



ORACLE

SAP NetWeaver® Application Server ABAP/Java with Oracle Exadata Database Service

Based on Oracle Exadata Cloud Infrastructure X11M

May 2025, Version 1.0

Copyright © 2025, 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, timing, and pricing 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 paper.

Date	Revision
May 2025	Initial certification and publication

Table of Contents

Scope and Assumptions	5
Terminology	6
SAP Application Servers and Database Services	8
The Role of SAP Software Provisioning Manager (SWPM)	8
Overview of Oracle Exadata Database Service	9
SAP Database on Oracle Exadata Database Service	9
Regions and Availability Domains	9
Oracle Cloud Infrastructure Services	10
Virtualization and Databases with Exadata Cloud Infrastructure X11M	11
VM Cluster Scaling Options	12
Oracle Database for SAP with Exadata Cloud Infrastructure X11M	12
Exadata Cloud Infrastructure X11M	12
Planning Your Exadata Database Service for SAP NetWeaver Application Server ABAP/Java	15
Prerequisites	15
Deployment Restrictions with SAP for Oracle Exadata Database Service	15
Support	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 OCI Account	19
Deploy the OCI Resources	20
Prepare Your Environment	23
Prepare VM Cluster Nodes	24
Install SAP NetWeaver Application Server ABAP/Java	29
Finalize the Installation	35
Revert Changes on the VM Cluster Nodes and the SAP Application Servers	49
SAP Bundle Patch for Oracle Exadata Database Service: Lifecycle Management for SAP Databases	49
Installation of Patches for Oracle RDBMS Software	50
Installation of Patches for the Oracle Exadata Database Service Tooling	50
Installation of OS Patches	50
Finish SAP Monitoring Setup	50
Set Up SAP Diagnostic Agent for SAP Solution Manager	51

Using SAP Transaction DB13	51
Local Update Dispatching	52
Migrating Databases	53
Installing Additional SAP Application Servers	53
VM Cluster Node Subsetting	53
Adding Nodes	54
Set Up SAP Diagnostic Agent for SAP Solution Manager on the New VM Cluster Node	66
Removing Nodes	66
High Availability with Oracle Data Guard	69
Support of SAP Standalone Enqueue Server 2 and Enqueue Replicator 2	69
References	70
SAP	70
Oracle	71

Scope and Assumptions

This document is a reference guide for deploying Oracle databases of the SAP NetWeaver® Application Server ABAP/Java platform on the Oracle Exadata Database Service, based on Oracle Exadata Cloud Infrastructure X11M running Oracle Linux 8. Customers who want to ensure that they have an SAP-supported installation should strictly follow this installation guide.

Note: For historical reasons, the terms *Oracle Exadata Database Service* (or *Exadata Database Service*) and *Oracle Exadata Cloud Service* (or *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 add-on. This is an *optional* method for implementing SAP HA. Customers who do not require SAP HA can *skip* the steps related to installing and configuring SAP ABAP central services (ASCS) and enqueue replication server (ERS).

Note: *Application-specific virtual IP addresses* defined for VM clusters running on Oracle Exadata Cloud Infrastructure X11M are required only by customers who plan to implement SAP HA using SAPCTL directly on Oracle Exadata Cloud Infrastructure X11M. Customers who do not want to implement SAP HA or want to implement a supported SAP HA solution other than SAPCTL do not need application-specific virtual IP addresses.

Additionally, this document describes how to configure a separate compute instance running Oracle Linux 8 as an SAP NetWeaver primary application server connected to a VM cluster running on Oracle Exadata Cloud Infrastructure X11M.

Notes: 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 compute instances (bare metal or virtual machine compute instances) within OCI connected to the client network of the VM cluster.

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 a network file system (NFS) mount not exported by VM cluster nodes.

Customers *must* use SAP Software Provisioning Manager (SWPM) to install and configure all SAP-related components, regardless of whether or not they want to implement SAP HA with SAPCTL. This means that customers must use SWPM to perform SAP host preparation on each VM cluster node as well as to install the SAP database instance and the SAP Diagnostic Agent. The *only* optional installation steps are those related to configuring SAP HA with SAPCTL. For more information, see “The Role of SAP Software Provisioning Manager (SWPM).”

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 landscape in Oracle Cloud in a supported and verified way.

This document requires the following knowledge:

- You are familiar with the fundamentals of Oracle Exadata Database Service and OCI. For information, see the following resources:
 - [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. For more information, see the following resources:
 - [SAP NetWeaver Platform](#)
 - [SAP on Oracle](#)

- [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 the VM cluster nodes
 - Oracle Linux 9 and Oracle Linux 8 as the OS for running a bastion host or SAP Application Server
 - SAP NetWeaver 7.x

Important: 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 (see SAP Note [2799970](#) for more details). Never apply patches to Oracle Database homes by using the Oracle Cloud Console or CLI. Furthermore, never apply SAP-provided Bundle Patches that have been built for environments other than Exadata Database Service, for example, for Oracle Exadata on-premises systems. 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 Oracle Exadata Database Service: Lifecycle Management for SAP Databases.”

Terminology

The following tables define some of the terms used within the context of Oracle Exadata Database Service and this document. They also define terms related to Oracle-specific and SAP-specific 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 Oracle Exadata Database Service and This Document

TERM	DEFINITION
Exadata Cloud Infrastructure	The specific physical infrastructure that powers an Oracle 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 virtual machines (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. This includes the presence of 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, KVM guest, DomU, or virtual compute node.
Virtual cloud network (VCN)	A virtual, private network that you set up in Oracle data centers. It 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 Oracle 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.

Table 2: Oracle-Specific Work Areas

ORACLE WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
Oracle Cloud Console	Used to manage all the resources in OCI, for example, the deployment of VCNs, VM clusters for Oracle Exadata Database Service, or independent compute nodes for SAP application servers.	Register at cloud.oracle.com .
Oracle Linux	Used on all VM cluster nodes. Other compute instances might run different OSs.	Oracle Linux 8 on VM cluster nodes is managed from the Oracle Cloud Console. Oracle Linux or other OSs on other compute instances must be managed by the customer. For Oracle Linux 8, see SAP Note 2936683 . For Oracle Linux 9, see SAP Note 3399081 .
Oracle Grid Infrastructure	Manage the HA of all database services, ACFS, and NFS by using tools like SRVCTL and CRSCTL, as well as application virtual IP addresses.	For HA-NFS, 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)	Back up, restore, and recover your SAP database.	None

Table 3: SAP-Specific Work Areas

SAP WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
SAP Maintenance Planner	Create a <code>stack.xml1</code> file for SAP Software Provisioning Manager (SWPM) and choose the SAP software components that you want to install.	See Maintenance Planner – User Guide .
SAP Software Provisioning Manager (SWPM)	Use for <i>mandatory</i> tasks such as SWPM-based host preparation, installation of the SAP database instance and the primary application server (PAS), and for optional tasks such as the installation of ABAP system central services (ASCS), enqueue replication server (ERS), and additional application servers (AAS).	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 WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
SAP NetWeaver software stack	Modify SAP instance profiles, configure RFC connections, and schedule database backups and checks in SAP transaction 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 .

SAP Application Servers and Database Services

With Oracle RAC, it is absolutely essential that each SAP NetWeaver Application Server instance connects to its own dedicated database service that has been specifically created for that instance. Furthermore, this database service must be configured in such a way that it is provided on only one single RAC database instance at any given time. This guarantees that database connections from an SAP NetWeaver Application Server instance are not spread across multiple RAC instances, which can lead to significant performance and availability issues in the SAP system. All database connections from an SAP NetWeaver Application Server instance must be directed to one RAC instance. See SAP Note [527843](#).

The Role of SAP Software Provisioning Manager (SWPM)

Customers must use the SAP Software Provisioning Manager (SWPM) for all SAP installation and configuration tasks to ensure their systems end up in a defined and tested state. Not using SWPM results in a system that does not meet the requirements for certification. In the worst case, this could result in the customer not receiving technical support for this system from either SAP or Oracle.

Most customers do not typically plan a completely new installation but instead migrate their SAP system or the existing database to Oracle Cloud Infrastructure and onto Exadata Cloud Infrastructure. However, just like with a fresh installation, a completely new installation is required beforehand for compliance reasons. Afterwards, the newly created database is fully replaced with the source database to be migrated.

SWPM must be used for the following tasks:

- Host preparation of all VM cluster nodes (mandatory)
- Installation of the SAP database instance on the VM cluster nodes (mandatory)
- Installation of the SAP Primary Application Server (PAS) on an OCI compute instance (mandatory)
- Installation of additional SAP Application Servers (AAS) on additional OCI compute instances (optional)
- Installation of the SAP Diagnostic Agent on all VM cluster nodes (mandatory)
- Installation of the SAP Diagnostic Agent on the PAS host (mandatory)
- Installation of the SAP Diagnostic Agent on the AAS hosts (mandatory on AAS hosts)
- Installation of the SAP ABAP System Central Services and Enqueue Replication Services (mandatory)

Note: During the installation of the PAS or an AAS, SWPM configures a dedicated database service in Oracle Clusterware to which the PAS or the AAS will connect. The database services must be reconfigured manually if VM cluster nodes are added to or removed from the VM cluster.

Overview of Oracle Exadata Database Service

Oracle Exadata Database Service enables full-featured Oracle databases to run on dedicated Oracle Exadata Cloud Infrastructure X11M in the Oracle public 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 for SAP Central Services

SAP Database on Oracle Exadata Database Service

All options and features certified for on-premises deployments of SAP NetWeaver of Oracle Database Release 19c are supported and certified for Oracle Exadata Database Service.

The support includes Oracle RAC, Oracle Automatic Storage Management (ASM), and Oracle Database In-Memory for on-premises deployments of SAP NetWeaver.

Note: Versions earlier than Oracle Database Release 19c are not supported with SAP for Oracle Exadata Database Service.

Regions and Availability Domains

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, compute nodes that are deployed as SAP 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 are not supported because of network latency.

Oracle Cloud Infrastructure Services

The following OCI services are relevant to deploying Oracle databases of SAP NetWeaver Application Server ABAP/Java on the Oracle Exadata Database Service.

Identity and Access Management (IAM)

The [Identity and Access Management](#) (IAM) service is relevant for deploying Oracle databases of SAP NetWeaver Application Server ABAP/Java on Exadata infrastructure-based VM clusters. IAM lets you control who has access to your cloud resources and what type of access they have. You can manage complex organizations and rules with logical groups of users and resources, and defining policies is straightforward. IAM helps you set up administrators, users, and groups, and specify their permissions. It lets you use a single model for authentication and authorization to securely control access and easily manage IT resources across all of OCI.

Networking

[OCI Networking](#) helps you set up virtual versions of traditional network components. You can extend your IT infrastructure with highly customizable VCNs and connectivity services that provide predictable and consistent performance, isolation, and availability.

