing the performance required to meet critical business needs. It prioritizes processing resources for key services when the system is under heavy load. With the use of FlexShare, administrators can confidently consolidate different applications and data sets on a single storage system. FlexShare makes it possible for administrators to prioritize applications based on how critical they are for the business. For example, production SAP systems are configured with a higher priority than test and development systems.

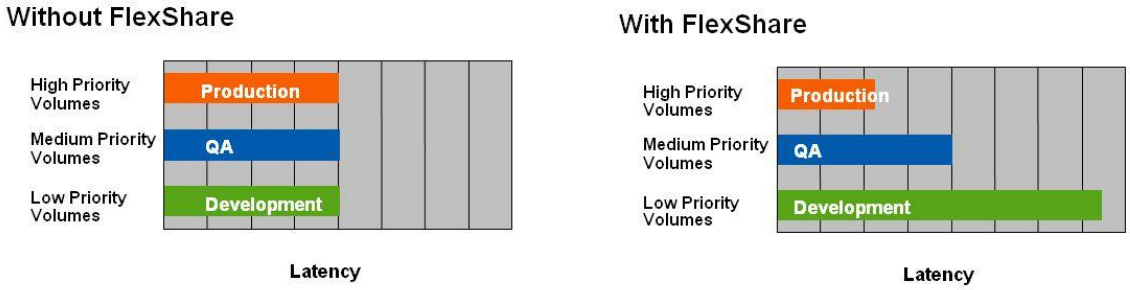

Figure 2) FlexShare.

STORAGE MANAGEMENT

SnapDrive® software offers a rich set of capabilities to virtualize and enhance storage management for Microsoft® Windows®, Linux®, and UNIX environments. SnapDrive tightly integrates with the native file system and provides a layer of abstraction between application data and physical storage associated with that data.

Business does not have to stop every time the IT organization has to add more storage. With SnapDrive, adding, deleting, mapping, unmapping, and mirroring virtual disks can be done online. Expanding capacity can be done with limited or no effect on application or system performance.

SnapDrive integrates Snapshot technology to give near-instantaneous point-in-time images of application and user data. SnapDrive also gives access to Snapshot copies by mounting them as virtual disks. These virtual disks can be used for routine administrative tasks such as online backup, testing of new applications, or population of data marts with limited or no downtime to your business-critical information. Restoring data can be done in minutes with SnapRestore and SnapDrive.

SnapDrive makes management simple and intuitive in the Windows environment by allowing administration through the Microsoft Management Console or a command line. Interactive wizards and easy to-use interfaces guide your administrators through all management tasks and create automatic schedules of operations. SnapDrive comes with an intuitive command line for script-based automation in the UNIX environment.

2.2 STORAGE LAYOUT

AGGREGATE LAYOUT

NetApp recommends using a few, large aggregates for all data of all SAP systems. The use of a large aggregate provides the performance benefits of all available disk spindles in the aggregate to every FlexVol volume in that aggregate. Adding a second aggregate is recommended only if the maximum capacity of the first aggregate is reached.

The aggregates are configured with RAID-DP®, which offers a high level of data protection. Only if three disks within the same RAID group fail at the same time data loss will occur. For information on RAID-DP, see NetApp Data Protection: Double-Parity RAID for Enhanced Data Protection with RAID-DP.

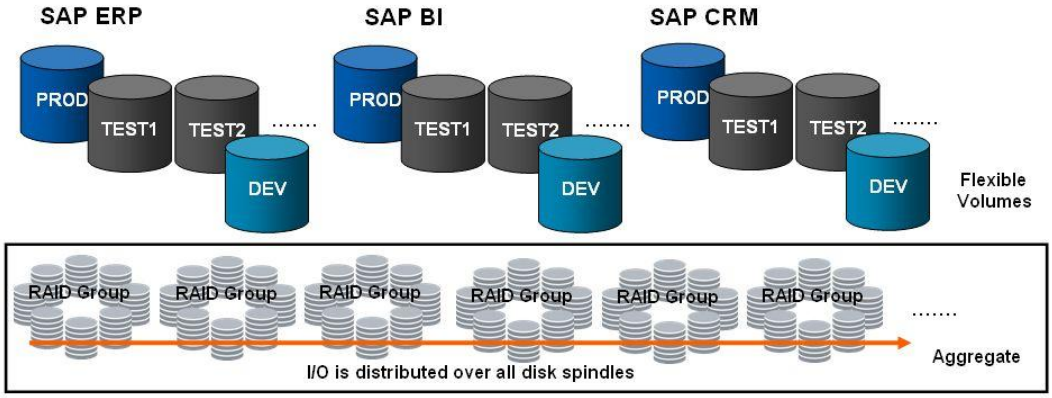

Figure 3) Aggregate layout.

The design of the physical disk layout is very simple because it is created as a physical storage resource pool, and storage resources are assigned on a logical, virtualized level with FlexVol volumes. The FlexVol volumes can be created or deleted instantly and their size can be easily increased or decreased during online operation without any reconfiguration of the underlying physical disk structure or impact to the host storage configuration. This allows optimal utilization of the storage resources.

During normal operations the production systems need the highest performance and therefore the highest number of disk spindles compared to test and development systems. Based on the resource-sharing concepts with disk aggregates, the production systems will benefit from the disk spindles of the test and development systems, which are also necessary because of capacity requirements.

FLEXVOL VOLUME LAYOUT

The best practices described in this document always try to keep the number of volumes as low as possible in order to simplify the storage administration.

Each SAP system is configured with at least two FlexVol volumes:

- One volume for the database data files
- One volume for the online redo log files, the archived log files and the SAP and Oracle binaries

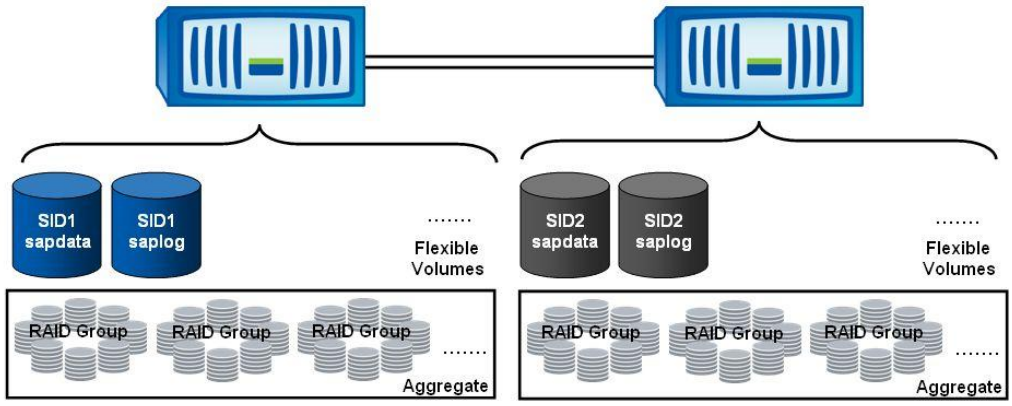

Figure 4) Volume layout.

Storing the database data files and the redo logs in two different FlexVol volumes is important to allow usage of Snapshot copies, SnapRestore, FlexClone, and other Data ONTAP features that work on the volume level.

Configuring more volumes and separating, for example, the binaries, the archive logs, or the temp table space can make sense in specific customer configurations.

In addition to data protection provided by RAID-DP, Oracle data and mirrored log files can be stored physically separated from the archive log files and the online redo logs in two different aggregates:

- One volume for the database data files
- One volume for the online redo log files, the archived log files, and the SAP and Oracle binaries
- One volume for the mirrored redo log files

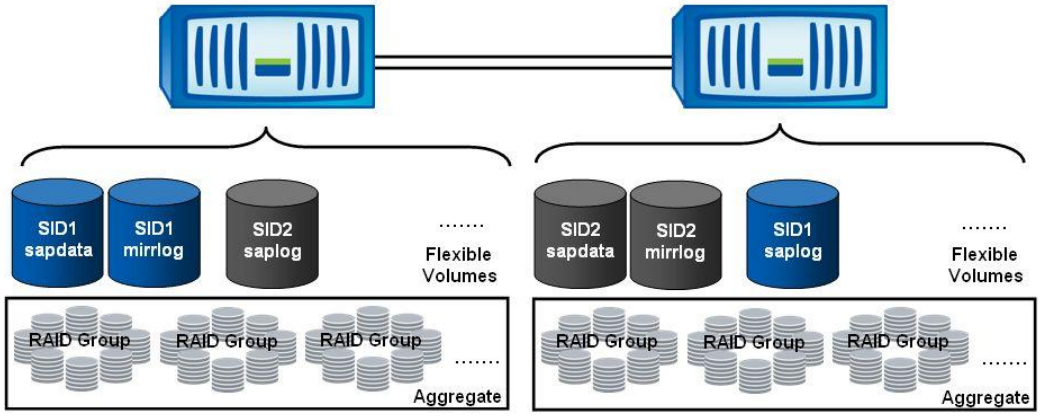

Figure 5) Volume layout with separated mirror logs.

The layout recommendations in the rest of the document cover a three-volume configuration with separated mirror logs. With a two-volume configuration the mirrored logs would be also stored in the “saplog” volume on the same aggregate as the “sapdata” volume.

The following table shows the distribution of the file systems of a single SAP instance to the FlexVol volumes.

Table 1) Storage volume layout

Aggregate 1		Aggregate 2
FlexVol “sapdata”	**FlexVol “mirrlog”**	**FlexVol “saplog”**
/oracle/SID/sapdata1	/oracle/SID/mirrlogA	/oracle/SID/origlogA
/oracle/SID/sapdata2	/oracle/SID/mirrlogB	/oracle/SID/origlogB
/oracle/SID/sapdata3		/oracle/SID/oraarch
/oracle/SID/sapdata4		/oracle/SID
		Oracle binaries
		SAP binaries

FLEXVOL VOLUME LAYOUT WITH METROCLUSTER

The synchronous mirroring with MetroCluster works on the aggregate level. Therefore the layout for MetroCluster is the same as described in the paragraph before.

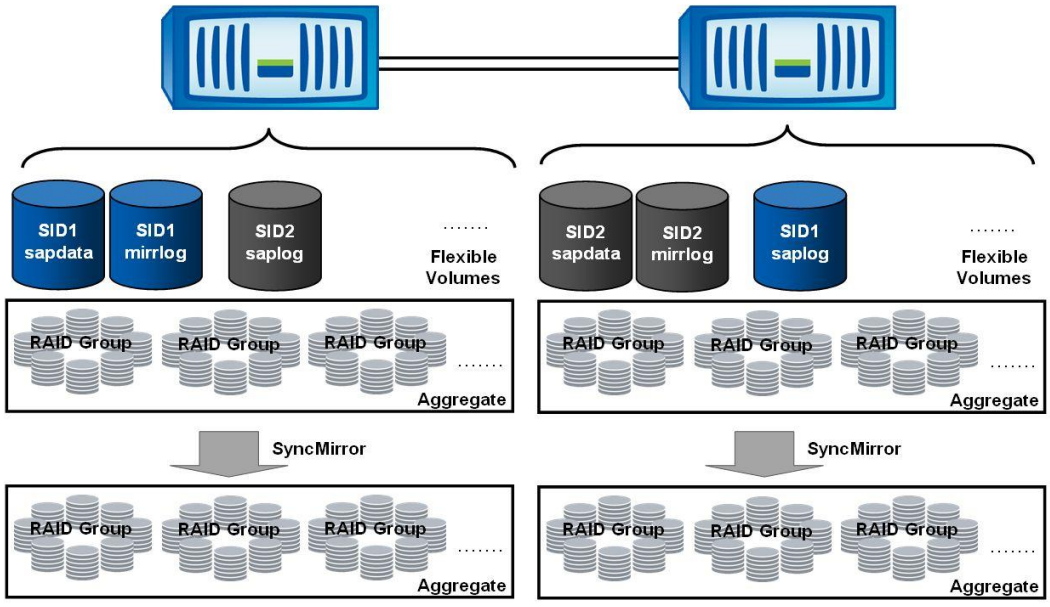

Figure 6) Storage layout with NetApp MetroCluster.

LUN AND VOLUME MANAGER LAYOUT

The following description assumes that a Volume Manager is being used on the UNIX host.

The size of the database is used to determine the number and size of the LUNs required. The goal is to find a balance between the performance advantages of a large number of smaller LUNs and the ease of management that comes with a smaller number of large LUNs.

The following table gives some guidelines for a reasonable number of LUNs, based on the size of the database.

Table 2) Number of LUNs, based on database size.

	Database Size	Number of LUNs for Data Files	Size of LUNs for Data Files	Number of LUNs for Log Files	Number of LUNs for Archive Log Files
Small System	< 200GB	2-4	50GB – 100GB	1	1
Medium System	200GB–1TB	4-8	100GB – 200GB	2-4	2-4
Large System	> 1TB	> 8	100GB – 200GB	>4	>4

The following figure shows the LUN configuration for a small or medium SAP system. From the host point of view, three disk groups need to be configured with the host Volume Manager:

- The “Data Disk Group” contains all LUNs for the database data files.
- The “Log Disk Group” contains all LUNs for the redo logs, the archived logs, and the binaries.
- The “Mirrlog Disk Group” contains all LUNs for the mirrored redo logs.

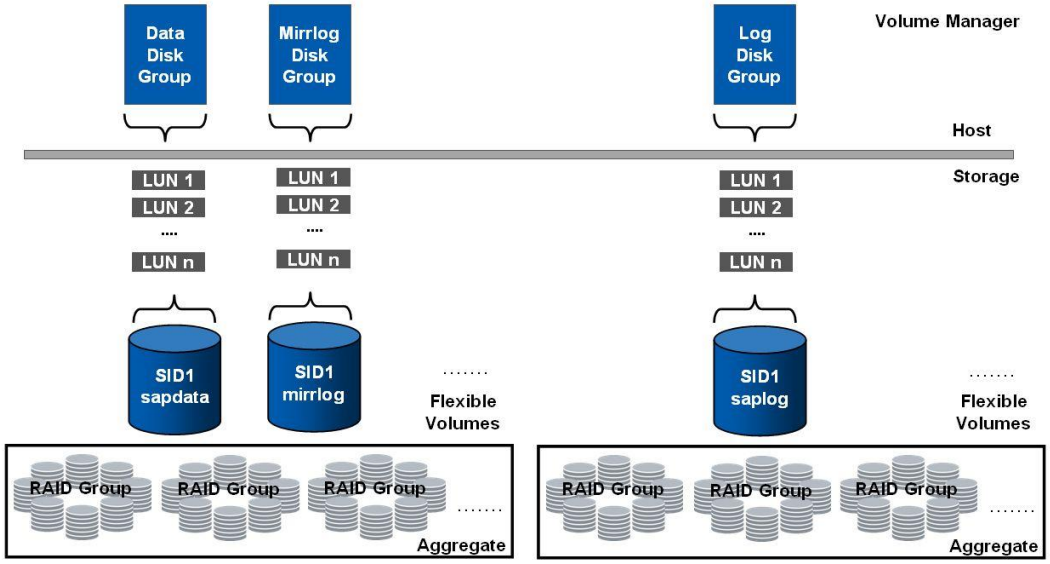

Figure 7) Standard LUN layout.

The following table shows the logical volumes and file systems, that need to be configured within the Volume Manager disk groups at the host.

Table 3) Logical volumes and file systems.

Data Disk Group	Log Disk Group	Mirrlog Disk Group
/oracle/SID/sapdata1	/oracle/SID/origlogA	/oracle/SID/origlogA
/oracle/SID/sapdata2	/oracle/SID/origlogB	/oracle/SID/origlogB
/oracle/SID/sapdata3	/oracle/SID/oraarch	
/oracle/SID/sapdata4	/oracle	
	/usr/sap/trans	
	/usr/sap/SID	
	/sapmnt/SID	

FLEXVOL VOLUME LAYOUT FOR LARGE SAP SYSTEMS WITH HIGH THROUGHPUT REQUIREMENTS

SAP systems with high throughput requirements should be distributed to both storage controllers in the cluster. It can also be beneficial to distribute data from small or medium production systems across both storage controllers to account for future growth. Taking this step during the initial installation will prevent costly downtime in the future as the production system's throughput requirements grow beyond the performance capabilities of a single storage controller.

In addition to data protection provided by RAID-DP, Oracle data and mirrored log files can be stored physically separated from the archive log files and the online redo logs.

