



SAP NetWeaver[®] Application Server ABAP/Java with Oracle Database@Azure



Based on Oracle Exadata Cloud Infrastructure X9M

July 2024 | Version 1.0
Copyright © 2024, Oracle and/or its affiliates
Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Revision History

The following revisions have been made to this document.

DATE	REVISION
July 2024	Initial version - For testing purposes only - Request SAP to certify (BC-DB-ORA)

Table of Contents

Introduction	5
Scope and Assumptions	5
Terminology	7
Planning Network File Systems	9
Overview of Oracle Database@Azure	10
SAP NetWeaver Database on Oracle Exadata Database Service	10
Regions and Availability Domains in OCI	10
Regions and Availability Zones in Azure	11
Virtualization and Databases with Exadata Cloud Infrastructure X9M	11
VM Cluster Scaling Options	12
Oracle Database for SAP NetWeaver with Exadata Cloud Infrastructure X9M	12
Exadata Cloud Infrastructure X9M	12
Exadata Database Service Licenses	13
Planning Your Oracle Database@Azure Service for SAP NetWeaver Application Server ABAP/Java	15
Prerequisites	15
Deployment Restrictions with SAP NetWeaver for Oracle Database@Azure	15
Documentation	16
Workload Size	16
Planning the SAP Deployment	16
Exadata Database Service Requirements	18
General Installation Overview	18
Implementing the Deployment	19
Get Your Azure and OCI Accounts	19
Deploy the Cloud Resources	20
Prepare Your Environment	35
Prepare VM Cluster Nodes	37
Install SAP NetWeaver Application Server ABAP/Java	43
Finalize the Installation	48
Revert Changes on VM Cluster Nodes and SAP NetWeaver Application Servers	60
SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases	61
Installation of Patches for Oracle RDBMS Software	61
Installation of Patches for Exadata Database Service Tooling	62
Installation of OS Patches	62
Finish SAP Monitoring Setup	62
Set Up SAP Diagnostic Agent for SAP Solution Manager	62
Using SAP Transaction DB13	63
Local Update Dispatching	64

Migrating Databases	64
VM Cluster Node Subsetting	64
Recommendation for SAP HA with SAPCTL	65
Add Nodes	65
Remove Nodes	77
Delete an Entire VM Cluster	80
High Availability with Oracle Data Guard	81
Support of SAP Standalone Enqueue Server 2 and Enqueue Replicator 2	81
References	81
SAP	81
Oracle	83
Azure	83

Introduction

Oracle Database@Azure is a service that runs Oracle Exadata Database Service on Oracle Exadata Cloud Infrastructure directly in a Microsoft data center. Oracle Database@Azure enables customers to leverage the benefits of Exadata Cloud Infrastructure technology with Microsoft Azure-based SAP NetWeaver Application Server ABAP/Java landscapes without sacrificing top-tier performance. SAP NetWeaver Application Servers operate on Azure compute instances that are located in the same region and availability zone and in the same virtual network (VNet) to which the VM cluster nodes of the Exadata Cloud Infrastructure are connected. This configuration ensures excellent network latency times, which are ideally suited for the operation of SAP NetWeaver Application Server ABAP/Java. Simultaneously, Exadata Cloud Infrastructure is fully integrated with Oracle Cloud Infrastructure (OCI), allowing access to all relevant OCI services, such as Object Storage.

SAP NetWeaver Application Server ABAP/Java landscapes based on Oracle Database@Azure are managed from both the Azure portal and the Oracle Cloud Console.

- The Azure portal is used for deploying Exadata Cloud Infrastructure, the VM clusters, the virtual networks, and the Azure compute instances.
- The Oracle Cloud Console is used for more specialized tasks, such as upgrading the Oracle VM clusters, patching the Oracle Grid Infrastructure software, or scaling. For this purpose, the Azure portal provides a link directly to the corresponding cloud resource in Oracle Cloud Console. The direct transition from the Azure portal to the relevant resource in the Oracle Cloud Console is facilitated by a federated single-sign-on mechanism in which the customer logs in to the Console by using Microsoft Entra ID.

This document refers to functions, input screens, and dialogs in the Azure portal as well as in the Oracle Cloud Console, as they were available at the time the document was created.

Scope and Assumptions

This document is a reference guide for deploying Oracle databases of the SAP NetWeaver Application Server ABAP/Java platform on Oracle Database@Azure, based on Oracle Exadata Cloud Infrastructure X9M running Oracle Linux 8 on the VM cluster nodes.

Note: For historic reasons, the terms *Exadata Database Service* and *Exadata Cloud Service* are used interchangeably.

This document also describes how to implement SAP high availability (HA) by using Oracle Grid Infrastructure with the SAPCTL addon, which is an optional method for implementing SAP HA. Customers who do not require SAP HA or want to implement SAP HA in another SAP supported way can skip the steps related to installing and configuring SAP ABAP central services (ASCS) and enqueue replication server (ERS). The SAPCTL software package is attached to SAP Note 1496927 and contains the installation and configuration manual.

Note: *Application-specific virtual IP addresses (APP-VIPs)* in OCI are required only by customers who plan to implement SAP HA by using SAPCTL directly on Exadata Cloud Infrastructure X9M. Customers who do not want to implement SAP HA by using SAPCTL do not require APP-VIPs.

Additionally, this document describes how to configure an Azure VM running Oracle Linux 8 as an SAP NetWeaver primary application server connected to a VM cluster running on Exadata Cloud Infrastructure X9M.

Note: Using a VM cluster node for SAP NetWeaver Application Server ABAP/Java instances is supported only for SAP central services. SAP NetWeaver Application Server ABAP/Java instances must be installed on separate Azure compute instances connected to the same virtual network as the VM cluster but in its own subnet.

Customers generally have multiple options for configuring a network file system (NFS). Consider in advance which option to use, according to your requirements for availability, cost, architecture, and scenario. The most common scenarios are discussed in the “Planning Network File Systems” section.

Note: Configuring Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is always mandatory for storing shared logs and traces on VM cluster nodes, even if you choose to place `/sapmnt` on an NFS mount not exported by the VM cluster nodes. Furthermore, we recommend Oracle ACFS for sharing your installation media for setting up necessary SAP software (for example SAP Host Agent) on the VM cluster nodes.

This document is *not* a full reference for SAP NetWeaver Application Server ABAP/Java. Rather, it is a description of how to plan and implement an SAP NetWeaver-based landscape with Oracle Database@Azure in a supported and verified way.

This document requires the following knowledge:

- You are familiar with the fundamentals of Azure, Exadata Database Service, and OCI. For information, see the following resources:
 - [About Oracle Database@Azure](#)
 - [Overview - Oracle Database@Azure](#)
 - [Exadata Database Service on Dedicated Infrastructure](#)
 - [Getting Started with OCI](#)
- You have advanced administrative skills in SAP NetWeaver Application Server ABAP/Java using Oracle Database and Oracle Linux on Azure. For more information, see the following resources:
 - [Azure Virtual Machines deployment for SAP NetWeaver](#)
 - [SAP NetWeaver product page](#)
 - [SAP on Oracle community page](#)
 - [Oracle Linux](#)
- You are familiar with the documentation for the following products:
 - Oracle Grid Infrastructure 19c and Oracle Database Release 19c
 - Oracle Linux 8 running on Exadata VM cluster nodes
 - Oracle Linux 8 as the OS for running a bastion host or an SAP NetWeaver application server (Oracle Linux is the only supported and certified OS.)
 - SAP NetWeaver 7.x

Warning: Take extra caution when *applying patches to the Oracle Database home directories*. The only approved way to patch Oracle Database homes used for SAP is by using the SAP-provided SAP Bundle Patch for Exadata Database Service. Never apply patches to Oracle Database homes by using the Oracle Cloud Console or CLI. Doing so can result in an unusable and unsupported configuration and might cause an unplanned outage of your SAP environment. For more information, see “SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases.”

Terminology

The following tables define some of the terms used within the contexts of Azure Cloud, Oracle Cloud, Exadata Database Service, and this document. They also define terms related to Azure, Oracle, and SAP work areas that you will likely interact with during deployment and while running SAP NetWeaver Application Server ABAP/Java on a VM cluster.

Table 1: Terms Used Within the Context of Exadata Database Service and This Document

TERM	DEFINITION
Exadata Cloud Infrastructure	The specific physical infrastructure that powers Exadata Database Service. The infrastructure typically consists of two, four, or eight DB servers or Exadata compute nodes, a number of storage servers, plus more components for networking and power supply.
DB server	A database (DB) server is one physical node of the Exadata Cloud Infrastructure, often referred as an Exadata compute node. It is configured as a KVM-based virtualization host (Dom0) and can run multiple VMs.
Storage server	A storage server provides shared disk storage and runs the Exadata Storage Server Software. This software provides unique and powerful technology such as Smart Scan, Smart Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes, and Hybrid Columnar Compression.
VM cluster	A cluster that consists of two or more VMs running on different DB servers.
VM cluster node	A VM with all the components to run Oracle Real Application Clusters (RAC) databases, including the operating system (OS), Oracle Grid Infrastructure software (Clusterware, ASM, ACFS, and so on), and Oracle Database software. A VM cluster node is also referred as an Oracle Database compute node, and might also be referred to as a VM, DomU, or virtual compute node.
Virtual cloud network (VCN)	A virtual, private network set up in Oracle or Microsoft data centers. Although a VCN might be (physically) located in Microsoft data centers, it is still embedded and managed by OCI. A VCN closely resembles a traditional network, with firewall rules and specific types of communication gateways that you can choose to use. A VCN resides in a single OCI region and covers one or more CIDR blocks. For more information, see Networking Overview .
VM Cluster Node Subsetting	Allows a VM cluster to span two or more DB servers (up to the maximum number of DB servers available in your Exadata Infrastructure). VM Cluster Node Subsetting lets you add new VM cluster nodes to an existing VM cluster if your workload or high availability requirements demand additional resources. VM Cluster Node Subsetting also lets you reduce the number of VM cluster nodes in a VM cluster down to a minimum of two.
Node-IP	The fixed IP address of a VM cluster node. It is not a virtual IP address and cannot move to another VM cluster node.
Node-VIP	A virtual IP address of a VM cluster node. It belongs to its VM cluster node and can only move to another VM cluster node if the owning node is unavailable.
SCAN hostname SCAN-VIP	Oracle Clusterware ensures high availability for client connections through three Single Client Access Name (SCAN) listeners bound to three virtual IP addresses (VIPs) distributed across the available VM cluster nodes. One virtual (SCAN) hostname in DNS is configured to return three IP addresses. For more information, see Understanding SCANS .

Table 2: Terms Specific to Azure Work Areas

AZURE WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
Azure portal	Use to deploy Oracle Database@Azure and VM clusters as well as native Azure cloud resources.	See the Azure portal documentation .
Virtual Network (VNet)	A virtual, private network that you set up in Azure data centers. It closely resembles a traditional network, with firewall rules and specific types of communication gateways that you can choose to use.	See the Virtual Network documentation .
Virtual Machine (VM) or Azure VM	In the context of this document, Azure VMs are used as bastion hosts or to run SAP NetWeaver application servers.	See Virtual machines in Azure .

Table 3: Terms Specific to Oracle Work Areas

ORACLE WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
Oracle Cloud Console	Use to manage all the resources in OCI, for example, the deployment of VCNs and VM clusters for Exadata Database Service.	Register at cloud.oracle.com .
Oracle Linux	OS used on all VM cluster nodes and Azure compute instances.	Oracle Linux 8 on VM cluster nodes is managed from the Oracle Cloud Console. Oracle Linux on Azure compute instances must be managed from the Azure portal. For Oracle Linux 7 and 8, see SAP Notes 2069760 and 2936683 .
Oracle Grid Infrastructure	Manage the HA of all database services, ACFS, and NFS by using tools like SRVCTL and CRSCCTL, as well as application virtual IP addresses (APP-VIPs).	For HA-NFS with Oracle Grid Infrastructure, see My Oracle Support Note 1934030.1: ODA (Oracle Database Appliance): HowTo export ACFS (cloudfs) using HANFS .
Oracle Database software	Use MOPatch or OPatch and SAP Bundle Patches to patch existing Oracle Database homes.	SAP Bundle Patch installation is documented in the readme file.
Oracle Automatic Storage Management (ASM)	Manage ASM disk groups by using ASMCMD and SQL*Plus.	None
Oracle Transparent Data Encryption (TDE)	Manage encryption wallets and encryption keys.	See SAP Notes 2591575 and 2799991 .
Oracle Database instances	Manage the SAP database and Oracle initialization parameters recommended by SAP.	For SAP-required Oracle initialization parameters, see SAP Notes 2799900 , 2470660 , and 2470718 .
Oracle Recovery Manager (RMAN)	Backup, restore, and recover your SAP database.	None

Table 4: Terms Specific to SAP Work Areas

SAP WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
SAP Maintenance Planner	Create a <code>stack.xml</code> file for SAP Software Provisioning Manager (SWPM) and choose the SAP software components that you want to install.	See Maintenance Planner – User Guide .
(Mandatory) SAP Software Provisioning Manager (SWPM)	Use for SWPM-based host preparation and to install your ABAP system central services (ASCS), enqueue replication server (ERS), primary application server (PAS), and SAP database instance.	Always use the latest version of SWPM to avoid issues with new versions of Oracle Database software and new versions of Oracle Linux not supported in older versions of SWPM.
SAP NetWeaver software stack	Modify SAP instance profiles and configure RFC connections and SAP transaction code DB13.	SAP instance profiles must be adjusted to configure the correct number of work processes and SAP HA components if using SAPCTL.
SAProuter	Set up and configure SAProuter.	Customers must configure SAProuter at least for SAP EarlyWatch.
SAP GUI	Install SAP GUI components.	None
SAP HA using SAPCTL (optional)	Configure SAP HA for SAP ASCS and ERS by integrating with Oracle Clusterware.	See SAP Note 1496927 to download the SAPCTL software package (includes documentation).
BR*Tools (optional)	Back up, restore, and recover your SAP database.	See SAP Notes 1598594 , 113747 , and 776505 .

Planning Network File Systems

One of the most important aspects to consider when choosing a network file system (NFS) implementation is whether you are planning to implement SAP HA.

- If you are *not* planning to implement SAP HA managed by a clustering solution such as Oracle Clusterware, you can either use an Azure-provided NFS service or deploy an Azure VM and configure your own NFS server for your shared SAP-specific directories, for example, `/sapmnt`.
- If you are planning to implement SAP HA, you need a highly available shared file system to share data from SAP central services (ASCS/SCS) instances between multiple cluster nodes where the ASCS/SCS and ERS instances can run. You can achieve this either by using Oracle ACFS in combination with application-specific virtual IPs (APP-VIPs) for ASCS/SCS and ERS managed by Oracle Clusterware, or by implementing one of the other certified SAP HA-solutions within Azure.

This document describes only the implementation of SAP HA using Oracle Clusterware. It also describes how to implement an additional HA-NFS share for `/sapmnt` and for sharing installation media across the VM cluster nodes as well as the Azure compute nodes, which is completely optional. For other implementations of SAP HA and HA-NFS, see the Azure documentation.

The recommended NFS services available in Azure are Azure Files and Azure NetApp Files. Azure Files-based NFS exports cannot be mounted on Exadata Cloud Infrastructure VM cluster nodes, but Azure NetApp Files-based NFS exports can be mounted from Azure compute instances and from Exadata Cloud Infrastructure VM cluster nodes.

If you want to use Azure NetApp Files on your Exadata Cloud Infrastructure VM cluster nodes, you must follow these steps:

1. Configure an additional subnet in your VNet with a delegation to service `Microsoft.Netapp/volumes`.
2. Create the necessary capacity pool and volumes.
3. Mount them on all VM cluster nodes under the exact same mountpoint by using the following mount options `rw,hard,noatime,rsize=262144,wspace=262144,vers=3,tcp <IP>:/<export>/<mountpoint>`.

Although Azure NetApp Files is a good option for RMAN database backups and other tasks (for example, for database reorgs), be aware that all network traffic will use the client network of the VM cluster and not the backup network. If you want to use the backup network of a VM cluster to perform database backups, you must use Object Storage buckets or NFS in OCI.

Overview of Oracle Database@Azure

Oracle Exadata Database Service enables full-featured Oracle databases to run on dedicated Oracle Exadata Cloud Infrastructure X9M in the Azure cloud, all preconfigured according to best practices that have been proven at thousands of mission-critical Exadata sites around the world.

Depending on the subscription model, you can have full access to the features and operations available in the following products, but with Oracle owning and managing the Exadata infrastructure:

- Oracle Grid Infrastructure
- Oracle Database, including Oracle Real Application Clusters (RAC)
- Oracle Automatic Storage Management Cluster File System (Oracle ACFS)
- SAP Infrastructure components for installation (Software Provisioning Manager)
- Database Administration (BR*Tools)
- SAP High Availability (HA) for SAP Central Services (SCS)

SAP NetWeaver Database on Oracle Exadata Database Service

All options and features of Oracle Database 19c certified for on-premises deployments of SAP NetWeaver are also available for Oracle Database@Azure. These features include Oracle RAC, Oracle Automatic Storage Management (ASM), and Oracle Database In-Memory for on-premises deployments of SAP NetWeaver.

Note: Oracle Database versions 11.2, 12.1, 12.2, and 18c are not supported with SAP NetWeaver and Oracle Database@Azure.

Regions and Availability Domains in OCI

OCI is physically hosted in [regions and availability domains](#). A *region* is a localized geographic area. *Availability domains* are one or more data centers located within a region. Most OCI resources are bound either to a particular region, such as a virtual cloud network (VCN), or to an availability domain, such as a compute instance.

Availability domains are isolated from each other, fault tolerant, and unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain is unlikely to impact the availability of the others.

All the availability domains in a region are connected to each other by a low-latency, high-bandwidth network. This connection makes it possible to provide high-availability connectivity to the internet and customer premises, and to build replicated systems in multiple availability domains for both high availability and disaster recovery.

Regions are independent of other regions and can be separated by vast distances. Generally, an application should be deployed in the region where it is most heavily used, because using nearby resources is faster than using distant resources.

For all SAP environments, Azure VMs that are deployed as SAP NetWeaver application servers must be located in the same region, and preferably in the same availability domain, as the corresponding Exadata infrastructure.

Hybrid deployments between on-premises and cloud infrastructures are not supported because of network latency.

Regions and Availability Zones in Azure

A *region* in Azure corresponds to a *region* in OCI, and an *availability zone* in Azure corresponds to an *availability domain* in OCI. Similarly, regions are understood to be distant locations, while availability zones are data centers within the same region that are geographically very close to each other.

Note: To guarantee optimal network latency, all Azure VMs deployed for SAP must be in the same region, availability zone, and VNet as the VM cluster where the SAP database is running. For this reason, running those Azure VMs in other regions or availability zones than their corresponding VM cluster is not supported. To ensure the lowest latency possible, see [SAP Note 1100926 FAQ: Network Performance](#).

Virtualization and Databases with Exadata Cloud Infrastructure X9M

Exadata Cloud Infrastructure X9M enables full-featured Oracle databases to run on VM cluster nodes. A VM cluster consists of two or more VM cluster nodes. A VM cluster can span two or more DB servers of the underlying Exadata Cloud Infrastructure X9M. For a quarter rack, a VM cluster always spans two nodes. For a scalable Exadata Cloud Infrastructure X9M, a VM cluster can span two to eight nodes. In a VM cluster node (DomU), the Oracle Grid Infrastructure and Oracle Database instances are deployed. Multiple databases can run within each of the VM clusters.

Consider the following important points for virtualization of SAP NetWeaver with Exadata Cloud Infrastructure X9M:

- The available free space of the file systems on the internal disk drives should be monitored constantly. We recommend making no changes to the structure and the size of the internal file systems (such as `/`, `/var`, and `/tmp`) that are created during VM cluster deployment. If changes are needed, contact Oracle Cloud operations.
- The amount of local storage per VM determines how many Oracle Database homes (`ORACLE_HOME`) for a database installation can be deployed within a VM. Oracle Database homes grow over time, and for ongoing maintenance using the required SAP Bundle Patches, multiple copies of an Oracle Database home are needed ([SAP Note 2799959](#)). For SAP NetWeaver deployments, 200 GB of disk space is required for an Oracle Database home over a time period of three years, excluding the fixed overhead for the VM of 184 GB required for the OS and Grid Infrastructure.

The maximum number of SAP NetWeaver databases on a VM cluster is determined by how the Oracle Database homes are configured. If an SAP NetWeaver database uses its own Oracle Database home, then only the maximum number of four SAP NetWeaver databases can be used. Using a shared Oracle Database home for multiple SAP NetWeaver databases of the same Oracle Database version allows the use of more SAP NetWeaver databases on a VM cluster. The number of SAP NetWeaver databases using a shared Oracle Database home that can be configured on a VM cluster depends on the sizes of the databases and the amount of main memory required for the database instances running on the VM cluster nodes. Proper careful sizing is required to determine which Exadata Cloud Infrastructure X9M based system is required for running an SAP NetWeaver landscape.

- Deploying SAP and non-SAP databases on the same VM cluster is not supported because of the specific configuration, administration, and maintenance aspects of SAP databases.
- Deploying SAP NetWeaver and non-SAP NetWeaver databases (such as SAP Business Objects BI Platform) on the same VM cluster is not supported because of the different configuration, administration, and maintenance aspects of SAP NetWeaver and non-SAP NetWeaver databases.

VM Cluster Scaling Options

Exadata Cloud Infrastructure offers several options for scaling VM clusters up or down to match your requirements for performance, storage, and high availability:

- Change the number of OCPUs allocated for each VM cluster node.
- Increase or decrease the amount of memory (RAM) allocated for each VM cluster node.
- Resize the local file system storage allocated for each VM cluster node.
- Increase or decrease the amount of Exadata storage (ASM) allocated to the VM cluster.
- Add a complete VM cluster node to, or remove one from, the VM cluster (within the given limits).

Note: Adding or removing complete VM cluster nodes is the most complex option for allocating or reducing compute resources. This option always requires numerous manual configuration steps and might also require downtime of your SAP system. Consider it only after scaling OCPUs and memory.

Scaling the physical Exadata Cloud Infrastructure (for example, by adding or removing DB servers or Exadata Storage servers) is not supported.

Oracle Database for SAP NetWeaver with Exadata Cloud Infrastructure X9M

All options and features certified for on-premises deployments of SAP NetWeaver with Oracle Database 19c are available for Oracle Database@Azure using Exadata Cloud Infrastructure X9M. These features include Oracle RAC, Oracle ASM, and Oracle Database In-Memory for on-premises deployments of SAP NetWeaver.

We do *not* recommend using single-instance databases for SAP NetWeaver on Exadata Cloud Infrastructure X9M. Use only Oracle RAC databases.

Note: Oracle Database versions 11.2, 12.1, 12.2, and 18c are not supported with SAP NetWeaver on Exadata Cloud Infrastructure X9M. Oracle Autonomous Database is not certified and supported with SAP NetWeaver.

Exadata Cloud Infrastructure X9M

Each Exadata Cloud Infrastructure X9M system configuration contains a predefined number of DB servers (or KVM-based virtualization hosts) and Exadata storage servers. These servers are connected by a high-speed, low-latency, 200-Gbit aggregated bandwidth of active-active Remote Direct Memory Access over Converged Ethernet (RoCE) fabric and intelligent Exadata software. Multiple, full-fledged Oracle RAC instances can be configured to run on VM clusters hosted on Exadata Cloud Infrastructure X9M.

Currently, Oracle offers the following configurations for Exadata Cloud Infrastructure X9M.

Table 4: Exadata Cloud Infrastructure X9M Configurations

SYSTEM	NUMBER OF COMPUTE NODES	NUMBER OF EXADATA STORAGE SERVERS
Quarter Rack X9M	2 nodes (4 to 252 OCPUs and total 2,780 GB RAM)	3
Half Rack X9M	4 nodes (8 to 504 OCPUs and total 5,560 GB RAM)	6
Full Rack X9M	8 nodes (16 to 1,008 OCPUs and total 11,120 GB RAM)	12
Elastic X9M	2 to 32 nodes (4 to 4,032 OCPUs and 2,780 to 44,480 GB RAM)	3 to 64

At the time this document was published, SAP might not have certified all the preceding systems. For up-to-date information about the certification status of all Exadata Cloud Infrastructure systems, see [SAP Note 2614028](#).

Exadata Cloud Infrastructure allows elastic scaling of the deployed VM cluster, which enables flexibility in the allocation of compute (CPU, memory, local storage) resources.

Note: In an Exadata Cloud Infrastructure elastic scaling configuration, a *minimum* of two OCPUs per VM per compute node running an SAP NetWeaver-based database workload is required.

Each VM cluster node (DomU) is running on a DB server (KVM-based virtualization host—Dom0). You have root privileges for the VM cluster nodes (DomU) and DBA privileges on the Oracle databases. You can configure the VM cluster nodes according to your requirements. You can also run additional agent software and SAP infrastructure components on the VM cluster nodes to conform to business standards or security monitoring requirements.

