Technical Report

SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP

Nils Bauer, NetApp
December 2013 | TR-4250

Abstract

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.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

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 datasets and without affecting production operations often determines whether SAP projects are delivered on time and within budget. SAP systems are often used globally, resulting in a 24/7 operation. Nondisruptive operation is therefore a key requirement.

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 that customers 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, BW, 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. Database technologies such as Oracle Real Application Clusters have introduced additional complexity into the database layer. Virtualization or cloud technologies have become more dominant 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 will 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 Clustered Data ONTAP Architecture

This section describes the architecture of clustered Data ONTAP®, with an emphasis on the separation of physical resources and virtualized containers. Virtualization of storage and network physical resources is the basis for scale-out and nondisruptive operations.

2.1 Hardware Support and Basic System Overview

As shown in Figure 1, a clustered Data ONTAP system consists of NetApp storage controllers (including V-Series) with attached disks. The basic building block is the high-availability (HA) pair, a concept familiar from Data ONTAP 7G and 7-Mode environments. An HA pair consists of two identical nodes, or instances of clustered Data ONTAP. Each node actively provides data services and has redundant cabled paths to the other node's disk storage. If either node is down for any reason, planned or unplanned, its HA partner can take over its storage and maintain access to the data. When the downed system rejoins the cluster, the partner node gives back the storage resources.

The minimum cluster size starts with two matching nodes in an HA pair. Using nondisruptive technology refresh, a two-node, entry-level cluster can evolve to the largest cluster size and most powerful hardware. At the time of writing this technical report, clusters with SAN protocols are supported up to eight nodes with midsized and high-end controllers. NAS-only clusters of high-end controllers scale up to 24 nodes and over 69PB of data storage.

Note: Clustered Data ONTAP 8.2 offers the additional option of a single-node cluster configuration. This is intended for smaller locations that replicate to a larger data center.

Note: Historically, the term "cluster" has been used to refer to an HA pair running Data ONTAP 7G or 7-Mode. This usage has been discontinued, and "HA pair" is the only correct term for this configuration. The term "cluster" now refers only to a configuration of one or more HA pairs running clustered Data ONTAP.

Figure 1) Data ONTAP cluster overview.

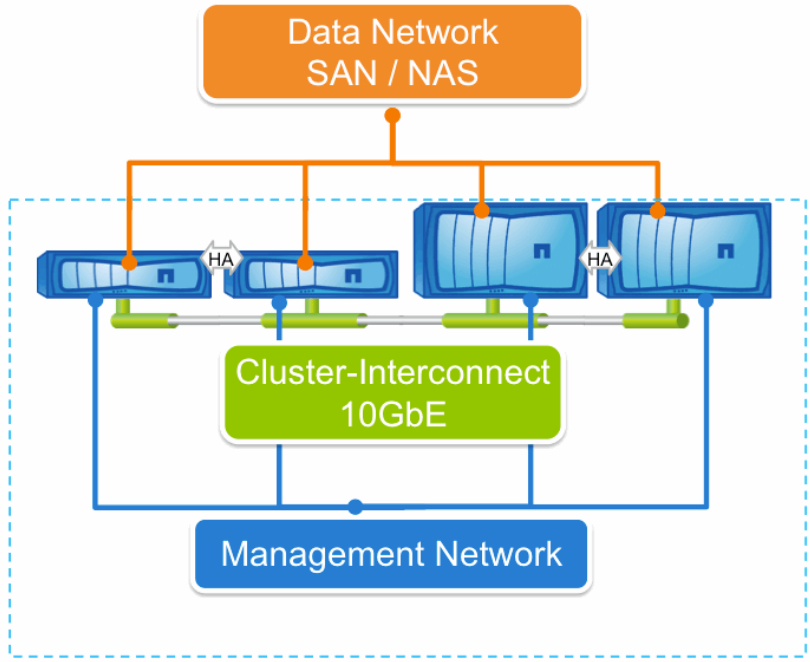

One of the key differentiators in a clustered Data ONTAP environment is that the storage nodes are combined into a cluster to form a shared pool of physical resources that are available to applications, SAN hosts, and NAS clients. The shared pool appears as a single system image for management purposes, providing a single common point of management, through GUI or CLI tools, for the entire cluster.

2.2 Scalability

Clustered Data ONTAP allows the inclusion of different controller types in the same cluster, protecting the initial hardware investment and providing the flexibility to adapt resources to meet business demands of the workloads. Similarly, support for different disk types, including SAS, SATA, and solid-state disk (SSD), makes it possible to deploy integrated storage tiering for different data types, together with the transparent DataMotion™ capabilities of clustered Data ONTAP. Flash Cache™ cards can also be used to provide accelerated read performance for frequently accessed data. Starting with version 8.1.1, clustered Data ONTAP supports Flash Pool™ intelligent caching, which combines SSD with traditional hard drives for optimal performance and efficiency. The highly adaptable clustered Data ONTAP architecture is key to delivering maximum on-demand flexibility for the shared IT infrastructure, offering flexible options to address needs for performance, price, and capacity.

Clustered Data ONTAP can scale both vertically and horizontally through the addition of nodes and storage to the cluster. This scalability, combined with protocol-neutral, proven storage efficiency, can meet the needs of the most demanding workloads.

2.3 Clustered Data ONTAP Networking

Figure 2 also shows the underlying network architecture of clustered Data ONTAP 8.1.2. Three networks are shown:

- **Cluster interconnect.** A private, dedicated, redundant network used for communication between the cluster nodes and for DataMotion data migration within the cluster. The cluster interconnect infrastructure is provided with every clustered Data ONTAP configuration to support this network. This infrastructure takes the form of redundant, high-performance, high-throughput 10Gb/sec enterprise-

class switch hardware in clusters of four or more nodes. Clusters of two nodes can optionally be configured without switches, with point-to-point connections used for the cluster interconnect. This configuration, available for the first time in clustered Data ONTAP 8.2, is known as a switchless cluster. This entry-level configuration provides all of the benefits of clustered Data ONTAP with a simpler infrastructure. Switchless clusters can be nondisruptively upgraded to include a switched cluster interconnect when the cluster grows beyond two nodes.

- **Management network.** All management traffic passes over this network. Management network switches can be included as part of a clustered Data ONTAP configuration, or customer-provided switches can be used. NetApp OnCommand® System Manager, OnCommand Unified Manager, and other NetApp applications are available for management, configuration, and monitoring of clustered Data ONTAP systems. System Manager provides GUI management, including a number of easy-to-use wizards for common tasks. Unified Manager provides monitoring and alerts. A powerful CLI is included, and management APIs are packaged and distributed in the Manage ONTAP® Software Developer's Kit.
- **Data networks.** These networks provide data access services over Ethernet or FC to the SAN hosts and NAS clients. They are provided by the customer according to requirements and could also include connections to other clusters acting as volume replication targets for data protection.

Figure 2 shows a larger cluster with different disk types and a mix of native NetApp FAS and V-Series controllers. V-Series makes it possible to use third-party storage with a NetApp controller as the front end, so that it can run clustered Data ONTAP and participate in a cluster. It also shows the client/host connections and the virtualized storage and network layer.

Figure 2) Clustered Data ONTAP large cluster.

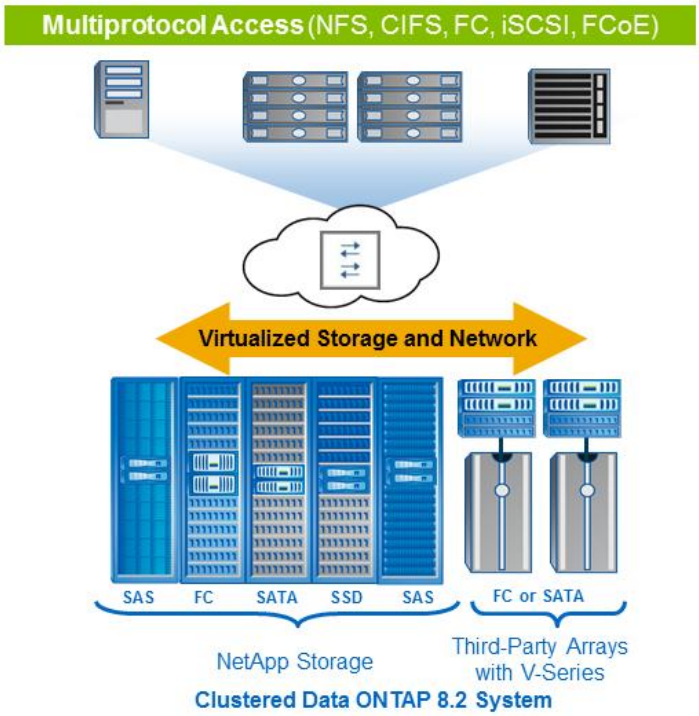

2.4 Storage Efficiency and Data Protection

Storage efficiency built into clustered Data ONTAP offers substantial space savings, allowing more data to be stored at lower cost. Data protection provides replication services so that valuable data is backed up and recoverable:

- **Thin provisioning.** Volumes are created by using virtual sizing. Thin provisioning is the most efficient way to provision storage because, although the clients see the total storage space assigned to them, the storage is not preallocated up front. In other words, when a volume or LUN is created by using thin provisioning, no space on the storage system is used. The space remains unused until data is written to the LUN or the volume, at which time only enough space to store the data is used. Unused storage is shared across all volumes, and the volumes can grow and shrink on demand.
- **NetApp Snapshot copies.** Automatically scheduled point-in-time copies that take up no space and incur no performance overhead when created. Over time, Snapshot copies consume minimal storage space, because only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- **NetApp FlexClone volumes.** Near-zero space, exact, writable virtual copies of datasets. They offer rapid, space-efficient creation of additional data copies ideally suited for test/dev environments.
- **NetApp FlexCache® volumes.** Provide the ability to cache volumes on other nodes in the cluster, thus balancing the read load to a frequently accessed volume across the cluster. The cache volumes are space efficient, because only the blocks accessed are cached. Data consistency is maintained through read delegation for files.
- **Deduplication.** Removes duplicate data blocks in primary and secondary storage, storing only unique blocks, which results in storage space and cost savings. Deduplication runs on a customizable schedule.
- **Compression.** Compresses data blocks by replacing repeating patterns within a subset of a file. Compression is complementary with deduplication; depending on the workload, compression only, deduplication only, or deduplication and compression together may provide the maximum space and cost savings.
- **NetApp SnapMirror.** Asynchronous replication of volumes, independent of protocol, either within the cluster or to another clustered Data ONTAP system for data protection and disaster recovery.
- **NetApp SnapVault®.** Volumes can be copied for space-efficient, read-only, disk-to-disk backup, either within the cluster or to another clustered Data ONTAP system.

2.5 Cluster Virtualization and Multi-Tenancy Concepts

A cluster is composed of physical hardware, including storage controllers with attached disk shelves; NICs; and, optionally, Flash Cache cards. Together, these create a physical resource pool, which is virtualized as logical cluster resources to provide data access. Abstracting and virtualizing physical assets into logical resources provide the flexibility and potential multi-tenancy in clustered Data ONTAP as well as the DataMotion ability at the heart of nondisruptive operations.

Physical Cluster Components

Storage controllers, independent of the model, are considered equivalent in the cluster configuration, in that they are all presented and managed as cluster nodes. Clustered Data ONTAP is a symmetrical architecture, with all nodes performing the same data-serving function.

Individual disks are managed by defining them into aggregates: groups of disks of a particular type that are protected by using NetApp RAID-DP®, similar to 7G and 7-Mode.

NICs and HBAs provide physical ports (Ethernet and FC) for connection to the management and data networks shown in Figure 3.

The physical components are visible to cluster administrators but not directly to the applications and hosts that use the cluster. The physical components provide a pool of shared resources from which the logical cluster resources are constructed. Applications and hosts access data only through storage virtual machines (SVMs, formerly known as Vservers), which contain volumes and logical interfaces.

Logical Cluster Components

The primary logical component of a cluster is the SVM; all client and host data access is through an SVM. Clustered Data ONTAP supports a minimum of one and up to hundreds of SVMs in a single cluster. Each SVM is configured for the client and host access protocols it supports, in any combination of SAN and NAS. Each SVM contains at least one volume and at least one logical interface. The administration of each SVM can optionally be delegated so that separate administrators are responsible for provisioning volumes and other SVM-specific operations. This is particularly appropriate for multi-tenant environments or when workload separation is desired. An SVM-delegated administrator would have visibility only to that administrator's specific SVM and would have no knowledge of any other hosted SVM.

For NAS clients, the volumes in each SVM are junctioned together into a namespace for CIFS and NFS access. For SAN hosts, LUNs are defined within volumes and mapped to hosts. A special type of SVM, known as Infinite Volume, is described later in this section.

The accessing hosts and clients connect to the SVM through a logical interface (LIF). LIFs present either an IP address (used by NAS clients and iSCSI hosts) or a World Wide Port Name (WWPN, for FC and FCoE access). Each LIF has a home port on a NIC or an HBA. LIFs are used to virtualize the NIC and HBA ports rather than mapping IP addresses or WWPNs directly to the physical ports, because there are almost always many more LIFs than physical ports in a cluster. Each SVM requires its own dedicated set of LIFs, and up to 128 LIFs can be defined on any cluster node. A LIF defined for NAS access can be temporarily migrated to another port on the same or a different controller to preserve availability, to rebalance client performance, or to evacuate all resources on a controller for hardware lifecycle operations.