Each SAP system uses five FlexVol volumes:

- Two volumes for the database data files distributed to both storage controllers
- One volume for the online redo log files, the SAP, and Oracle binaries
- One volume for the archived log files
- One volume for the mirrored redo log files

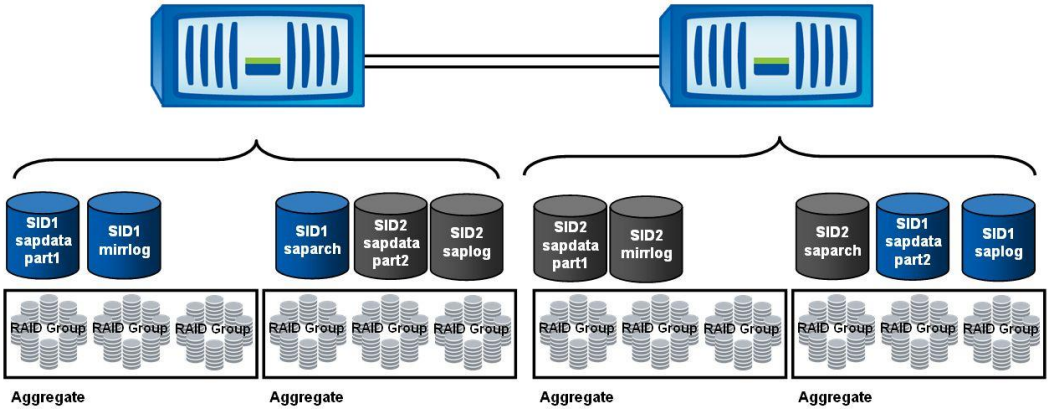

Figure 8) Volume layout for large SAP systems.

The following table shows the distribution of the file systems of a single SAP instance to the FlexVol volumes.

Table 4) FlexVol layout for large SAP systems.

Aggregate 1		Aggregate 2	Aggregate 3	
FlexVol "sapdata_part1"	**FlexVol "mirrlog"**	**FlexVol "saparch"**	**FlexVol "saplog"**	**FlexVol "sapdata_part1"**
/oracle/SID/sapdata1	/oracle/SID/mirrlogA	/oracle/SID/oraarch	/oracle/SID/origlogA	/oracle/SID/sapdata3
/oracle/SID/sapdata2	/oracle/SID/mirrlogB		/oracle/SID/origlogB	/oracle/SID/sapdata4
			/oracle/SID	
			Oracle binaries	
			SAP binaries	

LUN AND VOLUME MANAGER LAYOUT FOR LARGE SAP SYSTEMS WITH HIGH THROUGHPUT REQUIREMENTS

The following figure shows the LUN configuration for a large SAP system. From the host point of view, four disk groups need to be configured with the host Volume Manager:

- The “Data Disk Group” contains all LUNs for the database data files.These LUNs are distributed to both storage controllers to provide load balancing.
- The “Log Disk Group” contains all LUNs for the redo logs and the binaries.
- The “Mirrlog Disk Group” contains all LUNs for the mirrored redo logs.
- The “Arch Disk Group” contains all LUNs for the archive logs.

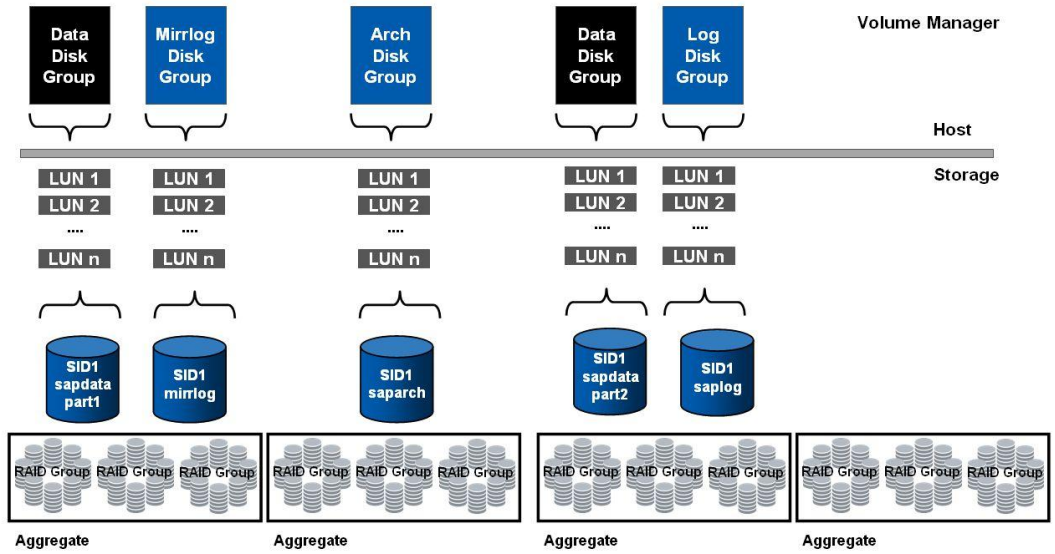

Figure 9) LUN layout for large SAP systems.

The following table shows the logical volumes and file systems that need to be configured within the Volume Manager disk groups at the host.

Each host volume for the database data files must be configured that it uses LUNs from both storage controllers. With this approach, the load distribution to both storage controllers is done for each sapdata file system.

Table 5) Logical volumes and file systems for large SAP systems.

Data Disk Group	Log Disk Group	Mirrlog Disk Group	Arch Disk Group
/oracle/SID/sapdata1	/oracle/SID/origlogA	/oracle/SID/origlogA	/oracle/SID/oraarch
/oracle/SID/sapdata2	/oracle/SID/origlogB	/oracle/SID/origlogB	
/oracle/SID/sapdata3	/oracle		
/oracle/SID/sapdata4	/usr/sap/trans		
	/usr/sap/SID		
	/sapmnt/SID		

2.3 INSTALLATION

This section describes the requirements and the configuration for installing a SAP Business Suite or SAP NetWeaver system with Oracle Database on a UNIX server using the FCP protocol.

GENERAL REQUIREMENTS

NetApp strongly recommends the use of SnapDrive for UNIX, a NetApp host-based software product that simplifies storage management and provisioning in SAP Fibre Channel storage environments. It integrates with NetApp Snapshot and SnapRestore to simplify the process of creating error-free, host-consistent data Snapshot copies. For information on SnapDrive system requirements, see SnapDrive Overview.

For additional recommendations on operating system configuration and tuning, see:

- NetApp Best Practices for Oracle
- Oracle10g Performance: Protocol Comparison on Sun Solaris
- AIX Performance with NFS, iSCSI, and FCP using an Oracle Database on NetApp Storage

NETAPP STORAGE CONTROLLER CONFIGURATION

Snapshot backups for database applications won't be consistent from the database point of view without shutting down the database or putting the Oracle Database in hot backup mode. Therefore automatically scheduled Snapshot copies on the storage level should be turned off on database volumes using the following Data ONTAP command:

```
filer> vol options <volname> nosnap on
```

FLEXVOL AND QTREE LAYOUT

For an SAP Business Suite or NetWeaver system installation on UNIX with Oracle at least three FlexVol volumes are set up. Within these FlexVol volumes one qtree is configured. Qtrees are necessary to allow disk-to-disk backup with SnapVault on a more granular level. All the LUNs are set up within the corresponding qtree.

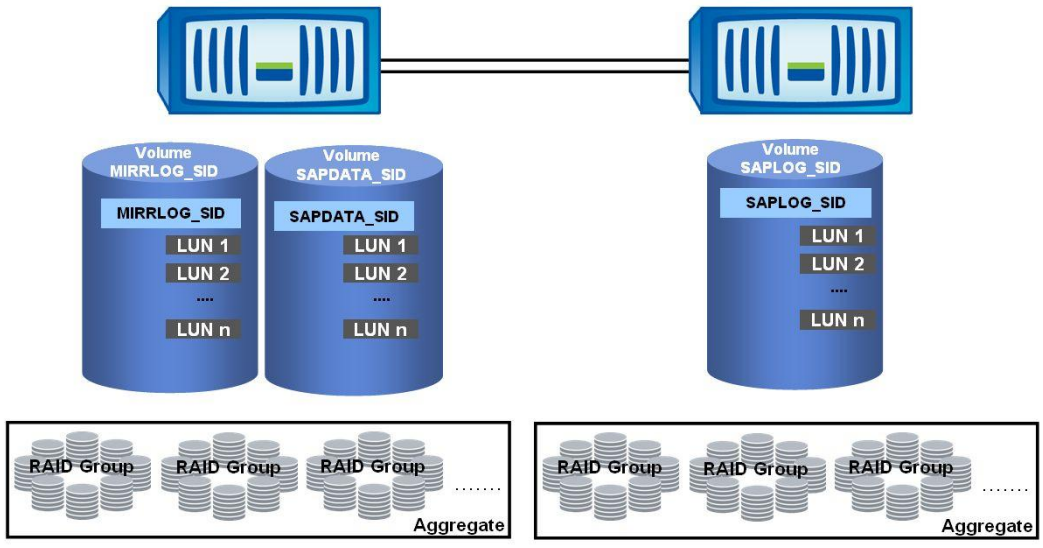

Figure 10) FlexVol and qtree layout.

SAP SYSTEM INSTALLATION

The easiest way to setup the necessary disk groups and LUNs is using SnapDrive for UNIX.

The following SDU command creates four LUNs with 25GB each within the volume sapdata_S10 and creates a Volume Manager Disk Group called S10DataDg.

```
bash-3.00# snapdrive storage create -dg S10DataDg -lun \
sapfiler1:/vol/sapdata_S10/sapdata_S10/lun1 \
sapfiler1:/vol/sapdata_S10/sapdata_S10/lun2 \
sapfiler1:/vol/sapdata_S10/sapdata_S10/lun3 \
sapfiler1:/vol/sapdata_S10/sapdata_S10/lun4 \
-lunsize 25g
        LUN sapfiler1:/vol/sapdata_S10/sapdata_S10/lun1 ... created
        LUN sapfiler1:/vol/sapdata_S10/sapdata_S10/lun2 ... created
        LUN sapfiler1:/vol/sapdata_S10/sapdata_S10/lun3 ... created
        LUN sapfiler1:/vol/sapdata_S10/sapdata_S10/lun4 ... created
        mapping new lun(s) ... done
        discovering new lun(s) ... done
        LUN to device file mappings:
        - sapfiler1:/vol/sapdata_S10/sapdata_S10/lun1 => /dev/vx/dmp/FAS31600_1
        - sapfiler1:/vol/sapdata_S10/sapdata_S10/lun2 => /dev/vx/dmp/FAS31600_2
        - sapfiler1:/vol/sapdata_S10/sapdata_S10/lun3 => /dev/vx/dmp/FAS31600_3
        - sapfiler1:/vol/sapdata_S10/sapdata_S10/lun4 => /dev/vx/dmp/FAS31600_4
        disk group S10DataDg created
bash-3.00#
```

The following SDU command creates the Log Disk Group and the necessary LUNs.

```
bash-3.00# snapdrive storage create -dg S10LogDg -lun \
sapfiler2:/vol/saplog_S10/saplog_S10/lun1 \
sapfiler2:/vol/saplog_S10/saplog_S10/lun2 \
sapfiler2:/vol/saplog_S10/saplog_S10/lun3 \
sapfiler2:/vol/saplog_S10/saplog_S10/lun4 \
-lunsize 5g
        LUN sapfiler2:/vol/saplog_S10/saplog_S10/lun1 ... created
        LUN sapfiler2:/vol/saplog_S10/saplog_S10/lun2 ... created
        LUN sapfiler2:/vol/saplog_S10/saplog_S10/lun3 ... created
        LUN sapfiler2:/vol/saplog_S10/saplog_S10/lun4 ... created
        mapping new lun(s) ... done
        discovering new lun(s) ... done
        LUN to device file mappings:
        - sapfiler2:/vol/saplog_S10/saplog_S10/lun1 => /dev/vx/dmp/FAS31600_5
        - sapfiler2:/vol/saplog_S10/saplog_S10/lun2 => /dev/vx/dmp/FAS31600_6
        - sapfiler2:/vol/saplog_S10/saplog_S10/lun3 => /dev/vx/dmp/FAS31600_8
        - sapfiler2:/vol/saplog_S10/saplog_S10/lun4 => /dev/vx/dmp/FAS31600_7
        disk group S10LogDg created
bash-3.00#
```

The following SDU command creates the Mirrlog Disk Group and the necessary LUNs.

```
bash-3.00# snapdrive storage create -dg S10MirrDg -lun \
sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun1 \
sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun2 \
-lunsize 500m
        LUN sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun1 ... created
        LUN sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun2 ... created
        mapping new lun(s) ... done
        discovering new lun(s) ... done
        LUN to device file mappings:
        - sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun1 => /dev/vx/dmp/FAS31600_5
        - sapfiler1:/vol/mirrlog_S10/mirrlog_S10/lun2 => /dev/vx/dmp/FAS31600_6
        disk group S10MirrDg created
bash-3.00#
```

The configuration of the Volume Manager host volumes and file systems within the configured disk groups is done using the host operating system Volume Manager commands and tools.

The SAP installation is accomplished as described in the corresponding SAP Installation Guide.

2.4 SIZING

This section gives an overview of the storage sizing for an SAP environment using NetApp storage. The goal is to provide a basic understanding of what kind of information is important in performing a storage sizing and how these requirements influence the storage landscape.

NetApp can provide storage sizings to SAP customers, based on a sizing questionnaire filled in by the customer.

Storage sizing for an SAP landscape is based on several conditions that are defined by customer requirements. All of these requirements together define the needed storage infrastructure:

- Throughput requirements
- Capacity requirements
- Backup and recovery requirements (mean time to recover, backup window, retention policy)
- Cloning requirements (FlexClone copies or full copies)
- Disaster recovery requirements (synchronous or asynchronous mirroring)
- High-availability requirements (storage system clustering)

For existing SAP systems, the throughput load is measured using database or operating system tools. Independent of which tools are used, it is important that the measurement is done during peak loads on the SAP system. When database tools are used for the measurement, a suitable time frame such as one hour must be chosen, because these tools calculate an average value, and the throughput sizing must be based on peak values.

For new SAP systems, where a throughput measurement is not possible, the SAPS *(*SAP Application Performance Standard*)* values for the systems, which are provided by the SAP Quick Sizer, can be used to estimate the throughput requirements. The storage sizing is much more accurate when real throughput values are measured. SAPS-based sizing should only be done if no other data is available.

Based on the throughput requirements, the type and number of disk spindles and storage controllers are determined.

In order to determine the needed capacity, the following information must be available:

- Size of each database
- Growth rate
- Daily change rate
- Number and retention policy of Snapshot copies
- Number and durability of FlexClone volumes
- Synchronous or asynchronous mirroring

Based on the capacity requirements, the type and number of disks and the storage controller supporting the capacity are determined.

The results of the throughput sizing and the capacity sizing are compared in a final step to define the right storage system supporting both the throughput and capacity requirements.

2.5 STORAGE MIGRATION

The decision about which migration approach fits best in a specific environment depends heavily on the acceptable downtime of the business application. Furthermore, the downtime depends on the amount of data that needs to be migrated. In general, there are different approaches to storage migration of the SAP data:

- Migration on the operating system level
- Migration on the volume manager level
- Migration on the database level

MIGRATION ON THE OPERATING SYSTEM LEVEL

In addition to the existing storage system, the NetApp storage system is connected to the database server. The NetApp storage system is configured and the storage objects are connected to the server. Before the data migration is started, the database and the SAP system must be shut down. The data is then copied using the server from the old storage system to the NetApp system using operating system commands. When all data is copied, the old storage system is disconnected from the database server. If the file system structure remains the same, the database can be started immediately. If there is a change in the file system structure, the new structure must be configured within Oracle by creating a new control file.

A migration on the operating system level can be done for an FCP-to-FCP, FCP-to-NFS or FCP-to-iSCSI migration. The disadvantage of this approach is that the SAP system won't be available while the database files are copied. Depending on the database size, the downtime could be several hours.

MIGRATION ON THE VOLUME MANAGER LEVEL

The migration is done using a host-based mirror provided by the host volume manager software. The NetApp storage system must be configured and attached to the host. All data that needs to be migrated must be mirrored to the NetApp storage by adding an additional plex to the disk/volume group. The synchronizing of the new mirror should be scheduled during a low-activity period, because there is a high load on the server during the synchronization. When the synchronization is finished, all new data is synchronously mirrored to the NetApp storage system. At this point in time the volume manager must be reconfigured so that the plex stored at the old storage system is not used any more.

