



Technical Report

# SAP with Oracle on UNIX and NFS with NetApp Clustered Data ONTAP and SnapManager for SAP 3.4

Nils Bauer, NetApp  
March 2016 | 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 of this document is the common storage infrastructure design, deployment, operation, and management challenges faced by business and IT leaders who use the latest generation of SAP solutions. The recommendations in this document are generic; they are not specific to an SAP application or to the size and scope of the SAP implementation. This document assumes that the reader has a basic understanding of the technology and operation of NetApp® and SAP products. The document was developed based on the interaction of technical staff from NetApp, SAP, Oracle, and our customers.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction.....</b>	<b>5</b>
1.1	Business Challenges Facing SAP Customers.....	5
1.2	Technology Challenges of Expanding SAP Landscape .....	5
1.3	NetApp Solutions for SAP .....	5
<b>2</b>	<b>Business Continuance.....</b>	<b>6</b>
2.1	Backup and Recovery .....	6
2.2	SAP Repair System.....	9
2.3	Disaster Recovery .....	10
<b>3</b>	<b>System Management and Maintenance .....</b>	<b>12</b>
3.1	SAP System Copy .....	12
3.2	SAP Testing Cycle .....	14
<b>4</b>	<b>Storage Virtualization with Clustered Data ONTAP .....</b>	<b>15</b>
4.1	Hardware Support and Basic System Overview.....	15
4.2	Scalability .....	16
4.3	Storage Efficiency and Data Protection.....	16
4.4	Cluster Virtualization and Multitenancy Concepts .....	17
4.5	Physical Cluster Components .....	17
4.6	Logical Cluster Components .....	17
<b>5</b>	<b>Storage Setup Considerations.....</b>	<b>18</b>
5.1	SVM Configuration .....	18
5.2	Volume Layout and LIF Configuration .....	18
5.3	Multiple SAP Systems Sharing One LIF .....	18
5.4	One or Multiple LIFs per SAP System.....	19
5.5	LIF Configuration Summary.....	20
<b>6</b>	<b>Volume and LIF Configuration Examples .....</b>	<b>21</b>
6.1	Minimum Configuration.....	22
6.2	Separating Oracle Mirrored Redo Logs.....	22
6.3	Large SAP Systems with High-Performance Requirements.....	23
<b>7</b>	<b>SAP System Installation .....</b>	<b>24</b>
7.1	Storage Network.....	24
7.2	Operating System Configuration .....	24
7.3	Snapshot Configuration.....	25
7.4	SAP Installation Process .....	25
<b>8</b>	<b>SnapManager for SAP 3.4 Setup and Configuration .....</b>	<b>27</b>
8.1	SAP Systems and SMSAP Lab Setup.....	27

8.2	Configure SMSAP Repository .....	28
<b>9</b>	<b>Back Up Using SnapManager for SAP .....</b>	<b>31</b>
9.1	Install and Configure SAP Source System P02.....	31
9.2	Configure Data Protection to Secondary Storage System.....	43
<b>10</b>	<b>SAP System Copies with SnapManager for SAP .....</b>	<b>49</b>
10.1	Install and Configure SAP Target System Q02 .....	49
10.2	SAP System Copy on Primary Storage .....	52
10.3	SAP System Copy on Secondary Storage .....	61
<b>11</b>	<b>Configure Br*Tools.....</b>	<b>63</b>
11.1	Configure Backint File .....	63
11.2	Configure SMSAP Credentials .....	64
11.3	Execute Brbackup Using CLI.....	64
11.4	Execute Backup Using SAP DBA Planning Calendar .....	67
11.5	SAP System Copy Using Backup Created by Brbackup .....	68
	<b>Version History .....</b>	<b>69</b>

## LIST OF TABLES

Table 1)	LIF configuration summary. ....	21
Table 2)	Minimum configuration, LIF, and volume layout.....	22
Table 3)	LIF and volume layout for separated mirrored redo logs. ....	23
Table 4)	LIF and volume layout for large SAP systems. ....	23
Table 5)	Volumes and mount points. ....	25

## LIST OF FIGURES

Figure 1)	Backup solution overview. ....	7
Figure 2)	Time requirements for different backup methods. ....	8
Figure 3)	Time requirements for restore and recovery operations. ....	9
Figure 4)	SAP repair system.....	10
Figure 5)	DR with SnapMirror. ....	11
Figure 6)	Synchronous storage replication. ....	11
Figure 7)	Traditional SAP system copy.....	12
Figure 8)	SAP system copy: NetApp approach.....	13
Figure 9)	SAP system copy: standard approach.....	14
Figure 10)	SAP system copy: NetApp approach.....	14
Figure 11)	SAP testing cycle.....	15
Figure 12)	Data ONTAP cluster overview. ....	16
Figure 13)	Configuration with one LIF per storage node.....	18
Figure 14)	Migration of multiple systems. ....	19
Figure 15)	Migration of a single system. ....	19

Figure 16) Configuration with one LIF per SAP system. ....	20
Figure 17) Migration of multiple single SAP systems. ....	20
Figure 18) Volume and LIF configuration examples. ....	22
Figure 19) Dedicated 10GbE storage network. ....	24
Figure 20) SMSAP lab setup. ....	28

# 1 Introduction

This document addresses the challenges of designing storage solutions to support SAP business suite products using an Oracle database. The primary focus of this document is the common storage infrastructure design, deployment, operation, and management challenges faced by business and IT leaders who use the latest generation of SAP solutions. The recommendations in this document are generic; they are not specific to an SAP application or to the size and scope of the SAP implementation. This document assumes that the reader has a basic understanding of the technology and operation of NetApp and SAP products. The document was developed based on the interaction of technical staff from NetApp, SAP, Oracle, and our customers.

## 1.1 Business Challenges Facing SAP Customers

Corporations that deploy 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 is usually larger than a single production instance of SAP. Business process owners and project managers must coordinate with IT managers 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. Therefore, nondisruptive operations are a key requirement.

## 1.2 Technology Challenges of 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 that 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 for quality assurance, testing, or perhaps a final staging system before 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 be 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 the 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 current data.

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

As SAP landscapes have expanded, the technology has also changed. SAP has evolved to take advantage of the latest technology trends. Virtualization and cloud technologies have become predominant as corporations seek to leverage efficient computing methods to maximize their investment and reduce data center expenses. Without a storage infrastructure that can adapt to the needs of the changing technology, IT organizations would be unable to meet the business needs of the company.

## 1.3 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 clustered Data ONTAP® 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 meets the high availability (HA) and performance requirements of SAP applications so that stringent service-level agreements (SLAs) can be achieved. In addition, NetApp helps to reduce the administration and management costs associated with deploying business applications and processes.

## 2 Business Continuity

### 2.1 Backup and Recovery

Corporations today require their SAP applications to be available 24 hours a day, 7 days a week. Consistent levels of performance are expected regardless of the ever-increasing data volumes and need for 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. With backup windows shrinking and the amount of data that needs to be backed up still increasing, it is difficult to define a time when backups can be performed with minimal impact on the business processes. The time needed to restore and recover SAP systems is of particular concern. That is because the downtime of SAP production and nonproduction systems must be minimized to minimize both the data loss and the cost to the business.

The following descriptions summarize the SAP backup and recovery challenges:

- **Performance impact on production SAP systems.** Conventional backups typically lead to a significant performance impact on the production SAP system. That is because there is a heavy load on the database server, the storage system, and the storage network during traditional copy-based backups.
- **Shrinking backup windows.** Conventional backups can be taken only during times when little dialog or batch activities take place on the SAP system. The scheduling of backups becomes more and more difficult to define when the SAP systems are in use 24/7.
- **Rapid data growth.** Rapid data growth, together with shrinking backup windows, results in ongoing investments into the backup infrastructure. These investments include additional tape drives, newer tape drive technology, faster storage networks, and the ongoing cost of storing and managing those tape assets. Incremental or differential backups can address these issues, but this option results in a very slow, cumbersome, and complex restoration process that is harder to verify. The option also usually leads to increased or elongated recovery time objective (RTO) or recovery point objective (RPO) times that are not acceptable to the business.
- **Increasing cost of downtime.** Unplanned downtime of an SAP system always has a financial impact on the business. A significant part of the unplanned downtime is the time that is needed to restore and recover the SAP system after a failure. The backup and recovery architecture must be designed based on an acceptable 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 dramatically cuts down the total available time for the upgrade process. The go/no-go decision is generally based on the amount of time required to restore and recover the database from the backup that was previously created. The option to restore very quickly allows more time to solve problems that might occur during the upgrade process rather than just restore the system back to its previous state.

NetApp Snapshot® technology can be used to create either online or offline database backups within minutes. Because a Snapshot copy does not move any physical data blocks on the storage platform, the time needed to create a Snapshot copy is independent of the size of the database. The use of Snapshot technology also has no performance impact on the live SAP system. That is because the NetApp Snapshot copy does not move or 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

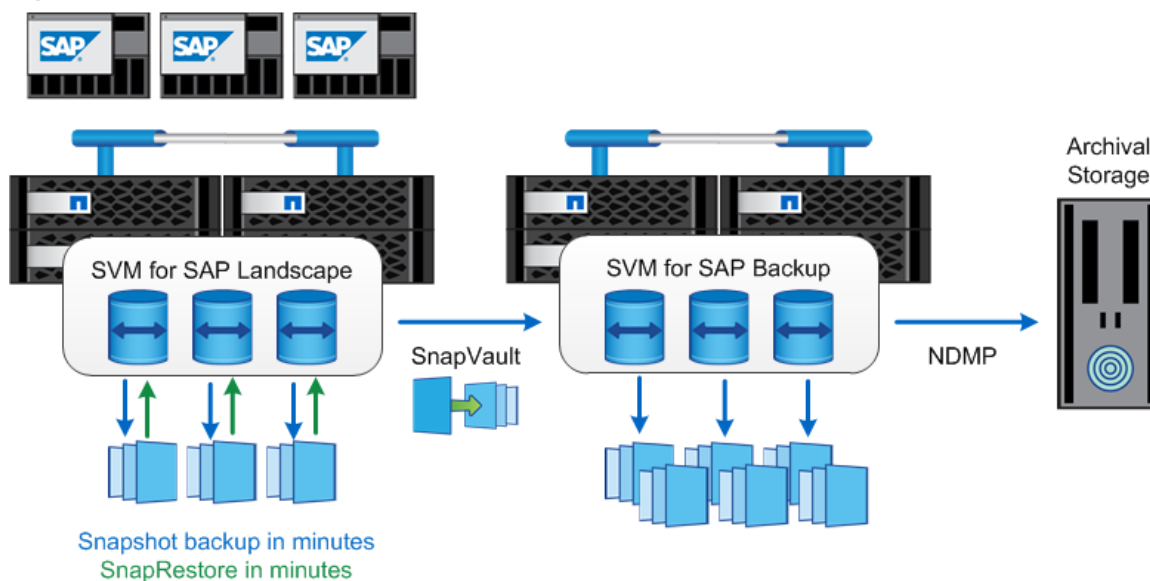
multiple online Snapshot backups during the day; for example, scheduling backups every four hours is common. These Snapshot backups are typically kept for three to five days on the primary storage system before being removed.

Snapshot copies also provide key advantages for the restore and recovery operation. NetApp SnapRestore® functionality allows restoration of the entire database, or even just parts of the database, to any point in time based on the available Snapshot copies. This restore process is performed in a few minutes, independent of the size of the database. Because several online Snapshot backups are created during the day, the actual time needed for the recovery process is dramatically reduced, as opposed to a traditional backup approach. A restore operation can be performed using a Snapshot copy that is only a few hours old (rather than up to 24 hours old); therefore, fewer transaction logs need to be applied. As a result, the mean time to recover, which is the time needed for restore and recovery operations, is reduced to just several minutes compared to multiple hours with conventional single-cycle 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 to, not a replacement for, backups to a secondary location. Most restore and recovery operations 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 month-end backup.

Figure 1 illustrates the backup solution overview.

Figure 1) Backup solution overview.



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

- A backup and restore operation using Snapshot and SnapRestore technology
- A backup and restore operation to and from a secondary location

A backup to a secondary location is always based on the 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 NetApp SnapVault® software.** The primary storage virtual machine (SVM) communicates directly with the secondary SVM and sends the backup data to the destination. The SnapVault functionality offers significant advantages compared to those of tape backups. After an initial data transfer, in which all the data has to be synchronized from the source to the destination, all subsequent backups copy only the changed blocks to the secondary storage using compressed blocks to reduce network traffic. The typical block change rate for an Oracle SAP system is around 2% per day. Therefore, the load on the primary storage system and the time needed for a full backup

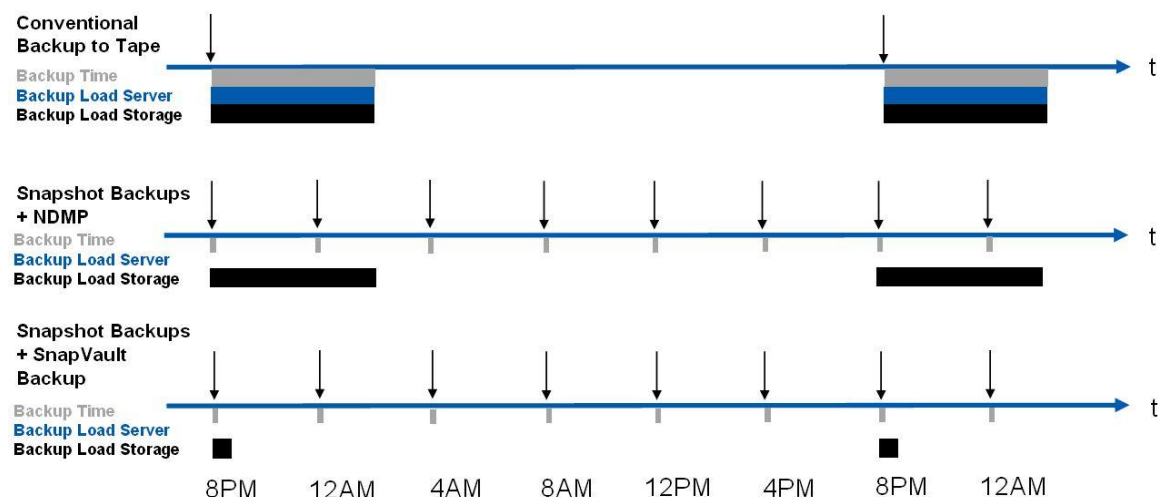
are significantly reduced. Because SnapVault stores only the changed blocks at the destination, a full database backup requires significantly less disk space.

If backing up data to tape for longer-term storage is still required, for example, a monthly backup that is kept for a year, the tape infrastructure is directly connected to the secondary SVM. The data is written to tape using NDMP, with no performance impact on the production SAP systems.

- **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 2 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 2) Time requirements for different backup methods.



### Snapshot Backups 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 technology consumes disk space only for the changed blocks. Snapshot backups are typically scheduled more often, for example, every four hours. A more frequent backup schedule enables more flexible restore options and reduces the number of logs that must be applied during the 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 the conventional tape backup process.

### Snapshot Backups with Disk-to-Disk Backup and SnapVault


The Snapshot backups described in this section are used 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 required 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 that for full tape backups.

Figure 3 compares the time required to perform restore and recovery operations.



Figure 3) Time requirements for restore and recovery operations.



### Restore from Tape or SnapVault Restore

The time required to restore a database from tape or disk depends on the size of the database and the tape or disk infrastructure that is being used. In either case, several hours are required to perform a restore operation. Because the backup frequency is typically one per day, a certain number of transaction logs must be applied after the restore operations are complete.

### Restore with SnapRestore

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

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

The combination of Snapshot and SnapRestore disk-to-disk backups (a 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 ideal for running database consistency checks. NetApp SnapManager® software offers the ability 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.

## 2.2 SAP Repair System

More and more companies face the challenges of addressing logical errors in a more complex SAP environment in which several SAP systems constantly 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 was recovered to a point in time in the past
- Inconsistency between the system that was restored and recovered to a point in time that is considered in the past by the other systems with which it exchanges data

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

Figure 4) SAP repair system.

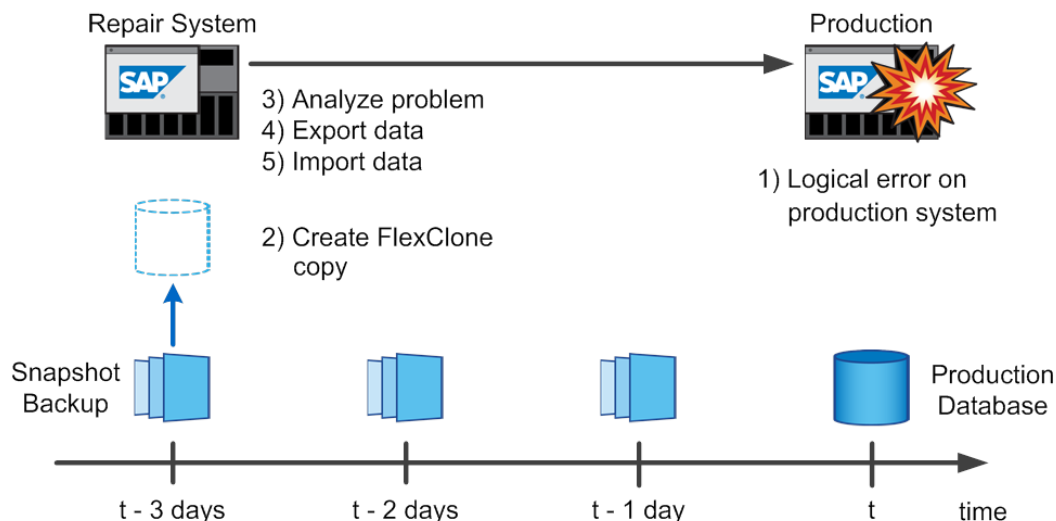


Figure 4 shows the general process for creating and using the repair system:

1. A logical error is discovered on the production system. Depending 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 can be used to create an SAP system copy on which the production will be based. 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 or copied from the repair system.
5. The data is imported to the production system.

In this example, there is dramatically less or even no impact on the production system, no data loss, and no inconsistency within the SAP landscape.

The described scenario is quite simple, and it should be obvious that not all logical errors can be solved in this way. However, the repair system approach also helps in more complex scenarios, because there is more flexibility, and there are more options available to analyze and to recover from logical errors.

## 2.3 Disaster Recovery

Organizations recognize the importance of having a business continuance plan in place to deal with disasters. The costs of not having one—lost productivity, revenue, customer loyalty, and possibly even business failure—make it mandatory to have a plan that results in minimum downtime and rapid recovery from disaster.

### Asynchronous Replication

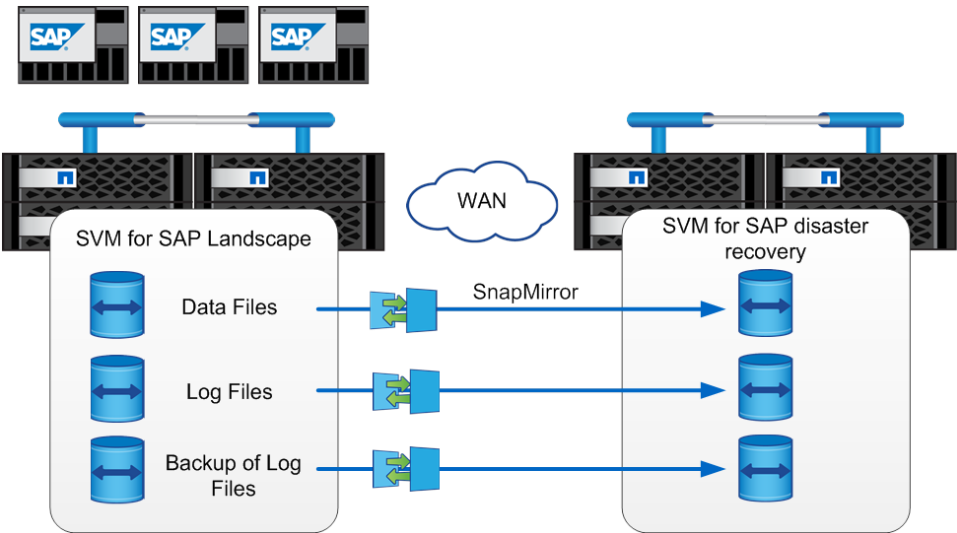
NetApp SnapMirror® software delivers a disaster recovery (DR) solution that today's global SAP systems need. By replicating data at high speeds over a LAN or a WAN, SnapMirror software provides the highest possible data availability and the fastest recovery.

SnapMirror technology copies data to one or more SVMs. It continually updates the mirrored data to keep it current and is therefore available for use in DR, backup to tape, read-only data distribution, testing, online data migration, and more.

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

Figure 5 illustrates DR with SnapMirror technology.


Figure 5) DR with SnapMirror.



### Synchronous Replication

The synchronous DR solution for SAP on Oracle is based on NetApp MetroCluster™ software. Figure 6 shows a high-level overview of the solution. The storage cluster at each site provides local HA and is used for production workloads. The data on each site is synchronously replicated to the other location and is available immediately in case of disaster or failover.

Figure 6) Synchronous storage replication.



## 3 System Management and Maintenance

### 3.1 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 simple 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 or QAS systems allows SAP customers to run more test or project systems and refresh those systems more often. Doing so enables project teams to reduce project cycles by running parallel testing and improves quality of testing and training with more data from production.

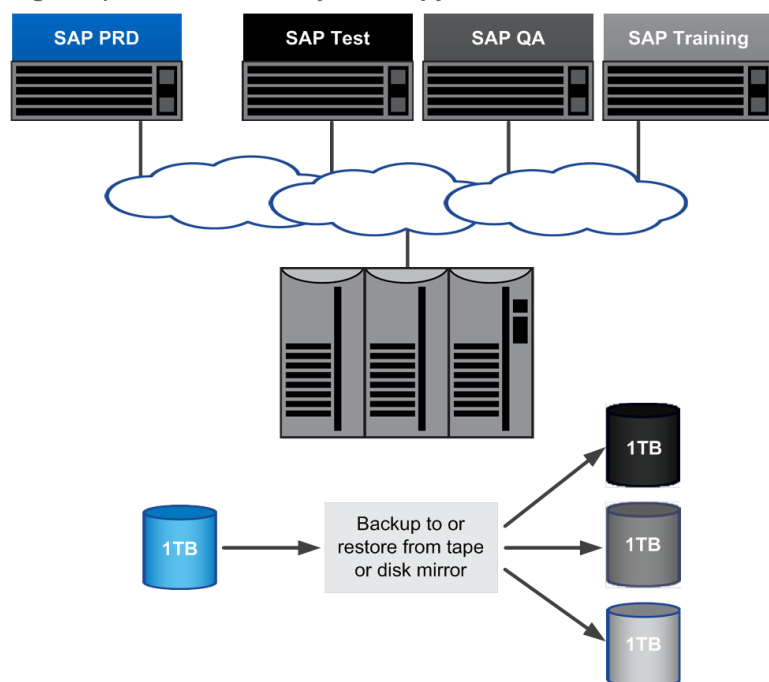
#### Capacity Requirements

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

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

Figure 17 illustrates a traditional SAP system copy.