Figure 3 shows a single SVM in a two-node cluster providing data services to SAN hosts and NAS clients. Each volume, shown by the orange circles, is provisioned on an aggregate on a cluster node, and the combination of all of the volumes constitutes the entire namespace or resource pool for LUNs. Volumes can be moved nondisruptively at any time from any aggregate to any aggregate as required. Delegated SVM administrators can provision volumes only in their own SVMs; these administrators have no visibility to any other SVM, or even awareness that other SVMs exist. The delegated SVM administrator cannot perform volume moves around the cluster, because this operation affects the capacity of aggregates shared by other SVMs. For this reason, only a cluster administrator can move volumes.

If SVM administration has been delegated, the cluster administrator must explicitly specify the aggregates available to the SVM administrator for provisioning volumes. This offers a mechanism for SVMs to provide different classes of service; for example, an SVM could be restricted to using only aggregates with SSD or SATA drives or only aggregates on a particular subset of controllers.

Figure 3) Cluster with a single SVM.

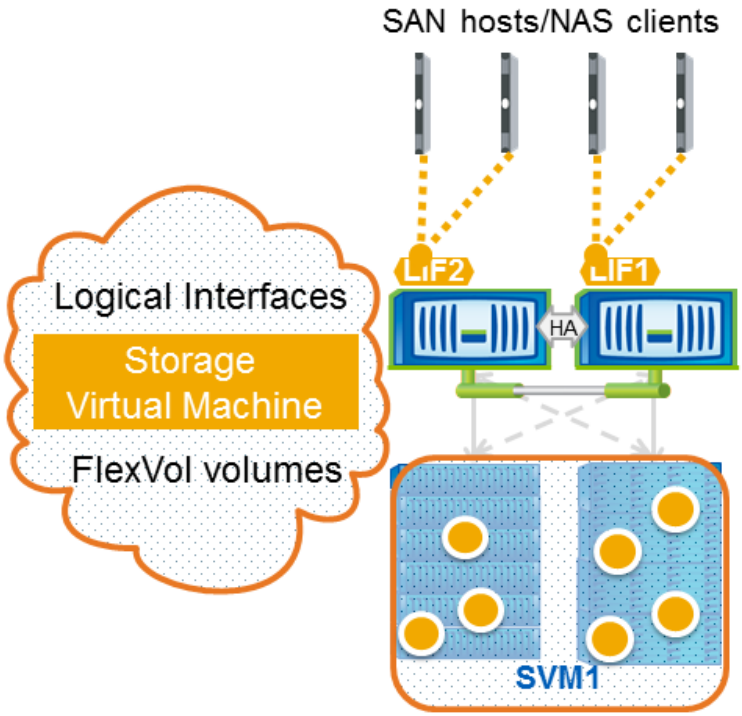

Figure 3 shows a more complex environment. The cluster here consists of four nodes, with two SVMs providing data access. Each SVM consists of different volumes and LIFs, for secure compartmentalized access. Although the volumes and LIFs in each SVM share the same physical resources (network ports and storage aggregates), a host or client can access the data in SVM1 only through a LIF defined in that SVM, and similarly for SVM2. Administrative controls make sure that a delegated administrator with access to SVM1 has visibility only to the logical resources assigned to that SVM, and an SVM2-delegated administrator similarly sees only the SVM2 resources.

Figure 4) Cluster with multiple SVMs.

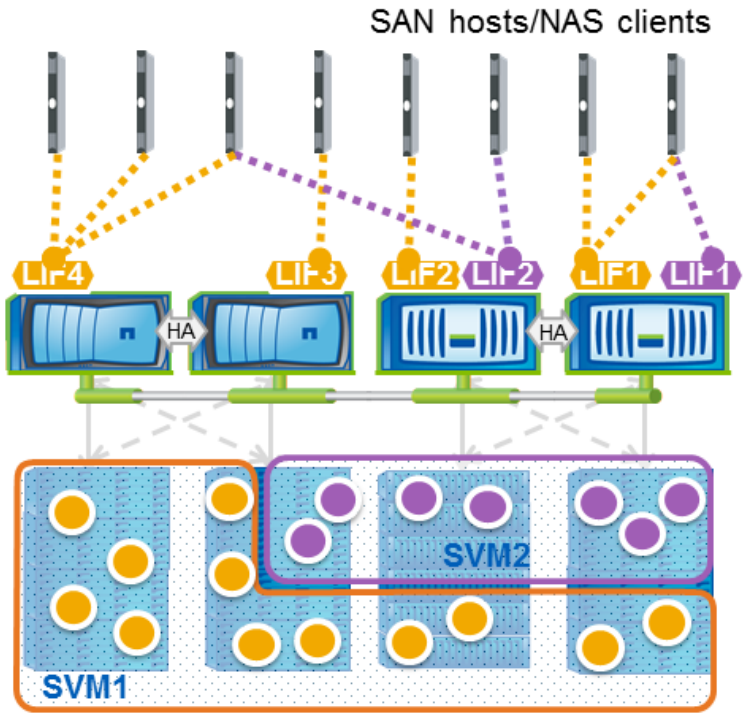

By virtualizing physical resources into the virtual server construct, Data ONTAP implements multi-tenancy and scale-out and allows a cluster to host many independent workloads and applications.

Storage QoS

Clustered Data ONTAP 8.2 provides storage QoS (quality of service) policies on cluster objects. An entire SVM, or a group of volumes or LUNs within an SVM, can be dynamically assigned to a policy group, which specifies a throughput limit, defined in terms of IOPS or MB/sec. This can be used proactively or reactively to throttle rogue workloads and prevent them from affecting the rest of the workloads. QoS policy groups can also be used by service providers to prevent tenants from affecting each other, as well as to avoid performance degradation of the existing tenants when a new tenant is deployed on the shared infrastructure.

3 Storage Setup Considerations

3.1 Storage Virtual Machine Configuration

A storage virtual machine (SVM) is a logical component of the storage cluster. The administration of an SVM can be delegated to separate administrators. In multi-tenancy environments, one or multiple SVMs are typically assigned to each tenant to allow each tenant to operate an own environment separated from those of other tenants.

Multiple SAP landscapes can use a single SVM, or an SVM could be assigned to each SAP landscape if they are managed by different teams within a company. When SnapManager® for SAP is used to create system copies, the source system (production) and the target system (QA, test) have to run within the same SVM.

3.2 Volume Layout and Logical Interface Configuration

Clustered Data ONTAP enables migration of SAP systems nondisruptively to any other storage node within the storage cluster. Migrations are, for example, used to rebalance the I/O load to new storage nodes that have been added to the cluster. Systems can also be migrated to another storage node if storage hardware is renewed and the old storage node is removed from the cluster.

The volume and logical interface (LIF) configuration with clustered Data ONTAP has a direct impact on the effective data path when a SAP system is migrated to a different storage node.

3.3 Multiple SAP Systems Sharing One Logical Interface

A configuration with the smallest number of LIFs would be one LIF per storage node in the cluster. With such a configuration, multiple SAP systems share a common LIF. Figure 5 shows a configuration with one LIF per storage node. In the example there are three SAP systems running on one storage node.

Figure 5) Configuration with one LIF per storage node.

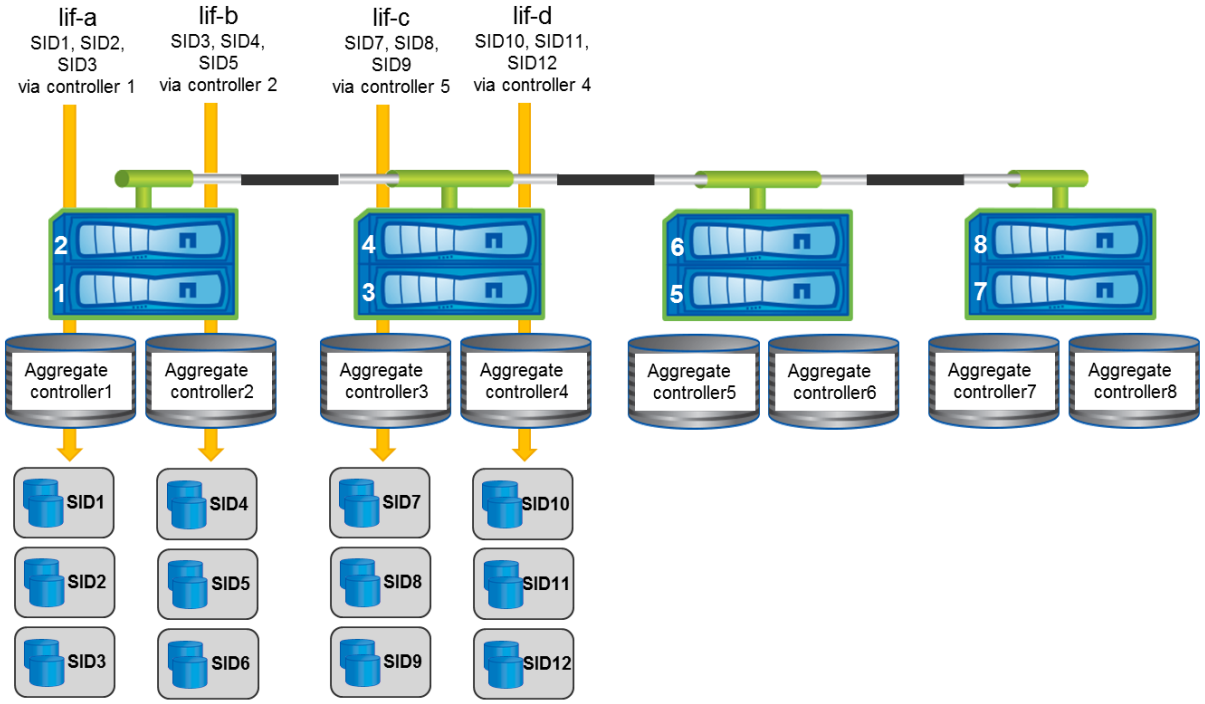

All SAP systems that share a common LIF can be migrated as one entity. The LIF is also migrated to the new storage node as shown in Figure 6. In this case no cluster interconnect traffic occurs.

Figure 6) Migration of multiple systems.

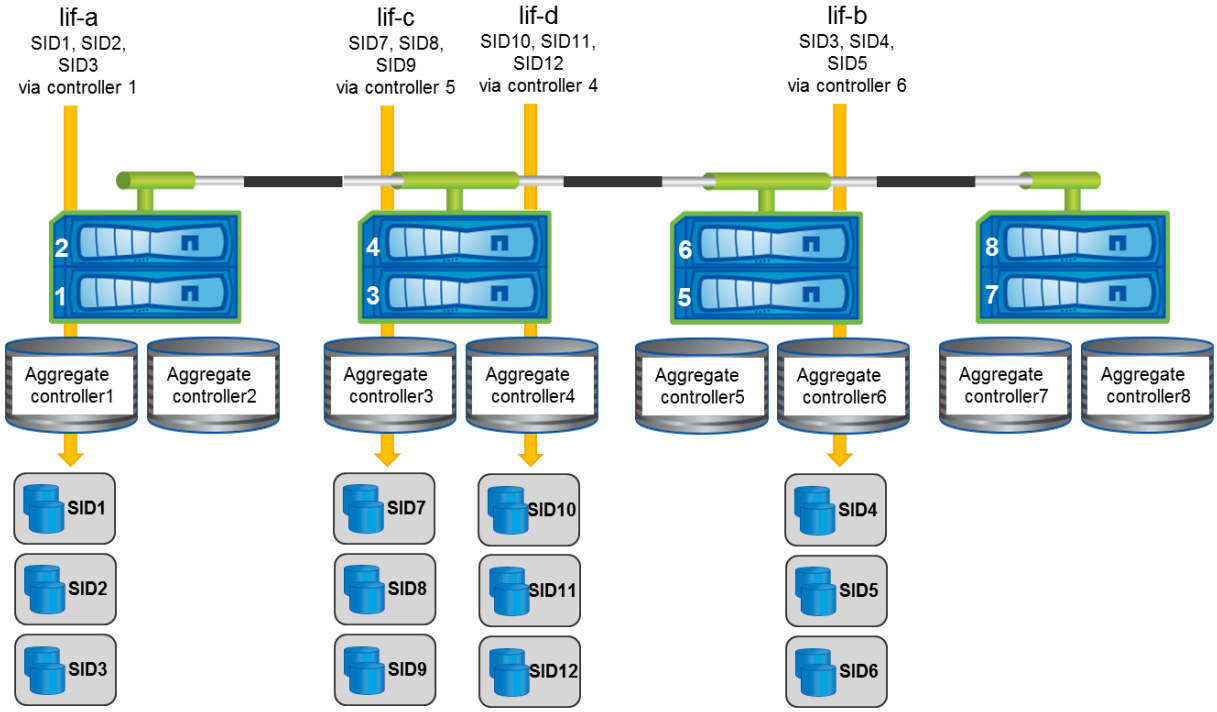

If a single SAP system is migrated, for example, SID4 as shown in Figure 7, the LIF can't be migrated together with the SAP system, because other SAP systems are still running on the original storage node. Therefore the SAP system SID4 will still be accessed through lif-a on storage node 1, and the data will be routed through the cluster interconnect. For SAP systems with low throughput requirements, for example, development and test systems, cluster interconnect traffic is typically acceptable. For production systems with high throughput requirements, this type of design should be avoided.

Figure 7) Migration of a single system.