On a VM cluster, you can create numerous database deployments for different SAP NetWeaver based applications. The number of production databases for SAP NetWeaver systems is determined by the Exadata Cloud Infrastructure system type, the number of VMs deployed, and how many different database versions are being used.

Note: We do not support using the VM cluster nodes for SAP NetWeaver Application Server ABAP/Java instances *except for* SAP central services.

However, you do not have administrative access to the following items, which are managed entirely by Oracle:

- Exadata Cloud Infrastructure components, including the physical compute node hardware, network switches, power distribution units (PDUs), and integrated lights-out management (ILOM) interfaces
- Exadata storage servers

Exadata Database Service Licenses

Exadata Database Service is available through two flexible subscription offerings for SAP customers:

- License included
- Exadata Database Service Bring Your Own License (BYOL), which is also valid for SAP ASFU licenses

License Included

This subscription model includes all the features of Oracle Database Enterprise Edition, plus all Oracle Database Enterprise Manager Packs and all Database Enterprise Edition Options. The following industry-leading capabilities are included:

- Database In-Memory
- Real Application Clusters (RAC)
- Active Data Guard
- Automatic Storage Management (ASM)
- Partitioning
- Advanced Compression
- Advanced Security
- Database Vault
- Real Application Testing
- OLAP
- Advanced Analytics
- Spatial and Graph

Oracle Multitenant is also included in an Exadata Database Service PaaS subscription. Oracle Multitenant enables high consolidation density, rapid provisioning and cloning, efficient patching and upgrades, and simplified database management.

Note: Oracle Multitenant is not supported for databases of SAP NetWeaver systems.

This subscription model is ideal for customers who don't have existing Oracle Database licenses or who want to use Oracle Database features beyond their current licenses.

Exadata Database Service Bring Your Own License (BYOL)

If you have bought Oracle Database licenses from SAP (ASFU), you can transfer them to Exadata Database Service. Notify SAP that you intend to bring your own license.

The same applies for licenses that you have bought from Oracle (Full Use, FU). If you have enough licenses, you can also transfer them from on-premises to Exadata Database Service. To ensure that the number of shapes, processors, and cores is correct, we recommend that you check with your Oracle sales manager or local license sales contact. They can help you set up the correct licensing.

Also, you can benefit from Oracle's BYOL to PaaS program for Oracle Exadata Database Service, with both ASFU and FU licenses. For more information about Oracle's BYOL to PaaS program, see the following resources:

- [Frequently asked questions: Oracle Bring Your Own License \(BYOL\)](#)
- [Oracle PaaS and IaaS Universal Credits Service Descriptions](#)

Exadata Database Service BYOL is designed to minimize costs when migrating to the cloud. In a BYOL model, customers can deploy their existing Oracle Enterprise Edition and Database Option licenses on Exadata Database Service. Standard Edition is not supported on any Exadata Database Service.

When a customer brings a Database Enterprise Edition license entitlement to Exadata Database Service, they are granted the rights to use the following database options without having on-premises license entitlements for those options:

- Oracle Transparent Data Encryption (TDE)
- Diagnostics Pack
- Tuning Pack
- Data Masking and Subsetting Pack
- Real Application Testing

The Exadata system software is also included in a BYOL subscription, so BYOL customers do not have to bring a license entitlement for the Exadata system software.

Planning Your Oracle Database@Azure Service for SAP NetWeaver Application Server ABAP/Java

Use the information in this section to plan your use of Oracle Database@Azure for your SAP NetWeaver Application Server ABAP/Java deployment.

Prerequisites

Work with your Oracle account team to get an offer for Exadata Cloud Infrastructure in the Azure region and availability zone of your choice. Follow the official [Oracle Database@Azure](#) documentation to find out how onboarding with Oracle Database@Azure works and how to deploy your claimed Exadata Cloud Infrastructure in Azure. Deploying Exadata Cloud Infrastructure in Azure is a generic task and well described at [Provisioning Exadata Infrastructure](#). The following sections assume that Exadata Cloud Infrastructure has been deployed accordingly.

Deployment Restrictions with SAP NetWeaver for Oracle Database@Azure

The following restrictions apply:

- Oracle Database 19c is the only supported database release.
- Oracle Autonomous Database is not certified with SAP NetWeaver and is not supported.
- Oracle Multitenant is not certified with SAP NetWeaver and is not supported.
- Non-Unicode deployments of SAP NetWeaver Application Server ABAP/Java are not supported.
- Oracle Linux is the only supported OS for SAP NetWeaver Server ABAP/Java on Azure VMs.
- SAP NetWeaver application servers on VM cluster nodes are not supported.
- Although SAP NetWeaver application servers on VM cluster nodes are not supported, SAP central services for SAP NetWeaver Application Server ABAP/Java can be deployed on VM cluster nodes.
- All databases must be created by using SAP Software Provisioning Manager (SWPM) either as a new database or as an SAP system copy of an existing database. Alternatively, existing SAP databases can be migrated to a VM cluster, for example, by using Oracle Recovery Manager (RMAN). Migrations of existing databases are discussed in a later section.
- Data encryption is mandatory for all Oracle databases in OCI. Not using data encryption causes certain management operations, such as adding tablespaces, to fail. Oracle 19c databases require SWPM with patch level 35 or later for Oracle Database 19c and TDE support.

- The strong password policy on VM cluster nodes must be changed for an SAP installation and then reverted after the installed is finished.
- Hostnames, including hostnames for virtual IPVMs, SCANs, and backup networks, must not exceed 13 characters.
- The Oracle Grid Infrastructure home is owned by the `grid` OS user. All Oracle Database homes being created must be owned by the `oracle` OS user. This is also important for patching.

Documentation

Determine the supported combination of Oracle Linux and Oracle Database for your planned SAP product by using the [SAP Product Availability Matrix](#) (PAM). Ensure that you are familiar with the relevant SAP NetWeaver master and installation guides and the referenced SAP notes within. To find planning, installation, patching, and operation documentation for your task, see the [SAP NetWeaver Guide Finder](#).

Become familiar with the product documentation for all the components of your stack: Azure, OCI, Oracle Linux, Oracle Database, and SAP NetWeaver Application Server ABAP/Java.

Workload Size

Estimate the size needed for your SAP installation by using the SAP Quick Sizer tool, and calculate the Exadata Database Service configuration needed for your SAP workload. For SAPS numbers, see [SAP Note 2614028](#).

Note: The SAPS numbers listed in the note are only for a performance indication and have not been achieved by using a high-performance benchmark.

The following table shows the current certified instance type by SAP. Check [SAP Note 2614028](#) for the most current certification status.

Table 5: Certified Exadata Database Service Instance Types

INSTANCE TYPE	OCPU	MEMORY	TOTAL USABLE STORAGE CAPACITY	NUMBER OF EXADATA STORAGE SERVERS
Exadata.X9M (Elastic, up to 4 database servers/nodes)	4 to 504	2,780 to 5,560 GB	190 to 4,070 TB	3 to 64

To run an SAP NetWeaver workload, minimum numbers of OCPUs are needed.

Presales and consulting teams from Oracle can help you size your planned SAP landscape in the cloud.

Planning the SAP Deployment

Running SAP NetWeaver Application Server ABAP/Java with Exadata Database Service requires deployment of numerous cloud resources, preparation of hosts and services, and installation and configuration of the SAP NetWeaver stack. The following illustration gives an overview of the major steps, which are covered in this document.

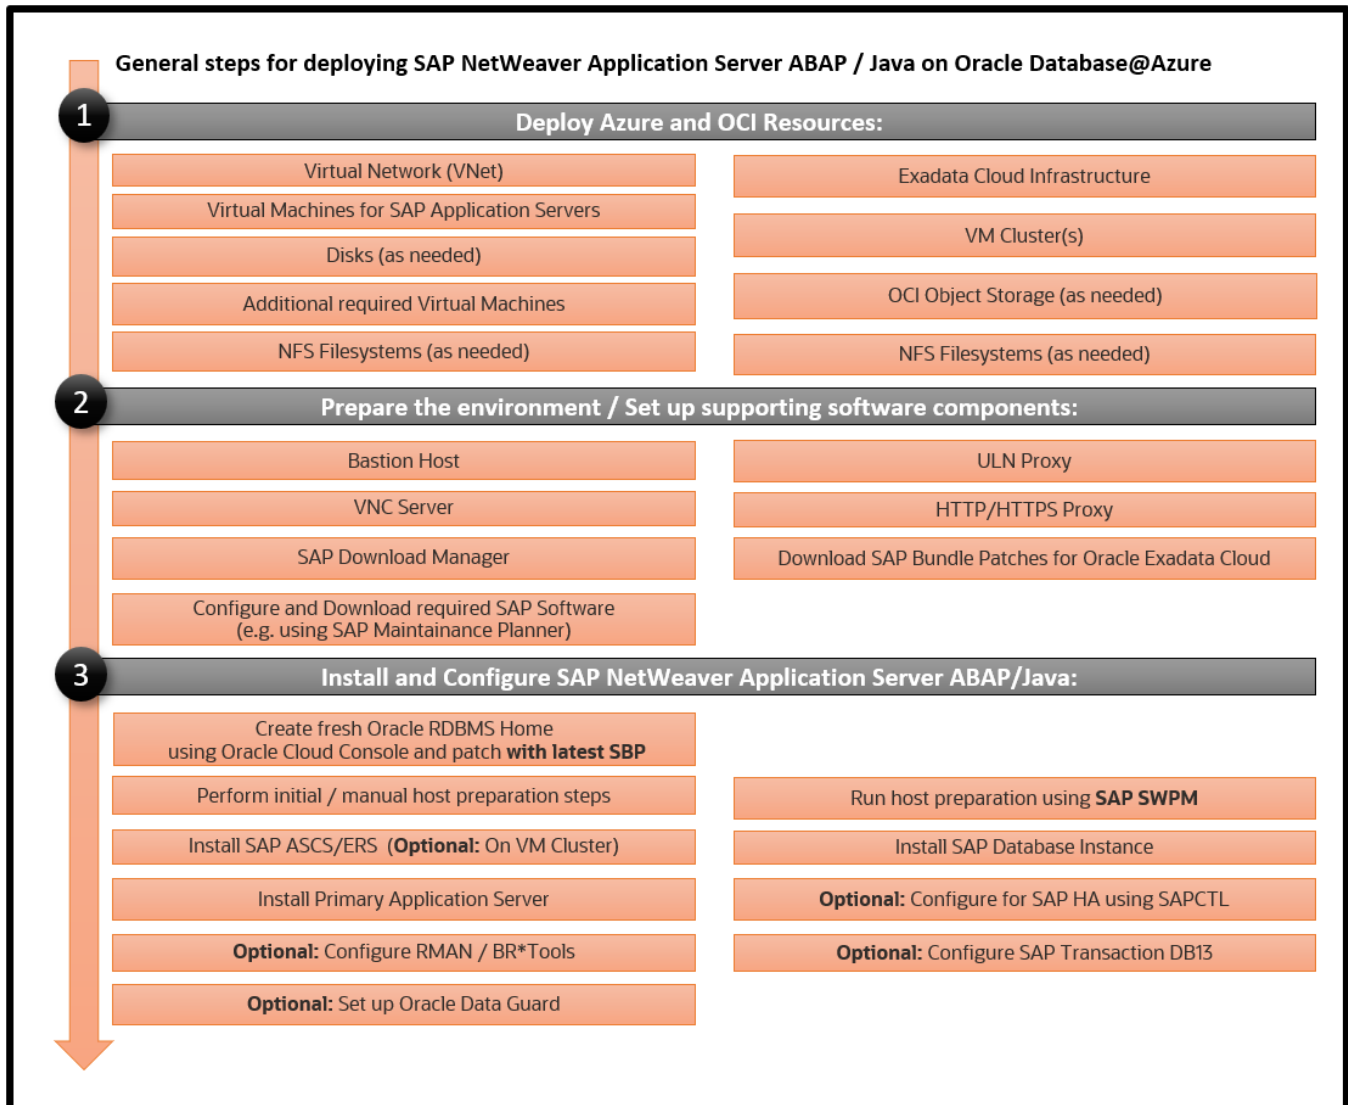


Figure 1: General Steps for Deploying NetWeaver Application Server ABAP/Java with Exadata Database Service

Use the information in this section to plan your SAP NetWeaver Application Server ABAP/Java deployment on Oracle Database@Azure based on Exadata Cloud Infrastructure X9M.

This document designs a minimal SAP landscape that consists of an SAP NetWeaver database using a virtual Oracle RAC database cluster (also referred as a VM cluster within this document) running on Exadata Cloud Infrastructure X9M that is connected to a single SAP NetWeaver application server (primary application server, or PAS) on a separate Azure host.

A real-world SAP landscape is considered more complex and might consist of multiple SAP NetWeaver application servers from one or more SAP systems connected to multiple instances of one or more RAC databases deployed on one or more VM clusters. Additionally, it would implement a VPN Connect setup and internet access through SAProuter.

The proper setup of the virtual network (VNet) with its address range, subnets, security, and Oracle delegation requires detailed planning for how customer requirements can be addressed within Azure. The landscape described in this document uses a manually configured VNet, to help customers deploy SAP on Oracle Database@Azure in a real-world scenario.

Exadata Database Service Requirements

The following items are required:

- **Sign-in credentials to the Azure portal:** The Azure portal is used to manage resources in Azure.
- **Sign-in credentials to the [Oracle Cloud Console](#):** The Oracle Cloud Console is used to manage resources in OCI and perform tasks not supported in the Azure portal.
- **SSH key pairs:** You need SSH key pairs to deploy and access Azure and OCI resources, VM cluster nodes, and at least one certified Azure host. See [Creating a Key Pair](#). We recommend password protection of key pairs.
- **Oracle SAP Bundle Patches for Exadata:** For information, see the “SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases” section later in this document.
- **SAP NetWeaver installation media:** You need the required media and versions of SAP SWPM, SAP Kernel, and Installation Exports, depending on your installation scenario. You might need access to SAP Marketplace to download SAP software. However, this document assumes that you have already selected and downloaded all the required software components and have them available, for example, on a network file system.
- **Passwords for various resources**

Additionally, you must have completed the following actions:

- Ensure that you have completed onboarding with Oracle Database@Azure and that you have subscribed and deployed an Exadata Cloud Infrastructure as mentioned earlier in this document and at [About Oracle Database@Azure](#).
- Decide whether to use OCI default domain naming in the form <hostname>.<subnet-DNS-label>.<VCN-DNS-label>.oraclevcn.com or your own domain-naming schema (for example, <hostname>.acme.com) for all the VM cluster nodes and the relevant Azure compute instances.

General Installation Overview

This section provides an overview of the installation and configuration steps outlined in this document. The detailed steps follow in the later sections.

After you deploy Exadata Cloud Infrastructure in Azure, you deploy a VNet and a VM cluster. Additionally, you will likely deploy Azure compute instances to run your SAP NetWeaver application servers, and if necessary, create network file systems (NFSs) or disks that you want to attach to your Azure compute instances. Finally, you might want to use the [Oracle Cloud Console](#) to create Object Storage buckets for your database backups.

When these steps are completed, the following resources are in place:

- A VNet that meets the requirements for deploying a VM cluster
- IP addresses (Node-IPs, Node-VIPs, SCAN-VIPS and APP-VIPs)
- Hostnames (node specific, as well as SCAN)
- IP addresses of all compute nodes, for example, for the SAP NetWeaver application servers
- SSH-based access to all the deployed compute resources

On the VM cluster nodes, you must make changes to local and shared file systems. The necessary steps vary among the root, oracle, and opc user accounts. On the Azure compute instances, you must make changes to hostname settings, time zones, local file systems, and network file systems, plus a few more, before SWPM can be used to install SAP software.

Before any of the VM cluster nodes in the VM cluster can be used for SAP software—for example, SAP Host Agent, ASCS, or ERS—or for an SAP NetWeaver database installation with SWPM, you must also perform some configuration steps to be able to run SAP SWPM host preparation.

Host preparation has the following steps:

1. Manual host preparation, which consists of some manual steps that must be completed by the administrator to prepare the host so that SWPM can be started on the host
2. SWPM host preparation, which is a functionality of SAP SWPM that prepares the host by making further changes, for example, by creating SAP-specific groups and users and by installing the SAP Host Agent
SWPM must be run on each VM cluster node of the VM cluster, so we recommend using shared media to simplify installation.

The next major steps are as follows:

3. Install Oracle RAC database on the VM cluster nodes. Optionally, to set up SAP HA on the VM cluster, also install the SAP ASCS instance and SAP ERS instance on the VM cluster nodes.
 - Manually prepare all VM cluster nodes for running SWPM.
 - Run SWPM host preparation on all VM cluster nodes.
 - (Optional) To implement SAP HA by using Oracle Grid Infrastructure with SAPCTL, install the SAP ASCS instance on a shared location on the VM cluster nodes of the VM cluster and complete the required post configuration steps.
 - Create an Oracle RAC database on the VM cluster by using SWPM or by migrating an existing database.
 - (Optional) To implement SAP HA by using Oracle Grid Infrastructure with SAPCTL, install SAP ERS on each VM cluster node of the VM cluster.
 - Verify that the database is Oracle RAC.
4. Install the primary application server (PAS) instance on an Azure compute instance.
 - Prepare the host for running SWPM.
 - Run SWPM to install PAS.
5. Configure the SAP GUI to verify that the SAP system is accessible.
6. Finalize the installation.

Implementing the Deployment

This section provides the steps for implementing your planned deployment of SAP NetWeaver Application Server ABAP/Java on Oracle Database@Azure by using the Azure portal and, if necessary, the Oracle Cloud Console. You can also use automation, which provides the advantage of repeatability, while the Console provides immediate provisioning and a human-friendly UI. Using automation to implement the deployment is not part of this document.

Get Your Azure and OCI Accounts

To get your accounts, work with your Microsoft and Oracle account teams. Ensure that federated single sign-on is enabled by using Microsoft Entra ID to be able to navigate to a resource in OCI whenever you need to perform tasks not supported in the Azure portal.

Deploy the Cloud Resources

This section provide steps for the following tasks:

- (Optional) Deploy a private view with an empty private DNS zone in OCI
- Deploy the VNet and a VM cluster in Azure
- (Optional) Deploy application-specific virtual IP addresses (APP-VIPs) and hostnames in OCI
- Deploy Azure compute instances and disks as needed in Azure
- Configure and deploy DNS forward and reverse lookup zones

Deploy a Private View with an Empty Private DNS Zone in OCI

As mentioned earlier in this document, you must decide whether to use OCI default domain naming in the form of `<hostname>.<subnet-DNS-label>.<VCN-DNS-label>.oraclevcn.com` or your own domain naming scheme (for example, `<hostname>.sap.acme.com`) for VM cluster nodes and Azure compute instances.

If you prefer to use your own domain naming scheme, you must deploy a *private view* with an empty *zone* for the domain of your choice in OCI as a first step. In this example, we use `sap.acme.com` as the domain name. If you want to use OCI default domain naming, you can skip this step. OCI automatically creates a default private view named according to the VCN with an automatically generated name.

Perform the following steps to deploy a private view and a zone in the private view in OCI *only* if you choose to use your own domain naming scheme.

1. Sign in to the Azure portal, navigate to your Exadata Infrastructure, and then click **Go to OCI**.
The Oracle Cloud Console opens with in the correct tenancy and compartment selected.
2. In the main navigation menu, select **Networking**, and then select **Private views**.
3. Click **Create private view**.
4. Enter a name for the view (for example, `SAP_ACME_COM`), keep the preselected compartment, then click **Create**.
5. Click on the name of the new view, and then click **Create zone**.
6. Enter a domain name for the zone (for example, `sap.acme.com`), keep the preselected compartment, and click **Create**.

OCI generates a private DNS zone that contains all the A records for DNS and additional separate private DNS zones for the Node-IPs and Node-VIPs of the VM cluster during deployment.

Deploy a VNet in Azure

1. In the Azure portal, click **Create Resource**, search for **Virtual Network**, and then click **Create**.
2. On the **Basics** tab, perform the following steps:
 - A. Ensure that the correct Azure subscription is selected.
 - B. Select a resource group for the VNet.
 - C. Enter a name, for example, `vn-sap-acme-com`, for the VNet.
 - D. Select the region to create the VNet in.
3. On the **Security** tab, keep the defaults.

4. On the **IP Addresses** tab, define the IP range of the VNet.

A best practice is to define three subnets in a VNet, so ensure that your total address range is large enough. In this document, we name those subnets `sn-bastion`, `sn-app`, and `sn-client`.

Add three private regional subnets of the appropriate size:

- `sn-bastion` to provide access from a bastion host to the VNet
- `sn-app` for the SAP NetWeaver application servers
- `sn-client` for the Exadata client network

The Exadata client subnet is required to launch a new VM cluster. Although you can deploy multiple VM clusters on the same VNet, we strongly recommend creating an additional VNet for each new VM cluster. If you decide to deploy more than one VM cluster per VNet, remember that a VM cluster allocates numerous IP addresses and ensure that you have enough free addresses left on the subnet (`sn-client`).

For the Exadata client subnet, select the delegation `Oracle.Database/networkAttachments` so that the IP addresses and hostnames on this subnet can be managed by OCI.

Note: Only one delegation of type `Oracle.Database/networkAttachments` is supported per VNet, and the IP addresses on the `sn-client` subnet are exclusively managed by OCI. That is, you cannot assign any IP addresses for other resources from the Azure side, for example, for Azure VMs.

5. On the **Tags** tab, add resource tags as needed.
6. On the **Review and Create** tab, review the data entered, and then click **Create**.

Note: Azure automatically routes network traffic between the subnets of a VNet. Do not add any Azure compute instances to the Exadata client subnet (`sn-client`). Additional Azure compute instances for SAP NetWeaver application servers should be attached to the subnet dedicated to SAP NetWeaver application servers (`sn-app`).

Deploy a VM Cluster in Azure

For detailed instructions, see the [Provisioning an Exadata VM Cluster](#).

1. In the Azure portal, select the Exadata Cloud Infrastructure in which you want to create the VM cluster.
2. Select **Settings**, select **Exadata VM Clusters**, and then click **Create VM Cluster**.
3. On the **Basics** tab, provide the necessary values. Select **Grid Infrastructure 19.0.0.0** and the correct time zone for your deployment, and provide your SSH public key.

Note: All Azure compute instances used for SAP NetWeaver application servers *must* be configured to use the same time zone as the VM cluster; otherwise, the SAP system cannot run.

4. On the **Configuration** tab, select the required number of OCPUs, the amount of memory and local disk space, and the amount of usable Exadata storage. If you want to perform backups to ASM, select **Use local backups**.
5. On the **Networking** tab, select the VNet, the client subnet, and the CIDR of the backup subnet. Ensure that the hostname prefix has a maximum length of 2 characters, so that automatically generated hostnames do not exceed the maximum length allowed for SAP.

If you want to specify your own custom DNS domain for the cluster, select **Custom DNS** and then select the private view and private DNS zone that you created earlier in OCI. Enter your own domain name, for example, `sap.acme.com`. If you do not select **Custom DNS**, OCI automatically generates a domain for the cluster.

6. On the **Diagnostics Collection** tab, keep the defaults.
7. On the **Consent** tab, accept the terms of service.
8. On the **Tags** tab, provide tags and values as needed.
9. On the **Review+Create** tab, review the entered information, and click **Create** to start VM cluster deployment.

Note: Deploying a VM cluster can take several hours.

Deploy Application-Specific Virtual IP Addresses (APP-VIPs)

You can configure APP-VIPs if you want to configure and export an ACFS file system as highly available NFS (HA-NFS) for the mountpoint `/sapmnt`, or you want to configure SAP high availability (SAP HA) for ASCS and ERS with SAPCTL. For those services, APP-VIPs are required.

You must define these APP-VIPs in the Oracle Cloud Console and then transfer them to the VNet in Azure. At the time this document was written, transferring APP-VIPs is a manual step, and you must file a ticket with Oracle Cloud operations team to start the process. Later, these APP-VIPs are added as cluster managed resources to Oracle Clusterware to make them highly available.

1. In the Azure portal, select the VM cluster that you want to add the APP-VIPs to.
2. Click **Go to OCI** and continue your login to the Oracle Cloud Console by choosing Microsoft Entra ID for authentication. The details page of the VM cluster is displayed.
3. Select **Virtual IP Addresses**, click **Attach Virtual IP Address**, and then enter the following values:
 - The “primary” subnet (`sn-client`).
 - The hostname of the virtual IP address.
 - (Optional) The virtual IP address to use. If you don’t select an IP address, OCI selects a free one.
 - The VM cluster node to which the APP-VIP is attached first.
4. Click **Attach**.
5. Repeat steps 3 and 4 for all APP-VIPs that you want to configure.

The following example shows three configured APP-VIPs (with their corresponding virtual hostnames), one for HA-NFS, one for SAP ASCS, and one for SAP ERS.