Figure 7) 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 and doesn't need to be preprovisioned. 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 space required is limited to only data that is changed at the source and the target system. Therefore, the space needed for SAP system copies is significantly decreased.

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 well as the system change rates. As longer test systems are kept, more block changes will take place 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.

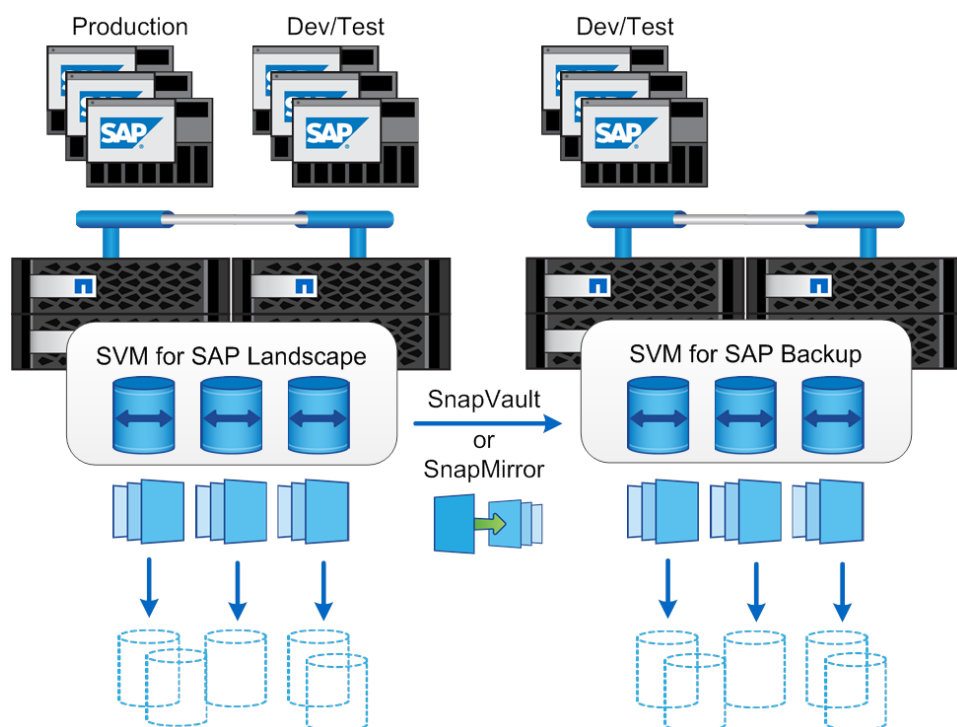
On the source system, a database-consistent Snapshot copy of the data files is created. This is done during online operation and does not affect performance on the source system. Therefore, this step can be carried out at any time.

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 even as a DR solution. The backup or DR replication images can be accessed for reading and writing using FlexClone technology. Existing backup or DR images will be used for test environments, leveraging unused, expensive DR assets. As a side effect, the backup and recovery or DR solution is tested without any additional effort and without any interruption.

Figure 8 illustrates the NetApp approach to creating an SAP system copy.

**Figure 8) SAP system copy: NetApp approach.**



## Time Requirements

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

- Time to create a backup of the source system
- Time to restore the backup to the target system
- Time to perform OS and database-specific postprocessing
- Time to perform SAP application postprocessing


**Note:** The SAP postprocessing time depends on the customer's SAP environment. Some customers can complete postprocessing in a few hours, while other customers need several days.

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,

performance will be affected on the source system during the backup. Because of the large number of logs that potentially 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.


Figure 9 and Figure 10 illustrate the differences between the amount of time spent creating an SAP system copy using a standard approach versus the time spent using the NetApp approach.

**Figure 9) SAP system copy: standard approach.**



All of the steps leading up to the point when the SAP system is started on the target host can be accomplished in only a few minutes using the NetApp solution, compared with several hours using the standard approach. With both approaches, the SAP postprocessing must be done as an additional step, but it can also be automated using third-party tools.

**Figure 10) SAP system copy: NetApp approach.**



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

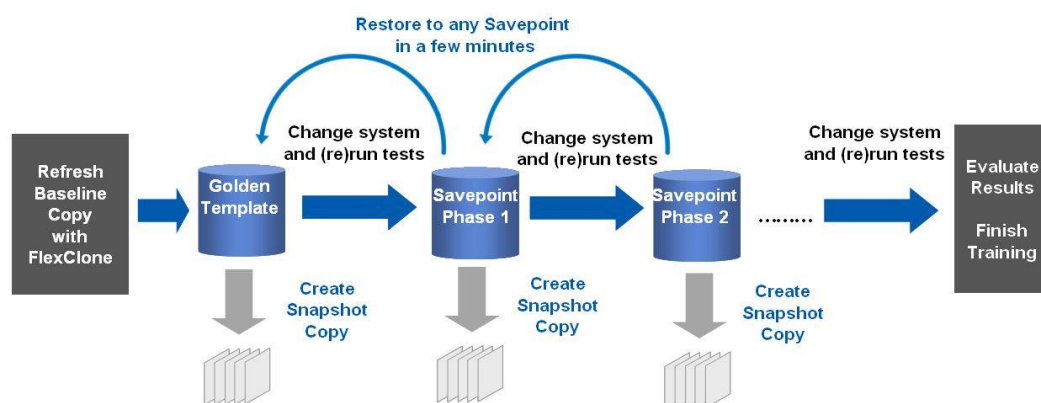
## 3.2 SAP Testing Cycle

The ability to create a backup in seconds and to use that backup to restore to another SAP system at any point in time through Snapshot copies is a turning point for SAP development and test environments. Projects such as data loads, SAP upgrades, and support package installations can be accelerated by using fast backup and restore functionalities. During these projects, backups can be performed at specific phases or on an ad hoc basis, allowing the systems to be easily and quickly reset to any previous point in time. Doing so enables you to repeat that phase as needed. Test runs can be easily repeated or even run in parallel with different code or configurations to make sure that the results are valid.

Figure 11 illustrates the SAP testing cycle.



Figure 11) SAP testing cycle.



Carrying out SAP upgrades 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 state can always be restored. Being able to create a backup or restore the system in minutes instead of hours allows business teams to do more testing before going live, reducing risk to the business.

The specific system changes are usually made first 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 be solved only 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. If an error occurs, 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 state. 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, a normal two-day weekend might not be sufficient for the upgrade. NetApp SnapManager software offers Snapshot technology as a backup method and SnapRestore technology to restore the system to its former release status. These technologies allow greater flexibility in 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.

## 4 Storage Virtualization with Clustered Data ONTAP

This section describes the architecture of NetApp 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.


### 4.1 Hardware Support and Basic System Overview

As shown in Figure 12, a clustered Data ONTAP system contains NetApp storage controllers. The basic building block is the HA pair. An HA pair employs two identical nodes of clustered Data ONTAP. Each node provides active data services and has redundant cable paths to the other node's storage. If either node is down for any reason, planned or unplanned, the HA partner can take over the failed node's storage and maintain access to the data. When the downed system rejoins the cluster, the partner node returns the storage resources back to the original node.

The minimum cluster size is two matching nodes for an HA pair. Using NetApp nondisruptive technology refresh, a simple two-node, entry-level cluster can evolve into a much larger cluster by adding additional nodes of the same type. The cluster can also evolve by adding nodes of a more or less powerful controller model while the cluster is still online. At the time of writing, clusters with SAN protocols support

up to eight nodes using both midsize and high-end controllers. NAS-only clusters made up of high-end controllers scale up to 24 nodes.

Figure 12) Data ONTAP cluster overview.



One of the key differentiators for 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 both SAN hosts and NAS clients. This clustered shared pool appears as a single system image for management purposes and provides a common point of management through either GUI or CLI tools.

## 4.2 Scalability

Clustered Data ONTAP supports different controller types within the same cluster, protecting initial hardware investments and providing the flexibility to adapt resources to meet business demands and 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 NetApp DataMotion™ data migration capabilities of clustered Data ONTAP.

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

## 4.3 Storage Efficiency and Data Protection

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

- **Thin provisioning.** Volumes are created by using virtual sizing. Thin provisioning is the most efficient way to provision storage, because storage is not preallocated up front, even though the clients see the total storage space assigned to them. In other words, when a volume or LUN is created with 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 that time, only the required space to store the data is consumed. Unused storage is shared across all volumes, and each of the volumes can grow and shrink on demand.
- **NetApp Snapshot copies.** Automatically scheduled, point-in-time Snapshot copies 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.** These near-zero-space, exact, writable virtual copies of datasets offer rapid, space-efficient creation of additional data copies that are well suited for test and development environments.



- **NetApp SnapMirror data replication software.** Asynchronous replication of volumes is supported, independent of the protocol, either within the cluster or to another clustered Data ONTAP system for data protection and DR.
- **NetApp SnapVault backup software.** Volumes can be copied for space-efficient, read-only, disk-to-disk backup, either within the cluster or to another clustered Data ONTAP system.
- **NetApp MetroCluster software.** Continuous data availability is supported beyond the data center or the cluster. MetroCluster is native within the Data ONTAP operating system (OS). It provides a synchronous mirroring relationship between two distinct but identically configured two-node clusters over distances up to 200km apart.

## 4.4 Cluster Virtualization and Multitenancy Concepts

A cluster is composed of physical hardware, including storage controllers with attached disk shelves, network interface cards (NICs), and, optionally, Flash Cache cards. Together these components create a physical resource pool that is virtualized as a logical cluster resource to provide data access. Abstracting and virtualizing physical assets into logical resources provides flexibility and, potentially, multitenancy in clustered Data ONTAP. These processes also enable the DataMotion capabilities, which are at the heart of nondisruptive operations.

## 4.5 Physical Cluster Components

Storage controllers, independent of the model deployed, are considered equal 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 are protected with NetApp RAID DP® technology. NICs and host bus adapters (HBAs) provide physical ports (Ethernet and FC) for connections to the management and data networks. The physical components of a system 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 defined SVMs that contain volumes and logical interfaces (LIFs).

## 4.6 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 requires a minimum of one and can be scaled up to hundreds of SVMs within a single cluster. Each SVM is configured for the client and host access protocols that it supports in any combination of SAN and NAS. Each SVM contains at least one volume and at least one LIF.

The administration of each SVM can optionally be delegated so that separate administrators are responsible for provisioning volumes and other SVM-specific operations. This capability is particularly appropriate for multitenant environments or when workload separation is desired. SVM-delegated administrators have visibility to only their specific SVM and have no knowledge of any other hosted SVM.

For NAS clients, the volumes in each SVM are joined together into a namespace for CIFS and Network File System (NFS) access. For SAN hosts, LUNs are defined within volumes and mapped to hosts.

The accessing hosts and clients connect to the SVM through a LIF. LIFs present either an IP address (used by NAS clients and iSCSI hosts) or a worldwide port name (WWPN, for FC and FCoE access). Each LIF has a home port on an NIC or HBA. LIFs are used to virtualize the NIC and HBA ports rather than for mapping IP addresses or WWPNs directly to the physical ports. That is 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, rebalance client performance, or evacuate all resources on a controller for hardware lifecycle operations.

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

For more information, see [NetApp Clustered Data ONTAP 8.3 and 8.2.x: An Introduction](#).

## 5 Storage Setup Considerations

### 5.1 SVM Configuration

An SVM is a logical component of the storage cluster. The administration of an SVM can be delegated to separate administrators. In multitenancy environments, one or multiple SVMs are typically assigned to each tenant to allow each tenant to operate as its own environment separated from those of other tenants.

Multiple SAP landscapes can use a single SVM, or an individual SVM can be assigned to each SAP landscape if they are managed by different teams within a company.

**Note:** When SnapManager for SAP (SMSAP) is used to create system copies, the source system and the target system have to run within the same SVM.

### 5.2 Volume Layout and LIF Configuration

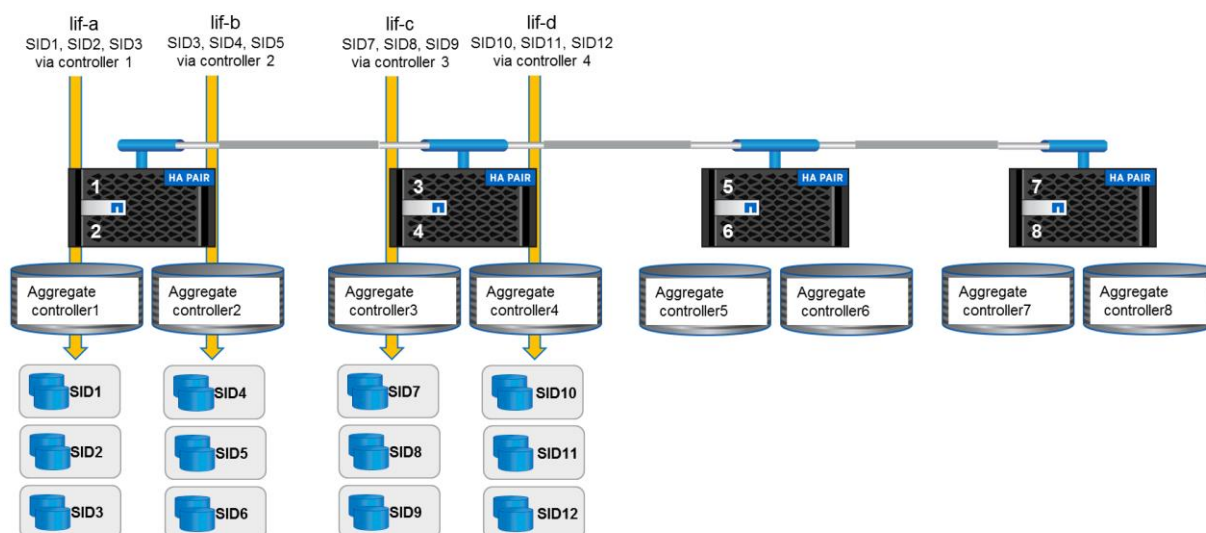
Clustered Data ONTAP enables you to migrate SAP systems nondisruptively to any of the other storage nodes within the storage cluster. A migration can be used to rebalance the I/O load of a controller within the cluster by moving load to any of the other storage nodes within the cluster. SAP systems can also be migrated to another storage node if, for example, the storage hardware is renewed and the old storage node is to be removed from the cluster.

The volume and LIF configuration within clustered Data ONTAP has a direct impact on the active data path when an SAP system is migrated to a different storage node.

### 5.3 Multiple SAP Systems Sharing One LIF

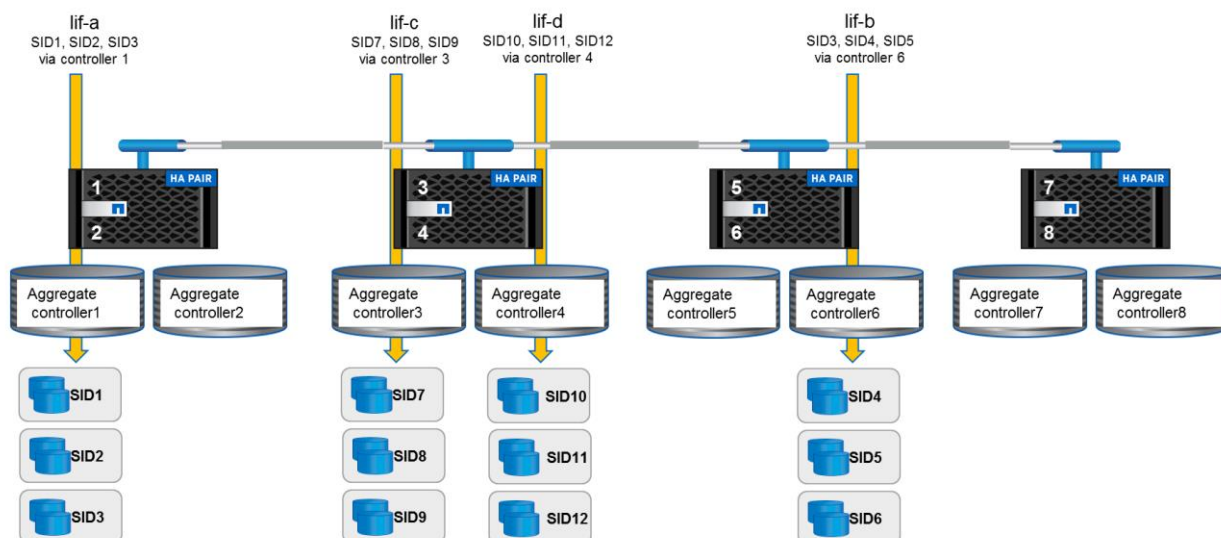
A configuration that is based on using the smallest number of LIFs means that there would be one LIF per storage node in the cluster. With this configuration, multiple SAP systems would share a common LIF. Figure 13 shows a configuration with one LIF per storage node. In this example, there are three SAP systems running per LIF across each of the first four storage nodes.

Figure 13) Configuration with one LIF per storage node.



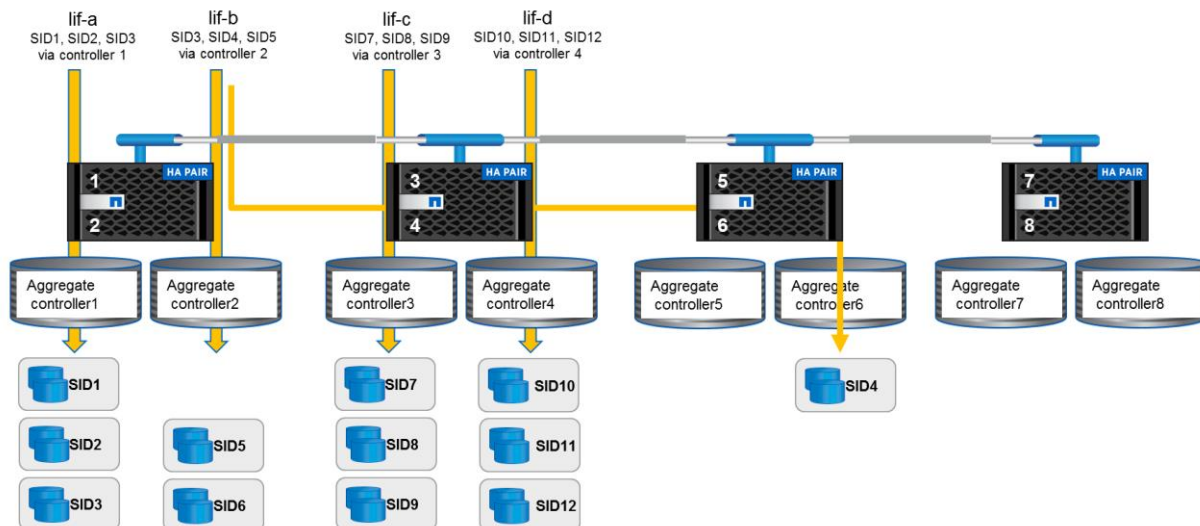
All SAP systems that share a common LIF can be migrated as one entity. The LIF is also migrated to the new storage node, as shown in Figure 14. In this example, no cluster interconnect traffic occurs.

Figure 14) Migration of multiple systems.



If a single SAP system is migrated, for example, SID4, as shown in Figure 15, the LIF cannot be migrated together with the SAP system because other SAP systems still run on the original storage node. Therefore, the SAP system SID4 is still accessed through lif-a on storage node 1, and the data is 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 15) 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.

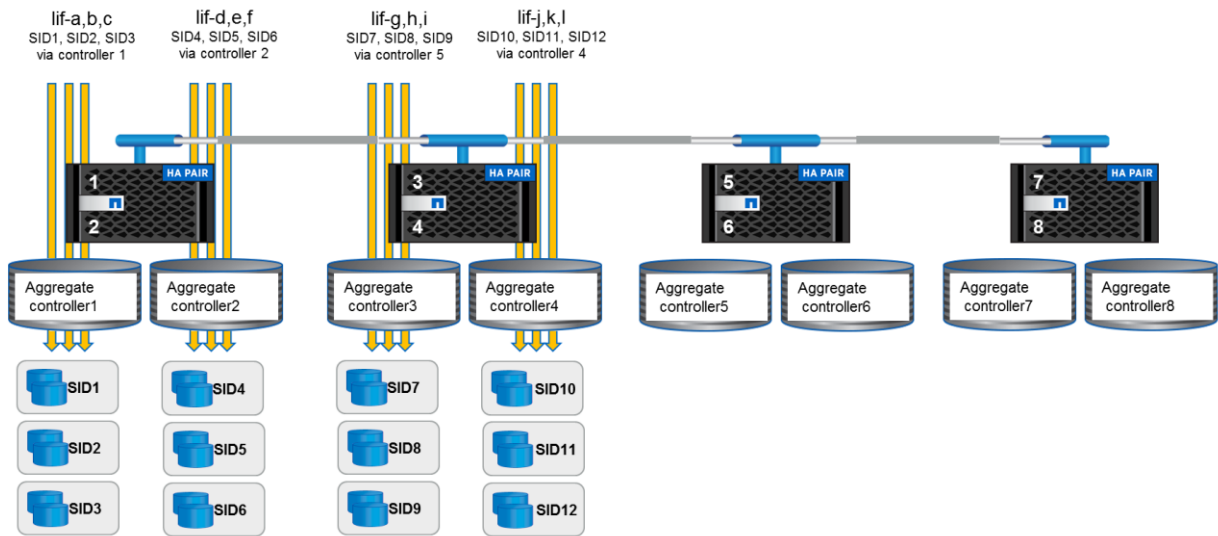
Sharing LIFs among multiple SAP systems can make sense for development and test systems that don't require scalability and nondisruptive operation, and cluster interconnect traffic is acceptable because of low throughput requirements.

## 5.4 One or Multiple LIFs per SAP System

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

Figure 16 shows a configuration with one LIF per SAP system.

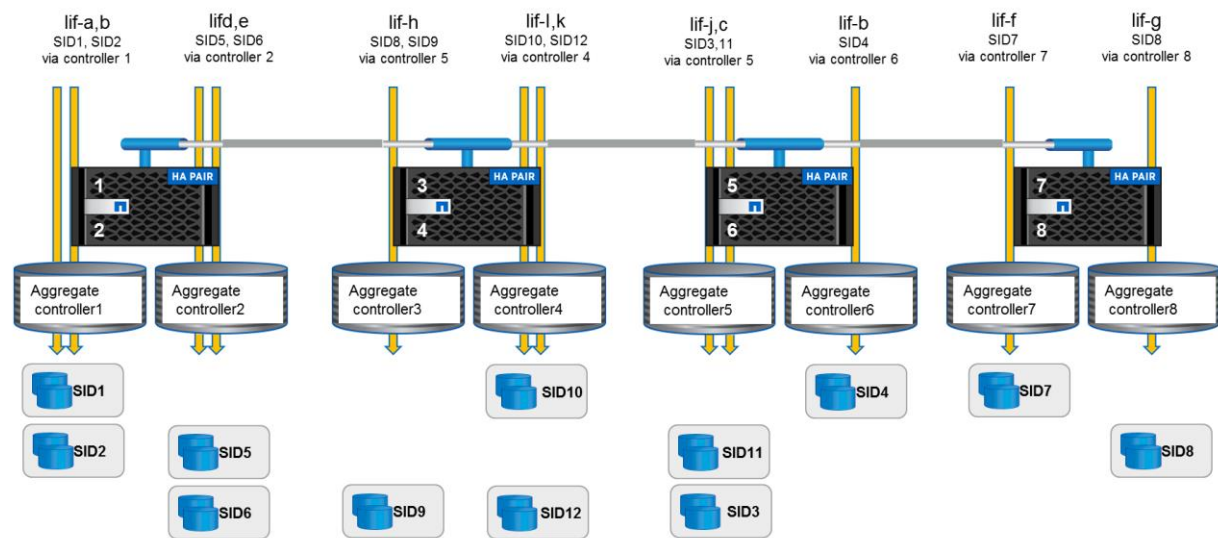
Figure 16) Configuration with one LIF per SAP system.