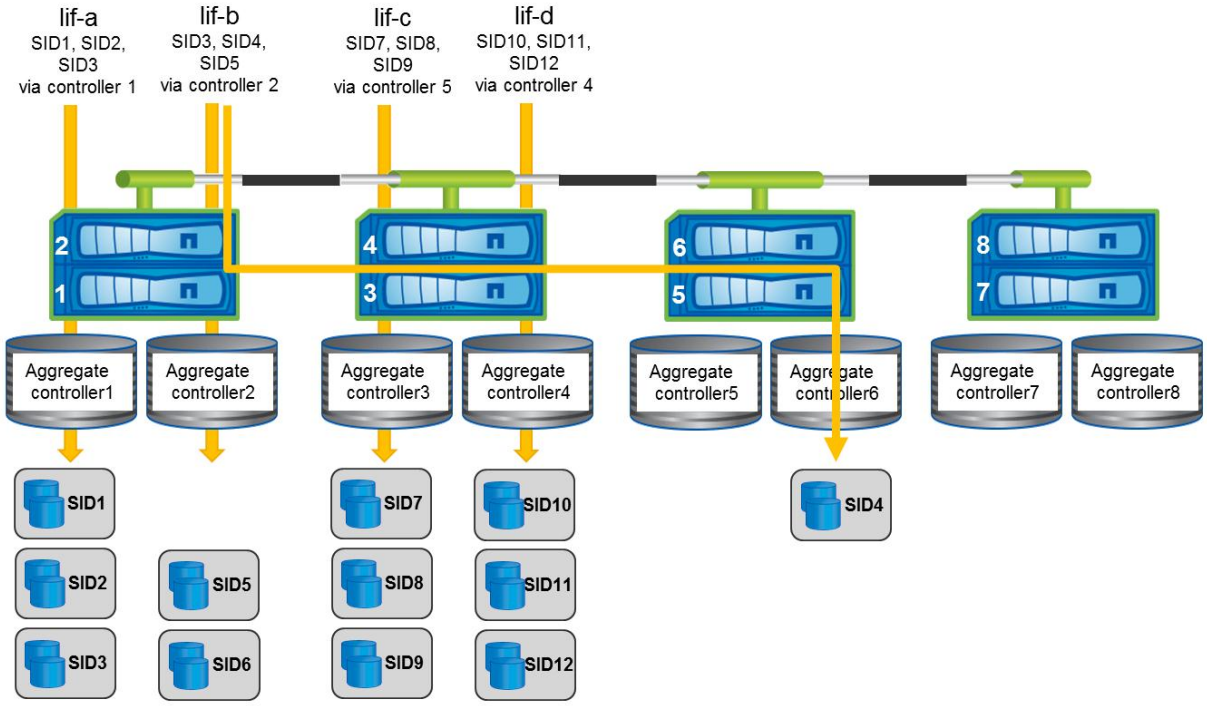

To avoid the interconnect traffic, a new LIF can be configured at the new storage node, the SAP system SID4 would need to be stopped, and the file systems would need to be remounted using the new LIF.

Sharing LIFs among multiple SAP systems can make sense for development and test systems that don't require scalability and nondisruptive operation, and cluster interconnect traffic is acceptable due to low throughput requirements.

3.4 One or Multiple Logical Interfaces per SAP System

One or multiple LIFs per SAP system are required for production systems that demand nondisruptive operations and scalability.

Figure 8 shows a configuration with one LIF per SAP system.

Figure 8) Configuration with one LIF per SAP system.

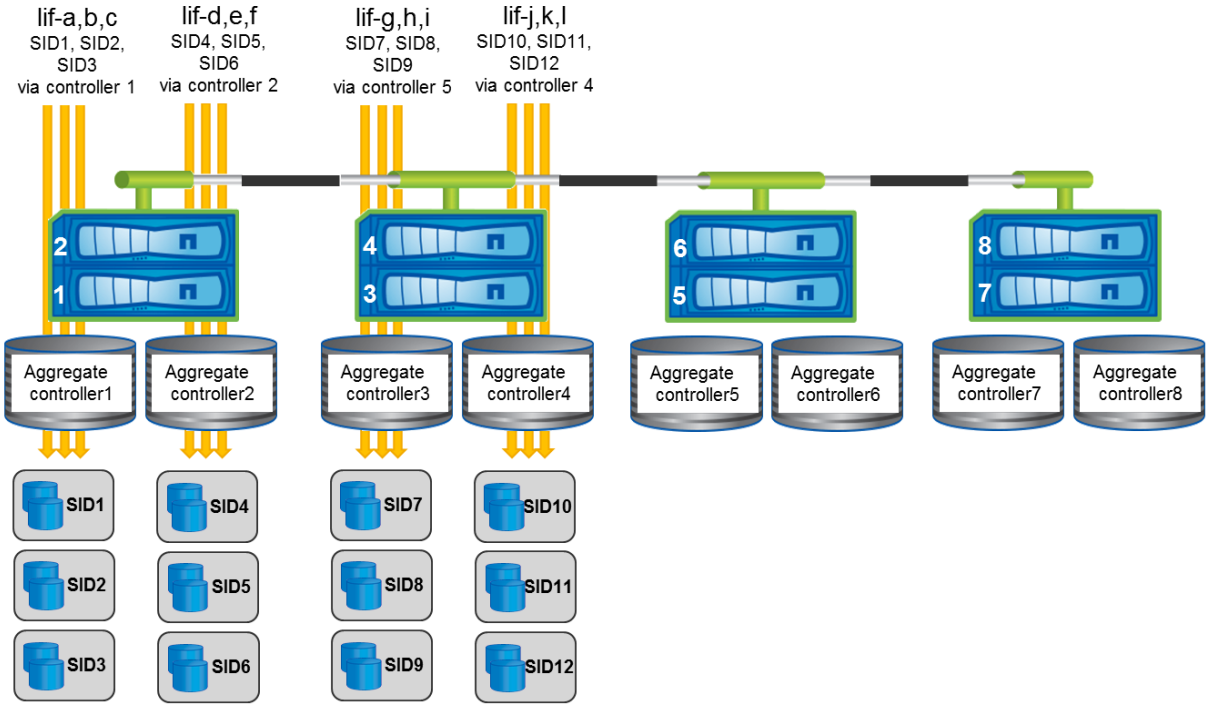

With this configuration, each single SAP system can be migrated nondisruptively to another storage node together with its own LIF.

Figure 9) Migration of multiple single SAP systems.

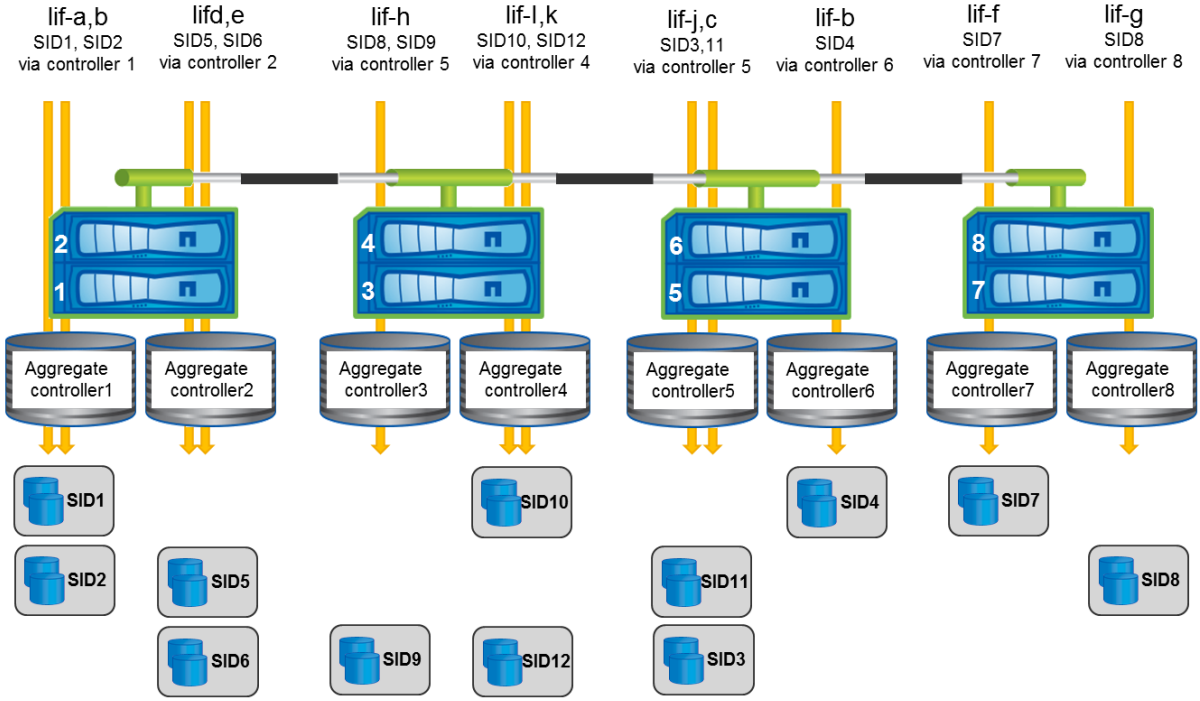

3.5 Logical Interface Configuration Summary

The highest flexibility in regard to migration of SAP systems within the storage cluster is achieved when each SAP system is associated with its own LIF or even multiple LIFs. Multiple LIFs are configured to allow nondisruptive migration for specific volumes of an SAP system. Because each LIF needs its own IP address, it might not be possible to configure a LIF for each volume, due to the amount of needed IP addresses.

Table 1 summarizes the SAP system characteristics and the configuration options.

Table 1) LIF configuration summary.

Characteristics / Configuration Options	Development and Test Systems	Small Production Systems	Large Production Systems
SAP System Characteristics			
Scalability and nondisruptive operation required	No Not business critical	Yes Business critical	Yes Business critical
Cluster interconnect traffic acceptable	Yes Typical low throughput requirements	Yes/no Depending on throughput requirements	No High throughput requirements
Downtime acceptable to reconfigure mounts and LIFs	Yes	No	No
Configuration Options			
Multiple SAP systems sharing one LIF	Reasonable configuration	No, doesn't comply with requirements	No, doesn't comply with requirements
One or multiple LIFs per SAP system	Reasonable configuration	Required to comply with requirements	Required to comply with requirements

4 Volume and Logical Interface Configuration Examples

No single configuration fits all customer environments. The decision regarding which configuration makes most sense must be based on customer-specific requirements. The following section provides configuration examples that should help to get a basic understanding of how configurations could look.

The following configuration examples are based on a setup with at least one LIF per SAP system. Multiple LIFs are required depending on the amount of volumes and how these volumes should be distributed to aggregates and storage nodes:

1. Minimum configuration with a single LIF and two volumes. Oracle data files in one volume; Oracle log files and SAP and Oracle binaries in a second volume.
2. Three volumes, two LIFs configuration to store Oracle mirrored redo logs on a different hardware than the online redo logs.
3. Five volumes, five LIFs configuration to distribute sapdata file systems to multiple storage nodes based on performance requirements.

Configuration examples 1 and 2 are reasonable configurations for development and test as well as production systems. Configuration example 3 can be used for larger production systems that require scalability and the ability to distribute load among multiple storage nodes.

Figure 10 shows configuration examples with different numbers of volumes and LIFs per SAP system.

Figure 10) Volume and LIF configuration examples.

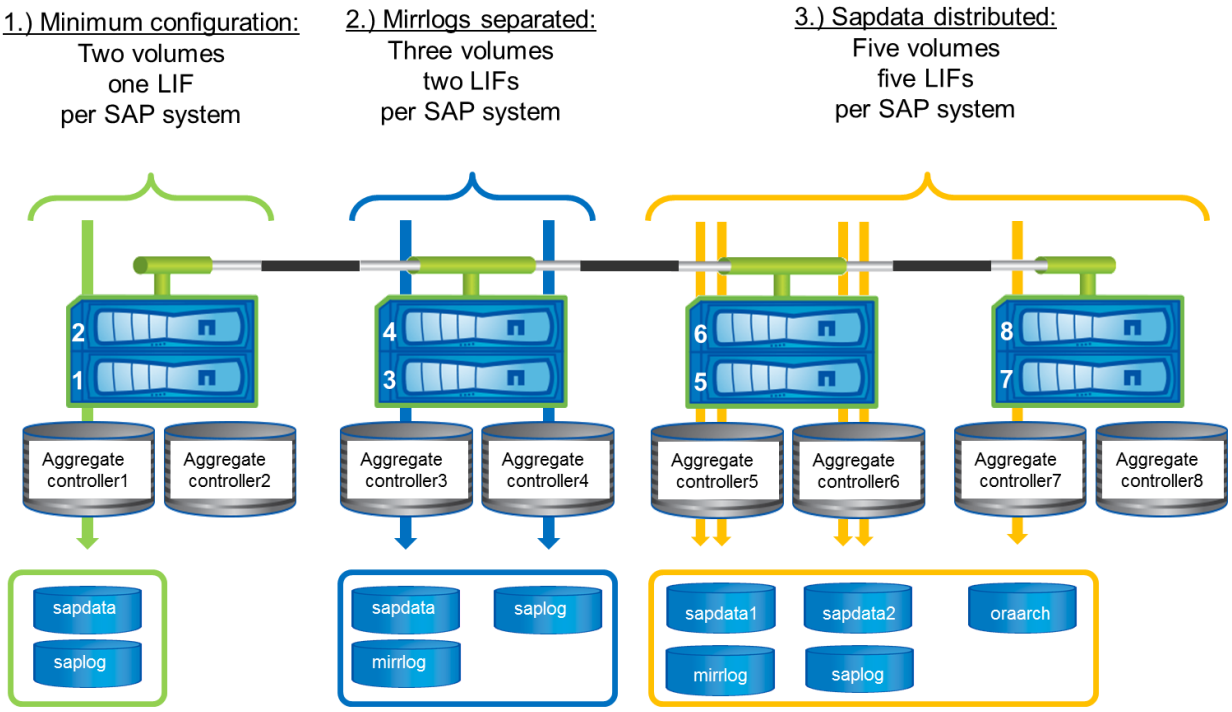

4.1 Minimum Configuration

Configuration 1 in Figure 10 shows the minimum configuration with two volumes and one LIF per SAP system:

- 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

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.