A migration on the volume manager level can be done for any block protocol migration. The data migration is done during online operation, but has a performance effect on the database host.

MIGRATION ON THE DATABASE LEVEL

The migration on database level is done using an Oracle standby database. A database backup of the production system is restored to a standby database server connected to the NetApp storage system. In addition, the archive logs are continuously copied to and applied at the standby database server. Before the final migration is started, the SAP database and the SAP system must be shut down and the remaining archive logs need to be applied to the standby database. The NetApp storage is then connected to the original production database server and the storage objects are connected to the server. The old storage system is disconnected from the database server.

A migration on the database level can be done for an FCP-to-FCP, FCP-to-NFS or FCP-to-iSCSI migration. This approach reduces downtime during the migration but requires an additional server during the migration process.

The following table summarizes the migration processes.

Table 6) Summary of different migration processes.

	Supported Protocols	Downtime	Additional Hardware
Migration on operating system level	Any to any	**Long** During copy process	None
Migration on volume manager level	All block protocols	**No downtime**	None
Migration on database level	Any to any	**Short** During switch of storage system and final forward recovery	Standby database server

3 SAP LIFECYCLE MANAGEMENT

3.1 SAP SYSTEM COPIES

CAPACITY REQUIREMENTS

When creating SAP system copies with most storage architectures, space must be allocated to accommodate the entire size of the source database. This can drastically increase the amount of storage required to support a single production SAP instance.

During a typical project a 1TB SAP production system will be copied to a quality assurance (QA) system, a test system, and a training system. With conventional storage architectures, this requires an additional 3TB of storage. Furthermore, it requires a significant amount of time to first back up the source system and then to restore the data to the three target systems.

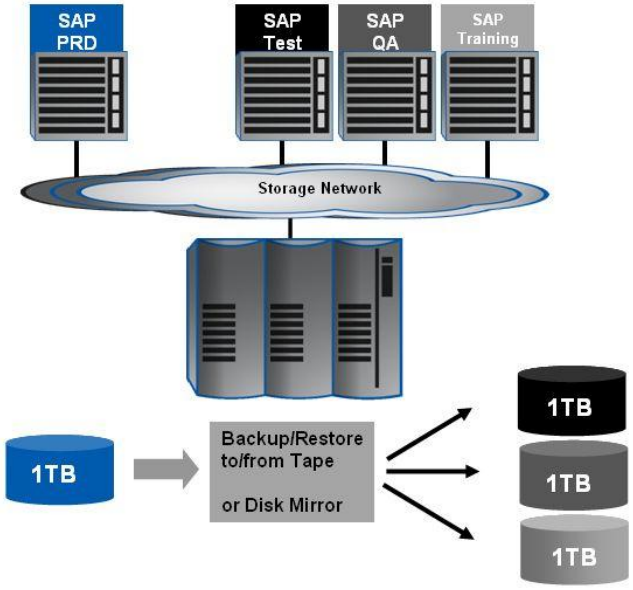

Figure 11) SAP system copy: standard approach.

In contrast, when using NetApp FlexClone technology to create SAP system copies, only a fraction of the storage space is required. NetApp FlexClone technology uses Snapshot copies, which are created in a few seconds without interrupting the operation on the source system, to perform SAP systems copies. Because the data is not copied but referenced in place, the amount of storage required is limited to only data that is changed at the source and the target system and therefore significantly decreases the disk space needed for SAP system copies.

As a result, the capacity requirements for a system copy in a NetApp storage environment depends on the refresh cycle of the target systems. As longer test systems are kept, more block changes will happen from the source and the target system. Storage requirements also depend on the number of copies that are made from the same source. Of course more copies of the same source system will result in higher storage savings.

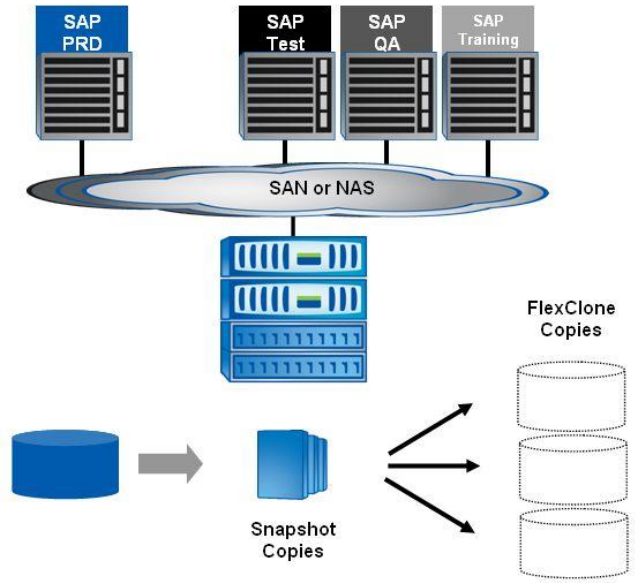

Figure 12) SAP system copy: NetApp approach at primary storage.

On the source system a database-consistent Snapshot copy of the data files is created. This is done during online operation and has no performance effect on the source system. Therefore this step can be carried out at any time.

The FlexClone copy can be created at the same storage system or at a secondary storage system.

The secondary storage system could be already in place and used as a disk-to-disk backup device or a disaster recovery solution. The backup or disaster recovery replication images can be accessed for reading and writing using FlexClone technology. Existing backup or disaster recovery images will be utilized for test environments, turning expenses into assets. As a side effect the backup and recovery or disaster recovery solution is tested without any additional effort and without any interruption.

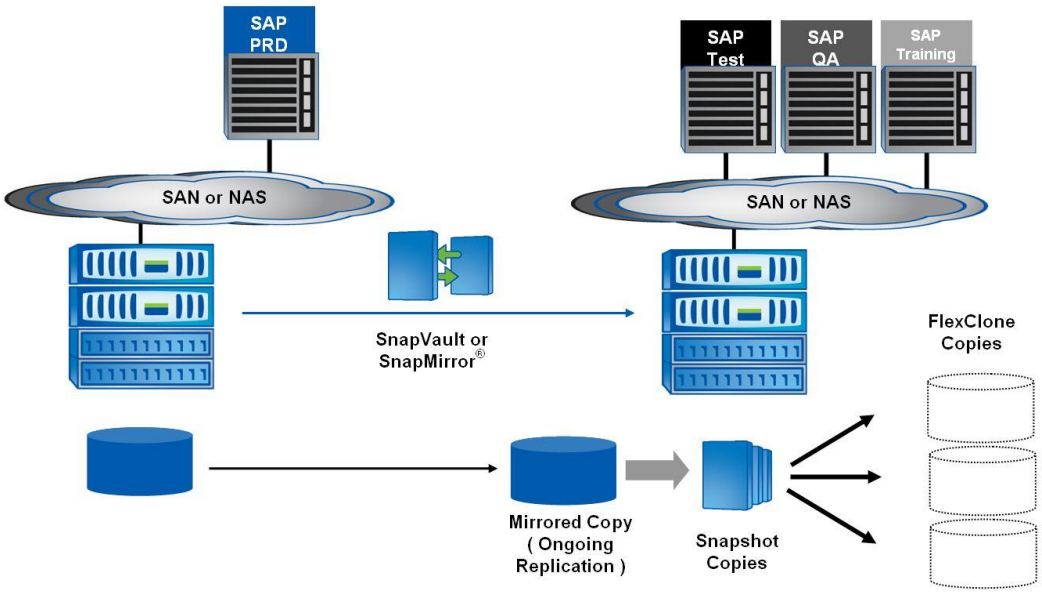

Figure 13) SAP system copy: NetApp approach at secondary storage.

TIME REQUIREMENTS

The time required to create an SAP system copy can be subdivided into three parts:

- Time to restore the backup to the target system.
- Time to perform OS and database-specific postprocessing.
- Time to perform SAP application postprocessing.
  The SAP postprocessing depends on the customer SAP environment. Some customers can finish the postprocessing in a few hours, while other customers need several days to accomplish this task.

In a conventional system copy process, the data is backed up to tape and then restored, which takes a great deal of time. If an online backup is used, there is no downtime for the source system; however, there might be a performance effect on the source system during the backup. Because of the large number of logs that need to be applied, the time required to recover the database and make it consistent is greatly increased, possibly adding hours to the system copy process. If an offline backup is used, the source system is shut down, resulting in a loss of productivity.

The following figures show an example describing the difference between the amount of time spent creating an SAP system copy with NetApp storage versus the time spent using a conventional approach.

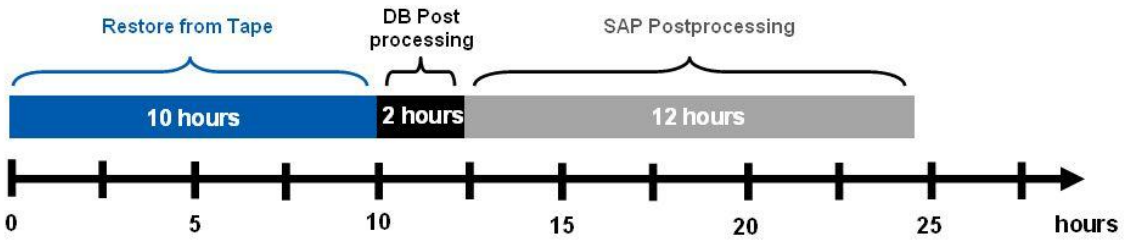

Figure 14) SAP system copy: Standard approach.

All steps up to the point when the SAP system can be started on the target host can be accomplished in a few minutes using the NetApp solution compared to several hours with the standard approach. With both approaches the SAP postprocessing needs to be done as an additional step.

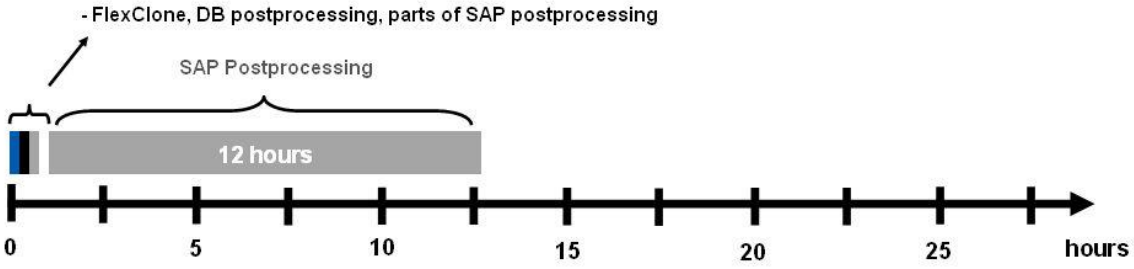

Figure 15) SAP system copy: NetApp approach.

A key requirement to successfully manage an SAP environment is the ability to create copies of production data for use in testing, quality assurance, or training. NetApp Snapshot and FlexClone technologies allow a fast and space-efficient creation of SAP systems.

3.2 SAP UPGRADE

A typical example of a complex customer project within an SAP landscape is an upgrade scenario. Similar to upgrade projects, the described issues also occur during unicode conversions, when applying SAP enhancement packages and during custom developments. The procedures to overcome those issues are similar.

Running a new version of the SAP applications is often a prerequisite for reengineering or introducing new business processes. Over time, SAP versions will also reach the end of their maintenance period. Therefore SAP customers need to go through an SAP upgrade project at regular intervals.

Customers face several challenges in an SAP upgrade project:

- **Costs:** The SAP upgrade project consumes large amounts of employee time. To cut the cost, the project time needs to be reduced.
- **Delayed innovation:** Business processes are affected during the upgrade project time period, because all development needs to be stopped, and SAP support packages can't be imported. Therefore it is very important to minimize the overall time for the upgrade project.
- **Risk:** An upgrade causes change, which introduces risk that the business processes might not work as expected after the upgrade. The risk needs to be reduced by increasing the frequency and quality of testing and training.
- **Production system downtime:** During the upgrade of the production system, the system will not be available. Production system downtime has to be minimized.

In complex environments with large databases, a normal two-day weekend might not be sufficient for upgrading the production SAP system. Every hour that can be saved while running the upgrade is important. Database backups consume a large portion of total time. Optimizing backup and restore functionality is therefore critical. During an SAP upgrade project, SAP basis administrators need to create several system copies to test the upgrade process with current data from the development or production SAP system. The creation of an SAP system copy usually takes several days and might negatively affect the production environment. In addition, many manual steps must be performed, consuming valuable IT staff time.

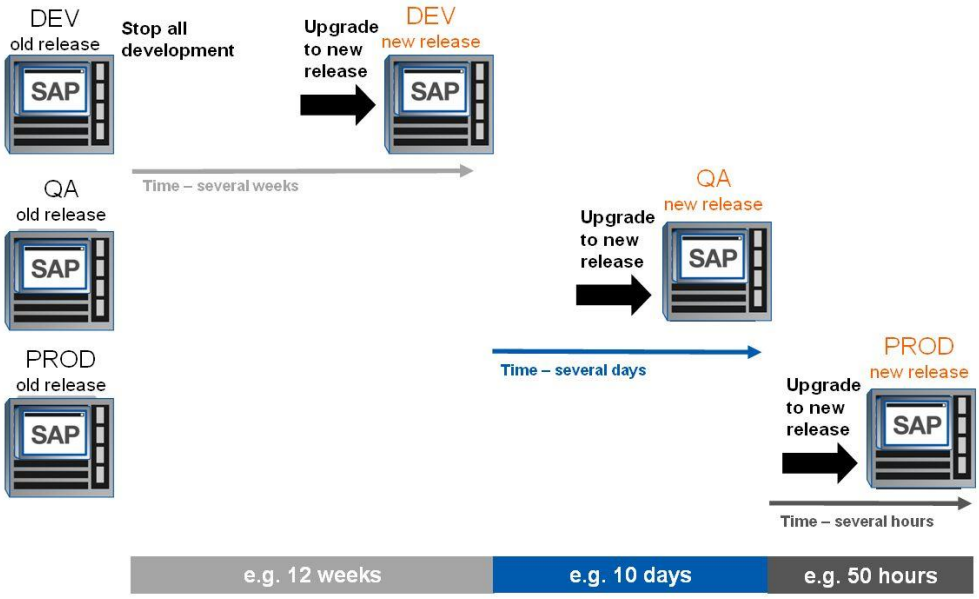

Figure 16) SAP upgrade overview.

UPGRADING THE DEVELOPMENT SYSTEM

The upgrade of the development system is usually carried out on a copy of the current development system running on separate hardware. During the upgrade process, the functionality of the upgrade is tested in the specific customer environment. In almost all cases, the upgrade of the development system is carried out more than once in order to define the necessary actions for all upgrade phases.

The setup of the separate SAP system is done based on a system copy of the original development system. This system copy can be provided using the NetApp system copy solution. Using this solution will significantly reduce the time and resources needed for the system copy. Reducing the time is critical because in most cases the copy is created several times.

During the upgrade process and during the modification adjustment, Snapshot backups are very helpful, allowing the system to be reset to any Snapshot copy and to restart the upgrade phase.

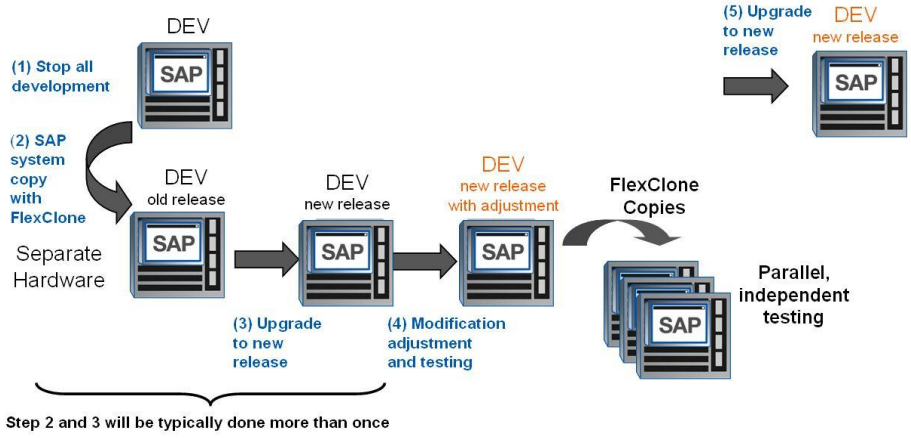

Figure 17) SAP upgrade: development system.