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

Figure 17 illustrates the migration of multiple single SAP systems.

Figure 17) Migration of multiple single SAP systems.



### 5.5 LIF Configuration Summary

The highest flexibility with 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 because of the number of needed IP addresses.

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

Table 1) LIF configuration summary.

System Characteristics and 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

## 6 Volume and LIF Configuration Examples

No single configuration fits all customer environments. The most logical configuration is based on customer-specific requirements. This section provides configuration examples.

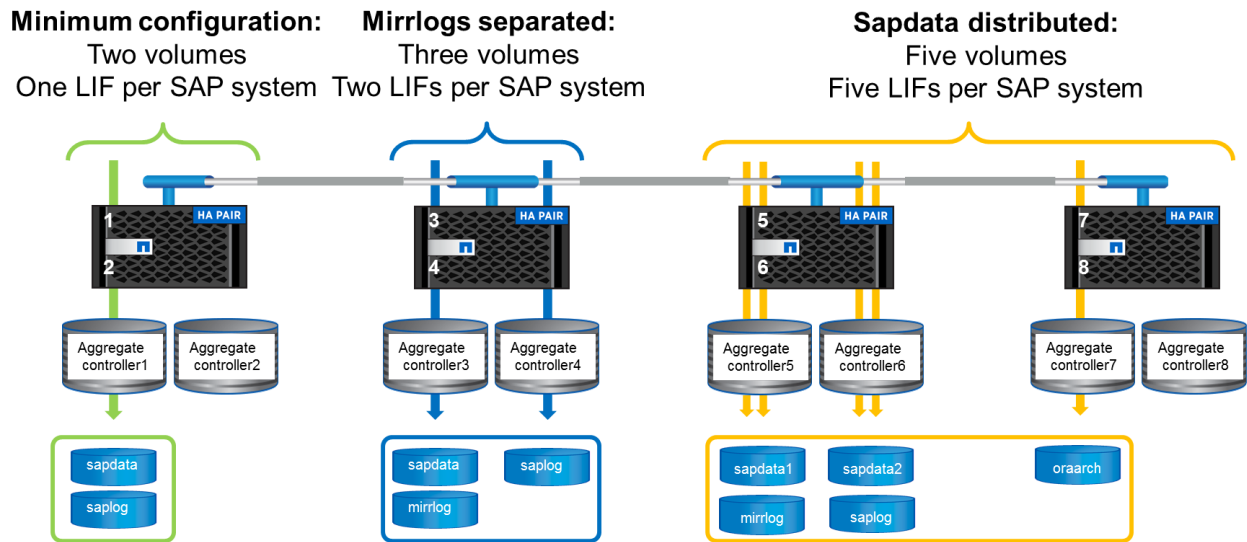
The following configuration examples are based on a setup with at least one LIF per SAP system. Multiple LIFs are required depending on the number of volumes and how these volumes should be distributed to aggregates and storage nodes:

- Example 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.
- Example 2: Three volumes, two-LIFs configuration to store Oracle mirrored redo logs on different hardware than the online redo logs.
- Example 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 that can be used for development and test as well as for production systems. Configuration example 3 can be used for larger production systems that require scalability and the capability to distribute load among multiple storage nodes.

Figure 18 shows the three configuration examples with different numbers of volumes and LIFs per SAP system.

Figure 18) Volume and LIF configuration examples.



6.1 Minimum Configuration

The first configuration in Figure 18 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 use of Snapshot copies, SnapRestore, FlexClone, and other Data ONTAP features that work on the volume level.

Table 2 lists the minimum configuration, LIF, and volume layout.

Table 2) Minimum configuration, LIF, and volume layout.

Aggregate on One Storage Node	
LIF 1	
Directories in data volume	Directories in log volume
/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
	/usr/sap/SID
	/sapmnt/SID

6.2 Separating Oracle Mirrored Redo Logs

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



setup is shown in the second configuration in Figure 18. 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 lists the LIF and volume layout for separated mirrored redo logs.

**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
Directories in data volume	Directories in mirrlog volume	Directories in log volume
/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
		/usr/sap/SID
		/sapmnt/SID

### 6.3 Large SAP Systems with High-Performance Requirements

SAP systems with very high throughput requirements should be distributed evenly across multiple storage nodes. It might 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 process could prevent costly downtime in the future as the production system's throughput requirements grow beyond the performance capabilities of a single storage node.

In the third configuration in Figure 18, 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 lists the LIF and volume layout for large SAP systems.

**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
Directories in data volume 1	Directories in data volume 2	Directories in log volume	Directories in mirrlog volume	Directories in oraarch volume
/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		

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
		/usr/sap/SID		
		/sapmnt/SID		

## 7 SAP System Installation


This section describes the requirements and the configuration for installing an SAP business suite or SAP NetWeaver system with Oracle Database under UNIX or Linux using the NFS protocol.

### 7.1 Storage Network

A dedicated, redundant 10 Gigabit Ethernet (10GbE) 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 requires two GbE cards connected to the switching infrastructure.

Figure 19 illustrates a dedicated 10GbE storage network.

Figure 19) Dedicated 10GbE storage network.



### 7.2 Operating System Configuration

Configuring the correct NFS mount options is important to provide optimal performance and system stability. The following common mount options are valid for each of the OS platforms.

- rw, bg, hard, vers=3, proto=tcp, timeo=600, rsize=65536, wsize=65536, nointr, and suid
- Additional mount options for the different OS platforms include:
  - Linux: <common>
  - Solaris: <common>, llock
  - HP/UX: <common>, forcedirectio
  - AIX: <common>

For additional information about OS-specific tuning, see [TR-3633: Oracle Databases on Data ONTAP](#).

**Note:** NetApp recommends that you turn off NFS server locking because if the Oracle database server crashes, you won't need to manually delete NFS locks at the storage system level. You can turn



off NFS server locking by using the `nolock` mount option for Linux or the `llock` mount option for AIX, HP/UX, and Solaris platforms.

## 7.3 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 first put into hot backup mode. Therefore, automatically scheduled Snapshot copies of the SAP data volumes at the storage level should be turned off.

**Note:** During the SAP installation, the visibility of the Snapshot directory has to be turned off for all volumes containing any of the file systems for the SAP system. Otherwise, the SAP installation tool (SoftWare Provisioning Manager [SWPM]) will try to change the permissions and the ownership on the Snapshot subdirectories. Because the Snapshot data is read-only, this change causes the SAP installation tools to fail and abort. After the installation of the SAP system is complete, the volume Snapshot option can be switched on again.

## 7.4 SAP Installation Process

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

Table 5 shows a typical example of a volume configuration.

Table 5) Volumes and mount points.

Volume	Subdirectory to Be Mounted	Mount Point at SAP System
sapdata_SID	sapdata1	/oracle/SID/sapdata1
	sapdata2	/oracle/SID/sapdata2
	sapdata3	/oracle/SID/sapdata3
	sapdata4	/oracle/SID/sapdata4
saplog_SID	oracle	/oracle
	sapusr_SID	/usr/sap/SID
	sapmnt_SID	/sapmnt/SID
	saptrans_SID	/usr/sap/trans
	saphome_SID	/home/sidadm
	orasid_home	/home/orasid
	oracle_home	/home/oracle

To set up the necessary file systems for the SAP installation, complete the following steps:

1. Create an 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:** The following examples use `SID = P02`.

3. Create the following directories at the SAP host:

```
sap-ora-p02:~ # mkdir -p /usr/sap/P02
sap-ora-p02:~ # mkdir -p /sapmnt/P02
sap-ora-p02:~ # mkdir /oracle
sap-ora-p02:~ # mkdir /home/p02adm
sap-ora-p02:~ # mkdir /home/oracle
sap-ora-p02:~ # mkdir /home/orap02
sap-ora-p02:~ # mkdir /usr/sap/trans
```

```

sap-ora-p02:~ # chmod -R 777 /usr/sap/P02
sap-ora-p02:~ # chmod -R 777 /sapmnt/P02
sap-ora-p02:~ # chmod -R 777 /oracle
sap-ora-p02:~ # chmod -R 777 /home/p02adm/
sap-ora-p02:~ # chmod -R 777 /home/oracle/
sap-ora-p02:~ # chmod -R 777 /home/orap02/
sap-ora-p02:~ # chmod -R 777 /usr/sap/trans

```

#### 4. Create the following directories within the saplog\_p02 volume:

```

sap-ora-p02:~ # mkdir /mnt/tmp1
sap-ora-p02:~ # mount 172.20.101.15:/saplog_p02 /mnt/tmp1

sap-ora-p02:/mnt/tmp1 # mkdir -p /mnt/tmp1/oracle/P02
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/oracle/P02/sapdata1
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/oracle/P02/sapdata2
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/oracle/P02/sapdata3
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/oracle/P02/sapdata4
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/sapusr_P02
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/sapmnt_P02
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/saptrans
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/saphome_P02
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/orap02_home
sap-ora-p02:/mnt/tmp1 # mkdir /mnt/tmp1/oracle_home

sap-ora-p02:/mnt/tmp1 # umount /mnt/tmp1

```

#### 5. Create the following directories within the sapdata\_p02 volume:

```

sap-ora-p02:~ # mkdir /mnt/tmp2
sap-ora-p02:~ # mount 172.20.101.15:/sapdata_p02 /mnt/tmp2

sap-ora-p02:/mnt/tmp2 # mkdir /mnt/tmp2/sapdata1
sap-ora-p02:/mnt/tmp2 # mkdir /mnt/tmp2/sapdata2
sap-ora-p02:/mnt/tmp2 # mkdir /mnt/tmp2/sapdata3
sap-ora-p02:/mnt/tmp2 # mkdir /mnt/tmp2/sapdata4

sap-ora-p02:~ # umount /mnt/tmp2

```

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

```

172.20.101.15:/saplog_p02/oracle /oracle nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/sapusr_P02 /usr/sap/P02 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/sapmnt_P02 /sapmnt/P02 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/saphome_P02 /home/p02adm nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/oracle_home /home/oracle nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/orap02_home /home/orap02 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/saplog_p02/saptrans /usr/sap/trans nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/sapdata_p02/sapdata1 /oracle/P02/sapdata1 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/sapdata_p02/sapdata2 /oracle/P02/sapdata2 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/sapdata_p02/sapdata3 /oracle/P02/sapdata3 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0
172.20.101.15:/sapdata_p02/sapdata4 /oracle/P02/sapdata4 nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,nolock 0 0

```

#### 7. Mount the new file systems by running the mount -a -t nfs command.

#### 8. Verify the file system structure.

```

sap-ora-p02:~ # df -a

172.20.101.15:/saplog_p02/sapusr_P02 - - - - /usr/sap/P02
172.20.101.15:/saplog_p02/oracle 199229440 22024000 177205440 12% /oracle
172.20.101.15:/saplog_p02/saphome_P02 - - - - /home/p02adm
172.20.101.15:/sapdata_p02/sapdata3 199229440 23485312 175744128 12% /oracle/P02/sapdata3

```

172.20.101.15:/saplog_p02/sapmnt_P02	-	-	-	-	/sapmnt/P02
172.20.101.15:/sapdata_p02/sapdata2	-	-	-	-	/oracle/P02/sapdata2
172.20.101.15:/saplog_p02/saptrans	-	-	-	-	/usr/sap/trans
172.20.101.15:/sapdata_p02/sapdata4	-	-	-	-	/oracle/P02/sapdata4
172.20.101.15:/sapdata_p02/sapdata1	-	-	-	-	/oracle/P02/sapdata1
172.20.101.15:/saplog_p02/oracle_home	-	-	-	-	/home/oracle
172.20.101.15:/saplog_p02/orap02_home	-	-	-	-	/home/orap02
sap-ora-p02:/home #					

The SAP installation tool (SWPM) fully supports NFS mounts. Therefore, the SAP installation can be completed as described in the corresponding version of the “SAP Installation Guide.”

**Note:** By default, the SAP installation tool stores one of the Oracle control files in `/oracle/SID/sapdata1`. If you use SMSAP, this control file must be stored outside the sapdata volume to allow fast volume restores with SMSAP. You can place the Oracle control file in a different volume either during the installation process with SWPM or when the installation process is complete, as described in the “Relocate Control Files in Sapdata1 File System” section.

## 8 SnapManager for SAP 3.4 Setup and Configuration

The configuration guidelines in this section are based on the following software versions:

- SMSAP 3.4
- SDU 5.3P2

### 8.1 SAP Systems and SMSAP Lab Setup

Figure 20 shows the lab setup that was used as the basis for this document.

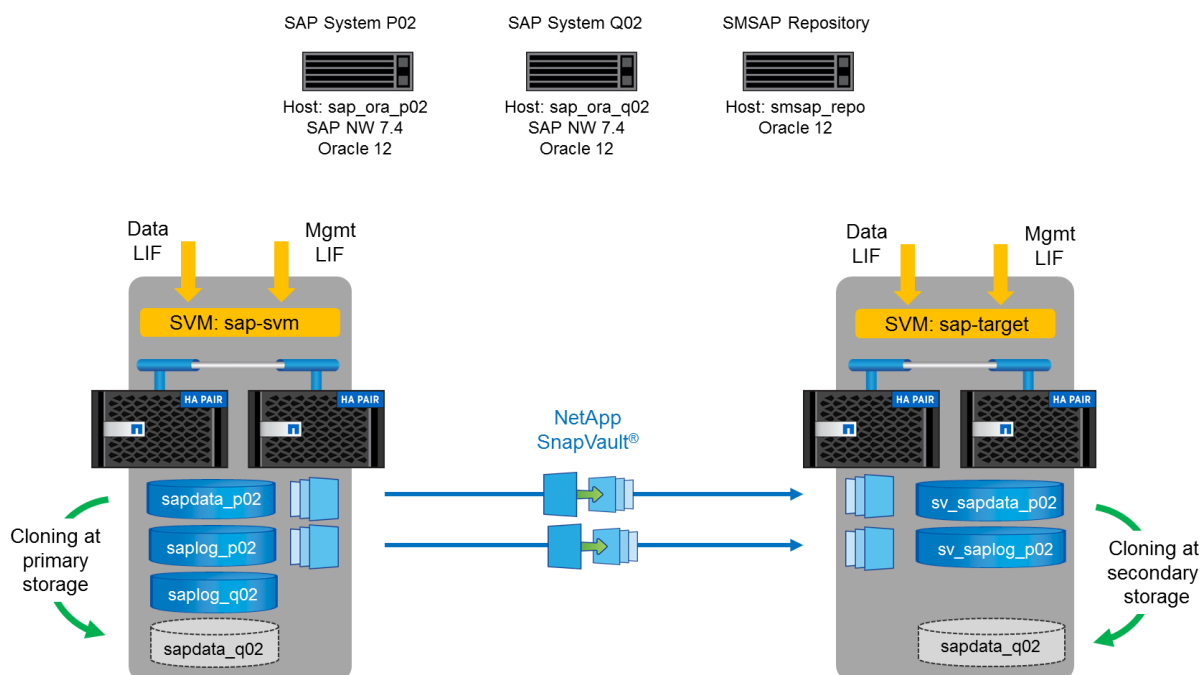
Two SAP systems (NetWeaver 7.4 with Oracle 12 on SuSE Linux) were installed using SAP Software Provisioning Manager (SWPM). The system Q02 was used as the system copy target for the data from system P02.

Both SAP systems were connected through NFS to the SVM `sap-svm`. The file systems were mounted using the data LIF of the `sap-svm` SVM. Each SAP system used a two-volume configuration at the storage level: one sapdata and one saplog volume.

For the SAP system copy use case, a clone of the source data volume of P02 can either be created on the primary storage or on the secondary storage.

A second SVM `sap-target` was used to configure data protection with SnapVault software.

**Figure 20) SMSAP lab setup.**



## 8.2 Configure SMSAP Repository

Within the lab setup, an Oracle 12 database was installed with SID=REP, which was used for the SMSAP repository.

To configure the SMSAP repository, complete the following steps:

1. Create a tablespace, which is used to access the repository database.

```
SQL> create tablespace repdata datafile '/oracle/REP/oradata/repdata.dbf' size 1000m autoextend on;
```

Tablespace created.

2. Create a user within the database, which is used to access the repository database.

```
SQL> create user smrepo identified by "Netappl23" default tablespace repdata;
```

User created.

```
SQL> grant connect, resource to smrepo;
```

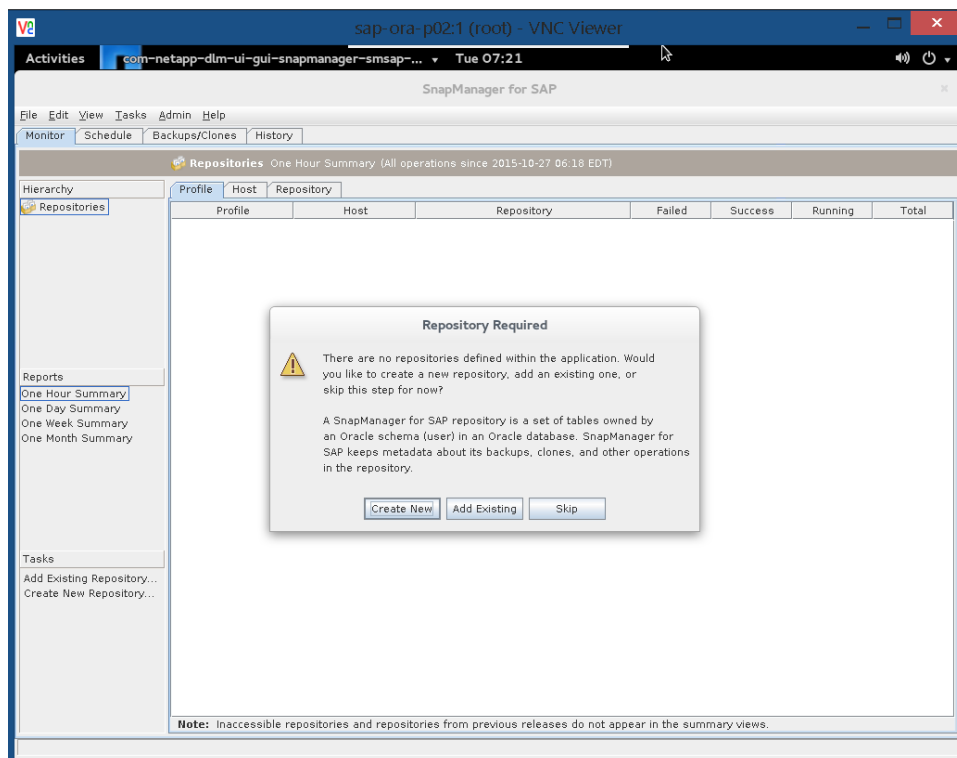
Grant succeeded.

```
SQL> grant unlimited tablespace to smrepo;
```

Grant succeeded.

```
SQL>
```

3. Launch the SMSAP GUI and create the SMSAP repository by using the SnapManager for SAP Repository wizard.



4. On the SnapManager for SAP Repository Wizard welcome page, click Next.



5. Enter the repository database configuration information and click Next.

**SnapManager for SAP Repository Wizard**

**Repository Database Configuration Information**

User Name	smrepo
Password	*****
Host	smsap-repo
Port	1521
Service Name	REP

Buttons: ? Help, < Back, Next >, Cancel

6. Review the configuration summary and click Create.

**SnapManager for SAP Repository Wizard**

**Perform Repository Create Operation**

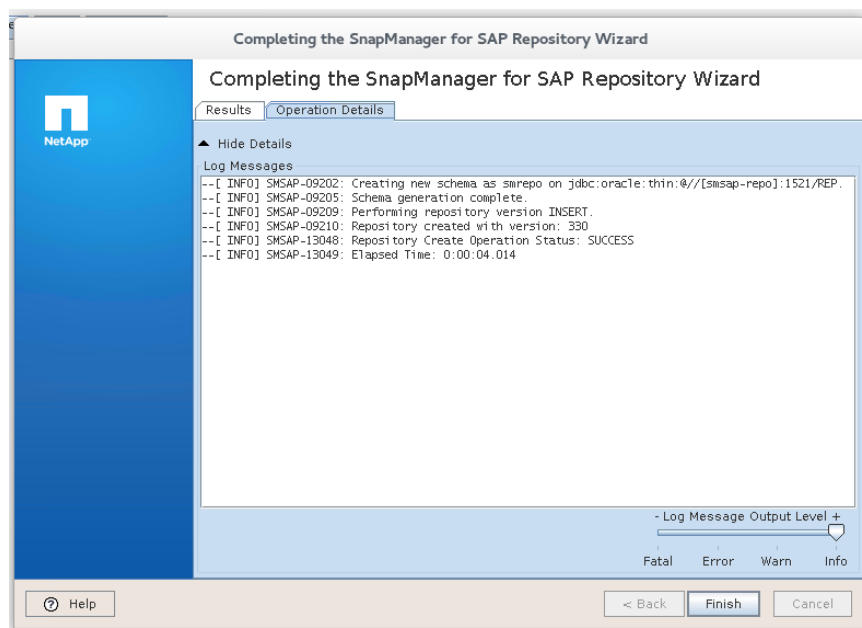
**Below is a summary of your operation configuration.**

Creating repository.

Repository Name: smrepo@REP/smsap-repo:1521  
 User Name: smrepo  
 Password: \*\*\*\*\*  
 Host: smsap-repo  
 Port: 1521  
 Service Name: REP

Buttons: ? Help, < Back, Create, Cancel

7. After the repository database is created, click Finish.



## 9 Back Up Using SnapManager for SAP

### 9.1 Install and Configure SAP Source System P02

SMSAP and SnapDrive for UNIX (SDU) must be installed on the host.