Table 2) Minimum configuration, LIF and volume layout.

Aggregate on One Storage Node LIF 1	
FlexVol "sapdata"	**FlexVol "saplog"**
/oracle/SID/sapdata1	/oracle/SID/origlogA
/oracle/SID/sapdata2	/oracle/SID/origlogB
/oracle/SID/sapdata3	/oracle/SID/mirrlogA
/oracle/SID/sapdata4	/oracle/SID/mirrlogB
	/oracle/SID/oraarch
	/oracle/SID
	Oracle binaries
	SAP binaries

4.2 Separating Oracle Mirrored Redo Logs

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. This configuration is shown as configuration 2 in Figure 10. With this configuration, two LIFs per SAP system are required:

- 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

Table 3) LIF and volume layout for separated mirrored redo logs.

Aggregate on Storage Node 1 LIF 1		Aggregate on Storage Node 2 LIF 2
FlexVol "sapdata1"	**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

4.3 Large SAP Systems with High-Performance Requirements

SAP systems with very high throughput requirements should be distributed to multiple storage nodes. It can also be beneficial to distribute data from small or medium production systems across multiple 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 node.

With this configuration, shown as configuration 3 in Figure 10, five volumes and five LIFs are used:

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

Table 4) LIF and volume layout for large SAP systems.

Aggregate on Storage Node 1 LIF 1	Aggregate on Storage Node 2 LIF 2	Aggregate on Storage Node 2 LIF 3	Aggregate on Storage Node 1 LIF 4	Aggregate on Storage Node 3 LIF 5
FlexVol "sapdata1"	**FlexVol "sapdata2"**	**FlexVol "saplog"**	**FlexVol "mirrlog"**	**FlexVol "oraarch"**
/oracle/SID/ sapdata1	/oracle/SID/ sapdata3	/oracle/SID/ origlogA	/oracle/SID/ mirrlogA	/oracle/SID/ oraarch
/oracle/SID/ sapdata2	/oracle/SID/ sapdata4	/oracle/SID/ origlogB	/oracle/SID/ mirrlogB	
		/oracle/SID		
		Oracle binaries		
		SAP binaries		

5 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
- High-availability requirements

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 SAP Application Performance Standard (SAPS) 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.

6 SAP System Installation

This section describes the requirements and the configuration for installing a SAP Business Suite or SAP NetWeaver system with Oracle Database under UNIX® using the NFS protocol.

6.1 Storage Network

A dedicated, redundant 10 Gigabit Ethernet storage network is required to attach the servers to the storage nodes. This dedicated storage network should be used exclusively for the storage traffic and not for any other purposes. Each server will therefore need two 10 Gigabit Ethernet cards connected to the switching infrastructure.

Figure 11) Storage network.

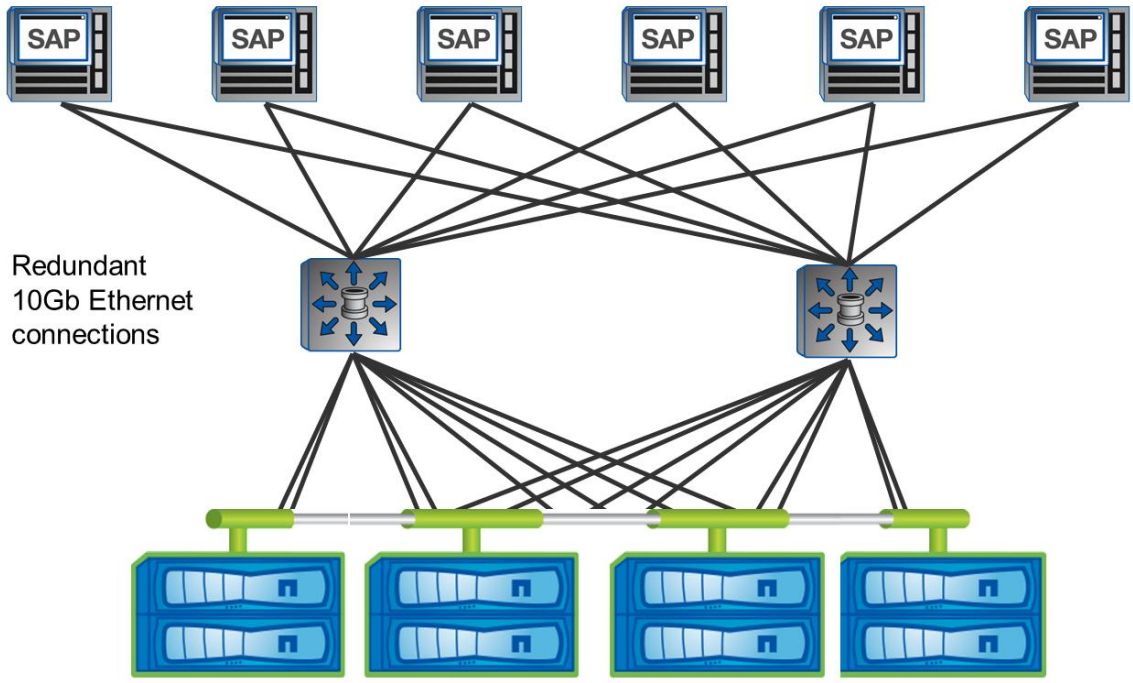

6.2 Operating System Configuration

Configuring the correct NFS mount options is important to provide optimal performance and system stability. There are common mount options that are valid for all operating system platforms.

Common mount options:

rw, bg, hard, vers=3, proto=tcp, timeo=600, rsize=65536, wsize=65536, nointr, suid

Additional mount options for the different operating system platforms:

Linux®: <common>

Solaris: <common>, llock

HP/UX: <common>, forcedirectio

AIX: <common>

Refer to NetApp Best Practices for Oracle for additional Information on operating system–specific tuning.

Note: It is recommended to switch off NFS server locking to avoid the need to manually delete NFS locks at the storage system in case of an Oracle database server crash. NFS server locking can be switched off with the “nolock” mount option for Linux or the “llock” mount option for AIX, HP/UX, and Solaris.

6.3 Snapshot Configuration

Snapshot backups on the storage level for database applications won’t be consistent from the database point of view unless the database is shut down or the Oracle Database is put in hot backup mode. Therefore, automatically scheduled Snapshot copies on the storage level should be turned off on database volumes.

Note: During the SAP installation, the visibility of the Snapshot directory has to be switched off for all volumes containing any of the file systems of the SAP system. Otherwise, SAP Software

Provisioning Manager (formerly SAPINST) will try to change permissions and ownership on Snapshot subdirectories. Because the Snapshot data is read-only, SAP Software Provisioning Manager will fail, and the installation will abort. After the installation of the SAP system, the volume option can be switched on again.

6.4 SAP Installation Process

The following description of the installation process assumes that the virtual storage machine is already available. The description is based on a two-volume, single LIF per SAP system configuration.

Table 5 shows an example of a volume configuration.

Table 5) Volumes and mount points.

FlexVol Volume	Qtree	Subdirectory to Be Mounted	Mount Point at SAP System
sapdata_SID	sapdata	sapdata1	/oracle/SID/sapdata1
		sapdata2	/oracle/SID/sapdata2
		sapdata3	/oracle/SID/sapdata3
		sapdata4	/oracle/SID/sapdata4
saplog_SID	saplog	oracle	/oracle
		sapusr_SID	/usr/sap/SID
		sapmnt_SID	/sapmnt/SID
		saptrans_SID	/usr/sap/trans
		saphome_SID	/home/SIDadm

The necessary file systems for the SAP installation are set up with the following steps:

1. Create a SID-specific LIF for the SAP system, for example, “nfs_sap_sid.”
2. Create the volumes, assign the appropriate export policy, and configure the junction path, for example, /sapdata_SID and /saplog_SID.

Note: If it is planned to use SMSAP with data protection, the junction path should be the same as the volume name. Otherwise, the postbackup script won’t work.

3. Create the following directories at the SAP host:
   /usr/sap/SID
   /sapmnt/SID
   /usr/sap/trans
   /oracle
   /home/sidadm
4. Edit the file system configuration file (for example, /etc/fstab for Linux) to mount the preceding file systems from the NetApp storage using the discussed mount options.

Note: The sapdata file systems will be mounted in a later step.

5. Mount the preceding file systems.
6. Create directories at the SAP host within the already mounted “/oracle” file system:
   /oracle/SID
   /oracle/SID/sapdata1
   /oracle/SID/sapdata2

/oracle/SID/sapdata3
/oracle/SID/sapdata4

7. Edit the file system configuration file /etc/fstab (Linux) and mount the sapdata file systems from the NetApp storage using the discussed mount options.

The following command output shows a file system configuration example for SID=PA0:

```
t002-lnx-60:~ # df
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/sda2             19125332  13477760   4676044  75% /
devtmpfs               4065348       104   4065244   1% /dev
tmpfs                  4065348         0   4065348   0% /dev/shm
nfs_sap_pa0:/saplog_PA0/saplog/oracle
                      52428800  14888128  37540672  29% /oracle
nfs_sap_pa0:/sapdata_PA0/sapdata/sapdata1
                     157286400  29363616 127922784  19% /oracle/PA0/sapdata1
nfs_sap_pa0:/sapdata_PA0/sapdata/sapdata2
                     157286400  29363616 127922784  19% /oracle/PA0/sapdata2
nfs_sap_pa0:/sapdata_PA0/sapdata/sapdata3
                     157286400  29363616 127922784  19% /oracle/PA0/sapdata3
nfs_sap_pa0:/sapdata_PA0/sapdata/sapdata4
                     157286400  29363616 127922784  19% /oracle/PA0/sapdata4
nfs_sap_pa0:/saplog_PA0/saplog/sapusr_PA0
                      52428800  14888128  37540672  29% /usr/sap/PA0
nfs_sap_pa0:/saplog_PA0/saplog/sapmnt_PA0
                      52428800  14888128  37540672  29% /sapmnt/PA0
nfs_sap_pa0:/saplog_PA0/saplog/saphome_PA0
                      52428800  14888128  37540672  29% /home/pa0adm
nfs_sap_pa0:/saplog_PA0/saplog/saptrans_PA0
                       5242880      6656   5236224   1% /usr/sap/trans
t002-lnx-60:~ #
```

The SAP installation tool SAP Software Provisioning Manager fully supports NFS mounts. The SAP installation can therefore be accomplished as described in the corresponding SAP Installation Guide.

Note: SAP Software Provisioning Manager will store one of the Oracle control files in /oracle/SID/sapdata1. If it is planned to use SnapManager for SAP, this control file has to be stored outside of the sapdata volume to allow fast volume restore with SMSAP. Replacing the Oracle control file in a different volume can be done either during the installation with SAP Software Provisioning Manager or when the installation is finished using Oracle SQL Server® commands.

7 Business Continuance

7.1 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 24/7.
- **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 12) Backup solution overview.

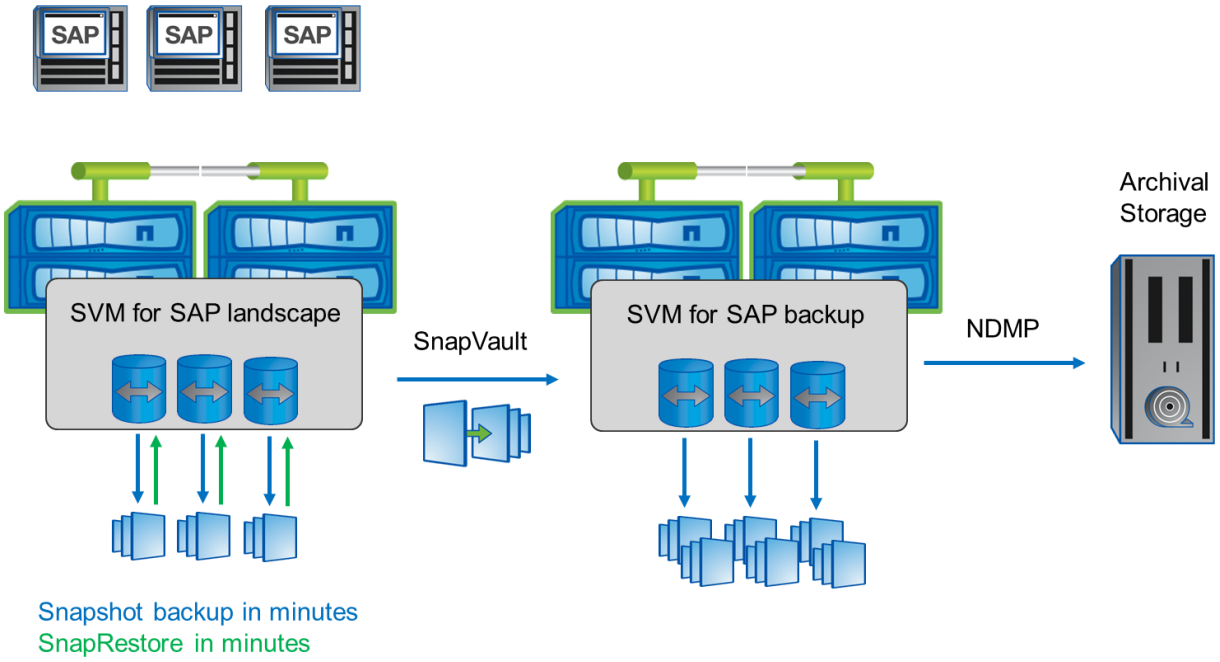

A backup and recovery solution using a NetApp storage system always consists of two parts:

- 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 SnapVault software.** The primary storage virtual machine communicates directly with the secondary storage virtual machine 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 virtual machine, and the data will be written to tape using NDMP.

