Redbooks

ibm.com/redbooks

Migrating SAP Workloads to IBM Power Virtual Server

A Practical Guide for Hybrid Cloud Transformation

Henry Vo

Tim Simon

Graham Hardy

Thorsten Hesemeye

Ralf Schmidt-Dannert

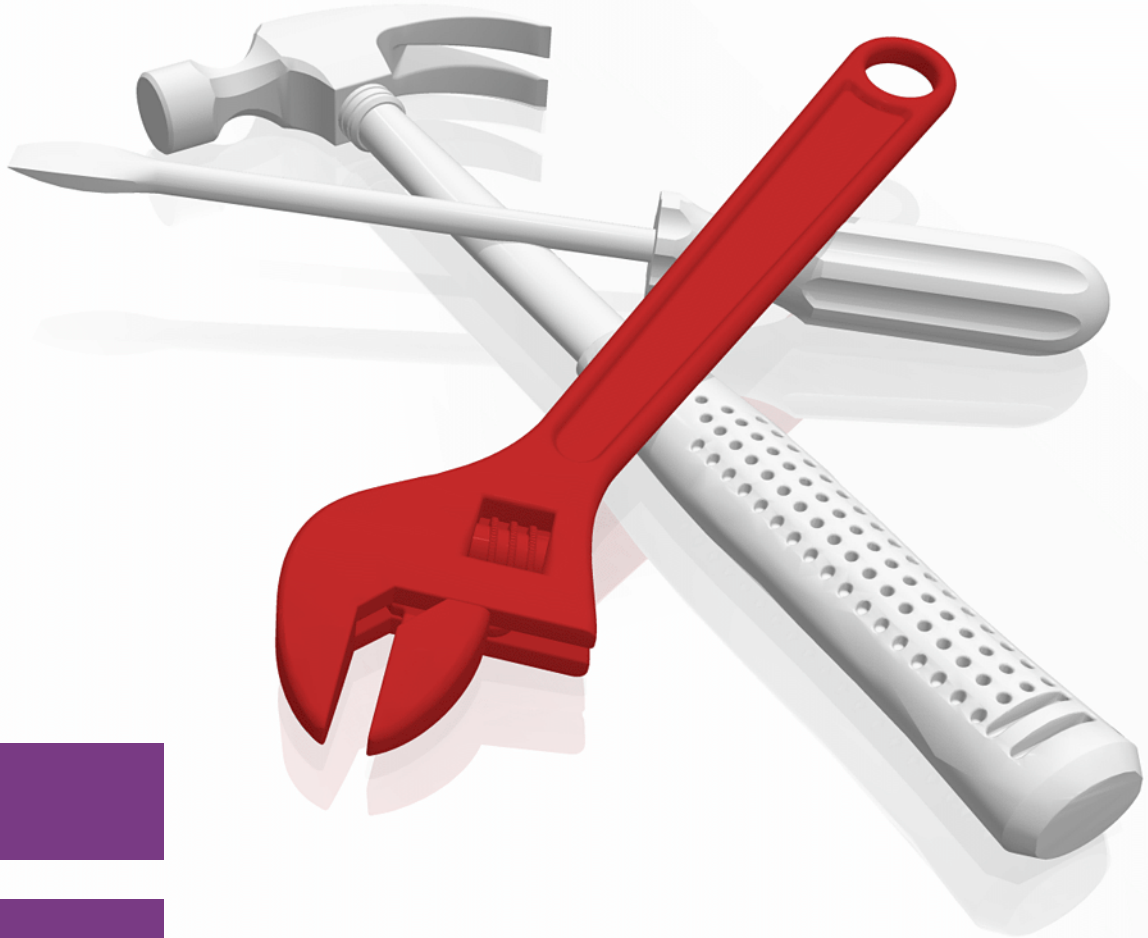

IBM Power

Infrastructure Solutions

IBM

Redbooks

IBM Redbooks

Migrating SAP Workloads to IBM Power Virtual Server: A Practical Guide for Hybrid Cloud Transformation

November 2025

SG24-8577-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (November 2025)

This edition applies to the following product versions:
- IBM Power Virtual Server
- IBM Power10 and IBM Power11
- IBM AIX 7.2 TL5 and 7.3 TL1
- IBM i 7.4, 7.5
- IBM Db2 11.5
- IBM Db2 HADR 11.5
- IBM Red Hat Enterprise Linux 8.x
- IBM HCMT 2.00.062.00.1650891137

© Copyright International Business Machines Corporation 2025. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Now you can become a published author, too! viii
Comments welcome viii
Stay connected to IBM Redbooks ix

Chapter 1. Overview for IBM Virtual Server and SAP HANA 1
1.1 IBM Power Virtual Server 2
1.1.1 Power Virtual Server and SAP 2
1.1.2 Deployment options 3
1.1.3 IBM Power Virtual Server for SAP 3
1.2 SAP on PowerVS 4
1.2.1 SAP HANA 5
1.2.2 SAP ERP solutions 6
1.3 SAP migration tools 7

Chapter 2. Migrating SAP S/4HANA to IBM Power Virtual Server 11
2.1 Checking the SAP S/4HANA pre-migration steps 12
2.1.1 Checking and confirming database health 12
2.1.2 Reorganizing the database row store 13
2.1.3 SQL mini checks 13
2.1.4 Scheduling an SAP HANA sizing report on the source system 14
2.1.5 Using the EarlyWatch Alert (EWA) reports as an early indicator 14
2.1.6 Source database credentials 15
2.2 Creating the target SAP HANA System on PowerVS 15
2.2.1 Planning the PowerVS deployment 15
2.2.2 Comparing the required CPU, cores, and storage for the target system 15
2.2.3 Creating the software repository file system and transferring packages 16
2.2.4 Unpacking the files 17
2.2.5 Installing SAP HANA on the target system 19
2.3 Migrating SAP S/4HANA by using SAP HANA system replication 24
2.3.1 Pre-checks before configuring SAP HANA system replication 24
2.3.2 SAP HANA landscape steps to activate SAP HANA system replication 25
2.3.3 Checking the database trace files 34
2.3.4 SAP HANA system replication resources 34
2.4 Performing an SAP HANA system replication takeover 35

Chapter 3. Migrating SAP ECC on Oracle to IBM Power Virtual Server 37
3.1 Preparation steps on the source system 38
3.1.1 Target audience and intent 38
3.1.2 Scope and coverage 38
3.1.3 Disclaimer 39
3.1.4 Discovery and migration option selection 39
3.1.5 Technical details and considerations 41
3.1.6 Preparation 42
3.2 Backing up the source Oracle database by using RMAN 44

© Copyright IBM Corp. 2025.

3.2.1 RMAN option 1: Offline backup and duplicate database 46
3.2.2 RMAN option 2: online backups 49
3.3 Restoring the Oracle database on the target system 57
3.3.1 RMAN restore options 57
3.3.2 Assumptions 57
3.3.3 Connectivity and preparation: Oracle on IBM PowerVS 58
3.3.4 Option 1: Creating an RMAN duplicate database 58
3.3.5 Checking the archive log status 61
3.3.6 Restarting the database in archive mode 61
3.3.7 Option 2: RMAN restore and recover database 61

Chapter 4. Migrating SAP ERP 6 on IBM Db2 to IBM Power Virtual Server 67
4.1 Overview 68
4.2 Migration option 1: Backup and restore 68
4.2.1 Preparing for migration 69
4.2.2 Shutting down the SAP application and deactivating the database 69
4.2.3 Using the IBM Db2 offline backup on the source system 71
4.2.4 Transferring the backup files 71
4.2.5 Restoring the database on the target system 72
4.2.6 Starting the SAP target system 73
4.3 Migration option 2: IBM Db2 HADR 74
4.3.1 Special case: Migrating a source IBM Db2 cluster 75
4.3.2 Preparing for migration 75
4.3.3 Ensuring that archive logging is enabled 76
4.3.4 Creating an online backup 78
4.3.5 Transferring backup files to the target system 78
4.3.6 Restoring the target system 79
4.3.7 Defining HADR local and remote service ports 80
4.3.8 Configuring HADR on the target and source system 80
4.3.9 Starting HADR and checking the synchronized data 82
4.3.10 Running the core migration 82

Related publications 85
IBM Redbooks 85
Help from IBM 85

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright IBM Corp. 2025.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at https://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM Cloud®	PowerHA®
Aspera®	IBM Cloud Pak®	Redbooks®
BLU Acceleration®	IBM FlashSystem®	Redbooks (logo) ®
Db2®	IBM Spyre™	Spyre™
IBM®	IBM watsonx®	watsonx®
IBM BLU®	IBM Z®	z/OS®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM Redbooks® publication provides a comprehensive guide for enterprises planning to migrate their SAP workloads to IBM Power Virtual Server on IBM Cloud®. It outlines the strategic benefits of adopting PowerVS as a cloud platform for SAP, including its performance, scalability, and enterprise-grade security. The book presents a structured approach to migration, covering architectural considerations, planning methodologies, and technical processes for migrating from on-premises or alternative cloud environments. Through reference architectures, best practices, and real-world examples, this guide equips IT architects, SAP administrators, and cloud engineers with the knowledge and tools that are needed to run a smooth, efficient, and future-ready migration to IBM's secure and high-performance cloud infrastructure.

Authors

This book was produced by a team of specialists from around the world working with IBM Redbooks professionals.

Henry Vo brings over 12 years of experience at IBM®, building deep expertise as a test specialist with a strong focus on technical analysis, strategic planning, and test architecture. With a background as a back-end programmer on IBM z/OS® systems, Henry has contributed to numerous enterprise-scale solutions. Today, he leads the IBM Redbooks team and guides the development of technical content that connects innovation with practical implementation for global clients.

Tim Simon is an IBM Redbooks Project Leader who is based in Tulsa, Oklahoma, US. With over 40 years of experience at IBM, primarily in technical sales, Tim built a distinguished career helping clients design and implement IBM solutions that align with their business goals. He holds a Bachelor of Science degree in Mathematics from Towson University in Maryland. Tim's expertise spans a broad range of IBM technologies, with deep experience in developing customer solutions that use IBM Power, IBM Storage, and IBM Z® platforms.

Graham Hardy is an IBM Systems Solution Architect with the IBM Technology Engineering for SAP team who is based in St. Leon-Rot, Germany. He is an SAP Certified Basis consultant and HANA Administrator with 18 years of experience in SAP solutions architecture and 10 years in multinational projects with GBS/GTS. His primary areas of expertise include SAP product migrations, SAP product testing with G2M, R&D functions for SAP banking services to SAP TRBK migrations, and cloud solution design. He currently leads the development of digital solutions and demos, including video collection demos for training purposes hosted on IBM TechZone. He is developing digital assets for the IBM MediaCenter.

Thorsten Hesemeyer is an IT Specialist working in Technical Field Enablement for SAP on IBM Power in St. Leon-Rot, Germany. He is an LPIC-3-certified Linux expert with 30 years of onsite customer experience. His expertise includes data center migrations, server virtualization, and container orchestration with Red Hat products for many IBM customers. Thorsten holds a degree in Physics from Ruhr-University Bochum.

© Copyright IBM Corp. 2025.

Ralf Schmidt-Dannert has over 33 years of experience in IT and works as an Executive IT Specialist with the IBM Advanced Technology Group, ISV on Power and Oracle team. He has focused most of his career on large database environments, designing, testing, and troubleshooting solutions for high performance, high availability, and near-zero data loss disaster recovery. He has helped customers in the financial, telecommunications, utility, retail, and manufacturing industries choose database and infrastructure technologies for databases up to a 100 terabytes. Most recently, he has evaluated and implemented technologies to provide Database as a Service for customers running databases on IBM Power servers (on-premises) or in IBM Power Virtual Server environments (off-premises). This work includes Oracle Database on IBM AIX® and open source databases on Linux on Power.

Thanks to the following people for their contributions to this project:

Hans-Juergen Zeltwanger
SAP on IBM Db2® Development, Boeblingen, BW, Germany

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise while honing your experience by using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Obtain more about the residency program, browse the residency index, and apply online at:

`ibm.com/redbooks/residencies.html`

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  `ibm.com/redbooks`

- Send your comments in an email to:

  `redbooks@us.ibm.com`

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  https://www.linkedin.com/groups/2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

  https://www.redbooks.ibm.com/rss.html

1

Overview for IBM Virtual Server and SAP HANA

This chapter provides a foundational overview of IBM Power Virtual Server (PowerVS) and SAP HANA to prepare for the technical discussions in this publication. This section reviews PowerVS features, including its cloud-native infrastructure, support for AIX, IBM i, and Linux, and alignment with enterprise-grade requirements such as high availability, security, and scalability.

This chapter explores key architectural concepts, platform capabilities, and integration points that make PowerVS a strong option for running SAP workloads in a hybrid cloud environment. It also compares the capabilities of both platforms to highlight the technical and business value of deploying SAP HANA on PowerVS. Finally, the chapter reviews performance optimization, workload consolidation, and hybrid cloud deployment models.

This chapter explains SAP HANA, its in-memory computing architecture, and the performance demands it places on the underlying infrastructure. By comparing the capabilities of both platforms, this chapter highlights the technical and business value of deploying SAP HANA on PowerVS.

The remaining chapters build on this foundation by exploring SAP migration strategies in detail. This section explains migration tools, methodologies, prerequisites, and IBM Cloud services that support a smooth and efficient transition to SAP HANA on IBM Power.

This chapter describes the following topics:

- “IBM Power Virtual Server”
- “SAP on PowerVS”
- “SAP migration tools”

© Copyright IBM Corp. 2025.

1.1 IBM Power Virtual Server

PowerVS is a cloud-based infrastructure offering that delivers the performance and reliability of IBM Power in a flexible, scalable virtual environment. It runs AIX, IBM i, and Linux workloads in a hybrid cloud setup that combines the strengths of on-premises Power servers with the agility of the cloud.

Built on the same architecture as IBM Power10 and Power11 systems, PowerVS enables modernization without disrupting existing applications by can ensure workload compatibility and seamless migration.