#### Configure Pluggable Authentications Modules

To configure Pluggable Authentication Modules (PAMs) authentication for SMSAP, complete the following step:

1. Create the file `snapmanager` in the `/etc/pam.d` directory using the following configurations:

```
sap-ora-p02:/ # cat /etc/pam.d/snapmanager
#%PAM-1.0
auth required pam_unix.so
account required pam_unix.so
```

#### Configure SnapDrive for UNIX

To configure SDU, complete the following steps:

1. The host name of the interface that is used for management access must match the name of the SVM. In a lab setup, the file systems are mounted using `sap-svm-data-path` and the SVM is managed using `sap-svm`. `Sap-target` is the management interface of the SVM that is used as a target for the SnapVault replication. `sap-target-data-path` is the data LIF of the target SVM.

```
sap-ora-p02:/ # cat /etc/hosts
172.20.101.14    sap-svm
172.20.101.15    sap-svm-data-path
172.20.101.44    sap-target
172.20.101.45    sap-target-data-path
```

2. Configure the required SDU credentials for the management interfaces by running the following commands:

```
sap-ora-p02:/ # snapdrive config set vsadmin sap-svm
Password for vsadmin:
Retype password:

sap-ora-p02:/ # snapdrive config set vsadmin sap-target
Password for vsadmin:
```

```
Retype password:
```

### 3. Add the data path interface to the management interface:

```
sap-ora-p02:/ # snapdrive config set -mgmtpath sap-svm sap-svm-data-path
sap-ora-p02:/ # snapdrive config set -mgmtpath sap-target sap-target-data-path
```

### 4. Verify the configuration by running the following commands:

```
sap-ora-p02:~ # snapdrive config list
username      appliance name  appliance type
-----
vsadmin       sap-target      StorageSystem
vsadmin       sap-svm         StorageSystem
sap-ora-p02:~ # snapdrive config list -mgmtpath
system name    management interface  datapath interface
-----
sap-svm        172.20.101.14          172.20.101.15
sap-target     172.20.101.44          172.20.101.45
```

### 5. Test access of SDU to the primary SVM sap-svm by running the following commands:

```
sap-ora-p02:/var/log/smsap # snapdrive storage list -all

WARNING: This operation can take several minutes
         based on the configuration.
NFS device: 172.20.101.15:/saplog_p02/saptrans mount point: /usr/sap/trans (persistent)
NFS device: 172.20.101.15:/inst_media mount point: /mnt/inst_media (persistent)
NFS device: 172.20.101.15:/saplog_p02/sapmnt_P02 mount point: /sapmnt/P02 (persistent)
NFS device: 172.20.101.15:/saplog_p02/orap02_home mount point: /home/orap02 (persistent)
NFS device: 172.20.101.15:/saplog_p02/saphome_P02 mount point: /home/p02adm (persistent)
NFS device: 172.20.101.15:/saplog_p02/sapusr_P02 mount point: /usr/sap/P02 (persistent)
NFS device: 172.20.101.15:/saplog_p02/oracle_home mount point: /home/oracle (persistent)
NFS device: 172.20.101.15:/saplog_p02/oracle mount point: /oracle (persistent)
NFS device: 172.20.101.15:/sapdata_p02/sapdata3 mount point: /oracle/P02/sapdata3 (persistent)
NFS device: 172.20.101.15:/sapdata_p02/sapdata1 mount point: /oracle/P02/sapdata1 (persistent)
NFS device: 172.20.101.15:/sapdata_p02/sapdata4 mount point: /oracle/P02/sapdata4 (persistent)
NFS device: 172.20.101.15:/sapdata_p02/sapdata2 mount point: /oracle/P02/sapdata2 (persistent)
sap-ora-p02:/var/log/smsap #
```

### 6. Test access to the secondary storage sap-target by running the following command:

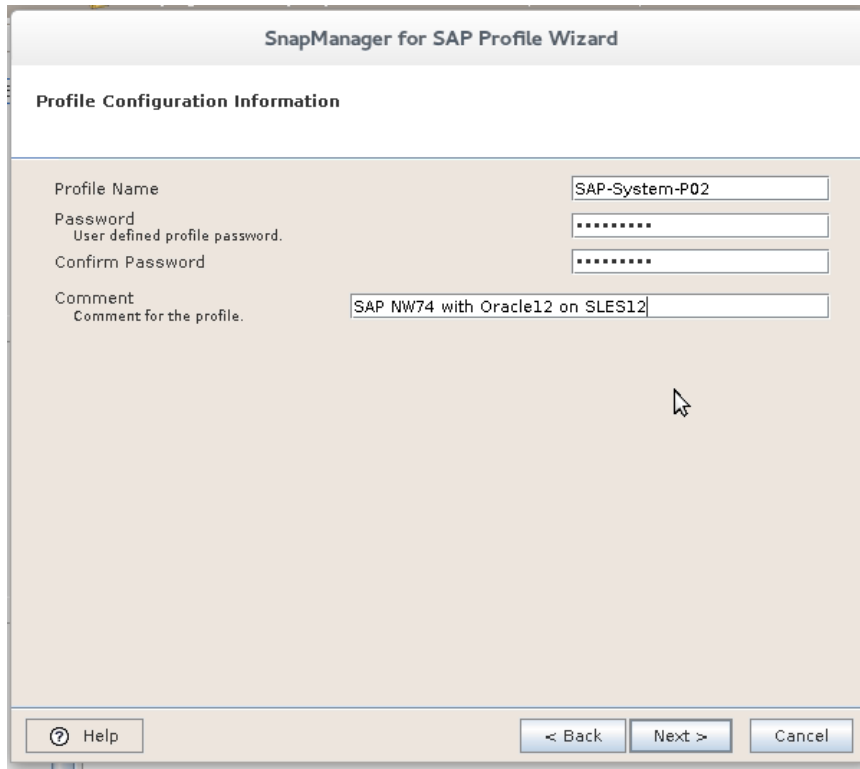
```
sap-ora-p02:/var/log/smsap # snapdrive snap list -filer sap-target
```

## Create SMSAP Profile

To create the SMSAP profile, complete the following steps by using the SnapManager for SAP Profile wizard:

1. Select the repository and then select Create Profile from the resulting menu.
2. Enter a name, password, and comment for the profile. Click Next.





**SnapManager for SAP Profile Wizard**

**Profile Configuration Information**

Profile Name: SAP-System-P02

Password:  (User defined profile password.)

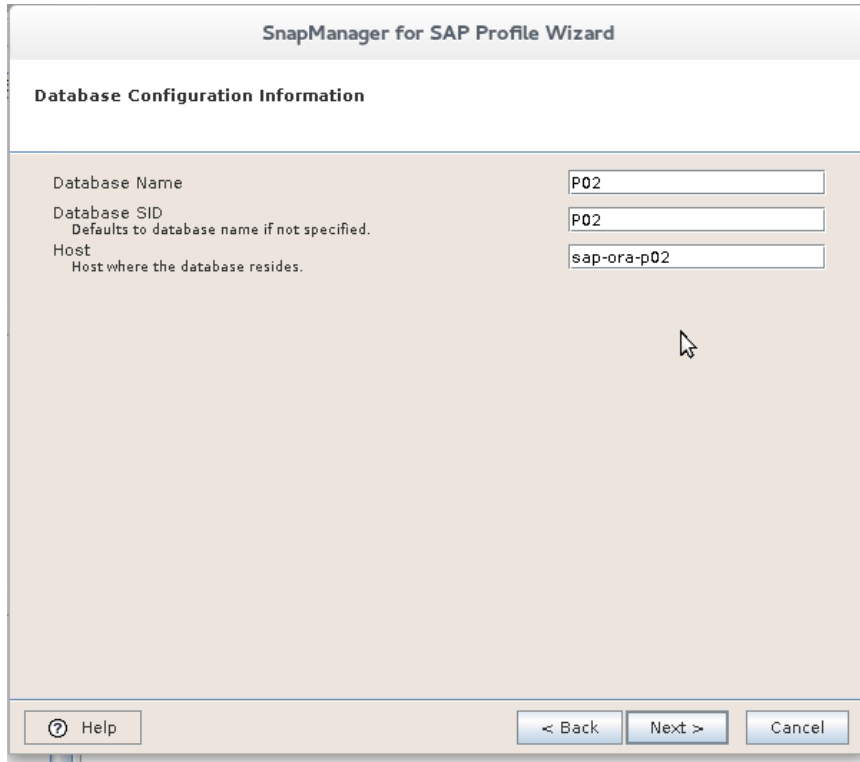
Confirm Password:

Comment: SAP NW74 with Oracle12 on SLES12 (Comment for the profile.)

Buttons: ? Help, < Back, Next >, Cancel

3. Enter the database name, database SID, and host name. Click Next.

**Note:** Make sure that the host name is resolved correctly within your DNS configuration.



**SnapManager for SAP Profile Wizard**

**Database Configuration Information**

Database Name: P02

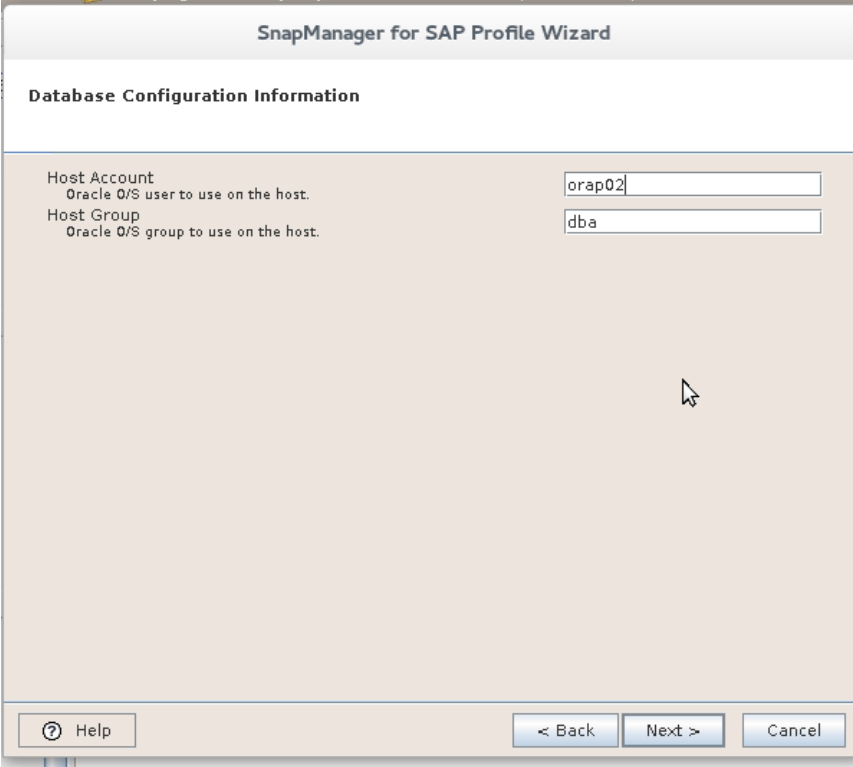
Database SID: P02 (Defaults to database name if not specified.)

Host: sap-ora-p02 (Host where the database resides.)

Buttons: ? Help, < Back, Next >, Cancel

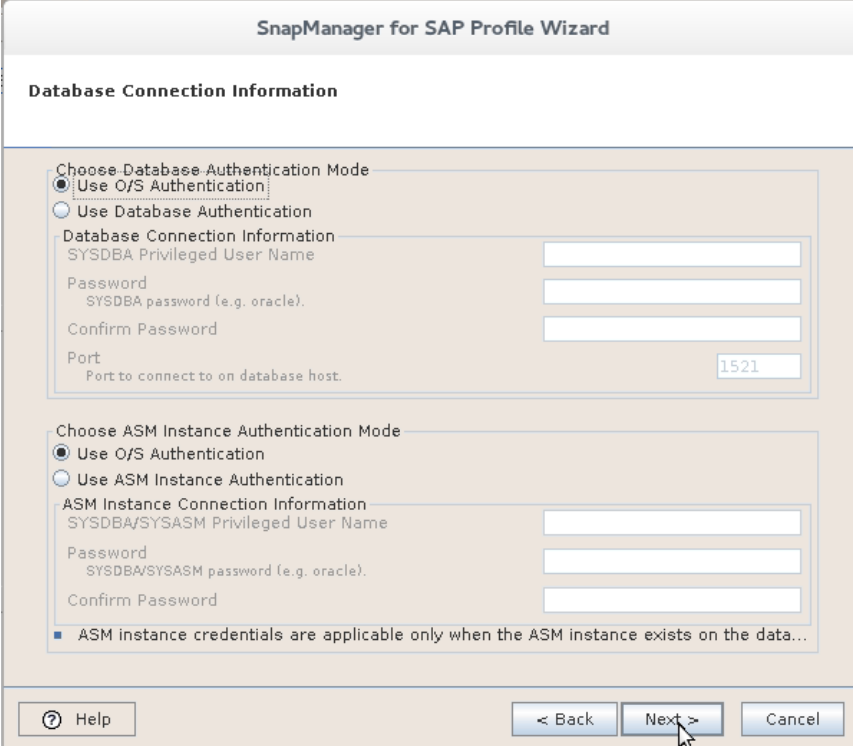
4. Enter the user and group information, which will be used for backup and restore/recovery operations. Click Next.

**Note:** With SAP on Oracle 12, either the user `oraSID` (group `dba`) or the user `oracle` (group `oinstall`) can be used to access the database.



The screenshot shows the 'Database Configuration Information' screen of the SnapManager for SAP Profile Wizard. It contains two input fields: 'Host Account' with the value 'orap02' and 'Host Group' with the value 'dba'. Below the fields are navigation buttons: '< Back', 'Next >', and 'Cancel'. A mouse cursor is visible over the 'Next >' button.

5. Select the Use O/S Authentication option for both modes and click Next.



The screenshot shows the 'Database Connection Information' screen of the SnapManager for SAP Profile Wizard. It has two sections. The first section, 'Choose Database Authentication Mode', has the 'Use O/S Authentication' radio button selected. The second section, 'Choose ASM Instance Authentication Mode', also has the 'Use O/S Authentication' radio button selected. Below these are input fields for 'SYSDBA Privileged User Name', 'Password', 'Confirm Password', and 'Port' (set to 1521). At the bottom, there is a checkbox labeled 'ASM instance credentials are applicable only when the ASM instance exists on the data...'. Navigation buttons '< Back', 'Next >', and 'Cancel' are at the bottom, with a mouse cursor clicking on 'Next >'.

6. Select Do Not Use RMAN and click Next.

**SnapManager for SAP Profile Wizard**

**RMAN Configuration Information**

☒ Do not use RMAN  
☐ Use RMAN via Control File  
☐ Use RMAN via Recovery Catalog

**RMAN Credentials**

User Name   
RMAN user name.  
 Password   
RMAN password.  
 Confirm Password   
 TNS Name

7. Choose the Snapshot naming convention and click Next.

To change the naming convention used for this profile's snapshots, select a variable in the table and then select "Add", "Remove", "Left", or "Right", or edit the text directly.

Variable Token	Description
{smid}	SnapManager uniq...
{profile}	profile name
{db-name}	database name
{db-sid}	database sid
{db-host}	fully qualified data...
{date}	backup date (yyy...

**Format:**

**Example:** smsap\_sap-system-p02\_p02\_f\_h\_1\_8abc01e915a55ac50115a55acc8d0001\_0

8. Configure the required retention policy settings. From the Protection Policy drop-down menu, select None. Click Next.

**SnapManager for SAP Profile Wizard**

**Policy Settings**

---

**Retention and Protection Policy**  
Specify the retention settings for hourly, daily, weekly, and monthly backups.

Retention Class	Count	Duration	
Hourly	4	0	hour(s)
Daily	5	7	day(s)
Weekly	0	4	week(s)
Monthly	0	12	month(s)

Protection Policy: None

☐ Backup Archivelogs Separately

---

**Archivelog Retention and Protection Policy**  
Primary Retention Duration:  hours

Protection Policy: None

☒ Always backup archivelogs along with online backups to enable cloning

9. Configure an e-mail notification, if needed. Click Next.

**SnapManager for SAP Profile Wizard**

**Configure Notification Settings**

---

☐ Add to Summary Notification

☐ Enable Profile Notification

**Profile Notification Settings**

☐ Successful Operation

E-mail ID(s)  
Comma-separated e-mail ID(s):

Subject  
Use "Advanced" button to configure the Subject:  Advanced...

☐ Failed Operation

E-mail ID(s)  
Comma-separated e-mail ID(s):

Subject  
Use "Advanced" button to configure the Subject:  Advanced...

**Note : Configure the Notification Settings to enable**

10. Configure retention for operation history and click Next.

**SnapManager for SAP Profile Wizard**

**History Configuration Information**

Select Operation to Configure retention

☒ Do Not Retain History

☐ Retain History For All Operations      Retention Period  Counts ▼

☐ Retain History For Selective Operations

Retention Configuration

Operation Type	Retention Period
<input type="checkbox"/> Backup	<input type="text"/> Counts ▼
<input type="checkbox"/> Backup Verify	<input type="text"/> Counts ▼
<input type="checkbox"/> Clone Create	<input type="text"/> Counts ▼
<input type="checkbox"/> Clone Split	<input type="text"/> Counts ▼
<input type="checkbox"/> Restore	<input type="text"/> Counts ▼

Help      < Back      Next >      Cancel

11. Review the configuration summary and click Create.

**SnapManager for SAP Profile Wizard**

**Perform Profile Create Operation**

**Below is a summary of your operation configuration.**

Create a new profile in repository "smrepo@REP/smsap-repo:1521".

**Profile Configuration Information**  
 Profile Name: SAP-System-P02  
 Password: \*\*\*\*\*  
 Comment: SAP NW74 with Oracle12 on SLES12  
 Host Account: orap02  
 Host Group: dba  
 Snapname Pattern: smsap\_{profile}\_{db-sid}\_{scope}\_{mode}\_{smid}

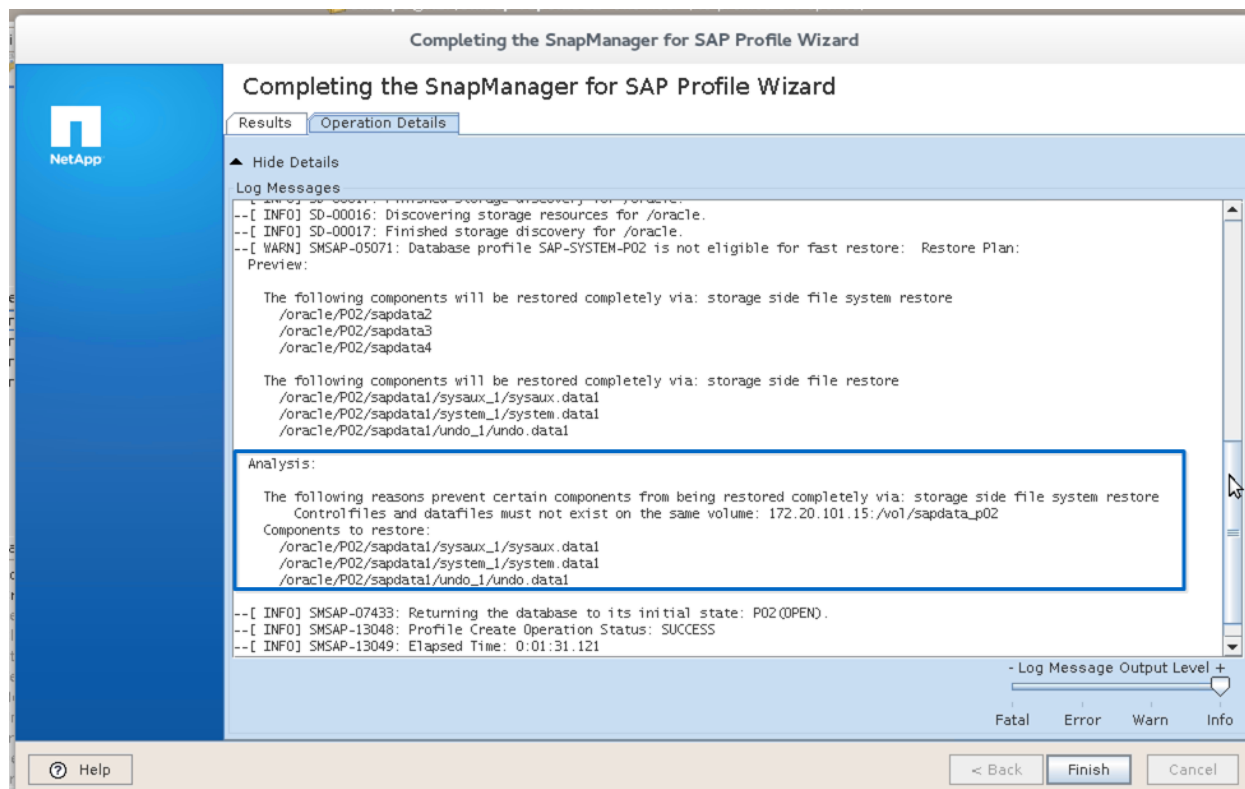
**Database Configuration Information**  
 Database Name: P02  
 Database SID: P02  
 Host: sap-ora-p02  
 Database Authentication Type: O/S Authentication  
 ASM Instance Authentication Type: O/S Authentication

**Retention Settings**

Help      < Back      Create      Cancel

12. In this example, the log output shows that a volume-based restore operation would not be possible because one of the Oracle control files is stored within the `sapdata1` file system. If a restore operation of the Oracle data files would be executed using volume-based restore, the control file would be overwritten with an older version. A restore operation would still be possible using a file-based restore.

If a volume-based restore operation is possible, the control file must be relocated.



## Relocate Control Files in Sapdata1 File System

To relocate the control files in the `sapdata1` file system, complete the following steps:

1. For the standard SAP on Oracle installation, the control files are located in the following directories:

```
SQL> select name from v$controlfile;
```

NAME

```
-----
/oracle/P02/origlogA/cntrl/cntrlP02.dbf
/oracle/P02/origlogB/cntrl/cntrlP02.dbf
/oracle/P02/sapdata1/cntrl/cntrlP02.dbf
SQL>
```

2. Relocate the control file within the `sapdata1` file system.

```
SQL> alter system set control_files='/oracle/P02/origlogA/cntrl/cntrlP02.dbf',
2  '/oracle/P02/origlogB/cntrl/cntrlP02.dbf', '/oracle/P02/mirrlogA/cntrl/cntrlP02.dbf'
scope=spfile;
```

System altered.

3. Shut down the database.

```
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
```

4. Relocate the control file.

```
sap-ora-p02:/ # mv /oracle/P02/sapdata1/cntrl/cntrlP02.dbf
/oracle/P02/mirrlogA/cntrl/cntrlP02.dbf
```

5. Start the database.

```
SQL> startup
ORACLE instance started.
```

```
Total System Global Area 4563402752 bytes
Fixed Size 2933352 bytes
```

Variable Size	2281704856 bytes
Database Buffers	2264924160 bytes
Redo Buffers	13840384 bytes
Database mounted.	
Database opened.	

## 6. Verify the location of the control files.

```
SQL> select name from v$controlfile;
```

```
NAME
```

```
-----  
/oracle/P02/origlogA/cntrl/cntrlP02.dbf  
/oracle/P02/origlogB/cntrl/cntrlP02.dbf  
/oracle/P02/mirrlogA/cntrl/cntrlP02.dbf
```

## Create Database Backup Without Data Protection to Secondary Storage

To create a database backup without data protection to secondary storage, complete the following steps by using the SnapManager for SAP Backup wizard:

1. On the SnapManager for SAP Backup Wizard welcome page, click Next.



2. Specify a label and click Next

**SnapManager for SAP Backup Wizard**

**SnapManager for SAP Backup Wizard**

Label  
Defaults to unique label if not specified.

Comment  
Optional.