A VCN is a customizable and private network in OCI. Just like a traditional data center network, the VCN provides you with complete control over your network environment. You can assign your own private IP address space, create subnets and route tables, and configure security lists (firewall rules). A single tenancy can have multiple VCNs, thereby providing grouping and isolation of related resources.

Security can be configured at several levels within a VCN. A subnet can be designated as public or private. A private subnet cannot have a public IP address. Security lists or the assignment of network security groups can control packet-level traffic into and out of a subnet or an instance. Customers can decide whether they prefer security lists or network security groups. Also, at the instance level, firewall rules can be implemented. Gateways and route tables provide control over traffic flow between the VCN and outside destinations. Finally, IAM policies provide control over who can access and configure which resources.

For naming, each subnet can resolve names to the internet or within a VCN. In addition, an on-premises DNS server can be added to the search scope. A description of the choices for using DNS in your VCN is provided in the [documentation](#).

OCI FastConnect is an alternative to using the public internet to connect your network to OCI. FastConnect provides an easy, elastic, and economical way to create a dedicated and private connection with higher bandwidth options. It also provides a more reliable and consistent networking experience compared to internet-based connections.

OCI's flat and fast network provides the latency and throughput of rack adjacency across the whole network, which allows synchronous replication and constant uptime. No network oversubscription also provides predictable bandwidth and performance. High-bandwidth network interconnect provides less than 100- μ s latency between any two hosts within an availability domain, and less than 1 ms between availability domains in a region.

Object Storage

[OCI Object Storage](#) helps you manage data as objects stored in containers, or *buckets*. Object Storage offers an unlimited amount of capacity, automatically replicating and healing data across multiple fault domains for high durability and data integrity. You can enhance the scale and performance of content-rich, analytic, and backup applications to serve more customers and achieve results faster.

Virtualization and Databases with Exadata Cloud Infrastructure X11M

Exadata Cloud Infrastructure X11M 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 spans 2 or more DB servers of the underlying Exadata Cloud Infrastructure X11M. In a VM cluster node (KVM guest), 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 X11M:

- Dependencies exist between the expansion stage of an Exadata Cloud Infrastructure X11M system and the maximum number of VM clusters available for SAP production environments.
 - 2 DB servers (fully equipped) support up to 5 VM clusters
 - 4 or more DB servers support 8 VM clusters
- 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 X11M based system is required for running an SAP NetWeaver landscape.
- When you use an Exadata Cloud Infrastructure system for SAP and non-SAP databases, you must use separate VM clusters for the SAP and non-SAP databases, because of specific configuration, administration, and maintenance aspects of SAP databases. Deploying SAP and non-SAP databases within the same VM cluster is not supported.
- Databases for SAP NetWeaver systems such as ERP, SCM, CRM, and BW must be separated from databases for other SAP products such as Business Objects by using different VM clusters for those different SAP solutions, because of different configuration, administration, and maintenance requirements. Deploying multiple databases of different SAP NetWeaver systems within the same VM cluster is supported.

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 ECPUs allocated for each VM cluster node.
- Increase or decrease the amount of memory (RAM) allocated for each VM cluster node.
- Expand or shrink the local file system /u02 allocated for each VM cluster node.
- Expand local file systems /, /var, /home, /tmp, /u01, /var/log, and /var/log/audit allocated for each VM cluster node.

Note: Expanding local file systems may reduce the maximum number of VM clusters.

- 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 after scaling ECPUs 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 with Exadata Cloud Infrastructure X11M

All options and features certified for on-premises deployments of SAP NetWeaver on Oracle Database Release 19c are supported and certified for Exadata Cloud Infrastructure X11M. The support includes Oracle Real Application Clusters (RAC), Oracle Automatic Storage Management (ASM), and Oracle Database In-Memory for on-premises deployments of SAP NetWeaver.

Use only RAC databases with SAP NetWeaver. Configuration, deployment, and maintenance of an SAP NetWeaver system is tailored to RAC databases. All tests and QA actions are performed with RAC databases.

Note: Oracle Database Release 19c is the only version of Oracle Database supported with SAP on Exadata Cloud Infrastructure X11M. Oracle Autonomous Database is not certified and supported with SAP.

Exadata Cloud Infrastructure X11M

Each Exadata Cloud Infrastructure X11M system configuration contains two or more DB servers (KVM-based virtualization hosts) and Exadata storage servers. These servers are connected by a high-speed, low-latency, 200-Gbit aggregated bandwidth of the active-active 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 X11M.

Currently, Oracle offers the following configurations for Exadata Cloud Infrastructure X11M.

Table 4: Exadata Cloud Infrastructure X11M Configurations

SYSTEM	NUMBER OF COMPUTE NODES	NUMBER OF EXADATA STORAGE SERVERS
X11M Elastic	2 to 32 nodes (16 to 24,320 ECPUs and total 2,780 to 44,480 GB RAM)	3-64

Note: At the time this document was published, SAP might not have certified all the preceding systems or expansion stages. For up-to-date information about the certification status of all Exadata Cloud Infrastructure systems and expansion stages, see SAP Note [2956661](#).

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 eight ECPU per VM per compute node running an SAP NetWeaver-based database workload is required.

Each VM cluster node (KVM guest) is running on a DB server (KVM-based virtualization host). You have root privileges for the VM cluster nodes (KVM guest) 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 applications. The number of production databases for SAP 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 administered 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 Exadata Database Service for SAP NetWeaver Application Server ABAP/Java

Use the information in this section to plan for using Exadata Database Service for SAP NetWeaver Application Server ABAP/Java deployment on OCI.

Prerequisites

Request and create an Exadata Cloud Infrastructure in the region and availability domain of your choice.

Deployment Restrictions with SAP for Oracle Exadata Database Service

The following *important* restrictions apply:

- Only Oracle Database Release 19c is supported.
- Oracle Autonomous Database is not certified with SAP and is not supported.
- Oracle Multitenant is not certified with SAP and is not supported.
- Non-Unicode deployments of SAP NetWeaver Application Server ABAP/Java are not supported.
- SAP application servers on VM cluster nodes are not supported.
- Although SAP application servers on VM cluster nodes are not supported, SAP central services for SAP NetWeaver Application Server ABAP/Java can be deployed on them.
- 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 38 or later for Oracle Database 19c and TDE support. For all other databases, encryption *must* be enabled as a separate step after successful SAP installation.
- The strong password policy on VM cluster nodes must be changed for an SAP installation and then reverted after the installation is finished.
- Hostnames, including hostnames for virtual IP, SCAN, and backup network, 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.

Support

If you need technical support or help with Exadata Database Service, you can go to [My Oracle Support](#) and file a service request. If you encounter any problem with the SAP NetWeaver Application Server ABAP/Java deployment with Exadata Database Service, log a support message with SAP support and assign it to the support queue BC-DB-ORA.

Customers must purchase Exadata Database Service directly from Oracle to use the service and get support for it. For details, see [Oracle Cloud Hosting and Delivery Policies](#).

In addition to support for technical issues, use [My Oracle Support](#) if you need to perform the following tasks:

- Reset the password or unlock the account for the tenancy administrator.
- Add or change a tenancy administrator.
- Request a service limit increase.

Note: SAP Note [2520061](#) describes the support prerequisites that are needed to run SAP NetWeaver Application Server ABAP/Java with Exadata Cloud Infrastructure with Oracle Linux.

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: OCI, Oracle Linux, Oracle Database, and SAP NetWeaver Application Server ABAP/Java.

Note: SAP Note [2614028](#) defines all the technical prerequisites for deploying an SAP NetWeaver Application Server ABAP/Java system with Exadata Database Service. This note is updated regularly, so read it before you start any deployment. Information in the note takes precedence over information in this paper.

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, however, that 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 types by SAP. Check SAP Note [2614028](#) for the most current certification status.

Table 5: Certified Exadata Database Service Instance Types

INSTANCE TYPE	ECPUs	MEMORY	TOTAL USABLE STORAGE CAPACITY	NUMBER OF EXADATA STORAGE SERVERS
Exadata.X11M (Elastic, up to 4 DB Servers)	16–3,040	2,780–5,560 GB	240–5,120 TB	3–64

Note: To run an SAP NetWeaver workload, the minimum numbers of ECPUs 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.