PowerVS supports mission-critical workloads with high availability and performance. Dynamic resource allocation, secure multitenancy, and integration with IBM Cloud services allow scaling on demand while maintaining enterprise-grade security and compliance.

Automation and orchestration tools are supported to simplify management of complex environments and accelerate DevOps practices.

PowerVS aligns well with hybrid cloud strategies. On-premises Power environments can be extended to the cloud for disaster recovery (DR), development and testing, or temporary capacity expansion. With support for SAP HANA, Oracle, and other enterprise applications, PowerVS provides a robust foundation for digital transformation. It integrates with IBM Cloud Pak® solutions and Red Hat OpenShift, enabling deployment of modern, containerized applications alongside traditional workloads.

1.1.1 Power Virtual Server and SAP

Designed for mission-critical workloads, PowerVS enable modernization of SAP environments with flexibility, speed, and confidence. PowerVS supports SAP ECC and SAP S/4HANA and is certified for use with RISE with SAP. With IBM Power11, PowerVS delivers higher performance, availability, and AI-readiness for SAP transformation initiatives.

RISE with SAP on PowerVS combines the flexibility of SAP's cloud transformation offering with the performance and reliability of IBM Power infrastructure in a virtualized cloud environment. SAP workloads on AIX, IBM i, or Linux can be modernized without migrating to a different platform and benefit from the scalability, high availability, and security.

Deploying SAP S/4HANA and other SAP applications on PowerVS enables adoption of a hybrid cloud model that supports business continuity, operational efficiency, and faster innovation, and maintains compatibility with existing enterprise systems.

The following benefits apply to SAP workloads:

- Hybrid cloud flexibility

  Integrates on-premises Power with IBM Cloud without redesigning SAP landscapes.

- SAP certification

  Provides SAP-certified configurations for SAP HANA and SAP NetWeaver to can ensure compatibility and performance.

- Zero planned downtime

  Uses Live Partition Mobility and automated maintenance orchestration to eliminate planned downtime and support continuous operations.

- Enterprise-grade security

  Applies built-in encryption, Secure Boot, and support for quantum-safe cryptography to help meet regulatory requirements.

- AI-ready infrastructure

  Uses Power11 systems in PowerVS with the IBM Spyre™ Accelerator and integrates with IBM watsonx® to enable AI-driven insights and automation.

- Flexible consumption models

  Offers pay-as-you-go, committed expense, or hybrid billing models to optimize cost and resource usage.

1.1.2 Deployment options

PowerVS is available in more than 22 IBM Cloud data centers globally or can be deployed as a private cloud solution on premises. This flexibility helps meet data residency, compliance, and latency requirements and enables benefits from cloud economics.

SAP clients can take advantage of PowerVS in several strategic scenarios that support modernization and resilience. The following use cases apply:

- RISE with SAP enablement

  PowerVS acts as a certified infrastructure provider for RISE with SAP to accelerate cloud transformation.

- Disaster recovery and high availability

  Built-in DR capabilities and integrate with IBM FlashSystem® and IBM PowerHA® to support business continuity.

- Development and test environments

  Environments can be provisioned quickly on a consistent architecture for SAP development, testing, and sandbox use cases.

- Global expansion

  SAP workloads can be deployed closer to users and customers by using the global PowerVS footprint.

1.1.3 IBM Power Virtual Server for SAP

PowerVS enables critical SAP workloads to run on a platform built for performance, trusted for security, and designed for flexibility. Whether migrating to SAP S/4HANA, continuing with ECC, or adopting RISE with SAP, PowerVS gives the freedom to modernize without compromise.

The following capabilities illustrate why PowerVS is a strong choice for SAP deployments:

- Running SAP where it makes sense

  SAP workloads can be deployed in the cloud, on premises, or in a hybrid model without changing the existing architecture. The same trusted IBM Power infrastructure with cloud agility.

- Certified and ready for SAP

  SAP-certified configurations are available for SAP HANA and SAP NetWeaver. PowerVS is also a supported infrastructure provider for RISE with SAP, offering a fast path to transformation.

- Built for performance

  IBM Power11 delivers up to 25% better performance for SAP workloads compared to previous generations, supporting higher consolidation and reduced total cost of ownership (TCO).

- Security that you can count on

  Quantum-safe encryption, Secure Boot, and strong isolation protect regulated industries and sensitive data.

- No downtime or disruption

  Live Partition Mobility and automated maintenance orchestration to eliminate planned downtime and keep SAP systems available.

- AI-ready and built for long-term adaptability

  IBM watsonx and IBM Spyre enable AI into the SAP landscape without migrating to a different platform.

- Global reach with local control

  More than 22 IBM Cloud data centers are available, or deployment can occur as a private cloud on premises to meet compliance, latency, and data residency requirements.

PowerVS is a purpose-built environment for SAP that combines the best of IBM Power with the flexibility of the cloud. Whether migrating, modernizing, or scaling, PowerVS helps achieve these goals faster, more securely, and more efficiently.

Note: For more information about PowerVS, see the following publications:

- *IBM Power Virtual Server Guide for IBM AIX and Linux*, SG24-8512
- *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513

1.2 SAP on PowerVS

IBM Power supports multiple SAP deployment scenarios on PowerVS. These scenarios vary based on the underlying SAP system architecture and database platform, such as migrating to SAP S/4HANA, continuing with SAP ECC on Oracle or IBM Db2, or adopting a fully managed RISE with SAP model, PowerVS provides the flexibility and performance for mission-critical workloads.

The following examples illustrate common configurations:

- SAP HANA-based systems

  SAP S/4HANA is deployed natively on PowerVS.

- SAP ECC on Oracle

  SAP ECC systems running on Oracle are migrated from on premises environments to PowerVS (see Chapter 3, “Migrating SAP ECC on Oracle to IBM Power Virtual Server” on page 37).

- SAP ECC on IBM Db2

  SAP ECC systems running on IBM Db2 are migrated to PowerVS (see Chapter 4, “Migrating SAP ERP 6 on IBM Db2 to IBM Power Virtual Server” on page 67).

These options allow modernization at a flexible pace. SAP HANA can be adopted immediately, or existing ECC systems can be migrated to the cloud first and with a future SAP HANA conversion planned.

IBM Power servers are designed for data-intensive applications such as SAP HANA and SAP ECC. They provide high memory bandwidth, strong I/O performance, and built-in virtualization, making them suited for hybrid cloud environments.

The benefits of running SAP workloads on PowerVS include the following capabilities:

- Flexible infrastructure

  Supports both SAP HANA and traditional databases.

- Consistent platform

  Provides a unified platform for SAP applications across on-premises and cloud environments.

- High performance and availability

  Meets the performance and reliability requirements of critical business processes.

- Deployment freedom

  Runs SAP where it makes the most sense, on premises, in the cloud, or in a hybrid model.

For a business-value perspective, see How can your organization benefit from running SAP HANA on IBM Power.

Note: Chapter 3, “Migrating SAP ECC on Oracle to IBM Power Virtual Server” on page 37 describes the migration of Oracle-based systems to PowerVS and Chapter 4, “Migrating SAP ERP 6 on IBM Db2 to IBM Power Virtual Server” on page 67 explains the migration of Db2-based systems to PowerVS.

1.2.1 SAP HANA

SAP HANA is an in-memory database platform that supports real-time data processing and advanced analytics. It includes two groups of components: mandatory components and additional components.

SAP HANA requires specific components to function properly across supported platforms. The following components are mandatory:

- SAP HANA server

  Also called the SAP HANA database, it runs on certified Linux systems, including IBM Power.

- SAP Hana clients

  Provide connectivity to an SAP HANA database and run on multiple operating systems, such as IBM i and AIX.

In addition to the mandatory components, SAP HANA includes optional tools that enhance database management and data integration capabilities. The following additional components are available:

- SAP HANA Studio

  An integrated development environment (IDE) for managing SAP HANA databases. It is used to create and manage database objects, monitor system performance, run SQL queries, and develop data models by using graphical or script-based approaches.

- SAP HANA Smart Data Access

  Provides connectivity to external data through SQL, OData, JDBC, and ODBC for real-time analytics and transactional processing. It offers fast, secure, and scalable connectivity to structured and unstructured data, and integrates with external systems and tools for data exchange and advanced data modeling.

- SAP HANA XS

  A lightweight, built-in application server within the SAP HANA platform that enables development and deployment of web-based applications directly on SAP HANA. It supports data access through various interfaces such as SQL, OData, JDBC, and ODBC, and integrates with front-end tools and external systems.

  SAP HANA XS includes the following versions:

  - XS Classic

    A lightweight, built-in application server that runs within the SAP HANA database and supports server-side JavaScript (XSJS). It suits simple, monolithic applications.

  - XS Advanced

    Based on a microservices architecture inspired by Cloud Foundry. It decouples application logic from the database, supports multiple programming languages (such as Node.js and Java) and offers enhanced scalability, advanced security (OAuth2), and CI/CD support for enterprise-grade, cloud-native applications.

- SAP HANA Application Function Library (AFL)

  Provides prebuilt, high-performance algorithms for advanced analytics. These functions run natively within the SAP HANA in-memory engine to enable predictive analytics, statistical computations, and machine learning without moving data.

1.2.2 SAP ERP solutions

SAP offers two generations of enterprise resource planning (ERP) systems, SAP ERP Central Component (ECC) and SAP S/4HANA.

ECC is the traditional ERP suite that supports core business functions such as finance, logistics, and human resources. It can run on various databases, including Oracle, IBM Db2, or SAP HANA.

In contrast, SAP S/4HANA is the modern successor, built exclusively for the SAP HANA in-memory database. It offers a simplified data model, real-time analytics, and an intuitive user experience through SAP Fiori. Whereas ECC often relies on batch processing and complex data structures, S/4HANA is designed for speed and agility. It helps organizations operate more efficiently and make faster, data-driven decisions.

SAP S4/HANA

SAP S/4HANA is SAP's next-generation ERP suite, which is designed to run exclusively on the SAP HANA in-memory database. It provides a simplified data model, improved user experience through SAP Fiori, and real-time analytics.

SAP S/4HANA is tightly integrated with the SAP HANA database, which enables real-time reporting, predictive analytics, and embedded AI. It reduces data redundancy and batch processing, and creates a more responsive environment. S/4HANA can be deployed on-premises or in the cloud.

SAP ERP Central Component

SAP ERP Central Component (SAP ECC) is a foundational ERP solution that supports a wide range of core business processes, including finance, logistics, human resources, and supply chain management. As part of SAP's on-premises ERP offering, SAP ECC runs on multiple relational database management systems, including Oracle and IBM Db2.

Deploying SAP ECC on Oracle Database provides advantages through Oracle's mature, enterprise platform. Oracle offers robust performance, high scalability, and features such as Real Application Clusters (RAC), which provide high availability and distribute workloads across multiple nodes. These capabilities make Oracle a common choice for enterprises that run SAP workloads.

Historically, Oracle has been widely adopted as the database back end for SAP ECC because it integrates strongly with the SAP NetWeaver platform and efficiently manages large-scale, mission-critical workloads.

Running SAP ECC on IBM Db2 provides advantages, especially in environments that already use IBM infrastructure. Db2 is optimized for high-performance transaction processing and analytics, and it integrates well with IBM Power, which is certified for SAP workloads.

SAP ECC on Db2 benefits from features such as IBM BLU® Acceleration® for in-memory processing and advanced compression, which can reduce storage costs and improve query performance. IBM and SAP maintain a long-standing partnership, which helps ensure integration and strong support for SAP applications on Db2.

Choosing between Oracle and Db2 for SAP ECC often depends on the existing IT landscape, licensing agreements, and performance requirements. Oracle offers broad industry adoption and a rich feature set. Db2 provides strong performance on IBM hardware and can be cost-effective in certain scenarios. SAP fully supports both platforms and recognizes them as reliable foundations for mission-critical ERP operations, especially in environments that are not ready to migrate to SAP S/4HANA.

1.3 SAP migration tools

SAP provides several tools to support system migrations, especially when moving from earlier systems such as SAP ECC to SAP S/4HANA. One of the core tools is Software Provisioning Manager (SWPM), which is uses for system installation, system copy, and migration tasks.

SWPM simplifies and automates many technical steps that are involved in setting up or moving SAP systems, including database export and import, configuration, and post-processing. It supports both homogeneous and heterogeneous migrations. It integrates with tools such as the Database Migration Option (DMO) of the Software Update Manager (SUM), which combines system upgrade and database migration into a single process.

These tools help ensure a smooth, consistent, and efficient migration path, reduce downtime, and minimize risk during complex SAP landscape transitions.

SWPM handles software provisioning processes including system installation, system copy and migration, system rename, and dual-stack split[1]. SWPM is delivered with the SAP Software Logistics Toolset. It calls the R3load process at the operating system level to migrate SAP ABAP stack database objects and the Jload process to migrate SAP Java stack database objects.

SWPM plays a vital role in on-premises and hybrid cloud environments because it streamlines the setup, migration, and management of SAP systems. In on-premises deployments, use SWPM for system installations, system copies, and database migrations, which helps help ensure consistency and reduce manual effort during complex provisioning tasks.

In hybrid cloud scenarios, SWPM supports migrating SAP workloads from on-premises infrastructure to cloud platforms or between cloud regions, which enables flexible deployment models. SWMP integrates with cloud-specific tools and services to enable efficient provisioning in environments such as IBM PowerVS or SAP-certified public clouds. It serves as a key component for organizations shifting to hybrid IT landscapes and helps maintain control, compliance, and operational continuity.

Additional migration tools and methods include the following tools:

- SAP HANA Migration Tool

  Included in the SAP HANA 2.0 installation package for IBM Power, this tool migrates and converts all database objects from SAP HANA 1.0 Big Endian for Power to SAP HANA 2.0 Little Endian for IBM Power.

- Database Migration Option DMO

  Integrated into the SUM tool, which is delivered with the SAP Software Logistics Toolset. Use DMO to combine a system upgrade with a database migration in a single process.