☐ Allow startup or shutdown of database if necessary

3. Specify a backup operation parameter and click Next.

**SnapManager for SAP Backup Wizard**

**Database, Tablespaces or Datafiles to Backup**

☒ Full Backup  
☐ Partial Backup  
☐ Prune Archivelogs  
☐ Protect the backup  
☐ Protect now

**Datafile Backup Options**  
Select either tablespaces or datafiles to backup.

P02

Tablespaces

- PSAPSR3740
- PSAPSR3
- PSAPSR3USR
- SYSTEM
- PSAPUNDO
- SYSAUX

Datafiles

Type

Retention Class

☐ Verify backup using the Oracle DBVERIFY utility

☐ Allow startup or shutdown of database if necessary

4. Select the archive log destination backup option and click Next.



## SnapManager for SAP Backup Wizard

### ArchiveLog Destination

Choose Archive log Destinations

Sl.No	Destination	Binding	Target	Backup	Prune
1	/oracle/P02/oraarch/P02...	OPTIONAL	PRIMARY	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Help < Back Next > Cancel

5. Select No to disable the pre- and post-backup tasks and click Next.

## SnapManager for SAP Backup Wizard

### Task Enabling Page

Do you wish to enable Pre/Post Tasks?

☒ No

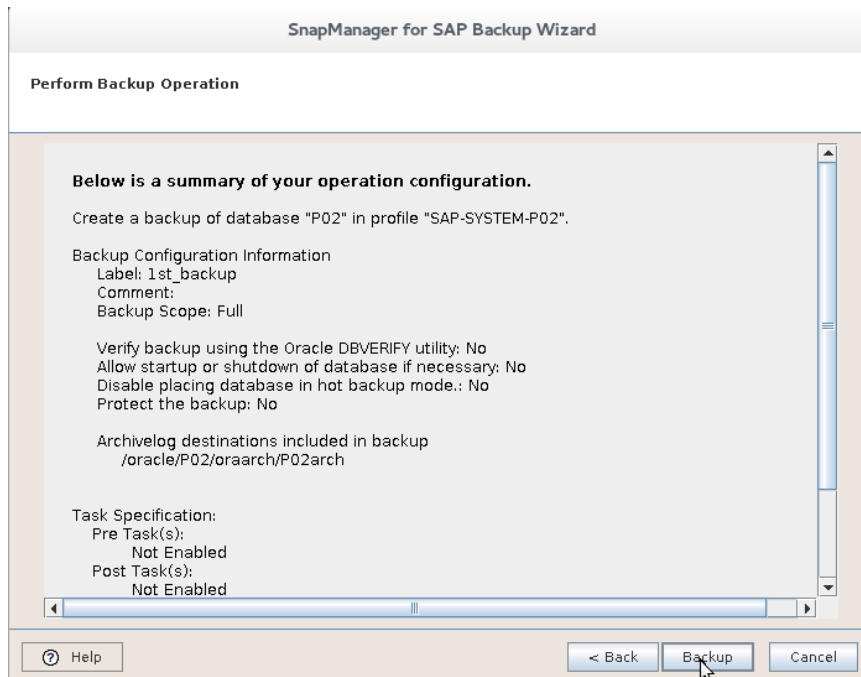
☐ Yes. Automatically initialize the task specification xml file.

☐ Yes. Load an existing task specification xml file

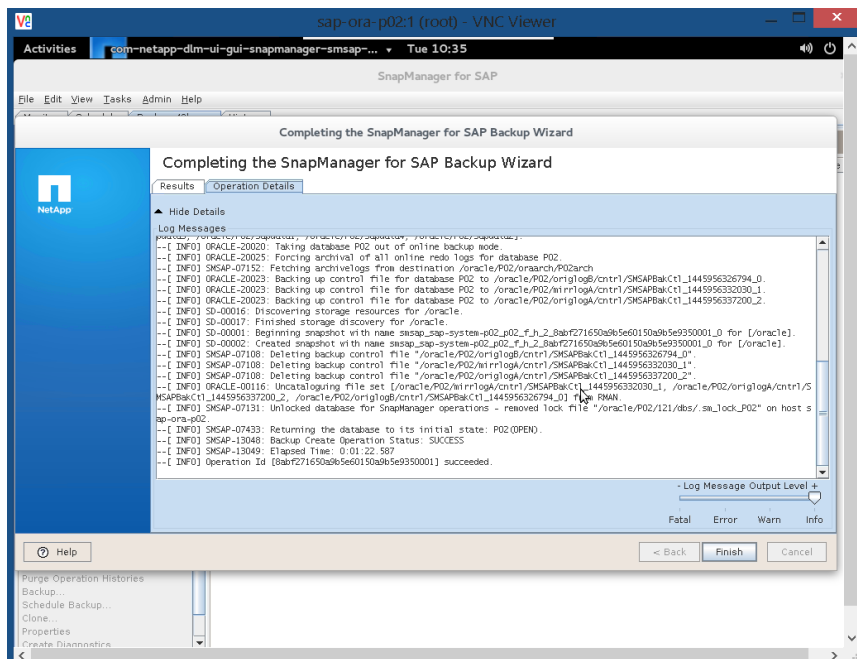
Load

Help < Back Next > Cancel

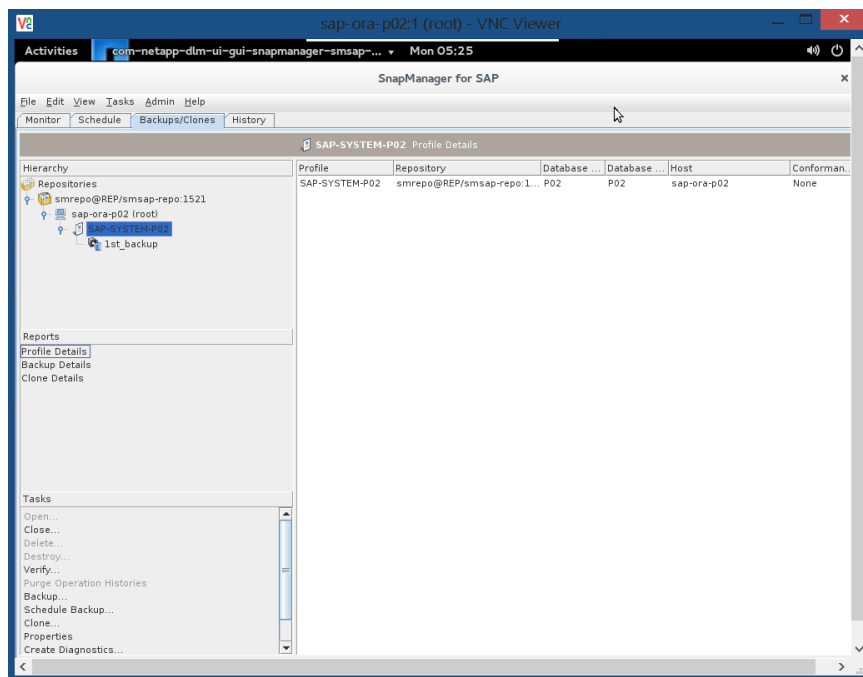
6. Review the configuration summary and click Backup.



7. After the database backup is complete, click Finish.



8. Verify that the backup that was created is now listed in the SMSAP GUI.



## 9.2 Configure Data Protection to Secondary Storage System

To configure data protection to a secondary storage system, the following tasks must be completed:

**Note:** Detailed steps are provided in the sections that follow.

1. Peer the source and target SVMs or clusters to allow SnapVault replication.
2. Create the volumes on the target SVM.
3. Configure the SnapVault replication:
  - a. Create a SnapMirror policy and rule that define the SnapMirror label and backup retention.
  - b. Configure the SnapMirror relationship between source and target volumes.
  - c. Initialize the SnapMirror relationship.
4. Configure SnapDrive.
5. Add data protection to the existing SMSAP profile.
6. Create a protected backup.

**Note:** The retention of protected backups is not managed by SMSAP. Retention is controlled by Data ONTAP using the retention defined within the SnapMirror policy.

### Create Volumes at Target SVM

To create volumes at the target SVM, complete the following step:

1. In the lab setup, the SAP system P02 was stored on two volumes: `sapdata_p02` and `saplog_p02`. For both of these volumes, create a target volume on the target SVM.

```
FAS8040SAP:> volume create -vserver sap-target -volume sv_sapdata_p01 -aggregate data_02 -size
200g -state online -type DP
[Job 3190] Job succeeded: Successful

FAS8040SAP:>
FAS8040SAP:> volume create -vserver sap-target -volume sv_saplog_p01 -aggregate data_02 -size
200g -state online -type DP
[Job 3191] Job succeeded: Successful
```

### Configure SnapVault Replication

To configure SnapVault replication, complete the following steps:

### 1. Create a SnapMirror policy.

```
FAS8040SAP:> snapmirror policy create -vserver sap-target -policy SV4_SMSAP
```

### 2. Add a rule that defines the SnapMirror label and the backup retention policy.

```
FAS8040SAP:> snapmirror policy add-rule -vserver sap-target -policy SV4_SMSAP -snapmirror-label P02_SMSAP -keep 14
```

### 3. Create and initialize the SnapMirror relationship for both volumes (data and log).

```
FAS8040SAP:> snapmirror create -source-path sap-svm:sapdata_p02 -destination-path sap-  
target:sv_sapdata_p01 -type XDP -policy SV4_SMSAP  
Operation succeeded: snapmirror create for the relationship with destination "sap-  
target:sv_sapdata_p01".
```

```
FAS8040SAP:> snapmirror initialize -destination-path sap-target:sv_sapdata_p01 -type XDP  
Operation is queued: snapmirror initialize of destination "sap-target:sv_sapdata_p01".
```

```
FAS8040SAP:> snapmirror create -source-path sap-svm:saplog_p02 -destination-path sap-  
target:sv_saplog_p01 -type XDP -policy SV4_SMSAP  
Operation succeeded: snapmirror create for the relationship with destination "sap-  
target:sv_saplog_p01".
```

```
FAS8040SAP:> snapmirror initialize -destination-path sap-target:sv_saplog_p01 -type XDP  
Operation is queued: snapmirror initialize of destination "sap-target:sv_saplog_p01".
```

### 4. Monitor the progress of the SnapMirror initialization by running the snapmirror show command.

```
FAS8040SAP:> snapmirror show
```

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Progress Healthy	Last Updated
sap-svm:sapdata_p02	XDP	sap-target:sv_sapdata_p01	Snapmirrored	Idle	-	true	-
sap-svm:saplog_p02	XDP	sap-target:sv_saplog_p01	Snapmirrored	Idle	-	true	-

2 entries were displayed.

```
FAS8040SAP:>
```

## Configure SnapDrive for UNIX

To configure SDU, complete the following steps:

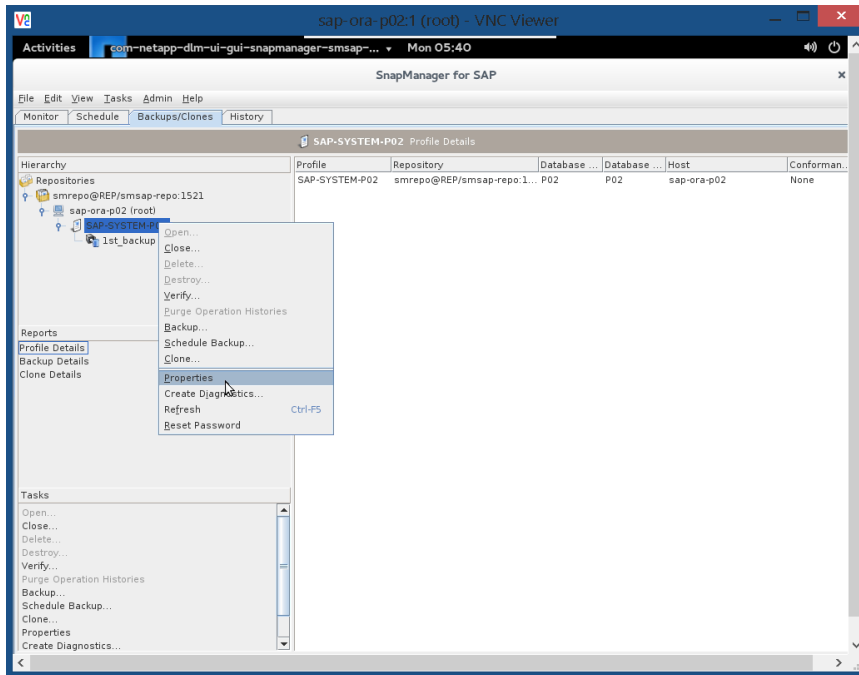
1. Configure SDU using the steps described in the section “Install and Configure SAP Source System P02.”
2. Verify the SDU configuration by running the following commands:

```
sap-ora-p02:~ # snapdrive protection show sap-svm:/vol/sapdata_p02  
sap-svm:/vol/sapdata_p02 sap-target:/vol/sv_sapdata_p01 snapmirrored,vault  
sap-ora-p02:~ #  
sap-ora-p02:~ # snapdrive protection show sap-svm:/vol/saplog_p02  
sap-svm:/vol/saplog_p02 sap-target:/vol/sv_saplog_p01 snapmirrored,vault  
sap-ora-p02:~ #
```

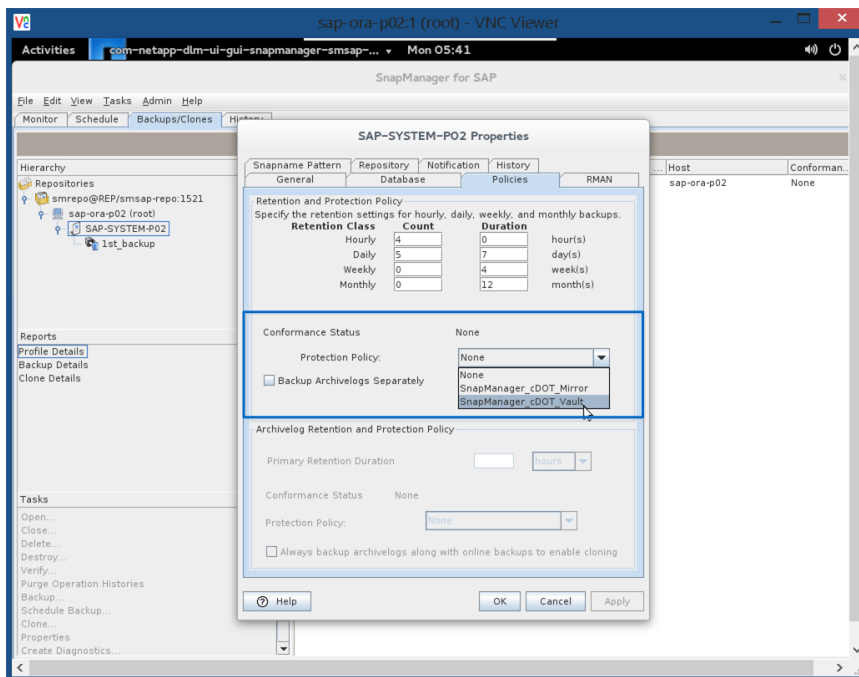
## Add Data Protection to Existing SMSAP Profile

To add data protection to the existing SMSAP profile, complete the following steps:

1. Right-click the SMSAP profile and select Properties.



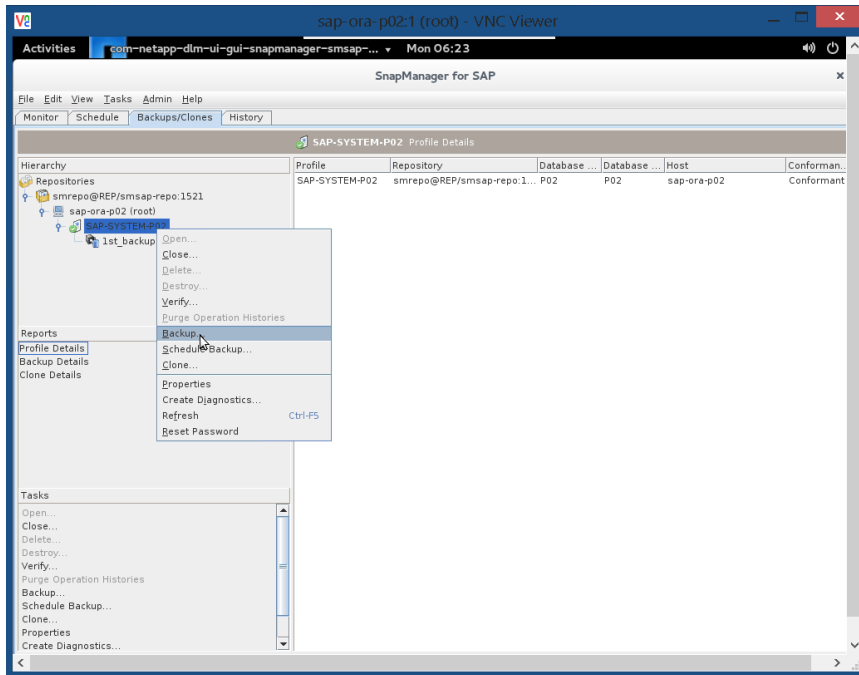
2. From the Protection Policy drop-down menu, select SnapManager-cDOT\_Vault. Click OK.



## Create Protected Backup

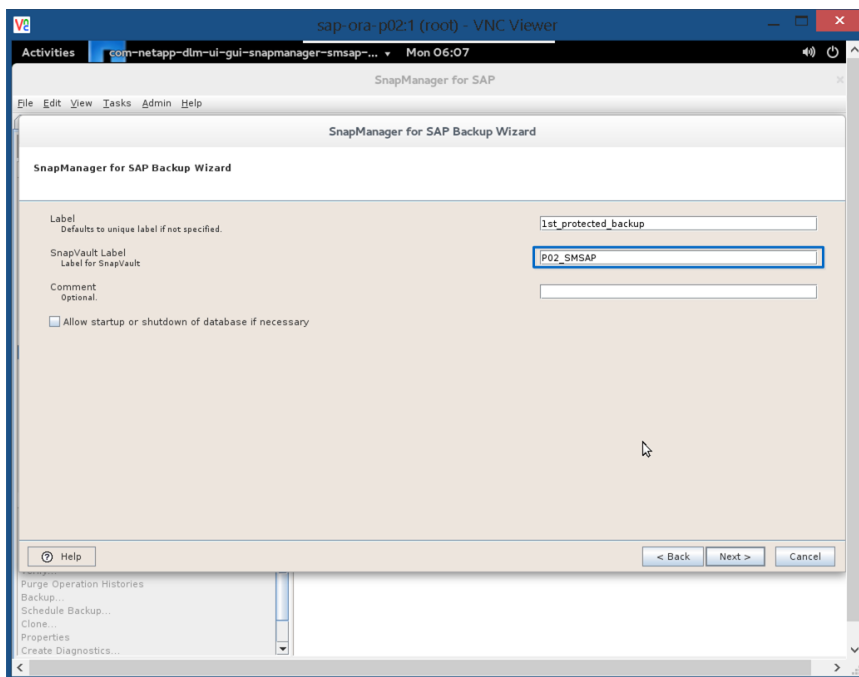
To create a protected backup, complete the following steps:

1. Right-click the SMSAP profile and select Backup.

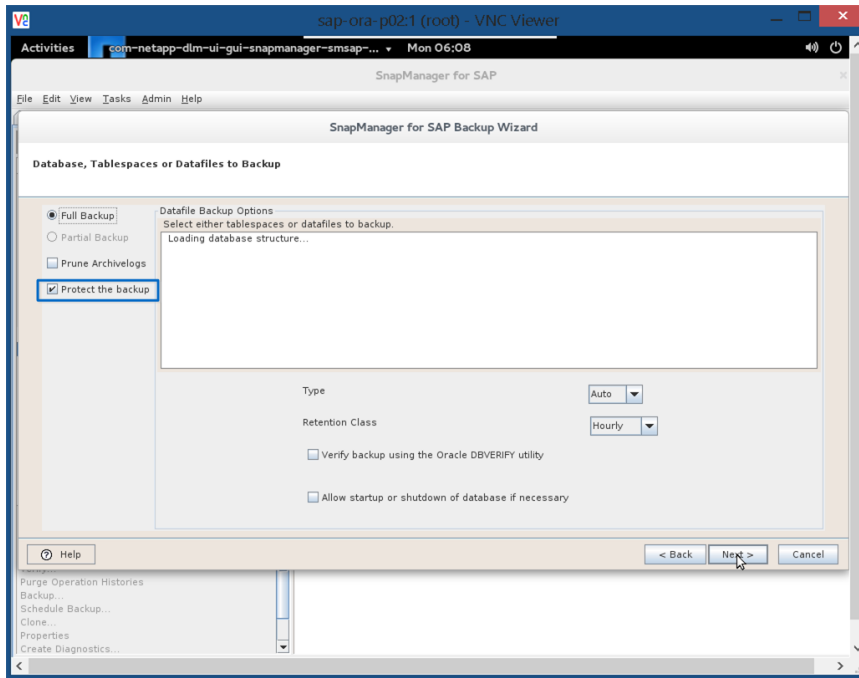


2. Enter the SnapVault label and click Next.


**Note:** This label must match the label that was configured in the “Configure SnapVault Replication” section.



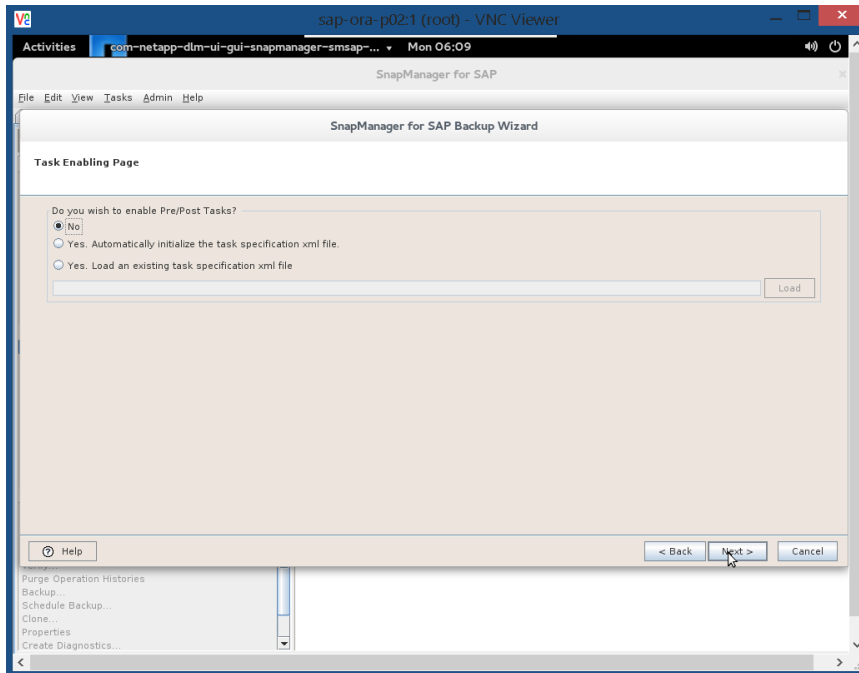
3. Select Protect the Backup and click Next.