UPGRADING THE QUALITY ASSURANCE SYSTEM

The quality assurance system is upgraded using a fresh system copy of the production SAP system. One important result of this upgrade is testing the upgrade with production data. The NetApp SAP system copy solution allows efficient refreshing of the quality assurance system. Reducing the time necessary to create this copy is also critical when upgrading the quality assurance system because the copy is usually made more than once to support multiple tests of the upgrade process and also multiple test cycles. Snapshot backups are helpful during the upgrade process and before the modification adjustments are imported. These Snapshot copies allow restore of the system to any specific Snapshot copy, allowing restarting an upgrade phase or restarting the import.

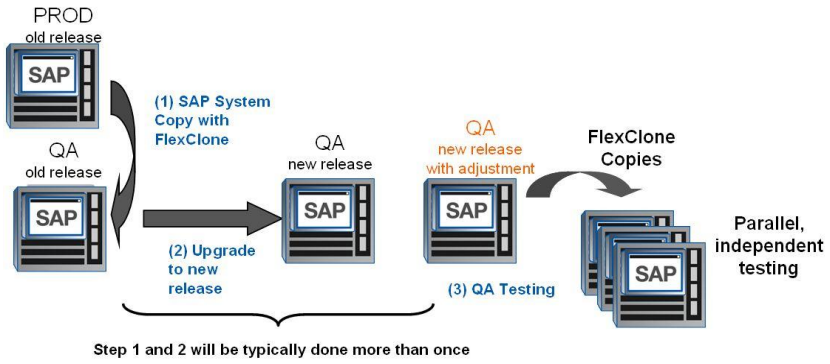

Figure 18) SAP upgrade: quality assurance system.

UPGRADING THE PRODUCTION SYSTEM

Scheduling is extremely important when upgrading the production system, because the system is not available at various stages during the upgrade. The schedule has to provide time to restore the system to its former release status. Depending on the size of the database and the time and effort required for the functional test and for importing the transports for the modification adjustment, a 48-hour weekend might not be enough time to complete the upgrade.

The production system upgrade includes at least three backups of the database. The first backup must be done immediately before the upgrade is started. After the upgrade is finished, a second backup is required before the modification adjustments are imported. After importing the adjustments and finishing the functionality tests, a third backup is required. If functionality testing fails, the system must be restored to the previous release level.

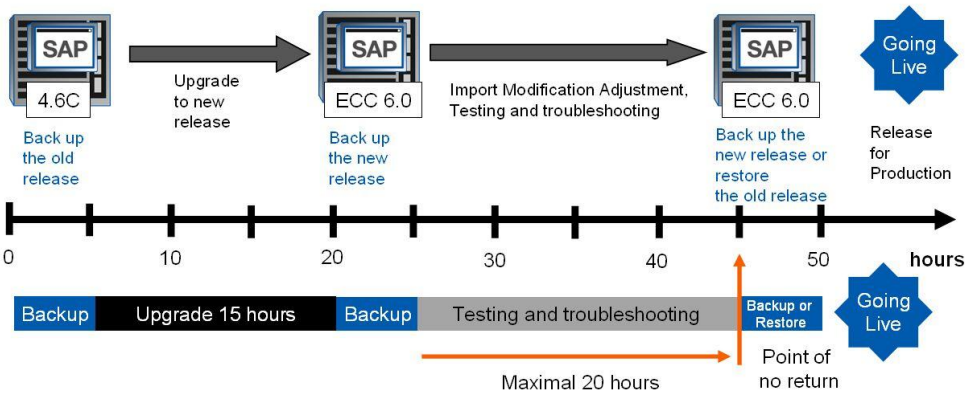

Figure 19) SAP upgrade: production system.

Using Snapshot copies as a backup method and SnapRestore for restoring the system to its former release status assures a higher level of flexibility with regard to scheduling. Normal tape backups take several hours, which must be considered when planning the upgrade schedule. This time is reduced to several minutes when using Snapshot and SnapRestore features. Therefore there will be more time for testing and troubleshooting since the “point of no return” moves closer to “going live.”

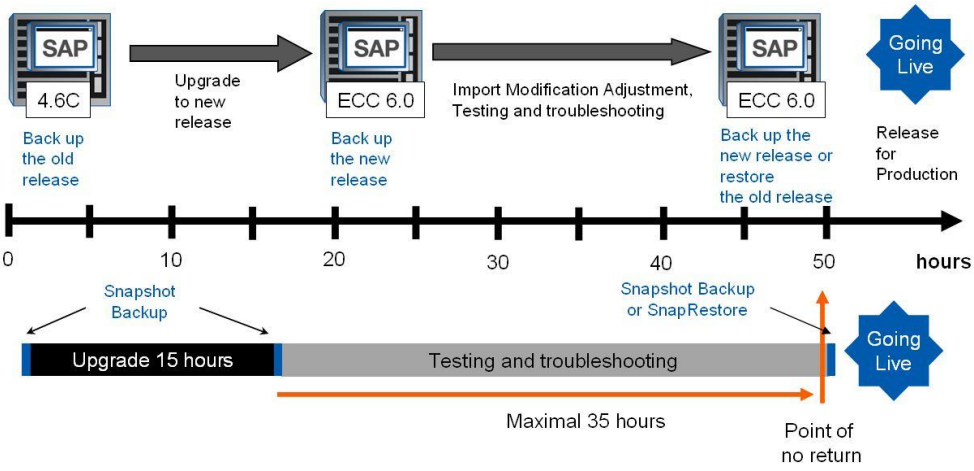

Figure 20) SAP upgrade: production system with NetApp.

Reducing the time needed for backup and restore will allow minimizing the upgrade downtime of the production SAP system. One option is to go live earlier with the upgraded release.

The other option is to use more time for testing the upgraded release before going live. More testing reduces the risk and provides more time to fix any issues that arise.

3.3 NETAPP AND SAP TDMS INTEGRATION

SAP Test Data Migration Server (TDMS) transfers data from a source system into a test or development system. Within TDMS transfer criteria can be defined. The criteria could be, for example, “only data from the last year.” TDMS can significantly reduce the amount of data that is duplicated in SAP development and test systems.

A requirement of the TDMS solution is that the production source system must be inactive while the TDMS processes are reading the data. If the source system is active, data inconsistency at the target system is possible. Therefore the production source system must be locked for dialog users, and no batch jobs are allowed to run. This requires production downtime.

The NetApp integration into TDMS allows data integrity without downtime for production. A clone of the production system is created on the NetApp storage by using NetApp FlexClone functionality. The cloned system is used as the source for the TDMS processing. Because the FlexClone copy is created based on a Snapshot copy of the source production system, almost no additional storage is needed. Also, the clone is deleted after the TDMS process is complete, eliminating the need to maintain a complete copy of the production database.

TDMS allows data selection from a certain date to the current date. For some customers, this could cause a regulatory conflict. Data from the current financial period could be imported into the development and test systems, giving employees access to sensitive information. NetApp technology allows customers to choose a Snapshot copy of the production system prior to the current reporting period to use as a source for the data extraction so that sensitive data is not part of the extraction. If sensitive data is not an issue, FlexClone copies can be created on a regular basis to make sure that the latest data from the production system is migrated to the test system by using the TDMS process.

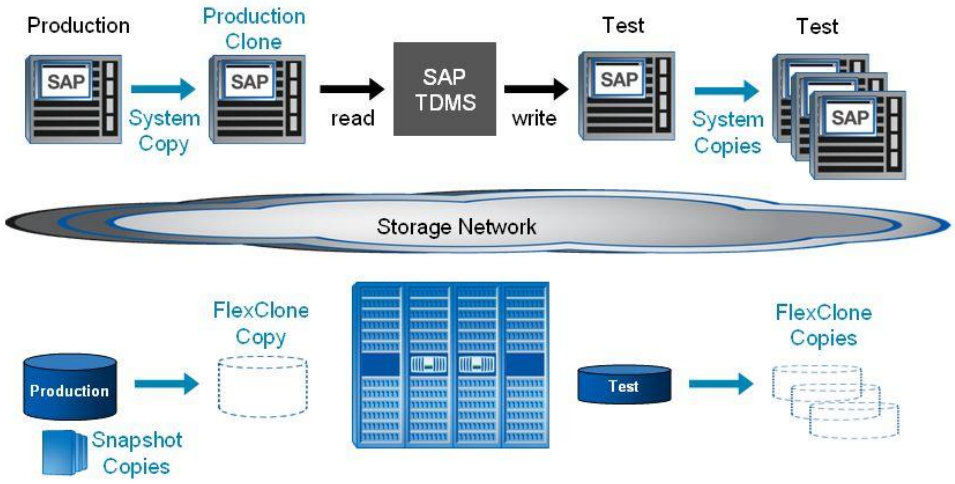

Figure 21) NetApp storage in an SAP TDMS environment.

The above figure illustrates the integrated process of using NetApp storage in an SAP TDMS environment. Several steps are performed to transfer the data from the source system to the target system:

1. The NetApp plug-in is called by TDMS to request a copy of the production system.
2. The NetApp plug-in creates an SAP system copy based on Snapshot and FlexClone technology and starts the production clone so that TDMS can access the cloned system.
3. TDMS reads data from the production clone and writes to the target system.
4. After the data extraction is finished, TDMS calls the NetApp plug-in to delete the production clone.
5. Additional target systems can be created as FlexClone copies.

A key requirement for successfully managing an SAP environment is the ability to have current data for test, development, training, and sandbox systems. However, because of the explosive growth of data in SAP production systems, it is not always feasible to create complete copies for these purposes. SAP offers TDMS to overcome this obstacle. When combined with NetApp Snapshot and FlexClone technologies to allow the fast and efficient creation

of SAP systems, TDMS extractions can happen without limiting production availability. Once the data has been extracted, multiple copies can be made to allow parallel testing and development activities without changing the data in the original target.

4 BUSINESS CONTINUANCE

4.1 BACKUP AND RECOVERY

Corporations today require their SAP applications to be available 24 hours a day, seven days a week. Consistent levels of performance are expected, regardless of increasing data volumes and routine maintenance tasks such as system backups. Performing backups of SAP databases is a critical task, and can have a significant performance effect on the production SAP system. Because backup windows are shrinking and the amount of data that needs to be backed up is increasing, it is difficult to define a point in time when backups can be performed with minimal effect on the business process. The time needed to restore and recover SAP systems is of particular concern because the downtime of SAP production and nonproduction systems must be minimized.

The following summarizes SAP backup and recovery challenges:

- **Performance effect on production SAP systems:** Backups typically have a significant performance impact on the production SAP system because there is a heavy load on the database server, the storage system, and the storage network during backups.
- **Shrinking backup windows:** Conventional backups have a significant performance effect on the production SAP system; backups can be made only during times with little dialog or batch activities taking place on the SAP system. It becomes more and more difficult to define a backup window when the SAP system is used 24x7.
- **Rapid data growth:** Rapid data growth together with shrinking backup windows results in ongoing investments in the backup infrastructure: more tape drives, new tape drive technology, faster storage networks. Growing databases also result in more tape media or disk space for backups. Incremental backups can address these issues, but result in a very slow restore process, which is usually not acceptable.
- **Increasing cost of downtime:** Unplanned downtime of an SAP system always causes a financial effect on the business. A significant part of the unplanned downtime is the time that is needed to restore and recover the SAP system in the case of a failure. The backup and recovery architecture must be designed based on an acceptable recovery time objective (RTO).
- **Backup and recovery time included in SAP upgrade projects:** The project plan for an SAP upgrade always includes at least three backups of the SAP database. The time needed to perform these backups' cuts down the total available time for the upgrade process. The go/no-go decision is based on the amount of time required to restore and recover the database from the backup that was created previously. The option to restore very quickly allows more time to solve problems with the upgrade rather than to restore the backup.

NetApp Snapshot technology can create an online or offline database backup in minutes. The time needed to create a Snapshot copy is independent of the size of the database, because a Snapshot copy does not move any data blocks. The use of Snapshot technology has no performance effect on the production SAP system because the NetApp Snapshot implementation does not copy data blocks when the Snapshot copy is created or when data in the active file system is changed. Therefore, the creation of Snapshot copies can be scheduled without having to consider peak dialog or batch activity periods. SAP and NetApp customers typically schedule several online Snapshot backups during the day, for instance, every four hours. These Snapshot backups are typically kept for three to five days on the primary storage system.

Snapshot copies also provide key advantages for the restore and recovery operation. The NetApp SnapRestore functionality allows restore of the entire database or parts of the database to the point in time when any available Snapshot copy was created. This restore process is done in a few minutes, independent of the size of the database. Because several online Snapshot backups were created during the day, the time needed for the recovery process is also dramatically reduced. Because a restore can be done using a Snapshot copy that is at most eight hours old, fewer transaction logs need to be applied. The mean time to recover, which is the time needed for restore and recovery, is therefore reduced to several minutes compared to several hours with conventional tape backups.

Snapshot backups are stored on the same disk system as the active online data. Therefore, NetApp recommends using Snapshot backups as a supplement, not a replacement for backups to a secondary location such as disk or tape. Although backups to a secondary location are still necessary, there is only a slight probability that these backups will be needed for restore and recovery. Most restore and recovery actions are handled by using

SnapRestore on the primary storage system. Restores from a secondary location are only necessary if the primary storage system holding the Snapshot copies is damaged or if it is necessary to restore a backup that is no longer available from a Snapshot copy, for instance, a two-week-old backup.

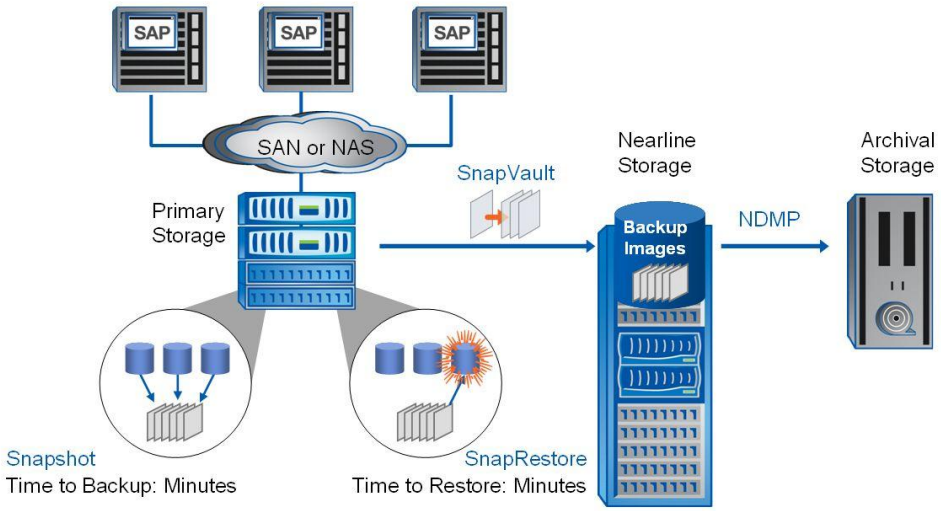

Figure 22) General backup solution overview.

A backup and recovery solution using a NetApp storage system always consists of two parts:

- Backup and restore using Snapshot and SnapRestore
- Backup and restore to/from a secondary location

A backup to a secondary location is always based on Snapshot copies created on the primary storage. Therefore, the data is read directly from the primary storage system without generating load on the SAP database server. There are two options to back up the data to a second location:

- **Disk-to-disk backup using a NetApp near-line or primary storage system and SnapVault software**: The primary storage communicates directly with the secondary storage and sends the backup data to the destination. The NetApp SnapVault functionality offers significant advantages compared to tape backups. After an initial data transfer, in which all the data has to be transferred from the source to the destination, all following backups copy only the changed blocks to the secondary storage. The typical block change rate for a SAP system is around 2% per day. Therefore the load on the primary storage system and the time needed for a full backup are significantly reduced. Because SnapVault stores only the changed blocks at the destination, a full database backup requires significantly less disk space.
  Backing up data to tape as a long-term backup might still be required. This could be, for example, a monthly backup, that is kept for a year. In this case the tape infrastructure can be directly connected to the secondary storage, and the data will be written to tape using NDMP.
- **Backup to tape using third-party backup software such as NDMP backup (server-less backup):** The tape is connected directly to the primary storage system. The data is written to tape using NDMP.

The following figure compares the different backup approaches with regard to the performance effect of a backup and the time in which the database must be in hot backup mode or offline.