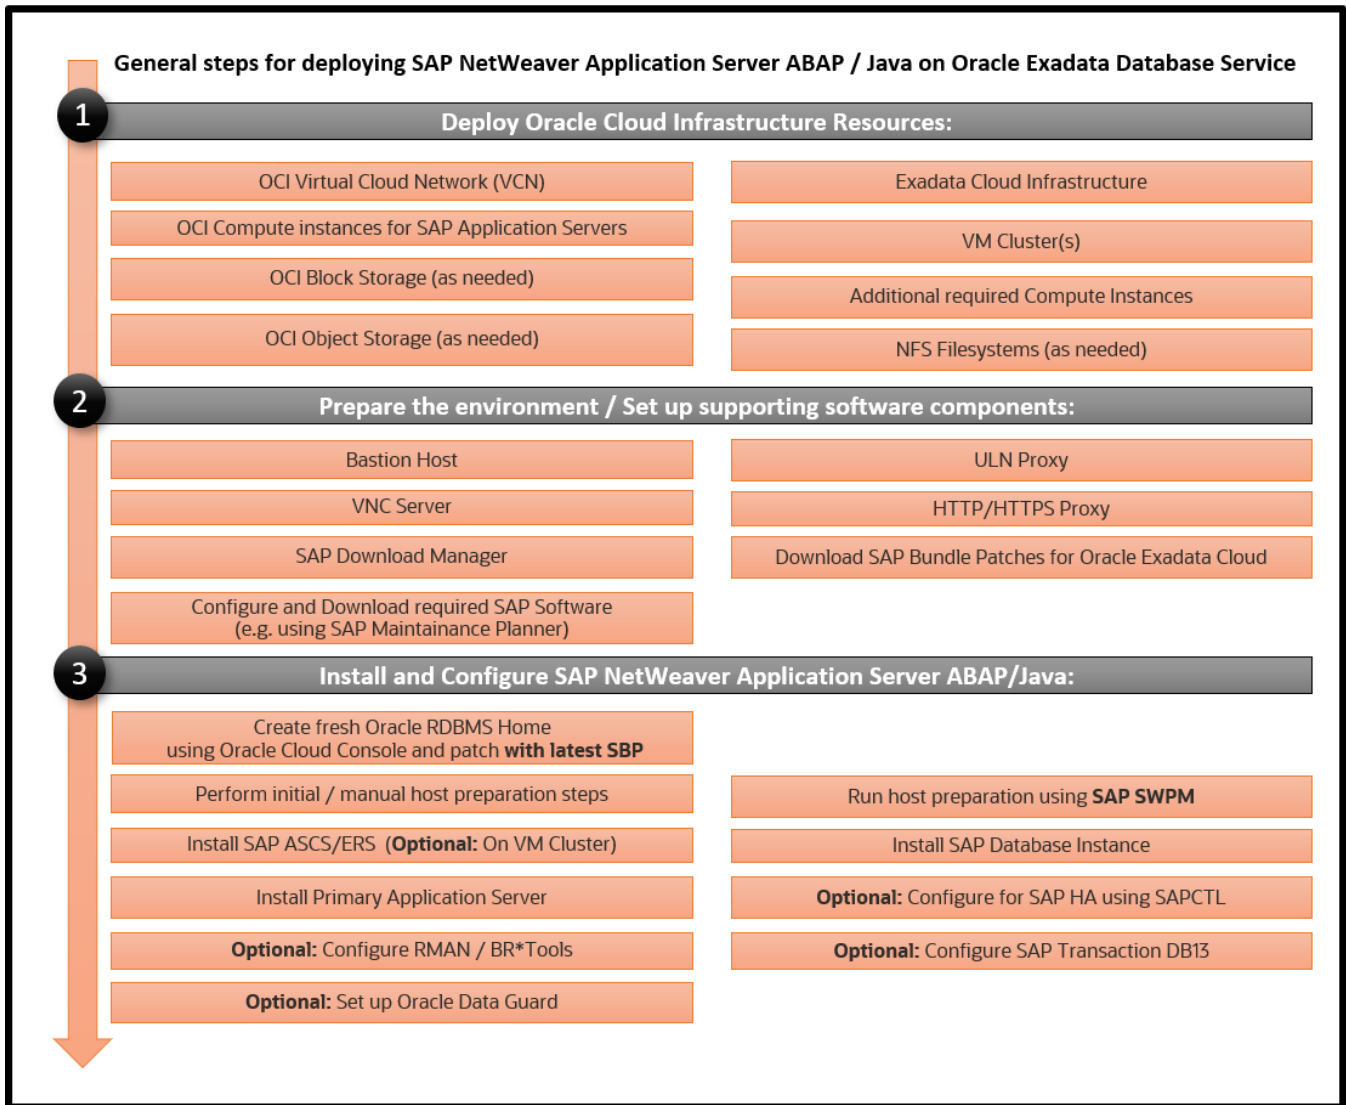


Figure 3: General Steps for Deploying NetWeaver Application Server ABAP/Java with Exadata Database Service

Use the information in this section to plan your SAP NetWeaver Application Server ABAP/Java deployment on Oracle Exadata Cloud Infrastructure X11M.

This document designs a minimal SAP landscape that consists of an SAP database using a virtual Oracle RAC database cluster (also referred as a VM cluster within this document) running on Exadata Cloud Infrastructure X11M that is connected to a single SAP application server (primary application server, or PAS) on a separate OCI host.

A real-world SAP landscape is considered fairly more complex and might consist of multiple SAP 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 cloud network (VCN) with its subnets, security (security lists or network security groups), and routing also requires detailed planning for how customer requirements can be addressed within OCI.

In OCI, you can use a quick network configuration wizard to set up a VCN and associated resources, or you can configure the VCN manually. Using the wizard simplifies the setup of the internet gateway or VPN and NAT, but numerous manual configuration steps are still required to finalize VCN configuration, such as adding more subnets,

security, and routing. Depending on your requirements (for example for SAP HA), you might need to consider additional resources.

The landscape described in this document uses a manually configured VCN, to help customers deploy SAP on Oracle Exadata Database Service in a real-world scenario.

Consider using [network security groups](#) instead of security lists.

Exadata Database Service Requirements

- Sign-in credentials to the [Oracle Cloud Console](#)

The Console is used to manage OCI resources. Your account must be enabled to create resources.

- SSH key pairs

You need SSH key pairs to deploy and access OCI resources, VM cluster nodes, and at least one certified OCI 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 Oracle Exadata Database Service: Lifecycle Management for SAP Databases” section later in this paper.

- 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

General Installation Overview

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

The first step is deploying resources in the [Oracle Cloud Console](#). In this use case, you use the Console to create an Exadata Cloud Infrastructure, a VCN that suits your needs, and a VM cluster. Additionally, you will likely deploy SAP application server hosts (virtual machine or bare metal compute instances), and, if necessary, create Object Storage buckets, network file systems (NFS), or block volumes.

When this step is completed, you have the following items:

- At least one VCN that meets the requirements for deploying a VM cluster, plus route tables and security lists
- IP addresses (node, virtual, and SCAN)
- Hostnames and SCAN names
- IP addresses of all compute nodes, for example, for the SAP application servers
- SSH key-based access to all the deployed compute resources

On the VM cluster nodes, you must make changes to local and shared file systems. The tasks to be done vary among the root, oracle, and opc user accounts. On the OCI compute nodes, 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 database installation with SWPM, you must also perform some configuration steps to be able to run SAP SWPM host preparation.

There are two types of host preparation:

- The first host preparation (or manual host preparation) 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.
- The second host preparation (or SWPM host preparation) is a function 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.

Because SWPM must be run on each VM cluster node of the VM cluster, we recommend using shared media to simplify installation.

The next major steps are as follows:

1. 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.
 - A. Manually prepare all VM cluster nodes for running SWPM.
 - B. Run SWPM host preparation on all VM cluster nodes.
 - C. (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. These steps are described in later sections of this document and include, for example, the configuration of the BR*Tools, the integration of SAP ASCS and ERS with SAPCTL, and the configuration of HA-NFS for /sapmnt.
 - D. Create an Oracle RAC database on the VM cluster by using SWPM or by migrating an existing database.
 - E. (Optional) To implement SAP HA by using Oracle Grid Infrastructure with SAPCTL, install SAP ERS on each VM cluster node of the VM cluster.
 - F. Verify that the database is Oracle RAC.
2. Install the Primary Application Server (PAS) instance on an SAP application server compute node.
 - G. Prepare the host for running SWPM.
 - H. Run SWPM to install PAS.
3. Configure the SAP GUI to verify that the SAP system is accessible.
4. Finalize the installation.

Implementing the Deployment

This section provides the steps for implementing your planned deployment of SAP NetWeaver Application Server ABAP/Java on Exadata Database Service by using the Oracle Cloud Console. You can also use automation, which might provide the advantage of repeatability, while the Console provides immediate provisioning and a human-friendly user interface. Using automation to implement the deployment is not part of this technical brief. Refer to SAP Note [3248703](#) for more information about automation.

Get Your OCI Account

To get your OCI account, work with your Oracle account team.

Deploy the OCI Resources

Deploy the following resources on OCI:

- VCN
- Exadata Cloud Infrastructure
- VM Cluster
- (Optional) Application specific virtual IP addresses and hostnames
- OCI Compute Instance

Deploy the VCN

1. Sign in to the Oracle Cloud Console.
2. Select the region in which you want to deploy your SAP application and provision the related resources. This region needs to be the same as the region in which you deploy Exadata Cloud Infrastructure (next section).
3. Open the navigation menu, go to **Networking**, and then select **Virtual Cloud Networks**.
4. Select **Create Virtual Cloud Network** to create an initial VCN of the necessary size (for example, 10.0.0.0/16). For detailed instructions, see [To create a VCN and Network Setup for Exadata Cloud Infrastructure Instances](#).
5. Add private regional subnets of the appropriate size for “client” access (SAP application servers) and a “backup” of the appropriate size (for example, 10.0.0.0/24 and 10.0.1.0/24). They are required to launch a VM cluster.
6. Configure network security to allow ICMP and TCP between all IP addresses on the same subnet for the client and backup subnets.
7. Add a service gateway to the VCN that allows access to the Oracle Services Network. Because Oracle services are region specific, add a route rule that targets the service gateway for “All Services in Oracle Services Network” in your specific region.
8. Add a public regional subnet (for example, 10.0.2.0/24), an internet gateway, and an associated security list for the bastion host to allow access from the internet.

Subsequent resources (for example, databases and compute instances) need to be deployed in that particular availability domain and client subnet. For more information, see [Overview of Networking](#) in the OCI documentation.

Deploy Exadata Cloud Infrastructure

For detailed instructions, see the [documentation](#).

1. Sign in to the Oracle Cloud Console.
2. Select the region in which you deployed the VCN.
3. Open the navigation menu, go to **Oracle Database**, select **Oracle Exadata Database Service on Dedicated Infrastructure**, and then select **Exadata infrastructure**.
4. Select **Create Exadata infrastructure** and enter the required values. For the Exadata Cloud Infrastructure model, select **X11M**.
5. Select **Create Exadata infrastructure**.

Deploy the VM Cluster

For detailed instructions, see the [documentation](#).

1. Sign in to the Oracle Cloud Console.
2. Select the region in which you deployed the VCN.
3. Open the navigation menu, go to **Oracle Database**, select **Oracle Exadata Database Service on Dedicated Infrastructure**, and then select **Exadata VM clusters**.
4. Select **Create Exadata VM cluster**, and then perform the following actions:
 - A. Select the compartment in which to create the VM cluster.
 - B. Enter a display name for the VM cluster.
 - C. Select the Exadata infrastructure that you just created (ends with **X11M**).
 - D. For the Oracle Grid Infrastructure version, select **19c** (23ai is not supported with SAP).
 - E. Select the Exadata guest image.
 - F. Provide the initial ECPU count for a single VM. The dialog box displays the total ECPU count for the VM cluster.
 - G. Select **Show additional local file system configuration options** and change the layout of the local file system according to your requirements.
 - H. For the Exadata storage configuration, select **Allocate storage for local backups** if you want to allocate more space to the RECO disk group and less to the DATA diskgroup.
 - I. Paste your public SSH keys or upload your public key files.
5. In the **Network Information** section, select your new VCN along with the client and backup subnets. If you want to use network security groups (NSGs), assign the NSGs for the client and backup subnet. Enter a maximum of 2 characters as the hostname prefix; otherwise, the hostnames generated by OCI will exceed the maximum of 13 characters.
6. Select your license type.
7. Keep the diagnostics notifications selected.
8. Expand **Show Advanced Options** and select an appropriate time zone for your Exadata VM cluster. Choose the one that matches your SAP application servers. The time zone for a VM cluster cannot be changed in the future.
9. Select **Create Exadata VM cluster**.

The creation process might take multiple hours to complete.
10. Apply the latest Grid Infrastructure updates by using the Oracle Cloud Console.

Deploy Application-Specific Virtual IP Addresses

You can *optionally* configure application-specific virtual IP addresses if you want to configure and export an ACFS file system as highly available NFS (HA-NFS) for the mountpoint /sapmnt or configure SAP high availability (SAP HA) for ASCS and ERS with SAPCTL. These virtual IP addresses must be defined in the client subnet of the VCN either through the Oracle Cloud Console or the OCI API, and then created in Oracle Clusterware to bring them online and make them highly available.