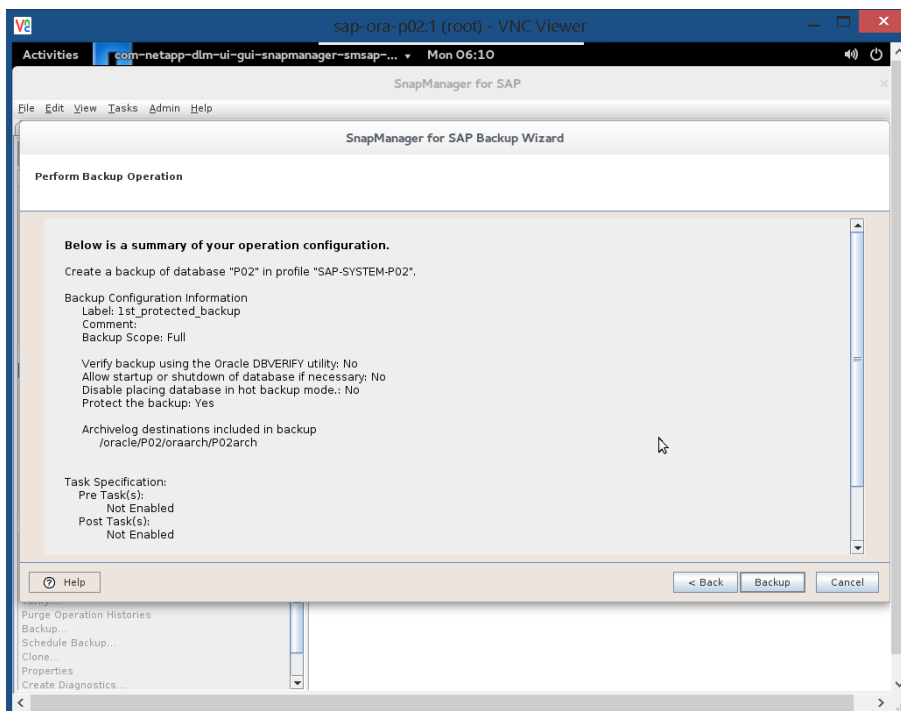
4. Select the archive log destination and click Next.



5. Select No to disable the pre- and post-backup tasks and click Next.



6. Review the configuration summary and click Backup.



7. After the operation log shows that the plug-in for data protection was successfully executed, click Finish.





## Required File System Changes

To refresh the SAP system Q02 based on data from the SAP system P02, the following tasks must be completed:

1. Stop the SAP system and the Oracle database by running the `stopsap` command.
2. Unmount all `sapdata` file systems.

**Note:** After installing SAP with SWPM, the file systems of the target system Q02 should look like the following example output. The `sapdata` file systems must be removed from the host, because these file systems will be available during the system copy process with SMSAP.

```
sap-ora-q02:~ # df -a
172.20.101.15:/saplog_q02/oracle_home      -          -          -          - /home/oracle
172.20.101.15:/saplog_q02/oracle          199229440  21926912  177302528  12% /oracle
172.20.101.15:/saplog_q02/sapmnt_Q02      -          -          -          - /sapmnt/Q02
172.20.101.15:/saplog_q02/sapusr_Q02      -          -          -          - /usr/sap/Q02
172.20.101.15:/saplog_q02/saphome_Q02     -          -          -          - /home/q02adm
172.20.101.15:/saplog_q02/oraq02_home     -          -          -          - /home/oraq02
172.20.101.15:/saplog_q02/saptrans        -          -          -          - /usr/sap/trans
172.20.101.15:/sapdata_q02/sapdata3       199229440  23318464  175910976  12% /oracle/Q02/sapdata3
172.20.101.15:/sapdata_q02/sapdata4      -          -          -          - /oracle/Q02/sapdata4
172.20.101.15:/sapdata_q02/sapdata2      -          -          -          - /oracle/Q02/sapdata2
172.20.101.15:/sapdata_q02/sapdata1      -          -          -          - /oracle/Q02/sapdata1
```

3. Unmount all of the `sapdata` file systems.

```
sap-ora-q02:~ # umount /oracle/Q02/sapdata1 /oracle/Q02/sapdata2 /oracle/Q02/sapdata4
/oracle/Q02/sapdata3
```

4. Delete the entries for `sapdata` file systems within `/etc/fstab` or add a comment sign to those lines containing `sapdata` file system entries.

```
sap-ora-q02: # cat /etc/fstab
172.20.101.15:/inst_media /mnt/inst_media nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/oracle /oracle nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/sapusr_Q02 /usr/sap/Q02 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/sapmnt_Q02 /sapmnt/Q02 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/saphome_Q02 /home/q02adm nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/oracle_home /home/oracle nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/oraq02_home /home/oraq02 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
172.20.101.15:/saplog_q02/saptrans /usr/sap/trans nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
#172.20.101.15:/sapdata_q02/sapdata1 /oracle/Q02/sapdata1 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
#172.20.101.15:/sapdata_q02/sapdata2 /oracle/Q02/sapdata2 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
#172.20.101.15:/sapdata_q02/sapdata3 /oracle/Q02/sapdata3 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
#172.20.101.15:/sapdata_q02/sapdata4 /oracle/Q02/sapdata4 nfs
rw,vers=3,hard,timeo=600,rsiz=65536,wsiz=65536,intr,nolock 0 0
```

5. Delete the following files and directories:

```
sap-ora-q02:/oracle/Q02 # rm -fr /oracle/Q02/origlogB/cntrl
sap-ora-q02:/oracle/Q02 # rm -fr /oracle/Q02/mirrlogA/cntrl
sap-ora-q02:/oracle/Q02 # rm -fr /oracle/Q02/oraarch
sap-ora-q02:/oracle/Q02 # rmdir /oracle/Q02/sapdata1
sap-ora-q02:/oracle/Q02 # rmdir /oracle/Q02/sapdata2
sap-ora-q02:/oracle/Q02 # rmdir /oracle/Q02/sapdata3
sap-ora-q02:/oracle/Q02 # rmdir /oracle/Q02/sapdata4
sap-ora-q02:/oracle/Q02 # rm /oracle/Q02/121/dbs/initQ02.ora
sap-ora-q02:/oracle/Q02 # rm /oracle/Q02/121/dbs/spfileQ02.ora
```

## Prepare Precloning and Postcloning Scripts

To prepare precloning and postcloning scripts, complete the following steps:

1. Copy the cleanup script `cleanup.sh` from the example directory to the preclone directory.

```
sap-ora-q02:~ # cp /opt/NetApp/smsap/plugins/examples/clone/create/pre/cleanup.sh
/opt/NetApp/smsap/plugins/clone/create/pre/
```

2. Copy the follow-up activities script from the example directory to the postclone directory.

```
sap-ora-q02:~ # cp
/opt/NetApp/smsap/plugins/examples/clone/create/post/sap_follow_up_activities.sh
/opt/NetApp/smsap/plugins/clone/create/post/
```

3. Adapt the lines of the cleanup script (highlighted in yellow) to align with the paths of the control files and the `initSID.ora` file based on your environment.

```
*****cleanup.sh*****
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/${SM_TARGET_SID}/origlogA/cntrl/cntrl${SM_TARGET_SID}.dbf:N"
"/oracle/${SM_TARGET_SID}/mirrlogA/cntrl/cntrl${SM_TARGET_SID}.dbf:N"
"/oracle/${SM_TARGET_SID}/origlogB/cntrl/cntrl${SM_TARGET_SID}.dbf: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/${SM_TARGET_SID}/121/dbs/init${SM_TARGET_SID}.ora:Y"
    )

    IFS=^
    for entry in ${files_to_cleanup[@]} ; do
        file=$(echo "$entry" | awk -F':' '{ print $1 }')
        save=$(echo "$entry" | awk -F':' '{ print $2 }')

        rm -fr $file
    done

    _exit 0
} *****cleanup.sh*****
```

## Configure Pluggable Authentication Modules

To configure PAM authentication for SMSAP, complete the following step:

1. Create the file `SnapManager` in the `/etc/pam.d` directory.

```
sap-ora-q02:/# cat /etc/pam.d/snapmanager
#%PAM-1.0
auth required pam_unix.so
account required pam_unix.so
```

## Configure `/etc/oratab` File

To configure the `/etc/oratab` file, complete the following step:

1. Delete or comment out the entries for the Q02 database in the `/etc/oratab` file.

```
sap-ora-q02:/etc/pam.d # cat /etc/oratab

## This file is used by ORACLE utilities.  It is created by root.sh
## and updated by either Database Configuration Assistant while creating
## a database or ASM Configuration Assistant while creating ASM instance.
```

```
# A colon, ':', is used as the field terminator. A new line terminates
# the entry. Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
# $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively. The third field indicates
# to the dbstart utility that the database should , "Y", or should not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
# Q02:/oracle/Q02/121:N # added by SAP RUNINSTALLER 2015_10_13-06-01-15
# Q02:/oracle/Q02/121:N # added by SWPM
```

## Configure SnapDrive for UNIX

To configure SDU, complete the following steps:

1. The host name of the interface that is used for management access must match the name of the SVM. In the lab setup, the file systems are mounted using `sap-svm-data-path` and the SVM is managed using `sap-svm`. `Sap-target` is the management interface of the SVM that was used as a target for the SnapVault replication; `sap-target-data-path` is the data LIF of the target SVM.

```
sap-ora-q02:/ # cat /etc/hosts
172.20.101.14    sap-svm
172.20.101.15    sap-svm-data-path
172.20.101.44    sap-target
172.20.101.45    sap-target-data-path
```

2. Configure the required SDU credentials for the management interfaces.

```
sap-ora-q02:/ # snapdrive config set vsadmin sap-svm
Password for vsadmin:
Retype password:

sap-ora-q02:/ # snapdrive config set vsadmin sap-target
Password for vsadmin:
Retype password:
```

3. Add the data path interface to the management interface.

```
sap-ora-q02:/ # snapdrive config set -mgmtpath sap-svm sap-svm-data-path

sap-ora-q02:/ # snapdrive config set -mgmtpath sap-target sap-target-data-path
```

4. Verify the configuration.

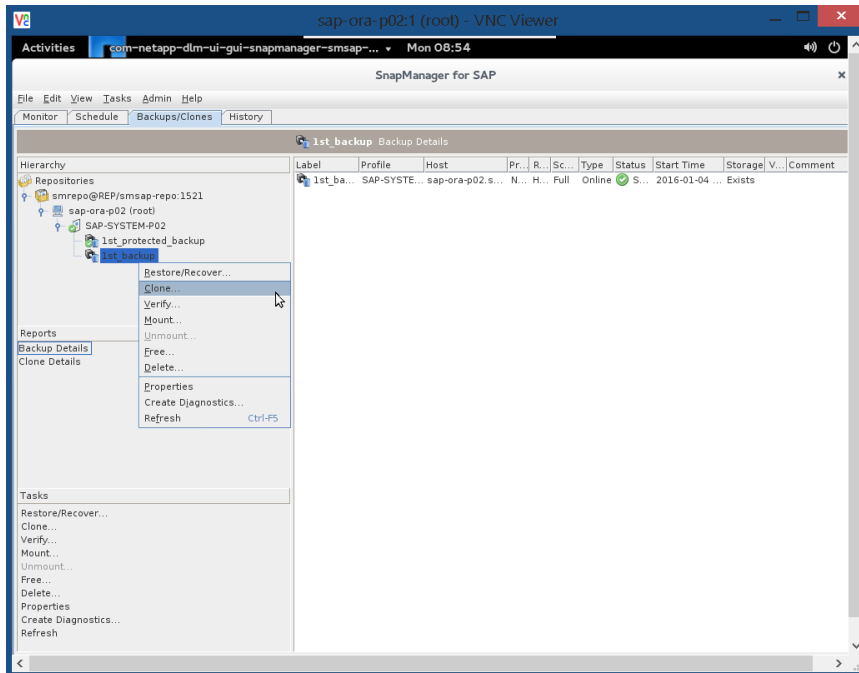
```
sap-ora-q02:~ # snapdrive config list
username      appliance name  appliance type
-----
vsadmin       sap-target      StorageSystem
vsadmin       sap-svm         StorageSystem
sap-ora-q02:~ # snapdrive config list -mgmtpath
system name   management interface  datapath interface
-----
sap-svm       172.20.101.14         172.20.101.15
sap-target    172.20.101.44         172.20.101.45
```

## 10.2 SAP System Copy on Primary Storage

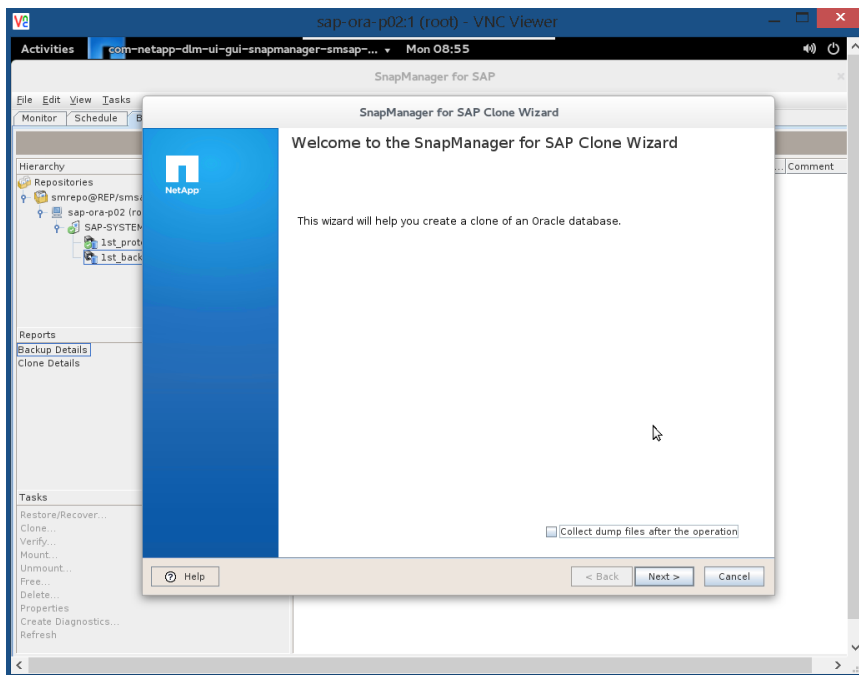
### Execute Cloning Process and Specify Initial Clone

To execute the cloning process and specify the initial clone, complete the following steps:

1. Right-click on the backup and select Clone.

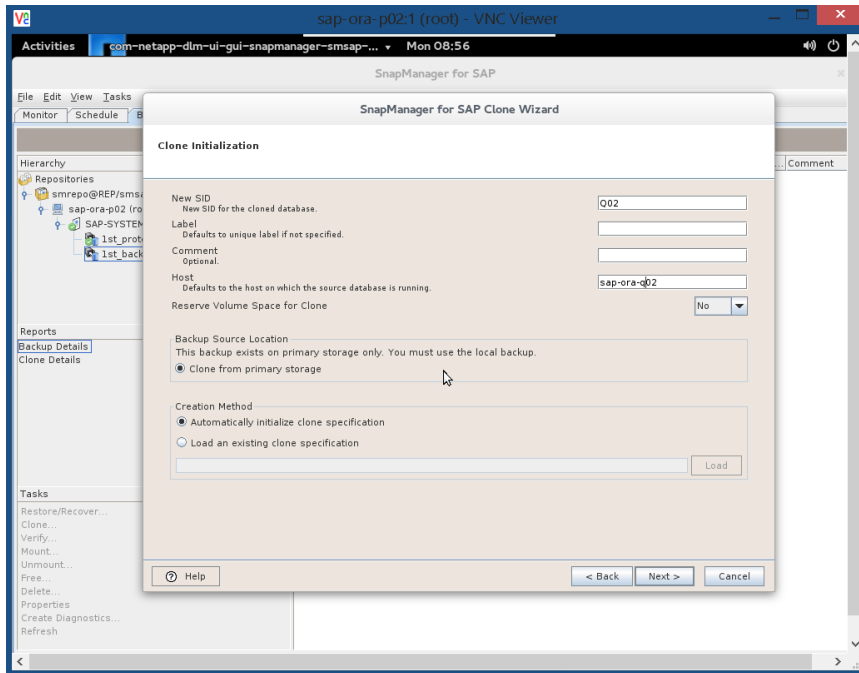


2. On the SnapManager for SAP Clone Wizard welcome page, click Next.

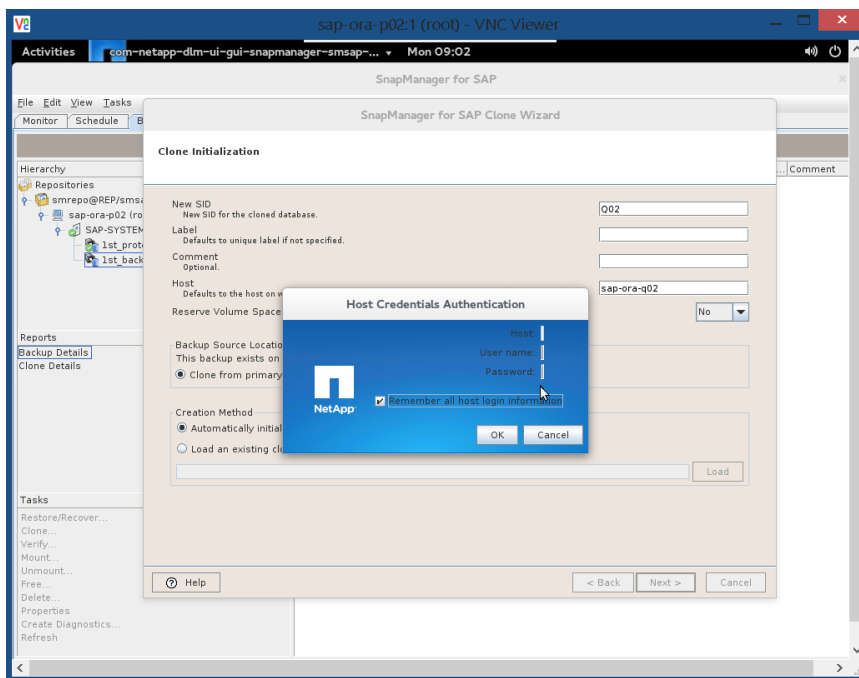


3. Enter the Target SID and target host, and then select Automatically Initialize Clone Specification. Click Next.

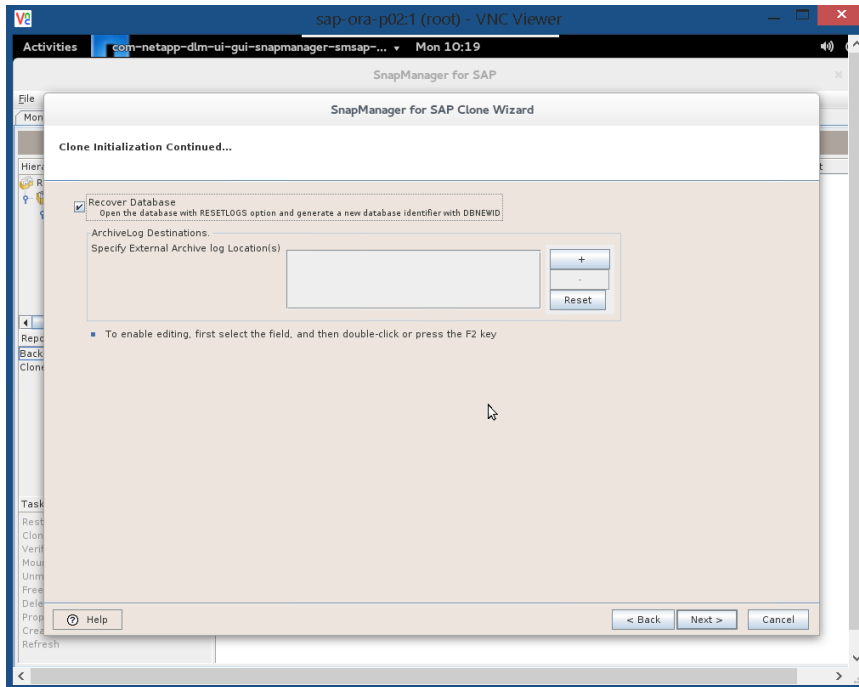
**Note:** In a later step, the cloning specification will be saved to a file so that it can be reused with every subsequent cloning process.



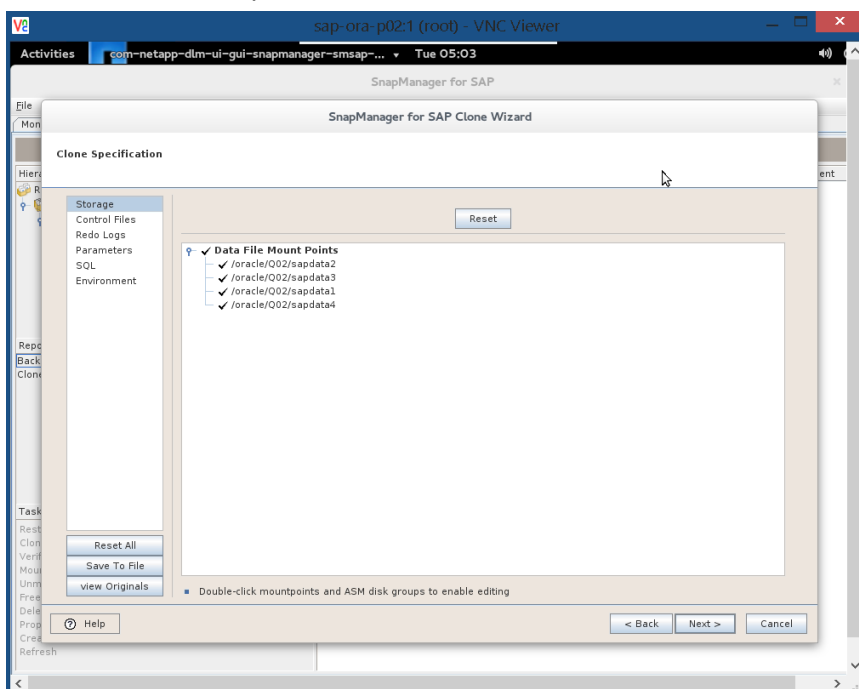
4. Enter the target host credentials. Select Remember All Host Login Information to store the host credentials for all subsequent operations. Click OK.