Attach virtual IP address						
Name ⓘ ▲	State	Virtual IP address	Subnet	Virtual Machines	Fully qualified domain name	
ascsmf	● Available	172.19.20.208	primary-subnet-1715883932982	bb-xuwwj1	...sap.acme.com	Show Copy
ersmf	● Available	172.19.20.69	primary-subnet-1715883932982	bb-xuwwj1	...sap.acme.com	Show Copy
nfsmf	● Available	172.19.20.236	primary-subnet-1715883932982	bb-xuwwj1	...sap.acme.com	Show Copy

Figure 2: Configured Application-Specific Virtual IP Addresses

Deploy an Azure Compute Instance

Deploy an Azure compute instance, of a type supported by SAP, that will host the SAP primary application server (PAS). Ensure that it is attached to the same VNet as the VM cluster nodes but to the subnet created for SAP NetWeaver application servers (for example, `sn-app`).

1. In the Azure portal, click **Create Resource** and then click **Create Virtual Machine**.
2. On the **Basics** tab, perform the following steps:
 - A. Select the region and availability zone where your VM cluster is located.
 - B. Select Oracle Linux 8.x as the installation image.
 - C. Select the required size of the compute instance.
 - D. Enter the username (we recommend `opc`).
 - E. Provide your SSH public key.
3. On the **Disks** tab, select the OS disk and add data disks as needed.
4. On the **Networking** tab, select the correct VNet and the subnet reserved for the SAP NetWeaver application servers. Keep the other default values.
5. On the **Management** tab, keep the default values.
6. On the **Monitoring** tab, keep the default values.
7. On the **Advanced** tab, keep the default values.
8. On the **Tags** tab, assign tags and values as needed.
9. Review all data entered and click **Create** to start deployment of the new virtual machine.

After all deployments succeed, you can use SSH to log in to all three compute nodes (for an Exadata Quarter Rack deployment). You can log in with the `opc` user and gain root access by using `sudo su -`. Also, with the same key as for the `opc` user, you can reach the Oracle account on the VM cluster nodes.

- `ssh opc@node0` takes you to the `opc` user on VM cluster `node0`.
- `ssh oracle@node0` (with the same key as for the `opc` user) takes you to the Oracle user account.

Azure compute instances must be connected to the subnet created for the SAP NetWeaver application servers (for example, `sn-app`). Be sure to follow these design rules:

- As mentioned earlier, always configure a VNet with three subnets: One for the bastion or jump host (for example, `sn-bastion`), one for SAP NetWeaver application servers (for example, `sn-app`), and one for the Exadata client network of the VM cluster (for example, `sn-client`).
- Use separate VNets for the different SAP landscapes.
- Use separate VNets for test, QA, and production.

A local firewall for each compute instance that comes from the OS, and security lists that are part of the Azure networking service, allow and deny specific network traffic. For an SAP deployment, the local firewall must be disabled (which you will do later in this document). If you plan to implement VNet-specific security rules, you can get an overview about the required ports for an SAP system from [SAP Help Ports](#).

Configure IP Resolution and Reverse Lookup

The IP addresses and hostnames of all the hosts in your SAP landscape must be properly resolvable from all hosts involved. A *forward DNS lookup* resolves an FQDN hostname or a short hostname to the corresponding IP address, and a *reverse DNS lookup* resolves a given IP address to an FQDN hostname.

Azure compute instances and OCI VM cluster nodes use different DNS resolvers to resolve DNS lookups. As a result, they must be kept in sync so that DNS lookups return equivalent data from both DNS resolvers.

With Oracle Database@Azure, you must use and maintain two private DNS zones in Azure, one for forward DNS lookups and one for reverse DNS lookups, for example `sap.acme.com` and `19.172.in-addr.arpa`. These zones must contain all the required A records and PTR records of all the hosts involved.

In OCI, more than just two zones are necessary to fulfill this requirement for forward and reverse DNS lookups. The following two sections (one for the OCI default domain naming scheme and one for the custom domain naming scheme) provide an overview of the items that you need to create in OCI *after* preparing the private DNS zones in Azure portal as the first step. In summary:

Review the following sections:

- Overview When Using the OCI Default Domain Naming Scheme
- Overview When Using the Custom DNS Domain Naming Scheme

Perform the steps in the following sections:

1. Deploy Private DNS Zones in Azure (Both Domain Naming Schemes)
2. Deploy Private DNS Zones in OCI
 - Steps When Using the OCI Default Domain Naming Scheme
 - Steps When Using the Custom DNS Naming Scheme


Overview When Using the OCI Default Domain Naming Scheme

If you decided to use OCI's default domain naming scheme, you need private views and private DNS zones with their respective A or PTR records. Some of the items are created automatically and some must be created manually as described in detail later in this section.

Overview of the items to create in OCI:

- In the default private view (usually named `VCN-multicloudnetworklinkXXXXXXXXXX`), create the following reverse lookups zones:
 - A reverse lookup zone for each SCAN-VIP in `.in-addr.arpa` format (for example, `245.20.16.172.in-addr.arpa`) and a corresponding PTR record for the SCAN hostname
 - A reverse lookup zone for each Azure compute instance in `.in-addr.arpa` format (for example, `4.10.16.172.in-addr.arpa`) and a corresponding PTR record for the Azure compute instance
- Because the default view is protected, create an additional private view (usually named `VCN-multicloudnetworklinkXXXXXXXXX_AZURE`) with one private DNS zone for each Azure compute node with the FQDN hostname as the name of the zone (for example, `sapapp1.ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com`) and a corresponding A record.
- Associate this additional private view with the DNS resolver of the VM cluster.

The following figures show examples of the items that you will create.

Zone name	Zone type	Protected 
245.20.16.172.in-addr.arpa	Primary	No
200.20.16.172.in-addr.arpa	Primary	No
172.20.16.172.in-addr.arpa	Primary	No
4.10.16.172.in-addr.arpa	Primary	No
4.0.16.172.in-addr.arpa	Primary	No

SCAN-VIPs (next to 200.20.16.172.in-addr.arpa)

Azure compute instances (next to 4.10.16.172.in-addr.arpa)

Figure 3: Manually Added Private DNS Zones for Reverse DNS Lookup

Zone name	Zone type	Protected 
sapapp1.ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com	Primary	No
sapjump1.ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com	Primary	No

Figure 4: Manually Added Private DNS Zones for DNS Lookup for the Azure Compute Instances

Associated private views

You can add private views to your resolver to manage how DNS queries are answered.

A resolver provides responses by checking zones in your custom private views, then in its default view, then by checking rules, and finally by using internet DNS.

Manage private views Remove

<input type="checkbox"/>	Order	Private view
<input type="checkbox"/>	1	VCN-multicloudnetworklink20240418100612_AZURE

0 selected

Figure 5: Additional Private View Associated with the DNS Resolver

Overview When Using the Custom DNS Domain Naming Scheme

If you decided to use your own custom DNS domain naming scheme, a private view (named SAP_ACME_COM in our example) with numerous private DNS zones is already associated with your DNS resolver, in addition to the default private view, which is always there. You must still add several items manually as described in detail later in this section.

Overview of the items to create in OCI:

- Add the missing A records for the Azure compute instances to the forward DNS lookup zone (sap.acme.com in our example).
- Create additional private DNS zones in .in-addr.arpa format with their respective PTR records for reverse lookup for each Azure compute instance and each SCAN-VIP.

The following figure shows an example of the items that you will create.

Zone name	Zone type	Protected ⓘ
147.20.19.172.in-addr.arpa	Primary	No
47.20.19.172.in-addr.arpa	Reverse-Lookup Zones for SCAN-IPs	No
19.20.19.172.in-addr.arpa	Primary	No
4.10.19.172.in-addr.arpa	Reverse-Lookup Zones for Azure compute instances	No
4.0.19.172.in-addr.arpa	Primary	No
159.255.168.192.in-addr.arpa	Primary	No
26.254.168.192.in-addr.arpa	Primary	No
92.20.19.172.in-addr.arpa	Primary	No
87.20.19.172.in-addr.arpa	Reverse-Lookup Zones for Node-IPs and Node-VIPs (auto-created)	No
65.20.19.172.in-addr.arpa	Primary	No
17.20.19.172.in-addr.arpa	Primary	No
sap.acme.com	Forward-Lookup Zone auto-generated + Azure compute instances added	No

Figure 6: All Private DNS Zones in View SAP_ACME_COM

Note: Although it is possible to add all the hosts and IP addresses to all `/etc/hosts` files on all hosts involved instead of using DNS, we strongly recommend *not* using `/etc/hosts` for production systems. For example, `/etc/hosts` cannot resolve the hostname for the SCAN listener to all three SCAN-VIPs and returns only the first one. This could cause problems if that SCAN-VIP is relocated and unavailable for a moment. With DNS, all three SCAN-VIPs are returned and Oracle Database clients can always try the remaining ones if one SCAN-VIP is unavailable for a moment.

Deploy Private DNS Zones in Azure (Both Domain Naming Schemes)

During the deployment of a VM cluster, a private DNS zone for forward DNS lookups is automatically created in Azure. The name of the private DNS zone is either the name of the custom domain that you provided when you selected **Custom DNS**, or the domain name that OCI has automatically created. You can find it in your list of resources within the Azure portal. On the **Overview** page, you can see that it already contains all the A records that belong to the VM cluster, such as Node-IPs, Node-VIPs, and SCAN-VIPs. If you created APP-VIPs for SAP HA and HA NFS in OCI and issued a ticket with support to enable those VIPs on your VNet in Azure, they should also be listed.

This zone is only for forward DNS lookups and does not contain any A records for the Azure compute instances or PTR records for reverse DNS lookups. You can add A records for Azure compute instances to the zone. For reverse DNS lookups, you must create another private DNS zone and add all PTR records for all hosts involved. This zone must be linked to the VNet by creating a corresponding virtual network link.

To add the missing A records to the private DNS zone in Azure, follow these steps:

1. In the Azure portal, select the private DNS zone created during VM cluster deployment (for example `sap.acme.com`). The **Overview** page shows the list of entries already created.
2. Click **+Record Set** and create an additional A record in which you specify the name of the host and its IP address. Repeat this step for all your Azure compute instances.

The following figure shows an example of a private DNS zone named `sap.acme.com` in Azure for a VM cluster deployment using custom DNS. The APP-VIPs (`ascsmfc`, `ersmfc`, and `nfsmfc`) were first added in OCI and then transferred to Azure by Oracle Cloud Support. The Azure compute instances (`sapapp3` and `sapjump3`) were added manually in the Azure portal.

Name	Type	TTL	Value	Auto registered
@	SOA	3600	Email: azureprivatedns-host.microsoft.com Host: azureprivatedns.net Refresh: 3600 Retry: 300 Expire: 2419200 Minimum TTL: 10 Serial number: 1	False
ascsmfc	A	3600	172.19.20.208	False
bb-xuwwj-scan	A	3600	172.19.20.47 172.19.20.147 172.19.20.19	False
bb-xuwwj1	A	3600	172.19.20.87	False
bb-xuwwj1-vip	A	3600	172.19.20.17	False
bb-xuwwj2	A	3600	172.19.20.92	False
bb-xuwwj2-vip	A	3600	172.19.20.65	False
ersmfc	A	3600	172.19.20.69	False
nfsmfc	A	3600	172.19.20.236	False
sapapp3	A	3600	172.19.10.4	False
sapjump3	A	3600	172.19.0.4	False

Figure 7: Forward DNS Lookup Zone in Azure

To create the private DNS zone for reverse lookups in Azure, follow these steps:

1. In the Azure portal, click **Create a Resource**, select **Private DNS Zone** in Azure Marketplace, and then click **Create**.
2. On the **Basics** tab, select a resource group if you want to add the zone to an existing resource group. Then, enter a name for your reverse DNS lookup zone in the `.in-addr.arpa` format, for example, `19.172.in-addr.arpa`, which defines the scope for all IP addresses starting with `172.19`.
3. On the **Tags** tab, assign tags and values as needed.
4. Review your entries and click **Create**.
5. Navigate to the new zone. The **Overview** page shows only one SOA record.
6. Click **+Record Set** and add a PTR record for each IP address-hostname combination. In the **IP address** field, enter only the remaining part of the IP address in reverse order because, as shown in our example, `19.172` is already defined for the zone.
7. After adding all required PTR records, link the private DNS zone to the VNet. Under **Settings, Virtual network links**, click **Add** to link the zone to the VNet.

The following figure shows an example of a private DNS zone named 19.172.in-addr.arpa in Azure created for reverse DNS lookups. All entries except the SOA record have been added manually.

Name	Type	TTL	Value	Auto registered
4.0	PTR	3600	sapjump3.sap.acme.com	False
4.10	PTR	3600	sapapp3.sap.acme.com	False
147.20	PTR	3600	bb-xuwwj-scan.sap.acme.com	False
17.20	PTR	3600	bb-xuwwj1-vip.sap.acme.com	False
19.20	PTR	3600	bb-xuwwj-scan.sap.acme.com	False
208.20	PTR	3600	ascsmfc.sap.acme.com	False
236.20	PTR	3600	nfsmfc.sap.acme.com	False
47.20	PTR	3600	bb-xuwwj-scan.sap.acme.com	False
65.20	PTR	3600	bb-xuwwj2-vip.sap.acme.com	False
69.20	PTR	3600	ersmfc.sap.acme.com	False
87.20	PTR	3600	bb-xuwwj1.sap.acme.com	False
92.20	PTR	3600	bb-xuwwj2.sap.acme.com	False
@	SOA	3600	Email: azureprivatedns-host.microsoft.com Host: azureprivatedns.net Refresh: 3600 Retry: 300 Expire: 2419200 Minimum TTL: 10 Serial number: 1	False

Figure 8: Reverse DNS Lookup Zone in Azure

Deploy Private DNS Zones in OCI

On the OCI side, VM cluster deployment automatically adds multiple private DNS zones to the private view created earlier in our example (SAP_ACME_COM) or to the system-generated private view (VCN-xxxxxx) attached to the VM cluster. These zones include one forward DNS lookup zone that corresponds to the forward DNS lookup zone in Azure and has the same name (sap.acme.com in our example), and multiple reverse lookup zones. The forward DNS lookup zone contains all the A records of the VM cluster, and the reverse DNS lookup zones all contain a single PTR record for the Node-IPs and Node-VIPs.

The following two figures show an example of which private DNS zones the VM cluster deployment process will create for you and which DNS entries are being created when you attach APP-VIPs. Note that the forward DNS lookup and reverse DNS lookup zones do not contain the A records and PTR records of the Azure compute instances. There is also no zone for reverse DNS lookups of the SCAN-VIPs and the APP-VIPs (ascsmfc, ersmfc and nfsmfc).

Zone name	Zone type	Protected ⓘ
159.255.168.192.in-addr.arpa	Primary	No
26.254.168.192.in-addr.arpa	Primary	No
92.20.19.172.in-addr.arpa	Primary	No
87.20.19.172.in-addr.arpa	Primary	No
65.20.19.172.in-addr.arpa	Primary	No
17.20.19.172.in-addr.arpa	Primary	No
sap.acme.com	Primary	No

Figure 9: Private View SAP_ACME_COM After Initial Deployment of the VM Cluster

Manage records				
Domain ▲	Type	TTL ⓘ	RDATA	
ascsmfc.sap.acme.com	A	3600	172.19.20.208	
bb-xuwwj-scan.sap.acme.com	A	3600	172.19.20.19 172.19.20.47 172.19.20.147	
bb-xuwwj1-backup.sap.acme.com	A	3600	192.168.254.26	
bb-xuwwj1-vip.sap.acme.com	A	3600	172.19.20.17	
bb-xuwwj1.sap.acme.com	A	3600	172.19.20.87	
bb-xuwwj2-backup.sap.acme.com	A	3600	192.168.255.159	
bb-xuwwj2-vip.sap.acme.com	A	3600	172.19.20.65	
bb-xuwwj2.sap.acme.com	A	3600	172.19.20.92	
ersmfc.sap.acme.com	A	3600	172.19.20.69	
nfsmfc.sap.acme.com	A	3600	172.19.20.236	
sap.acme.com	NS ⓘ	86400	vcn-dns.oraclevcn.com.	
sap.acme.com	SOA ⓘ	86400	vcn-dns.oraclevcn.com. hostmaster.oracle.com. 39 3600 3600 3600 10	

Figure 10: Private DNS Zone sap.acme.com After Initial Deployment of the VM Cluster and After Attaching Additional APP-VIPs

Steps When Using the OCI Default Domain Naming Scheme

If you decided to use OCI's default domain naming scheme, you must create a new private view with a private DNS zone and the respective A record for each Azure compute instance for forward DNS lookups, as described earlier in this section.

To create the private view and the private zones for forward DNS lookups, follow these steps:

1. Sign in to Azure portal and navigate to your VM cluster.
2. Under **Overview**, click **Go to OCI** to navigate to the Oracle Cloud Console and display the details page of the VM cluster. Note the name of the private view.

Navigating from Azure to OCI selected the right compartment for the next steps.

3. In the main navigation menu, select **Network, DNS Management** and then select **Private Views** to display a list of private views that already exist.
4. Click **Create private view**, and then enter the name that you noted in step 2, adding a suffix (for example, `_AZURE`) to clarify that the view complements the default view already present and is intended for the new zones you add in the next step. Click **Create**.
5. In the new private view, click **Create zone** and enter the name for the forward DNS lookup zone. Enter the FQDN hostname of your Azure compute instance (for example, `sapapp3.sap.acme.com`) and click **Create**.
6. In the new zone, click **Manage records** and add an A record for the Azure compute instance. Then, click **Publish records** and **Confirm publish records**.
7. Repeat steps 5 and 6 for all relevant Azure compute instances.
8. Navigate back to the VM cluster details page, click **Virtual cloud network** and then click **DNS Resolver**.
9. Select the **Associated private views** pane on the left side and then click **Manage private views**.
10. In the **Manage private views** dialog box, select the private view created in step 4 and click **Save changes**.

To create the private DNS zones for reverse DNS lookups, follow these steps:

1. Sign in to Azure portal and navigate to your VM cluster.
2. Under **Overview**, click **Go to OCI** to navigate to the Oracle Cloud Console and display the details page of the VM cluster.
3. Click **Private view** to navigate to the default view of the DNS resolver.
4. Click **Create zone** and create a reverse DNS lookup zone in the private view by entering the full IP address in reverse order in `.in-addr.arpa` format as the name of the zone (for example, `4.10.19.172.in-addr.arpa` for `sapapp3.sap.acme.com`). Click **Create**.
5. In the new private zone, click **Manage records** and add a PTR record. Leave the name empty, select the PTR record type, enter the FQDN hostname of the host that you want to add, and save your changes. Click **Publish changes** and **Confirm publish changes**.
6. Repeat step 4 and 5 for all relevant Azure compute instances, APP-VIPS and for the SCAN-VIPS.

To apply the required changes for DNS on the Azure compute nodes, follow these steps:

1. Log in to your Azure compute instance that is planned to be part of the SAP system (at least the host deployed for the SAP primary application server), and switch to root.
2. Adjust the NetworkManager configuration to prevent NetworkManager from overwriting file `/etc/resolv.conf`.
 - A. Edit the `/etc/NetworkManager/NetworkManager.conf` file, ensure that it contains the following line directly after the `[main]` section, and then save the changes.


```
dns=none
```
 - B. Edit the `/etc/resolv.conf` file with the correct search domain. For example, replace `search e13gbcr0fxhuvasra5dstkqukf.frax.internal.cloudapp.net` with `search ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com` and then save the changes.

The following figures illustrate the preceding configuration steps.

Zone name	Zone type	Protected ⓘ
245.20.16.172.in-addr.arpa	Primary	No
200.20.16.172.in-addr.arpa	Primary	No
172.20.16.172.in-addr.arpa	Primary	No
4.10.16.172.in-addr.arpa	Primary	No
4.0.16.172.in-addr.arpa	Primary	No
ocibackupvnxlo.ocivnxlobsapfr.oraclevcn.com	Primary	Yes
ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com	Primary	Yes
255.168.192.in-addr.arpa	Primary	Yes
254.168.192.in-addr.arpa	Primary	Yes
253.168.192.in-addr.arpa	Primary	Yes
252.168.192.in-addr.arpa	Primary	Yes
20.16.172.in-addr.arpa	Primary	Yes

Figure 11: Default Private View with Additional Reverse DNS Lookup Zones Added for Azure Compute Instances and SCAN-VIPs

Networking > DNS management > Private views > VCN-multicloudnetworklink20240418100612_AZURE

VCN-multicloudnetworklink20240418100612_AZURE

ACTIVE

Private view information

OCID: ...zeutua [Show](#) [Copy](#)

Protected: No ⓘ

Private zones in 1e36c776-fc0a-4d93-b8e7-e2c1215a

Private zones contain DNS data only accessible from within a VCN, such as private IP addresses.

Create zone

Zone name
sapapp1.ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com
sapjump1.ocivnxlobsapfr.ocivnxlobsapfr.oraclevcn.com

Figure 12: Additional Private View Required to Define the Forward DNS Lookup Zones for the Azure Compute Instances

Networking > Virtual cloud networks > VCN-multicloudnetworklink20240418100612 > Private resolver details

VCN-multicloudnetworklink20240418100612

Edit Move resource Add tags

Private resolver information Tags

OCID: ...c3rkma [Show](#) [Copy](#)

Dedicated virtual cloud network: [VCN-multicloudnetworklink20240418100612](#)

Protected: Yes ⓘ

A private DNS resolver handles DNS queries within your VCN based on private views and the private zones they contain. [Learn more](#)

A private DNS zone has similar capabilities to an internet DNS zone, but provides responses only for clients that can reach it through a VCN. VCN creation includes a dedicated DNS resolver and a default private view with system-managed zones.

You will have to create views and zones that the resolver can use to direct internal traffic. [Learn more](#)

Resources

- Associated private views (1)
- Rules (0)
- Endpoints (0)
- Work requests (0)

Associated private views

You can add private views to your resolver to manage how DNS queries are answered.

A resolver provides responses by checking zones in your custom private views, then in its default view, then by checking rules, and finally by using internet DNS.

Manage private views Remove

<input type="checkbox"/>	Order	Private view
<input type="checkbox"/>	1	VCN-multicloudnetworklink20240418100612_AZURE

Figure 13: Additional Private View Attached to the DNS Resolver

Steps When Using the Custom DNS Naming Scheme

If you decided to use your own custom DNS domain naming scheme, you already created a private view before you deployed your VM cluster (SAP_ACME_COM in our example). You do not need to create an additional view, and all additional private DNS zones that you must create are solely for reverse DNS lookups. Missing A records in the private DNS zone intended for forward DNS lookups (sap.acme.com in our example) are directly added to that zone.

To create the private DNS zones for reverse DNS lookups for the Azure compute instances and the SCAN-VIPs, follow these steps:

1. Sign in to Azure portal and navigate to your VM cluster.
2. Under **Overview**, click **Go to OCI** to navigate to the Oracle Cloud Console and display the details page of the VM cluster.
3. Click **Virtual cloud network**, and then click **DNS Resolver** to navigate to the Private resolver details page.
You should see the private view that you created before you deployed the VM cluster (SAP_ACME_COM in our example), associated with the DNS resolver.
4. Click the name of the private view.
5. Click **Create zone** and create a reverse DNS lookup zone in the private view by entering the full IP address of the host in reverse order in `.in-addr.arpa` format as the name of the zone (for example, `4.10.19.172.in-addr.arpa` for `sapapp3.sap.acme.com`). Click **Create**.
6. In the new private zone, click **Manage records** and add a PTR record. Leave the name empty, select the PTR record type, enter the FQDN hostname of the host that you want to add, and then save your changes. Click **Publish changes** and **Confirm publish changes**.
7. Repeat step 5 and 6 for all relevant Azure compute instances, APP-VIPs, and SCAN-VIPs.

To add the missing A records to the private DNS zone for forward DNS lookups (sap.acme.com in our example), follow these steps:

1. Select the private DNS zone to open the list of DNS records, and then click **Manage records**.
2. Click **Add**, enter the hostname of the host that you want to add, select an A record type record, and enter the IP address of the host. Click **Add record**.
3. Repeat step 2 and create an A record for each host (Azure compute instances and SCAN-VIPs).
4. Click **Publish changes** and **Confirm publish changes**.

To apply the required changes for DNS on the Azure compute nodes, follow these steps:

1. Log in to your Azure compute instance that is planned to be part of the SAP system (at least the host deployed for the SAP primary application server) and switch to root.
2. Adjust the NetworkManager configuration to prevent NetworkManager from overwriting file /etc/resolv.conf. Edit the /etc/NetworkManager/NetworkManager.conf file and ensure that it contains the following line directly after the [main] section, and then save the changes.

```
dns=none
```

3. Edit the /etc/resolv.conf file with the correct search domain. For example, replace search e13gbcr0fxhuvasra5dstkqukf.frax.internal.cloudapp.net with search sap.acme.com and save the changes.

The following figures illustrate the preceding configuration steps.

Private zones contain DNS data only accessible from within a VCN, such as private IP addresses.