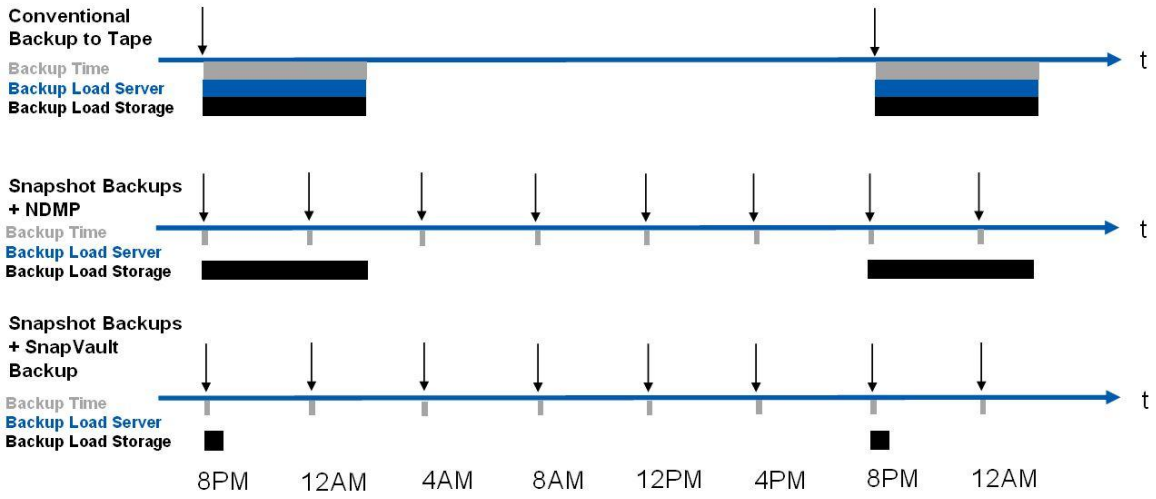

Figure 23) Comparison of time required for different backup methods.

SNAPSHOT BACKUPS TOGETHER WITH NDMP BACKUPS

Snapshot backups do not generate any load on the database server or the primary storage system. A full database backup based on Snapshot consumes disk space only for changed blocks. Snapshot backups are typically scheduled more often, for example, every four hours. A higher backup frequency allows a more flexible restore process and reduces the number of logs that must be applied during forward recovery. In addition, a full NDMP backup to tape is scheduled once a day. This backup still creates a heavy load on the primary storage system and takes the same amount of time as conventional tape backup.

SNAPSHOT BACKUPS TOGETHER WITH DISK-TO-DISK BACKUP AND SNAPVAULT

Snapshot backups are used here in the same way as described in the previous section.

Because SnapVault runs at the storage level, there is no load on the database server. SnapVault transfers only the changed blocks with each backup. Therefore, the load on the primary storage is significantly reduced. For the same reason, the time needed to perform a full database backup is short. In addition, each full backup stores only the changed blocks at the destination. Therefore, the amount of disk space that is needed for a full backup is very small compared to full tape backups.

The following figure compares the time required to perform restore and recovery.

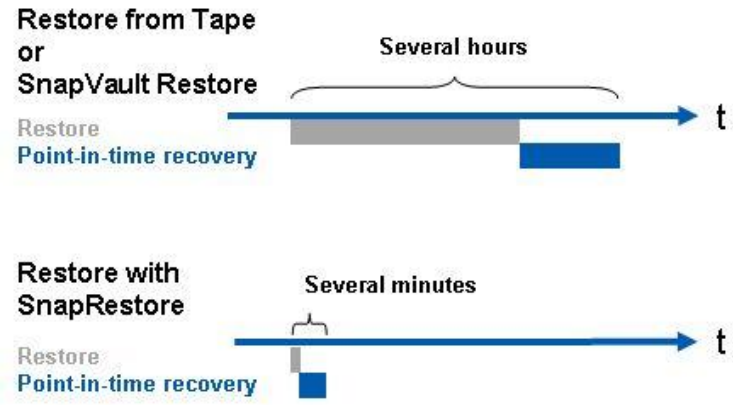

Figure 24) Comparison of time needed for restore and recovery using NetApp solutions.

RESTORE FROM TAPE OR SNAPVAULT RESTORE

The time needed to restore the database from tape or disk depends on the size of the database and the tape or disk infrastructure that is used. In either case, several hours are required for performing a restore. Because the backup frequency is typically one backup a day, a certain number of transaction logs need to be applied after the restore is finished.

RESTORE WITH SNAPRESTORE

The database restore time with SnapRestore is independent of the database size. A SnapRestore process is always finished in a few minutes. Snapshot backups are created with a higher frequency, such as every four hours, so the forward recovery is much faster, because fewer transaction logs need to be applied.

If Snapshot backups are used in combination with tape or SnapVault backups, most restore cases are handled with SnapRestore. A restore from tape or disk is only necessary if a Snapshot copy is no longer available.

The combination of Snapshot and SnapRestore with a disk-to-disk backup concept based on SnapVault offers significant improvement over conventional tape backups:

- Negligible effect of backups on the production SAP system
- Dramatically reduced RTO
- Minimum disk space needed for database backups on the primary and the secondary storage systems

ACCELERATING TEST AND TRAINING CYCLES

Backups based on Snapshot copies offer the possibility to create several consistent images/golden images of an SAP system, which can be used for test runs or user trainings.

If a test run has to be repeated, SnapRestore reverts the system back to a consistent image within minutes to repeat the test or to rerun the test with other settings.

This also applies for user trainings. After the training has been finished, the system will be reverted back to a golden image within minutes using SnapRestore and the next training can be started immediately.

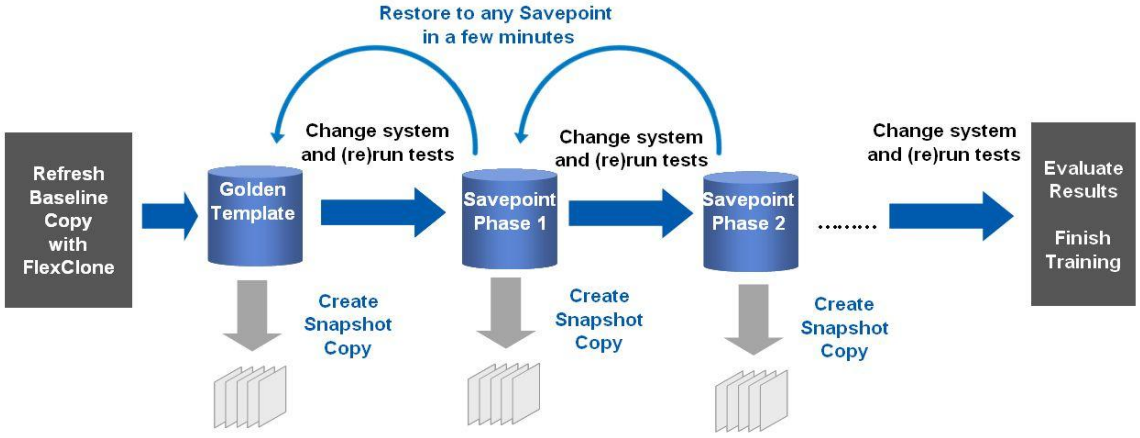

.

Figure 25) SAP testing cycle.

4.2 SAP REPAIR SYSTEM

More and more companies are facing the challenge of addressing logical errors in a more complex SAP environment, where several SAP systems exchange data with each other.

A logical error can be addressed by restoring the system using the last backup and doing a forward recovery up to the point before the logical error occurred. This approach has several disadvantages:

- Downtime for the analysis when the logical error occurred and for the restore and recovery process
- Data loss, because the system got recovered to a point in time in the past
- Inconsistency between the system that got restored and recovered to point in time in the past and the other systems that exchange data with that system

Therefore SAP customers are looking for a more efficient and flexible solution to address logical errors. The NetApp Snapshot and FlexClone technology helps to provide a solution that allows recovery from logical errors without the need to restore and recover the affected system.

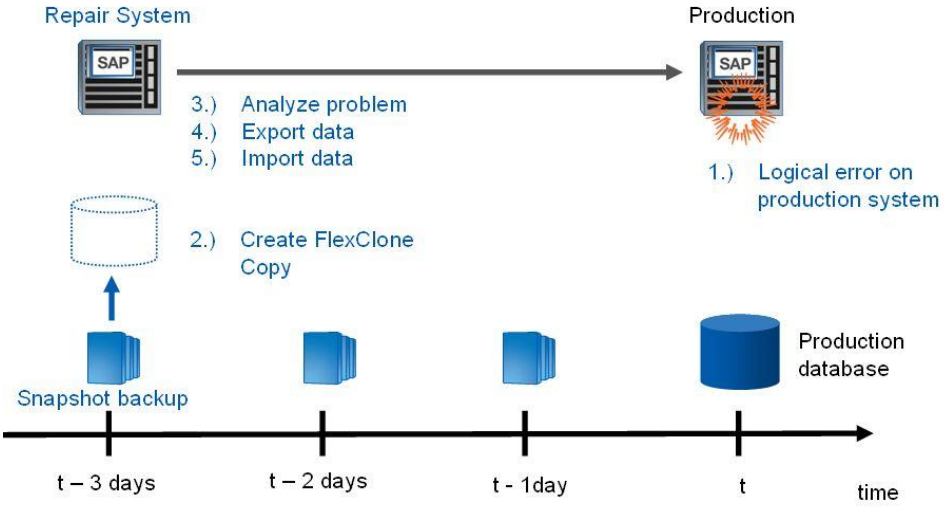

Figure 26) SAP repair system.

The figure above shows the general process of creating and using the repair system.

1. A logical error is discovered on the production system. Dependent on the kind of logical error the decision can be made to shutdown the production system, or to keep it online and only parts of the business processes are affected.
2. Several Snapshot backups of the production system are available and any of these backups in the past can be chosen to create a SAP system copy of the production. The SAP system copy is created using a FlexClone copy of the Snapshot copy.
3. The repair system is used to analyze the problem.
4. The appropriate data is exported from the repair system.
5. The data is imported to the production system.

In the described example there is less or no effect on the production system, no data loss, and no inconsistency within the SAP landscape.

The described scenario is quite simple, and it is obvious that not all logical errors can be solved in such an easy way. However the repair system approach will also help in more complex scenarios, because there is more flexibility, and there are more options to analyze and to recover from the logical error.

4.3 HIGH AVAILABILITY

Production SAP systems are business-critical applications that require 24x7 availability. Meeting these requirements requires an infrastructure that does not have any single point of failure.

NetApp clustered failover delivers a robust and highly available data service for business-critical environments. Installed on a pair of NetApp storage controllers, NetApp clustered failover provides data availability by transferring the data service of an unavailable storage controller to the other storage controller in the cluster.

The following figure shows a sample clustered failover configuration. A cluster is created with two storage controllers by connecting the storage controllers using a cluster interconnect. This connection is fully redundant and is used to exchange cluster heartbeats and to synchronize the NVRAM on both storage controllers. The disk shelves of the cluster partner are connected to the second storage controller using a second Fibre Channel loop. If the first storage controller fails, the second storage controller is able to access its partner's disk shelves. The MAC and IP addresses and the WWPN of the first storage controller are also adopted. Such a failover is transparent for the connected servers.

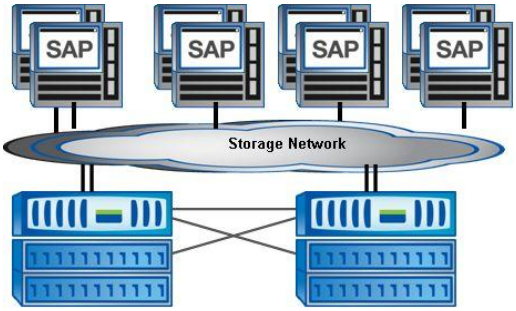

Figure 27) NetApp clustered storage system solution.

4.4 DISASTER RECOVERY

Organizations recognize the importance of having a bulletproof business continuance plan in place to deal with a disaster. The costs of not having one—lost productivity, revenue, customer loyalty, and possibly even business failure—makes it mandatory to have a plan that makes sure of an absolute minimum of downtime and rapid recovery from a disaster, with minimal or no loss of data. NetApp offers several solutions that can be configured to meet your corporation's specific recovery point objective (RPO) and recovery time objective (RTO).

SNAPMIRROR

NetApp SnapMirror software delivers a disaster recovery solution that today's global SAP systems need. By replicating data at high speeds over a LAN or a WAN, SnapMirror software provides the highest possible data availability and the fastest recovery. SnapMirror is integrated within SnapManager® for SAP using Protection Manager.

SnapMirror technology mirrors data to one or more storage controllers. It updates the mirrored data to keep it current and is now available for disaster recovery, tape backup, read-only data distribution, testing, online data migration, and more. SnapMirror supports synchronous, semi- synchronous, and asynchronous mirroring. This chapter will describe the asynchronous version of SnapMirror.

SnapMirror performs an initial transfer to initialize the disaster recovery site. After the initial transfer, incremental changes are passed to the disaster recovery site asynchronously. The SnapMirror disaster recovery solution is based on the NetApp backup and recovery solution: Snapshot backups are mirrored to the disaster recovery site. Additionally, the volumes where the archive logs and the archive log backups are stored are mirrored using SnapMirror. The frequency of SnapMirror updates of the archive logs and archive log backups determines the amount of data lost in the event of a disaster.

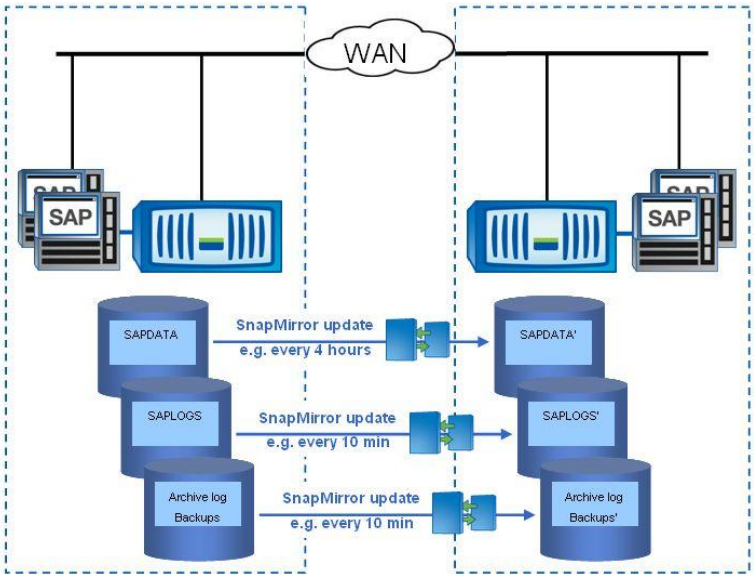

Figure 28) Disaster Recovery with SnapMirror.

METROCLUSTER

NetApp MetroCluster is an integrated high-availability and business continuance solution that provides disaster recovery with no data loss. MetroCluster extends failover capability from within a data center to a site located many miles away. It also replicates data from the primary site to the remote site to make sure that data there is completely current. The combination of failover and data replication makes sure that you can recover from disaster—with no loss of data—in minutes rather than hours or days.

MetroCluster is much like NetApp clustered failover but with the added benefit of disaster recovery. Clustered failover creates a cluster of NetApp storage appliances in one location with access to both sets of disk. MetroCluster extends this cluster configuration to remote locations up to 100 kilometers. Because there is no physical connection to the cluster appliance's disk in case of a site failure, MetroCluster requires the use of SyncMirror® to make sure that both storage controllers in the cluster have copies of the other storage controller's data. Compared to other storage disaster recovery solutions MetroCluster offers the functionality to recover from a disaster by just executing one command at the surviving controller. Even this can be automated to recover from a disaster transparently.

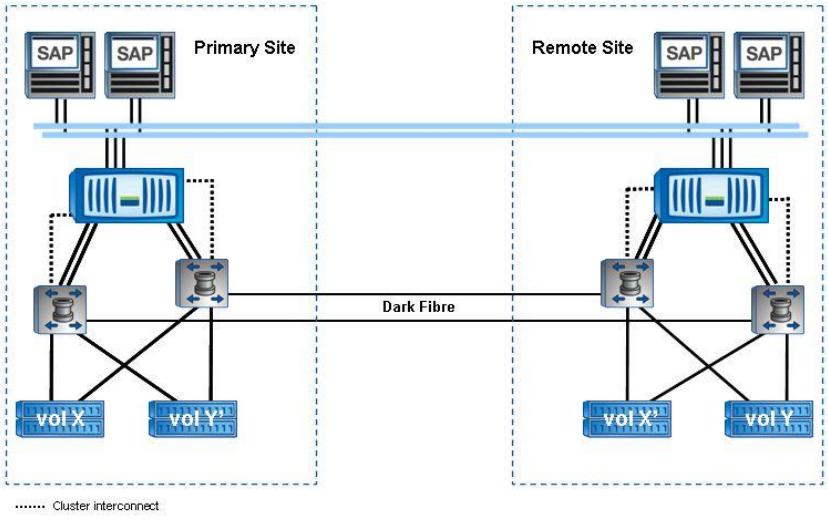

Figure 29) NetApp MetroCluster.