5. Select the Recover Database and click Next.



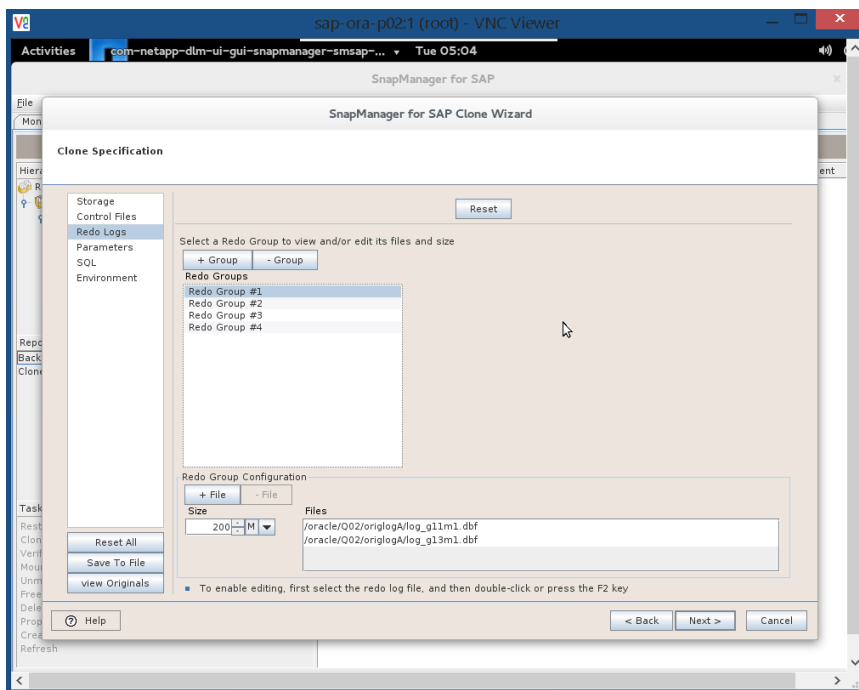
6. Enter the mount points for the data files and click Next.



7. Enter the control file locations and click Next.

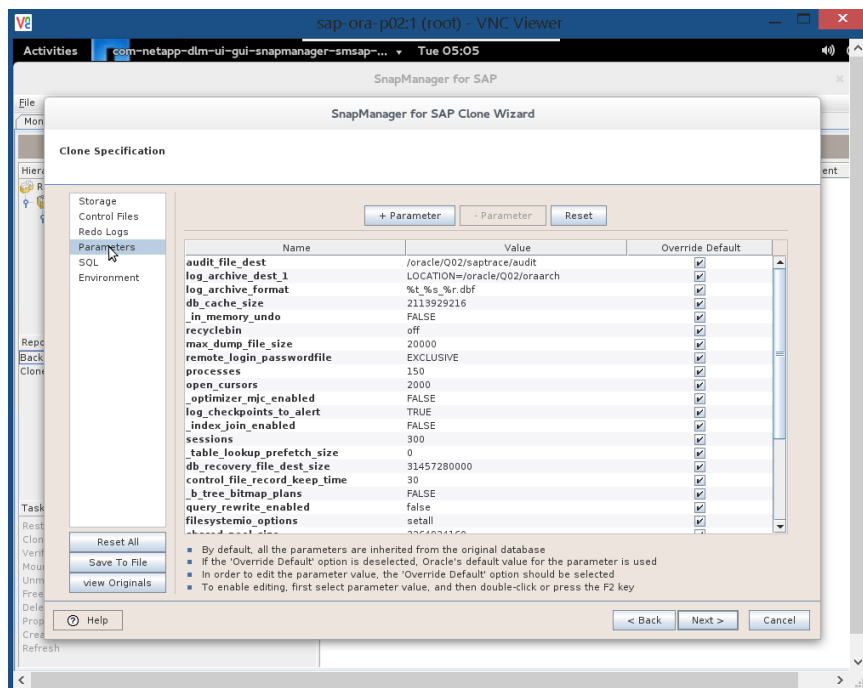


8. Enter the redo log file locations and click Next.



9. Enter the values for the audit\_file\_dest and log\_archive\_dest\_1 parameters and click Next.

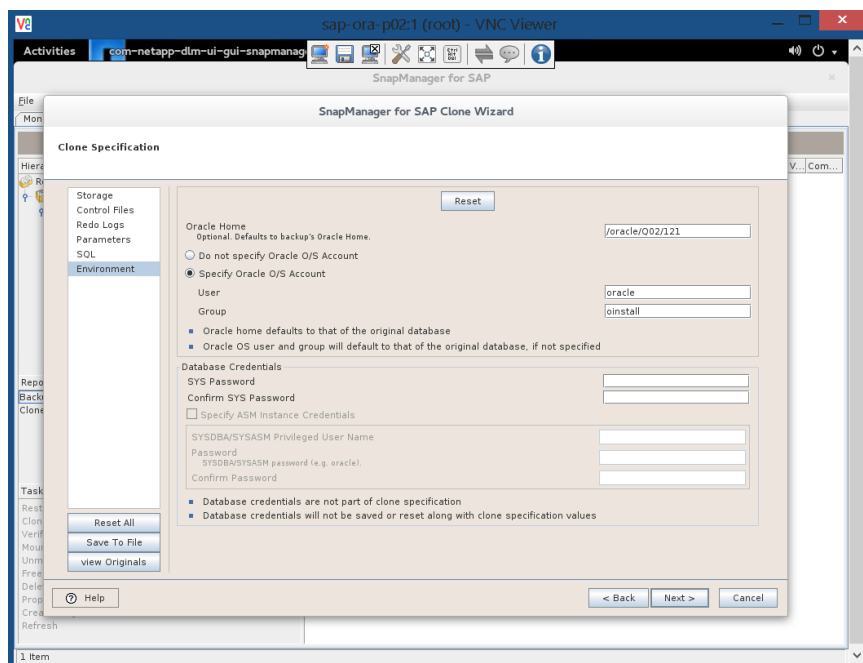




10. Enter the Oracle Home and Oracle OS accounts.

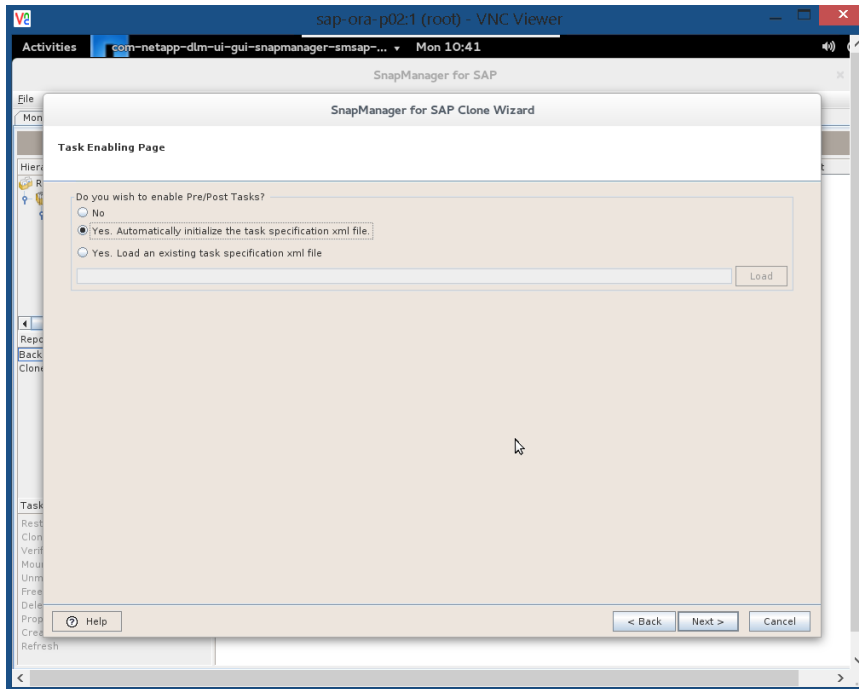
**Note:** With SAP on Oracle 12, the user `oracle` and the group `oinstall` must be used for the cloning process.

Click **Save to File** to save the cloning parameter to use for subsequent cloning operations and click **Next**.

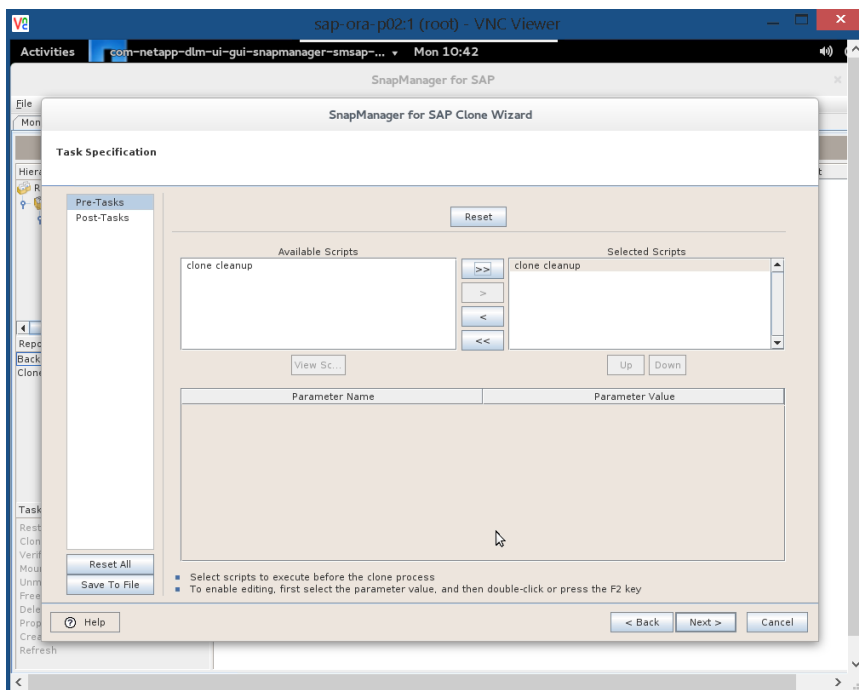


11. Select the **Yes. Automatically Initialize the Task Specification** option.

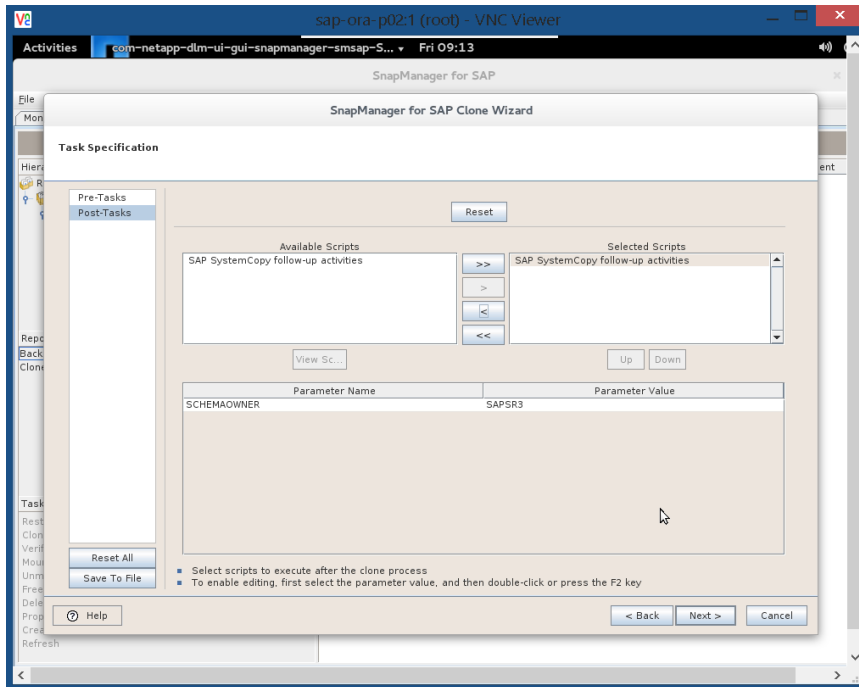
**Note:** In a later step, the task specification will be saved to a file and the file can then be reused for every subsequent cloning process between these two systems. This step saves you from having to enter the parameters again.



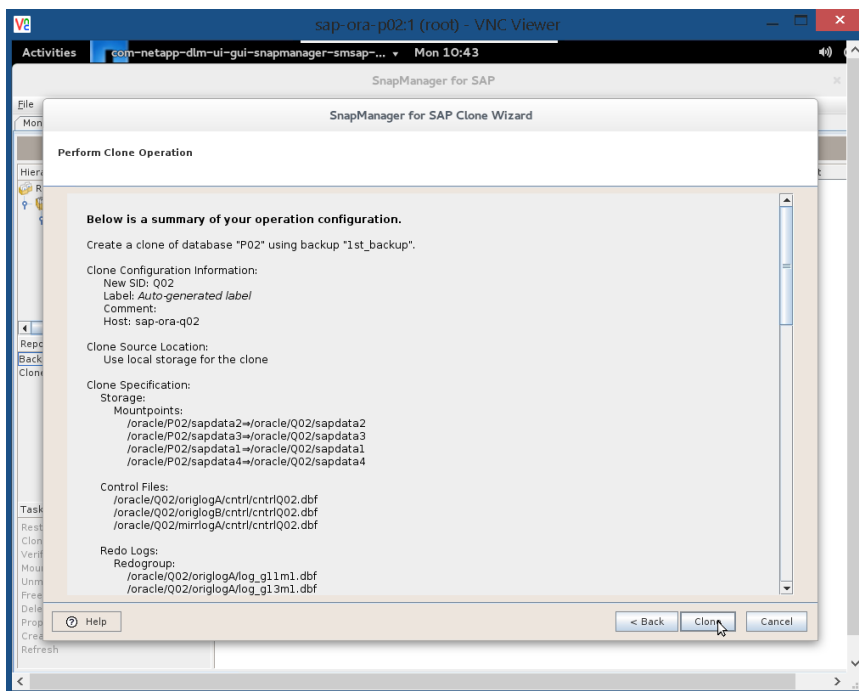
12. Select Clone Cleanup as a Pre-Task and click Next.



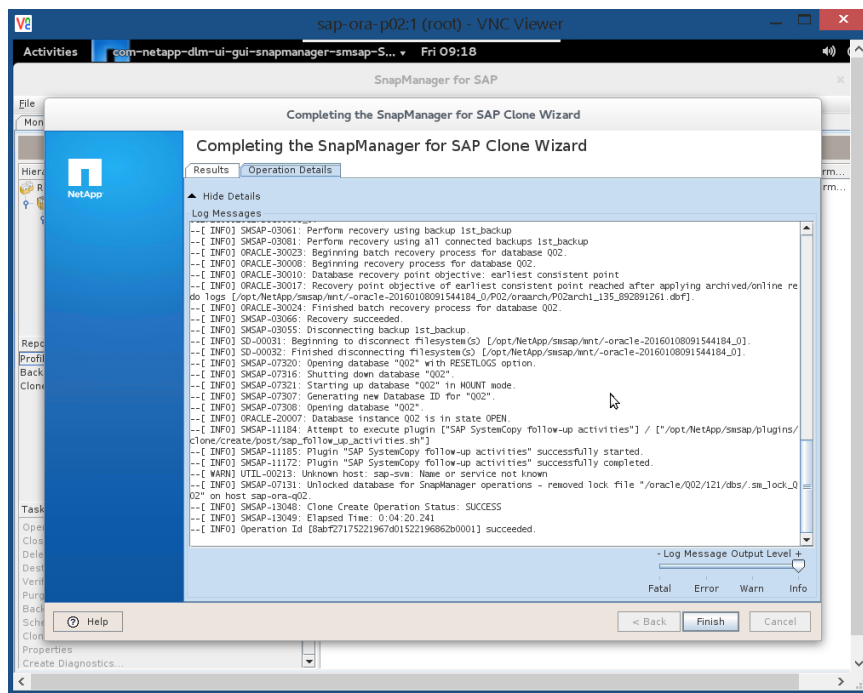
13. Select SAP SystemCopy Follow-up Activities as a post task and enter the parameter name SCHEMAOWNER. Click Save to File to save the task specification for subsequent cloning operations.



14. Review the configuration summary and click Clone.



15. After the clone is created, click Finish.



## Configure OPS\$ and Start SAP

For SAP on Oracle 12, a secure store is required for the OPS\$ configuration. The password within the database and the secure store in the file system must be set using Brconnect.

1. Run the Brconnect command with the user SIDadm.

```
sap-ora-q02:/ # su - q02adm
sap-ora-q02:q02adm 53>
sap-ora-q02:q02adm 53> brconnect -u / -c force -f chpass -o SAPSR3 -p donald
BR0801I BRCONNECT 7.40 (12)

BR0280I BRCONNECT time stamp: 2016-01-08 09:24:01
BR0828I Changing password for database user SAPSR3 ...

BR0280I BRCONNECT time stamp: 2016-01-08 09:24:01
BR0829I Password changed successfully in database for user SAPSR3
BR0831I Table SAPUSER not found for user SAPSR3

BR0280I BRCONNECT time stamp: 2016-01-08 09:24:01
BR1525I Setting password for database user SAPSR3 in secure storage
/usr/sap/Q02/SYS/global/security/rsecssfs/data/SSFS_Q02.DAT ...

BR0280I BRCONNECT time stamp: 2016-01-08 09:24:08
BR1526I Password set successfully for database user SAPSR3 in secure storage
/usr/sap/Q02/SYS/global/security/rsecssfs/data/SSFS_Q02.DAT

BR0280I BRCONNECT time stamp: 2016-01-08 09:24:08
BR0802I BRCONNECT completed successfully
```

2. Verify the configuration by running the R3trans command.

```
sap-ora-q02:q02adm 53> R3trans -d
This is R3trans version 6.24 (release 742 - 18.11.14 - 20:14:09).
unicode enabled version
R3trans finished (0000).
```

3. Start the SAP system.

```
sap-ora-q02:q02adm 54> startsap
Checking Q02 Database
Database is running
-----
Starting Startup Agent sapstartsrv
```

```

OK
Instance Service on host sap-ora-q02 started
-----
starting SAP Instance ASCS01
Startup-Log is written to /home/q02adm/startsap_ASCS01.log
-----
/usr/sap/Q02/ASCS01/exe/sapcontrol -prot NI_HTTP -nr 01 -function Start
Instance on host sap-ora-q02 started
Starting Startup Agent sapstartsrv
OK
Instance Service on host sap-ora-q02 started
-----
starting SAP Instance DVEBMGS00
Startup-Log is written to /home/q02adm/startsap_DVEBMGS00.log
-----
/usr/sap/Q02/DVEBMGS00/exe/sapcontrol -prot NI_HTTP -nr 00 -function Start
Instance on host sap-ora-q02 started
sap-ora-q02:q02adm 55>

```

## 10.3 SAP System Copy on Secondary Storage

### Mount and Export Volumes at Target SVM

Before an SAP system copy can be executed, the SnapVault target volumes must be mounted within the target SVM's namespace and then exported. This process allows the SAP system copy target host to mount the volumes.

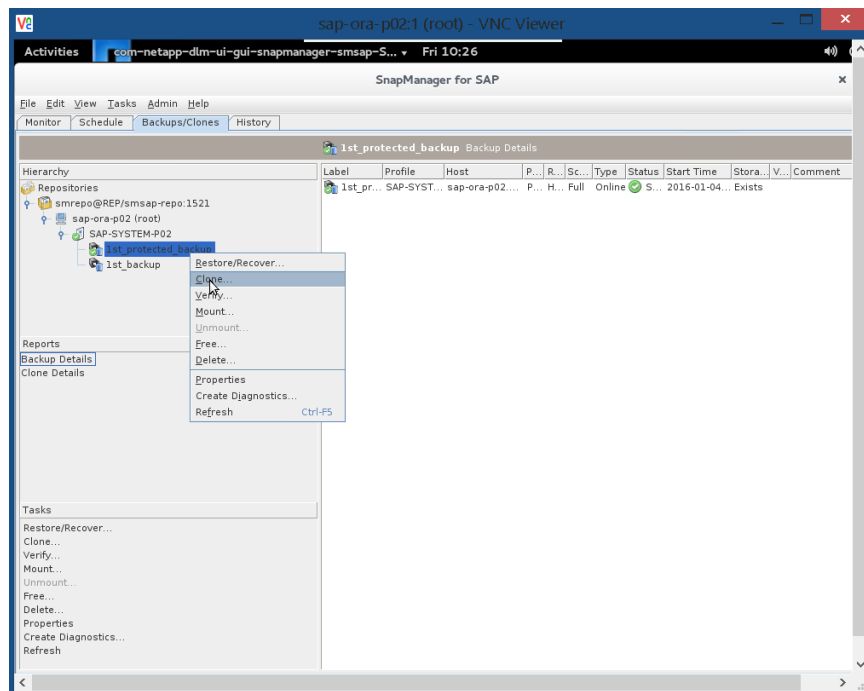
### Execute Cloning Process

The clone and task specification files, which were saved in the process described in the “SAP System Copy on Primary Storage” section, are used during the clone execution.

To execute the cloning process, complete the following steps by using the SMSAP wizard.

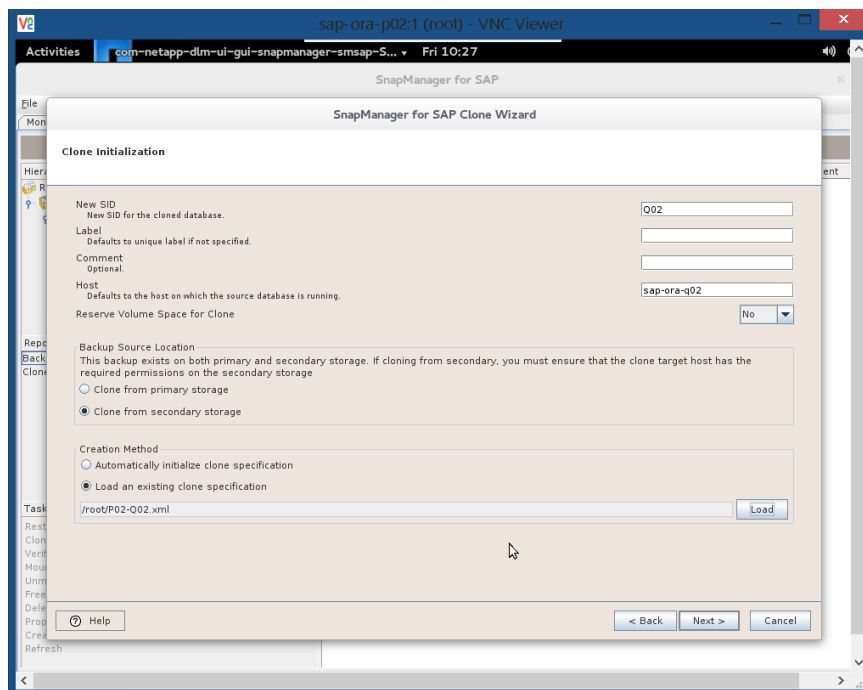
**Note:** The following workflow description only shows steps that are different from the workflow described in the “SAP System Copy on Primary Storage” section.