- SAP HANA System Replication (HSR)

  An SAP HANA platform feature that replicates one source SAP HANA database to a target SAP HANA database. It is primarily used for DR and HA, but it can also be used to migrate the SAP HANA database.

Table 1-1 focuses on scenarios where the target environment is SAP HANA running on IBM Power.

Table 1-1 Tools and methods for migrating to SAP HANA on IBM Power

Source database	Target SAP HANA on IBM Power	SAP application?	Method and tool
Non-SAP HANA DB type	SAP HANA 1.0 Big Endian	Yes	SWPM (R3load and Jload), DMO
SAP HANA 1.0 Big Endian	SAP HANA 1.0 Big Endian	Yes	SWPM (R3load and Load), HSR, backup and recovery
SAP HANA 1.0 Big Endian	SAP HANA 1.0 Big Endian	No	HSR Backup and recovery
SAP HANA 1.0 Little Endian	SAP HANA 1.0 Big Endian	Yes	SWPM (R3load and Jload)

[1] SAP Software Provisioning Manager (dual-stack split case), found at: https://support.sap.com/en/tools/software-logistics-tools/software-provisioning-manager.html

Source database	Target SAP HANA on IBM Power	SAP application?	Method and tool
SAP HANA 1.0 Big Endian	SAP HANA 2.0	Yes	SWPM (R3load and Jload), SAP HANA Migration Tool, Combined SWPM and SAP HANA Migration Tool
SAP HANA 1.0 Big Endian	SAP HANA 2.0	No	SAP HANA Migration Tool
SAP HANA 1.0 Little Endian	SAP HANA 2.0	Yes	SWPM (R3load and Jload), backup and recovery, HSR
SAP HANA 1.0 Little Endian	SAP HANA 2.0	No	Backup and recovery, and HSR
SAP HANA 2.0	SAP HANA 2.0	Yes	SWPM (R3load and Jload), HSR, and backup and recovery
SAP HANA 2.0	SAP HANA 2.0	No	Backup and recovery, and HSR

Note: Endianness refers to the order in which a sequence of bytes is stored in computer memory. For more information, see SAP Note 552464.

Warning:

- For backup and recovery and HSR methods, see SAP Note 1948334. The same update path must be used for backup and recovery and for System Replication.
- For more information about backup and recovery, see SAP Note 1642148.
- For more information about homogeneous system copy, see SAP Note 1844468.
- Before installing SAP HANA 1.0 or 2.0 on IBM Power, see SAP Note 2188482.
- For an overall compatibility check, see the SAP Product Availability Matrix (PAM).

2

Migrating SAP S/4HANA to IBM Power Virtual Server

This chapter introduces the concepts and technical components for migrating an existing SAP environment to SAP HANA on IBM Power, hosted on IBM Power Virtual Server (PowerVS). It guides readers through the tools, methods, and prerequisites required for a successful migration.

This chapter describes the following topics:

- “Checking the SAP S/4HANA pre-migration steps”
- “Creating the target SAP HANA System on PowerVS”
- “Migrating SAP S/4HANA by using SAP HANA system replication”
- “Performing an SAP HANA system replication takeover”

© Copyright IBM Corp. 2025.

2.1 Checking the SAP S/4HANA pre-migration steps

Before starting any SAP HANA database migration or replication activity, ensure that the source system is operating normally and maintains data consistency. Skipping this step can cause data loss application inconsistencies, and extended downtime during migration.

These SAP-recommended checks help identify and resolve issues that can interfere with the migration or result in inaccurate sizing and performance assumptions.

When preparing for SAP HANA database migration or replication, be aware of the following potential issues:

- Interrupted or failed garbage collection
- Entries, tables, or data that is left from client deletion that inflate the reported database size
- Reported inconsistencies in database checks
- Extreme load or unload actions that lead to orphaned entries
- Hardware issues during delta merge from memory to disk
- Excessive page memory dumps (possible page corruption)
- Alerts that are reported during SAP HANA Mini Check

2.1.1 Checking and confirming database health

To avoid transferring existing issues to the target SAP HANA system, check the health of the source database. Health checks help prevent problems such as consistency or block corruption from affecting the target system. SAP HANA system replication does not correct existing issues, so resolve them before migration.

Review the following SAP Notes:

- SAP Note 2116157 - FAQ: SAP HANA Consistency Checks and Corruptions
- SAP Note 2272121 - How To: Analyzing Physical Corruptions with the SAP HANA Persistence Diagnosis Tool
- SAP Note 2380176 - FAQ: SAP HANA Database Trace

Check the database trace files. The system writes database traces to service-specific files at the operating system level. The trace files are in the following directory:

`/usr/sap/<SID>/HDB<inst>/<host>/trace/DB_<SID>/`

Access trace files by using one of the following methods:

- At the operating system level
- By selecting **SAP HANA Studio** → **Administration** → **Diagnosis Files**
- By selecting **DBACOCKPIT** → **Diagnostics** → **Diagnosis Files**

For dynamic tiering, trace files can include the following names:

- `esserver_console_<host>.<port>.<counter>.trc`
- `indexserver_saphana01.30003.024.trc`

Tip: Define the alias **cdtrace** in the <sid>Application Development Manager user environment to quickly switch to the trace directory at the OS level. Database trace files use this naming convention:

<service>_<host>.<port>.<counter>.trc

Important: Fix known issues in the source database before migration or replication to avoid data corruption or application failure. Review and implement the relevant SAP Notes and correction instructions before proceeding. IBM Cloud is not responsible for data loss or application integrity issues that result from improper migration procedures.

2.1.2 Reorganizing the database row store

This section explains how to reorganize the row store to reduce fragmentation and improve database performance.

Run a row store reorganization if the database is heavily fragmented:

- SAP HANA 2.0 SPS04 and later

  Online row store reorganization runs automatically for large row stores (allocated size >= 3.2 GB) when the utilization ratio is less than the threshold (default 60%). The system evaluates this ratio every hour in the background. See SAP Note 2789255 - Automatic Online Row Store Reorganization.

- Earlier than SAP HANA 2.0 SPS04

  Run a row store reorganization if the database is heavily fragmented. For more information, see the following SAP Notes:

  - SAP Note 1813245 - SAP HANA Row Store Reorganization
  - SAP Note 1977584 - Technical Consistency Checks for SAP HANA Databases (includes SQL statements for CATALOG, DEPENDENCY, and TABLE consistency)

These SQL statements show whether the database requires reorganization and how much space can be saved afterward. Use this check to accomplish the following goals:

- Determine whether the SAP HANA database requires reorganization.
- If reorganization is required, estimate the size after space-saving actions finish.

Important: Complete the steps in the SAP Notes that apply exactly to the SAP HANA version.

2.1.3 SQL mini checks

SQL mini checks provide a quick way to verify key database information without running full diagnostic queries.

Run the following SQL statement to display the current SAP HANA database size:

```
SELECT HOST, PORT,
TO_DECIMAL(SUM(FREE_SIZE)*100 / SUM(ALLOCATED_SIZE), 10,2) "Free Space Ratio in
%",
TO_DECIMAL(SUM(ALLOCATED_SIZE)/1048576, 10, 2) "Allocated Size in MB",
TO_DECIMAL( SUM(FREE_SIZE)/1048576, 10, 2) "Free Size in MB"
FROM M_RS_MEMORY
```

```
WHERE ( CATEGORY = 'TABLE' ) and ( ALLOCATED_SIZE > 0 )
GROUP BY HOST, PORT;
```

For more SQL statements, see SAP Note 11969700 - SQL Statement Collection for SAP HANA. To identify critical issues, run the SQL statements at the CLI by using the **hdbsql** command or in the SQL Console in SAP HANA Studio.

To identify critical technical issues, locate the SQL collection named `HANA_Configuration_MiniChecks`. This script runs several mini checks and returns `C = 'X'` if it detects a potentially critical situation. Use SAP Note 1999993 - How-To: Interpreting SAP HANA Mini Check Results to interpret the results.

Important: Complete the steps in the SAP Notes that apply exactly to the SAP HANA version.

Important: Commands from `SQLStatements_Internal.zip` can cause instabilities, such as crashes or terminations. If these commands are run, test them first on a development (DEV) or proof-of-concept (POC) system.

2.1.4 Scheduling an SAP HANA sizing report on the source system

If migrating from an on-premises system to IBM Cloud is planned, schedule and run the SAP HANA memory Sizing report (advanced correction version 17 at the time of writing). For more information, see the following SAP Notes:

- SAP Note 3338309 HANA memory Sizing report - Advanced correction 17
- SAP Note 1872170 - ABAP on HANA sizing report (S/4HANA, Suite on HANA...)

Use the most current advanced correction version and include a growth forecast for the SAP HANA database. The report recommends CPU, memory, and storage for the target PowerVS instance. Select a current SAP-certified PowerVS profile.

For more sizing and configuration guidance, review the following resources:

- SAP 2947579 - SAP HANA on IBM Power Virtual Servers
- SAP 2188482 - SAP HANA on IBM Power Systems: Supported hardware and features
- SAP Note 2363248 - SAP BW/4HANA Hardware Sizing
- SAP SD benchmark results
- Measuring in SAPS (SAP Application Performance Standard)

2.1.5 Using the EarlyWatch Alert (EWA) reports as an early indicator

If the on-premises landscape uses SAP Solution Manager, generate the EWA report for the source system. Address SQL performance indicators and urgent performance KPIs immediately. Resolve any Red or Severe items as soon as possible. Review the EWA findings for the source SAP HANA database, especially in Service summary and Alert overview.

For guidance on EarlyWatch Alert usage and best practices, review these resources:

- SAP Note 1257308 - FAQ: Using EarlyWatch Alert
- SAP Note 1958910 - EarlyWatch Alert For HANA Database
- SAP Note 1911180 - HANA EarlyWatch Alerts (EWA) Issues
- SAP Solution Manager - Help Portal Documentation

2.1.6 Source database credentials

When an SAP HANA system is added to system replication, replication overwrites the MDC user tables `SAP${sid}.USR02` on the secondary system. In SAP HANA Studio, confirm which user with SYS privileges registered the MDC on the source system.

Ensure that the database user credentials and the source system password are available before starting the migration or replication process. Valid credentials are required to register the system in SAP HANA Studio or SAP HANA Cockpit. If these credentials are forgotten and replication continues, test the secondary target by swapping the primary and secondary servers.

2.2 Creating the target SAP HANA System on PowerVS

This section explains how to create the target SAP HANA system before planning the deployment.

2.2.1 Planning the PowerVS deployment

A PowerVS workspace in the IBM Cloud account and a hybrid cloud network connection are required. For planning details, see Hybrid Cloud Network Considerations for SAP on IBM Power Virtual Server and Planning your deployment in IBM Cloud documentation.

Size the target system based on the SAP HANA sizing report on the source system and EWA findings. For details, see Sizing process for SAP Systems.

2.2.2 Comparing the required CPU, cores, and storage for the target system

Confirm that certified PowerVS profiles match (or closely match) the sizing report and EWA summary:

- IBM Power VS Certified Profiles SAP HANA
- SAP Note 2947579 - SAP HANA on IBM Power Virtual Servers

Ensure that the target server has storage capacity equal to or greater than the source system, with headroom for growth. Use the IBM System Storage Architecture and Configuration Guide for SAP HANA TDI to design storage on IBM Cloud.

Create a file system mount point for software programs and the initial SAP HANA system backup. Depending on the PowerVS architecture, an NFS mount can be created and exported to other systems in the landscape.

2.2.3 Creating the software repository file system and transferring packages

This section explains how to create the software repository file system and transfer the required SAP HANA packages to the target system.

Complete the following steps to create the software repository and download the required SAP HANA packages:

1. Create the mount point `/swrepo` and assign ownership:

   ```
   sudo mkdir /swrepo
   sudo chown $USER: /swrepo
   ```

2. Download SAP HANA software and use versions that match the source system:
   a. Go to the Software Downloads main page and complete the following steps:
      i. Select **SAP for Me** → **Software Downloads**.
      ii. Enter the SAP S user ID and password.
   b. In the Software Center catalog view, complete the following steps:
      i. Select **Support Packages & Upgrades**.
      ii. Select **By Alphabetical Index (A-Z)**.
      iii. Select **H** → **SAP HANA PLATFORM EDITION** → **SAP HANA PLATFORM EDITION 2.0** → **SAP HANA DATABASE 2.0**.
      iv. Confirm that the **Selection** box shows LINUX ON POWER LE 64BIT.
      v. Select the IMDB_SERVER20 package and download it to a notebook.
      vi. Return to SAP HANA PLATFORM EDITION 2.0 and select **SAP HANA CLIENT 2.0**.
      vii. Confirm that the **Selection** box shows **LINUX ON POWER LE 64BIT**.
      viii. Download the release installed on the source system, or one version higher if the installed version is not listed.
3. Download SAPSCAR by completing to the following steps:
   a. Go to the Software Center Catalog view.
   b. Select **Support Packages & Upgrades**.
   c. Use the search box to find SAPCAR.
   d. From the results list, select **SAPCAR 7.53 Maintenance Software Component**.
   e. Download `SAPCAR_1200-70007726.EXE` to ensure that the selection box shows LINUX ON POWER LE 64BIT, and save the file to a notebook or jump host.
4. Create the `/swrepo` directory on the target system and make sure that the user owns it:

   ```
   sudo mkdir /swrepo
   sudo chown $USER: /swrepo
   ```

5. Transfer the installation archives and SAPSCAR to `/swrepo`. Then grant executable permissions to SAPCAR to enable unpacking of .SAR archive files:

   ```
   chmod -R 755 /swrepo/SAPCAR_1200-70007726.EXE
   ```

6. (Optional) Add a convenience alias to the bash profile:

   ```
   echo "alias SAPCAR='/swrepo/SAPCAR_1200-70007726.EXE'" >>$HOME/.bash_profile
   source $HOME/.bash_profile
   SAPCAR -v
   ```

2.2.4 Unpacking the files

This section explains how to unpack the SAP HANA installation archives by using SAPCAR.