5 SNAPMANAGER FOR SAP

5.1 OVERVIEW CONFIGURATION SCENARIOS

SnapManager for SAP (SMSAP) provides the following functionalities supporting backup, recovery and cloning of SAP systems:

- Backups based on Snapshot copies (local backup)
- Backups restored on storage level using SnapRestore (from local backup)
- Database verification with Oracle dbv
- Data protection using SnapVault or SnapMirror with Protection Manager (remote backup)
- Restoring from remote backups
- Cloning from local backup
- Cloning from remote backup

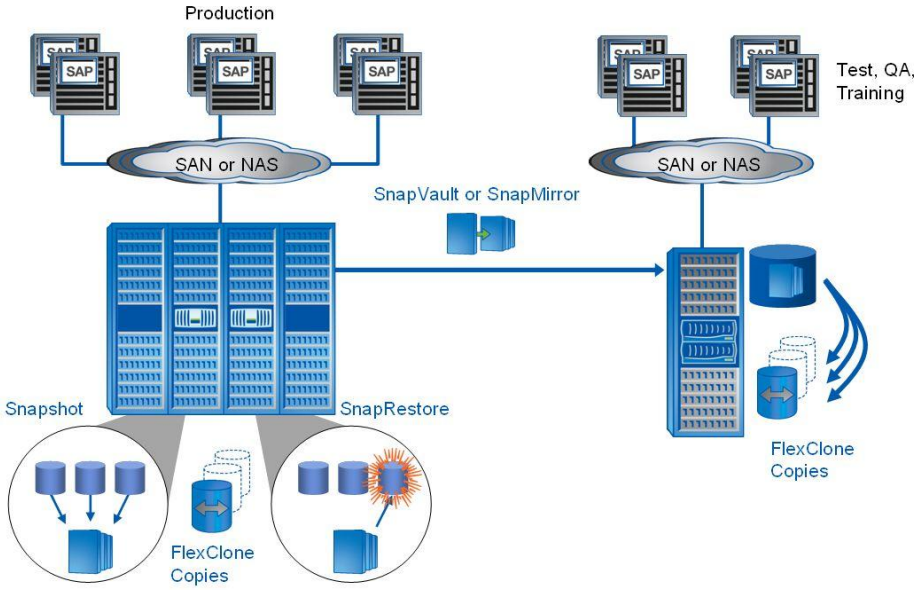

Figure 30) SMSAP Overview.

SMSAP supports two different methods to manage backup, recovery and cloning of SAP systems.

- Using the SAP BR*Tools together with (SMSAP) BACKINT
- SMSAP GUI or CLI

Both methods can be used in combination, but can't be mixed. For example, a backup created with BRBACKUP can only be restored with BRRESTORE and a backup created with SMSAP can only be restored with SMSAP GUI or CLI.

The following table shows which features are supported with the two different methods.

Table 7) SMSAP supported features.

	Local Backup Recovery from Local Backup	Data Protection Recovery from Remote Backup	Cloning from Local Backup Cloning from Remote Backup
BRBACKUP BRRESTORE BRRECOVER	Yes	No	No
SMSAP GUI SMSAP CLI	Yes	Yes	Yes

With the current release of SMSAP 3.0.3 archive log management is not included using the SMSAP GUI or CLI. Archive logs could be managed using the BACKINT implementation together with the BR*Tools.

Table 8) SMSAP archive log management.

	Local Backup Restore from Local Backup	Data Protection Recovery from Remote Backup
BRARCHIVE BRRESTORE BRRECOVER	Yes	No
SMSAP GUI SMSAP CLI	No	No

Using the BACKINT interface for the archive log management has limitations that need to be considered:

- No data protection, only Snapshot backups at primary storage
- Snapshot backups in the archive log volume need disk space based on the retention policy, even when archive logs get deleted by BRARCHIVE.

Therefore it is not recommended to use BACKINT for the archive log backups. The recommendation is to use BRARCHIVE without BACKINT to manage the archive log backup as shown in the figure below.

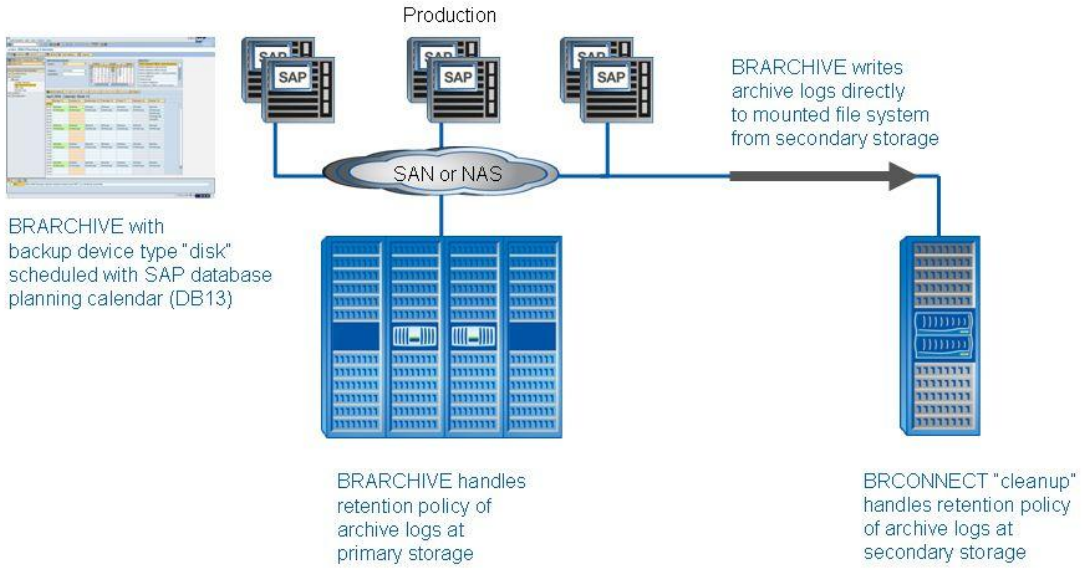

Figure 31) Archive log backup overview.

BRARCHIVE is configured with “backup_dev_type = disk.” The “archive_copy_dir” is a mountpoint from a secondary storage system. The storage access protocol can be NFS, iSCSI, or FCP.

The retention policy for archive logs at the primary storage is controlled by executing BRARCHIVE with the corresponding options “save,” “delete saved,” or “save delete.”

BRCONNECT is used with the option “cleanup” to handle the retention policy of the archive logs at the secondary storage.

It is optional to mirror the archive logs from the secondary storage to a third location in order to make sure that two copies of the archive logs are always available.

The forward recovery using BRRECOVER will also be much faster using this approach, because BRRECOVER will not need to restore the archive logs from the secondary storage. Since the archive logs are accessible from the database host, BRRECOVER will apply the logs directly from the secondary storage without the need of restoring the logs to /`oracle/<SID>/oraarch` in a first step.

RECOMMENDED SCENARIO WHEN USING THE BR*TOOLS

The BR*Tools together with SMSAP BACKINT are used for backup, restore, and recovery of local backups.

Data protection and backups as source for SAP system copies need to be done by creating additional backups with SMSAP GUI or CLI.

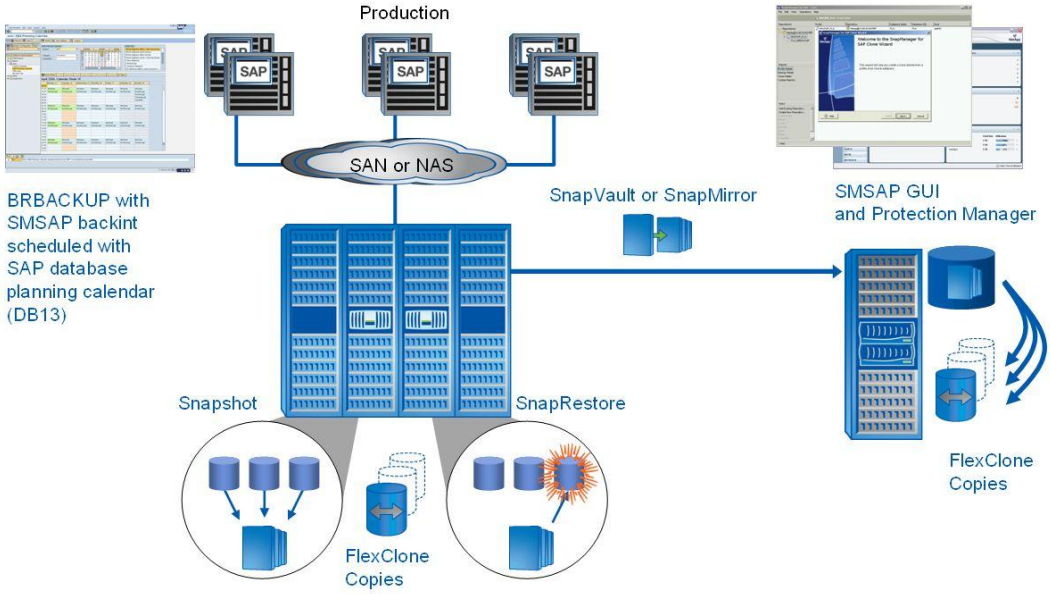

Figure 32) Scenario BR*Tools.

Archive log backups are managed with BRARCHIVE without SMSAP BACKINT. BRARCHIVE is configured to save the archive logs directly to a mountpoint at a secondary storage system.

Table 9) Scenario BR*Tools.

Task	
Local backup	BRBACKUP with SMSAP BACKINT
Restore from local backup	BRRESTORE, BRRECOVER with SMSAP BACKINT
Archive log backup	BRARCHIVE with backup_dev_type = disk without SMSAP BACKINT
Archive log restore	BRRESTORE, BRRECOVER with backup_dev_type = disk without SMSP BACKINT
Data protection (remote backup)	SMSAP GUI or CLI; additional daily or on demand backups
Restore from remote backup	SMSAP GUI or CLI
Cloning from local backup	SMSAP GUI or CLI; on demand backups for cloning purposes
Cloning from remote backup	SMSAP GUI or CLI; based on backups that have been created with SMSAP GUI or CLI

RECOMMENDED SCENARIO WHEN USING SMSAP GUI OR CLI

SMSAP GUI or CLI can be used to backup, restore and recover local or remote backups. The local and remote backups can also be used as a source for SAP system copies.

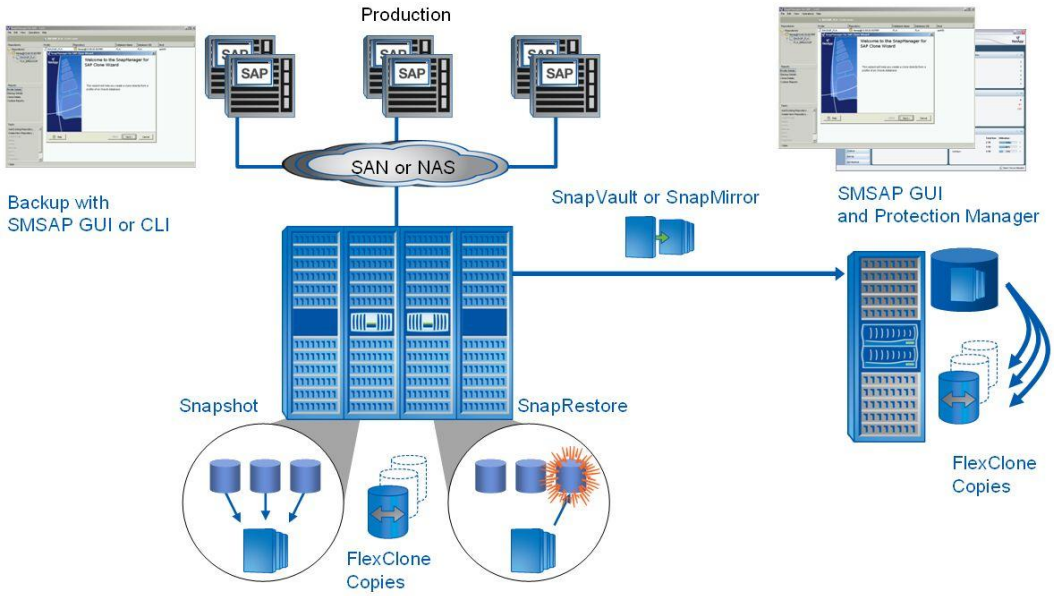

Figure 33) Scenario SMSAP GUI or CLI.

Archive log backups are managed with BRARCHIVE without SMSAP BACKINT. BRARCHIVE is configured to save the archive logs directly to a mountpoint at a secondary storage system.

If archive logs need to get restored for a restore and recovery of the database, it has to be done with BRRESTORE before the SMSAP GUI or CLI is used to run the database restore and recovery.

Table 10) Scenario SMSAP GUI or CLI.

Task	
Local backup	SMSAP GUI or CLI
Restore from local backup	SMSAP GUI or CLI
Archive log backup	BRARCHIVE with backup_dev_type = disk without SMSAP BACKINT
Archive log restore	BRRESTORE with backup_dev_type = disk without SMSAP BACKINT
Data protection (remote backup)	SMSAP GUI or CLI
Restore from remote backup	SMSAP GUI or CLI
Cloning from local backup	SMSAP GUI or CLI
Cloning from remote backup	SMSAP GUI or CLI

5.2 CONFIGURATION SCENARIO FOR BR*TOOLS

SMSAP CONFIGURATION

Two SMSAP profiles need to be configured.

One profile is used for the backups that are created by BRBACKUP. The other profile is used for the backups that are created by SMSAP GUI or CLI for data protection and SAP system copies.

The difference between these two profiles is the retention policy and the data protection option.

Table 11) SMSAP Profile configuration.

	Retention Policies for Local Backups	Data Protection	Remark
Profile for BRBACKUP	For example: six Hourly five Daily	No	Hourly Snapshot copy every six hours, kept for two days. Daily Snapshot copy once per day, kept for five days.
Profile for SMSAP GUI	For example: one Daily	Yes	Keep only one Snapshot copy locally; more backups are kept at secondary storage to be configured with Protection Manager.

The users ora<SID> and <SID>adm need access to the repository and the profiles because both users interact with the repository. The credentials for the repository and profile must be set, and the profile must be synced:

- smsap credential set –repository –dbname <repository> –host <host> –port <port> –login –username <username> –password <password>
- smsap profile sync –repository –dbname <repository> –host <host> –port <port> –login –username <username>
- smsap credential set –profile –name <profile name> –password <password>

CONFIGURATION BR*TOOLS

In order to be able to use different retention classes with BRBACKUP backups, it is necessary to create two different init<SID>.sap parameter files as shown in the table below.

For archive log backups with BRARCHIVE a separate parameter file is used in order to allow configuring the disk destination for the backups instead of using BACKINT.

Table 12) BR*TOOLS configuration.

	init<SID>.sap file	init<SID>.utl file
BRBACKUP Retention class hourly	**init<SID>_Hourly.sap** backup_dev_type = util_file util_par_file = init<SID>_Hourly.utl	**init<SID>_Hourly.utl** profile_name = <SID>_BRBACKUP fast = override retain = hourly protect = no
BRBACKUP Retention class daily	**init<SID>_Daily.sap** backup_dev_type = util_file util_par_file = init<SID>_Daily.utl	**init<SID>_Daily.utl** profile_name = <SID>_BRBACKUP fast = override retain = daily protect = no
BRARCHIVE	**Init<SID>_BRARCHIVE.sap** backup_dev_type = disk archive_copy_dir = /mnt/backup2disk	NA

CONFIGURATION DATA PROTECTION

Data protection is configured for the second profile (for example, <SID>_SMSAP). Typically a protection policy is configured within Protection Manager before this policy will get linked to the SMSAP profile.

Within the protection policy the retention policy of the backups at the secondary storage gets defined. When creating the profile with the SMSAP GUI, data protection and the data protection policy can be chosen. Within the SMSAP profile the retention policy of the backups at the primary storage is configured. This retention policy should be set to “daily count = 1” and “daily duration = 1.” With this configuration there is always only one Snapshot copy at the primary storage which will not be older than one day.

The retention policy of backups at the secondary storage is defined with Protection Manager and is typically one to four weeks.