Zone name	Zone type	Protected
147.20.19.172.in-addr.arpa	Primary	No
47.20.19.172.in-addr.arpa	Primary	No
19.20.19.172.in-addr.arpa	Primary	No
4.10.19.172.in-addr.arpa	Primary	No
4.0.19.172.in-addr.arpa	Primary	No
159.255.168.192.in-addr.arpa	Primary	No
26.254.168.192.in-addr.arpa	Primary	No
92.20.19.172.in-addr.arpa	Primary	No
87.20.19.172.in-addr.arpa	Primary	No
65.20.19.172.in-addr.arpa	Primary	No
17.20.19.172.in-addr.arpa	Primary	No
sap.acme.com	Primary	No

Figure 14: All the Private DNS Zones in the SAP_ACME_COM View

Domain	Type	TTL (i)	RDATA
ascsmf.sap.acme.com	A	3600	172.19.20.208
bb-xuwwj-scan.sap.acme.com	A	3600	172.19.20.19 172.19.20.47 172.19.20.147
bb-xuwwj1-backup.sap.acme.com	A	3600	192.168.254.26
bb-xuwwj1-vip.sap.acme.com	A	3600	172.19.20.17
bb-xuwwj1.sap.acme.com	A	3600	172.19.20.87
bb-xuwwj2-backup.sap.acme.com	A	3600	192.168.255.159
bb-xuwwj2-vip.sap.acme.com	A	3600	172.19.20.65
bb-xuwwj2.sap.acme.com	A	3600	172.19.20.92
ersmf.sap.acme.com	A	3600	172.19.20.69
nfsmf.sap.acme.com	A	3600	172.19.20.236
sap.acme.com	NS	86400	vcn-dns.oraclevcn.com.
sap.acme.com	SOA	86400	vcn-dns.oraclevcn.com. hostmaster.oracle.com. 48 3600 3600 3600 10
sapapp3.sap.acme.com	A	3600	172.19.10.4
sapjump3.sap.acme.com	A	3600	172.19.0.4

Figure 15: All the DNS A Records in the sap.acme.com Domain

Networking » Virtual cloud networks » VCN-multicloudnetworklink20240516182437 » Private resolver details

VCN-multicloudnetworklink20240516182437

R
ACTIVE

Edit Move resource Add tags

Private resolver information Tags

OCID: ...c4ykrq [Show](#) [Copy](#)

Dedicated virtual cloud network: [VCN-multicloudnetworklink20240516182437](#)

Protected: Yes (i)

A private DNS resolver handles DNS queries within your VCN based on private views and the private zones they contain. [Learn more](#)
 A private DNS zone has similar capabilities to an internet DNS zone, but provides responses only for clients that can reach it through a VCN. VCN creation includes a dedicated DNS resolver and a default private view with system-managed zones.
 You will have to create views and zones that the resolver can use to direct internal traffic. [Learn more](#)

Resources

- Associated private views (1)**
- Rules (0)
- Endpoints (0)
- Work requests (0)

Associated private views

You can add private views to your resolver to manage how DNS queries are answered.
 A resolver provides responses by checking zones in your custom private views, then in its default view, then by checking rules, and finally by us

Manage private views Remove

<input type="checkbox"/>	Order	Private view
<input type="checkbox"/>	1	SAP_ACME_COM

Figure 16: Private DNS Resolver with SAP_ACME_COM View Associated with It

Now your DNS setup is complete from both the Azure and OCI perspectives. Other host-specific changes are discussed in later sections.

Set Up /etc/hosts (Not Recommended)

If you do not want to set up DNS as described in the preceding sections, you can add the necessary entries to the /etc/hosts file of each Azure compute instance and each VM cluster node involved.

Note: Using /etc/hosts instead of configuring DNS is not recommended, especially for production systems. Only DNS-based hostname resolution delivers all three SCAN-VIPs for the SCAN hostname, which is important during failovers, where a single SCAN-VIP can be unavailable for a short time period.

Edit the /etc/hosts file of each relevant host with the IP addresses and hostnames for the following items:

- Node-IPs, Node-VIPs, SCAN-VIPs, and their corresponding hostnames or virtual hostnames
- APP-VIPs for HA-NFS and SAP HA
- Hostname for the PAS host and potential additional Azure compute instances

Ensure that the FQDN hostname is at the first position with the short hostname following.

The following example shows the same entries as in the previous example with private DNS zones but in /etc/hosts.

```
172.19.0.4 sapjump3.sap.acme.com sapjump3 #Bastion
172.19.10.4 sapapp3.sap.acme.com sapapp3 #Appserver1
172.19.20.87 bb-xuwwj1.sap.acme.com bb-xuwwj1 #ClusterNode IP
172.19.20.17 bb-xuwwj1-vip.sap.acme.com bb-xuwwj1-vip #ClusterNode VIP
172.19.20.92 bb-xuwwj2.sap.acme.com bb-xuwwj2 #ClusterNode IP
172.19.20.65 bb-xuwwj2-vip.sap.acme.com bb-xuwwj2-vip #ClusterNode VIP
172.19.20.47 bb-xuwwj-scan.sap.acme.com bb-xuwwj-scan #ScanListener
172.19.20.147 bb-xuwwj-scan.sap.acme.com bb-xuwwj-scan #ScanListener
172.19.20.19 bb-xuwwj-scan.sap.acme.com bb-xuwwj-scan #ScanListener
172.19.20.208 ascsmfc.sap.acme.com ascsmfc #AppVIP ASCS
172.19.20.69 ersmfc.sap.acme.com ersmfc #AppVIP ERS
172.19.20.236 nfmfc.sap.acme.com nfmfc #AppVIP HANFS
```

Prepare Your Environment

You can set up all the resources by using the Azure portal and the Oracle Cloud Console.

- Apply the latest Grid Infrastructure patch
- Set up the bastion host
- Set up the ULN proxy
- Set up the VNC server
- Set up the SAP Download Manager
- Download your SAP software

Apply the Latest Grid Infrastructure Patch

Use the Oracle Cloud Console to apply the latest Grid Infrastructure patch.

Set Up the Bastion Host

We recommend that you use Oracle Linux 8 on the bastion host. A bastion host typically has one network interface connected to the public internet and a second network interface connected to the VNet, typically to the subnet reserved for the bastion host (`sn-bastion`). This configuration enables you to implement a good set of security measures. The bastion host *can* have the following roles:

- Provide a VNC server for graphical access and an SSH server for outside access. VNC or X-Windows is not required for newer versions of SAP SWPM or SAP SUM because they can be accessed from a web browser.
- Provide a graphical workspace for any related operations (for example, download, install, and access).
- Work as a ULN proxy.
- Act as HTTP/HTTPS proxy, for example, for automatic lifecycling (automatic updates) of SAP enhanced monitoring. VM cluster nodes should never have direct access to the internet.

To store your SAP installation media, SAP patches, and Oracle SAP Bundle Patches (SBPs), create an NFS file system in Azure that can be shared between the bastion host and other Azure compute instances planned for your SAP landscape. You will have to copy all the files to shared ACFS on the VM cluster later.

Note: It is not possible to share Azure-based NFS file systems between Azure compute instances and VM cluster nodes; they can be mounted only on the Azure compute instances. The same is true for OCI-based NFS file systems; they can be mounted only on the VM cluster nodes. The only exceptions are ACFS-based NFS file systems exported by the VM cluster and made highly available through Oracle Clusterware using an APP-VIP (HA-NFS) made available in the Azure VNet. With this method, a highly available NFS can be shared across all involved Azure compute instances as well as all VM cluster nodes.

Set Up the ULN Proxy

To ensure that you have the latest OS updates for Oracle Linux 8 available from Oracle, register the system with the Oracle Unbreakable Linux Network (ULN) and set up a ULN proxy. A proxy enables you to update compute instances with the latest packages even if the instance is not connected to the internet. A requirement for maintaining the proxy is to ensure that sufficient disk space is available to hold all the updates.

Register your Oracle Linux 8 system to ULN, and follow the [ULN User's Guide](#) to configure a ULN proxy to mirror the needed local channels. Provide a block volume after you approximate the size of your needed channels.

Set Up the VNC Server

GUI access at the OS level is needed to run any graphical tools. You can access the native GUI by enabling a VNC server on the bastion host. Ensure that security lists are maintained to allow access to only approved sources.

Configure a VNC server on the bastion host as described in the [Install the VNC Remote Access Server on Oracle Linux tutorial](#). Implement local firewall rules or entries in the security lists to allow access to the VCN from your outside network.

Set Up the SAP Download Manager

SAP Download Manager helps you download software from the SAP Software Download Center that you have put in the download basket. Install the SAP Download Manager on the bastion host and set the needed S-User and password credentials to download SAP software from the [SAP Software Download Center](#) (SWDC).

Download Your SAP Software

From the SWDC, download the required installation software for your specific SAP product. With your S-User permissions, you can download the installation media directly or you can use the SAP Download Manager. We recommend storing the software on a shared file system.

We also recommend using the SAP Maintenance Planner to compose the required installation and upgrade media and push them to the download basket. You can generate a `stack.xml` file to use with SWPM to provide a consistent set of installation media that matches the contents of your download basket. You can then add more Oracle RDBMS and Oracle Client media from the SAP marketplace before downloading all the media.

Prepare VM Cluster Nodes

This section provides the necessary steps for preparing the VM cluster nodes. The key steps are shown in the following diagram.

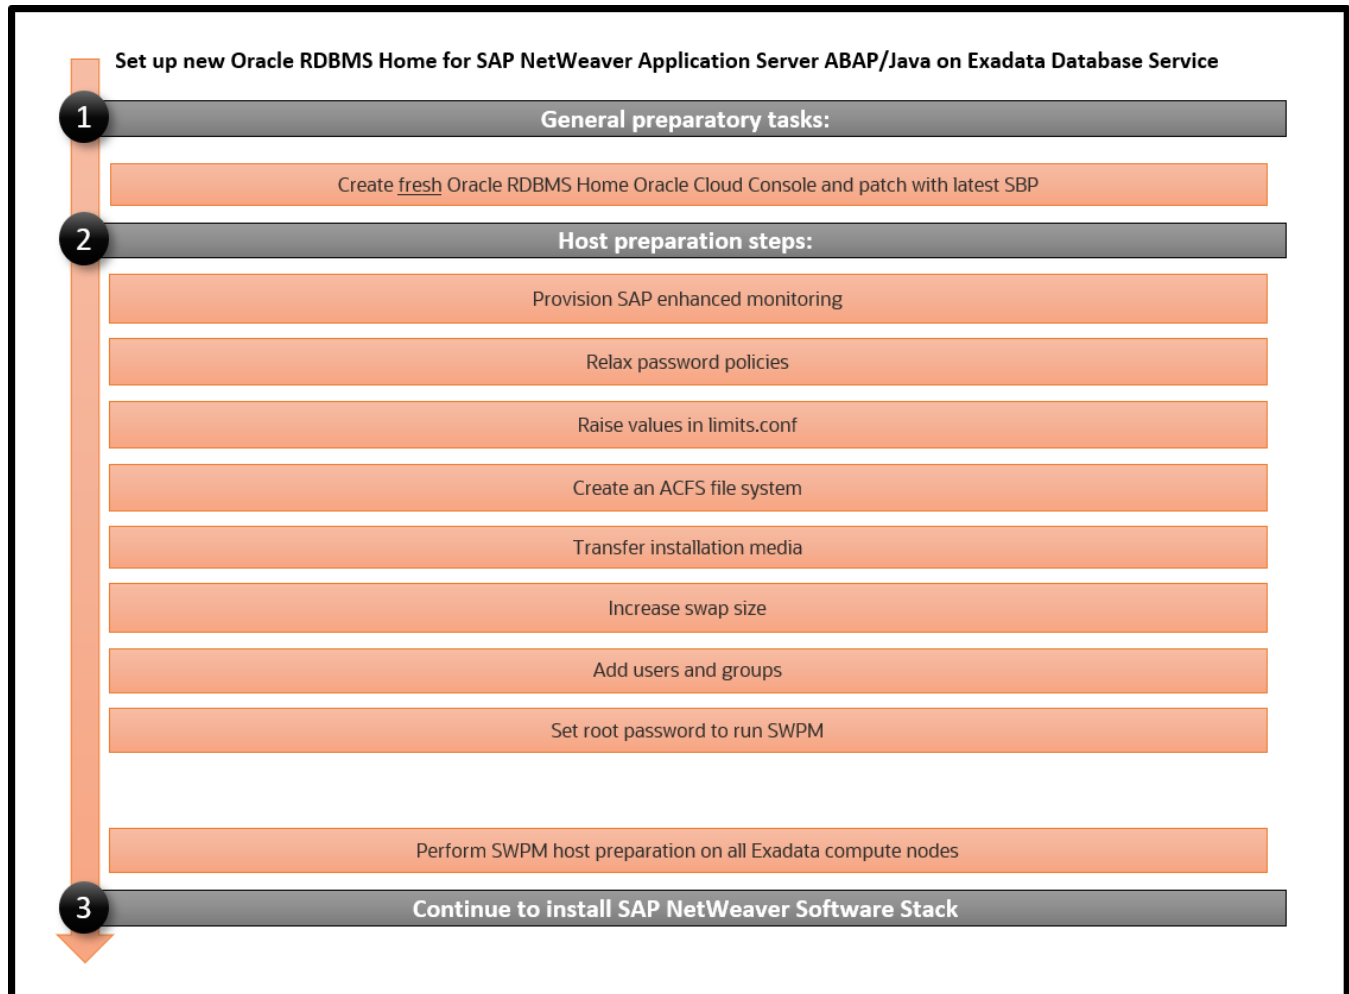


Figure 17: Steps for Setting Up an Oracle Database Home for SAP NetWeaver Application Server ABAP/Java with Oracle Exadata Database Service

Create a Fresh Oracle Database Home

1. In the Azure portal, navigate to the VM cluster in which you want to create the Oracle Database home.
2. Click **Go to OCI** to go to the details page for the VM cluster in Oracle Cloud Console.
3. Under **Resources**, select **Database Homes** and then click **Create Database Home**.
4. Enter a display name for the Oracle Database home.
5. Click **Change database image**, click **Display all available versions**, select **19.21.0.0**, and then click **Select**.

Because of a patching problem starting with 19.22.0.0 when applying the latest SAP Bundle Patch (SBP), you must start with 19.21.0.0 and then apply the latest SBP. You patch the new Oracle Database home in a later step after transferring the latest SAP Bundle Patch to a shared file system location.

6. Click **Create**.

Provision SAP Monitoring

For every cloud solution, SAP requires the collection of configuration and performance data for the cloud platform being used.

With Exadata Database Service, the SAP Host Agent needs to run on all the VM cluster nodes. Installing the SAP Host Agent either by using SAP SWPM or manually is described in the [SAP Host Agent Installation](#) SAP documentation topic.

The required version and patch level of the SAP Host Agent are described in [SAP Note 2614080](#).

For SAP monitoring, the SAP Host Agent consumes Exadata Database Service–specific configuration and performance metrics collected by a Linux service called `oraescscol`. The `oraescscol` service must be installed and started on each VM cluster node. It is shipped as a Linux RPM called `oraescscol.rpm`.

Part of the `oraescscol.rpm` package is a Python script called `oraescswatcher`. This script ensures that updates of the package are applied automatically and that the `oraescscol` service is started if it is not running. Schedule `oraescswatcher` as a cron job to run periodically.

For automatic lifecycling of `oraescscol.rpm`, the database compute nodes need either direct access to the internet or an HTTPS proxy server with access to the internet (for example, on the bastion host) to be able to find and download new versions of the package from Oracle. For example, if you are using a bastion host, the proxy server software “squid” could be configured to forward HTTPS requests from the database compute nodes to the internet, as follows:

```
[root@bastion-host] # yum -y install squid
[root@bastion-host]# systemctl enable squid
Created symlink from /etc/systemd/system/multi-user.target.wants/squid.service to
/usr/lib/systemd/system/squid.service.
[root@bastion-host] # systemctl start squid
```

Install and Configure the oraescscol Package

The required version of `oraescscol.rpm` is already available on the VM cluster nodes after initial deployment. To install and configure it, perform the following steps:

1. As the root user, copy `oraescscol.rpm` from `/u02/opt/dbaas_images/oraescscol.rpm` to `/tmp` and make `/tmp` your current working directory:

```
[root@nodeN] # cp /u02/opt/dbaas_images/oraescscol.rpm /tmp ; cd /tmp
```

2. Install `oraescscol.rpm`:

```
rpm -i oraescscol.rpm
```

3. Enable the service:

```
systemctl enable oraescol.service
```

4. Start the service:

```
systemctl start oraescol.service
```

5. As the root user, add the following cron job into the root user's crontab:

- A. Edit crontab:

```
[root@nodeN] # crontab -e
```

- B. If the compute nodes have public internet access, add the following line:

```
*/15 * * * * sudo /usr/bin/python /opt/oracle.oraescol/oraecswatcher
```

- C. If the compute nodes do not have public internet access and require an HTTPS proxy to allow public internet access, add the following lines:

```
*/15 * * * * sudo https_proxy=<proxy_address> /usr/bin/python
/opt/oracle.oraescol/oraecswatcher
```

For example:

```
*/15 * * * * sudo https_proxy=https://<IP_bastion_host:3128> /usr/bin/python
/opt/oracle.oraescol/oraecswatcher
```

- D. Save crontab:

```
:wq
```

6. Wait two minutes, and then check whether metrics collection works as expected by running the following command as the root user:

```
[root@nodeN] # curl http://127.0.0.1:18181
```

This command should return the XML document for consumption by the SAP Host Agent. For example:

```
.
.
<!-- Provider Health Description #113 -->
<metric category="config" context="vm" last-refresh="1519899668" refresh-interval="60"
type="string" unit="none">
<name>
    Provider Health Description
</name>
<value>
OK
</value>
</metric>
.
.
```

Logs for oraescol are written to `/opt/oracle.oraescol`.

Note: If this test does not return an XML document at all (for example, you get a “connection refused” error) or returns a status other than OK for Provider Health Description, open a ticket with Oracle Support and ask for the Dom0 part of the SAP metrics collector.

7. Follow SAP Note [3475554 - Updating SAP Metrics Collector \(oraecscol\) for Oracle Exadata Cloud Infrastructure and Oracle Exadata Cloud@Customer](#) and ensure that you are running the latest version of oraecscol.

An additional step, discussed later in this document, is to complete the SAP monitoring setup by registering the VM cluster nodes in SAP transaction RZ21.

Relax the Password Policies

An SAP NetWeaver installation does not work with the Exadata Database Service strong password policy, so you must modify the policy.

Run the following command as root on each cluster node:

```
[root@nodeN] # /opt/oracle.cellos/host_access_control pam-auth --deny 10 --lock 60 --pwquality 6 --remember 0
```

With this approach, you still cannot run `su - <someuser>` from the `opc` account because of PAM policy, and you must still switch to root via `sudo su - first`.

Increase Values in the limits.conf File

On each compute node, perform the following steps as the root user:

1. Open `/etc/security/limits.conf` for editing.
2. Add the following lines under the `oracle` entries:

```
root    soft    memlock    unlimited
root    hard    memlock    unlimited
```

3. If you plan to perform offline backups by using SAP BR*Tools, add the following entries for each SAPSID that you plan to install. These entries are required because SAP BR*Tools startup mounts a database instance under one of these OS users to perform the offline backup.

```
ora<sid>    soft    memlock    unlimited
ora<sid>    hard    memlock    unlimited
<sid>adm    soft    memlock    unlimited
<sid>adm    hard    memlock    unlimited
```

4. Save and exit the file.

```
:wq
```

Create a Decently Sized Oracle ACFS

ACFS will be used within the VM cluster for numerous of purposes, for example to share the installation media, backup logs of BR*Tools or `/sapmnt` between the VM cluster nodes.

1. Follow the [documentation](#) to create an Oracle ACFS on one node only. For example:

```
[opc@node0] $ sudo su -
[root@node0] # su - grid
[grid@node0] $ asmcmd
ASMCMD> volcreate -G DATA1 -s 1024G sapshare_v1
ASMCMD> volinfo -G DATA1 sapshare_v1
Diskgroup Name: DATA1

Volume Name: SAPSHARE_V1
Volume Device: /dev/asm/sapshare_v1-128
State: ENABLED
Size (MB): 1048576
Resize Unit (MB): 512
```



```

Redundancy: HIGH
Stripe Columns: 8
Stripe Width (K): 1024
Usage:
Mountpath:

```

```

[grid@node0] $ /sbin/mkfs -t acfs /dev/asm/sapshare_v1-128
mkfs.acfs: version                = 19.0.0.0.0
mkfs.acfs: on-disk version        = 46.0
mkfs.acfs: volume                 = /dev/asm/sapshare_v1-128
mkfs.acfs: volume size           = 1099511627776 ( 1.00 TB )
mkfs.acfs: Format complete.
[grid@node0] $ exit
[root@node0] # /sbin/acfsutil registry -a /dev/asm/sapshare_v1-128 /sapshare

```

The preceding commands create a 1-TB cluster file system out of the +DATA1 disk group and mount it to /sapshare. This change is persistent.

2. On a single compute node, run the following commands as root:

```

[root@node0] # mkdir /sapshare/trans
[root@node0] # mkdir /sapshare/sapmnt
[root@node0] # mkdir /sapshare/sapbins

```

3. On each compute node, run the following commands as root (which avoids the /usr/sap directory being in the boot partition):

```

[root@nodeN] # mkdir /u02/sap ; ln -s /u02/sap /usr/sap
[root@nodeN] # ln -s /sapshare/sapmnt /sapmnt
[root@nodeN] # ln -s /sapshare/trans /usr/sap/trans
[root@nodeN] # chmod 777 /sapshare/sapbins

```

Transfer the Latest SAP Bundle Patch for Exadata Database Service

After the Oracle Database home is created, as described in an earlier section, transfer the latest SAP Bundle Patch for Exadata Database Service to a shared directory, for example under /sapshare/sapbins. The installation of the SAP Bundle Patch is discussed in “Patch Oracle RDBMS Homes Before Installing the Database.”

Note that patching is done *before* SAP database instance installation using SWPM, while catsbp must be run *after* SAP database instance installation. For more information about SAP Bundle Patches, see “SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases.”

Transfer the SAP Installation Media

Transfer your SAP installation media. The media must include SAP NetWeaver, SWPM, DBA Tools, the most up-to-date SAP Host Agent, Oracle Client software, and SAPCAR to extract the SAR archives. Refer to SAP PAM for suitable installation media. Unpack the archives.

If you put the media in a shared location, such as /sapshare/sapbins, you have to transfer only once. If you have composed the media required for installation by using the SAP Maintenance Planner and plan to use the stack.xml option (SAPINST_SAPINST_STACK_XML=<stack.xml>), you should already have all the necessary components for installation.

Increase the Swap Size

Current deployments have 16 GB of swap space by default. If you need more swap space, you can get some space from /u02.

Add Users and Groups

SWPM expects the oper group, which is not present. On each compute node, run the following commands as the root user:

```
[root@nodeN] # groupadd --gid 504 oper
[root@nodeN] # usermod oracle -a -G oper
[root@nodeN] # usermod oracle -a -G asmadmin
```

Change the Root Password on VM cluster nodes

As root, change the password for yourself on each compute node. SWPM asks for authentication.

```
[root@nodeN ] # passwd root
```

Run SAP SWPM on VM Cluster Nodes

During the following steps, you run SWPM several times to perform configurations. Current versions of SWPM run in browser mode. Ensure that the VM cluster nodes (IP addresses and ports) can be reached from the computer where you run your browser. If the hostnames of the VM cluster nodes cannot be resolved from the computer where you run your browser, replace the hostnames with the corresponding IP addresses.

When prompted, confirm the security exception, and enter the credentials for root in the login dialog box.

Prepare All VM Cluster Nodes Using SWPM

Perform this step sequentially, node-by-node.

1. Ensure that SWPM temporary files are placed in a directory with enough space:

```
[root@nodeN] # mkdir -p /usr/sap/tmp ; export TMP=/usr/sap/tmp
```

2. Run SWPM host preparation. For example, if you want to run preparations for Kernel 7.50, navigate to **Generic Options**, then **Oracle**, then **Database Tools**, and select **RAC/ASM/Exadata Database Instance Preparation ABAP – Kernel 7.50**.

Note the following guidelines:

- Do *not* provide a `stack.xml` file for host preparations.
 - At any SWPM instance, never use the FQDN option.
 - All hostnames must be short, with a maximum length of 13 characters.
 - At the local listener configuration, keep the default values.
 - At the Oracle Client selection page, select **Oracle Client 19c**.
 - For the grid installation, choose `/u01/app/19.0.0.0/grid`, and for the ASM instance, choose their respective ASM instance name, such as `+ASM1`, `+ASM2`, and so on.
 - You can obtain the name of the SCAN listener from Oracle Cloud Console (use the short name).
 - Verify and, if needed, adjust the hostnames of the VM cluster node.
3. After completing instance preparation on all the VM cluster nodes, on the first node, as the root user, remove all files and directories under the `/usr/sap/<SAPSID>/SYS/exe/uc/linuxx86_64` directory and keep the empty directory. For example, for SAPSID MFC, run the following command:

```
[root@node0] # rm -rf /usr/sap/MFC/SYS/exe/uc/linuxx86_64/*
```

Install SAP NetWeaver Application Server ABAP/Java

This section describes the steps for installing SAP NetWeaver Application Server ABAP/Java, which includes installing SAP NetWeaver instances such as SAP ABAP central services (ASCS) and installing the database instance by using the latest available version of the SAP Software Provisioning Manager (SWPM). At the time that this document was published, this version was patch level 36. If your SCAN listener is running on a port other than 1521, you need patch level 36 or later. The key steps are shown in the following illustration.