1. Sign in to the Oracle Cloud Console.
2. Select the region in which you deployed the VCN.
3. Open the navigation menu, go to **Oracle Database**, and select **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. Select **Exadata VM clusters**, and then select the name of the VM cluster where you want to define the virtual IP addresses.
5. Under **Resources**, select **Virtual IP Addresses**, and then select **Attach Virtual IP Address**.
6. Provide the following values:
 - A. Select the client subnet.
 - B. Enter the hostname of the virtual IP address. The length of the hostname must not exceed 13 characters.
 - C. (Optional) Enter the virtual IP address that should be used or let OCI select a free one.
 - D. Select the VM cluster node where the virtual IP address should be attached first.
7. Repeat steps 5 and 6 for all virtual hostnames and virtual IP addresses that you want to configure.

In the following example, three virtual hostnames and IP addresses are configured, one for HA-NFS, one for SAP ASCS and one for SAP ERS.

Name ⓘ ▲	State	Virtual IP address	Subnet	Virtual Machines
ascsmfg	● Available	10.0.1.223	exa_client	mb-q0qzt1
ersmfg	● Available	10.0.1.190	exa_client	mb-q0qzt1
hanfs	● Available	10.0.1.123	exa_client	mb-q0qzt1

Deploy the OCI Compute Instance

Deploy an OCI Compute instance of a supported shape that will host the SAP primary application server (PAS) in the same VCN as the VM cluster nodes.

In the Oracle Cloud Console, navigate to the **Compute** page. For detailed instructions, see the [documentation](#). Supported shape types are referenced in SAP Note [2474949](#). Customers can choose between Oracle Linux 8 and Oracle Linux 9 images. This paper uses Oracle Linux 8 images.

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 Exadata Database Service compute nodes.

To summarize:

- `ssh opc@node0` takes you to the `opc` user on VM cluster node0.
- `ssh orac1e@node0` (with the same key as for the `opc` user) takes you into the Oracle user account.

As a consequence, managing SSH key pairs properly is an important measure.

SAP application server instances must be located either in the client subnet of the VM cluster or in a fully separated subnet that has a routing to the client subnet of the VM cluster. Be sure to follow these design rules:

- Separate the client subnet and backup subnet
- Separate the different SAP landscapes in different VCNs
- Separate test, QA, and production VCNs

A local firewall for each compute instance that comes from the OS, and security lists that are part of the OCI 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 paper), and only the [security lists](#) for the subnets must be managed. You can get an overview about the required ports for an SAP system from [SAP Help Ports](#).

Prepare Your Environment

You can set up all the resources by using the Oracle Cloud Console or by using automation.

- 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 for Oracle Grid Infrastructure 19c.

Set Up the Bastion Host

We recommend that you use Oracle Linux 9 on the bastion host. A bastion host typically has a network interface connected to the public internet and a second network interface connected to the client network of the VM cluster. The bastion host *can* have the following roles:

- Provide a VNC server for graphical access and SSH server from outside. 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
- If the compute nodes cannot access the internet on their own, serve as an HTTPS proxy for automatic lifecycling (automatic updates) of SAP monitoring

Note: To store your SAP installation media, SAP patches, and Oracle SAP Bundle Patches (SBPs), create a network file system in the Oracle Cloud Console in the respective availability domain. Then, mount this file system on your compute nodes. For more information, see [Overview of File Storage](#).

Set Up the ULN Proxy

To ensure that you have the latest OS updates for Oracle Linux 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 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 [Install the VNC Remote Access Server on Oracle Linux](#). 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 illustration.

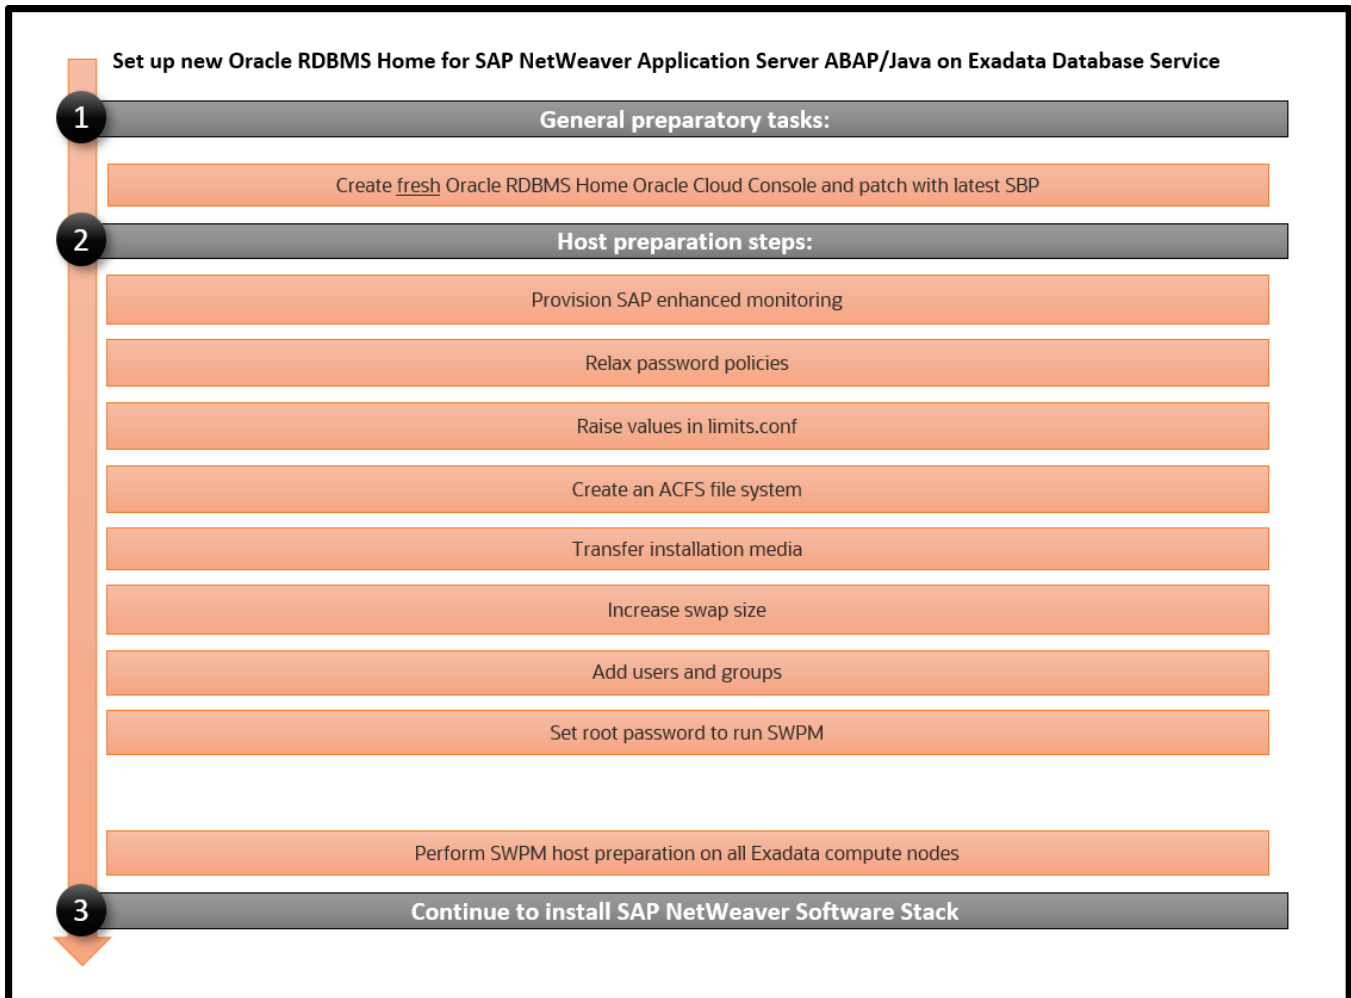


Figure 4: Steps for Setting Up an Oracle Database Home for SAP NetWeaver Application Server ABAP/Java with Oracle Exadata Database Service

Create a Fresh Oracle Database Home

1. In the Oracle Cloud Console, navigate to the VM cluster where you want to create the Oracle Database home.
2. Under **Resources**, select **Database Homes** and then select **Create Database Home**.
3. In the dialog box, select a display name and the software version for the Oracle Database home, and then select **Create Database Home**.

Do not select a database image newer than the latest SAP Bundle Patch version available. You will patch the new Oracle Database home in a later step after transferring the latest SAP Bundle Patch to a shared file system location.

Provision SAP Monitoring on Exadata VM Cluster Nodes

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

With Exadata Database Service, the SAP Host Agent needs to run on all the VM cluster nodes. Installation of 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 oraecscol. The oraecscol service must be installed and started on each VM cluster node. It is shipped as a Linux RPM called oraecscol.rpm.

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

For automatic lifecycling of oraecscol.rpm, the VM cluster 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 VM cluster nodes to the Internet:

```
[root@bastion-host] # dnf -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 oraecscol Package

The required version of oraecscol.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 oraecscol.rpm from /u02/opt/dbaas_images/oraecscol.rpm to /tmp and make /tmp your current working directory:

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

2. Install oraecscol.rpm:

```
rpm -i oraecscol.rpm
```

3. Enable the service:

```
systemctl enable oraecscol.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 VM cluster nodes have public internet access, add the following line:

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

- C. If the VM cluster nodes don't 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. As the root user, run the following command:

```
[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.

An additional step, discussed later in this document, is to complete 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 VM 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`.

Raise Values in the limits.conf File

On each VM cluster 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 will perform offline backups using SAP BR*Tools, add the following entries for each SAPSID that you are going to install. These entries are required because SAP BR*Tools will startup mount 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

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
ASMCMDS> volcreate -G DATA1 -s 1024G sapshare_v1
ASMCMDS> 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 +DATAC1 disk group and mount it to /sapshare. This change is persistent.

2. On a single VM cluster 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 VM cluster 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 Oracle Exadata Database Service: Lifecycle Management for SAP 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 VM cluster 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 VM cluster 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 MFG, run the following command:

```
[root@node0] # rm -rf /usr/sap/MFG/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).

The key steps are shown in the following illustration.