VOLUME LAYOUT AND SNAPSHOT COPIES

The following table shows the recommended layout for the BR*Tools scenario. By separating the archive logs in an additional volume, no Snapshot copies are created for the archive log file system, which typically has a high change rate.

As soon as backups are created with SMSAP GUI or CLI for data protection or SAP system copies there are Snapshot copies on the archive log volume. In order to avoid that much disk capacity is needed for these Snapshot copies, it is recommended to only keep as few as possible Snapshot copies (backup) at the primary storage system.

More backups are kept at the secondary storage system, for example two weeks’ worth.

Table 13) Volume layout and Snapshot copies.

NetApp FlexVol	sapdata_<SID>	saplog_<SID>	mirrlog_<SID>	saparch_<SID>
Host Disk Group (SAN/iSAN only)	**sapdata_dg**	**saplog_dg**	**sapmirr_dg**	**saparch_dg**
File Systems	**/oracle/<SID>/sapdata1** **/oracle/<SID>/sapdata2** **/oracle/<SID>/sapdata3** **/oracle/<SID>/sapdata4**	**Controlfile1 in origlogA** **Controlfile2 in origlogB** **/usr/sap/<SID>** **/sapmnt/<SID>** **/oracle/<SID>** **/oracle/<SID>/origlogA** **/oracle/<SID>/origlogB**	**Controlfile3 in mirrlogA** **/oracle/<SID>/mirrlogA** **/oracle/<SID>/mirrlogB**	**/oracle/<SID>/oraarch**
BRBACKUP online backup (local backups) Recommended retention policy: 3-5 days	**First Snapshot copy** All data files get backed up	**First Snapshot copy** cntrl<SID>.dbf in /../sapbackup **Second Snapshot copy** init<SID>.ora, and so on in /../dbs space<SID>.log in /../sapreorg back.log, and so on in /../sapbackup	----------	----------
BRBACKUP offline backup (local backups) Recommended retention policy: 3-5 days	**First Snapshot copy** All data files get backed up	**First Snapshot copy** cntrl<SID>.dbf in /../origlogA All redologs in /../origlogA /../origlogB **Second Snapshot copy** init<SID>.ora, and so on in /../dbs space<SID>.log in /../sapreorg back.log, and so on in /../sapbackup	----------	----------
SMSAP GUI or CLI online backup (used for data protection and/or cloning) Recommended retention policy: 1 day at primary storage	**First Snapshot copy** All data files get backed up	**First Snapshot copy** Controlfile1 Controlfile2	**First Snapshot copy** Controlfile3	**First Snapshot copy** All archive logs (needed for cloning)
SMSAP GUI or CLI offline backup (Used for data protection and/or cloning) Recommended retention policy: 1 day at primary storage	**First Snapshot copy** All data files get backed up	**First Snapshot copy** Controlfile1 Controlfile2	**First Snapshot copy** Controlfile3	**First Snapshot copy** All archive logs

DATABASE VERIFICATION

The following table gives an overview of the different options to run the database verification.

Table 14) Database verification.

	Process	Advantages/Disadvantages
BRBACKUP –w use_dbv	• Backup with BACKINT • Restore with BACKINT to "compress_dir" • Brbackup runs Oracle dbv	Pros: • BRBACKUP integrated Cons: • Load on database server • Slow restore using host-copy • Disk space for restore
SMSAP GUI dbverify	• Backup with SMSAP • SMSAP mounts backup • SMSAP runs Oracle dbv	Pros: • SMSAP integrated • No restore, no disk space Cons: • Load on database server
Mount backup at separate server and run Oracle dbv	• Mount backup with SMSAP • Run Oracle dbv manually	Pros: • Verify decoupled from db server • No restore, no disk space Cons: • No product integration • Manual start of Oracle dbv or scripting necessary

SCHEDULING BACKUP JOBS WITHIN SAP CCMS AND THE SMSAP SCHEDULER

With the SAP transaction "dbacockpit" or "db13" all necessary backup and housekeeping jobs can now be scheduled:

- Database backups with BRBACKUP and retention class "Hourly"
- Database backups with BRBACKUP and retention class "Daily"
- Archive log backups with BRARCHIVE
- Deleting of already saved archive logs with BRARCHIVE at primary storage
- Deleting of archive logs with BRCONNECT at secondary storage

The jobs for protected backups and backups for SAP system copies are scheduled using the SMSAP scheduler.

The following table shows an example of a database backup, archive log backup and housekeeping jobs with the following assumptions:

- A full online backup should be created every six hours and kept for two days at the primary storage. The retention policy gets defined within the SMSAP profile for the retention class "Hourly."
- A full online backup should be created once per day and kept for five days. The retention policy gets defined within the SMSAP profile for the retention class "Daily."
- Archive logs should be backed up every hour and should not be kept longer than six hours at the primary storage:
  BRARCHIVE –s (save option) every hour
  BRARCHIVE –ds (delete saved) every six hours

- Archive logs are kept for 30 days at the secondary storage
  (default value for cleanup in init<SID>.sap).
  BRCONNECT –f cleanup is scheduled once per week
- A full online backup should be created and replicated to a secondary storage once per day.
  A backup is created using the SMSAP GUI scheduler once per day with data protection.
- A full online backup with database verification should be done once per week.
  A backup is created using the SMSAP GUI scheduler once per week with the dbverify option.

Table 15) Job scheduling overview

	00:00	01:00	06:00	07:00	12:00	13:00	18:00	19:00
Mon	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Tue	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Wed	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Thu	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Fri	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Sat	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap
Sun	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDDaily.sap BRCONNECT –f cleanup initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with database verification	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap BRBACKUP online initSIDHourly.sap	Brarchive –ds initSID_Brarchive.sap

.

5.3 CONFIGURATION SCENARIO FOR SMSAP GUI OR CLI

SMSAP CONFIGURATION

One SMSAP profile needs to be configured. The profile is used for the local backups, protected backups and SAP system copies.

VOLUME LAYOUT AND SNAPSHOT COPIES

The following table shows the recommended layout for this scenario.

When backups are created with SMSAP GUI or CLI there will be Snapshot copies on the archive log volume, which typically has a high change rate. In order to avoid that much disk capacity is needed for these Snapshot copies, it is recommended to keep only one to three days of backups at the primary storage system. More backups are kept at the secondary storage system, for example two weeks' worth.

Table 16) Volume layout and Snapshot copies.

NetApp FlexVol	**sapdata_<SID>**	**saplog_<SID>**	**mirrlog_<SID>**	**saparch_<SID>**
Host Disk Group (SAN/iSAN only)	**sapdata_dg**	**saplog_dg**	**sapmirr_dg**	**saparch_dg**
File Systems	**/oracle/<SID>/sapdata1** **/oracle/<SID>/sapdata2** **/oracle/<SID>/sapdata3** **/oracle/<SID>/sapdata4**	**Controlfile1 in origlogA** **Controlfile2 in origlogB** **/usr/sap/<SID>** **/sapmnt/<SID>** **/oracle/<SID>** **/oracle/<SID>/origlogA** **/oracle/<SID>/origlogB**	**Controlfile3 in mirrlogA** **/oracle/<SID>/mirrlogA** **/oracle/<SID>/mirrlogB**	**/oracle/<SID>/oraarch**
SMSAP GUI or CLI online backup (local backups also used for data protection and/or cloning) Recommended retention policy: 1-3 days	One Snapshot copy All data files get backed up	One Snapshot copy Controlfile1 Controlfile2	One Snapshot copy Controlfile3	One Snapshot copy All archive logs (needed for cloning)
SMSAP GUI or CLI offline backup (local backups also used for data protection and/or cloning) Recommended retention policy: 1-3 days	One Snapshot copy All data files get backed up	One Snapshot copy Controlfile1 Controlfile2	One Snapshot copy Controlfile3	One Snapshot copy All archive logs

CONFIGURATION DATA PROTECTION

Typically a protection policy is configured within Protection Manager before this policy will get linked to the SMSAP profile.

Within the protection policy the retention policy of the backups at the secondary storage gets defined. When creating the profile with the SMSAP GUI data protection and the data protection policy can be chosen. Within the SMSAP profile the retention policy of the backups at the primary storage is configured.

The retention policy of backups at the secondary storage is defined with Protection Manager and is typically one to four weeks.

DATABASE VERIFICATION

The following table gives an overview of the different options to run the database verification.

Table 17) Database verification.

	Process	Advantages/Disadvantages
SMSAP GUI dbverify	• Backup with SMSAP • SMSAP mounts backup • SMSAP runs Oracle dbv	Pros: • SMSAP integrated • No restore, no disk space Cons: • Load on database server
Mount backup at separate server and run Oracle dbv	• Mount backup with SMSAP • Run Oracle dbv manually	Pros: • Verify decoupled from db server • No restore, no disk space Cons: • No product integration • Manual start of Oracle dbv or scripting necessary

SCHEDULING BACKUP JOBS WITHIN SAP CCMS AND THE SMSAP SCHEDULER

With the SAP transaction “dbacockpit” or “db13” all necessary archive log backup and housekeeping jobs can now be scheduled:

- Archiving log backups with BRARCHIVE
- Deleting of already saved archive logs with BRARCHIVE at primary storage
- Deleting of archive logs with BRCONNECT at secondary storage

The jobs for protected backups and backups for SAP system copies are scheduled using the SMSAP scheduler.

The following table shows an example of a database backup, archive log backup and housekeeping jobs with the following assumptions:

- A full online backup should be created every six hours and kept for two days at the primary storage. The retention policy gets defined within the SMSAP profile for the retention class “Hourly”.
- A full online backup should be created once per day and kept for five days.
  The retention policy gets defined within the SMSAP profile for the retention class “Daily”.
- Archive logs should be backed up every hour and should not be kept longer than six hours at the primary storage:
  BRARCHIVE –s (save option) every hour
  BRARCHIVE –ds (delete saved) every six hours

- Archive logs are kept for 30 days at the secondary storage
  (default value for cleanup in init<SID>.sap).
  BRCONNECT –f cleanup is scheduled once per week
- A full online backup should be created and replicated to a secondary storage once per day.
  A backup is created using the SMSAP GUI scheduler once per day with data protection.
- A full online backup with database verification should be done once per week.
  A backup is created using the SMSAP GUI scheduler once per week with the dbverify option

Table 18) Job scheduling overview

	00:00	01:00	06:00	07:00	12:00	13:00	18:00	19:00
Mon	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Tue	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Wed	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Thu	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Fri	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Sat	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with data protection	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap
Sun	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup BRCONNECT –f cleanup initSIDDaily.sap	Brarchive –ds initSID_Brarchive.sap SMSAP scheduled online backup with database verification	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap	Brarchive –s every hour initSID_Brarchive.sap SMSAP scheduled online backup	Brarchive –ds initSID_Brarchive.sap

.

5.4 SAP SYSTEM COPY OVERVIEW

This chapter explains how SnapManager for SAP and the underlying Snapshot and FlexClone features can be used together with SAPinst to create an SAP system copy.

The chapter also describes how SAPinst can be used with minimal or no required user input in order to accelerate the system copy process.

A more detailed, step-by-step description can be found at SAP System Copy with SnapManager for SAP.

Within the chapter two scenarios are covered:

- **New installation:** Complete fresh setup of the target system based on a database copy of the source system
- **System refresh:** Refresh the database of the target system based on a database copy of the source system

The following table shows the differences between using SAPinst standalone and using SAPinst together with SMSAP.

Table 19) System copy overview.

	ABAP or Java Stack	New Installation		System Refresh	
		SAPinst Standalone	SAPinst with SMSAP	SAPinst Standalone	SAPinst with SMSAP
SAPinst needs to run on source system	ABAP	No	No	No	No
	Java	Yes	Yes	Yes	Yes
OraBrCopy needs to run on source system	ABAP	Yes	No	Yes	No
	Java	Yes	No	Yes	No
Backup of source system	ABAP	Offline	Online	Offline	Online
	Java	Offline	Online	Offline	Online
SAPinst needs to run on target system	ABAP	Yes	Yes	Yes	No
	Java	Yes	Yes	Yes	Yes

The main advantages of SMSAP are:

- Online backups based on Snapshot copies are used as a base for the system copy. The Snapshot backup can be done at any time and has no effect on the production source system.
- A FlexClone copy is used to attach the online Snapshot backup of the source system to the target system. The FlexClone copy is done in a few minutes.
- SMSAP handles all database-specific tasks of the system copy. Therefore there is no need to run OraBrCopy anymore.
- For systems with ABAP stack only a system refresh is done without the need to run SAPinst.

5.5 SAP SYSTEM COPY JAVA AND ABAP STACK

SAP STANDARD APPROACH

The following figure gives an overview of the system copy process of an SAP system with double stack or Java stack only. This process is described in the SAP Guide: “System Copies for SAP Systems based on SAP NetWeaver 7.0 SR3 ABAP & Java” in chapter “Database-specific system copy.”

There are several steps that need to be executed on the source system:

- Background jobs need to be set to “scheduled” in order to avoid that these jobs get started on the target system. Since SAPinst will start the SAP system on the target server during the installation the background jobs need to be handled on the source system.
- SAPinst is used to archive the Software Deployment Manager (SDM) data.
- The OraBrCopy tool is used to create a “control file to trace.”
- An offline backup needs to be created.
  SAPinst only supports offline backups as a source for the system copy.

The SDM archive and the CONTROL.SQL file are required for the SAPinst process on the target system.

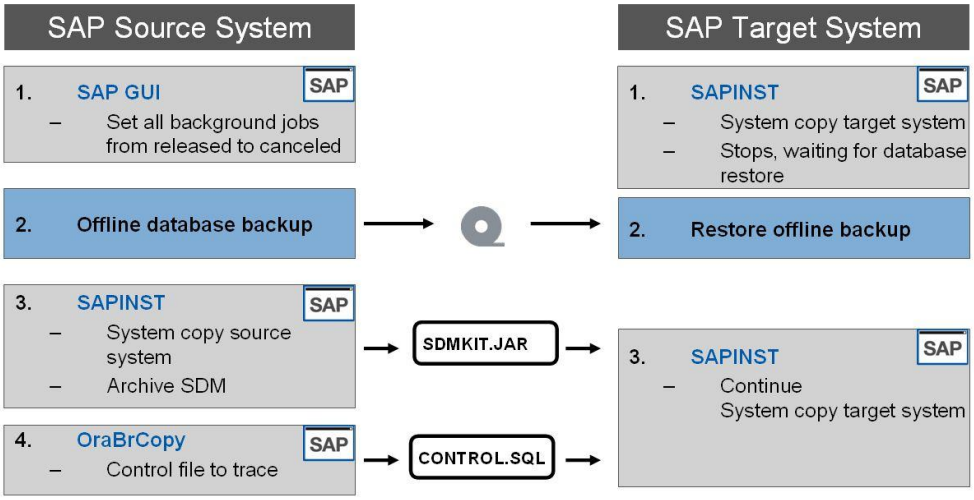

Figure 34) SAP system copy double stack, standard approach.

The SAP standard approach includes many manual tasks and can’t be automated.

SYSTEM COPY WITH SMSAP: NEW INSTALLATION

The approach of creating a SAP system copy for a new installation with SMSAP differs in the following tasks:

- There is no need to run OraBrCopy on the source system, since SMSAP will handle the change of the database SID on the target host.
- Online Snapshot backups can be used as the source of the system copy, because SMSAP will handle the recovery of the online backup on the target host. SAPinst won't see a difference.
- There is no need to handle the background jobs on the source system before creating the Snapshot backup. The handling of the background jobs is done within a postcloning plug-in of SMSAP.

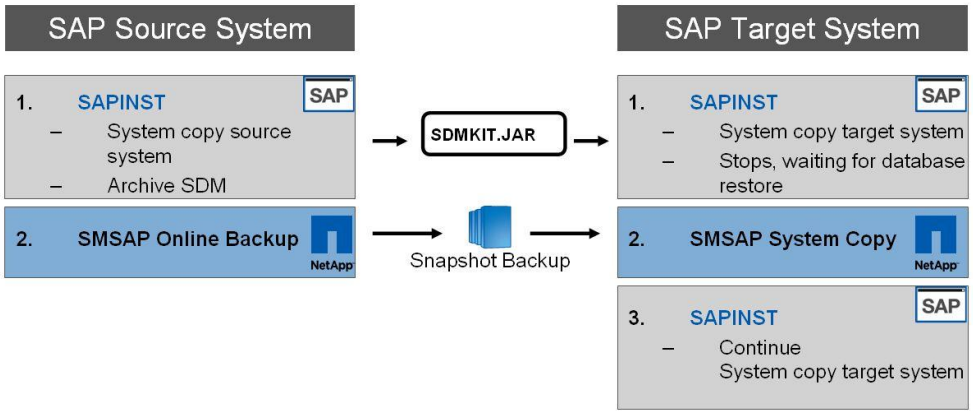

Figure 35) SAP system copy double stack, new installation, NetApp approach.