Use `SAPCAR` to unpack the archives. The following commands show how to extract the installation files:

```
SAPCAR -xvf IMDB_CLIENT20_XXX_XX-XXXXXXXX.SAR -manifest
/SAP_HANA_CLIENT/SIGNATURE.SMF
```

Example 2-1 shows the output that is generated when the `SAPSCAR` file is extracted.

Example 2-1 The SAPCAR file extraction

```
x SAP_HANA_CLIENT/SIGNATURE.SMF
SAPCAR: 98 files extracted
SAPCAR -xvf IMDB_SERVER20_XXX_XX-XXXXXXXX.SAR -manifest
/SAP_HANA_DATABASE/SIGNATURE.SMF
x /SAP_HANA_DATABASE/SIGNATURE.SMF
SAPCAR: 355 files extracted
```

The system creates `/swrepo/SAP_HANA_DATABASE` and `/swrepo/SAP_HANA_CLIENT` during extraction. These directories contain the installation files.

Verifying the target operating system and fix level

For productive systems, match the source OS and fix level to help ensure consistent installation and smooth migration. For nonproductive systems, such as a POC in IBM Power VS, a higher OS version can be acceptable.

Target OS

Use Red Hat Enterprise Linux 8 (RHEL) or SUSE Enterprise Linux.

To check the operating system version and fix level, run the following command:

```
cat /etc/os-release
```

On RHEL, run the following command:

```
cat /etc/redhat-release
```

On SUSE Enterprise Linux systems, run the following command:

```
lsb_release -a
```

Verifying the file system and mount points

Ensure that source and target systems use identical mapping for storage, logical volume management (LVM), and file systems. The target system can have larger storage capacity if required for the migration. Match mount points and user/group ownership, including user ID (UID) and group ID (GID) to avoid permission conflicts during migration. Example 2-2 shows how to use the same <SID> and <instance number> on both the source and target systems to help ensure consistency.

Example 2-2 Mount points on source and target systems

```
export SID=<SID>                          # SAP HANA System ID (uppercase)
export sid=<sid>                          # SAP HANA System ID (lowercase)
export INSTNO=<INSTNO>                    # SAP HANA Instance Number
```

```
export SiteOnPrem=<PrimarySiteName>        # HANA System Replication Site Name 1 -
Migration from On-Prem - Source
export SiteOnCloud=<secondarySiteName>     # HANA System Replication Site Name 2 -
Migration to On-Cloud - Target
export NODE1=<Hostname 1>                   # Hostname of On-Premises Server
export NODE2=<Hostname 2>                   # Hostname of IBM Cloud Power Virtual
Server Instance
```

Maintaining /etc/hosts entries

Ensure that the `/etc/hosts` file includes entries for the source system and any dependent application servers, such as SAP NetWeaver or SAP S/4HANA Front-End Server (FES). Although Domain Name System (DNS) resolution is acceptable, include the IP address, short hostname, fully qualified domain name (FQDN), and a short description. These entries help troubleshoot connectivity issues if DNS becomes unavailable.

Preparing and tuning the OS for SAP HANA

Use these SAP Notes to prepare the target OS for SAP HANA:

- SAP Note 2777782 - SAP HANA DB: Recommended OS Settings for RHEL 8
- SAP Note 2772999 - Red Hat Enterprise Linux 8.x: Installation and Configuration
- SAP Note 3018133 - Linux: Running SAP applications compiled with GCC 10.x

Complete all recommended OS settings for RHEL 8. These steps include essential performance tuning parameters that optimize SAP HANA operation on PowerVS.

Pre-checks for SAP HANA by using HCMT

Use the SAP HANA Hardware and Cloud Measurement Tools (HCMT) to measure and analyze hardware or cloud systems before deploying SAP HANA or applying for certification. For details, see SAP Note 2493172 - SAP HANA Hardware and Cloud Measurement Tools.

Running HCMT *in server-client mode*

HCMT operates in the following session roles:

- Server session

  Run the HCMT server on a jump host or similar system. Go to the installation directory and run the following commands:

  ```
  cd <hcmt_install_dir>
  sudo ./hcmt -v -S
  ```

  The expected output appears in the following example:

  “`Server started, listening on port 50000 …`**”**

- Client session

  Run the HCMT client on each target host. Go to the directory where HCMT is installed and run the following commands:

  ```
  cd <hcmt_install_dir>
  sudo ./hcmt -v -p /swrepo/HCMT/setup/config/full_executionplan.json
  ```

  Keep the Hosts field blank when prompted.

The HCMT server collects data from the HCMT clients. A jump host is typically used to run the server session. The expected output appears in Example 2-3 on page 19.

Example 2-3 Output on the server

```
hcmt-2.00.062.00.1650891137 (2022-04-25 15:12:20)
Server started, listening on port 50000 ...
```

The client run generates a file that is named `hcm_result-YYYYMMDDHHMMSS.zip` in the setup directory. Upload this file to the HCMT SAP site and review the results. If issues occur, see SAP Note 1943937 – Hardware Configuration Check Tool.

Tip: If a port issue occurs, open a second terminal session. Go to the HCMT setup directory and start a session that keeps the required port open.

Running HCMT in client mode

This section describes how to run the HCMT tool in client mode on the SAP HANA server by using the full execution plan.

On the SAP HANA server, run the client by using the full execution plan:

```
sudo ./hcmt -v -p /swrepo/HCMT/setup/config/full_executionplan.json
```

Example 2-4 shows how to run the **`hcmt`** command on a client system.

Example 2-4 Client system output

```
hcmt-2.00.062.00.1650891137 (2022-04-25 15:12:20)
Loading executionplan
LogVolume (/hana/log):
DataVolume (/hana/data):
Hosts: <`Leave Blank!!!!`>  Keep this field blank, otherwise it affects the test.
Start execution of plan
Running Test C9C9F832-854F-492D-8E7EFB4609AC435C
Note: CPU Micro Benchmark
```

Plan variant: CPU performance

This command generates `hcm result-YYYYMMDDHHMMSS.zip` in the setup directory. Upload this file to the HCMT SAP website and review the results. For the SAP HCMT Cloud Server URL, see the HCMT documentation. If issues occur, use SAP Note 1943937 – Hardware Configuration Check Tool.

2.2.5 Installing SAP HANA on the target system

This section explains how to install SAP HANA on the target system by using predefined environment variables. Use these variables shown in Example 2-5.

Example 2-5 SID variables

```
export SID=<SID>                          # SAP HANA System ID (uppercase)
export sid=<sid>                          # SAP HANA System ID (lowercase)
export INSTNO=<INSTNO>                    # SAP HANA Instance Number
```

During installation, the process pauses and prompts for confirmation. Enter Y to continue. Go to the `SAP_HANA_DATABASE/server` directory and start the lifecycle manager:

```
cd /swrepo/SAP_HANA_DATABASE/server
sudo ./hdblcm
```

The expected output appears in Example 2-6.

Example 2-6 SID output

```
SAP HANA Lifecycle Management - SAP HANA Database 2.00.061.00.1644229038
************************************************************************
```

Example 2-7 illustrates the directory scan for the required software.

Example 2-7 Software scan

```
Scanning software locations...
Detected components:
    SAP HANA Database (2.00.061.00.1644229038) in
/swrepo/HANA/SAP_HANA_DATABASE/server
    SAP HANA Database Client (2.11.20.1644165757) in
/swrepo/HANA/SAP_HANA_CLIENT/client

Do you want to specify additional components location? (y/n) [n]: `n`
```

When prompted for additional component locations, enter `N` for no.

In Example 2-8, which illustrates the Action panel, enter 1 for `install`, then press Enter to install a new system. Provide the same <SID> and <instance number> as the source system.

Example 2-8 Action panel

```
Choose an action

Index | Action                    | Description
-------------------------------------------------------------
1     | install                   | Install new system
2     | extract_components        | Extract components
3     | print_detected_components | Print detected components
4     | Exit (do nothing)         |

Enter selected action index [4]: `1`
```

The output continues with Example 2-9.

Example 2-9 SAP HANA components

```
SAP HANA Database version '2.00.061.00.1644229038' will be installed.

Select additional components for installation:

Index | Components | Description
-------------------------------------------------------------------------------
1     | all        | All components
2     | server     | No additional components
3     | client     | Install SAP HANA Database Client version 2.11.20.1644165757

Enter comma-separated list of the selected indices [3]: `1`
```

Key prompts and recommended values

Example 2-10 shows a series of defaults in the next line of the output.

Example 2-10 More outputs on the SAP HANA component panel

```
Enter Installation Path [`/hana/shared`]:
Enter Local Host Name [`Yourhostname`]:
Do you want to add hosts to the system? (y/n) [`n`]:
```

Then check the source SAP HANA database system parameters:

- Source SAP HANA SID
- Source SAP HANA Instance Number

Select the appropriate system usage (for example, `2` for `test`) for the target system, as shown in Example 2-11.

Example 2-11 SAP HANA on the target system

```
Enter SAP HANA System ID: `<Needs to match the source system>`
Enter Instance Number [00]: `<Needs to match the source system>`
Enter Local Host Worker Group [default]:

Index | System Usage | Description
-------------------------------------------------------------------------------
1     | production   | System is used in a production environment
2     | test         | System is used for testing, not production
3     | development  | System is used for development, not production
4     | custom       | System usage is neither production, test nor development

Select System Usage / Enter Index [4]: 2
```

Example 2-12 shows more default values. Accept all defaults.

Example 2-12 Continuing on the SAP HANA target system

```
Do you want to enable data and log volume encryption? [n]:
Enter Location of Data Volumes [/hana/data/<SID>]:
Enter Location of Log Volumes [/hana/log/<SID>]:
Restrict maximum memory allocation? [n]:
Apply System Size Dependent Resource Limits? (SAP Note 3014176) [y]:
```

Accept the default paths unless a change is required:

- Data volumes. `/hana/data/<SID>/`
- Log volumes. `/hana/log/<SID>/`
- Consider enabling system size-dependent resource limits (SAP Note 3014176).

Set these passwords to match the source system, as shown in Example 2-13:

- The `sapadm` password
- The `<sid>adm` password
- The System Database User `SYSTEM` password

Example 2-13 Password on the target system

```
Enter SAP Host Agent User (sapadm) Password: <Use the same password used on the
source system>
Confirm SAP Host Agent User (sapadm) Password: <Use the same password used on the
source system>
Enter System Administrator (<sid>adm) Password:  <Use the same password used on
the source system>
Confirm System Administrator (<sid>adm) Password: <Use the same password used on
the source system>
Enter System Administrator Home Directory [/usr/sap/<SID>/home]:
Enter System Administrator Login Shell [/bin/sh]:
Enter System Administrator User ID [1001]: <check that the user ID number matches
the source system>
Enter ID of User Group (sapsys) [79]:  <Check that the GUID number matches the
source system>
Enter System Database User (SYSTEM) Password: <Use the same password used on the
source system>
Confirm System Database User (SYSTEM) Password: <Use the same password used on the
source system>

Restart system after machine reboot? [n]:
Summary before execution
```

Review the summary, confirm the home directory and shell, and proceed with the installation. The installation typically ends in about 20 minutes.

Example 2-14 indicates that the SAP HANA database system is successfully installed.

Example 2-14 SAP HANA installed

```
Registering SAP HANA Database Components on Local Host...
- Deploying SAP Host Agent configurations...
Creating Component List...
SAP HANA Database System installed
Log file written to xxxxxx
```

After installation finishes successfully, verify the all services start.

To check that SAP HANA is running and to determine the version, run these commands to verify that all services started on the primary and secondary SAP HANA system:

```
sudo -i -u ${sid}adm -- HDB proc
```

To verify the SAP HANA database version, run the following command on both nodes:

```
sudo -i -u ${sid}adm -- HDB version
```

Backing up the SYSTEMDB and MDC tenant databases

Before migration, back up the system database and the MDC tenant database to help ensure data consistency. Use the **Backup & Recovery** option in SAP HANA Studio or SAP HANA Cockpit.

Add the SYSTEMDB and the MDC tenant `DB_${sid}` to SAP HANA Studio or SAP HANA Cockpit. If the landscape includes SAP HANA Cockpit, the target system can be added there instead. After both systems are added, perform an initial system backup.

Backing up SYSTEMDB

To back up the system database, use SAP HANA Studio or SAP HANA Cockpit and perform these steps:

1. Select **Backup & Recovery** → **Back up System Database**.
2. Set these options:
   - Backup type: Complete data backup
   - Destination: File
   - Backup destination: `/swrepo/backup/data/SYSTEMDB`

     (Ensure that the path exists and is writable with user ${sid}adm.)
   - Backup prefix: `COMPLETE_DATA_BACKUP_INITIAL_DDMMYYYY`
3. Select **Next**, review backup setup, and select **Finish**.
4. Ensure that the backup of `SYSTEMDB@<SID>` finishes successfully.

Backing up MDC tenant (DB_$(sid))

To back up the MDC tenant database, use SAP HANA Studio or SAP HANA Cockpit and complete the following commands:

1. Select **Backup & Recovery** → **Back up Tenant Database**.
2. Select database ${sid}, and select **Next**.
3. Set these options
   - Backup type: Complete data backup
   - Destination: File
   - Backup destination: `/swrepo/backup/data/DB_${sid}`

     (Ensure that this directory exists and is writable by user `${sid}`adm.)
   - Backup prefix: `COMPLETE_DATA_BACKUP_INITIAL_DDMMYYYY`
4. Select **Next,** review the backup setup, and click **Finish**.
5. Ensure that the backup of `DB_${sid}` finishes without errors.

Checking the backup status

To check the backup status, use SAP HANA Studio or SAP HANA Cockpit to perform the following steps:

1. Right-click the **SYSTEMDB** entry.
2. Go to the GUI and select **Backup & Recovery** → **Open Backup Console** → **Backup catalog**.
3. In Database, select `${sid}` (tenant) to confirm the backup finished successfully.
4. In Database, select `${sid} <SYSTEMDB>` to confirm the backup finished successfully.

Optional: Checking trace log files

To check the trace log files, use the operating system or SAP HANA tools. The database trace is written to service-specific files at that operating system level.