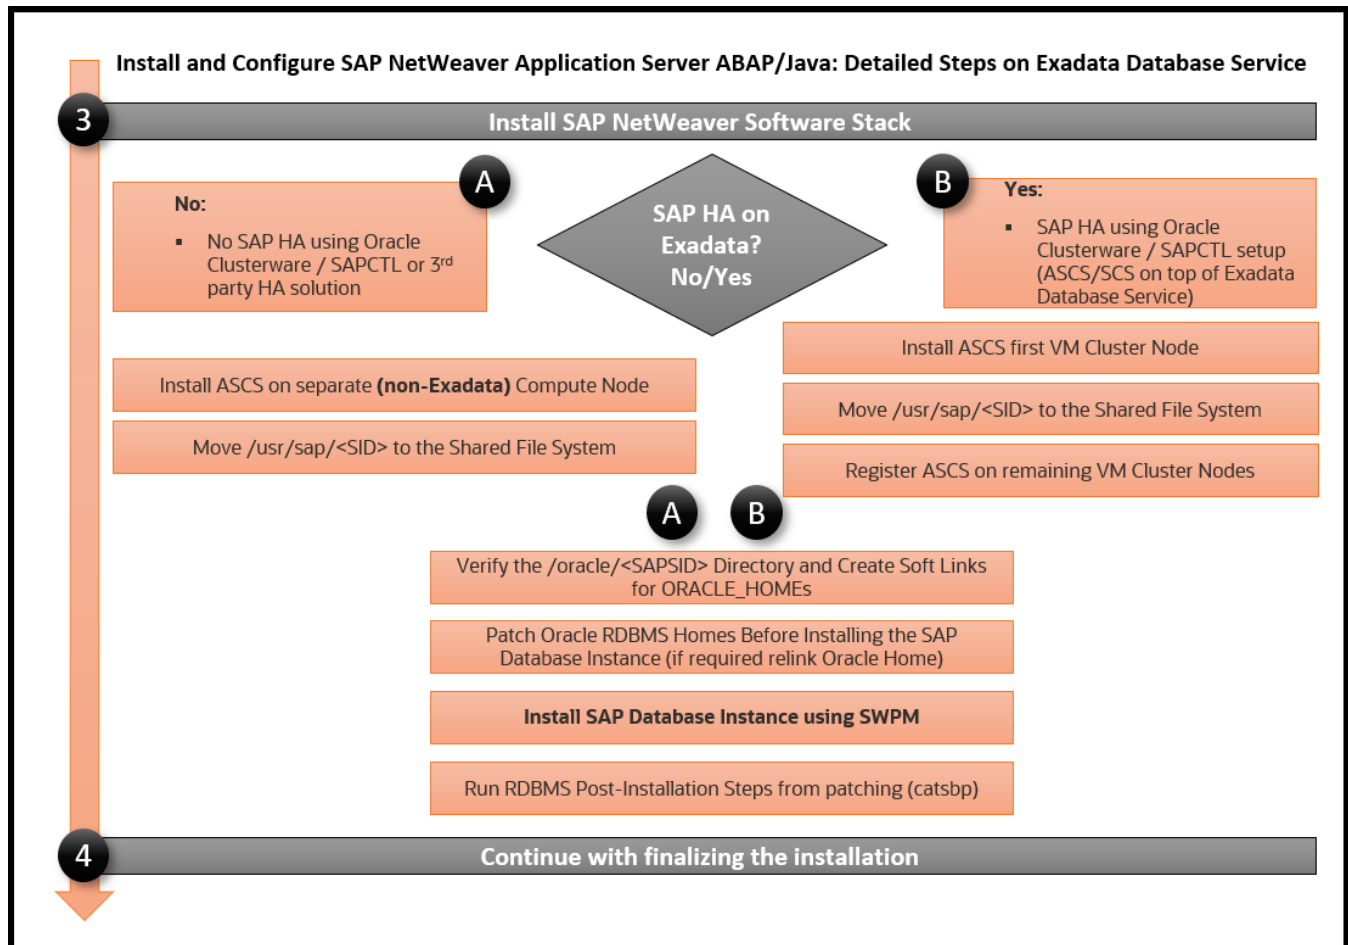


Figure 18: Steps to Install and Configure the SAP NetWeaver Software Stack

Install the ASCS Instance (Optional)

Note: This step is required only if you want to implement SAP HA by using SAPCTL on the VM cluster.

Install the ASCS instance on the first VM cluster node. If you plan to use SAP transaction DB13, select the installation of the **ASCS integrated gateway** during the installation process for all VM cluster nodes. The ASCS instance installation is on a shared resource for later SAPCTL preparation. Previously, in the “Deploy Application-Specific Virtual IP Addresses (APP-VIPs)” section, you deployed the ASCS APP-VIP (172.19.20.208). Now you need to start it.

- Log in to the first node as root, and run the following commands:

```
[root@node0] # . oraenv [+ASM1]
[root@node0] # appvipcfg create -network=1 -ip=172.19.20.208 -vipname=tmp -user=root
[root@node0] # crsctl start res tmp ; crsctl stat res -t
```

2. Check the host location of the resource tmp file. If it's not on the first node, relocate it to the first node.

```
[root@node0] # crsctl stat res -t
[root@node0] # crsctl relocate res tmp -n <current_node_name>
```

Note: Before you can create a database instance, SAP requires you to have an ASCS instance. For later HA awareness of the ASCS instance, follow [SAP Note 1877857](#). HA-aware ASCS installations must be in a shared location, namely /usr/sap/<SAPSID>. Here, you use /sapshare and put the ASCS instance under /sapshare from the first host. Ideally, you would use a separate Oracle ACFS or other shared file system.

3. Ensure that SWPM temporary files are placed in a directory with enough space:

```
[root@node0] # mkdir -p /usr/sap/tmp ; export TMP=/usr/sap/tmp
```

4. Invoke SWPM on first the compute node using the virtual hostname for ASCS by running ./sapinst SAPINST_USE_HOSTNAME=ascsmfc and install the ASCS instance. Choose instance ID 00. You can provide your stack.xml file if you want to use it during installation.

Note the following guidelines:

- At any SWPM instance, never use the FQDN option.
- All hostnames must be short with a maximum length of 13 characters.
- Verify and, if needed, adjust the hostnames of the cluster member.

Move /usr/sap/<SID> to the Shared File System (Optional)

Note: This step is required only if you want to implement SAP HA by using SAPCTL on the VM cluster.

In this step, you move the local /usr/sap/<SID> of the first VM cluster node to the ACFS shared file system that you created earlier. Then, you create local symbolic links to this shared location on all the VM cluster nodes.

1. On the first node only, run the following commands:

```
[root@node0] # cd /usr/sap ; tar -cvf MFC.tar MFC
[root@node0] # cp MFC.tar /sapshare ; cd /sapshare ; tar -xvf MFC.tar
```

2. On all compute nodes, run the following commands:

```
[root@nodeN] # cd /usr/sap ; rm -f MFC.tar ; mv MFC was.MFClocal
[root@nodeN] # ln -s /sapshare/MFC MFC ; chown mfcadm:sapsys MFC
```

Install the Enqueue Replication Server (Optional)

Note: This step is required only if you want to implement SAP HA by using SAPCTL on the VM cluster.

Still in the context of [SAP Note 1877857](#), install the enqueue replication server (ERS) on all compute nodes locally (run sapinst without arguments or using only the option for a stack.xml file). On all nodes, choose the same instance number for ERS, for example, 01.

On the first node where the ASCS instance is running, perform the following steps:

1. Install the ERS instance. When SWPM asks for the hostname for ERS, enter the hostname of the VM cluster node, not the virtual hostname, which is ersmfc in our example configuration.
2. After ERS installation is completed on the first node, run the following commands as SIDADM:

```
node0: mfcadm> sapcontrol -nr 00 -function Stop
node0: mfcadm> sapcontrol -nr 00 -function StopService
```

On all subsequent nodes, perform the following steps for each node:

1. Relocate the temporary ASCS resource to the current node:

```
[root@nodeN] # . oraenv
+ASM2 (and +ASM3 and so forth)
[root@nodeN] # crsctl relocate res tmp -n <currentnodename>
```

2. Register the ASCS instance with saphostctl locally:

```
[root@nodeN] # /usr/sap/hostctrl/exe/saphostctl -function RegisterInstanceService -sid MFC -nr
00 -saplocalhost ascsmfc
```

3. Start saphostctl and the ASCS instance:

```
[root@nodeN] # su - mfcadm
nodeN: mfcadm> sapcontrol -nr 00 -function StartService MFC
nodeN: mfcadm> sapcontrol -nr 00 -function Start
```

4. Install the ERS instance. When SWPM asks for the hostname for ERS, enter the hostname of the VM cluster node, not the virtual hostname, which is ersmfc in our example configuration.

5. Stop the ASCS instance:

```
[root@nodeN] # su - mfcadm
nodeN: mfcadm> sapcontrol -nr 00 -function Stop
nodeN: mfcadm> sapcontrol -nr 00 -function StopService
```

After completing installation of ERS on all nodes, perform the following steps:

1. Relocate the temporary ASCS resource back to the first compute node:

```
[root@node0] # . oraenv
+ASM1
[root@node0] # crsctl relocate res tmp -n <currentnodename>
```

2. Start the ASCS services manually for the subsequent SAP database instance installation:

```
node0: mfcadm> sapcontrol -nr 00 -function StartService MFC
node0: mfcadm> sapcontrol -nr 00 -function Start
node0: mfcadm> sapcontrol -nr 00 -function GetProcessList
```

Now you are finished with the contents of [SAP Note 1877857](#). You will finish HA integration later.

Verify the /oracle/<SAPSID> Directory and Create Soft Links for ORACLE_HOMES

SAP distinguishes between an *installation* ORACLE_HOME (called IHRDBMS in this context) and a *runtime* ORACLE_HOME (called OHRDBMS in this context). The OHRDBMS is usually defined as a soft link named <VERSION> (for example, 19) under /oracle/<SAPSID>/ that points to the IHRDBMS where the Oracle Database software has been installed.

On an engineered system, you do not install Oracle RDBMS software; instead, you use the existing and preinstalled RDBMS software by creating a soft link to the IHRDBMS intended for use by SAP. The ORACLE_HOME environment variable is usually set to OHRDBMS. Wherever SWPM asks for the location of ORACLE_HOME for the Oracle Database, use the *runtime* ORACLE_HOME (OHRDBMS).

For example:

- OHRDBMS is at /oracle/MFC/19 (where 19 is a symbolic link to IHRDBMS).
- IHRDBMS is at /u02/app/oracle/product/19.0.0.0/dbhome_1.
- The ORACLE_HOME environment variable is set to /oracle/MFC/19.

Perform the following steps:

1. On each compute node, verify whether the following path exists:

```
/oracle/<SAPSID>/19
```

This might come from the host preparation and needs to be fixed.

2. As root on each compute node, run the following commands:

```
[root@nodeN] # mv /oracle/<SAPSID>/19 /oracle/<SAPSID>/was.19
```

3. Check your `/etc/oratab` file and ensure that it has an entry for the OHRDBMS that you want to use for the SAP database instance. For example:

```
MFC:/oracle/MFC/19:N
```

4. As the `oracle` user on each compute node, create a soft link for the respective `ORACLE_HOME`:

```
[oracle@nodeN] $ ln -s /u02/app/oracle/product/19.0.0.0/dbhome_1 /oracle/<SAPSID>/19
```

Patch Oracle RDBMS Homes Before Installing the Database

Before you can install the SAP database instance, you must patch the Oracle RDBMS software. For instructions, see “Installation of Patches for Oracle RDBMS Software.” This applies to all compute nodes.

Note: The only approved way to patch Oracle Database homes used for SAP is by using the SAP-provided SAP Bundle Patch for Exadata Database Service. Never apply patches to Oracle Database homes by using the Oracle Cloud Console or CLI. Doing so can result in an unusable and unsupported configuration and might cause an unplanned outage of your SAP environment. For more information, see “SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases.”

You also need to perform the post-installation steps in “Run the RDBMS Post-Installation Steps from Patching.” This applies to all compute nodes.

Install the Database Instance

Obtain the SCAN listener name (short hostname) and Node-VIPs through the Oracle Cloud Console or by running the following commands:

```
[opc@bb-xuwwj1 ~]$ sudo su -
[root@bb-xuwwj1 ~]# . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@bb-xuwwj1 ~]# srvctl config scan | egrep "IPv4 VIP|SCAN name"
SCAN name: bb-xuwwj-scan.sap.acme.com,
Network: 1
SCAN 1 IPv4 VIP: 172.19.20.47
SCAN 2 IPv4 VIP: 172.19.20.147
SCAN 3 IPv4 VIP: 172.19.20.19
[root@bb-xuwwj1 ~]# srvctl config vip -node `hostname -s` | egrep "VIP
Name|IPv4 Address"
VIP Name: bb-xuwwj1-vip.sap.acme.com
VIP IPv4 Address: 172.19.20.17
[root@bb-xuwwj1 ~]# srvctl config scan_listener | grep TCP
Endpoints: TCP:1521/TCP:2484
[root@bb-xuwwj1 ~]#
```

The Node-VIPs are required during the next run of SWPM, in which the database is created and loaded.

Adjust SAPDBHOST

Edit `/sapmnt/<SAPSID>/profile/DEFAULT.PFL` to adjust the `SAPDBHOST` parameter to the node where SWPM will run.

```
SAPDBHOST = bb-xuwwj1
j2ee/dbtype = ora
j2ee/dbname = MFC
j2ee/dbhost = bb-xuwwj1
```

Run the orabtt Script

1. With the correct environment variable set, run the `orabtt` script on each VM cluster node as the `oracle` user. In this example, `MFC` is used as the `SAPSID`. Adjust the `SAPSID` according to your environment.

```
[oracle@nodeN] $ export ORACLE_HOME=/oracle/MFC/19
[oracle@nodeN] $ $ORACLE_HOME/sap/orabtt/orabtt.sh -add -dbsid MFC
```

2. Verify as follows:

```
[oracle@nodeN] $ $ORACLE_HOME/bin/orabase
```

The output should return `/u02/app/oracle`.

Check Entries in /etc/oratab

Oracle CRS activity in the patching phase might destroy entries in `/etc/oratab`. For each VM cluster node, verify that the Grid Home is present in `/etc/oratab`. If it is not, add it as follows:

```
+ASM1:/u01/app/19.0.0.0/grid:N for the first node
+ASM2:/u01/app/19.0.0.0/grid:N for the second node, and so on
```

Ensure That ASCS Is Up on the First Node (Optional)

Note: This step is required only if you want to implement SAP HA by using `SAPCTL` on the VM cluster.

1. As the `root` user, ensure that the `tmp` resource is at `node0`:

```
[root@node0] # . oraenv [+ASM1]
[root@node0] # crsctl relocate res tmp -n <node0>
```

2. As the `SIDADM` user, run the following command:

```
node0: mfcadm> sapcontrol -nr 00 -function GetProcessList
```

3. If an error occurs or the connection is refused, run the following commands:

```
node0: mfcadm> sapcontrol -nr 00 -function StartService MFC
node0: mfcadm> sapcontrol -nr 00 -function Start
node0: mfcadm> sapcontrol -nr 00 -function GetProcessList
```

Invoke SWPM as the Root User

1. Run the following command:

```
[root@node0] # export TMP=/usr/sap/tmp
[root@node0] # </path/to/SWPM/>sapinst
```

Provide the `stack.xml` file as an option if required.

- For disk groups, choose +DATA<n> and +RECO<n>.
- At the Oracle RAC Parameters screen, Init.ora RAC parameters, adjust the following values (MFC is used as an example; adjust accordingly):

```
MFC001.local_listener = node0-vip:<port>
MFC002.local_listener = node1-vip:<port>
```

Also perform this step for any additional compute nodes that are being used.

- When SWPM asks for encryption of tablespaces, choose TDE for each tablespace listed by SWPM. This step is mandatory; all data in OCI must be encrypted.

At the end of this step, SWPM has finished creating your SAP database and created an `init<SIDxxx>.ora` and password file under `$ORACLE_HOME/dbs` on the first compute node. These files must be created on all subsequent VM cluster nodes later on.

- Start the database on one instance only before running `catsbp` in the next step. For example:

```
[root@nodeN ] # su - oracle
[oracle@nodeN ] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@nodeN ] $ export ORACLE_SID=MFC001

[oracle@nodeN ] $ srvctl stop database -db MFC -stopoption immediate -force
[oracle@nodeN ] $ sqlplus / as sysdba

Connected to an idle instance.

SQL> startup
ORACLE instance started.

Total System Global Area 2.7492E+10 bytes
Fixed Size                  18349312 bytes
Variable Size               1.3824E+10 bytes
Database Buffers           1.3623E+10 bytes
Redo Buffers                26054656 bytes
Database mounted.
Database opened.
SQL> exit
```

Run the RDBMS Post-Installation Steps from Patching

From the SAP Bundle Patch README, run `catsbp` after setting the required environment variables.

Finalize the Installation

Perform the following steps to finish the database installation, configure RMAN and SAP BR*Tools to perform backups, set up high availability for SAP central services, and install the primary application server.

Move the `saptrace` Directory and SAP BR*Tools Directories to a Shared File System

Because space in the local file systems of VM cluster nodes is limited, you must move the Oracle diagnostic destination, defined by the Oracle initialization parameter `diagnostic_dest`, to a shared file system location on ACFS. With SAP, this parameter typically points to the `/oracle/<DBSID>/saptrace` directory, for example, `/oracle/MFC/saptrace`. Moving the diagnostic destination to a shared location is also important for database-specific SAP transactions, where Oracle trace information is checked or viewed in SAP, or for special functions such as end-to-end tracing or monitoring.

If you want to use SAP BR*Tools on all VM cluster nodes—for example, for backup and restore, reorganizations, or database checks—you must also move the SAP BR*Tools-specific directories to a shared file system location. If brbackup logs are not in a shared file system location, you cannot restore and recover your database from a VM cluster node other than the one where the backup was taken.

To perform these actions, create new directories on the shared file system, and replace the original directories under /oracle/<DBSID> with symbolic links that point to the new directories.

The following example uses the ACFS file system /sapshare to store all these directories.

1. Shut down the database:

```
[root@nodeN ~] # . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@nodeN] # srvctl stop database -d MFC -stopoption immediate
```

2. From one VM cluster node, create the new shared directories and set the owner and permissions:

```
[root@nodeN] # mkdir -p /sapshare/MFC/saptrace
[root@nodeN] # mkdir -p /sapshare/MFC/saparch
[root@nodeN] # mkdir -p /sapshare/MFC/sapreorg
[root@nodeN] # mkdir -p /sapshare/MFC/sapbackup
[root@nodeN] # mkdir -p /sapshare/MFC/sapcheck
[root@nodeN] # cd /sapshare/MFC
[root@nodeN] # chown oracle:oinstall saptrace saparch sapreorg sapbackup sapcheck
[root@nodeN] # chmod 775 saptrace saparch sapreorg sapbackup sapcheck
```

3. On each VM cluster node, as the oracle user, rename the original directories to keep them:

```
[root@nodeN] # su - oracle
[oracle@nodeN] $ cd /oracle/MFC
[oracle@nodeN] $ mv saptrace saptrace.waslocal
[oracle@nodeN] $ mv saparch saparch.waslocal
[oracle@nodeN] $ mv sapreorg sapreorg.waslocal
[oracle@nodeN] $ mv sapbackup sapbackup.waslocal
[oracle@nodeN] $ mv sapcheck sapcheck.waslocal
```

4. On each VM cluster node, as the oracle user, create the following symbolic links:

```
[root@nodeN] # su - oracle
[oracle@nodeN] $ cd /oracle/MFC
[oracle@nodeN] $ ln -s /sapshare/MFC/saptrace saptrace
[oracle@nodeN] $ ln -s /sapshare/MFC/saparch saparch
[oracle@nodeN] $ ln -s /sapshare/MFC/sapreorg sapreorg
[oracle@nodeN] $ ln -s /sapshare/MFC/sapbackup sapbackup
[oracle@nodeN] $ ln -s /sapshare/MFC/sapcheck sapcheck
```

Configure Database Linux Huge Pages

Huge pages configuration is provided for the out-of-the-box database but not for custom sizing of the SAP database that is created from SWPM. Therefore, you must run the hugepages script provided by My Oracle Support Note 401749.1 when all required databases are running and all nonrequired databases are stopped. Adjust the value for vm.nr_hugepages in the /etc/sysctl.conf file for each compute node accordingly. To use the new huge pages configuration, shut down the SAP systems and all databases, and restart the nodes.

Check the database instance's alert_<DBSID>.log file to determine whether all huge pages have been allocated by the database instance.

Adjust Database Parameters

In addition to the parameters set during installation and patching, adjust your database parameters according to one of the following SAP Notes:

- [SAP Note 2470718 - Oracle Database Parameter 12.2 / 18c / 19c](#)
- [SAP Note 2378252 - Oracle Database Initialization Parameters for SAP NetWeaver Systems](#)

Configure HA-NFS for /sapmnt (Optional)

Note: This step is required only if you want to use the VM cluster for HA-NFS. If you are using an external NFS service, you need to mount your NFS exports (usually /sapmnt) on the relevant hosts of your SAP landscape. For more details how to export ACFS as a highly available NFS (HA-NFS) refer to MOS note 1934030.1.

1. On all VM cluster nodes, run the following commands as root, with a Grid Infrastructure environment set:

```
[root@nodeN] # systemctl enable rpcbind ; systemctl enable nfs-server
[root@nodeN] # systemctl start rpcbind ; systemctl start nfs-server
```

2. On the *first node only*, run the following commands:

```
[root@node0] # srvctl add havip -id hanfsmfc_id -address nfmfc -netnum 1 -description
"hanfsmfc"
[root@node0] # srvctl add exportfs -id hanfsmfc_id -path /sapmnt -name hanfsmfc -options
"rw, sync, no_root_squash" -clients <ip1, ip2, ip3, ...>
[root@node0] # srvctl start havip -id hanfsmfc_id
```

3. Check the status, location, and configuration of the APP-VIP for HA-NFS by running the following commands as root:

```
[root@node0] # srvctl status exportfs -id hanfsmfc_id
export file system hanfsmfc is enabled
export file system hanfsmfc is exported on node node0

[root@nodeN] # srvctl config havip
```

Configure RMAN and SAP BR*Tools to Perform Backups to Object Storage (Optional)

You can optionally configure RMAN and SAP BR*Tools to perform database backups either to OCI Object Storage or to a file system. The configuration of Object Storage or a file system as backup targets is optional, and you are free to choose a different way to perform backups.

Object Storage can be consumed as a durable, efficient, and fast destination for backups, and consequently, a restore and recovery source. In contrast to classic file systems, the interface to Object Storage is provided by an SBT_LIBRARY to Recovery Manager (RMAN).

Note: For instructions for using a file system instead of Object Storage, see “Configure RMAN and SAP BR*Tools to Perform Backups to the File System (Optional).”

Configure Object Storage

Step-by-step instructions for *non-SAP databases* are at [Backing Up a Database](#), and at least Java 7 is required to install. These instructions create the auto-open wallet, which we are configuring for the oracle OS user in \$ORACLE_HOME/dbs/<DBSID>_opc_wallet/cwallet.sso.

This auto-open wallet is different from and has nothing in common with the auto-open wallet used for TDE. It is used to establish a passwordless link between Object Storage and the compute node.