- **Backup to tape using third-party backup software such as NDMP backup (serverless backup).** The tape is connected directly to the primary storage system. The data is written to tape using NDMP.

Figure 13 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 13) Comparison of time required for different backup methods.

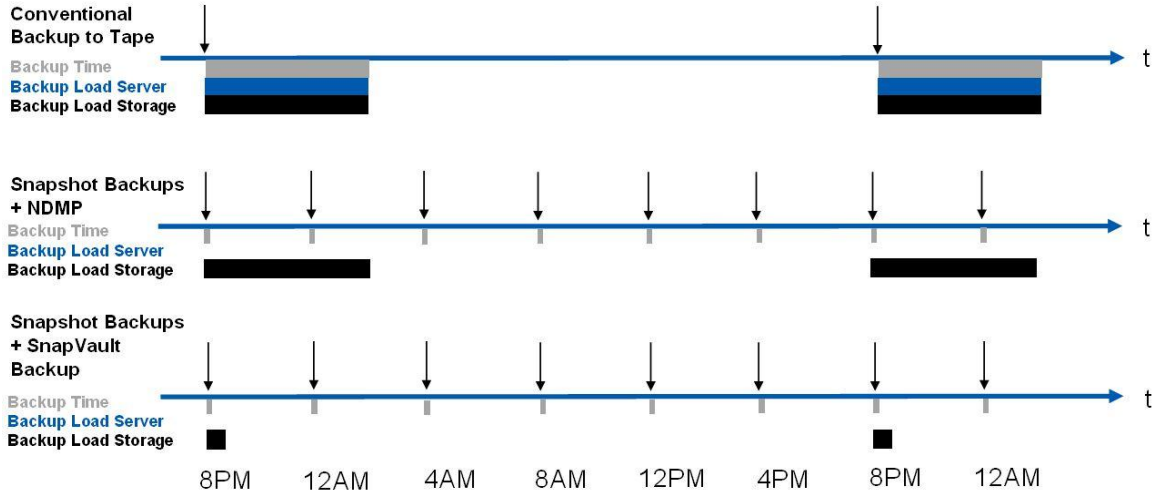

Snapshot Backups Together with NDMP Backups

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 system 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.

Figure 14 compares the time required to perform restore and recovery.

Figure 14) Comparison of time needed for restore and recovery.

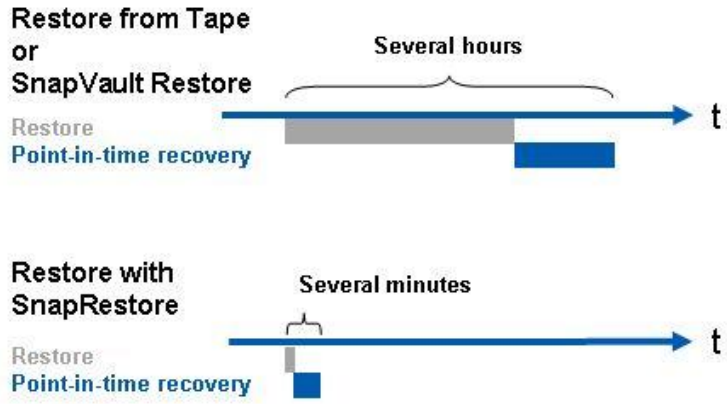

Restore from Tape or SnapVault Restore

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 secondary storage systems

Database Verification

Database verification is an important part of a backup concept. Snapshot backups are perfect for running database consistency checks. NetApp SnapManager software offers the possibility to run a database consistency check on a separate server automatically or manually after a backup without creating any load on the productive database system.

7.2 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 a 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 15) SAP repair system.

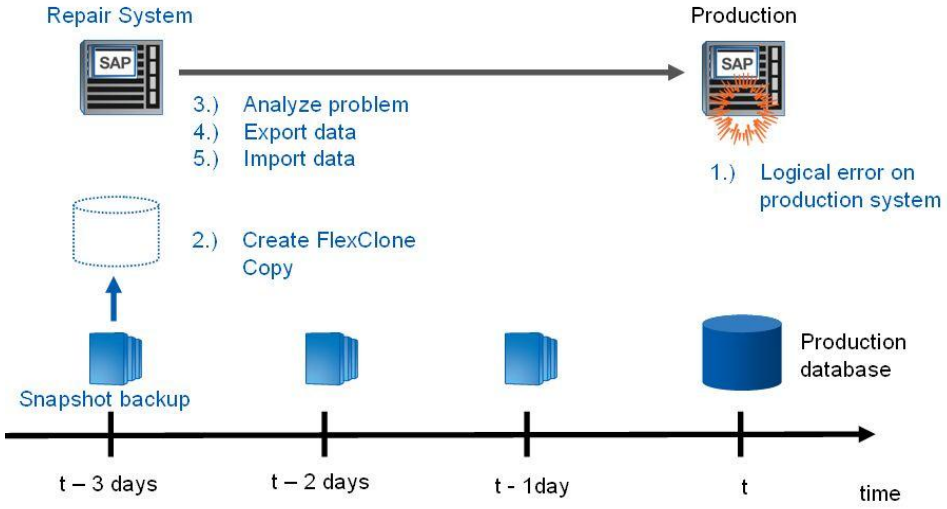

Figure 15 shows the general process of creating and using the repair system:

1. A logical error is discovered on the production system. Dependent on the kind of logical error, the decision can be made to shut down 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.

7.3 Disaster Recovery

Organizations recognize the importance of having a 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—make it mandatory to have a plan that makes sure of an absolute minimum of downtime and rapid recovery from a disaster.

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 technology mirrors data to one or more storage virtual machines. 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 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 log files and the log file backups are stored are mirrored using SnapMirror. The frequency of SnapMirror updates of the log files and log backups determines the amount of data lost in the event of a disaster.

Figure 16) Disaster recovery with SnapMirror.

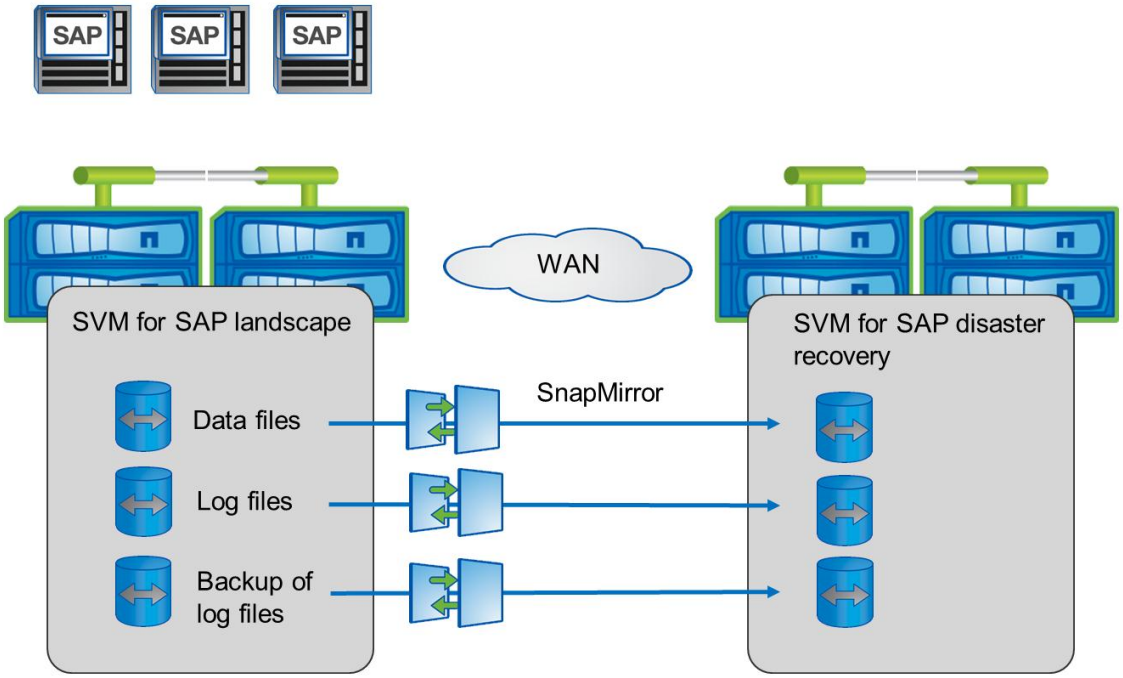

8 System Management and Maintenance

8.1 SAP System Copy

Business Challenges

A typical SAP customer environment today consists of different SAP Business Suite and SAP NetWeaver components. To be able to test application patches, run performance and data integrity tests, or provide user training environments, copies of SAP components are required. A typical SAP customer needs about 10 copies of different SAP components. These copies must be refreshed, often on a weekly or monthly basis.

Rapid and space-efficient provisioning of test systems would allow SAP customers to run more test or project systems and refresh the systems more often. This would enable project teams to reduce project cycles by running parallel testing and would improve quality of testing and training with more actual data from production.

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 restore the data to the three target systems.

Figure 17) Traditional SAP system copy.

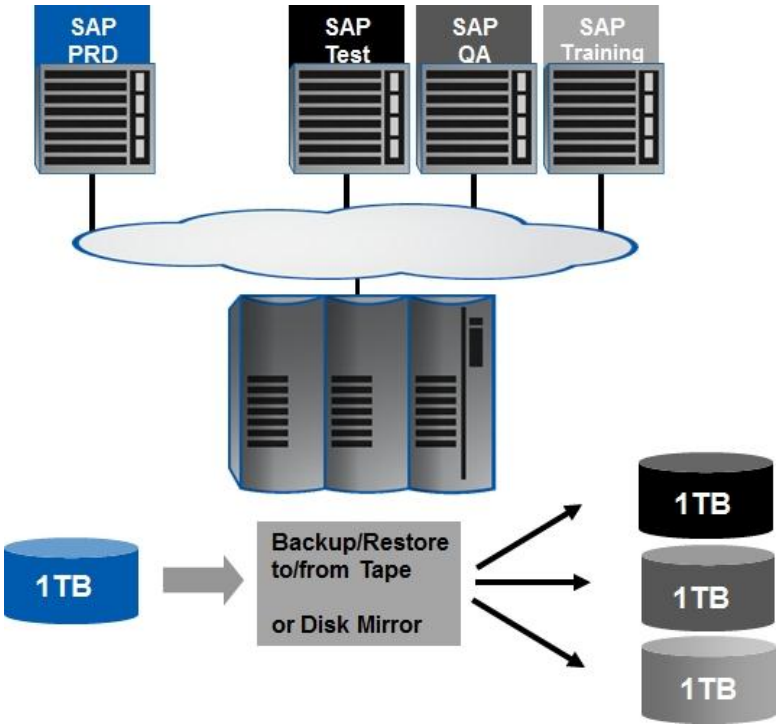

In contrast, when using NetApp FlexClone technology to create SAP system copies, only a fraction of the storage space is required. NetApp FlexClone technology uses Snapshot copies, which are created in a few seconds without interrupting the operation on the source system, to perform SAP system copies. Because the data is not copied but is 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 depend 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.

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 18) SAP system copy: NetApp approach.

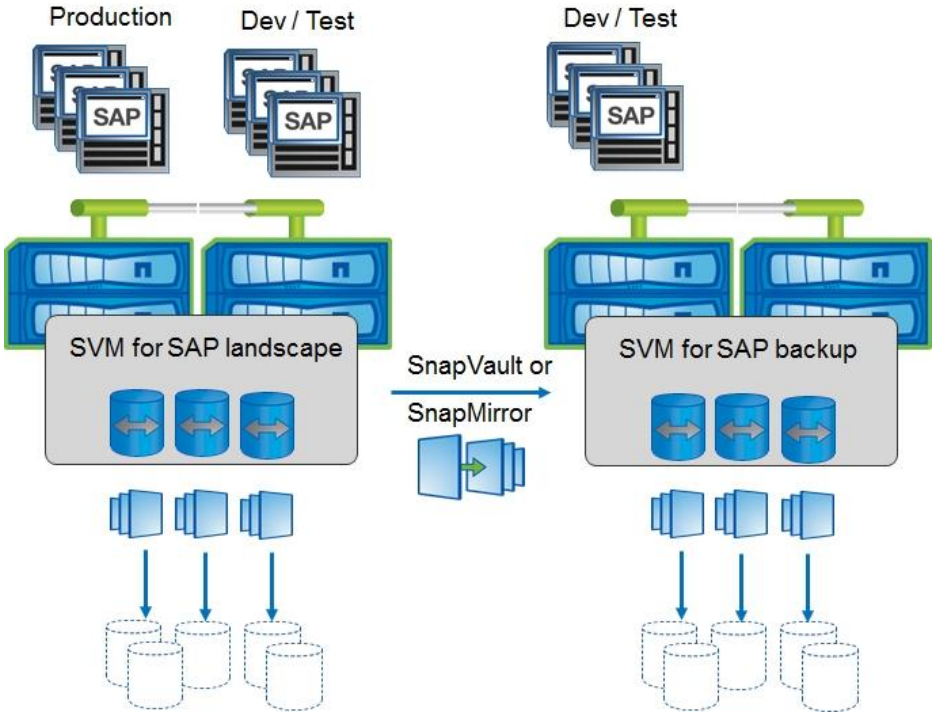

Time Requirements

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, whereas 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 19) SAP system copy: standard approach.