- Trace directory
  - The trace files are in the following directory:

    `/usr/sap/${sid}/HDB<inst>/<host>/trace/DB_${sid}/`
  - In the `${sid}adm user` environment, switch to the trace directory by running the following command:

    `cdtrace`
  - Alternatively, use the **`cdtrace`** alias, if defined, or go to the directory directly.
- File naming convention

  Trace files use the following formats:
  - Standard format:

    `<service>_<host>.<port>.<counter>.trc`
  - Dynamic tiering format (see SAP Note 2871785):

    `esserver_console_<host>.<port>.<counter>.trc`
- Access methods

  Access trace files by using any of the following methods:
  - OS level
  - By selecting **SAP HANA Studio** → **Administration** → **Diagnosis Files**.
  - By selecting **DBACOCKPIT** → **Diagnostics** → **Diagnosis Files**.

2.3 Migrating SAP S/4HANA by using SAP HANA system replication

SAP HANA system replication is used to migrate the SAP S/4HANA system to PowerVS. This section explains the steps and requirements for configuring replication.

2.3.1 Pre-checks before configuring SAP HANA system replication

Before configuring SAP HANA system replication, complete the following prerequisites (applies to RHEL and SUSE Enterprise Linux):

- Confirm the SAP HANA database user that is used to access the system (commonly SYSTEM or the SAP schema owner if the SAP Basis team applied security advisories).
- Ensure that the target SAP HANA version is equal to or greater than the source (identical for Active or Active read-enabled).
- Align system configuration and `ini` settings across both systems.
- In scale-out configurations, match the number of worker nodes and roles.
- Use the same <SID> and instance number on both systems.

2.3.2 SAP HANA landscape steps to activate SAP HANA system replication

This section describes the steps to activate SAP HANA system replication in the landscape.

Setting the environment variables on both systems

To simplify the setup, run these commands as `${sid}adm` on both nodes and keep values identical, as shown in Example 2-15.

Example 2-15 SAP HANA node on the source and the target

```
export SID=<SID>                                  # SAP HANA System ID (uppercase)
export sid=<sid>                                  # SAP HANA System ID (lowercase)
export INSTNO=<INSTNO>                            # SAP HANA Instance Number
export DIR_INSTANCE=/usr/sap/${SID}/HDB${INSTNO} # "${sid}adm" home directory

export SiteOnPrem=        # HANA System Replication Site Name 1 - Migration from On-Prem - Source
export SiteOnCloud=     # HANA System Replication Site Name 2 - Migration to On-Cloud - Target

export NODE1=                  # Hostname of On-Premises Server
export NODE2=                  # Hostname of IBM Cloud Power Virtual Server Instance
```

Verifying that SAP HANA is running on both systems

Log in as the operating system user `${sid}adm`, use the HDB `proc` command to verify that all services are started.

Run the following command on both the primary and secondary SAP HANA systems:

```
sudo -i -u ${sid}adm -- HDB proc
```

Important: The SAP HANA version on each node must be equal to or greater than the version on the primary server.

Run the following command on each node to determine the SAP HANA server version:

```
sudo -i -u ${sid}adm -- HDB version
```

The target system version must be equal to or greater than the source system version. The only exception applies to an `Active` or `Active read-enabled` configuration, where the HDB version must be identical on both the source and target systems.

Help ensure consistency across both systems:

- Ensure that the system configuration is identical on both the source and target servers. Then, compare the settings in the `INI` files on both systems.
- For a scale-out configuration, ensure that the number of worker nodes and their roles are identical on both the source and target servers.
- Use the same `${sid}` and instance numbers on both systems.
- Back up the `PKI SSFS.key` and the `.DAT` files from the primary and secondary systems.
- Copy the existing `PKI` keys from the primary system to the secondary system.

To help ensure that the system can be recovered to its original installed state if needed, back up the existing keys on both the primary and secondary systems, as shown in Example 2-16.

Example 2-16 Backup command

```
sudo -i -u ${sid}adm -- cp -p
/usr/sap/${SID}/SYS/global/security/rsecssfs/data/SSFS_${SID}.DAT
/usr/sap/${SID}/SYS/global/security/rsecssfs/data/SSFS_${SID}.DAT_<hostname>
```

After backing up the existing `PKI SSFS.key` and `.DAT` files, copy them from the primary system to the target system.

The SAP HANA 2.0 data and log transmission channels for the replication process require authentication by using the system `PKI` SSFS storage certificate files (see SAP Note 2369981 - Required configuration steps for authentication with HANA System Replication).

The system PKI SSFS storage certificate files are stored in the `/usr/sap/${SID}/SYS/global/security/rsecssfs/` directory in the data and key subdirectories.

Log in as the `${sid}adm` user on NODE2 and copy files `SSFS_${SID}.DAT` and `SSFS_${SID}.KEY` from NODE1 by running the following commands (see Example 2-17).

Example 2-17 Copy commands

```
scp ${NODE1}:/usr/sap/${SID}/SYS/global/security/rsecssfs/data/SSFS_${SID}.DAT
/usr/sap/${SID}/SYS/global/security/rsecssfs/data/SSFS_${SID}.DAT

scp ${NODE1}:/usr/sap/${SID}/SYS/global/security/rsecssfs/key/SSFS_${SID}.KEY
/usr/sap/${SID}/SYS/global/security/rsecssfs/key/SSFS_${SID}.KEY
```

The copied `PKI` SSFS storage certificates on NODE2 become active when the SAP HANA system starts.

Verifying that the log_mode *configuration parameter is set to normal*

Set the **`log_mode`** parameter to `normal` in the `persistence` section of the `global.ini` file on both the primary and secondary SAP HANA servers.

Run the following command on both systems and verify the output:

```
sudo -i -u ${sid}adm -- grep -i 'log_mode'
/usr/sap/${SID}/HDB${INSTNO}/exe/config/global.ini
log_mode=normal
```

Registering the primary server first

On the primary SAP HANA system, run the following command to register the node as the primary for SAP HANA System Replication:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_enable --name=${SiteOnPrem}
```

The expected output appears as shown in Example 2-18.

Example 2-18 Output status

```
Name server is active, proceeding ...
successfully enabled system as system replication source site
done.
```

Verifying that the primary system is registered

Before you configure system replication, verify that the primary system is registered.

To verify that the primary system is successfully registered, run the following command:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_state
```

The expected output appears as shown in Example 2-19.

Example 2-19 System register output

```
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~~


online: true

mode: primary
operation mode: primary
site ID: 1
site name: SiteOnPrem

is source system: true
is secondary/consumer system: false
has secondaries/consumers attached: false
is a takeover active: false
is primary suspended: false

Host Mappings:
~~~~~~~~~~~~~~


Site Mappings:
~~~~~~~~~~~~~~
SiteCloud (primary/)

Tier of SiteCloud: 1

Replication mode of SiteCloud: primary

Operation mode of SiteOnPrem :


Hint based routing site:
done.
```

Ensuring that SAP HANA is not active on the secondary site

The secondary site must remain inactive as an SAP HANA server.

To stop SAP HANA database services, run the following command:

```
sudo -i -u ${sid}adm -- HDB stop
```

The expected output appears as shown in Example 2-20.

Example 2-20 Second server stop confirmed

```
hdbdaemon will wait maximal 300 seconds for NewDB services finishing.
Stopping instance using: /usr/sap/${SID}/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr
10 -function Stop 400

10.08.2023 10:32:07
Stop
OK
Waiting for stopped instance using: /usr/sap/${SID}/SYS/exe/hdb/sapcontrol -prot
NI_HTTP -nr 10 -function WaitforStopped 600 2


10.08.2023 10:32:51
WaitforStopped
OK
hdbdaemon is stopped.
```

Registering the secondary system

Register the secondary system by running the following command:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_register
--name=<secondarySiteName>
--remoteHost=<primary_host>
--remoteInstance=<primary_systemnr>
--replicationMode=[sync|syncmem|async]
--operationMode=[delta_datashipping|logreplay|logreplay_readaccess]
```

For example, use the following values:

- `SiteOnCloud` as the secondary site name
- `syncmem` as replication mode and
- `logreplay` as operation mode

Example 2-21 shows how to run the command using the **`logreplay`** option.

Example 2-21 Command register second system

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_register
--name=${SiteOnCloud}
--remoteHost=${NODE1}
--remoteInstance=${INSTNO}
--replicationMode=syncmem     --operationMode=logreplay
```

Example 2-22 shows the result of the command.

Example 2-22 Sample logreplay output

```
Thu 10 Aug 10:36:13 CEST 2023
adding site ...
collecting information ...
updating local ini files ...
done.
```

Troubleshooting hdbnsutil errors with SELinux enabled

If Security-Enhanced Linux (SELinux) is enabled, the output of **hdbnsutil** might not appear as expected. The following list shows possible symptoms:

- Command is not recognized error message.
- Usage information displayed.

When SELinux is set to `enforcing`, it prevents the **hdbnsutil** command from restarting the `saphostagent` in the `${sid}adm` user context. Either add the appropriate SELinux security policies or as SAP recommends, disable SELinux.

To check the current SELinux status, run the following command and review the output:

```
sestatus
SELinux status:                 enforcing
```

If the **sestatus** command returns `enforcing`, commands can be blocked even when they run with root privileges, depending on the security policy.

To disable SELinux temporarily, run the following command:

```
sudo setenforce 0
```

SELinux is now temporarily disabled until the next restart.

Check the status again by running the **sestatus** command. The status shows `disabled`.

Next, check whether the `saphostagent` process is running by using this command:

```
sudo ps -ef | grep -i host
```

If the output does not display a process, restart the `saphostagent` manually:

```
sudo -i -u ${sid}adm -- /usr/sap/hostctrl/exe/saphostexec -restart
/usr/sap/hostctrl/exe/host_profile
```

Checking the state on both sides of the SAP HANA system replication

Verify the replication state on both the primary and secondary systems to help ensure synchronization.

Checking the primary system state

To verify the system replication state on the primary node, run the following command on the primary server:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_state
```

The expected output appears as shown in Example 2-23.

Example 2-23 Output on the prime node

```
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~


online: true

mode: primary
operation mode: primary
site ID: 1
site name: SiteOnPrem_hostname
```

```
is source system: true
is secondary/consumer system: false
has secondaries/consumers attached: true
is a takeover active: false
is primary suspended: false

Host Mappings:
~~~~~~~~~~~~~~~~

<SiteOnCloud_hostname> -> [SiteOnPrem_hostname] <SiteOnPrem_hostname_hostname>
<SiteOnCloud_hostname> -> [SiteOnCloud] <SiteOnPrem_hostname_hostname>


Site Mappings:
~~~~~~~~~~~~~~~~
SiteOnPrem_hostname (primary/primary)
    |---SiteOnCloud (syncmem/logreplay)

Tier of SiteCloud: 1
Tier of SiteOnPrem_hostname: 2

Replication mode of SiteOnPrem_hostname: primary
Replication mode of SiteOnCloud: syncmem

Operation mode of SiteOnPrem_hostname: primary
Operation mode of SiteOnCloud logreplay

Mapping: SiteOnPrem_hostname -> SiteOnCloud

Hint based routing site:
done.
```

Checking the secondary system state

After confirming the replication status on the primary system, verify the state of the secondary node to ensure that it is properly configured and ready for synchronization.

Run the following command on the secondary server:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_state
```

The expected output appears as shown in Example 2-24.

Example 2-24 Secondary system state

```
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~~~~~

online: false

mode: syncmem
operation mode: unknown
site ID: 2
site name: SiteOnCloud

is source system: unknown
is secondary/consumer system: true
```

```
has secondaries/consumers attached: unknown
is a takeover active: false
is primary suspended: false
is timetravel enabled: false
replay mode: auto
active primary site: 1

primary masters: <SiteOnPrem_hostname_hostname>
done.
```

Restarting the secondary server

After both SAP HANA servers are configured as replication partners, restart the secondary SAP HANA server to initiate the replication process by running this command:

```
sudo -i -u ${sid}adm -- HDB start
```

The expected output appears as shown in Example 2-25.

Example 2-25 Output on the second server

```
StartService
OK

Starting instance using: /usr/sap/${SID}/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr
10 -function StartWait 2700 2
OK


10.08.2023 10:38:47
Start
OK

10.08.2023 10:40:17
StartWait
OK
```

Verifying that SAP HANA is running on the secondary system

After confirming that the secondary server started successfully, verify that SAP HANA is active and that replication is functioning as expected. Use the **HDB info** or **HDB proc** command to confirm the system status, then check the replication state from the primary system.

Run the following command on the primary server:

```
sudo -i -u ${sid}adm -- hdbnsutil -sr_state
```

The expected output appears as shown in Example 2-26.

Example 2-26 Output from the HDB info command

```
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~

online: true

mode: primary
operation mode: primary
site ID: 1
```

```
site name: SitePrem

is source system: true
is secondary/consumer system: false
has secondaries/consumers attached: true
is a takeover active: false
is primary suspended: false

Host Mappings:
~~~~~~~~~~~~~~~

<SiteOnCloud_hostname> -> [SiteOnPrem] <SiteOnPrem>
<SiteOnCloud_hostname> -> [SiteOnCloud] <SiteOnPrem>


Site Mappings:
~~~~~~~~~~~~~~~
SiteOnPrem (primary/primary)
    |---SiteOnCloud (syncmem/logreplay)

Tier of SiteOnprem : 1
Tier of SiteOnCloud: 2

Replication mode of SiteOnPrem: primary
Replication mode of SiteOnCloud: syncmem

Operation mode of SiteOnPrem: primary
Operation mode of SiteOnCloud: logreplay

Mapping: SiteOnprem_hostname -> SiteOnCloud

Hint based routing site:
done.
```

Checking the replication status

After restarting the secondary system, verify that the replication process started and is synchronizing data correctly. Use the Python script on the primary server to monitor replication progress.

Run the following command:

```
sudo -i -u ${sid}adm -- python
${DIR_INSTANCE}/exe/python_support/systemReplicationStatus.py
```

The expected output appears as shown in Example 2-27.