This link has the following effects:

- If your `cwallet.sso` is lost and you can't restore it for any reason, you can re-create it with the API key.
- If you lose your API key, you can get a new one and re-create the `cwallet.sso`.
- If you lose both your `cwallet.sso` and your API key, you can create a new API key and re-create the wallet.
- You must delete old, unused, unknown API keys.
- You can back up multiple databases into a bucket.
- You can have multiple buckets configured. Consider changing the configuration file (`config_db_name`, the `/lib` storage, and the wallet directory). Before you perform any operation, you must adjust RMAN's configuration as follows:

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=/path/to/lib/libopc.so,
SBT_PARMS=(OPC_PFILE=/path/to/wallet_config)';
```

Using Exadata Database Service, you can use Object Storage as a backup target and source for restore and recovery. However, with *SAP databases*, you cannot use the method provided at [Backing Up an Exadata Database](#). Instead, you must follow this outline:

1. Get the [Oracle Database Cloud Backup module](#). The most current version of the backup module is also included in the latest SAP Bundle Patch.
2. Extract the `oci_install.jar` file from the `oci_installer.zip` file, and copy it to all compute nodes as the `oracle` OS user.
3. Use the Oracle Cloud Console to obtain an API key and fingerprint and store the API key in a file, for example, in `/home/oracle/.oci/oci_api_key.pem`.
4. Determine the relevant Object Storage endpoint. A list of endpoints is available in the [API documentation](#).
5. Configure an Object Storage bucket.
6. To allow the Oracle Cloud Backup Module Installer (`oci_install.jar`) to access Object Storage to download the backup module, we recommend using the same proxy server configured earlier in this document for lifecycleing of `oraecsc01`. Alternatively, ensure that your NAT gateway and associated route rule are configured properly.

Then, perform the following steps:

1. Perform the following actions on each compute node as the `oracle` user:
 - A. Install the RMAN driver components for OCI in the user home for the `oracle` OS user.
 - B. Install the wallet and configuration information in an SAP DBSID-aware format in `$ORACLE_HOME/dbs`.
2. From the Oracle Cloud Console, obtain the OCIDs for your tenancy, user, and compartment, and the name of the Object Storage bucket.
3. Check the installation options for `oci_install.jar` (java is in your `$PATH` at user `oracle`):

```
[oracle@nodeN] $ java -jar opc_install.jar
```

4. Create a location to store the RMAN driver components for OCI:

```
[oracle@node0] $ mkdir -p ~oracle/cloud_backup/libdir
```

- As the oracle user, invoke the installer with your tenancy, user, bucket, and private key file along with the Object Storage endpoint on each node:

```
[oracle@node0] $ java -jar oci_install.jar \
-host https://objectstorage.eu-frankfurt-1.oraclecloud.com \
-pvtKeyFile ~/.oci/oci_api_key.pem \
-pubFingerPrint 98:96:b2:ff:d9:83:7c:51:ce:ee:f5:53:ea:fb:89:55 \
-tOCID ocid1.tenancy.oc1.<unique_ID> \
-uOCID ocid1.user.oc1.<unique_ID> \
-cOCID ocid1.compartment.oc1.<unique_ID> \
-bucket <bucket_name> \
-walletDir $ORACLE_HOME/dbs/MFC_opc_wallet \
-configFile $ORACLE_HOME/dbs/opcMFC.ora \
-libDir ~oracle/cloud_backup/libdir
-proxyHost <proxyhost>
-proxyPort <proxyport>
```

Because your compute nodes do not have internet access, add the `-proxyHost`, `-proxyPort`, and if needed, the `-proxyId` and `-proxyPass` directives. The `$ORACLE_HOME/dbs/MFC_opc_wallet` directory is created.

Configure RMAN

Configure RMAN to consume the object storage (have the proper environment for your DBSID set), also in accordance with SAP Note [1598594](#).

- Run the following commands:

```
[oracle@node0] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@node0] $ export ORACLE_SID=MFC001
[oracle@node0] $ rman target /
RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS
'SBT_LIBRARY=/home/oracle/cloud_backup/libdir/libopc.so,
SBT_PARMS=(OPC_PFILE=/u02/app/oracle/product/19.0.0/dbhome_1/dbs/opcMFC.ora)';
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO SBT_TAPE;
RMAN> CONFIGURE BACKUP OPTIMIZATION ON;
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP ON;
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F';
RMAN> CONFIGURE SNAPSHOT CONTROLFILE NAME TO '+DATA1/MFC/snapcf_MFC.f';
RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON;
```

- Set parallelism, which speeds up backup and restore activities:

```
RMAN> CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 16;
```

You are ready to run a backup. The following step is a test.

- Run a backup:

```
RMAN> BACKUP INCREMENTAL LEVEL 0 SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;
RMAN> restore database validate check logical;
```

- On the remaining hosts, confirm their settings after connecting with RMAN, and adjust parameters if needed:

```
RMAN> show all;
```

Now you can perform RMAN backups to Object Storage.

Prepare for SAP BR*Tools

A previous section described how to integrate Oracle RMAN with OCI Object Storage. This section describes how to integrate BR*Tools. BR*Tools must be at version 7.40 PL 32 or later. Earlier versions are not supported.

1. On each compute node, edit the `/oracle/<DBSID>/sapprof/init<DBSID>.sap` file (for example, `/oracle/MFC/sapprof/initMFC.sap`), and make following changes on each compute node:

```
_file_mask=002
rman_channels = 16
backup_dev_type = rman_disk
rman_sectionsize = 512M
backup_type = online
rman_compress = no | yes
rman_parms =
"SBT_LIBRARY=/home/oracle/cloud_backup/libdir/libopc.so,SBT_PARAMS=(OPC_PFILE=/u02/app/oracle/pr
oduct/19.0.0.0/dbhome_1/dbs/opcMFC.ora)"
```

Note: The steps described here are for passwordless operation of BR*Tools. With Oracle 12.1 and later, the `remote_os_authent` parameter defaults to FALSE. Therefore, BR*Tools needs to connect either with a dedicated user or by using SAP Secure Store; connections made with / don't work.

On Exadata Database Service, the Oracle ASM instance is running as a user grid that cannot write to directories owned by `oracle:oinstall`. Because of this, BR*Tools requires a 775 umask instead of 755 to allow the Oracle ASM instance to copy files to directories owned by `oracle:oinstall`. This is ensured by setting parameter “`_file_mask=002`”, as shown in this step.

2. On each compute node as the `oracle` user, verify that `$ORACLE_HOME/dbs` contains the appropriate `init<DBSID><index>.ora` file. If it does not, create it.

Following the previous examples, this file needs to be `initMFC001.ora` on `node0`, `initMFC002.ora` on `node1`, and so on. The content of the `init<DBSID><index>.ora` file is always identical.

```
#Generate initsid.ora for ASM spfile
spfile = (+DATAAC1/MFC/spfileMFC.ora)
```

3. At the same location for each node, as the `oracle` user, you need a password file `orapw<DBSID>` and a symbolic link `orapw<DBSID><index>` pointing to it. If the password file is not present at a particular node, copy the password file as the `oracle` user. In our example, we are logged in to a target node `nodeN` and copy the password file from a remote node `node0` to the local node. Replace MFC with the correct DBSID and repeat this step for each node to ensure that each node has its own password file.

```
[oracle@nodeN] # scp oracle@node0:/oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs/orapwMFC
```

4. Create the corresponding symbolic link on each node, replacing MFC with the correct DBSID and `<index>` with the correct instance number, for example, `001`, `002`, and so on.

```
[oracle@nodeN] $ ln -s /oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs/orapwMFC<index>
```

5. Ensure that all database instances can be started properly by using `srvctl`.

```
[root@node0] # su - oracle
[oracle@node0] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@node0] $ srvctl stop database -db MFC
[oracle@node0] $ srvctl start database -db MFC
```

- Invoke sqlplus / as sysdba, and run the following commands on one host only:

```
SQL> create user brt$adm identified by "sometrstrongpassword";
```

- Invoke sqlplus / as sysdba, and run the following commands on all compute nodes to update each Oracle password file with the required grants:

```
SQL> grant sapdba to brt$adm;
SQL> grant sysdba, sysoper to brt$adm;
```

All the commands need to succeed. If they do not, fix the password file.

- On each host, run the following commands:

```
[oracle@nodeN] $ cd /oracle/MFC ; mkdir -p security/rsecsfs ; cd security/rsecsfs
[oracle@nodeN] $ mkdir key data ; cd /oracle/MFC ; chmod 700 -R security
```

- Switch to SIDADM from root, for example, su - mfcadm.

- On each compute node, run the following command:

```
nodeN: mfcadm> brconnect -u / -c -f chpass -o 'BRT$ADM' -p 'sometrstrongpassword'

BR0801I BRCONNECT 7.40 (46)

BR0828I Changing password for database user BRT$ADM ...

BR0829I Password changed successfully in database for user BRT$ADM

BR1525I Setting password for user BRT$ADM in secure storage
/oracle/MFC/security/rsecsfs/data/SSFS_MFC.DAT ...

BR1526I Password set successfully for user BRT$ADM in secure storage
/oracle/MFC/security/rsecsfs/data/SSFS_MFC.DAT

BR0802I BRCONNECT completed successfully
```

- (Optional) Verify the proper operation of other BR*Tools:

```
nodeN: mfcadm> brspace -u // -c force
nodeN: mfcadm> brbackup -u // -q
nodeN: mfcadm> brarchive -u // -q
```

Configure RMAN and SAP BR*Tools to Perform Backups to the File System (Optional)

You can optionally configure RMAN and SAP BR*Tools to perform database backups either to OCI Object Storage or to a file system. The configuration of Object Storage or a file system as backup targets is optional, and you are free to choose a different way to perform backups.

Note: For instructions for using Object Storage instead of a file system, see “Configure RMAN and SAP BR*Tools to Perform Backups to Object Storage (Optional).”

Create a Backup Directory and Mount the NFS Share

To back up your database, we recommend providing a dedicated, highly available network file system (NFS) and mounting it on each VM cluster node. Ensure that the NFS is mounted through the backup network. Each VM cluster node should be configured to be able to perform database backups.

The following example shows the basic steps required:

1. On your NFS, export the NFS share with the following options: “rw, async, no_acl, no_root_squash”
2. On each VM cluster node, perform the following steps:
 - A. Create a backup directory and mount the NFS share to it. For example:

```
[root@nodeN] # mkdir /backup
[root@nodeN] # mount -t nfs <NFSSHOST>:/backup /backup
[root@nodeN] # mkdir /backup/sapbackup
[root@nodeN] # chown oracle:oinstall /backup/sapbackup
```

- B. Make the mount point persistent by adding it to /etc/fstab:

```
<NFSSHOST>:/backup      /backup      nfs      rw,bg 0 0
```

Configure RMAN

Configure backup target directories on one of the VM cluster nodes:

```
[root@nodeN] # su - oracle
[oracle@nodeN] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@nodeN] $ export ORACLE_SID=MFC001
[oracle@nodeN] $ rman target /

connected to target database: MFC (DBID=1234578402)

RMAN> CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/sapbackup/%U' maxpiecesize 128G;

new RMAN configuration parameters:
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/sapbackup/%U' MAXPIECESIZE 128 G;
new RMAN configuration parameters are successfully stored

RMAN> CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '/backup/sapbackup/%F';

new RMAN configuration parameters:
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '/backup/sapbackup/%F';
new RMAN configuration parameters are successfully stored
```

Prepare for SAP BR*Tools

BR*Tools must be at version 7.40 PL 32 or later. Earlier versions are not supported.

1. On each compute node, edit the /oracle/<DBSID>/sapprof/init<DBSID>.sap file (for example, /oracle/MFC/sapprof/initMFC.sap), and make the following changes on each VM cluster nodes:

```
_file_mask=002
rman_channels = 16
backup_dev_type = disk
disk_copy_cmd = rman_set
rman_compress = no | yes
backup_type = online
rman_sectionsize = 512M
backup_root_dir = <your backup root directory>
stage_root_dir = <your stage root directory>
archive_copy_dir = <your archive copy directory>
archive_copy_dir2 = <your second archive copy directory>
```

```
archive_stage_dir = <your archive stage directory>
```

Notes: The steps described here are for passwordless operation of BR*Tools. With Oracle 12.1 and later, the `remote_os_authent` parameter defaults to `FALSE`. Therefore, BR*Tools needs to connect either with a dedicated user or by using SAP Secure Store because connections made with / don't work.

On Exadata Database Service, the Oracle ASM instance is running as a user grid that cannot write to directories owned by `oracle:oinstall`. Because of this, BR*Tools requires a 775 umask instead of 755 to allow the Oracle ASM instance to copy files to directories owned by `oracle:oinstall`. This is ensured by setting parameter `"_file_mask=002"`, as shown in this step.

2. On each VM cluster node as the `oracle` user, verify that `$ORACLE_HOME/dbs` contains the appropriate `init<DBSID><index>.ora` file. If it does not, create it.

Following the previous examples, this file needs to be `initMFC001.ora` on `node0`, `initMFC002.ora` on `node1`, and so on. The content of the `init<DBSID><index>.ora` files is always identical.

```
#Generate initsid.ora for ASM spfile
spfile = (+DATAC1/MFC/spfileMFC.ora)
```

3. At the same location for each newly added node, as the `oracle` user, you need a password file `orapw<DBSID>` and a symbolic link `orapw<DBSID><index>` pointing to it. If the password file is not present at a particular node, copy the password file as the `oracle` user. In our example, we are logged in to a new target node `nodeN` and copy the password file from a preexisting remote node `node0` to the local node. Replace MFC with the correct DBSID and repeat this step for each new node to ensure that each node has its own password file.

```
[oracle@nodeN] # scp oracle@node0:/oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs/orapwMFC
```

4. Create the corresponding symbolic link on each new node, replacing MFC with the correct DBSID and `<index>` with the correct instance number.

```
[oracle@nodeN] $ ln -s /oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs/orapwMFC<index>
```

5. Ensure that all database instances can be started up properly using `srvctl`:

```
[root@nodeN] # su - oracle
[oracle@nodeN] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@nodeN] $ srvctl stop database -db MFC
[oracle@nodeN] $ srvctl start database -db MFC
```

6. Invoke `sqlplus / as sysdba`, and run the following commands on one host only:

```
SQL> create user brt$adm identified by "sometrangepassword";
```

7. Invoke `sqlplus / as sysdba`, and run the following commands on all VM cluster nodes to update each Oracle password file with the required grants:

```
SQL> grant sapdba to brt$adm;
SQL> grant sysdba, sysoper to brt$adm;
```

All the commands need to succeed. If they do not, fix the password file.

8. On each host, run the following commands:

```
[oracle@nodeN] $ cd /oracle/MFC ; mkdir -p security/rsecssfs ; cd security/rsecssfs
[oracle@nodeN] $ mkdir key data ; cd /oracle/MFC ; chmod 700 -R security
```


9. Switch to SIDADM from root, for example, `su - mfcadm`.
10. On each VM cluster node, run the following command:

```
nodeN: mfcadm> brconnect -u / -c -f chpass -o 'BRT$ADM' -p 'sometrstrongpassword'

BR0801I BRCONNECT 7.40 (46)

BR0828I Changing password for database user BRT$ADM ...

BR0829I Password changed successfully in database for user BRT$ADM

BR1525I Setting password for user BRT$ADM in secure storage
/oracle/MFC/security/rsecssfss/data/SSFS_MFC.DAT ...

BR1526I Password set successfully for user BRT$ADM in secure storage
/oracle/MFC/security/rsecssfss/data/SSFS_MFC.DAT

BR0802I BRCONNECT completed successfully
```

11. (Optional) Verify the proper operation of other BR*Tools:

```
nodeN: mfcadm> brspace -u // -c force
nodeN: mfcadm> brbackup -u // -q
nodeN: mfcadm> brarchive -u // -q
```

Set Up High Availability for SAP Central Services (Optional)

Note: This step is required only if you want to implement SAP HA by using SAPCTL on the VM cluster.

Use the latest `sapctl` package from [SAP Note 1496927](#), which is version 10.0 Patch 1 or later. The path `/oracle/GRID/19` must be replaced with `/u01/app/19.0.0.0/grid`.

Before invoking the `sapctl create` command, you must remove the temporary resource for the ASCS HA IP, `tmp`, that you created earlier. To do this, shut down ASCS, as follows:

```
nodeN: mfcadm> sapcontrol -nr 00 -function Stop

[root@node0] # . oraenv [+ASM1]
[root@node0] # crsctl stop res tmp
[root@node0] # appvipcfg delete -vipname=tmp
```

After deploying `sapctl` according to the instructions in [SAP Note 1496927](#), run the following commands, replacing MFC with the correct SAPSID:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl create -sapsid MFC -if bondeth0 -nm 255.255.255.0 -net
172.19.20.0 -nodes node0,node1 -abapenq ASCS00 -abapvip 172.19.20.208 -abapmsport 3900 -abaprep ERS01
-aersvip 172.19.20.69 -nx 1

[root@node0] # /usr/sap/sapctl/bin/sapctl start all -sapsid MFC
```

Install the Primary Application Server

The SAP primary application server (PAS) must be installed *on a separate Azure compute instance* connected to the subnet reserved for SAP NetWeaver application servers (`sn-app`). This section provides instructions for installing the PAS on Oracle Linux 8.

Relax the Password Policy

Before performing the following steps, create a backup of the `/etc/pam.d/system-auth` file and keep it safe for later restoration.

Then, edit the `/etc/pam.d/system-auth` file to remove `use_authok` from the lines starting with `password requisite pam_pwhistory.so` and `password sufficient pam_unix.so`.

Ensure that `/usr/sap` Is Not in the Root File System

Earlier in this document, when setting up the VM cluster nodes, you created a file system `/u02`, and then linked `/usr/sap` to `/u02/sap`. Do the same for the PAS, as follows:

```
[root@pas] # mkdir /u02/sap ; ln -s /u02/sap /usr/sap
```

Stop the Local Firewall

Run the following command to stop the local firewall:

```
[root@pas] # systemctl stop firewalld ; systemctl disable firewalld
```

Mount `/sapmnt` from HA-NFS (Optional)

Note: This step is required only if you want to use the VM cluster for HA-NFS. If you are using an external NFS service, you need to mount your shared directories from there. For example, if you are using Azure NetApp Files, the recommended mount options for `/sapmnt` are “`rw,hard,noatime,rsize=262144,wsize=262144,vers=3,tcp<IP>:<export> /sapmnt`”.

1. Run the following command:

```
[root@pas] # mkdir /sapmnt ; mount -t nfs nfsmfc:/sapshare/sapmnt /sapmnt
```

2. Edit `/etc/fstab` to add `/sapmnt`:

```
nfsmfc:/sapshare/sapmnt /sapmnt nfs rw,bg 0 0
```

3. Save and exit the file.

Install Additional Software Packages on OCI Compute Shapes

1. Run the following command as root:

```
[root@pas] # dnf -y install libnsl csh uuid
```

2. If you plan to access SWPM from a browser on the same host where `sapinst` is being started or you plan to use other software that requires a GUI (for example, SAP PlatinGUI), run the following commands as root to install the relevant GUI components and to temporarily enable the VNC server:

```
[root@pas] # dnf -y groupinstall "Server with GUI"
```

```
[root@pas] # dnf -y install tigervnc-server
```

3. If you installed the components in step 2, edit `/etc/gdm/custom.conf` and ensure that the `WaylandEnable=False` line is *not* commented out.

Note: These steps install only the VNC server, not the VNC system daemon. If you want to use the VNC server as a permanent service, you must also install and configure the VNC system daemon. For more information, see the [Oracle Linux 8 documentation](#).

Start the uidd Daemon

1. Start the uidd daemon and ensure that it comes up at system startup:

```
[root@pas] # systemctl start uidd ; systemctl enable uidd
```

2. Adjust the SELinux setting. Edit `/etc/selinux/config` and change `SELINUX=enforcing` to `SELINUX=permissive`.

3. Adjust the time zone to match the VM cluster nodes. Link `/etc/localtime` to your OS time zone file (Europe/Berlin in the following example):

```
[root@pas] # ln -sf /usr/share/zoneinfo/Europe/Berlin /etc/localtime
```

4. Start `ntpd` and configure it to start automatically:

```
[root@pas] # systemctl start chronyd ; systemctl enable chronyd
```

5. Restart the uidd daemon.

Start VNC the Server (Optional)

If you plan to access SWPM from a browser on the same host where `sapinst` is being started or you plan to use other software that requires a GUI (for example, SAP PlatinGUI), you installed the required GUI components in an earlier step. Now, set a password for VNC and start the VNC server with the `opc` user.

```
[opc@pas ~] $ vncpasswd
```

Password:

Verify:

Would you like to enter a view-only password (y/n)? n

A view-only password is not used

```
[opc@pas ~] $ vncserver -geometry 1280x1024
```

WARNING: `vncserver` has been replaced by a `systemd` unit and is now considered deprecated and removed in upstream.

Please read `/usr/share/doc/tigervnc/HOWTO.md` for more information.

New 'pas:1 (opc)' desktop is pas:1

Starting applications specified in `/home/opc/.vnc/xstartup`

Log file is `/home/opc/.vnc/pas:1.log`

If you do not plan to use the VNC server as a permanent service, stop `vncserver` after installation of PAS by running the following command:

```
[opc@pas ~]$ vncserver -kill :1
```

Provide a Strong Password for the opc User

This step is required only if you plan to use X-Windows on the PAS compute node (for example, for SAP PlatinGUI). To work with X-Windows, the `opc` user should require a strong password to unlock the X screensaver.

Run the following command and set a strong password:

```
[opc@pas] $ passwd
```

Note the password and keep it accessible.

Install and Configure SAP PAS on an Azure compute instance

1. As root, change the password for yourself on an Azure compute instance:

```
[root@pas] # passwd root
```

SWPM will ask for authentication in the next step.

2. Run SWPM:

```
[root@pas] # mkdir -p /usr/sap/tmp ; export TMP=/usr/sap/tmp ;
[root@pas] # cd /path/to/extracted_SWPM ; ./sapinst SAPINST_GUI_HOSTNAME=<pas-hostname>
```

Optionally, you can provide a `stack.xml` file if you have created one.

Use the *FQDN hostname* for `SAPINST_GUI_HOSTNAME` when invoking SWPM, and provide the *short hostnames* when you are asked to enter hostnames within SWPM.

The URL for your browser is displayed, for example:

```
https://sapapp3.sap.acme.com:4237/sapinst/docs/index.html
```

3. Invoke an SAP recommended browser and open the URL. When prompted, confirm a security exception.
4. At the authentication dialog box, enter the root user and its password and start PAS installation.

Configure the SAP GUI

Install and configure SAP GUI for Java on the bastion host that is running Oracle Linux. With the unified SAP front end, you can connect to SAP NetWeaver ABAP installations. Details are described in [SAP Note 146505](#) and on the [SAP Community Wiki](#).

SAP GUI for Java needs configuration information about your SAP environment, such as the names and addresses of your SAP servers. Based on this information, a connection directory is created that contains all available connections that can be selected in the SAP logon list. This directory can be centrally stored on a web server, and only a URL needs to be configured in SAP GUI for Java. Preset configuration and options can be distributed as templates during the initial installation process, so that a manual configuration after a first installation of SAP GUI for Java is not required. Access to the SAP ports for the connection needs to be created in the security lists.

Revert Changes on VM Cluster Nodes and SAP NetWeaver Application Servers

Now that installation is complete, make the following configuration changes:

- Revert the changes to password policies on the VM cluster nodes and the SAP NetWeaver application servers. Namely, restore their backups, such as `/etc/pam.d/system-auth`.
- If you added SSH keys that were needed for installation only to the `authorized_keys` or `key` files, delete them to avoid leaving vulnerabilities for potential security problems.

SAP Bundle Patch for Exadata Database Service: Lifecycle Management for SAP NetWeaver Databases

An Exadata Database Service instance requires regular patching at several levels of its software stack:

- Oracle Database software
- Oracle Grid Infrastructure software for systems that do not have the most recent software installed
- Oracle Exadata Database Service tooling
- Exadata image (OS)

This section describes how to install patches for the preceding components of an Exadata Database Service instance running databases for SAP applications. For more general information about Exadata Database Service versions and patches, see [My Oracle Support Note 2333222.1](#) (Exadata Database Service Software Versions).

Oracle manages some components of the Exadata Database Service:

- Exadata Storage Server hardware (Cell), firmware, and software
- Exadata Database Server hardware, firmware, and the Oracle Virtualization Server (OVS) management domain (DOM0)
- Storage switch hardware and firmware
- Power distribution units (PDU)

In all but exceptional circumstances, you will receive advance communication about updates to these components through the Cloud Notification Portal to help you plan for them. If there are corresponding recommended updates for your compute node VM environment, Oracle will also provide notification about these. You cannot opt out of updates. For more information, see [My Oracle Support Note 2124174.1](#) (Oracle Database Cloud Exadata Service Supported Software Versions and Planning for Updates).

Any updates to the preceding components are done by Oracle in a rolling manner. However, because of limitations in Oracle's connection handling, such a rolling update is not transparent to the SAP applications connected to the databases running on Exadata Database Service. Therefore, we recommend shutting down any SAP applications during an update of these components, if the update requires the shutdown of nodes of your Exadata Database Service instance.

Installation of Patches for Oracle RDBMS Software

An Exadata Database Service instance requires the following SAP Bundle Patch for an update of its Grid Infrastructure and database components (the patch version shown was current when this document was published):

- SAP Bundle Patch 19.22.0.0.0 for Exadata Cloud Service

Oracle regularly tests and certifies these patches and their successors for SAP databases and makes them available for SAP customers on My Oracle Support or the SAP Service Marketplace.

You can find up-to-date release information about the patches and their download locations in [SAP Note 2799970 - Oracle Exadata Cloud Service: Patches for 19c](#).

Note: The only approved way to patch Oracle Database homes used for SAP is by using the SAP-provided SAP Bundle Patch for Exadata Database Service. Never apply patches to Oracle Database homes by using the Oracle Cloud Console or CLI. This can result in an unusable and unsupported configuration and might cause an unplanned outage of your SAP environment. Always follow the instructions of the readme that ships with the SAP Bundle Patch.