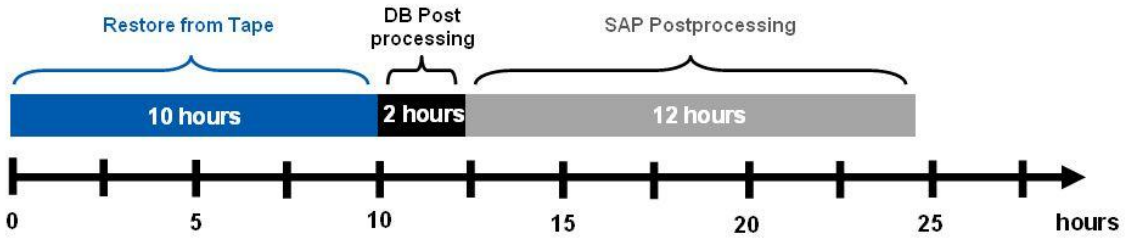

All steps up to the point when the SAP system can be started on the target host can be accomplished in a few minutes using the NetApp solution, compared to several hours with the standard approach. With both approaches, the SAP postprocessing needs to be done as an additional step.

Figure 20) SAP system copy: NetApp approach.

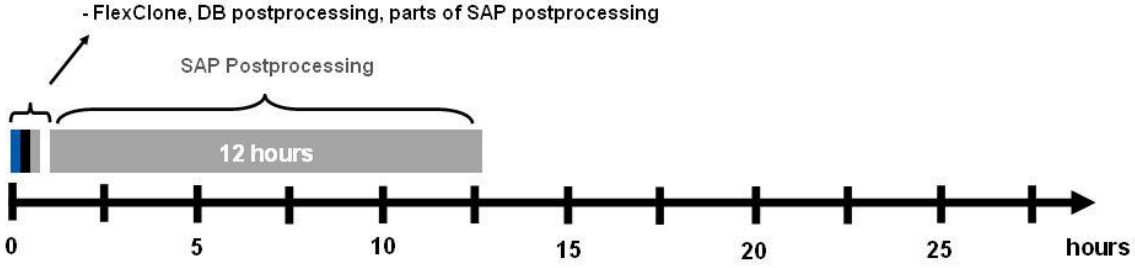

A key requirement to successfully managing an SAP environment is the ability to create copies of production data for use in testing, quality assurance, or training. NetApp Snapshot and FlexClone technologies allow a fast and space-efficient creation of SAP systems.

8.2 SAP Testing Cycle

The possibility of simply creating backups in seconds and being able to restore the SAP system to a point in time of any available Snapshot copy is also very helpful in SAP development and test environments. Projects such as data import, SAP upgrades, and installation of support packages can be accelerated using fast backup and restore functionalities. During these projects, backups can be done at specific phases, and the system can be easily and quickly reset to a starting point in order to be able to repeat that phase. Test runs can easily be repeated with different code or configurations to make sure that the results are valid.

Figure 21) SAP testing cycle.

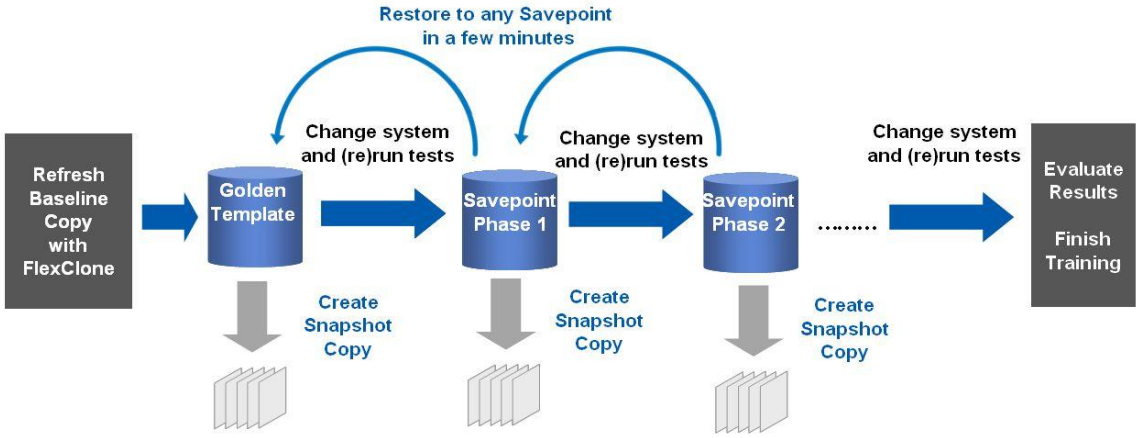

Carrying out a SAP upgrade or importing support packages and critical transports always involves SAP system downtime. It is important that this downtime be kept to a minimum and that the previous status can always be restored. The specified system changes are usually first made in the development system to test the general functionality and procedures. In many cases, test systems must be upgraded several times, because problems can occur that can only be solved by restoring the system and restarting the upgrade. In this respect, NetApp Snapshot copies and FlexClone functionality can save a considerable amount of time. A tape backup does not have to be made; a Snapshot copy can be created instead. In the event of an error, the system can be quickly restored to its original status, and the upgrade can be repeated.

Time management is extremely important when the production system is upgraded, because the system is not available at various stages during the upgrade. Scheduling must also include time for restoring 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 importing the transports for the modification adjustment, one normal weekend might not be sufficient for the upgrade. NetApp SnapManager software offers Snapshot as a backup method and SnapRestore for restoring the system to its former release status. This allows a higher level of flexibility with regard to scheduling. By creating several Snapshot copies at certain stages during the upgrade, it is possible to restart the upgrade without having to revert to the former release status.

9 SnapManager for SAP

The described example setup is based on SnapManager for SAP 3.3 and SnapDrive® for UNIX 5.2. It has been tested with SuSE SLES Linux.

9.1 Place Oracle Control Files

In order to use SMSAP for fast restore, the location of the control files (and its copies) is of eminent importance and must not be in the same volume as the SAP database files. If you use standard SAP installation procedures, SAP puts the control files in the /oracle/SID/origlogA, /oracle/SID/origlogB, and /oracle/SID/sapdata1 file systems.

The control file in the sapdata1 file system conflicts with the SnapManager requirements for separating the control files and data files into separate volumes and must be adjusted to allow for fast restore capability.

In the case of a new SAP install, you can adjust the location of the control files using SAP Software Provisioning Manager during the SAP installation process and move the control file normally placed in the sapdata1 file system (that is, the data volume) to the LOG volume. However, in the case of an SAP system that has already been installed, you must move the control file out of that file system to allow for fast restores using SnapManager.

9.2 SnapDrive for UNIX Configuration

Figure 22 shows the used example setup. The SVM is configured with two LIFs: one management LIF and one SAP system-specific LIF.

Figure 22) Example setup.

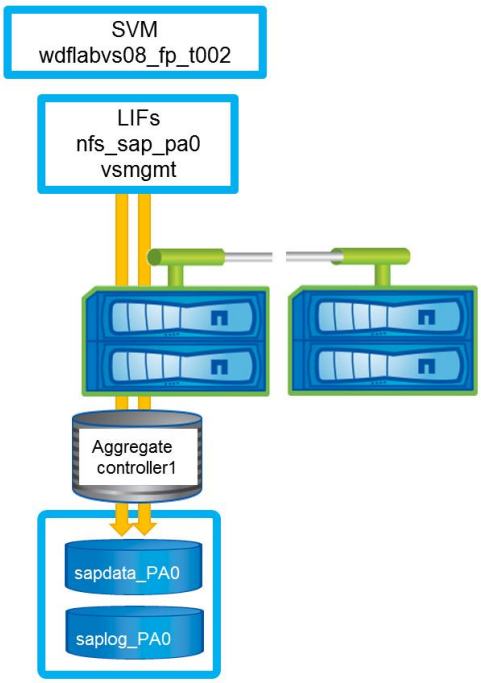

The SnapDrive for UNIX (SDU) credentials have to be set for the management interface of the SVM:

```
t002-lnx-60: # snapdrive config set vsadmin vsmmgmt
Password for vsadmin:
```

```
Retype password:
```

```
t002-lnx-60: # snapdrive config list
username    appliance name   appliance type
----------------------------------------------
vsadmin     192.168.102.5    StorageSystem
```

In addition, the management path for the LIF specific to SAP has to be configured:

```
t002-lnx-60:/mnt/software/scripts # snapdrive config set -mgmtpath vsmgmt nfs_sap_pa0
```

```
t002-lnx-60:/mnt/software/scripts # snapdrive config list -mgmtpath
system name          management interface   datapath interface
-------------------------------------------------------------
wdflabvs08_fp_t002   192.168.102.5          192.168.102.100
```

Note: SDU 5.2 requires that the volume attribute “Snapshot Directory Access Enabled” is set to true. Otherwise all Snapshot copies, even if they have been created by SDU, will be shown as “non SnapDrive Snapshot copies” and can’t be managed by SDU and SMSAP.

Note: SDU 5.2 does not support load-sharing mirrors of the root volume. SDU will fail to connect (mount) a FlexClone volume in a load-sharing mirror setup of the root volume. This will prevent SMSAP from running different tasks such as “mount a backup,” “restore a backup,” or “clone a backup.”

9.3 Overview Configuration Examples

SnapManager for SAP (SMSAP) provides the following functionalities:

- Backups based on Snapshot copies (local backup)
- Backups restored on storage level using SnapRestore (from local backup)
- Database verification with Oracle dbv
- SAP system copy based on a clone from a local backup

Note: SMSAP doesn't support data protection with clustered Data ONTAP. Replicating backups to a second location can only be accomplished with a postbackup script. The restore process from secondary can't be automated with SMSAP and is a manual process. Also, SAP system copies can't be done from secondary storage using SMSAP.

SMSAP supports two different methods to manage backup, restore, and recovery of SAP systems:

- Using the SAP BR*TOOLS together with BACKINT
- SMSAP GUI or CLI

SMSAP can create SAP system copies (clones) based on backups that have been created by the BR*TOOLS or SMSAP GUI or CLI backups.

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.

Table 6 shows which features are supported with the two different methods.

Table 6) SMSAP features.

Used interface	Local Backup Restore, Recovery from Local Backup	Cloning from Local Backup	Archive Log Management	Data Protection with Postbackup Script
BR*TOOLS	Yes	Yes	Yes - Based on Snapshot copies - Directly to a file system	No
SMSAP GUI SMSAP CLI	Yes	Yes	Yes - Based on Snapshot copies	Yes

Table 6 shows that data protection can only be done by using SMSAP backups and not with the BR*TOOLS. If the BR*TOOLS should be used for backup, a second schedule should be defined with SMSAP backups that are replicated to a second location using a postbackup script.

Archive log management can be done with either SMSAP or BRARCHIVE. BRARCHIVE with BACKINT or SMSAP archive log management is based on Snapshot backups. In both cases, Snapshot backups in the archive log volume need disk space based on the retention policy, even when archive logs get deleted by BRARCHIVE or SMSAP.

If BRARCHIVE is used, 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.

BRARCHIVE should be configured with "backup_dev_type = disk." The "archive_copy_dir" should be a mount point from a secondary storage system.

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. Because 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.

9.4 BR*TOOLS Configuration Example

The BR*TOOLS together with SMSAP BACKINT are used for backup, restore, and recovery of local backups. Data protection needs to be done by creating additional backups with 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 mount point at a secondary storage system.

Table 7 shows the different tasks.

Table 7) Tasks with BR*TOOLS scenario.

Task	Used interface
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 SMSAP BACKINT
Data protection (remote backup)	SMSAP GUI or CLI; additional daily backups with postbackup scripts
Restore from remote backup	Manual process
Cloning from local backup	SMSAP GUI or CLI
Cloning from remote backup	Not supported by SMSAP

SMSAP Configuration

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>

BR*TOOLS Configuration

BRBACKUP backups are configured with the init<SID>.sap parameter and init<SID>.utl files, as shown in Table 8.

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 8) BR*TOOLS configuration.

BR*TOOL	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 = <name_of_profile> fast = override retain = hourly protect = no
BRARCHIVE	**Init<SID>_BRARCHIVE.sap** backup_dev_type = disk archive_copy_dir = /mnt/backup2disk	N/A

9.5 Data Protection Configuration

Figure 23 shows an example setup. The further description is based on the example shown in this figure. The integration of data protection with SnapMirror is based on the concept of SnapMirror labels. The SnapMirror policy is configured to replicate all Snapshot copies with the labels SMSAP-HOURLY, SMSAP-DAILY, and so on. The defined SnapMirror schedule defines how often the backup SVM checks if there are new Snapshot copies to replicate.

Note: The restore process from secondary can't be automated with SMSAP and is a manual process. Also, SAP system copies can't be done from secondary storage using SMSAP.

The SMSAP postbackup script sets the SnapMirror label for all Snapshot copies that have been created with the backup. In order to be able to set the SnapMirror labels, the script must be able to execute commands at the source SVM. Therefore, public key authentication has to be configured.