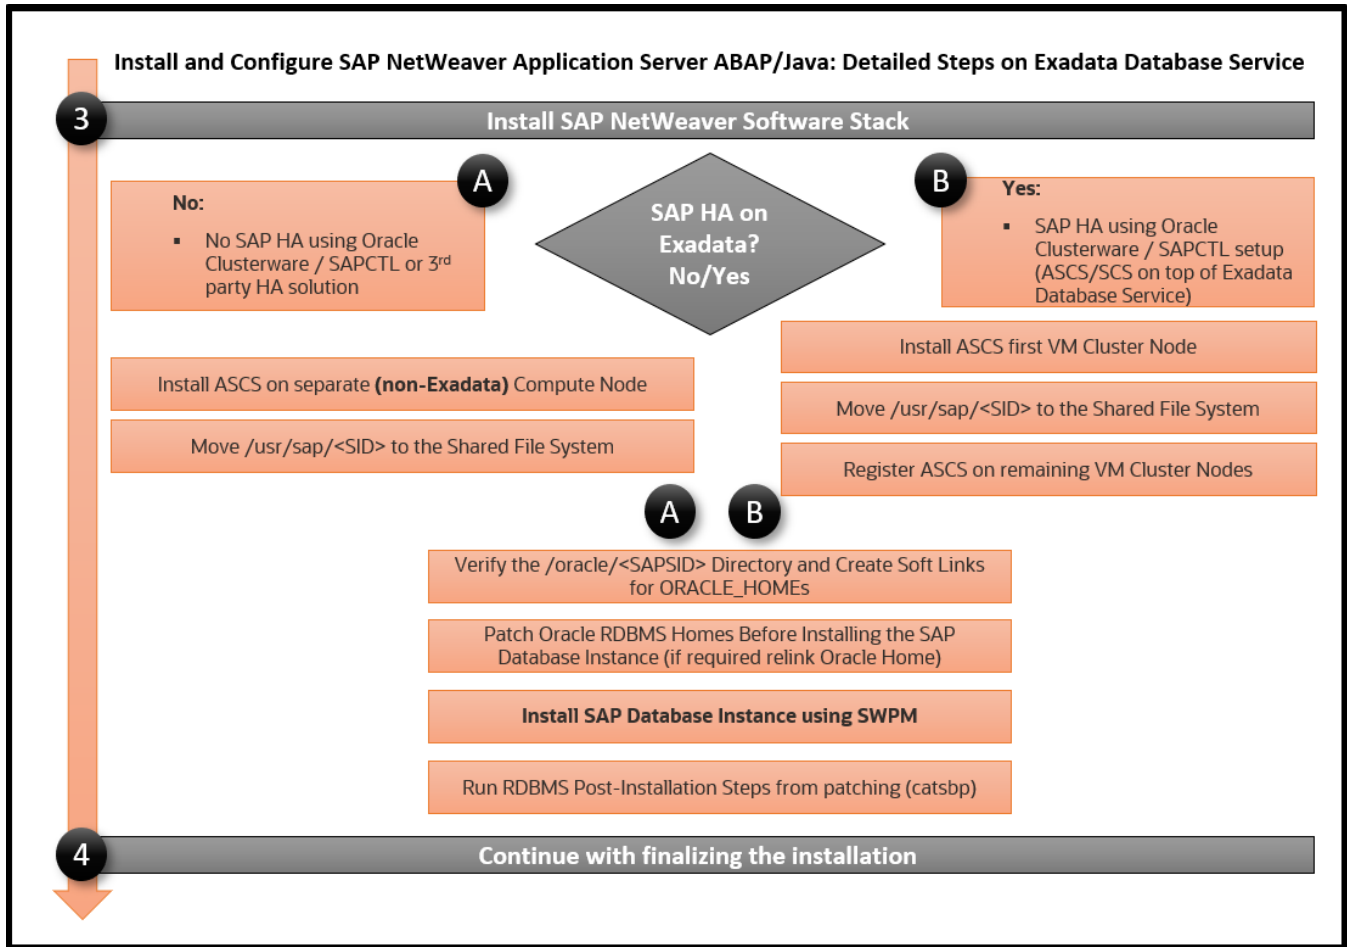


Figure 5: Steps to Install and Configure the SAP NetWeaver Software Stack

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 “Deploy Application-Specific Virtual IP Addresses” you deployed the ASCS virtual IP address (10.0.1.223). Now you need to start it.

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

```
[root@node0] # . oraenv [+ASM1]
[root@node0] # appvipcfg create -network=1 -ip=10.0.1.223 -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 appropriately:

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

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 need to 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 VM cluster node using the virtual hostname for ASCS by running `./sapinst` `SAPINST_USE_HOSTNAME=ascsmfg` 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 MFG.tar MFG
[root@node0] # cp MFG.tar /sapshare ; cd /sapshare ; tar -xvf MFG.tar
```

2. On all VM cluster nodes, run the following commands:

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

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 VM cluster 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 `ersmfg` in our example configuration.
2. After ERS installation is completed on the first node, run the following commands as `SIDADM`:

```
node0: mfgadm> sapcontrol -nr 00 -function Stop
node0: mfgadm> 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 saphostctrl locally:

```
[root@nodeN] # /usr/sap/hostctrl/exe/saphostctrl -function RegisterInstanceService -sid
MFG -nr 00 -saplocalhost ascsmfg
```

3. Start saphostctrl and the ASCS instance:

```
[root@nodeN] # su - mfgadm
nodeN: mfgadm> sapcontrol -nr 00 -function StartService MFG
nodeN: mfgadm> 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 ersmfg in our example configuration.
5. Stop the ASCS instance:

```
[root@nodeN] # su - mfgadm
nodeN: mfgadm> sapcontrol -nr 00 -function Stop
nodeN: mfgadm> 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 VM cluster 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: mfgadm> sapcontrol -nr 00 -function StartService MFG
node0: mfgadm> sapcontrol -nr 00 -function Start
node0: mfgadm> sapcontrol -nr 00 -function GetProcessList
```

At this point, 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/MFG/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/MFG/19.

Perform the following steps:

1. On each VM cluster 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 VM cluster 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:

```
MFG:/oracle/MFG/19:N
```

4. As the `oracle` user on each VM cluster 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 VM cluster 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 OCI 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 Oracle Exadata Database Service: Lifecycle Management for SAP Databases.”

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

Install the Database Instance

Obtain the SCAN listener name (short hostname) and node virtual IP addresses through the Oracle Cloud Console or by running the following commands:

```
[opc@mb-q0qzt1 ~] $ sudo su -
[root@mb-q0qzt1 ~] # . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@mb-q0qzt1 ~] # srvctl config scan | egrep "IPv4 VIP|SCAN name"
SCAN name: mb-q0qzt-scan.exaclient.sapx11m.oraclevcn.com, Network: 1
SCAN 1 IPv4 VIP: 10.0.1.75
SCAN 2 IPv4 VIP: 10.0.1.137
SCAN 3 IPv4 VIP: 10.0.1.200
[root@mb-q0qzt1 ~] # srvctl config vip -node `hostname -s` | egrep "VIP Name|IPv4 Address"
VIP Name: mb-q0qzt1-vip.exaclient.sapx11m.oraclevcn.com
VIP IPv4 Address: 10.0.1.201
```

The node virtual IP addresses are required during the next run of SWPM, in which the database will be created and loaded.

Adjust SAPDBHOST

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

```
SAPDBHOST = mb-q0qzt1
j2ee/dbtype = ora
j2ee/dbname = MFG
j2ee/dbhost = mb-q0qzt1
```

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, MFG is used as the SAPSID. Adjust the SAPSID according to your environment.

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

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: mfgadm> sapcontrol -nr 00 -function GetProcessList
```

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

```
node0: mfgadm> sapcontrol -nr 00 -function StartService MFG
node0: mfgadm> sapcontrol -nr 00 -function Start
node0: mfgadm> 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.

2. For disk groups, choose +DATAC<n> and +RECOC<n>.
3. At the Oracle RAC Parameters screen, Init.ora RAC parameters, adjust the following values (MFG is used an example; adjust accordingly):

```
MFG001.local_listener = node0-vip:<port>
MFG002.local_listener = node1-vip:<port>
```

Also perform this step for any additional VM cluster 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 VM cluster 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] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@nodeN ] $ export ORACLE_SID=MFG001

[oracle@nodeN ] $ srvctl stop database -db MFG -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 Optionally SAP BR*Tools Directories to a Shared File System Location

Because space in the local file systems of VM cluster nodes is limited, it is necessary to 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/MFG/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 MFG -stopoption immediate
```

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

```
[root@nodeN] # mkdir -p /sapshare/MFG/saptrace
[root@nodeN] # mkdir -p /sapshare/MFG/saparch
[root@nodeN] # mkdir -p /sapshare/MFG/sapreorg
[root@nodeN] # mkdir -p /sapshare/MFG/sapbackup
[root@nodeN] # mkdir -p /sapshare/MFG/sapcheck
[root@nodeN] # cd /sapshare/MFG
[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/MFG
[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/MFG
[oracle@nodeN] $ ln -s /sapshare/MFG/saptrace saptrace
[oracle@nodeN] $ ln -s /sapshare/MFG/saparch saparch
[oracle@nodeN] $ ln -s /sapshare/MFG/sapreorg sapreorg
[oracle@nodeN] $ ln -s /sapshare/MFG/sapbackup sapbackup
[oracle@nodeN] $ ln -s /sapshare/MFG/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 VM cluster 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:

- [2470718 - Oracle Database Parameter 12.2 / 18c / 19c](#)
- [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 only to mount the directories as shown in the next section.

Follow the instructions in [My Oracle Support Note 1934030.1 - ODA \(Oracle Database Appliance\): HowTo export ACFS \(cloudfs\) using HANFS](#).

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

3. Check the status, location, and configuration of the HA-NFS virtual IP address by running the following commands as root:

```
[root@node0 MFG] # srvctl status exportfs -id hanfsvip_id
export file system hanfs is enabled
export file system hanfs is exported on node node0
[root@nodeN] # srvctl config havip
```

Configure RMAN and SAP BR*Tools to Perform Backups

This section describes how to 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 customers are free to choose a different way to perform backups.

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

OCI 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).

Configure Object Storage

Step-by-step instructions for *non-SAP databases* are at [Oracle Database Backup Cloud Service](#), 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 VM cluster node.

This link has the following effects:

- If your `cwallet.sso` is lost and you can't restore it for any reason, re-create it with the API key.
- If you lose your API key, get a new one and re-create the `cwallet.sso`.
- If you lose both your `cwallet.sso` and your API key, 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 Oracle Exadata Database Service, you can consume 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 process:

1. Copy `oci_installer.zip` to `/tmp` and extract the `oci_install.jar` file.
2. Copy `oci_install.jar` to all VM cluster nodes as the `oracle` OS user.
3. Use the [Oracle Cloud Console](#) to obtain an API key and store it 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. Ensure that your NAT gateway and associated route rule are configured properly. Otherwise, the Oracle Cloud Backup Module Installer (`oci_install.jar`) might not be able to access the Oracle Cloud Storage to download the backup module.

Note: The latest version of the Oracle Database Cloud Backup Module is shipped with the latest SAP Bundle Patch and can be copied from `$ORACLE_HOME/lib/oci_installer.zip`.