Ensure that the OS environment requirements are fulfilled on all nodes of the Exadata Database Service instance as described in the “Operating System Environment Requirements” section of the SAP Bundle Patch Readme.

Then, follow the instructions in the “SAP Bundle Patch Installation” section to install the SAP Bundle Patch for Exadata Database Service. This process usually includes the following steps:

1. Install the latest OPatch and MOPatch utilities.
2. Install the database patches in the Oracle Database homes.
3. Run post-installation instructions, most notably the catsbp script.

Note: Depending on the content of the SAP Bundle Patch, its installation can be non-RAC-rolling. See the SAP Bundle Patch readme for more information.

Installation of Patches for Exadata Database Service Tooling

To patch the Oracle Exadata Database Service tooling, follow the [steps in the OCI documentation](#).

Installation of OS Patches

Use the tools and methods provided by the OS to prepare and populate the yum repository with Exadata channel content. The method is [outlined in the OCI documentation](#).

Finish SAP Monitoring Setup

If you install a new SAP system and do not replace the newly created SAP database with a migrated SAP database, you must apply numerous SAP-specific updates, support packages (SUM), and SAP notes. These contain bug fixes required to properly register VM cluster nodes with SAP monitoring.

If you replace the newly created SAP database with another, well-maintained SAP database that was migrated to Exadata Database Service, these bug fixes are most likely included in the migrated database.

For instructions on how to use SAP transaction RZ21 to register the nodes of the VM cluster for SAP central monitoring, see [Registering SAP NetWeaver Components and Hosts in CEN](#) in the SAP documentation.

To implement SAP enhanced monitoring for all Azure compute instances running SAP software, see the following Microsoft articles:

- [What is Azure Monitor for SAP solutions](#)
- [Configure SAP NetWeaver for Azure Monitor for SAP solutions](#)

Set Up SAP Diagnostic Agent for SAP Solution Manager

To integrate all the hosts of your SAP environment into SAP Solution Manager, install and configure an SAP Diagnostic Agent on each VM cluster node and on each compute node where SAP software is being run (for example, on the PAS).

To install the diagnostic agent, start SWPM and choose **Generic Options**, then **Diagnostics in SAP Solution Manager**, then **Install – Diagnostics Agent**.

All SAP diagnostic agents must be configured to run in “agent-on-the-fly” mode to handle the dynamics of a clustered database environment. More configuration steps might be required after installation; follow the relevant SAP documentation and SAP notes, for example, [SAP Note 1738351](#).

Using SAP Transaction DB13

If you want to schedule your database backups and database verification jobs with SAP transaction DB13, you must ensure that an SAP Gateway is running on each of the VM cluster nodes. The easiest way to do this is to install the ASCS integrated gateway, as described in the “Install the ASCS Instance (Optional)” section. With SAPCTL managing SAP HA, the SAP Gateway is always available as long as an ASCS instance can be started on one of the VM cluster nodes.

Before you run database-specific jobs from DB13, complete the following tasks:

1. Complete the steps in the “Prepare for SAP BR*Tools” section.
2. Install the primary application server (PAS) and possibly additional SAP NetWeaver application servers.
3. Configure and test the RFC connection to the SAP Gateway. This is done in SAP transaction SM59. You must use the virtual hostname. In the following example, the virtual hostname of the ASCS instance is `ascsmfc`. (Note that the following figures show `ascsmfg` instead of `ascsmfc`.)

The following two figures show an RFC example configuration for starting database-specific operations from SAP transaction DB13, configured within SAP transaction SM59.

The name of the RFC destination must be concatenated by the static text `SAPXPG_DBDEST_` plus the virtual hostname used for ASCS. In the example, the name of the RFC destination is `SAPXPG_DBDEST_ASCSMFC`.

On the **Technical Settings** tab, specify the following values:

- For **Activation Type**, select **Start on Explicit Host**.
- For **Program**, enter the full path plus the program name of the `sapxpg` executable. For example, `/sapmnt/sap/MFC/SYS/exe/uc/linuxx86_64/sapxpg`.
- For **Target Host**, enter the virtual hostname used for ASCS (for example, `ascsmfc`).
- For **Gateway Host**, enter the virtual hostname used for ASCS (for example, `ascsmfc`).
- For **Gateway service**, enter the name of the gateway service with the instance number used for ASCS. In this example, the service name is `sapgw00`. Refer to `/etc/services` to find the appropriate service name. The service name also contains the relevant SID number.

The screenshot shows the SAP SM59 configuration for an RFC destination named `SAPXPG_DBDEST_ASCSMFG`. The configuration is set to a TCP/IP Connection. The description is `ASCS-GW`. The **Technical Settings** tab is active, showing the following values:

- Program:** `/sapmnt/MFG/SYS/exe/uc/linuxx86_64/sapxpg`
- Target Host:** `ascsmfg`
- Save to Database as:** `ascsmfg` (selected as Host)
- Start Type of External Program:** `Default Gateway Value`

After successful connection testing, start DB13 and schedule a couple different jobs (for example, “Check database” and “Whole database online + redo log backup”) to see if they work correctly. For successful connection testing, you might need to adjust the gateway security files, for example by using SAP transaction SMGW.

If you experience a “logon error” in which brbackup, brarchive, or brconnect cannot connect to the database during jobs started from DB13, check [SAP Note 1764043](#).

Local Update Dispatching

SAP NetWeaver Application Server ABAP performs updates on the database asynchronously through *update processes*, also called UPD and UP2 processes. A dialog work process inserts an update job into the VBHDR, VBDATA, and VBMOD tables and posts the dispatcher process on the central instance to select an update process to actually perform the update job.

This update process may be connected to another RAC instance, similar to the dialog process. If so, the update process would read the data just written by the dialog process. To satisfy this read request, all required database blocks have to be shipped over the interconnect from the instance where the data was inserted to the instance where the data has to be read.

This approach can produce a *massive* amount of unnecessary additional data load and should be avoided. To avoid these “non-local updates,” we recommend the following actions:

- Have several update processes on *each* SAP instance
- Turn off update dispatching and use the local update processes

To turn off update dispatching, set SAP instance profile parameter `rdisp/vb_dispatching` to 0. Also, set the SAP instance profile parameter `rdisp/vb_name` to the name of the local instance (for example, `rdisp/vb_name = app1_MFC_00`).

Migrating Databases

Migration of existing SAP databases is discussed in a separate document, [Migrating SAP NetWeaver Based Systems Within the Scope of Oracle Databases](#). The recommended general way to migrate an existing SAP database is to set up a new SAP system on Exadata Database Service and then replace the new SAP database with a migrated one.

VM Cluster Node Subsetting

VM Cluster Node Subsetting allows VM clusters to span two or more DB servers in a flexible way by expanding or shrinking the VM cluster. VM cluster nodes can be added to a VM cluster if additional DB servers with enough resources are available, and nodes can be removed from a VM cluster if they are no longer needed, to free resources on DB servers. However, from an SAP perspective, `add-node` and `remove-node` operations are expensive in terms of the labor involved. Consider elastic scaling of compute resources (scaling OCPUs, local storage, or memory) as an alternative.

Recommendation for SAP HA with SAPCTL

Although VM Cluster Node Subsetting offers flexibility in the number of VM cluster nodes belonging to a VM cluster, adding or removing nodes always requires significant changes to the configuration of the VM cluster and is not a simple “one-click” operation.

Customers who use SAP HA with SAPCTL must eventually drop and re-create ASCS and ERS specific cluster resources to reflect the changed configuration. Unless you have good reasons to configure SAP HA with SAPCTL on *every* VM cluster node that may be available, we recommend that you configure it only on two or three VM cluster nodes to avoid reconfiguration of SAP HA during add-node or remove-node operations. Using two or three VM cluster nodes provides the best protection of SAP ASCS and ERS instances while keeping the probability low that SAP HA configuration must be changed.

Notes: Reconfiguration of SAP HA with SAPCTL requires downtime of the SAP central services ASCS and ERS to reconfigure where those services are configured to run.

The relocation of DB services disconnects all affected SAP NetWeaver application servers, which causes them to reconnect to the DB service. This action results in ended SAP transactions or ABAP short dumps. To avoid these unwanted interruptions, such actions should be issued during planned downtimes or by moving all SAP users and jobs to other SAP NetWeaver application servers, for example when they log on, so that the SAP NetWeaver application server whose DB service is being relocated has no users logged on and is not running any transactions.

Add Nodes

The add-node operation creates a VM cluster node by cloning it from an arbitrary existing VM cluster node and injecting specific parameters, such as IP addresses and hostnames, for the new VM cluster node. This cloning operation does not cover all aspects required to configure and run the new VM cluster nodes in a supported way, and numerous additional tasks must be performed after the add-node operation is completed.

Add VM Cluster Nodes

1. Open the Azure portal and navigate to the VM cluster.
2. Click **Settings** and then click **Virtual Machines**.
3. Click **Add**.
4. In the **Add Virtual Machine** dialog box, select the DB servers where you want to add VM cluster nodes and click **Submit**.

Note: By default, all available DB servers are selected. You must deselect the DB servers where you do not want to add new VM cluster nodes.

Note that this operation is long-running and can take several hours to complete.

Complete General Post Add-Node Tasks

After the add-node operation is completed, perform the following tasks.

Check the Private DNS Zones

After new nodes have been added to the VM cluster, you should first check all the private DNS zones for forward and reverse lookups in Azure and OCI. If A records or PTR records are missing (most likely for the Node-IPs and Node-VIPs of the new VM cluster node), add them and verify that resolution works as required.

Correct the Oracle Inventory on All VM Cluster Nodes

During the initial deployment of a VM cluster, you created the Oracle Database home. This creation process also created a corresponding Oracle inventory in which all the VM cluster nodes are listed. When a VM cluster node is added, the Oracle Database home is cloned from one of the existing cluster nodes. However, the Oracle inventories on all VM cluster nodes must be updated manually to reflect the new setup.

The simplest approach is to update the `/u01/app/oraInventory/ContentsXML/inventory.xml` file on one of the preexisting VM cluster nodes by adding the new VM cluster node, and then copy the contents of that file to each of the other nodes. Newly added VM cluster nodes usually do not contain any of the preexisting Oracle Database homes.

The following example shows adding the name of a new VM cluster node in the `/u01/app/oraInventory/ContentsXML/inventory.xml` file:

```
<?xml version="1.0" standalone="yes" ?>
<!-- Copyright (c) 1999, 2022, Oracle and/or its affiliates.
All rights reserved. -->
<!-- Do not modify the contents of this file by hand. -->
<INVENTORY>
<VERSION_INFO>
  <SAVED_WITH>12.2.0.7.0</SAVED_WITH>
  <MINIMUM_VER>2.1.0.6.0</MINIMUM_VER>
</VERSION_INFO>
<HOME_LIST>
<HOME NAME="OraGiHome19000" LOC="/u01/app/19.0.0.0/grid" TYPE="0" IDX="1" CRS="true"/>
<HOME NAME="OraHome1" LOC="/u02/app/oracle/product/19.0.0.0/dbhome_1" TYPE="0" IDX="2">
  <NODE_LIST>
    <NODE NAME="bb-xuwwj1"/>
    <NODE NAME="bb-xuwwj2"/>
    <NODE NAME="bb-xuwwj3"/>
  </NODE_LIST>
</HOME>
</HOME_LIST>
<COMPOSITEHOME_LIST>
</COMPOSITEHOME_LIST>
</INVENTORY>
```

Set Correct Permissions on \$GIHOME/network/admin

Because SWPM changes the owner of the `/u01/app/19.0.0.0/grid/network/admin` directory from `grid:oinstall` to `oracle:oinstall`, the permissions on this directory can cause problems writing to `listener.ora`.

On the new VM cluster nodes, change the owner to match the owner on the preexisting nodes and change permissions from 755 to 775.

```
[opc@nodeN] $ sudo chown oracle:oinstall /u01/app/19.0.0.0/grid/network/admin
[opc@nodeN] $ sudo chmod 775 /u01/app/19.0.0.0/grid/network/admin
```

Ensure that the owner and permissions are identical on each node.

Complete Post Add-Node Tasks on the New VM Cluster Nodes

Now perform the following tasks on the new VM cluster nodes.

Create an SAP-Specific /oracle/<SID>/19 Symbolic Link to IHRDBMS

Run the following commands:

```
[root@node2] # mkdir /oracle
[root@node2] # mkdir /oracle/MFC
[root@node2] # chown -R oracle:oinstall /oracle
[root@node2] # ln -s /u02/app/oracle/product/19.0.0.0/dbhome_1 /oracle/MFC/19
```

Adjust /etc/oratab

Add an entry for each database and each runtime Oracle Database home. For example:

```
MFC:/oracle/MFC/19:N
```

Install SAP-Enhanced Monitoring

Perform the steps described in the section “Provision SAP Monitoring.”

Relax the Password Policy

Run the following command:

```
[root@node2] # /opt/oracle.cellos/host_access_control pam-auth --deny 10 --lock 60 --pwquality 6 --remember 0
```

Adjust /etc/security/limits.conf

In the /etc/security/limits.conf file, add the following entries under the oracle entries:

```
root    soft    memlock    unlimited
root    hard    memlock    unlimited
ora<sid> soft    memlock    unlimited
ora<sid> hard    memlock    unlimited
<sid>adm soft    memlock    unlimited
<sid>adm hard    memlock    unlimited
```

Prepare SAP-Specific Shared Directories

If you configured an Oracle ACFS file system for sapshare containing your SAP binaries and /sapmnt, it is mounted automatically.

If /sapmnt was located on its own ACFS, expect it to be mounted.

If you are using an external sapmnt directory, mount it per node and make the mount persistent in /etc/fstab as you did on the preexisting nodes. Ensure that it's reboot safe.

If you are using an external NFS directory for backups, mount it per node and make the mount persistent in /etc/fstab as you did on the preexisting nodes. Ensure that it's reboot safe.

The add-node operation involves cloning from an unpredictable preexisting node's file system. However, the /u02 file system where you have the symbolic link set for /usr/sap, which usually directs to /u02/sap, is not cloned and must be created manually:

```
[root@node2] # mkdir /u02/sap
[root@node2] # ln -s /u02/sap /usr/sap
```

This demonstrates that `sapmnt` was on the ACFS file system `/sapshare`, and the link is lost. Rebuild it as follows:

```
[root@node2] # ln -s /sapshare/sapmnt /sapmnt
[root@node2] # ln -s /sapshare/trans /usr/sap/trans
```

Create the oper Group and Add the oracle User to It

Run the following commands:

```
[root@node2] # groupadd --gid 504 oper
[root@node2] # usermod oracle -a -G oper
```

Change the Root Password

Run the following command:

```
[root@node2] # passwd root
Changing password for user root.
New password:
BAD PASSWORD: The password is shorter than 15 characters
Retype new password:
passwd: all authentication tokens updated successfully.
```

Run SWPM Host Preparation

Run the following commands:

```
[root@node2] # mkdir -p /usr/sap/tmp ; export TMP=/usr/sap/tmp
[root@node2] # ./sapinst SAPINST_SAPINST_STACK_XML=/path/to/STACK.XML
```

On the SWPM dialog page, Database RAC Parameters, where the number of instances need to be entered, choose the new number of database instances (after the add-node operation).

On the SWPM dialog page where hostnames and instance numbers are associated, change the hostnames to match the order of the instances (the instances cannot be changed). For example:

	HOST NAME	INSTANCE NUMBER
1	node0	001
2	node1	002
3	node2	003

Copy Directories Under `/oracle/<SID>/` Not Created on the New VM Cluster Node

Copy the directories under `/oracle/<SID>/` that were not created on the new VM cluster node from one of the preexisting VM cluster nodes and adjust the ownership:

```
[oracle@node2] $ scp -r oracle@node0:/oracle/MFC/admin /oracle/MFC/
[oracle@node2] $ scp -r oracle@node0:/oracle/MFC/oraarch /oracle/MFC/
[oracle@node2] $ scp -r oracle@node0:/oracle/MFC/security /oracle/MFC/
[oracle@node2] $ chown -R oracle:asmadmin /oracle/MFC/admin
```

Copy the Database Password File from a Preexisting VM Cluster Node

As the oracle user in the \$ORACLE_HOME/dbs directory, you need a database password file named orapw<SID> and a symbolic link named orapw<SID><INSTANCENUMBER> pointing to it.

Copy the password file from another VM cluster node and create the symbolic link as shown in the following example. Replace MFC with the correct DBSID for your database.

```
[oracle@node2] $ scp oracle@node0: /oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs
[oracle@node2] $ ln -s /oracle/MFC/19/dbs/orapwMFC /oracle/MFC/19/dbs/orapwMFC003
```

Remove Obsolete Files from Oracle Database Home at \$ORACLE_HOME/dbs

Note that /u02 on the newly added VM cluster node was cloned from an already existing cluster node and contains several files in \$ORACLE_HOME/dbs are not required on the new node. Back up and then remove them from the Oracle Database home at \$ORACLE_HOME/dbs:

- orapw<SID>
- orapw<SID><INSTANCENUMBER>
- init<SID>.ora
- init<SID><###>.ora

Create an Instance-Specific pfile in Oracle Database Home at \$ORACLE_HOME/dbs

Run the following commands:

```
[oracle@node2 dbs] $ cp initMFC002.ora initMFC003.ora
[oracle@node2 dbs] $ cat initMFC003.ora
#Generate initsid.ora for ASM spfile
spfile = (+DATAC1/MFC/spfileMFC.ora)
```

Verify That orabase Returns the Correct Value

Run the following commands:

```
[root@node2] # su - oracle
[oracle@node2] # . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@node2] $ orabase
/u02/app/oracle
```

If orabase does not report /u02/app/oracle, run orabtt:

```
[oracle@node2] $ export ORACLE_HOME=/oracle/MFC/19 ; $ORACLE_HOME/sap/orabtt/orabtt.sh -add -dbid MFC
```

Set Owner, Group, and S-Bit for the User on the Oracle Binary in Oracle RDBMS Home \$ORACLE_HOME/bin

Run the following commands:

```
[oracle@node2] # ls -la /oracle/MFC/19/bin/oracle
-rwxr-x--x 1 oracle oinstall 449442792 Nov 17 10:59 /oracle/MFC/19/bin/oracle

[oracle@node2] # chown oracle:asmadmin /oracle/MFC/19/bin/oracle
[oracle@node2] # chmod gu+s /oracle/MFC/19/bin/oracle

[oracle@node2] # ls -la /oracle/MFC/19/bin/oracle
-rwsr-s--x 1 oracle asmadmin 449442792 Nov 17 10:59 /oracle/MFC/19/bin/oracle
```

Complete Post Add-Node Tasks on All VM Cluster Nodes

Now perform the following tasks on *all* the VM cluster nodes.

Adjust /oracle/MFC/sapprof/initSID.sap

Adjust the /oracle/MFC/sapprof/initSID.sap file so that all instances on all VM cluster nodes are included. For example, change the following lines as shown:

```
parallel_instances = (MFC001:/oracle/MFC/19@MFC001, MFC002:/oracle/MFC/19@MFC002) to
parallel_instances = (MFC001:/oracle/MFC/19@MFC001, MFC002:/oracle/MFC/19@MFC002,
MFC003:/oracle/MFC/19@MFC003)

asm_ora_home = (MFC001:/u01/app/19.0.0.0/grid, MFC002:/u01/app/19.0.0.0/grid) to
asm_ora_home = (MFC001:/u01/app/19.0.0.0/grid, MFC002:/u01/app/19.0.0.0/grid,
MFC003:/u01/app/19.0.0.0/grid)

asm_ora_sid = (MFC001:+ASM1, MFC002:+ASM2) to
asm_ora_sid = (MFC001:+ASM1, MFC002:+ASM2, MFC003:+ASM3)
```

Verify the Value of _file_mask in initSID.sap on All Nodes

Verify that the initSID.sap file on all VM cluster nodes has _file_mask set as follows:

```
_file_mask = 002
```

Add the New Instances to listener.ora on Preexisting VM Cluster Nodes

On each preexisting VM cluster node (here node0 and node1), add the new instances to listener.ora in GRID_HOME and reload the listener configuration:

```
[root@node0] # . oraenv
ORACLE_SID = [+ASM1] ?
The Oracle base remains unchanged with value /u01/app/grid
[root@node0] # vi $ORACLE_HOME/network/admin/listener.ora

SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = MFC001)
      (ORACLE_HOME = /oracle/MFC/19)
    )
    (SID_DESC =
      (SID_NAME = MFC002)
      (ORACLE_HOME = /oracle/MFC/19)
    )
    (SID_DESC =
      (SID_NAME = MFC003)
      (ORACLE_HOME = /oracle/MFC/19)
    )
  )

[root@node0] # lsnrctl reload

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER)))
The command completed successfully

[root@node0] # lsnrctl status
```

```

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER)))
STATUS of the LISTENER
-----
Alias                LISTENER
Version              TNSLSNR for Linux: Version 19.0.0.0.0 - Production
Start Date           14-JAN-2022 20:23:52
Uptime               2 days 17 hr. 9 min. 3 sec
Trace Level          off
Security             ON: Local OS Authentication
SNMP                 OFF
Listener Parameter File /u01/app/19.0.0.0/grid/network/admin/listener.ora
Listener Log File    /u01/app/grid/diag/tnslsnr/node0/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=10.0.1.56)(PORT=2484)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=10.0.1.32)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=10.0.1.56)(PORT=1521)))
Services Summary...
Service "MFC001" has 1 instance(s).
  Instance "MFC001", status UNKNOWN, has 1 handler(s) for this service...
Service "MFC002" has 1 instance(s).
  Instance "MFC002", status UNKNOWN, has 1 handler(s) for this service...
Service "MFC003" has 1 instance(s).
  Instance "MFC003", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully

```

Complete More Post Add-Node Tasks on the New VM Cluster Nodes

Now perform the following additional tasks on the new VM cluster nodes.

Configure HA-NFS for /sapmnt (Optional)

Note: This step is required only if you set up HA NFS on the VM cluster. If you are using an external NFS service, you need only to mount the directories as shown in the next section.

Run the following commands on the new VM cluster nodes:

```

[root@nodeN] # systemctl enable rpcbind ; systemctl enable nfs-server
[root@nodeN] # systemctl start rpcbind ; systemctl start nfs-server

```

Configure RMAN and SAP BR*Tools

Bring the new VM cluster nodes in line with the preexisting nodes and complete the steps required as shown in the section “Configure RMAN and SAP BR*Tools to Perform Backups to the File System.”

Move the saptrace Directory and Optionally SAP BR*Tools Directories to an Existing Shared File System Location

As you already did with the preexisting VM cluster nodes during their initial installation, you must move the Oracle diagnostic destination, defined by the Oracle initialization parameter `diagnostic_dest`, from the file system of the new VM cluster nodes to a shared file system location on ACFS. With SAP, this parameter typically points to the `/oracle/<DBSID>/saptrace` directory, for example, `/oracle/MFC/saptrace`. Moving the diagnostic destination to a shared location is also important for database-specific SAP transactions, where Oracle trace information is checked or viewed in SAP, or for special functions like end-to-end tracing or monitoring.

If you want to use SAP BR*Tools on all VM cluster nodes—for example, for backup and restore, reorganizations, or database checks—you must also move the SAP BR*Tools-specific directories to a shared file system location. If brbackup logs are not in a shared file system location, you cannot restore and recover your database from a VM cluster node other than the one where the backup was taken.

To perform this action, back up the original directories and replace them with symbolic links that point to the existing shared directories under `/oracle/<DBSID>`.

The following example uses the ACFS file system `/sapshare` for all these directories. Run these steps on each *new* VM cluster node.

As the `oracle` user, rename the original directories to keep them:

```
[root@node2] # su - oracle
[oracle@node2] $ cd /oracle/MFC
[oracle@node2] $ mv saptrace saptrace.waslocal
[oracle@node2] $ mv saparch saparch.waslocal
[oracle@node2] $ mv sapreorg sapreorg.waslocal
[oracle@node2] $ mv sapbackup sapbackup.waslocal
[oracle@node2] $ mv sapcheck sapcheck.waslocal
[oracle@node2] $ ln -s /sapshare/MFC/saptrace saptrace
[oracle@node2] $ ln -s /sapshare/MFC/saparch saparch
[oracle@node2] $ ln -s /sapshare/MFC/sapreorg sapreorg
[oracle@node2] $ ln -s /sapshare/MFC/sapbackup sapbackup
[oracle@node2] $ ln -s /sapshare/MFC/sapcheck sapcheck
```

For more information, see the following SAP notes:

- [2992680 - Managing shared and multiple Oracle Homes on Oracle Engineered Systems](#)
- [2884306 - Managing SAPDATA_HOME and ORACLE_BASE on Oracle Engineered Systems](#)

Adjust Linux Huge Pages

If you configured Linux huge pages according to the instructions in the section “Configure Database Linux Huge Pages,” can copy the determined value for `vm.nr_hugepages` from the preexisting VM cluster nodes and add it to `/etc/sysctl.conf`.