Example 2-27 Python script output

```
|Database |Host     |Port  |Service Name |Volume ID |Site ID |Site Name |Secondary |Secondary
|Secondary |Secondary  |Secondary     |Replication |Replication |Replication    |Secondary    |
|         |         |      |             |          |        |          |Host      |Port
|Site ID   |Site Name  |Active Status |Mode        |Status      |Status Details |Fully Synced |
|-------- |-------- |----- |------------ |--------- |------- |--------- |----------|---------
|--------- |---------- |------------- |----------- |----------- |-------------- |-------------|
|SYSTEMDB |<NODE1>  |31001 |name server   |       1 |     1 |SiteOnPrem|  <NODE2> |   31001 |
2 |SiteOnCloud|YES           |SYNCMEM     |ACTIVE      |               |         True|
```

```
|S4H     |<NODE1>  |31007 |xsengine    |       2 |     1 |SiteOnPrem|  <NODE2> |   31007 |
2 |SiteOnCloud|YES           |SYNCMEM     |ACTIVE      |               |          True|
|S4H     |<NODE1>  |31040 |docstore    |       5 |     1 |SiteOnPrem|  <NODE2> |   31040 |
2 |SiteOnCloud|YES           |SYNCMEM     |ACTIVE      |               |          True|
|S4H     |<NODE1>  |31003 |indexserver |       3 |     1 |SiteOnPrem|  <NODE2> |   31003 |
2 |SiteOnCloud|YES           |SYNCMEM     |ACTIVE      |               |          True|
|S4H     |<NODE1>  |31011 |dpserver    |       4 |     1 |SiteOnPrem|  <NODE2> |   31011 |
2 |SiteOnCloud|YES           |SYNCMEM     |ACTIVE      |               |          True|

status system replication site "1": ACTIVE
overall system replication status: ACTIVE

Local System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

mode: PRIMARY
site ID: 1
site name: SiteOnPrem_hostname
```

Four methods to check the system replication status

Verify SAP HANA system replication by using one of several methods. Each method provides a different level of detail and perspective on the replication state.

Option 1: landscapeHostConfiguration.py

Use the **`landscapeHostConfiguration.py`** Python script to display a status line for each SAP HANA server from a server-level perspective.

Run the following command:

```
sudo -i -u ${sid}adm -- python
${DIR_INSTANCE}/exe/python_support/landscapeHostConfiguration.py
```

Ensure that each server that is listed in the output shows `OK` in the Host Status column.

Option 2: systemReplicationStatus.py

Use the **`systemReplicationStatus.py`** Python script to view the SAP HANA system replication at the database level. The script displays one status line per database and includes an overall status summary.

Run the following command:

```
sudo -i -u ${sid}adm -- python
${DIR_INSTANCE}/exe/python_support/systemReplicationStatus.py
```

Ensure that each database listed shows `ACTIVE` in Replication Status column. The script must display the following output:

```
overall system replication status: ACTIVE
```

Option 3: hdbcons

Use the SAP HANA DB Management Client Console to check detailed system replication status per server and per service.

Run the following command:

```
sudo -i -u ${sid}adm -- hdbcons -e hdbindexserver "replication info"
```

Option 4: SQL script

Use this method to view the SAP HANA system replication by host and site. The SQL script can be run in SAP HANA Studio, SAP HANA Cockpit, or any SQL client.

Run the following SQL script:

```
select host, SECONDARY_HOST, PORT, SITE_NAME, SECONDARY_SITE_NAME,
REPLICATION_MODE, REPLICATION_STATUS,
REPLICATION_STATUS_DETAILS,SECONDARY_ACTIVE_STATUS from M_SERVICE_REPLICATION;
```

Verify the `REPLICATION_STATUS` and `REPLICATION_STATUS_DETAILS` columns.

After replication completes

Before disabling the replication setup, check the trace logs for any inconsistencies or anomalies. After the replication finishes, the database contains all active services on the primary system only. Review the trace logs for issues.

2.3.3 Checking the database trace files

This section explains how to check database trace files for diagnostic or monitoring purposes. The system writes database trace information to service-specific files at the OS level.

The trace directory is at the following path:

```
/usr/sap/<SID>/HDB<inst>/<host>/trace/DB_<SID>/
```

As the `${sid}adm` user, use the **`cdtrace`** alias to quickly go to the trace directory:

```
$ cdtrace
```

Database trace files use the following naming convention:

```
<service>_<host>.<port>.<counter>.trc
```

In the context of dynamic tiering, an extra file can exist with this naming convention (for more information, see SAP Note 2871785):

```
esserver_console_<host>.<port>.<counter>.trc
Example: indexserver_saphana01.30003.024.trc
```

Database trace files can be accessed in the following ways:

- Directly on the operating system level
- Select **SAP HANA Studio** → **Administration** → **Diagnosis Files**.
- Select **DBACOCKPIT** → **Diagnostics** → **Diagnosis Files**.

2.3.4 SAP HANA system replication resources

For more information, see the following resources:

- SAP Note 1999880 - HANA System Replication FAQ
- SAP Note 11969700 - SQL Statement Collection for SAP HANA
- SAP Note 3357978 - Configuring SAP HANA Multi Target System Replication
- SAP HANA Replication Setups
- SAP HANA System Replication
- SAP HANA multi-target System Replication

- SAP HANA System Replication
- SAP Help Portal Documentation HSR HANA 2.0 SP07

2.4 Performing an SAP HANA system replication takeover

The final step is to perform an SAP HANA system replication takeover. A takeover activates services on the target site and helps ensure that clients connect to the new location.

For procedural details and decision guidance, see the following resources:

- Performing a Takeover
- SAP Note 2063657 SAP HANA System Replication Takeover Decision Guideline

3

Migrating SAP ECC on Oracle to IBM Power Virtual Server

This chapter describes the end-to-end process of migrating an SAP ECC system that runs on an Oracle database to IBM Power Virtual Server (PowerVS). Keep the existing Oracle database platform and move the SAP application and infrastructure to the IBM Cloud environment.

Begin by preparing the source system for stability and readiness. Run a backup and restore procedure by using IBM-supported tools and services to help ensure a secure and efficient data transfer. To minimize business disruption, choose a migration strategy based on system size, complexity, and downtime tolerance. These strategies reduce cutover time and support a smooth transition.

Validate and test at the end to confirm data integrity, verify application functions, and ensure that the SAP system operates as expected in the PowerVS environment.

This chapter describes the following topics:

- “Preparation steps on the source system”
- “Backing up the source Oracle database by using RMAN”
- “Restoring the Oracle database on the target system”

© Copyright IBM Corp. 2025.

3.1 Preparation steps on the source system

This section describes the process of migrating an existing SAP ECC system on Oracle to PowerVS. Keep the Oracle database platform and move the SAP application and infrastructure to IBM Cloud. The source system can run on-premises or in another cloud. The target system runs on PowerVS, which provides enterprise-grade virtualized infrastructure for SAP workloads.

This migration does not require a database conversion to SAP HANA or IBM Db2. Maintain the existing Oracle database and move the SAP application and its configurations to the new platform.

The following key activities support the migration while preserving the existing Oracle database:

- Infrastructure transition

  Move the SAP ECC system to PowerVS.
- Database preservation

  Retain the Oracle database without conversion.
- Application migration

  Transfer the SAP ECC application, including custom configurations and interfaces.
- Downtime planning

  Select a migration method that minimizes business disruption.
- System validation

  Confirm data integrity, application consistency, and operational readiness after migration.

3.1.1 Target audience and intent

This section provides two Oracle RMAN-based options to migrate AIX-based Oracle databases from IBM Power to PowerVS. The target audience includes solution and infrastructure architects and DBAs. The goal is to provide representative steps for planning and performing Oracle database migrations. Each scenario has unique deployment, configuration, and resource considerations.

In this section, discovery refers to the customer DBA's responsibility to document the current Oracle and SAP infrastructure. Use this information to raise incidents or report issues.

3.1.2 Scope and coverage

Ensure that network connectivity sufficiently supports data transfer and replication. If required, transfer on-premises backups or database files to PowerVS by using Seagate Lyve Mobile, IBM Cloud Object Storage, or IBM Aspera® Connect.

For more information about data transfer options, see the following resources:

- Seagate Lyve Mobile Solution
- Getting started with IBM Cloud Object Storage
- IBM Aspera Connect

Review task variability and follow Oracle and SAP documentation for command-specific requirements. See the cited Oracle documentation for command-specific details. For this migration, use the User Concept - Oracle Standard as outlined in SAP Note 1915323 - OS User Concept for SAP NetWeaver for 12c and higher (SAP S-user required).

3.1.3 Disclaimer

Align these procedures with established processes for operating and maintaining production and non-production systems. Schedule maintenance and downtime as required. Review the representative procedures in the context of the specific environment and adjust as needed.

These Oracle Database migration options are not unique to PowerVS. Oracle technical members are familiar with the procedures, but not every technical detail is explicitly identified. The DBA who runs the procedures must understand the full scope of Oracle database backup and recovery methods, including details that are not explicitly stated.

3.1.4 Discovery and migration option selection

This section explains how to gather essential information about the source environment and choose the most suitable migration option. It introduces the discovery process and outlines how to evaluate migration options based on technical and business requirements.

Discovery process

Perform a comprehensive assessment of the source environment before selecting a migration option. This step helps ensure that all technical and business requirements are understood and that migration planning is accurate.

Complete a detailed discovery that addresses the following areas:

- Business, technical, functional, and nonfunctional requirements of the Oracle database (including access and availability requirements)
- Installation configuration, hardware and software inventory, and sizing information for the source database and the underlying platform
- Metrics for sizing, change rate, service-level requirements, downtime tolerance, and network capacity for data transfer

Evaluate and deploy the target environment to match the source system's capacity, performance, and availability.

Migration options covered by this chapter

This chapter describes options for migrating SAP systems to PowerVS:

- Option 1

  Use a consistent RMAN backup from the source system, with generated files for the RMAN **`DUPLICATE DATABASE`** procedure to restore to PowerVS.

- Option 2

  Use an inconsistent RMAN backup plus archived redo logs (or Level 0 plus incremental backups), with files for the `RMAN RESTORE/RECOVER DATABASE` procedure to restore to PowerVS.

- Option 3

  Use Oracle Data Guard. This option is under evaluation and is expected to become the recommended method for same-platform migrations. Create a mirror and switch over with minimal downtime. IBM will update documentation after testing and assessment are finished.

Figure 3-1 illustrates these options.

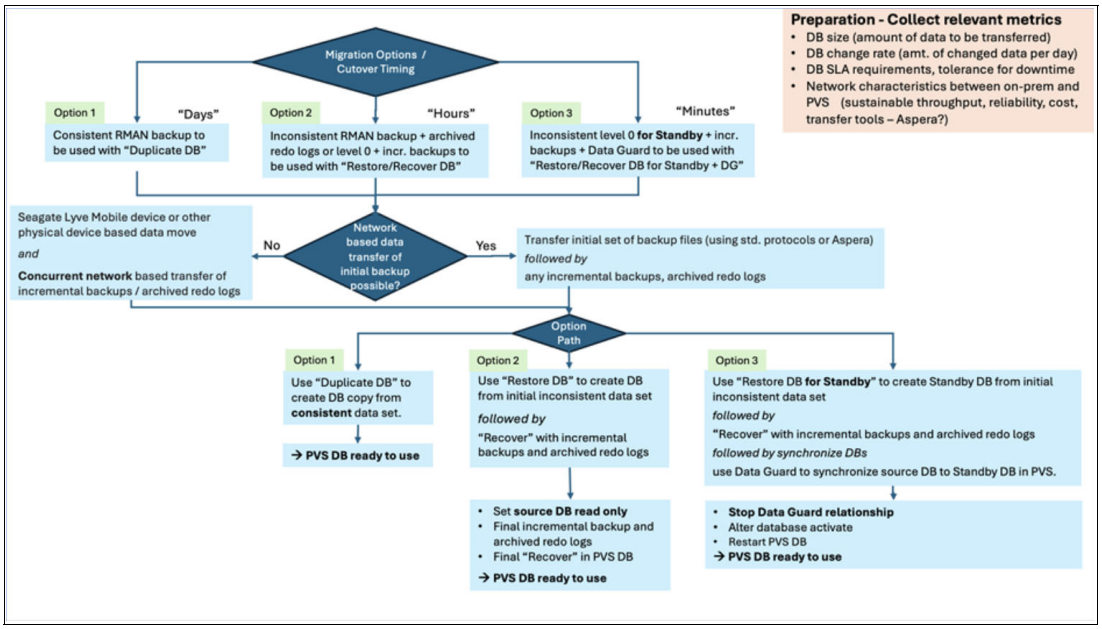

Figure 3-1 Preparation: Collecting relevant metrics

Selected considerations

When planning the migration, consider the following factors to help ensure a smooth transition and minimize risk:

- Service-level requirements, particularly Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO)
  - Option 1

    Use when a delay of several days before cutover from the source database to the target database is acceptable. This option is suitable for non-production workloads but rare for production databases.
  - Option 2

    Use when a service interruption of several hours is acceptable. This option is suitable for non-production workloads and some production databases.
  - Option 3 (under evaluation)

    Use Oracle Data Guard database synchronization to enable cutover in minutes. Documentation will be updated after testing and assessment are finished.
- Database Size, network throughput, and reliability

  Transferring large backups over the network takes time. For example, transferring 600 GB at 80 MBps takes about 125 minutes. Larger backups or slower connections require more time. Unreliable networks can disrupt transfers and cause delays.

Consider these recommendations:

- Use transfer tools that compress data during transfer, use high-throughput protocols, and decompress data after transfer.
- Use IBM Aspera accelerates data transfers from on-premises to PowerVS.

- Availability of skills to run

  These options require experienced DBAs and coordination with infrastructure teams to migrate database content to a new target database on PowerVS within a PowerVS workspace. Run the procedure as part of a detailed, rehearsed transfer and cutover plan.

3.1.5 Technical details and considerations

This section describes what RMAN backs up and what it does not to support accurate migration planning.