Perform the following steps:

1. Install the RMAN driver components for OCI in the user home for the `oracle` OS user and install the wallet and configuration information in an SAP DBSID-aware format in `$ORACLE_HOME/dbs`. Perform this action on each VM cluster node as the `oracle` user. Obtain the OCIDs for your tenancy, user, compartment, and the name of the Object Storage bucket, from the Oracle Cloud Console.

2. Check the installation options for `oci_install.jar` (java is in your `$PATH` at user `oracle`):

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

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

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

4. 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/MFG_opc_wallet \
-configFile $ORACLE_HOME/dbs/opcMFG.ora \
-libDir ~oracle/cloud_backup/libdir
```

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

Configure RMAN

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

```
[oracle@node0] $ . oraenv
ORACLE_SID = [oracle] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@node0] $ export ORACLE_SID=MFG001
[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/19.0.0.0/dbhome_1/dbs/opcMFG.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 '+DATAC1/MFG/snapcf_MFG.f';
RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON;
```

2. 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.

3. Run a backup:

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

4. 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 VM cluster node, edit the `/oracle/<DBSID>/sapprof/init<DBSID>.sap` file (for example, `/oracle/MFG/sapprof/initMFG.sap`), and make following changes on each VM cluster node:

```
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/product/19.0.0.0/dbhome_1/dbs/opcMFG.ora)"
```

Note: The steps described here are for passwordless operation of BR*Tools. Since Oracle 12.1, 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`.

Set the following parameter in `/oracle/<DBSID>/sapprof/init<DBSID>.sap`. For example, edit `/oracle/MFG/sapprof/initMFG.sap` and append `_file_mask = 002`.

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 `initMFG001.ora` on `node0`, `initMFG002.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 = (+DATA1/MFG/spfileMFG.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 `MFG` 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/MFG/19/dbs/orapwMFG /oracle/MFG/19/dbs/orapwMFG
```

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

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

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

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

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

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

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

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

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

BR0801I BRCONNECT 7.40 (49)
BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
BR0828I Changing password for database user BRT$ADM ...

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
BR0829I Password changed successfully in database for user BRT$ADM

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
BR1525I Setting password for user BRT$ADM in secure storage
/oracle/MFG/security/rsecssfs/data/SSFS_MFG.DAT ...

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:19
BR1526I Password set successfully for user BRT$ADM in secure storage
/oracle/MFG/security/rsecssfs/data/SSFS_MFG.DAT

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:19
BR0802I BRCONNECT completed successfully
```

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

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

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

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 <NFHOST>:/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`:

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

- C. Reload the new configuration.

```
[root@nodeN] # systemctl daemon-reload
```

Configure RMAN

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

```
[root@nodeN] # su - oracle
[oracle@nodeN] $ . oraenv
ORACLE_SID = [oracle] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@nodeN] $ export ORACLE_SID=MFG001
[oracle@nodeN] $ rman target /

connected to target database: MFG (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 VM cluster node, edit the `/oracle/<DBSID>/sapprof/init<DBSID>.sap` file (for example, `/oracle/MFG/sapprof/initMFG.sap`), and make the following changes on each VM cluster nodes:

```
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. Since Oracle 12.1, 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`.

Set the following parameter in `/oracle/<DBSID>/sapprof/init<DBSID>.sap`. For example, edit `/oracle/MFG/sapprof/initMFG.sap` and append `_file_mask = 002`.

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 initMFG001.ora on node0, initMFG002.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 = (+DATA1/MFG/spfileMFG.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 MFG 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/MFG/19/dbs/orapwMFG /oracle/MFG/19/dbs/orapwMFG
```

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

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

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

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

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

9. Switch to SIDADM from root, for example, su - mfgadm.

10. On each VM cluster node, run the following command:

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

```
BR0801I BRCONNECT 7.40 (49)
BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
BR0828I Changing password for database user BRT$ADM ...
```

```
BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
BR0829I Password changed successfully in database for user BRT$ADM
```

```
BR0280I BRCONNECT time stamp: 2025-02-20 17:13:13
```

```
BR1525I Setting password for user BRT$ADM in secure storage
/oracle/MFG/security/rsecssfss/data/SSFS_MFG.DAT ...

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:19
BR1526I Password set successfully for user BRT$ADM in secure storage
/oracle/MFG/security/rsecssfss/data/SSFS_MFG.DAT

BR0280I BRCONNECT time stamp: 2025-02-20 17:13:19
BR0802I BRCONNECT completed successfully
```

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

```
nodeN: mfgadm> brspace -u // -c force
nodeN: mfgadm> brbackup -u // -q
nodeN: mfgadm> 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 3 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: mfgadm> 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 MFG with the correct SAPSID:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl create -sapsid MFG -if bondeth0 -nm 255.255.255.0 -net
10.0.1.0 -nodes node0,node1 -abapenq ASCS00 -abapvip 10.0.1.223 -abapmsport 3900 -abaprep ERS01
-aersvip 10.0.1.190 -nx 1

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

Install the Primary Application Server

The SAP primary application server (PAS) must be installed *on a separate OCI Compute shape* connected to the client subnet of your Exadata Database Service instance. 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_authtok` from the lines starting with `password requisite pam_pwhistory.so` and `password sufficient pam_unix.so`.

Edit /etc/hosts

Edit the /etc/hosts file with the IP addresses for the following items:

- VM cluster nodes, node virtual IP addresses, and SCAN names
- Virtual IP address for HA-NFS
- Short hostname for the PAS host itself

For example:

```
10.0.1.84 mfgpas.exaclient.sapx11m.oraclevcn.com mfgpas
10.0.1.226 mb-q0qzt1.exaclient.sapx11m.oraclevcn.com mb-q0qzt1
10.0.1.51 mb-q0qzt2.exaclient.sapx11m.oraclevcn.com mb-q0qzt2
10.0.1.201 mb-q0qzt1-vip.exaclient.sapx11m.oraclevcn.com mb-q0qzt1-vip
10.0.1.127 mb-q0qzt2-vip.exaclient.sapx11m.oraclevcn.com mb-q0qzt2-vip
10.0.1.75 mb-q0qzt-scan.exaclient.sapx11m.oraclevcn.com mb-q0qzt-scan
10.0.1.137 mb-q0qzt-scan.exaclient.sapx11m.oraclevcn.com mb-q0qzt-scan
10.0.1.200 mb-q0qzt-scan.exaclient.sapx11m.oraclevcn.com mb-q0qzt-scan
10.0.1.223 ascsmfg.exaclient.sapx11m.oraclevcn.com ascsmfg
10.0.1.190 ersmfg.exaclient.sapx11m.oraclevcn.com ersmfg
10.0.1.123 hanfs.exaclient.sapx11m.oraclevcn.com hanfs
```

Use Block Storage or Local NVMe Storage

Supported OCI compute shapes need some file system storage in which to store SAP software. In the smallest supported shape, only [OCI Block Volume storage](#) is available. Create one or more file systems according to the Linux documentation. /u02 is a good candidate for the next section.

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
```

Provision SAP Monitoring

Note: This step is mandatory if you deploy your PAS on an OCI Compute VM shape, for example, on VM.Standard.E5.Flex or VM.Standard.E4.Flex. Customers who are using bare metal compute shapes such as BM.Standard.E5.192 must skip this step.

For every cloud solution, SAP requires the collection of configuration and performance data for the cloud platform being used. On all compute instances associated somehow with SAP, the SAP Host Agent must be installed. This includes the compute instance where the PAS runs or the compute instances where Additional Application Servers (AAS) run.

You can install the SAP Host Agent either by using SWPM or manually, as described in the “SAP Host Agent Installation” chapter in the SAP documentation. The required version and patch level of the SAP Host Agent is 7.21 PL35 or later. For details, see SAP Note [2655715](#).

For VM compute nodes, additional monitoring components must be installed to enable SAP enhanced monitoring. The SAP Host Agent consumes configuration and performance metrics. These metrics are collected by a Linux service called `oci-sap-metrics-collector`. This service must be installed and started on each VM compute node. Part of the package is a Python script that ensures that updates of the package are applied automatically and that the service is being started if it's not running. For this script, an entry in crontab of the root OS user is created during installation of the package. Ensure that the script can access the update URL to download new packages either directly or through an HTTP proxy.

1. Download and install the latest version of the `oci-sap-metrics-collector` package.

```
cd /tmp
curl https://sapcertify.objectstorage.eu-frankfurt-1.oci.customer-oci.com/n/sapcertify/b/oci_metrics_collector_for_sap/o/oci-sap-metrics-collector-1.0-15.noarch.rpm --output oci-sap-metrics-collector-1.0-15.noarch.rpm
rpm -i oci-sap-metrics-collector-1.0-15.noarch.rpm
```

2. Modify crontab of the user root and set an HTTP proxy if required.

```
*/15 * * * * sudo https_proxy=<ip>:<port> /usr/bin/python /opt/oci-sap-metrics-collector/auto_updater.py
```

3. Wait 1-2 minutes and then verify that metrics collection works as expected. As the root user, run the following command:

```
curl http://127.0.0.1:18181
```

This command should return the XML document for consumption by the SAP Host Agent. Logs for `oci-sap-metrics-collector` are written to `/var/log/oci-sap-metrics`.

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.