Reboot the New VM Cluster Nodes

Reboot the new VM cluster nodes and wait until they are fully back online by checking the status with the following command:

```
crsctl stat res -t
```

Complete Post Add-Node Tasks on One of the Preexisting Nodes

Now perform the following tasks on *one* of the preexisting VM cluster nodes.

Add Oracle Initialization Parameters for the New Instance to spfile

Run the following commands:

```
[root@node0] # su - oracle
[oracle@node0 ~]$ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@node0 ~]$ export ORACLE_SID=MFC001 ; sqlplus / as sysdba

alter system set instance_number=3 scope=spfile sid='MFC003';
```



```
alter system set instance_name='MFC003' scope=spfile sid='MFC003';
alter system set local_listener='node2vip:<port>' scope=spfile sid='MFC003';
alter system set log_archive_format='%t_%s_%r.dbf' scope=spfile sid='MFC003';
alter system set log_archive_trace=0 scope=spfile sid='MFC003';
alter system set thread=3 scope=spfile sid='MFC003';
alter system set undo_tablespace='PSAPUNDO003' scope=spfile sid='MFC003';
alter system set service_names='MFC','MFC003' scope=spfile sid='MFC003';
```

Add an Undo Tablespace

Add an undo tablespace with SAP naming convention PSAPUNDO### for the new instance:

```
CREATE UNDO TABLESPACE PSAPUNDO003 DATAFILE '+DATAC1' SIZE 700M AUTOEXTEND ON NEXT 20M MAXSIZE 10000M;
```

Add Online Redo Logs

Add online redo logs for the new instance and enable its thread:

```
alter database add logfile thread 3 group 51 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 52 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 53 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 54 '+DATAC1' size <sizeofyourredologs>;
alter database enable thread 3;
```

Add the New Instance to Oracle Cluster Repository and Start It

Run the following commands:

```
[root@node0] # su - oracle
[oracle@node0] $ . oraenv
ORACLE_SID = [oracle] ? MFC
The Oracle base has been set to /u02/app/oracle
[oracle@node0] $ srvctl add instance -d MFC -i MFC003 -n node2
[oracle@node0] $ srvctl start instance -d MFC -i MFC003
```

Check the Distribution of DB Services

After you add one or more VM cluster nodes to a VM cluster, DB services should usually be redistributed across all the VM cluster nodes. You should distribute DB services based on the workload, the corresponding SAP NetWeaver application servers, and the free resources available.

Important: Relocating DB services requires downtime of the SAP application because it causes broken connections, ORA-600 errors, and short dumps for the SAP application.

Modify the DB services by using the `srvctl modify service` command. Then, relocate them to their preferred node by using the `srvctl relocate service` command with the `-force` option. For example:

```
[root@node0] # srvctl modify service -db MFC -service MFC_D10 -preferred "MFC001" -available
"MFC002,MFC003" -modifyconfig

[root@node0] # srvctl relocate service -db MFC -service MFC_D10 -oldinst MFC003 -newinst MFC001 -force
```

Complete SAP HA-Specific Post Add-Node Steps (Optional)

Note: These steps are required only if you implemented SAP HA using SAPCTL and want to make the new VM cluster node available for SAP ASCS and ERS.

Preexisting VM Cluster Node Step

Stop ASCS and ERS on a preexisting cluster node, as follows:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl status all -sapsid MFC
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
SAP ABAP Enqueue service is ONLINE on node0
SAP ABAP Replication service is ONLINE on node1

[root@node0] # /usr/sap/sapctl/bin/sapctl stop all -sapsid MFC
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
Stopping SAP ABAP Replication service
Stopping SAP ABAP Enqueue service
SAP ABAP Enqueue service is OFFLINE
SAP ABAP Replication service is OFFLINE
Done
```

New VM Cluster Node Steps

Perform the following steps on a new VM cluster node:

1. Create a link for /usr/sap/<SID> to point to the shared directory:

```
[root@node2] # mv /usr/sap/MFC /usr/sap/was.MFC
[root@node2] # ln -s /sapshare/MFC /usr/sap/MFC
[root@node2] # chown mfcadm:sapsys /usr/sap/MFC
```

2. To register the instance service, relocate the ASCS APP-VIP to the new node:

```
[root@node2] # crsctl relocate res sap.MFC.abapvip -n node2
CRS-2673: Attempting to stop 'sap.MFC.abapvip' on 'node0'
CRS-2677: Stop of 'sap.MFC.abapvip' on 'node0' succeeded
CRS-2672: Attempting to start 'sap.MFC.abapvip' on 'node2'
CRS-2676: Start of 'sap.MFC.abapvip' on 'node2' succeeded
```

3. Register the ASCS instance service:

```
[root@node2] # /usr/sap/hostctrl/exe/saphostctrl -function RegisterInstanceService -sid MFC -nr
00 -saplocalhost ascsmf
Webmethod returned successfully
Operation ID: 525400123CE31EEC9E8CD7E3C62E9C8B

----- Log messages -----
Info: saphostcontrol: Executing 'sapstartsrv'
Info: saphostcontrol: exitcode=0
Info: saphostcontrol: 'sapstartsrv' successfully executed
Info: saphostcontrol: Executing 'sapstartsrv'
Info: saphostcontrol: exitcode=0
Info: saphostcontrol: 'sapstartsrv' successfully executed
```

4. Start the ASCS instance:

```
[root@node2] # su - mfcadm
node2:mfcadm 32> sapcontrol -nr 00 -function StartService MFC

StartService
OK
node2:mfcadm 33> sapcontrol -nr 00 -function Start

Start
OK
```

5. Run SAP SWPM to install an ERS instance.
6. Stop ASCS:

```
[root@node2] # su - mfcadm
node2:mfcadm 35> sapcontrol -nr 00 -function Stop

Stop
OK
node2:mfcadm 36> sapcontrol -nr 00 -function StopService MFC

18.01.2022 15:20:33
StopService
OK
```

7. Move the ASCS APP-VIP back to the previous node:

```
[root@node2] # . oraenv
ORACLE_SID = [root] ? +ASM3
The Oracle base has been set to /u01/app/grid
[root@node2] # crsctl relocate res sap.MFC.abapvip -n node0
CRS-2673: Attempting to stop 'sap.MFC.abapvip' on 'node2'
CRS-2677: Stop of 'sap.MFC.abapvip' on 'node2' succeeded
CRS-2672: Attempting to start 'sap.MFC.abapvip' on 'node0'
CRS-2676: Start of 'sap.MFC.abapvip' on 'node0' succeeded
```

8. Set up SAP HA with SAPCTL. For more information, see “Scope and Assumptions” and adjust the ERS instance profile for the newly added VM cluster node according to the SAPCTL documentation.
9. Re-create the SAP HA configuration with all VM cluster nodes having SAP ASCS and ERS installed:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl stop all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
Stopping SAP ABAP Replication service
Stopping SAP ABAP Enqueue service
SAP ABAP Enqueue service is OFFLINE
SAP ABAP Replication service is OFFLINE
Done

[root@node0] # /usr/sap/sapctl/bin/sapctl remove -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
CRS-2586: Deletion of a running resource 'ora.net1.network' requires the force option
CRS-4000: Command Delete failed, or completed with errors.
CRS-2586: Deletion of a running resource 'ora.net1.network' requires the force option
CRS-4000: Command Delete failed, or completed with errors.
Done

[root@node0] # /usr/sap/sapctl/bin/sapctl create -sapsid MFC -if bondeth0 -nm 255.255.255.0 -
net 172.19.20.0 -nodes node0,node1,node2 -abapenq ASCS00 -abapvip 172.19.20.208 -abapmsport
3900 -abaprep ERS01 -aersvip 172.19.20.69 -nx 1
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
Creating SAP ABAP VIP
Creating SAP ABAP ERS VIP
Creating SAP ABAP Enqueue resource
Creating SAP ABAP Replication resource
Done

[root@node0] # /usr/sap/sapctl/bin/sapctl start all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
Starting SAP ABAP Enqueue service
```

```
Starting SAP ABAP Replication service
SAP ABAP Enqueue service is ONLINE on node0
SAP ABAP Replication service is ONLINE on node0
Done
```

Note: You can ignore the error messages that say resource ora.net1.network cannot be deleted.

10. Confirm that SAP HA works as expected:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl relocate -sapsid MFC -abaprep ERS01 -to node2
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
Stopping SAP ABAP Replication service
Stopping SAP ABAP ERS VIP
Relocating SAP ABAP ERS VIP to node node2
Relocating SAP ABAP Replication service to node node2
SAP ABAP ERS VIP is ONLINE on node2
SAP ABAP Replication service is ONLINE on node2
Done

[root@node0] # /usr/sap/sapctl/bin/sapctl status all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
SAP ABAP Enqueue service is ONLINE on node0
SAP ABAP Replication service is ONLINE on node2

[root@node0] # /usr/sap/sapctl/bin/sapctl relocate -sapsid MFC -abapenq ASCS00 -to node2
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
Stopping SAP ABAP Enqueue service
Stopping SAP ABAP VIP
Relocating SAP ABAP VIP to node node2
Relocating SAP ABAP Enqueue service to node node2
SAP ABAP VIP is ONLINE on node2
SAP ABAP Enqueue service is ONLINE on node2
Done
```

Note: ERS should now move away after some minutes.

```
[root@node0] # /usr/sap/sapctl/bin/sapctl status all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
SAP ABAP Enqueue service is ONLINE on node2
SAP ABAP Replication service is ONLINE on node2
[root@node0] # /usr/sap/sapctl/bin/sapctl status all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
SAP ABAP Enqueue service is ONLINE on node2
SAP ABAP Replication service is OFFLINE
[root@node0] # /usr/sap/sapctl/bin/sapctl status all -sapsid MFC
sapctl version 10.0 Patch 1 Production Copyright 2019 Oracle. All rights reserved
SAP ABAP Enqueue service is ONLINE on node2
SAP ABAP Replication service is ONLINE on node0
```

Complete Final Post Add-Node Tasks on the New VM Cluster Nodes

Set up the SAP Diagnostic Agent for SAP Solution Manager on the new VM cluster node. For details, see “Set Up SAP Diagnostic Agent for SAP Solution Manager” earlier in this document.

Remove Nodes

This section describes the tasks associated with removing a VM cluster node from a VM cluster. Note that you must complete a couple of tasks *before* you remove the VM cluster node.

Reconfigure SAP HA (Optional)

Note: This step is required only if you implemented SAP HA using SAPCTL on the VM cluster. Reconfiguring SAP HA takes down SAP ASCS and ERS and requires downtime for the SAP application.

If SAP HA using SAPCTL is configured on the VM cluster node that you plan to remove, then SAP HA must be reconfigured to ensure that SAP ASCS and ERS are running on the VM cluster nodes that remain.

This reconfiguration is also required if SAP ASCS or SAP ERS is not actively running on the VM cluster node that you plan to remove but are configured in a way that they could run there.

1. Stop the ASCS and ERS instances:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl stop all -sapsid MFC
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
Stopping SAP ABAP Replication service
```

2. Remove all cluster resources associated with the ASCS and ERS instances:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl remove -sapsid MFC
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
CRS-2586: Deletion of a running resource 'ora.net1.network' requires the force option
CRS-4000: Command Delete failed, or completed with errors.
CRS-2586: Deletion of a running resource 'ora.net1.network' requires the force option
CRS-4000: Command Delete failed, or completed with errors.
Done
```

3. Re-create the cluster resources for the ASCS and ERS instances with the new list of VM cluster nodes (all the nodes that are *not* being removed):

```
[root@node0] # /usr/sap/sapctl/bin/sapctl create -sapsid MFC -if bondeth0 -nm 255.255.255.0 -
net 172.19.20.0 -nodes node0,node1 -abapenq ASCS00 -abapvip 172.19.20.208 -abapmsport 3900 -
abaprep ERS01 -aersvip 172.19.20.69 -nx 1
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
Creating SAP ABAP VIP
Creating SAP ABAP ERS VIP
Creating SAP ABAP Enqueue resource
Creating SAP ABAP Replication resource
Done
```

4. Restart the ASCS and ERS instances:

```
[root@node0]# /usr/sap/sapctl/bin/sapctl start all -sapsid MFC
sapctl version 10.0 Patch 1    Production Copyright 2019 Oracle.  All rights reserved
Starting SAP ABAP Enqueue service
Starting SAP ABAP Replication service
SAP ABAP Enqueue service is ONLINE on node0
SAP ABAP Replication service is ONLINE on node1
Done
```

Modify DB Service Definitions

Note: Relocate DB services *only* when the SAP application is stopped. Attempting to relocate DB services when the application is running breaks the connection between the application and the database, causes ORA-600 errors in the application, and causes the creation of short dumps by the application.

If you followed the SAP-specific guidelines for SAP NetWeaver on Oracle RAC databases, then you configured one dedicated database (DB) service for each SAP NetWeaver application server that connects to the SAP database. These DB services have one preferred database instance and one or more available database instances where they can run if the preferred database instance is not available.

If a VM cluster node (specifically, the DB instances running on the VM cluster node) will be removed from the VM cluster, then all DB services that can run on the VM cluster node must be configured to reflect that change. Distribute the DB services across the remaining available database instances to ensure that all the VM cluster nodes have a similar workload. Additionally, relocate the DB services to their preferred or available instances by using the `-force` option.

In the following example, DB service MFC_D10 (belonging to SAP NetWeaver application server MFC D10) has instance MFC003 configured as the preferred instance and instances MFC001 and MFC002 as the available instances. Because instance MFC003 will be removed, instance MFC001 is now configured as the preferred instance and instance MFC002 is configured as an available instance.

```
[root@node0] # srvctl relocate service -db MFC -service MFC_D10 -oldinst MFC003 -newinst MFC001 -force
[root@node0] # srvctl modify service -db MFC -service MFC_D10 -preferred "MFC001" -available "MFC002"
-modifyconfig
```

Remove a VM Cluster Node

1. Open the Azure portal and navigate to the VM cluster.
2. Click **Settings**, and then click **Virtual Machines**.
3. Click **Remove**.
4. In the **Remove Virtual Machines** dialog box, select the DB servers where you want to remove the VM cluster nodes. Note that two DB servers is the minimum for a VM cluster.
5. Click **Submit**.

Important: Be careful to remove the correct VM cluster node. This operation is long-running and may take several hours to complete. Wait until the remove-node operation has finished. Do *not* continue with the next steps until then.

Complete Post Remove-Node Steps

After the remove-node operation has completed, perform the following steps on one of the remaining nodes.

Note: Take care to disable the correct thread and drop the correct online redo logs and undo tablespace.

1. Disable obsolete redo log threads and drop unused redo logs. In this example, node 2 with redo log thread 3 and undo tablespace PSAPUNDO003 are removed.

```
SQL> select group#,thread#,members,status from v$log;
```

GROUP#	THREAD#	MEMBERS	STATUS
11	1	2	INACTIVE
12	1	2	INACTIVE
13	1	2	CURRENT

```

14      1      2 INACTIVE
21      2      2 INACTIVE
22      2      2 INACTIVE
23      2      2 INACTIVE
24      2      2 CURRENT
41      3      1 INACTIVE
42      3      1 INACTIVE
43      3      1 INACTIVE
44      3      1 INACTIVE

```

```

SQL> alter database disable thread 3;
SQL> alter database drop logfile group 41;
SQL> alter database drop logfile group 42;
SQL> alter database drop logfile group 43;
SQL> alter database drop logfile group 44;

```

- Drop the obsolete PSAPUNDO tablespaces:

```

SQL> select tablespace_name from dba_tablespaces where tablespace_name like 'PSAPUNDO%';

TABLESPACE_NAME
-----
PSAPUNDO001
PSAPUNDO002
PSAPUNDO003

SQL> drop tablespace PSAPUNDO003;

```

- Remove the obsolete database instance resource from the Oracle cluster repository:

```

[root@node0] # . oraenv
ORACLE_SID = [+ASM1] ?
The Oracle base remains unchanged with value /u01/app/grid

[root@node0] # srvctl remove instance -db MFC -instance MFC003
Remove instance from the database MFC? (y/[n]) y

```

- Remove all the remains of the deleted VM cluster node in `init<SID>.sap`.
- Remove the deleted VM cluster node from the Oracle central inventory of all remaining VM cluster nodes. Edit `/u01/app/oraInventory/ContentsXML/inventory.xml` on all VM cluster nodes.

```

<?xml version="1.0" standalone="yes" ?>
<!-- Copyright (c) 1999, 2023, Oracle and/or its affiliates.
All rights reserved. -->
<!-- Do not modify the contents of this file by hand. -->
<INVENTORY>
  <VERSION_INFO>
    <SAVED_WITH>12.2.0.7.0</SAVED_WITH>
    <MINIMUM_VER>2.1.0.6.0</MINIMUM_VER>
  </VERSION_INFO>
  <HOME_LIST>
    <HOME NAME="OraGiHome19000" LOC="/u01/app/19.0.0.0/grid" TYPE="0" IDX="1" CRS="true"/>
    <HOME NAME="OraHome1" LOC="/u02/app/oracle/product/19.0.0.0/dbhome_1" TYPE="0" IDX="2">
      <NODE_LIST>
        <NODE NAME="bb-xuwwj1"/>
        <NODE NAME="bb-xuwwj2"/>
        <NODE NAME="bb-xuwwj3"/>
      </NODE_LIST>
    </HOME_LIST>
  </INVENTORY>

```

```

</HOME>
</HOME_LIST>
<COMPOSITEHOME_LIST>
</COMPOSITEHOME_LIST>
</INVENTORY>

```

Note: Some cluster resources cannot be removed and remain in the Oracle cluster repository.

Delete an Entire VM Cluster

To delete a VM cluster that is no longer needed, follow these steps. At the time this document was written, it is not possible to create a new VM cluster with a database name (SID) that was already used if the prior VM cluster with its databases was not properly deleted. This applies to unregistered databases such as in SAP environments.

1. Shut down *all* database instances of all databases running on the VM cluster and remove the database resource:

```

[opc@bb-xuwwj1 ~]$ sudo su -
[root@bb-xuwwj1 ~]# . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@bb-xuwwj1 ~]# srvctl stop database -db MFC -stopoption IMMEDIATE -force
[root@bb-xuwwj1 ~]# srvctl remove database -db MFC

```

2. Stop all customer-created ACFS file systems and drop the underlying volumes:

```

[root@bb-xuwwj1 ~]# srvctl status volume
Volume SAPSHARE_V1 of diskgroup DATA1 for device /dev/asm/sapshare_v1-128 is enabled
Volume SAPSHARE_V1 of diskgroup DATA1 for device /dev/asm/sapshare_v1-128 is running

[root@bb-xuwwj1 ~]# srvctl stop filesystem -device /dev/asm/sapshare_v1-128 -f

srvctl remove filesystem -device /dev/asm/sapshare_v1-128
srvctl remove volume -volume SAPSHARE_V1 -diskgroup DATA1 -force

```

3. Remove database-related files on all ASM disk groups:

```

[root@bb-xuwwj1 ~]# asmcmd
ASMCMD> rm -rf DATA1/MFC
ASMCMD> rm -rf RECO1/MFC

```

4. To delete the VM cluster:
 - A. In the Azure portal, navigate to the VM cluster.
 - B. On the **Overview** page, click **Delete**.
5. Note that the Azure portal may remove the VM cluster from the list of VM clusters after one hour or show the VM cluster in “deleted” state. In some cases, deletion of a VM cluster may take longer. We recommend monitoring the process in the Oracle Cloud Console.

High Availability with Oracle Data Guard

Oracle Data Guard provides additional high availability for the Oracle database of an SAP installation with Exadata Database Service. Only physical standby is supported for an SAP environment. The physical standby database runs on a separate Exadata machine or OCI Compute instance that fulfils the same SAP system requirements as the primary database, for example, identical OS user and group IDs. The Oracle database software must be installed by using SWPM to the same location as the primary site (/oracle/<SID>) and run on the same release and patch level as the primary database. Ensure that you sufficiently test the reconnection of the SAP instances to the standby database.

For more information about using Oracle Data Guard in Exadata Database Service, see the [Exadata Cloud Service documentation](#).

Support of SAP Standalone Enqueue Server 2 and Enqueue Replicator 2

SAP Standalone Enqueue Server 2 and Enqueue Replicator 2 *are not the default* for SAP HA with SAPCTL. Additionally, their support with Oracle is limited to specific SAP kernel versions, and SWPM does not offer to install and configure them.

Customers who want to use SAP Standalone Enqueue Server 2 and Enqueue Replicator 2 for SAP HA with SAPCTL must first check the relevant SAP Notes and [documentation](#) to determine whether SAP kernel support is available and then configure them manually.

References

SAP

Most of the SAP links require SAP login credentials for access.

SAP Documentation

- [SAP Product Availability Matrix \(PAM\)](#)
- [SAP Software Logistics Toolset \(SL Tools\)](#)
- [SAP Download Manager](#)
- [SAP Software Download Center \(SWDC\)](#)
- [SAP NetWeaver Guide Finder](#)
- [SAP Community Network - Oracle Community](#)
- [SAP Help Portal: TCP/IP Ports of All SAP Products](#)

SAP Notes

- [2614028 - SAP NetWeaver on Oracle Database Exadata Cloud Service](#)
- [1928533 – SAP Applications on Microsoft Azure: Support Products and Azure VM Types](#)
- [3475554 - Updating SAP Metrics Collector \(oraecscol\) for Oracle Exadata Cloud Infrastructure and Oracle Exadata Cloud@Customer](#)
- [2992680 - Managing shared and multiple Oracle Homes on Oracle Engineered Systems](#)
- [1778431 - SAP Installations using Shared Oracle Homes](#)

- [2884306 - Managing SAPDATA HOME and ORACLE BASE on Oracle Engineered Systems](#)
- [2614080 - SAP on Linux with Oracle Database Exadata Cloud Service: Enhanced Monitoring](#)
- [2474949 - SAP NetWeaver on Oracle Cloud Infrastructure](#)
- [2470718 - Oracle Database Parameter 12.2 / 18c / 19c](#)
- [2997175 - Oracle: utlrp.sql has failed with ORA-4068 during execution of catsbp](#)
- [2378252 - Oracle Database Initialization Parameters for SAP NetWeaver Systems](#)
- [2520061 - SAP on Oracle Cloud Infrastructure: Support prerequisites](#)
- [1565179 - SAP software and Oracle Linux](#)
- [1597355 - Swap-space recommendation for Linux](#)
- [1770532 - HugePages on Linux for Oracle Database](#)
- [146505 - SAP GUI for the Java Environment](#)
- [2799900 - Central Technical Note for Oracle Database 19c](#)
- [1868094 - Overview: Oracle Security SAP Notes](#)
- [1996481 - Using correct hostnames for Oracle Exadata Database Nodes](#)
- [1496927 - Protection of SAP instances through Oracle Clusterware](#)
- [2591575 - Using Oracle Transparent Data Encryption \(TDE\) with SAP NetWeaver](#)
- [2799991 - TDE Encryption Conversions for Tablespaces and Databases](#)
- [1598594 - BR*Tools configuration for Oracle installation using user "oracle"](#)
- [113747 - Owners and authorizations of BR*Tools](#)
- [776505 - ORA-01017/ORA-01031 in BR*Tools on Linux and Solaris 11](#)
- [2799970 - Patches for 19c: Oracle Exadata Cloud Service](#)
- [2799900 - Central Technical Note for Oracle Database 19c](#)
- [2422996 - Oracle: OPatch Versions 12.2.0.1.8, 11.2.0.3.18 and Newer](#)
- [3002467 - Oracle: MOPatch 2.7 and 2.7.1 fail with error "Cannot escalate patch"](#)
- [3215426 - Engineered Systems: mod sapinit may fail for SAPHOSTAGENT PL 56 and beyond \(ASCS HA\)](#)
- [1858920 - Diagnostics Agent installation with Software Provisioning Manager](#)
- [2711036 - Usage of the Standalone Enqueue Server 2 in an HA Environment](#)
- [2630416 - Support for Standalone Enqueue Server 2](#)
- [3248703 - Automation Options with Exadata Cloud Infrastructure](#)
- [1738351 - Connecting Oracle RAC to Solution Manager 7.1/7.2](#)
- [129997 - Hostname and IP lookup](#)
- [1100926 - FAQ: Network Performance](#)
- [3499586 - Configuring Database Services and Local Update Dispatching for all SAP databases using Oracle RAC](#)

Oracle

- [Oracle Exadata Cloud Service X9M](#)
- [Exadata Database Service on Dedicated Infrastructure](#)
- [Oracle Cloud Infrastructure – Exadata DB Systems](#)
- [Oracle Cloud Hosting and Delivery Policies](#)
- [Oracle Database](#)
- [Oracle Linux](#)
- [Oracle-SAP Solutions site](#)
- [My Oracle Support Note 1934030.1 - ODA \(Oracle Database Appliance\): HowTo export ACFS \(cloudfs\) using HANFS](#)

Azure

- [Azure Virtual Machines deployment for SAP NetWeaver](#)

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2024, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners

Author: Markus Breunig

Contributing Authors: Torsten Grambs, Jan Klokckers, Ralf Klahr