Files backed up by RMAN

The RMAN **`BACKUP`** command supports backing up the following types of files:

- Data files and control files
- Server parameter file
- Archived redo logs
- RMAN backups

Important: RMAN does not back up all Oracle database files. It excludes network configuration files, password files, and Oracle home contents. Some Oracle features, such as external tables, depend on files other than data files, control files, and redo logs.

Other files to include

Before running the RMAN backup, ensure that the following files from the source database are included:

- Oracle database parameter file: `init<SID>.ora`
- TNS network configuration files: (for example, `listener.ora` and `tnsnames.ora`)
- Oracle database password file, if present

Include these files in the backup and recovery plan. A best practice is to copy the files to a directory on the same mount as the RMAN database backups.

Backup formats

Run the **`BACKUP`** command in RMAN to produce one or more backup sets or image copies**:**

- A *backup set* is an RMAN-specific proprietary format.
- An *image copy* is a bit-for-bit copy of a file.

By default, RMAN creates backup sets. This chapter covers backup sets only.

For definitions, see the following topics:

- Backup set explanation
- Image Copy explanation

Assumptions

The following assumptions apply to the backup process:

- If the source database is a production system, shutting down for a full backup is unlikely. In this case, use RMAN incremental backups.
- No changes to the source database (schema and configuration) occur during this procedure.
- Migrate the database from source to target by using the following Oracle RMAN options:
  - **DUPLICATE DATABASE**
  - **BACKUP DATABASE**
  - **RESTORE DATABASE**
  - **RECOVER DATABASE**
- The target environment has Oracle homes that are preinstalled with a version that matches the source instance.
- RMAN **DUPLICATE**, **RESTORE**, and **RECOVER** commands provide options to modify data file locations. This chapter does not cover that option. The assumption is that the location of database files, either AIX JFS2 file systems or Oracle ASM disk groups, is identical on the source and the target environments.
- The target OS uses one of the following configurations:
  - An exact copy (an OVA generated by using **mksysb**)
  - A freshly deployed version of AIX that IBM supports and Oracle certifies for the Oracle version to be instantiated

3.1.6 Preparation

Prepare the source and target environments before starting the migration. This preparation helps ensure that the target system meets performance and capacity requirements and that the migration process runs smoothly.

Documenting the current source database sizing

Document the current source database sizing to guide the configuration of the target Oracle instance in PowerVS. Perform the following tasks:

- Record compute, IOPS, and storage for one or more source Oracle instances.
- Construct the target PowerVS Oracle instance to match or exceed the source system.
- On the target system, select the disk tier and disk capacity to meet IOPS requirements.
- Run performance tests on the source database to capture metrics for future reference and comparison.

Confirming that a recent, successful backup exists

Use standard backup tools to confirm that a recent, successful, full backup of the source Oracle system exists. Ensure that the source system can be fully restored. Then complete the following tasks:

- If confirmation is not possible, review the RMAN backup procedures.
- Work with customer DBAs or SMEs to validate database integrity and identify any corrupted blocks or schema problems.
- Record the location of Oracle instances and credentials for DBA access.

Documenting the RMAN configuration

Document the current RMAN configuration for each Oracle database instance before starting the migration:

1. Connect to the Oracle server:

   ```
   ssh oracle@<host name>
   ```

2. Check that the variable `ORACLE_SID` is set:

   ```
   echo $ORACLE_SID
   ```

   If a value is not shown, set one manually:

   ```
   setenv ORACLE_SID <SID value>
   ```

   Note: Replace `<SID value>` with the appropriate SID value.

3. Connect to the source database by using the **`rman`** command:

   ```
   rman target /
   ```

   An example of a typical output appears in Example 3-1.

   Example 3-1 Output showing that the target DB is connected

   ```
   Recovery Manager: Release 19.0.0.0.0 - Production on Thu May 2 13:24:05 2024
   Version 19.22.0.0.0

   Copyright (c) 1982, 2019, Oracle and/or its affiliates.  All rights reserved.
   connected to the target database: Ex (DBID=xxxxxxxx)
   ```

4. Display the RMAN configuration settings:

   ```
   RMAN> SHOW ALL;
   ```

   Store and document the output. For reference, a typical `SHOW ALL` output appears as shown in Example 3-2.

   Example 3-2 Showing all output

   ```
   using target database control file instead of recovery catalog
   RMAN configuration parameters for database with db_unique_name EXX are:
   CONFIGURE RETENTION POLICY TO REDUNDANCY 1; # default
   CONFIGURE BACKUP OPTIMIZATION OFF; # default
   CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
   CONFIGURE CONTROLFILE AUTOBACKUP ON; # default
   CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '%F'; # default
   CONFIGURE DEVICE TYPE DISK PARALLELISM 8 BACKUP TYPE TO BACKUPSET; # default
   CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
   CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
   CONFIGURE MAXSETSIZE TO UNLIMITED; # default
   CONFIGURE ENCRYPTION FOR DATABASE OFF; # default
   CONFIGURE ENCRYPTION ALGORITHM 'AES128'; # default
   CONFIGURE COMPRESSION ALGORITHM 'BASIC' AS OF RELEASE 'DEFAULT' OPTIMIZE FOR
   LOAD TRUE ; # default
   CONFIGURE RMAN OUTPUT TO KEEP FOR 7 DAYS; # default
   CONFIGURE ARCHIVELOG DELETION POLICY TO NONE; # default
   CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/oracle/EXX/19/dbs/snapcf_EXX.f'; #
   default
   ```

5. Exit RMAN:

```
RMAN > exit;
```

Verifying the disk space

Confirm that sufficient local disk space (formatted, and mounted) exists to store the compressed migration database backup sets. Consider the following guidelines:

- For best performance, use storage that is local to the instance being backed up.
- Backing up directly to IBM Cloud Object Storage can introduce delays because of network latency.
- Compression planning:
  - MEDIUM compression reduces backup file size to about 25% of the database size.
  - Ensure that the local file system has enough capacity to store all backup files that are required to reach the wanted restore point.
- Ensure that the target system can access the backup sets by using one of these options:
  - Local file system where backups are stored.
  - IBM Aspera agent-assisted transfer.
  - NFS mount to a Seagate Lyve Mobile device or IBM Cloud Object Storage (for best performance, copy to local storage before restore).

File system for backup sets on the source

Use a dedicated local file system (for example `/backup/rman` for the migration backup sets. As an alternative, a direct restore or recovery from an NFS-mounted Seagate Lyve Mobile device was also tested successfully by using 10 GbE and 100 GbE.

Confirm the available space for backups by running the following command:

```
df -g /backup/rman
```

File system for backup sets on the target

Provide the backup sets to the target system by using one of these methods:

- A local file system (copied backup sets)
- An NFS mount to a Seagate Lyve Mobile transfer device
- An NFS mount to IBM Cloud Object Storage where backup files are already stored

Note: Depending on network speed, copy files from NFS mount to local storage for better performance.

3.2 Backing up the source Oracle database by using RMAN

This section explains how to create RMAN backup sets for database migration.

The following options are available:

- Option 1: Offline backup with RMAN **`DUPLICATE DATABASE`**
- Option 2: Online backups with RMAN incremental levels

For more information, see Oracle Database Backup and Recovery - February 2024.

These conditions apply to the backup options:

- Option 1 assumes a consistent offline backup of the source database.
- Option 2 requires the source database to be in ARCHIVELOG mode before the backup procedure.

When configuring backups, balance parallelism, compression, and resource allocation with database availability and performance requirements. Consider the following factors:

- Apply a section size to the backup configuration**.** Without it, large files might be created, which makes handling or restarting after a transfer failure difficult.
- Parallelism during restore is limited to the number of backup files available for concurrent processing. A good section size helps control the size of each backup file and influences how many files are generated. This approach allows parallel processing during restore in PowerVS.

  Here are the compression options:

  - MEDIUM compression requires the Oracle Advanced Compression license and provides a good balance of performance and size reduction.
  - BASIC compression is slower and achieves reduced compression rates.
  - HIGH compression is CPU-intensive and provides limited benefits compared to MEDIUM. HIGH also requires the Advanced Compression license.
- Incremental backup, without database block change tracking performs a full scan of all data files, creating a high read workload. Although enabling block change tracking is not expected to affect database performance, test it in the target environment.

Network transfer: Standard protocols versus IBM Aspera

This section compares standard network transfer protocols with IBM Aspera to help select the most efficient method for migrating backup files. Review the following options:

- Standard protocols (scp or sftp) provide slower data transfer. Backup files can be transferred either directly to an IBM AIX LPAR in PowerVS or to IBM Cloud Object Storage. Using scp or sftp with IBM Cloud Object Storage assumes that an IBM FileManage gateway service is in place. Ensure that a gateway service is available or that a sftp server is installed and configured within or next to the target PowerVS environment.
- IBM Aspera offers a faster option for data transfer. In many situations, IBM Aspera transfers data several times faster than traditional TCP-based protocols.
- For more details, see IBM Aspera Technologies - IBM Cloud. This reference also contains the Accelerated network transfer migration guide.

Figure 3-2 shows the following migration options:

- RMAN Option 1: Offline Backup / Duplicate Database
- RMAN Option 2: Online Backups

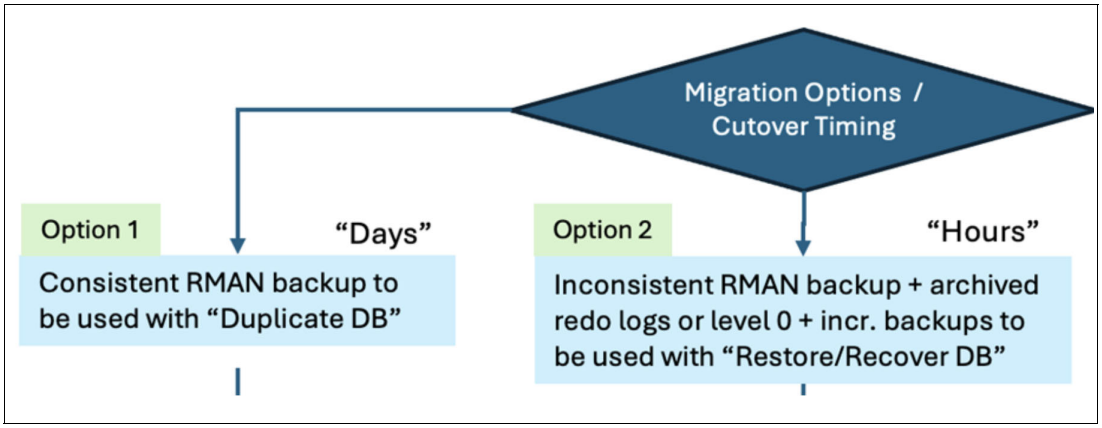

Figure 3-2 IBM Aspera backup procedure

3.2.1 RMAN option 1: Offline backup and duplicate database

This section explains how to run an offline backup and duplicate the database by using RMAN. Run this procedure with the database offline to create a full, consistent backup.

Application and database shutdown

Use standard operational procedures to shut down the SAP system and the Oracle database before performing the backup. If a multitenant architecture is in use, connect to the root container database (CDB). The backup includes all pluggable databases (PDBs).

For more information, see About Performing Operations on CDBs and PDBs.

Ensure that the RMAN configuration is documented

As described in Document RMAN Configuration, record the current RMAN configuration and associated parameters. Ensure that the RMAN configuration remains unchanged after completing special backups for migration so that scheduled backup operations continue as expected.

Backup option 1: Creating a target directory *(*RMAN Option 1)

This directory must match the backup script that is described in "Backup option 1: Backup script (RMAN Option 1)" on page 47. If the directory does not exist, the RMAN script fails. If an NFS mount is used, ensure that the `oracle` user has read and write (`rw`) permissions in the NFS export definition file.

Avoid using NFS for the RMAN backup process. Using JFS2 or a locally mounted file system achieves better results.

Create the backup target location and set `ORACLE_SID`. Replace `<SID>` with the correct value for the system:

```
mkdir -p /backup/rman/<sid>_option1
setenv ORACLE_SID <SID>
```

Use any target location if the `oracle` user can write to it, and modify the backup script to reference the target directory.

Backup option 1: Backup script *(*RMAN Option 1)

Use the RMAN script (**option1_backup.rman**) to perform the offline backup. Copy and paste the script and adjust it to match the target environment.

Example 3-3 shows the output of the backup script.

Example 3-3 Output of the backup script

```
connect target /
SHUTDOWN IMMEDIATE
STARTUP MOUNT
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT
'/backup/rman/<sid>_option1/option1_cmp_%d_%U';
CONFIGURE DEVICE TYPE DISK PARALLELISM 60;
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'/backup/rman/<sid>_option1/option1_cf_%F';
CONFIGURE COMPRESSION ALGORITHM 'HIGH';
CONFIGURE ENCRYPTION FOR DATABASE ON ;
CONFIGURE ENCRYPTION ALGORITHM 'AES256' ;
SET ENCRYPTION ON IDENTIFIED BY passwOrd ONLY ;
BACKUP AS compressed BACKUPSET section size 6000M DEVICE TYPE DISK DATABASE TAG
<YOUR TAG ID HERE> include current controlfile ;
alter database open;
quit;
```

Run the following script from the CLI as the `oracle` user:

```
rman @option1_backup.rman
```

By default, RMAN creates backups on disk. If a fast recovery area is enabled and the **FORMAT** parameter is not specified, RMAN creates backups in the recovery area and automatically assigns unique names. Modify the **FORMAT** parameter to direct output to the target location.

Main components of the RMAN script

The RMAN script includes several key components that helps ensure a consistent backup and restore process. The script includes the following components:

- Database shutdown and start in mount mode

  Shut down the database cleanly and start it in mount mode to help ensure a consistent backup. Use these commands:

  ```
  SHUTDOWN IMMEDIATE
  STARTUP MOUNT
  ```

- Control file backup

  The control file contains the RMAN catalog that is required to restore the backup pieces to a functioning database. Use these commands to configure the backup file system location:

  ```
  CONFIGURE CHANNEL DEVICE TYPE DISK
  CONFIGURE CONTROLFILE AUTOBACKUP FORMAT for DEVICE TYPE DISK TO
  '/backup/rman/<sid>_option1/option1_cf_%F';
  ```

Example 3-4 shows how the control file is included in the backup.