1. Run the following command:

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

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

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

3. Save the changes and reload the configuration:

```
[root@pas] # systemctl daemon-reload
```

Install Additional Software Packages on OCI Compute Shapes

1. Run the following command as root:

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

2. If you will access SWPM from a browser on the same host where `sapinst` is being started or if 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. 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 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. Enable the uidd daemon and ensure that it comes up at system startup:

```
[root@pas] # 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 time zone of the VM cluster nodes (Europe/Berlin in the following example):

```
[root@pas] # timedatectl set-timezone "Europe/Berlin" ; timedatectl
```

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

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

5. Start the uidd daemon:

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

Start VNC the Server (Optional)

If you will access SWPM from a browser on the same host where `sapinst` is being started or if 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 and if it is no longer required 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 OCI Compute Shape

1. As root, change the password for yourself on each OCI compute node:

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

SWPM asks for authentication.

Current versions of SWPM run in browser mode. To access the SWPM user interface with your browser from the bastion host, you must enable the respective port reported to you in the VCN's security list in the Oracle Cloud Console. For example:

- source=0.0.0.0/0 (restrict if you can for outside access)
- source ports=any
- IP protocol=TCP
- destination port=4237 (the port from the sapinst URL)

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.

The URL for your browser is displayed.

3. Invoke an SAP recommended browser by using the following URL:

```
https://<public_IP_of_PAS>:<port>/sapinst/docs/index.html
```

4. When prompted, confirm a security exception.
5. 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/Java installations. Details are described in SAP Note [146505](#) and on the [SAP Help Portal](#).

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 the VM Cluster Nodes and the SAP 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 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, and take extra care doing so.
- Modify your VCN's security lists to close routes that were needed only for installation. For example, remove destination port 4237 used for `sapinst` HTTPS installation.

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

An Oracle 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 virtual machine environment, Oracle will also provide notification about these. You cannot opt out of updates. For more information, see [My Oracle Support Note 2333222.1](#) (Exadata Database Service Software Versions).

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 Oracle 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 paper was published):

- SAP Bundle Patch 19.27.0.0.250421 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](#).

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 OCI 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 the Oracle 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 Oracle Cloud Infrastructure 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 will most likely be 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.

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:

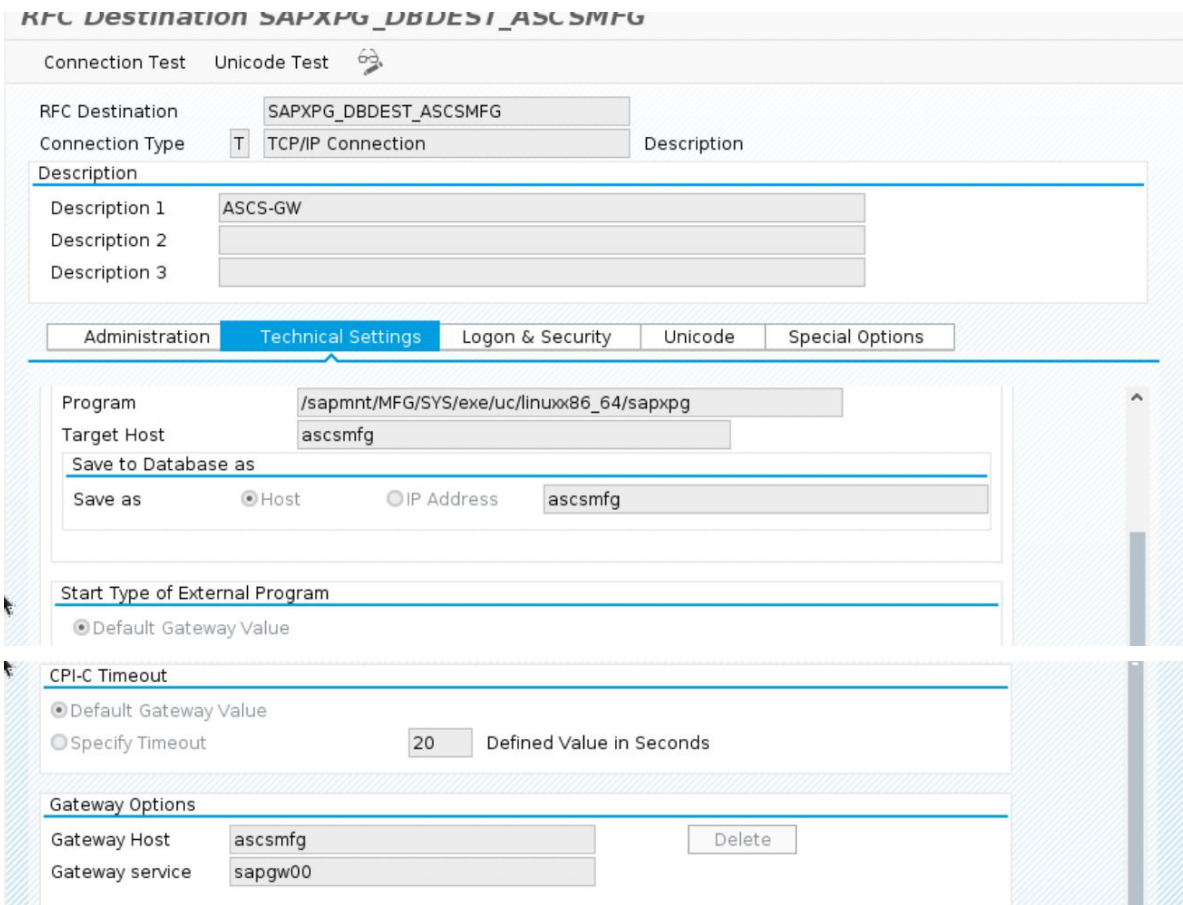
1. Complete the steps in the “Prepare for SAP BR*Tools” section.
2. Install the primary application server (PAS) and possibly additional SAP application servers.
3. Configure and test the RFC connection to the SAP Gateway. This is done in SAP transaction SM59. You must use the virtual hostname. In the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

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_ASCSMFG`.

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/MFG/SYS/exe/uc/linuxx86_64/sapxpg`.
- For **Target Host**, enter the virtual hostname used for ASCS (for example, `ascsmfg`).
- For **Gateway Host**, enter the virtual hostname used for ASCS (for example, `ascsmfg`).
- 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.



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_MFG_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](#). Basically, 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 the migrated one.

Installing Additional SAP Application Servers

The steps required to install additional SAP Application Servers (AAS) are basically the same as installing the Primary Application Server (PAS), discussed earlier in this document.

Important: For every SAP Application Server, a dedicated DB service must exist for the RAC database. Each SAP Application Server must use its dedicated DB service to connect to an RAC instance. A DB service may only be served by one RAC instance at any time.

Before you run SWPM, verify that `SAPDBHOST` in your SAP default profile (`DEFAULT.PFL`) is set to an existing and running host. If this is not the case, it is very likely that the installation of the SAP Host Agent will fail during the installation of the additional SAP Application Server, if you share your SAP profile directory across all SAP instances.

This verification is especially important for customers who use VM Cluster Node Subsetting, described in the next section. For example, if a VM cluster node (node2 in the following example) is removed from the VM cluster, it may leave `DEFAULT.PFL` with an invalid value for `SAPDBHOST`:

```
[root@node0 ~]# cat /sapmnt/MFG/profile/DEFAULT.PFL
SAPDBHOST = node2
```

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. VM cluster nodes can also 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 first whether elastic scaling of compute resources (scaling OCPUs, local storage, or memory) is an alternative.

Although VM Cluster Node Subsetting offers flexibility in the number of VM cluster nodes belonging to a VM cluster, adding or removing VM cluster 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 there are 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 application servers, which causes them to reconnect to the DB service. This action results in aborted 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 application servers, for example when they log on, so that the SAP application server whose DB service is being relocated has no users logged on and is not running any transactions.

Adding 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 Oracle Cloud Console and navigate to the VM Cluster Details page of the VM cluster.
2. Under **Resources**, select **Virtual Machines**.
3. Select **Add Virtual Machines**.
4. Select the DB servers where you want to add VM cluster nodes.

Note: By default, all available DB servers are selected. You must deselect the DB servers that you do *not* want in the VM cluster.

5. Verify that your selection is correct, and then select **Add Virtual Machines**.

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 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 VM cluster nodes and the new VM cluster node is added to all Oracle Inventories.

Check whether the Oracle Inventories on all VM cluster nodes correctly reflect the new setup.

The following example shows an example `/u01/app/oraInventory/ContentsXML/inventory.xml` file:

```
<?xml version="1.0" standalone="yes" ?>
<!-- Copyright (c) 1999, 2025, 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="node0"/>
    <NODE NAME="node1"/>
    <NODE NAME="node2"/>
  </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/MFG
[root@node2] # chown -R oracle:oinstall /oracle
[root@node2] # ln -s /u02/app/oracle/product/19.0.0.0/dbhome_1 /oracle/MFG/19
```

Adjust /etc/oratab

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

```
MFG:/oracle/MFG/19:N
```

Install SAP-Enhanced Monitoring

Perform the steps described in the section “Provision SAP Monitoring on Exadata VM Cluster Nodes.”

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
[2025-02-25 14:06:29 +0100] [INFO] [IMG-SEC-0806] Deny on login failure count set to 10
[2025-02-25 14:06:29 +0100] [INFO] [IMG-SEC-0808] Account unlock_time after {deny} failed login
attempts set to 60
[2025-02-25 14:06:29 +0100] [INFO] [IMG-SEC-0809] Password history (pam remember) set to 0
```

```
[2025-02-25 14:06:29 +0100] [INFO] [IMG-SEC-080A] Password strength set to pam_pwquality.so
minlen=6 minclass=0 dcredit=0 ucredit=0 lcredit=0 ocredit=0 difok=0 maxrepeat=0 maxclassrepeat=0
local_users_only retry=3 authtok_type=
```

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/MFG/admin /oracle/MFG/
[oracle@node2] $ scp -r oracle@node0:/oracle/MFG/oraarch /oracle/MFG/
[oracle@node2] $ scp -r oracle@node0:/oracle/MFG/security /oracle/MFG/
[oracle@node2] $ chown -R oracle:asmadmin /oracle/MFG/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 MFG with the correct DBSID for your database.

```
[oracle@node2] $ scp oracle@node0:/oracle/MFG/19/dbs/orapwMFG /oracle/MFG/19/dbs
[oracle@node2] $ ln -s /oracle/MFG/19/dbs/orapwMFG /oracle/MFG/19/dbs/orapwMFG003
```

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

Usually, you keep the following files and symbolic links: orapw<SID>, orapw<SID><INSTANCENUMBER>, init<SID>.ora, and init<SID><###>.ora. Back up the files before removing them to ensure that you can restore them if needed.

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

Run the following commands:

```
[oracle@node2 dbs] $ cp initMFG002.ora initMFG003.ora
[oracle@node2 dbs] $ cat initMFG003.ora
#Generate initsid.ora for ASM spfile
spfile = (+DATAC1/MFG/spfileMFG.ora)
```

Verify That orabase Returns the Correct Value

Run the following commands:

```
[root@node2] # su - oracle
[oracle@node2] # . oraenv
ORACLE_SID = [oracle] ? MFG
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/MFG/19 ; $ORACLE_HOME/sap/orabtt/orabtt.sh -add -
dbsid MFG
```

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/MFG/19/bin/oracle
-rwxr-x--x 1 oracle oinstall 449442792 Nov 17 10:59 /oracle/MFG/19/bin/oracle

[oracle@node2] $ chown oracle:asmadmin /oracle/MFG/19/bin/oracle
[oracle@node2] $ chmod gu+s /oracle/MFG/19/bin/oracle

[oracle@node2] $ ls -la /oracle/MFG/19/bin/oracle
-rwsr-s--x 1 oracle asmadmin 449442792 Nov 17 10:59 /oracle/MFG/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/MFG/sapprof/init<SID>.sap

Adjust the /oracle/MFG/sapprof/init<SID>.sap file so that all instances on all VM cluster nodes are included. For example, change the following lines as shown:

```
parallel_instances = (MFG001:/oracle/MFG/19@MFG001, MFG002:/oracle/MFG/19@MFG002) to
parallel_instances = (MFG001:/oracle/MFG/19@MFG001, MFG002:/oracle/MFG/19@MFG002,
MFG003:/oracle/MFG/19@MFG003)

asm_ora_home = (MFG001:/u01/app/19.0.0.0/grid, MFG002:/u01/app/19.0.0.0/grid) to
asm_ora_home = (MFG001:/u01/app/19.0.0.0/grid, MFG002:/u01/app/19.0.0.0/grid,
MFG003:/u01/app/19.0.0.0/grid)

asm_ora_sid = (MFG001:+ASM1, MFG002:+ASM2) to
asm_ora_sid = (MFG001:+ASM1, MFG002:+ASM2, MFG003:+ASM3)
```

Set the _file_mask Parameter in init<SID>.sap on All VM Cluster Nodes

Ensure that the init<SID>.sap file on all VM cluster nodes has the _file_mask parameter 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 = MFG001)
      (ORACLE_HOME = /oracle/MFG/19)
    )
    (SID_DESC =
      (SID_NAME = MFG002)
      (ORACLE_HOME = /oracle/MFG/19)
    )
    (SID_DESC =
      (SID_NAME = MFG003)
      (ORACLE_HOME = /oracle/MFG/19)
    )
  )

[root@node0] # lsnrctl reload

LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 25-FEB-2025 14:49:13

Copyright (c) 1991, 2024, Oracle. All rights reserved.