Note: The script has been tested on Linux and uses /proc/mounts to get the junction path. If the junction path is different than the volume name, the script won't work. The script also needs to be adapted for other UNIX environments than Linux.

Figure 23) Data protection configuration.

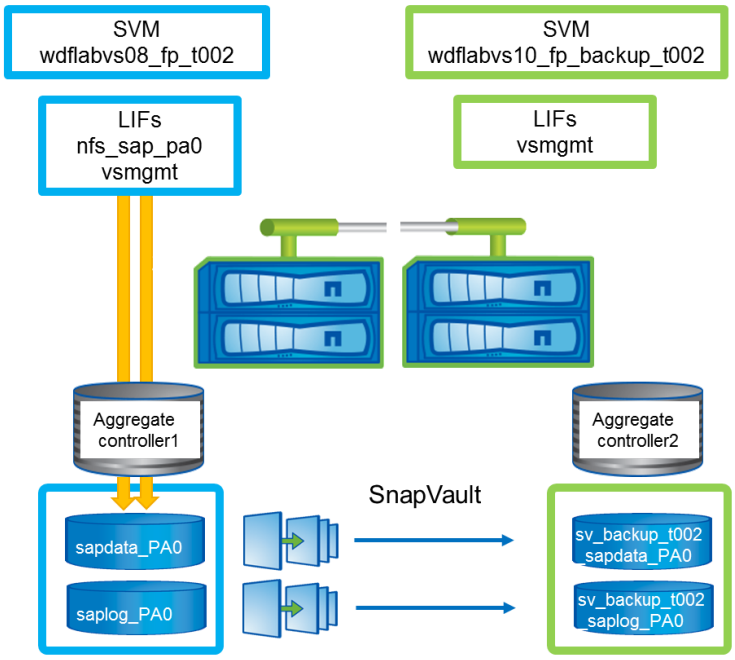

- Manual backup volume creation
- SnapMirror policy
  - Configured to use SnapMirror label SMSAP-HOURLY, SMSAP-DAILY, etc with specific retention policies
- Manual SnapMirror relation setup and initialization
- SMSAP Post Backup Script
  - Sets SnapMirror label of all Snapshots included in the backup to “SMSAP-HOURLY”, “SMSAP-DAILY”, etc.
- Backup SVM replicates all Snapshots with specific label and keeps them based on defined retention

Data protection is configured with the following steps:

1. Create SVM for backup volumes.
2. Create a peer relation between source and backup SVM.
3. Create backup volumes at backup SVM.
4. Create SnapMirror policy and add rules at backup SVM. With this example, all backups with the SnapMirror label “SMSAP_DAILY” will be replicated and kept for 20 days at the backup SVM:

   ```
   snapmirror policy create -policy SMSAP-Vault
   snapmirror policy add-rule -policy SMSAP-Vault -snapmirror-label SMSAP-Daily -keep 20
   ```

5. Create and initialize SnapMirror relation for the data and log volume at backup SVM. With this example, the predefined schedule “5min” is used. With this configuration, the backup SVM will check every 5 minutes if there are new Snapshot copies to be replicated:

   ```
   snapmirror create -source-path wdflabvs08_fp_t002:t002_sapdata_PA0 -destination-path wdflabvs10_fp_backup_t002:sv_backup_t002_sapdata_PA0 -type XDP -schedule 5min -policy SMSAP-Vault
   snapmirror initialize -destination-path wdflabvs10_fp_backup_t002:sv_backup_t002_sapdata_PA0
   snapmirror create -source-path wdflabvs08_fp_t002:t002_saplog_PA0 -destination-path wdflabvs10_fp_backup_t002:sv_backup_t002_saplog_PA0 -type XDP -schedule 5min -policy SMSAP-Vault
   snapmirror initialize -destination-path wdflabvs10_fp_backup_t002:sv_backup_t002_saplog_PA0
   ```

6. Configure SMSAP postbackup script.
7. Run backup or define schedule within SMSAP, including the postbackup script.

Postbackup Script Configuration

Copy the contents of the SMSAP postbackup script from the appendix in this document to your SAP system. The file must be stored as “Set-SM-Label.sh” in the directory “/opt/NetApp/smsap/plugins/backup/create/post.”

Within the SMSAP GUI, the posttasks script “Set SnapMirror label” needs to be selected. The script requires the name of the SVM as a parameter.

The postbackup task should be saved to a file in order to simplify further backup processes. An example xml file can be found in the appendix of this document.

Figure 24) Data protection configuration within SMSAP GUI.

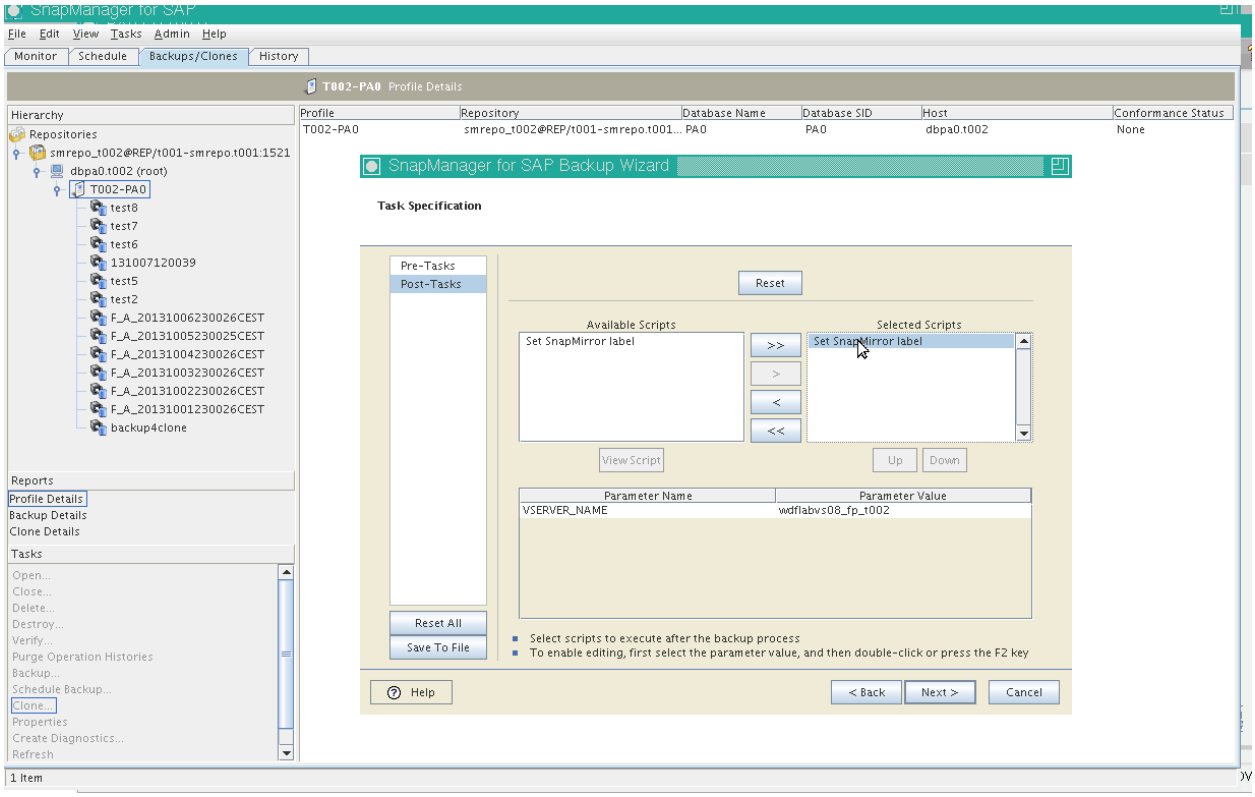

Depending on the selected retention policy of the backup, the postbackup script will set the SnapMirror label with the extension HOURLY, DAILY, or WEEKLY.

9.6 Scheduling Backup Jobs Within SAP CCMS and SMSAP Scheduler

With the SAP transaction “dbacockpit” or “db13,” all necessary backup and housekeeping jobs can now be scheduled. Backups with data protection can be scheduled with SMSAP.

Table 9 shows an example of a backup schedule.

Table 9) Backup schedule.

Task	Command	Retention Policy	Backup Replicated to Secondary
Local database backup	BRBACKUP with BACKINT	Retention class: HOURLY Run every 4 hours Keep 36 backups (3 days)	No
Backup with data protection	SMSAP GUI or CLI with postbackup scripts	Retention class: DAILY Run once per day Keep 42 backups (6 weeks)	Yes
Archive log backup	BRARCHIVE with “backup_dev_type = disk”	BRARCHIVE –s (save) every hour BRARCHIVE –ds (delete saved) every six hours	Yes, directly written to remote mount point by BRARCHIVE
Archive log management at secondary storage	BRCONNECT –f cleanup	Defined in init>SID>.sap	

.

9.7 Database Verification

Table 10 gives an overview of the different options to run the database verification.

Table 10) Database verification.

Used Interface	Process	Advantages/Disadvantages
BRBACKUP –w use_dbv	• Back up 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	• Back up 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

9.8 SAP System Copy

The documentation describes the setup of a new SAP system (ABAB-stack only) and the refresh of the new SAP system based on data from a source SAP system.

Note: SAP system copies can only be done at the primary storage using SMSAP. A system copy using a protected backup at a secondary storage can't be automated with SMSAP and would be a manual process.

Figure 25 shows the basic setup of the example environment.

Figure 25) SAP system copy setup.

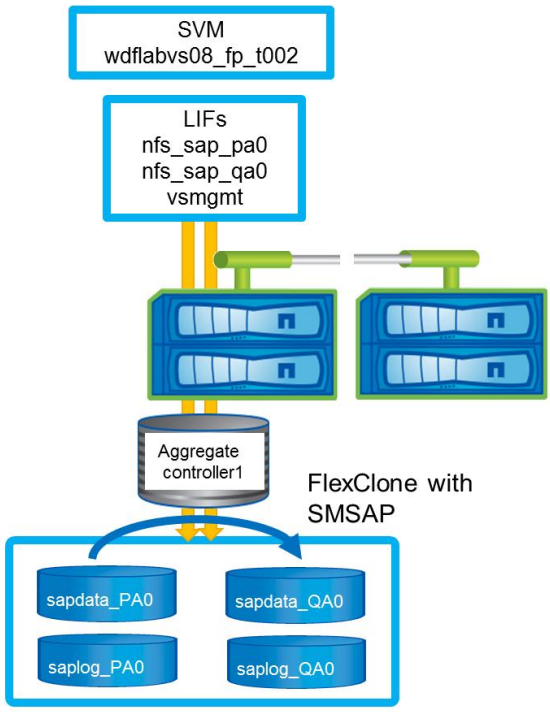

This example describes the installation of the new system, QA0, as described in chapter 6, "Sizing

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
- High-availability requirements

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 SAP Application Performance Standard (SAPS) 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.

SAP System Installation.” After the system installation is finished, the data volume sapdata_QA0 will be deleted, and the system copy will be executed using SMSAP.

SAP System Installation

The SAP system QA0 is installed as described in chapter 6, “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
- High-availability requirements

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 SAP Application Performance Standard (SAPS) 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.

SAP System Installation."

After the installation is finished, the SAP system needs to be stopped, and the sapdata file systems need to be unmounted. The volume sapdata_QA0 can now be deleted.

SnapDrive for UNIX Configuration

The SnapDrive for UNIX (SDU) credentials have to be set for the management interface of the SVM:

```
t002-lnx-61:# snapdrive config set vsadmin vsmgmt
Password for vsadmin:
Retype password:
```

In addition the management path for the SAP specific LIF has to be configured.

```
t002-lnx-61:# snapdrive config set -mgmtpath vsmgmt nfs_sap_qa0
```

SDU also requires access via the SAP specific LIF of the source system.

```
t002-lnx-61:# snapdrive config set -mgmtpath vsmgmt nfs_sap_pa0
```

```
t002-lnx-61:# snapdrive config list -mgmtpath
system name          management interface   datapath interface
-------------------------------------------------------------
wdflabvs08_fp_t002   192.168.102.5          192.168.102.102|192.168.102.100
```

File System Preparations

The following files and directories need to be deleted:

```
t002-lnx-61 # rm -fr /oracle/QA0/origlogA/cntrl
t002-lnx-61:# rm -fr /oracle/QA0/origlogB/cntrl
t002-lnx-61:# rm -fr /oracle/QA0/mirrlogA/cntrl
t002-lnx-61:# rm -fr /oracle/QA0/oraarch
t002-lnx-61:# rm -fr /oracle/QA0/sapdata1
t002-lnx-61:# rm -fr /oracle/QA0/sapdata2
t002-lnx-61:# rm -fr /oracle/QA0/sapdata3
t002-lnx-61:# rm -fr /oracle/QA0/sapdata4
t002-lnx-61:# rm /oracle/QA0/112_64/dbs/initQA0.ora
t002-lnx-61:# rm /oracle/QA0/112_64/dbs/spfileQA0.ora
```

Precloning and Postcloning Scripts

The script cleanup.sh is used to delete specific files and directories before the cloning process. The script must be copied from `/opt/NetApp/smsap/plugins/examples/clone/create/pre` to `/opt/NetApp/smsap/plugins/clone/create/pre/`.

Note: The function “execute” of the script has to be adapted. The corrected version of the function execute is available in the appendix of this document.

The script `os_db_authentication.sh` is used to configure the OPS$ mechanism after the cloning process. The script has to be copied from `/opt/NetApp/smsap/plugins/examples/clone/create/post` to `/opt/NetApp/smsap/plugins/clone/create/post/`.