1. Select the protected backup and then select Clone.



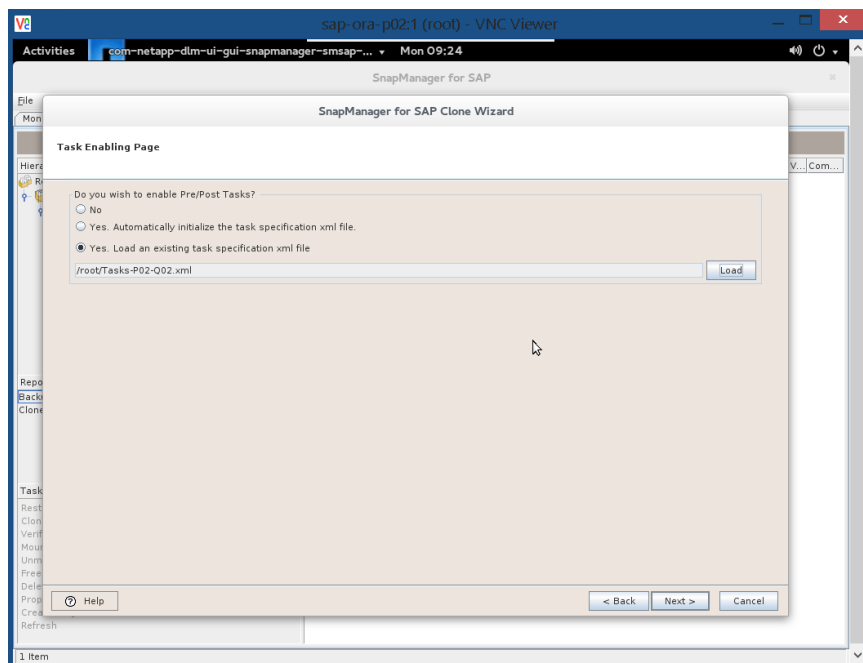
2. Select Clone from Secondary Storage, enter the target SID and target host, and then select Load an Existing Clone Specification. Click Next.

**Note:** The cloning specification that was saved in the “Execute Cloning Process and Specify Initial Clone” section can be reused for the cloning process.

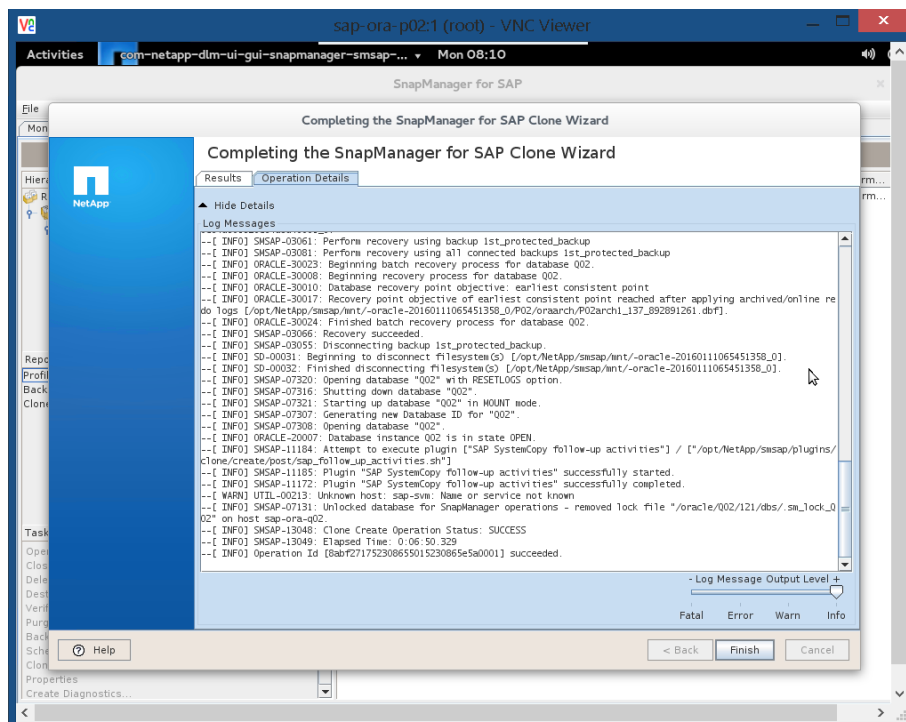


3. Select Yes. Load an Existing Task Specification xml File and click Next.

**Note:** The task specification that was saved in the “Execute Cloning Process and Specify Initial Clone” section can be reused for the cloning process.



4. After the clone is created, click Finish



## Configure OPS\$ and Start SAP

For SAP on Oracle 12, a secure store is required for the OPS\$ configuration. The password within the database and the secure store in the file system must be set using `Brconnect`.

For detailed steps about configuring OPS\$, see the “Configure OPS\$ and Start SAP” section.

## 11 Configure Br\*Tools

The current release of SMSAP using BR\*Tools does not support data protection to a secondary storage. Backups can be created only at the primary storage system.

### 11.1 Configure Backint File

To configure the `backint` file, complete the following steps:

1. Create a link of the `backint` executable in the SAP executable directory.

```
sap-ora-p02:~ # su - p02adm
sap-ora-p02:p02adm 57> ln -s /opt/NetApp/smsap/bin/backint /sapmnt/P02/exe/uc/linuxx86_64/backint
```

2. Create a `util` file for `backint`.

```
sap-ora-p02:/oracle/P02 # su - oracle
sap-ora-p02 /home/oracle% cd /oracle/P02/sapprof/
sap-ora-p02 P02/sapprof% vi initP02_hourly.utl
```

3. Configure the `util` file with the following information:

```
sap-ora-p02 P02/sapprof% cat initP02_hourly.utl
profile_name = SAP-SYSTEM-P02
fast = FALLBACK
protect = NO
retain = HOURLY
```

4. Copy the configuration template.

```
sap-ora-p02 P02/sapprof% cp initP02.sap initP02_hourly.sap
sap-ora-p02 P02/sapprof% vi initP02_hourly.sap
```

5. Change the following entries in `initP02_hourly.sap`:

```
backup_dev_type = util_file
util_par_file = initP02_hourly.utl
```

## 11.2 Configure SMSAP Credentials

With Oracle 12, the SMSAP credentials have to be set for the user “oracle” if the Br\*Tools are owned by the “oracle” user and the sticky bit is set accordingly. The owner of the Br\*Tools can be checked by running the following commands:

```
sap-ora-p02:p02adm 54> ls -al br*
-rwsrwsr-- 1 oracle oinstall 9927726 Nov 21 2014 brarchive
-rwsrwsr-- 1 oracle oinstall 10008774 Nov 21 2014 brbackup
-rwsrwsr-- 1 oracle oinstall 11975593 Nov 21 2014 brconnect
-rwsrwsr-- 1 oracle oinstall 10483567 Nov 21 2014 brrecover
-rwsrwsr-- 1 oracle oinstall 6183978 Nov 21 2014 brrestore
-rwsrwsr-- 1 oracle oinstall 12428837 Nov 21 2014 brspace
```

To configure SMSAP credentials, run the following commands:

### 1. Set the SMSAP repository credentials.

```
sap-ora-p02:/home/oracle # su - oracle

sap-ora-p02 /home/oracle% smsap credential set -repository -dbname REP -host smsap-repo -port
1521 -login -username smrepo
Enter password for database connection smrepo@smsap-repo:1521/REP: *****
[ INFO] SMSAP-20019: Set password for repository "smrepo@REP/smsap-repo:1521" in user credentials
for "oracle".
```

### 2. Synchronize the SMSAP repository.

```
sap-ora-p02 /home/oracle% smsap profile sync -repository -dbname REP -host smsap-repo -port 1521
-login -username smrepo
[ INFO] SMSAP-20010: Synchronizing mapping for profiles in repository "smrepo@REP/[smsap-
repo]:1521".
[ INFO] SMSAP-20011: Loaded mapping for profile "SAP-SYSTEM-P02".
```

### 3. Set the credentials for the SMSAP profile.

```
sap-ora-p02 /home/oracle% smsap credential set -profile -name SAP-SYSTEM-P02
Enter password for profile SAP-SYSTEM-P02: *****
[ INFO] SMSAP-20020: Set password for profile "SAP-SYSTEM-P02" in user credentials for "oracle".
```

## 11.3 Execute Brbackup Using CLI

To run a backup by using the CLI, complete the following steps:

### 1. Execute the backup as user SIDadm.

```
sap-ora-p02:/home/oracle # su - p02adm

sap-ora-p02:p02adm 54> brbackup -u / -t online -p /oracle/P02/sapprof/initP02_hourly.sap
BR0051I BRBACKUP 7.40 (12)
BR0055I Start of database backup: besgszlp.anf 2016-01-12 07:54:57
BR0484I BRBACKUP log file: /oracle/P02/sapbackup/besgszlp.anf
BR0477I Oracle pfile /oracle/P02/sapprof/initP02.ora created from spfile
/oracle/P02/121/dbs/spfileP02.ora
BR1805I Oracle spfile /oracle/P02/sapprof/spfileP02.ora created from pfile
/oracle/P02/sapprof/initP02.ora

BR0280I BRBACKUP time stamp: 2016-01-12 07:54:58
BR0319I Control file copy created: /oracle/P02/sapbackup/cntrlP02.dbf 20725760

BR0280I BRBACKUP time stamp: 2016-01-12 07:54:58
BR0057I Backup of database: P02
BR0058I BRBACKUP action ID: besgszlp
BR0059I BRBACKUP function ID: anf
BR0110I Backup mode: ALL
BR0077I Database file for backup: /oracle/P02/sapbackup/cntrlP02.dbf
BR0061I 16 files found for backup, total size 23199.883 MB
BR0143I Backup type: online
BR0130I Backup device type: util_file
```



```

BR0109I Files will be saved by backup utility at file level

BR0280I BRBACKUP time stamp: 2016-01-12 07:54:58
BR0256I Enter 'c[ont]' to continue, 's[top]' to cancel BRBACKUP:
c
BR0280I BRBACKUP time stamp: 2016-01-12 07:55:01
BR0257I Your reply: 'c'
BR0259I Program execution will be continued...

BR0280I BRBACKUP time stamp: 2016-01-12 07:55:01
BR0315I 'Alter database begin backup' successful

BR0280I BRBACKUP time stamp: 2016-01-12 07:55:01
BR0229I Calling backup utility with function 'backup'...
BR0278I Command output of '/usr/sap/P02/SYS/exe/run/backint -u P02 -f backup -i
/oracle/P02/sapbackup/.besgszlp.lst -t file -p /oracle/P02/sapprof/initP02_hourly.utl':
[ INFO] Backint Interface for SnapManager(R) 1.0
[ INFO] SMSAP-01208: Updating utility profile parameter profile_name, [P02] -> [SAP-SYSTEM-P02]
[ INFO] SMSAP-01208: Updating utility profile parameter fast, [FALLBACK] -> [require]
[ INFO] SMSAP-01208: Updating utility profile parameter protect, [] -> [no]
[ INFO] SMSAP-01208: Updating utility profile parameter retain, [DAILY] -> [HOURLY]
Backint Interface for SnapManager(R) 1.0
[ INFO] SMSAP-07822: Using directory /oracle/P02/sapbackup for writing the .switch files.

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:13
#FILE.... /oracle/P02/sapdata2/sr3_1/sr3.data1
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:13
#FILE.... /oracle/P02/sapdata2/sr3_2/sr3.data2
#SAVED.... 160112075501

Truncated ....

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:13
#FILE.... /oracle/P02/sapbackup/cntrlP02.dbf
#SAVED.... 160112075501
Operation Id [N81cal038aec59f052c2efa8fe75e307d] succeeded.

BR0232I 16 of 16 files saved by backup utility

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:13
BR0230I Backup utility called successfully

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:15
BR0317I 'Alter database end backup' successful

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:15
BR0340I Switching to the next online redolog file for database instance P02 ...
BR0321I Switch to the next online redolog file for database instance P02 successful

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:16
BR0229I Calling backup utility with function 'backup'...
BR0278I Command output of '/usr/sap/P02/SYS/exe/run/backint -u P02 -f backup -i
/oracle/P02/sapbackup/.besgszlp.lst -t file -p /oracle/P02/sapprof/initP02_hourly.utl':
[ INFO] Backint Interface for SnapManager(R) 1.0
[ INFO] SMSAP-01208: Updating utility profile parameter profile_name, [P02] -> [SAP-SYSTEM-P02]
[ INFO] SMSAP-01208: Updating utility profile parameter fast, [FALLBACK] -> [require]
[ INFO] SMSAP-01208: Updating utility profile parameter protect, [] -> [no]
[ INFO] SMSAP-01208: Updating utility profile parameter retain, [DAILY] -> [HOURLY]
Backint Interface for SnapManager(R) 1.0
[ INFO] SMSAP-07822: Using directory /oracle/P02/sapbackup for writing the .switch files.

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapprof/initP02.ora
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapprof/spfileP02.ora
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapprof/initP02_hourly.sap
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53

```

```
#PFLOG.... /oracle/P02/sapprof/initP02_hourly.utl
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapreorg/spaceP02.log
#SAVED.... 160112075501

BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapbackup/besgszlp.anf
#SAVED.... 160112075501

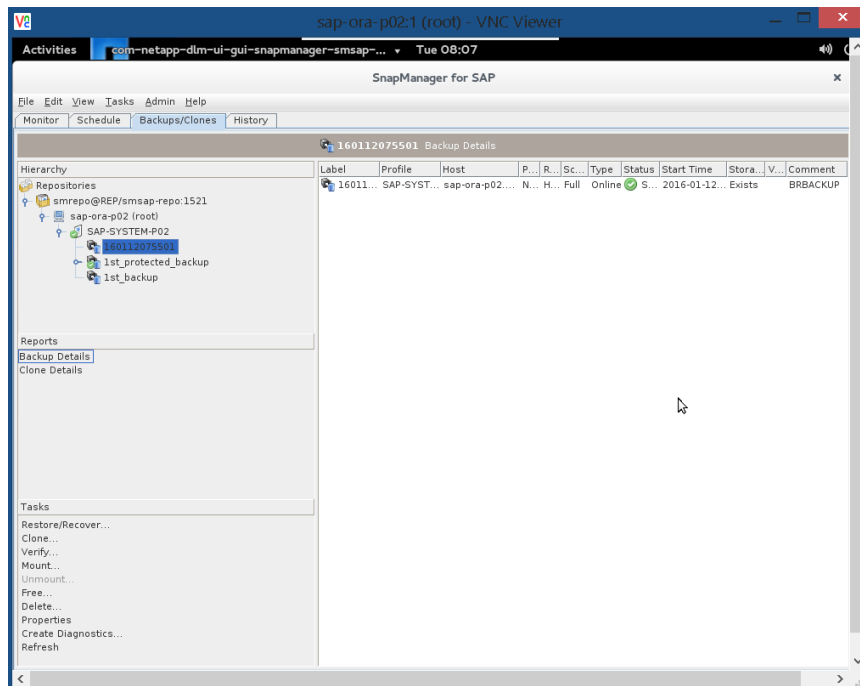
BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
#PFLOG.... /oracle/P02/sapbackup/backP02.log
#SAVED.... 160112075501
Operation Id [N34f54a0a8b9a6e1e2cdf623f7775f825] succeeded.

BR0232I 7 of 7 files saved by backup utility

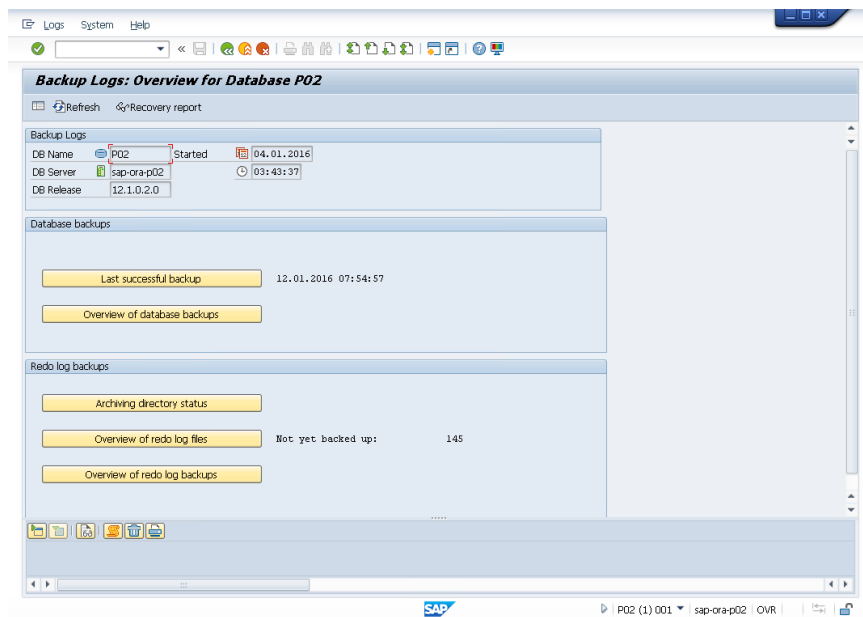
BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
BR0230I Backup utility called successfully

BR0056I End of database backup: besgszlp.anf 2016-01-12 07:56:53
BR0280I BRBACKUP time stamp: 2016-01-12 07:56:53
BR0052I BRBACKUP completed successfully
sap-ora-p02:p02adm 54>
```

## 2. The backup is visible from within the SMSAP GUI.



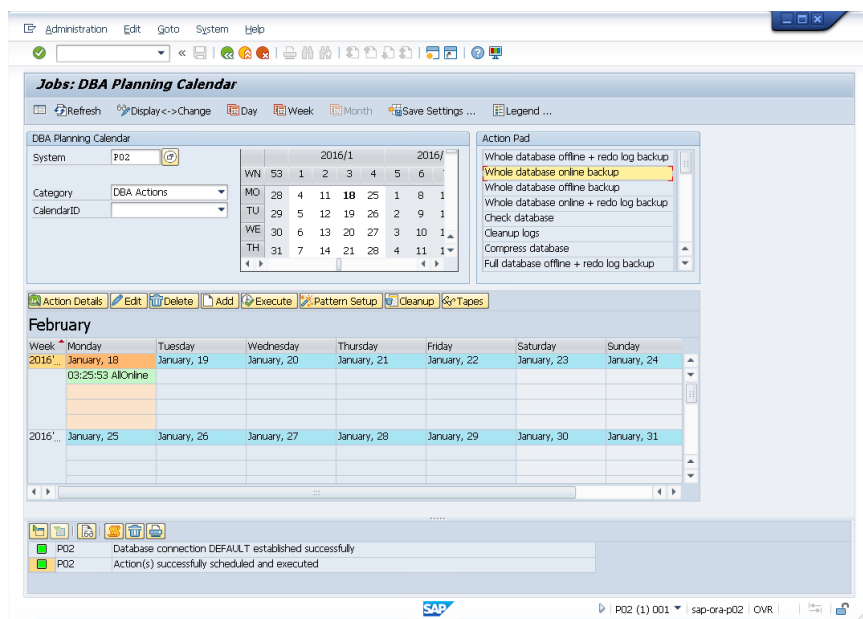
## 3. The backup is also visible from within the SAP GUI by using the transaction DB02.



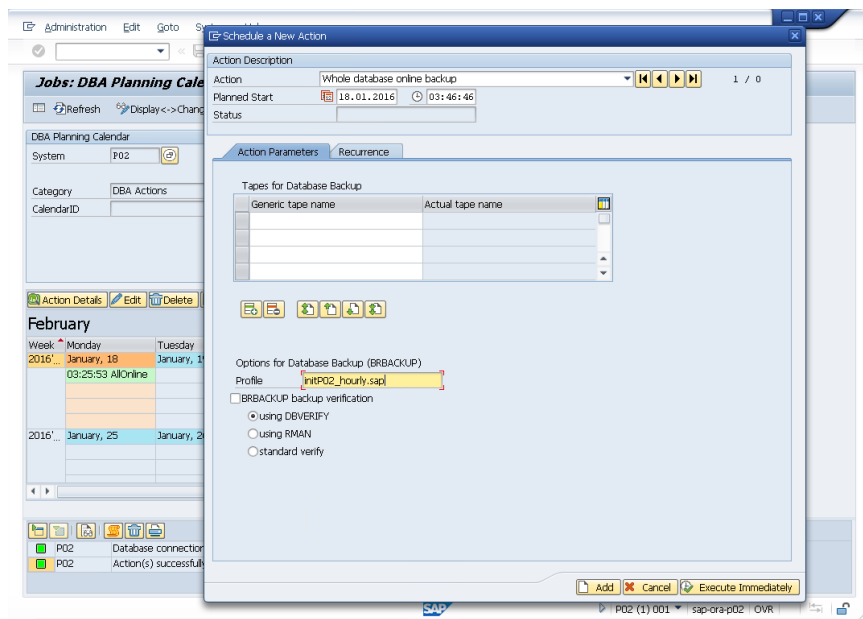
## 11.4 Execute Backup Using SAP DBA Planning Calendar

To execute a backup by using SAP's DBA planning calendar, complete the following steps:

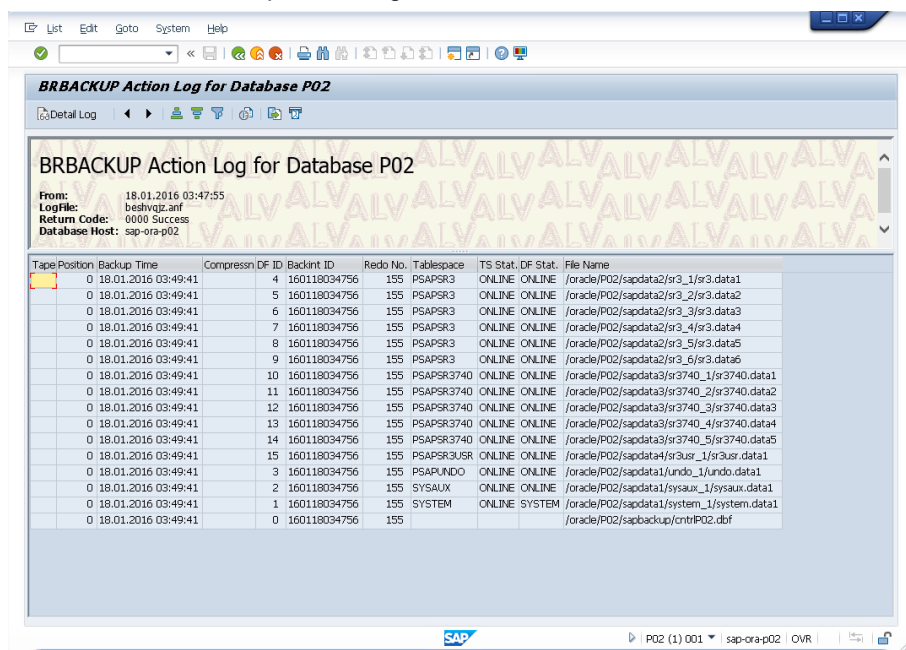
1. Execute transaction DB13 within the SAP GUI.



2. Select Whole Database Online Backup and enter the profile name InitP02\_hourly.sap. Click Execute Immediately.



### 3. Review the backup action log.



## 11.5 SAP System Copy Using Backup Created by Brbackup

To use a backup created by Brbackup as a source for an SAP system copy, complete the following steps:

1. Set the following parameter within the SMSAP configuration file:  
/opt/NetApp/smsap/properties/smsap.config.

```
brbackup.enable.cloneable.backups=true
```

2. After the configuration change, restart the SMSAP server.

```
sap-ora-p02:/oracle/P02/sapprof # smsap_server restart
SMSAP-17100: SnapManager Server started on secure port 27314 with PID 13261.
```

3. All backups that were created after the configuration change can now be used as a source for an SAP system copy.

**Note:** This system copy process is identical to the process described in “SAP System Copy on Primary Storage” except that the saved cloning and tasks specifications can be reused.

## Version History

Version	Date	Document Version History
Version 1.0	December 2013	Initial version
Version 2.0	March 2016	Update to cover SnapManager for SAP 3.4

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.

## Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4250-0316