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

[root@node0] # lsnrctl status

LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 25-FEB-2025 14:49:46

Copyright (c) 1991, 2024, Oracle. All rights reserved.

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           25-FEB-2025 12:48:53
Uptime               0 days 2 hr. 0 min. 53 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/mb-q0qzt1/listener/alert/log.xml
Listening Endpoints Summary...
```

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER)))
Services Summary...
Service "MFG001" has 1 instance(s).
  Instance "MFG001", status UNKNOWN, has 1 handler(s) for this service...
Service "MFG002" has 1 instance(s).
  Instance "MFG002", status UNKNOWN, has 1 handler(s) for this service...
Service "MFG003" has 1 instance(s).
  Instance "MFG003", 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 already done on the preexisting VM cluster nodes during their initial installation, it is necessary to 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/MFG/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/MFG
[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/MFG/saptrace saptrace
[oracle@node2] $ ln -s /sapshare/MFG/saparch saparch
[oracle@node2] $ ln -s /sapshare/MFG/sapreorg sapreorg
[oracle@node2] $ ln -s /sapshare/MFG/sapbackup sapbackup
[oracle@node2] $ ln -s /sapshare/MFG/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] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@node0 ~]$ export ORACLE_SID=MFG001 ; sqlplus / as sysdba

alter system set instance_number=3 scope=spfile sid='MFG003';
alter system set instance_name='MFG003' scope=spfile sid='MFG003';
alter system set local_listener='node2vip:<port>' scope=spfile sid='MFG003';
alter system set log_archive_format='%t_%s_%r.dbf' scope=spfile sid='MFG003';
alter system set log_archive_trace=0 scope=spfile sid='MFG003';
alter system set thread=3 scope=spfile sid='MFG003';
alter system set undo_tablespace='PSAPUNDO003' scope=spfile sid='MFG003';
alter system set service_names='MFG','MFG003' scope=spfile sid='MFG003';
```

Add an Undo Tablespace

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

```
CREATE UNDO TABLESPACE PSAPUNDO003 DATAFILE '+DATA1' 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 31 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 32 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 33 '+DATAC1' size <sizeofyourredologs>;
alter database add logfile thread 3 group 34 '+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] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@node0] $ srvctl add instance -d MFG -i MFG003 -n node2
[oracle@node0] $ srvctl start instance -d MFG -i MFG003
```

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 according SAP 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] # . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@node0] # srvctl modify service -db MFG -service MFG_D10 -preferred "MFG001" -available
"MFG002,MFG003" -modifyconfig
[root@node0] # srvctl relocate service -db MFG -service mfg_d10 -oldinst MFG003 -newinst MFG001
-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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 as the root user:

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

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

2. To register the instance service, relocate the ASCS virtual IP address to the new node:

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

3. Register the ASCS instance service:

```
[root@node2] # /usr/sap/hostctrl/exe/saphostctrl -function RegisterInstanceService -sid
MFG -nr 00 -saplocalhost ascsmfg
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 saphostctrl and the ASCS instance:

```
[root@node2] # /etc/init.d/sapinit start
### START ASCS IN ORDER TO INSTALL ERS
[root@node2] # su - mfgadm
node2:mfgadm 32> sapcontrol -nr 00 -function StartService MFG

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

Start
OK
```

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

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

Stop
```

```

OK
node2:mfgadm 36> sapcontrol -nr 00 -function StopService MFG

18.03.2025 17:20:33
StopService
OK

```

7. Move the ASCS virtual IP address 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.MFG.abapvip -n node0
CRS-2673: Attempting to stop 'sap.MFG.abapvip' on 'node2'
CRS-2677: Stop of 'sap.MFG.abapvip' on 'node2' succeeded
CRS-2672: Attempting to start 'sap.MFG.abapvip' on 'node0'
CRS-2676: Start of 'sap.MFG.abapvip' on 'node0' succeeded

```

8. Set up SAP HA with SAPCTL. For detailed information, see the “Scope and Assumptions” section 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG -if bondeth0 -nm
255.255.255.0 -net 10.0.1.0 -nodes node0,node1,node2 -abapenq ASCS00 -abapvip 10.0.1.223
-abapmsport 3900 -abaprep ERS01 -aersvip 10.0.1.190 -nx 1
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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
```

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 MFG -abaprep ERS01 -to node2
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG -abapenq ASCS00 -to node2
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 Oracle. All rights
reserved
```

```
SAP ABAP Enqueue service is ONLINE on node2
SAP ABAP Replication service is ONLINE on node0
```

Set Up SAP Diagnostic Agent for SAP Solution Manager on the New VM Cluster Node

See “Set Up SAP Diagnostic Agent for SAP Solution Manager” earlier in this document for details.

Removing Nodes

This section describes the tasks associated with removing a VM cluster node from a VM cluster.

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 will be removed, then SAP HA must be reconfigured to ensure that SAP ASCS and ERS are running on VM cluster nodes that will *not* be removed.

This reconfiguration is also required if SAP ASCS or SAP ERS is not actively running on the VM cluster node being removed 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 MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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
```

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

```
[root@node0] # /usr/sap/sapctl/bin/sapctl remove -sapsid MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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:

```
[root@node0] # /usr/sap/sapctl/bin/sapctl create -sapsid MFG -if bondeth0 -nm
255.255.255.0 -net 10.0.1.0 -nodes node0,node1 -abapenq ASCS00 -abapvip 10.0.1.223 -
abapmsport 3900 -abaprep ERS01 -aersvip 10.0.1.190 -nx 1
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 Oracle.  All rights
reserved
Creating SAP ABAP VIP
Creating SAP ABAP ERS VIP
Creating SAP ABAP Enqueue resource
Creating SAP ABAP Replication resource
Done
```

- Restart the ASCS and ERS instances:

```
[root@node0]# /usr/sap/sapctl/bin/sapctl start all -sapsid MFG
sapctl version 10.0 Patch 3    Production Copyright 2017 - 2024 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. Otherwise, doing so breaks the connection between the SAP application and the database, causes ORA-600 errors in the SAP applications, and causes the creation of short dumps by the SAP application.

If you followed the SAP-specific guidelines for SAP on Oracle RAC Databases, then you configured one dedicated database (DB) service for each SAP 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 wisely across the 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, assume that DB service MFG_D10 (belonging to SAP application server MFG D10) has instance MFG003 configured as the preferred instance and instances MFG001 and MFG002 as the available instances. Because instance MFG003 will be removed, instance MFG001 is now configured as the preferred instance and instance MFG002 is configured as an available instance.

```
[root@node0] # . oraenv
ORACLE_SID = [root] ? +ASM1
The Oracle base has been set to /u01/app/grid
[root@node0] # srvctl relocate service -db MFG -service mfg_d10 -oldinst MFG003 -newinst MFG001
-force
[root@node0] # srvctl modify service -db MFG -service MFG_D10 -preferred "MFG001" -available
"MFG002" -modifyconfig
```

Check the SAP Default Profile

Ensure that `SAPDBHOST` in your SAP default profile (`DEFAULT.PFL`) is not set to the hostname of the VM cluster node that you are going to remove, and set it to a hostname of one of the remaining VM cluster nodes.

For example:

```
[root@node0 ~]# grep DBHOST /sapmnt/MFG/profile/DEFAULT.PFL
SAPDBHOST = node0
```

Remove a VM Cluster Node

1. Open Oracle Cloud Console and navigate to the VM Cluster Details page of the VM cluster.
2. Under **Resources**, select **Virtual Machines**.
3. In the list of virtual machines, select the Actions menu for the VM cluster node that you want to remove and select **Terminate**.

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

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 the following example, node 2 with redo log thread 3 and undo tablespace PSAPUNDO003 are removed.

```
[root@node0 ] # su - oracle
[oracle@node0] $ . oraenv
ORACLE_SID = [oracle] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@node0 ] $ export ORACLE_SID=MFG001
[oracle@node0 ] $ sqlplus / as sysdba
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
31	3	1	INACTIVE
32	3	1	INACTIVE
33	3	1	INACTIVE
34	3	1	INACTIVE

```
SQL> alter database disable thread 3;
SQL> alter database drop logfile group 31;
SQL> alter database drop logfile group 32;
SQL> alter database drop logfile group 33;
SQL> alter database drop logfile group 34;
```

- Drop 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 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 MFG -instance MFG003
Remove instance from the database MFG? (y/[n]) y
```

Note: Some cluster resources cannot be removed and remain in Oracle Cluster Repository.

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 the 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, 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 Downloads](#)
- [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](#)
- [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](#)

- [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"](#)
- [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](#)
- [2936683 - Oracle Linux 8: SAP Installation and Upgrade](#)
- [3119751 - Linux Requirements for SAP Kernel 754 and for SAP Kernel 788 and higher](#)
- [3399081 - Oracle Linux 9: SAP Installation and Upgrade](#)
- [3582297 - Install the latest version of Oracle Metrics Collector for SAP for VM Shapes in Oracle Cloud Infrastructure \(OCI\)](#)
- [3564601 - SAP license missing when implementing SAP High Availability on Oracle Engineered Systems or Oracle RAC Clusters configured by customer](#)
- [1174136 - Oracle: End of Support Dates](#)
- [527843 - Oracle RAC support in the SAP environment](#)

Oracle

- [Oracle Exadata Cloud Service X11M](#)
- [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](#)

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 © 2025, 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.

SAP NetWeaver® Application Server ABAP/Java with Oracle Exadata Database Service Based on Oracle Exadata Cloud Infrastructure X11M
May 2025

Author: Markus Breunig

Contributing Authors: Torsten Grambs, Jan Klokkers, Timm Seitz