Note: The function “execute” of the script has to be adapted. The corrected version of the function execute is available in the appendix of this document.

The script `sap_follow_up_activities.sh` is used to delete batch jobs in the target SAP system after the cloning process. The script must be copied from `/opt/NetApp/smsap/plugins/examples/clone/create/post` to `/opt/NetApp/smsap/plugins/clone/create/post/`.

Download `oradbusr10.zip` as described in SAP note 50088. Extract the zip file and copy `ORADBUSR.SQL` to `/opt/NetApp/smsap/plugins/clone/create/post/`.

SAP System Copy with SMSAP

The SAP system copy process can now be started within the SMSAP GUI. The preceding pre- and postcloning scripts need to be selected within the SMSAP cloning wizard dialog box. The configuration should be saved in order to simplify the SAP system copy process when further system refreshes are executed.

Example xml files for the configuration can be found in the appendix of this document.

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.

Appendix

SMSAP Postbackup Script

```
#!/bin/bash
# Copyright (c) 2013 NetApp, Inc.
# All rights reserved.
#
#IMPORTANT NOTE: This script is provided for reference only. It has been tested with SnapDrive
5.2 for LINUX but may not work in all environments.  Please review and then customize based on
your secondary protection requirements.
#
name="Set SnapMirror label"
description="Sets the SnapMirror label SMSAP-<Backup Retention Class>, e.g. SMSAP-HOURLY"
context=
timeout="0"
parameter=("VSERVER_NAME             :Primary Vserver Name")

SM_LABEL="SMSAP-$SM_BACKUP_RETENTION"

EXIT=0

function _exit {

      rc=$1

        echo "Command complete."

        exit $rc
}

function usage {
        echo "usage: $(basename $0) { -check | -describe | -execute }"
        _exit 99
}

function describe {
        echo "SM_PI_NAME:$name"
        echo "SM_PI_DESCRIPTION:$description"
        echo "SM_PI_TIMEOUT:$timeout"
         IFS=^
        for entry in ${parameter[@]}; do
                echo "SM_PI_PARAMETER:$entry"
        done

        _exit 0
}

function check {
        _exit 0
}

function execute {
```

```
     echo "execute started"

     index=0
     IFS=" "
     for VALUE in ${SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS[@]}; do

        PRIMARY_SNAPSHOT_NAME[index]=`echo $VALUE | awk -F ":" '{print $1}'`

        MOUNT_POINT[index]=`echo $VALUE | awk -F ":" '{print $2}'`

        VOLUME_NAME[index]=`cat /proc/mounts | grep "${MOUNT_POINT[$index]} " | awk -F " "
'{print $1}' | awk -F "/" '{print $2}'`

       ssh -l vsadmin $VSERVER_NAME snap modify -volume ${VOLUME_NAME[$index]} -snapshot
${PRIMARY_SNAPSHOT_NAME[$index]} -snapmirror-label $SM_LABEL
       if [ $? -ne 0 ] ; then
          _exit 4
        fi

        ((index=index+1))
     done

     echo "execute ended"

    _exit 0



}

case $(echo $1 | tr [A-Z] [a-z]) in
        -check)    check
                   ;;
        -execute)  execute
                   ;;
        -describe) describe
                   ;;
        *)         echo "unknown option $1"
                   usage
                   ;;
esac
```

Function “execute” of cleanup.sh Script

```
function execute {
        echo "cleaning up the environment"

        [ -z "$SM_TARGET_SID" ] && echo "target SID [SM_TARGET_SID] not set" && _exit 4

        files_to_cleanup=("/oracle/${SM_TARGET_SID}/origlogA/cntrl:N"
                          "/oracle/${SM_TARGET_SID}/mirrlogA/cntrl:N"
                          "/oracle/${SM_TARGET_SID}/origlogB/cntrl:N"
                          "/oracle/${SM_TARGET_SID}/origlogA/log_g11m1.dbf:N"
                          "/oracle/${SM_TARGET_SID}/origlogB/log_g12m1.dbf:N"
                          "/oracle/${SM_TARGET_SID}/origlogA/log_g13m1.dbf:N"
                          "/oracle/${SM_TARGET_SID}/origlogB/log_g14m1.dbf:N"
                          "/oracle/${SM_TARGET_SID}/mirrlogA/log_g11m2.dbf:N"
                          "/oracle/${SM_TARGET_SID}/mirrlogB/log_g12m2.dbf:N"
                          "/oracle/${SM_TARGET_SID}/mirrlogA/log_g13m2.dbf:N"
                          "/oracle/${SM_TARGET_SID}/mirrlogB/log_g14m2.dbf:N"
                          "/oracle/${SM_TARGET_SID}/saptrace/usertrace:Y"
                          "/oracle/${SM_TARGET_SID}/saptrace/background:Y"
                          "/oracle/${SM_TARGET_SID}/102_64/dbs/init${SM_TARGET_SID}.ora:Y"
        )

        IFS=^
        for entry in ${files_to_cleanup[@]} ; do
                file=$(echo "$entry" | awk -F':' '{ print $1 }')

echo $file
rm -fr $file
        done

        _exit 0
}
```

Function “execute” of os_db_authentication.sh

```
function execute {
        EXIT=0
        [ -z "$SCHEMAOWNER" ] && EXIT=4 && echo "parameter [SCHEMAOWNER] not set"
        [ -z "$ORADBUSR_FILE" ] && EXIT=4 && echo "parameter [ORADBUSR_FILE] not set"
        [ -z "$SM_TARGET_SID" ] && EXIT=4 && echo "parameter [SM_TARGET_SID] not set"
        [ -z "$SM_ORIGINAL_SID" ] && EXIT=4 && echo "parameter [SM_ORIGINAL_SID] not set"

        [ $EXIT -ne 0 ] && echo "processing stopped due to missing parameters" && _exit $EXIT

        [ ! -f "$ORADBUSR_FILE" ] && echo "file [$ORADBUSR_FILE] is not a regular file" && _exit
4

        sqlplus /nolog @${ORADBUSR_FILE} $SCHEMAOWNER UNIX $SM_TARGET_SID x

        . ${ORACLE_HOME}/../.profile && . ${ORACLE_HOME}/../.dbenv.sh


        ${DIR_LIBRARY}/brconnect -u / -c force -f chpass -o SAPSR3 -p sap

        _exit $?
}
```

Backup Tasks Specification File Example

Example of the postbackup task specification to enable data protection with SMSAP:

```
<preposttask-specification xmlns="http://www.netapp.com">
  <task-specification>
    <post-tasks>
      <task>
        <name>Set SnapMirror label</name>
        <parameter>
          <name>VSERVER_NAME</name>
          <value>wdflabvs08_fp_t002</value>
        </parameter>
      </task>
    </post-tasks>
  </task-specification>
</preposttask-specification>
```

Clone Tasks Specification File Example

Example of the postcloning task specification to configure OPS$ mechanism and delete batch after the SAP system copy process with SMSAP:

```
<preposttask-specification xmlns="http://www.netapp.com">
  <task-specification>
    <pre-tasks>
      <task>
        <name>clone cleanup</name>
      </task>
    </pre-tasks>
    <post-tasks>
      <task>
        <name>Oracle Users for OS based DB authentication</name>
        <parameter>
          <name>SCHEMAOWNER</name>
          <value>SAPSR3</value>
        </parameter>
        <parameter>
          <name>ORADBUSR_FILE</name>
          <value>/opt/NetApp/smsap/plugins/clone/create/post/ORADBUSR.SQL</value>
        </parameter>
      </task>
      <task>
        <name>SAP SystemCopy follow-up activities</name>
        <parameter>
          <name>SCHEMAOWNER</name>
          <value>SAPSR3</value>
        </parameter>
      </task>
    </post-tasks>
  </task-specification>
</preposttask-specification>
```

Clone Specification File Example

Example of clone specification for an SAP system copy PA0 to QA with SMSAP:

```
clone-specification xmlns="http://www.netapp.com">
  <database-specification>
    <controlfiles>
      <file>/oracle/QA0/origlogA/cntrl</file>
      <file>/oracle/QA0/origlogB/cntrl</file>
      <file>/oracle/QA0/mirrlogA/cntrl</file>
    </controlfiles>
    <redologs>
      <redogroup>
        <file>/oracle/QA0/origlogA/log_g11m1.dbf</file>
        <file>/oracle/QA0/mirrlogA/log_g11m2.dbf</file>
```

```
        <number>1</number>
        <size unit="M">50</size>
      </redogroup>
      <redogroup>
        <file>/oracle/QA0/origlogB/log_g12m1.dbf</file>
        <file>/oracle/QA0/mirrlogB/log_g12m2.dbf</file>
        <number>2</number>
        <size unit="M">50</size>
      </redogroup>
      <redogroup>
        <file>/oracle/QA0/origlogA/log_g13m1.dbf</file>
        <file>/oracle/QA0/mirrlogA/log_g13m2.dbf</file>
        <number>3</number>
        <size unit="M">50</size>
      </redogroup>
      <redogroup>
        <file>/oracle/QA0/origlogB/log_g14m1.dbf</file>
        <file>/oracle/QA0/mirrlogB/log_g14m2.dbf</file>
        <number>4</number>
        <size unit="M">50</size>
      </redogroup>
    </redologs>
    <parameters>
      <parameter>
        <name>audit_file_dest</name>
        <value>/oracle/QA0/saptrace/audit</value>
      </parameter>
      <parameter>
        <name>log_archive_dest_1</name>
        <value>LOCATION=/oracle/QA0/oraarch</value>
      </parameter>
      <parameter>
        <name>_b_tree_bitmap_plans</name>
        <value>FALSE</value>
      </parameter>
      <parameter>
        <name>log_archive_format</name>
        <value>%t_%s_%r.dbf</value>
      </parameter>
      <parameter>
        <name>control_file_record_keep_time</name>
        <value>30</value>
      </parameter>
      <parameter>
        <name>star_transformation_enabled</name>
        <value>true</value>
      </parameter>
      <parameter>
        <name>remote_os_authent</name>
        <value>TRUE</value>
      </parameter>
      <parameter>
        <name>replication_dependency_tracking</name>
        <value>FALSE</value>
      </parameter>
      <parameter>
        <name>_optim_peek_user_binds</name>
        <value>FALSE</value>
      </parameter>
      <parameter>
        <name>shared_pool_size</name>
        <value>1124073472</value>
      </parameter>
      <parameter>
        <name>db_cache_size</name>
        <value>1124073472</value>
      </parameter>
      <parameter>
        <name>_index_join_enabled</name>
        <value>FALSE</value>
      </parameter>
```

```
      <parameter>
        <name>sessions</name>
        <value>260</value>
      </parameter>
      <parameter>
        <name>pga_aggregate_target</name>
        <value>1498624819</value>
      </parameter>
      <parameter>
        <name>parallel_execution_message_size</name>
        <value>16384</value>
      </parameter>
      <parameter>
        <name>parallel_threads_per_cpu</name>
        <value>1</value>
      </parameter>
      <parameter>
        <name>max_dump_file_size</name>
        <value>20000</value>
      </parameter>
      <parameter>
        <name>_optimizer_mjc_enabled</name>
        <value>FALSE</value>
      </parameter>
      <parameter>
        <name>filesystemio_options</name>
        <value>setall</value>
      </parameter>
      <parameter>
        <name>log_checkpoints_to_alert</name>
        <value>TRUE</value>
      </parameter>
      <parameter>
        <name>_in_memory_undo</name>
        <value>FALSE</value>
      </parameter>
      <parameter>
        <name>_table_lookup_prefetch_size</name>
        <value>0</value>
      </parameter>
      <parameter>
        <name>db_recovery_file_dest_size</name>
        <value>31457280000</value>
      </parameter>
      <parameter>
        <name>undo_tablespace</name>
        <value>PSAPUNDO</value>
      </parameter>
      <parameter>
        <name>processes</name>
        <value>130</value>
      </parameter>
      <parameter>
        <name>open_cursors</name>
        <value>2000</value>
      </parameter>
      <parameter>
        <name>recyclebin</name>
        <value>off</value>
      </parameter>
      <parameter>
        <name>query_rewrite_enabled</name>
        <value>false</value>
      </parameter>
      <parameter>
        <name>remote_login_passwordfile</name>
        <value>EXCLUSIVE</value>
      </parameter>
      <parameter>
        <name>_sort_elimination_cost_ratio</name>
        <value>10</value>
```

```
      </parameter>
    </parameters>
    <oracle-home>/oracle/QA0/112_64</oracle-home>
    <oracle-os-account>
      <username>oraqa0</username>
      <group>dba</group>
    </oracle-os-account>
    <sql-statements/>
  </database-specification>
<storage-specification>
    <storage-mapping>
      <mountpoint>
        <source>/oracle/PA0/sapdata1</source>
        <destination>/oracle/QA0/sapdata1</destination>
      </mountpoint>
      <mountpoint>
        <source>/oracle/PA0/sapdata3</source>
        <destination>/oracle/QA0/sapdata3</destination>
      </mountpoint>
      <mountpoint>
        <source>/oracle/PA0/sapdata2</source>
        <destination>/oracle/QA0/sapdata2</destination>
      </mountpoint>
      <mountpoint>
        <source>/oracle/PA0/sapdata4</source>
        <destination>/oracle/QA0/sapdata4</destination>
      </mountpoint>
    </storage-mapping>
  </storage-specification>
</clone-specification>
```

Version History

Version	Date	Document Version History
Version 1.0	December2013	First version

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com