During the new installation of the target system all required input and configuration should be done and prepared in order to be able to accelerate and automate any further refresh of the target system.

- Preparation of "unattended mode" for SAPinst on the source system.
  SAPinst will run without any user input.
- Preparation of SAPinst input files on the target host.
  SAPinst will run with minimal user input.
- Preparation of cloning specification for SMSAP.
  SMSAP will run in unattended mode.

The installation includes the following steps:

- (Source)
  Run SAPinst to archive SDM data.
  Preparation of unattended mode for all subsequent runs.
- (Source)
  Run SMSAP to create online backup of the SAP system.
- (Target)
  Install operating system; prepare file systems for SAP installation. Install SnapDrive and SMSAP.
- (Target)
  Run SAPinst with system copy service.
  Preparation of all input for subsequent runs.
  SAPinst will stop for the restore of the database.
- (Target)
  Run SMSAP to create a clone of the database.
  During this step the cloning specification parameter can be saved in a file for subsequent runs.

- (Target)
  Continue SAPinst.

SYSTEM COPY WITH SMSAP: REFRESH OF TARGET SYSTEM

A refresh of the target system includes the following steps.

- (Target)
  Stop SAP system only. Database needs to run when clone is deleted with SMSAP in the next step.
- (Target)
  Delete clone using SMSAP.
- (Source)
  Create online backup of source system with SMSAP (optional).
- (Target)
  Create clone using SMSAP and the Snapshot copy created before or any existing Snapshot copy. The SMSAP cloning process uses the clone specification that has been created during the new installation of the target system as described in the previous chapter. When SMSAP has finished the cloning process the database on the target system is up and running with the changed SID. The background jobs are already canceled by a SQL script that is executed as a postcloning plug-in of SMSAP.
- (Source)
  Run SAPinst on the source system in “unattended mode” using the environment that has been prepared during the new installation of the target system.
- (Target)
  Run SAPinst on the target system using the environment that has been prepared during the new installation of the target system. Only minimal user input is required. When SAPinst stops to ask for the restore of the database it can be continued immediately, since the database is already made available by SMSAP.

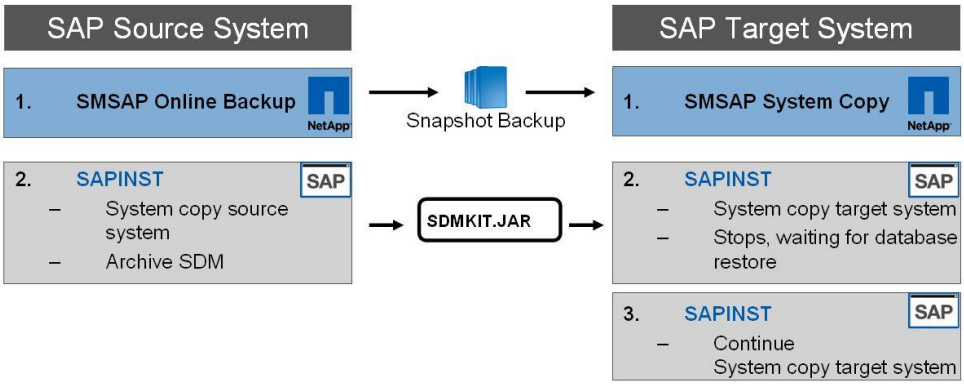

Figure 36) SAP system copy double stack, system refresh, NetApp approach.

Refreshing the target system only requires minimal user input when running SAPinst on the target system. All the other tasks can run unattended.

5.6 SAP SYSTEM COPY ABAP STACK ONLY

SAP STANDARD APPROACH

The following figure gives an overview of the system copy process of an SAP system with ABAP stack only. This process is described in the SAP Guide: “System Copies for SAP Systems based on SAP NetWeaver 7.0 SR3 ABAP” in chapter “Database-specific system copy.”

There are several steps that need to be executed on the source system:

- Background jobs need to be set to “scheduled” in order to avoid that these jobs get started on the target system. Since SAPinst will start the SAP system on the target server during the installation the background jobs need to be handled on the source system.
- The OraBrCopy tool is used to create a “control file to trace.”
- An offline backup needs to be created.
  SAPinst only supports offline backups as a source for the system copy.

The CONTROL.SQL file needs to be made available to the SAPinst process on the target system.

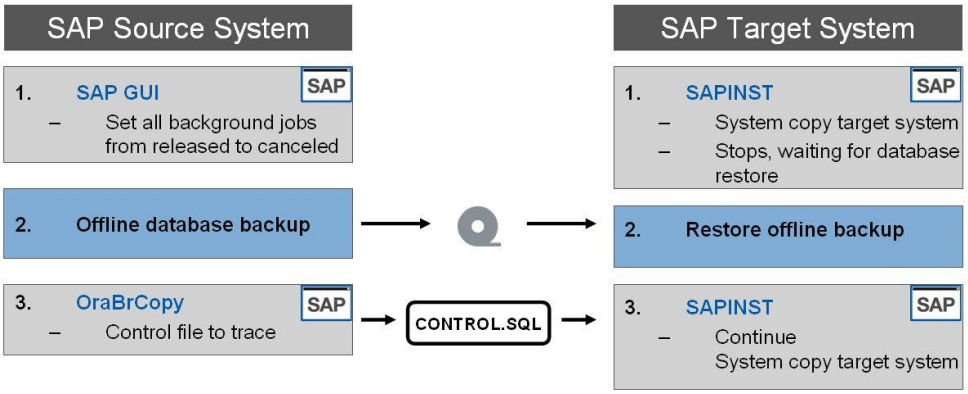

Figure 37) SAP system copy ABAP stack, standard approach.

The SAP standard approach includes many manual tasks and can’t be automated.

SYSTEM COPY WITH SMSAP: NEW INSTALLATION

The approach of creating a SAP system copy for a new installation with SMSAP differs in the following areas:

- There is no need to run OraBrCopy on the source system, since SMSAP will handle the change of the database SID on the target host.
- Online Snapshot backups can be used as the source of the system copy, because SMSAP will handle the recovery of the online backup on the target host. SAPinst won't see a difference.
- There is no need to handle the background jobs on the source system before creating the Snapshot backup. The handling of the background jobs is done within a postcloning plug-in of SMSAP.

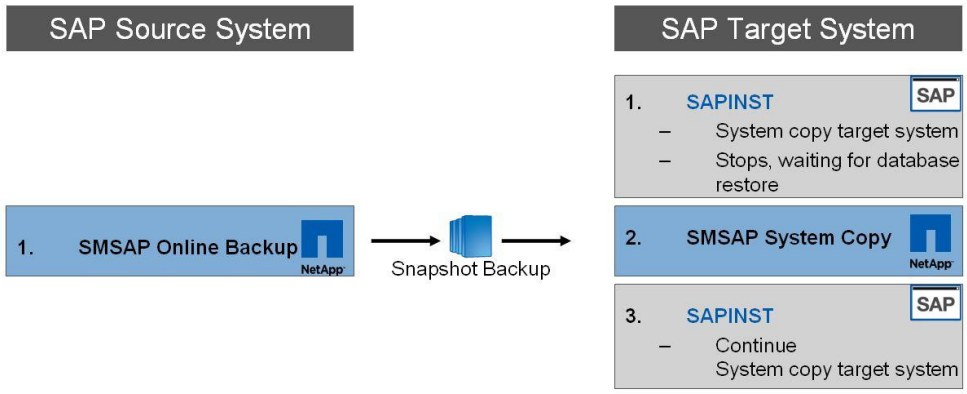

Figure 38) SAP system copy ABAP stack, new installation, NetApp approach.

The installation will include the following steps:

- (Source)
  Run SMSAP to create online backup of the SAP system.
- (Target)
  Install operating system; prepare file systems for SAP installation. Install SnapDrive and SMSAP.
- (Target)
  Run SAPinst with task system copy service.
- (Target)
  Run SMSAP to create a clone of the database. During this step the cloning specification parameters can be saved in a file for subsequent runs.
- (Target)
  Continue SAPinst.

SYSTEM COPY WITH SMSAP: REFRESH OF TARGET SYSTEM

A refresh of the target system will run with the following steps.

- (Target)
  Stop SAP system only. Database needs to run when clone is deleted with SMSAP in the next step.
- (Target)
  Delete clone using SMSAP.
- (Source)
  Create online backup of source system with SMSAP (optional).
- (Target)
  Create clone using SMSAP and the Snapshot copy created before or any existing Snapshotcopy. The SMSAP cloning process uses the clone specification that has been created during the new installation of the target system as described in the previous chapter. When SMSAP has finished the cloning process the database on the target system is up and running with the changed SID. The background jobs are already canceled by a SQL script that is executed as a postcloning plug-in of SMSAP.

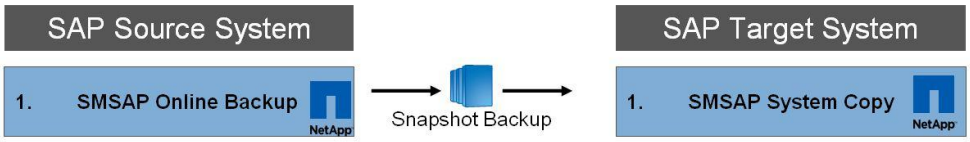

Figure 39) SAP system copy ABAP stack, system refresh, NetApp approach.

Refreshing the target system requires no user input. All tasks can run unattended.

.

6 ARCHIVING AND COMPLIANCE

The long-term accumulation of data in the SAP database affects the performance and availability of SAP applications. To keep the SAP systems and applications running at peak efficiency, it is important to implement a data archiving process to enhance availability while reducing performance and management overhead concurrently.

Choosing the media type and platform for archival storage requires companies to conform to many content retention mandates. IT organizations must respond by analyzing the business requirements and then choose the proper solution based on factors such as time to access data, risk, storage scalability, compatibility, and total cost of ownership (TCO). Current write once, read many (WORM) technologies such as optical disk and tape do not provide sufficiently rapid access or high reliability or have high TCO. What organizations need is a solution that easily and inexpensively integrates archived storage with corporate applications and enables them to comply with existing retention and security regulations for reference data.

In addition to managing system size and performance, SAP customers are keenly aware of increasing industry regulations that have introduced significant financial penalties for failing to comply with retention, indexing, auditing, privacy, and reporting requirements. These regulations span almost all public companies and industry sectors. Nearly every major corporation must put a regulatory compliance solution in place or face the risk of being exposed to litigation and fines. In most cases, this solution requires the purchase of new storage subsystem hardware and software.

Historically, most regulated data has been stored on optical devices, tape, paper, and/or microfiche/microfilm. Disks have not often been utilized in the past, due to a number of factors, including cost and the lack of necessity to retrieve information quickly. However, the Enterprise Strategy Group (ESG) estimates that moving forward, disks will be the fastest growing area for the storage of regulated data.

NetApp offers solutions for SAP Archive Development Kit (ADK)/ArchiveLink and SAP Information Lifecycle Management (ILM) that consists of three solutions: data archiving using the Web-Based Distributed Authoring and Versioning (WebDAV) protocol, Information Retention Manager, and Retention Warehouse.

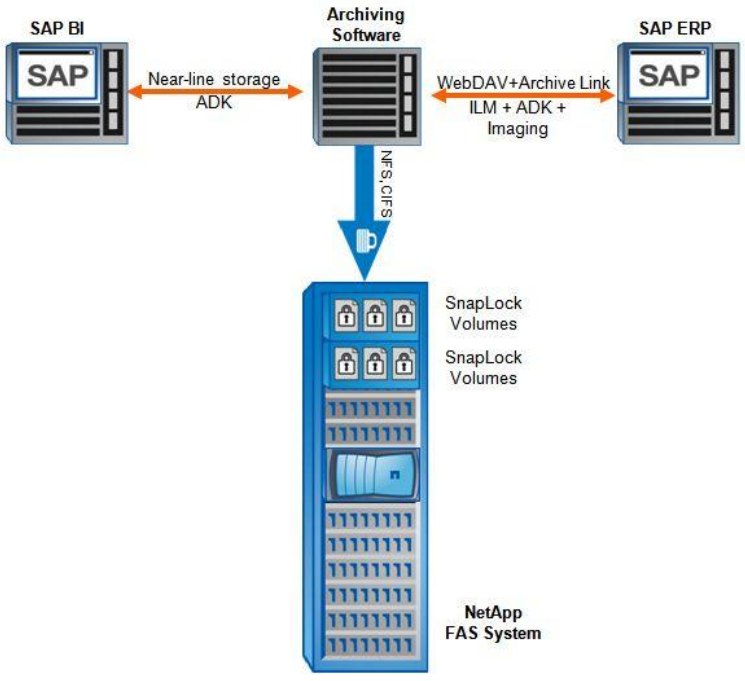

Figure 40) Archiving Overview.

The ADK is the software layer that encapsulates the technical aspects of data archiving programs and provides an application-programming interface that both customers and partners can use to develop their own archiving solutions. ArchiveLink is both an interface and a service for facilitating the process-driven management of business documents. Business-related documents are linked to and retrieved from application objects using workflow.

WebDAV is a set of extensions to the HTTP protocol that allows users to collaboratively edit and manage files on remote Web servers. The major features of the protocol are locking, metadata management, and namespace manipulation. NetApp storage is certified for solutions such as PBS ContentLink and OpenText Document Access which are certified for SAP ILM - WebDAV Storage Interface 2.0 (BC-ILM 2.0). More information can be found at the SAP Partner Information Center.

Once archive files have been created, the data marked for archive can be deleted from the source database. The archiving data will then be transferred directly from the primary storage system to the archive server. The SAP Retention Manager will set and manage the retention of the archived data based on legal requirements, including legal hold management. For system decommissioning of legacy systems SAP offers the Retention Warehouse. Detailed information can be obtained from SAP.

NetApp solutions for SAP archiving such as NetApp systems with SnapLock® work hand in hand with technologies from SAP and their archiving partners, such as OpenText, FileNet, PBS, or others. The result of effective SAP archiving is better-performing applications that cost less to operate and manage.

SnapLock is the NetApp implementation of high-performance disk-based magnetic WORM storage. SnapLock provides secure, storage-enforced data retention functionality using open file protocols such as CIFS and NFS while leveraging existing NetApp technologies. This implementation also includes significant efforts in securing Data ONTAP and its administrative interfaces to the degree that SnapLock can be deployed for protecting data in regulatory environments. Two SnapLock versions are available:

- SnapLock Compliance: Enables organizations to satisfy strict records retention regulations, such as SEC Rule 17a-4 (broker-dealers), HIPAA (health care), Sarbanes-Oxley (public companies), 21CFR Part 11 (life sciences), and DOD 5015.2 (government). Only an act of willful destruction, such as physically removing disks from a NetApp system, can result in record deletion or alteration prior to the specified retention date.
- SnapLock Enterprise: Enables adherence to rigorous organizational best practices through functionality similar to that of SnapLock Compliance, but allows administrators to delete entire SnapLock Enterprise volumes. It is not possible for any SnapLock Enterprise user or administrator to delete or modify individual SnapLock Enterprise WORM records.

NetApp provides a flexible, scalable, and secure solution for SAP ILM, compliance, and data archiving needs:

- SnapLock enables locking of some files without forcing WORM behavior for all data.
- There is no risk of software vendor lock-in. NetApp works well with existing document and content management packages such as OpenText, FileNet, PBS, and other software vendors.
- Data can be managed and backed up using the customer's current products and strategies.
- Better ROI and lower TCO are achieved through increased availability, enhanced system performance, lower administration overhead, and increased staff productivity.
- Compliance and archived data remain easily accessible on near-line storage, a more cost-effective alternative for archiving SAP data instead of adding database storage or processing power.

7 CONCLUSION

As SAP landscapes grow to support more and more business-critical applications, the job of maintaining those landscapes becomes increasingly complex. The NetApp solutions for SAP combine technologies that simplify and accelerate this process and align with the SAP application lifecycle.

The NetApp solutions for SAP accelerate upgrades and changes; enable fast backup, restore and SAP system copies; and provide simplified, economical, and highly available disk-based archiving. NetApp solutions help enterprises to reduce cost and complexity, minimize risk, and control change in their SAP environments.

www.netapp.com