Example 3-4 Command for configuring an automatic backup

```
CONFIGURE CONTROLFILE AUTOBACKUP ON
BACKUP AS compressed BACKUPSET section size 6000M DEVICE TYPE DISK DATABASE TAG
EXX_1OOK_INITV3 include current controlfile ;
```

- Backup parallelism

  Set disk device parallelism to speed up the backup and reduce the backup window. The optimal value depends on CPU resources, storage throughput, and available memory.

- Backup compression

  Set backup file compression. Although HIGH was tested in the script, MEDIUM provides better performance and efficiency. Both require the Oracle Advanced Compression license:

  - Tested

    ```
    CONFIGURE COMPRESSION ALGORITHM 'HIGH';
    ```

  - Recommended

    ```
    CONFIGURE COMPRESSION ALGORITHM 'MEDIUM';
    ```

- Backup encryption

  Enable backup encryption and specify the algorithm and password for each RMAN session. If Transparent Data Encryption (TDE) is used, ensure that the wallet is open; otherwise, the session fails with a “wallet not open” error.

  Example 3-5 shows these settings.

  Example 3-5 Backup encryption command

  ```
  CONFIGURE ENCRYPTION FOR DATABASE ON ;
  CONFIGURE ENCRYPTION ALGORITHM 'AES256' ;
  SET ENCRYPTION ON IDENTIFIED BY passw0rd ONLY ;
  ```

- Backup command

  The final command triggers the backup of the database to disk by using RMAN-specific backup sets. Each RMAN process reads up to 6000M of consecutive data from a data file as a backup piece, compresses and encrypts that data, and writes it to the destination file system.

  The last action alters the database to open. Creating a specific tag (including the ORACLE_SID <SID>) for identification. The goal of the section size is to limit the size of a backup piece so that retransmission of a failed file transfer to the destination environment is manageable. Using a TAG simplifies the management of multiple backups in an RMAN catalog, as shown in Example 3-6.

  Example 3-6 TAG command

  ```
  BACKUP AS compressed BACKUPSET section size 6000M DEVICE TYPE DISK DATABASE TAG
  <SID>_1OOK_INITV3 include current controlfile ;
  alter database open;
  ```

- Backup validation and cross-check

  After the backup finishes, use RMAN commands to validate and cross-check the results (see Example 3-7 on page 49).

Example 3-7 RMAN command

```
RMAN> REPORT SCHEMA;
RMAN> LIST BACKUP SUMMARY;
RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
RMAN> VALIDATE DATAFILE 10;
RMAN> VALIDATE BACKUPSET 3;
RMAN> CROSSCHECK BACKUP;
```

The following commands provide information about the database and its backups:

- **REPORT SCHEMA;**

  Lists and displays information about database files, table spaces, and related objects.

- **LIST BACKUP SUMMARY;**

  Lists all existing backups. A SUMMARY option can be used.

- **LIST BACKUP <TAG ID>;**

  Lists a specific backup by TAG ID.

- **BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;**

  Validates the contents of backup files.

- **VALIDATE DATAFILE 10;**

  Validates a specific data file.

- **VALIDATE BACKUPSET 3;**

  Validates a specific backup set.

- **CROSSCHECK BACKUP;**

  Synchronizes the physical reality of backups and copies with their logical records in the RMAN repository.

Creating a PFILE

Create a plain-text database parameter file (PFILE) from the binary server parameter file (SPFILE). Keep the database in mounted mode. Use the appropriate <SID> in the command.

As the Oracle user, run the following commands:

```
export ORACLE_SID=<SID>
sqlplus "/ as sysdba";
```

Inside SQL*Plus, create the parameter file:

```
SQL> create pfile='/backup/rman/init<SID>.ora' from spfile;
```

After completing the migration backup, restore all RMAN parameters to their original discovery settings.

For Option 1, create one backup to transfer and restore on the target system. After completing the backup and validation, shut down the source database to maintain consistency. The source system is no longer be used.

3.2.2 RMAN option 2: online backups

This option creates one RMAN Level 0 backup and one or more Level 1 incremental backups with the database remains online.

Requirements

The following requirements must be met before starting the backup process:

- The database must be in ARCHIVELOG mode.
- Include all required archived redo logs in the backups.

Creating target directories

An RMAN Incremental Level 0 backup is a full backup that includes the entire database. An RMAN Incremental Level 1 backup captures changes after the previous Level 0 or Level 1 backup. The cumulative option changes this behavior. For details, see Oracle RMAN backup concepts documentation.

If a multitenant architecture is used, connect to the root CDB. The backup includes all PDBs.

For more information, see Performing Operations on CDBs and PDBs.

Run the commands in Example 3-8 as the Oracle user. Replace `<SID>` with the correct value for the system.

Example 3-8 Creating the target directories

```
mkdir -p /backup/rman/<sid>_option2
mkdir -p /backup/rman/<sid>_option2_inc1

setenv ORACLE_SID <SID>
```

Checking the database size

Ensure that enough file system space exists to store the backup. Use this SQL statement to determine the current database size:

```
SELECT SUM (bytes)/1024/1024/1024 AS GB FROM dba_segments;
```

Ensuring that the RMAN configuration is documented

As described in the preparation section Document RMAN Configuration, record the current RMAN configuration and associated parameters. Restore the RMAN configuration to its original state after completing special backups for migration so that scheduled backup operations continue as expected.

Note: The full backup (Level 0), all incremental backups (Level 1) after the last full backup, and the archived redo logs are required to restore and recover the database successfully. When the database is moved to a new server and services are ended on the source system, the last incremental backup is an offline backup.

Backup option 2: Backup scripts (RMAN Option 2) Full online backup (Level 0)

The first RMAN script **`option2_backup_full.rman`** configures the RMAN environment and creates the initial Level 0 (full) online backup. The key command to start the full (level 0) backup in this script is shown in Example 3-9):

```
BACKUP ... incremental level 0 ...
```

Example 3-9 Backup script `option2_backup_full.rman` script

```
connect target /
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT '/backup/rman/<sid>_option2/lev0_%d_%U';
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE DEVICE TYPE DISK PARALLELISM 60;
```

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'/backup/rman/<sid>_option2/lev0_cf_%F';
CONFIGURE COMPRESSION ALGORITHM 'MEDIUM';
CONFIGURE ENCRYPTION FOR DATABASE ON ;
CONFIGURE ENCRYPTION ALGORITHM 'AES256' ;
SET ENCRYPTION ON IDENTIFIED BY passw0rd ONLY ;
BACKUP  tag '<Your TAG Here>' incremental level 0 AS compressed BACKUPSET section
size 6000M DEVICE TYPE DISK DATABASE INCLUDE CURRENT CONTROLFILE PLUS ARCHIVELOG;
quit;
```

Run the following script as the Oracle user:

```
rman @option2_backup_level0.rman
```

After the full (level 0) backup finishes successfully, the following message should display:

```
Recovery Manager complete.
```

Important: RMAN backs up data to the configured default device for the type of backup requested. By default, RMAN creates backups on disk. If a fast recovery area is enabled and the **FORMAT** parameter is not specified, RMAN creates backups in the recovery area and automatically assigns unique names. Modify the **FORMAT** parameter in the script to direct output to the target location.

Incremental backup: Level 1

Use the RMAN script **option2_backup_inc1.rman** to create the Level 1 incremental backup.

Replace <sid> with the system SID and ensure that the target backup directory matches the directory that was created in Backup Option 2 - Create Target Directories.

Example 3-10 The option2_backup_inc1.rman script

```
connect target /
CONFIGURE CHANNEL DEVICE TYPE DISK FORMAT
'/backup/rman/<sid>_option2_inc1/inc1_%d_%U';
CONFIGURE BACKUP OPTIMIZATION ON;
CONFIGURE DEVICE TYPE DISK PARALLELISM 60;
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO
'/backup/rman/<sid>_option2_inc1/inc1_cf_%F';
CONFIGURE COMPRESSION ALGORITHM 'MEDIUM';
CONFIGURE ENCRYPTION FOR DATABASE ON ;
CONFIGURE ENCRYPTION ALGORITHM 'AES256' ;
SET ENCRYPTION ON IDENTIFIED BY passw0rd ONLY ;
BACKUP TAG `<SID>`_INC1 incremental level 1 AS compressed BACKUPSET section size
6000M DEVICE TYPE DISK DATABASE INCLUDE CURRENT CONTROLFILE PLUS ARCHIVELOG;
quit;
```

Run the following script as the Oracle user:

```
setenv ORACLE_SID <SID>
rman @option2_backup_inc1.rman
```

Differences from the Level 0 script

The following differences apply to the Level 0 script:

- The backup directory changes to the incremental directory.
- The TAG identifies the backup set as incremental.
- The backup level changes from 0 to 1.

Configuration and validation steps

Complete these steps to help ensure that the backup and recovery process works as expected.

Key configuration details for RMAN option 2

Review the RMAN configuration to understand how archived redo logs are handled and how to avoid duplication. The following details are necessary:

- **`PLUS ARCHIVELOG`**

  RMAN includes all archived redo logs in a Level 0 or Level 1 backup if the **`PLUS ARCHIVELOG`** flag is specified. This action can result in duplicate redo logs across consecutive incremental backups.

- Avoiding duplicates

  Use **`CONFIGURE BACKUP OPTIMIZATION ON;`** to instruct RMAN to check whether a specific archived redo log was already included in a previous backup. If yes, RMAN does not include the file in a new incremental backup.

  Important: Do not use this option if other backups exist that are not being transferred to the destination. In that case, set the option to `OFF`**.**

- Control file backup

  Always include the control file because it contains the RMAN catalog that is required to restore the backup pieces into a functioning database, as shown in Example 3-11.

  Example 3-11 Configuring the control file automatic backup

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;

BACKUP AS compressed BACKUPSET section size 6000M DEVICE TYPE DISK DATABASE
PLUS ARCHIVELOG TAG ECOM_option1 include current controlfile;
```

- Parallelism

  Set disk device parallelism to speed up the backup and reduce the backup window. The optimal parallelism depends on the following items:

  - CPU resources for concurrent backup processes (each RMAN process typically uses all CPU cycles of a logical processor when compression and encryption are enabled)
  - Storage subsystem throughput for RMAN data file reads and writes.
  - Available memory.
  - Database size.

    In testing, parallelism of 60 was used, but 8 is a good starting point for most environments:

    `CONFIGURE DEVICE TYPE DISK PARALLELISM 8 BACKUP TYPE TO BACKUPSET;`

- Compression
  - Use MEDIUM compression for a balance of performance and size reduction.
  - HIGH compression is CPU-intensive and provides limited benefit.
  - Both require the Oracle Advanced Compression license:

    ```
    CONFIGURE COMPRESSION ALGORITHM 'MEDIUM' ;
    ```
- Section size

  The **SECTION SIZE** parameter limits the size of each backup piece (for example, 6000M). Limiting the size makes retransmission simpler if a file transfer fails.
- Encryption
  - Enable encryption and specify the algorithm and password, unless TDE is used. Otherwise, the session fails with a "wallet `not open`" error.
  - Example 3-12 shows you how to avoid this error.

    Example 3-12 Configuring encryption and avoiding Wallet Not Open

    ```
    CONFIGURE ENCRYPTION FOR DATABASE ON;
    CONFIGURE ENCRYPTION ALGORITHM 'AES256';
    SET ENCRYPTION ON IDENTIFIED BY passw0rd ONLY ;
    ```
- TAG usage

  Use a unique TAG for each backup to simplify management in the RMAN catalog.

  Important: Using a TAG simplifies management of multiple backups in an RMAN catalog. Incremental Level 0 specifies this backup as an incremental backup at Level 0, which means that all data is included in the backup. For more information, see Example 3-13.

  Example 3-13 Level 0 backup with PLUS ARCHIVELOG

  ```
  BACKUP  tag '<Your TAG here>' incremental level 0 AS compressed BACKUPSET
  section size 6000M DEVICE TYPE DISK DATABASE INCLUDE CURRENT CONTROLFILE PLUS
  ARCHIVELOG TAG
  <SID>_LEV0;
  ```

Incremental Level 1 differences

The incremental Level 1 script differs from the Level 0 script only in the following areas:

- The directory where the backup files are written.
- The TAG used to identify the backup set.
- Specifying Level 1 instead of Level 0 indicates that this backup is incremental, which by default includes all changes after the last Level 1 or Level 0 backup.

Final incremental backup preparation

The final incremental backup requires extra preparation to determine the date and timestamp for the final database recovery step on the target environment. The timestamp is the point to which the database is recovered. Any changes after that timestamp are discarded. To complete the migration by using RMAN Option 2, prepare for the final Level 1 incremental backup and determine the recovery point.

Inspecting the server time and date

Prepare the database environment by inspecting system time and configuring the Oracle instance for recovery.

After connecting to the database as `sysdba`, run the following commands:

1. Put the database in `MOUNT` mode and (not `OPEN`).
2. Set the Globalization Support date format to ensure correct `date-to-string` conversions for point-in-time recovery.
3. Force Oracle to write to a new redo log and capture the timestamp for recovery.

Final steps for option 2

Finish the final steps for RMAN option 2 to validate backups and prepare the parameter file:

1. Validate and cross-check backups by using the following RMAN commands:

```
RMAN> REPORT SCHEMA;
RMAN> LIST BACKUP SUMMARY;
RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
RMAN> VALIDATE DATAFILE 10;
RMAN> VALIDATE BACKUPSET 3;
RMAN> CROSSCHECK BACKUP;
```

2. Create a plain text parameter file (PFILE) from the SPFILE:

```
SQL> create pfile='/backup/rman/init<SID>.ora' from spfile;
```

3. Restore all RMAN parameters to their original settings after migration.
4. After completing the final incremental backup and validation, shut down the source database to prepare for migration. Figure 3-3 on page 55 shows this process.

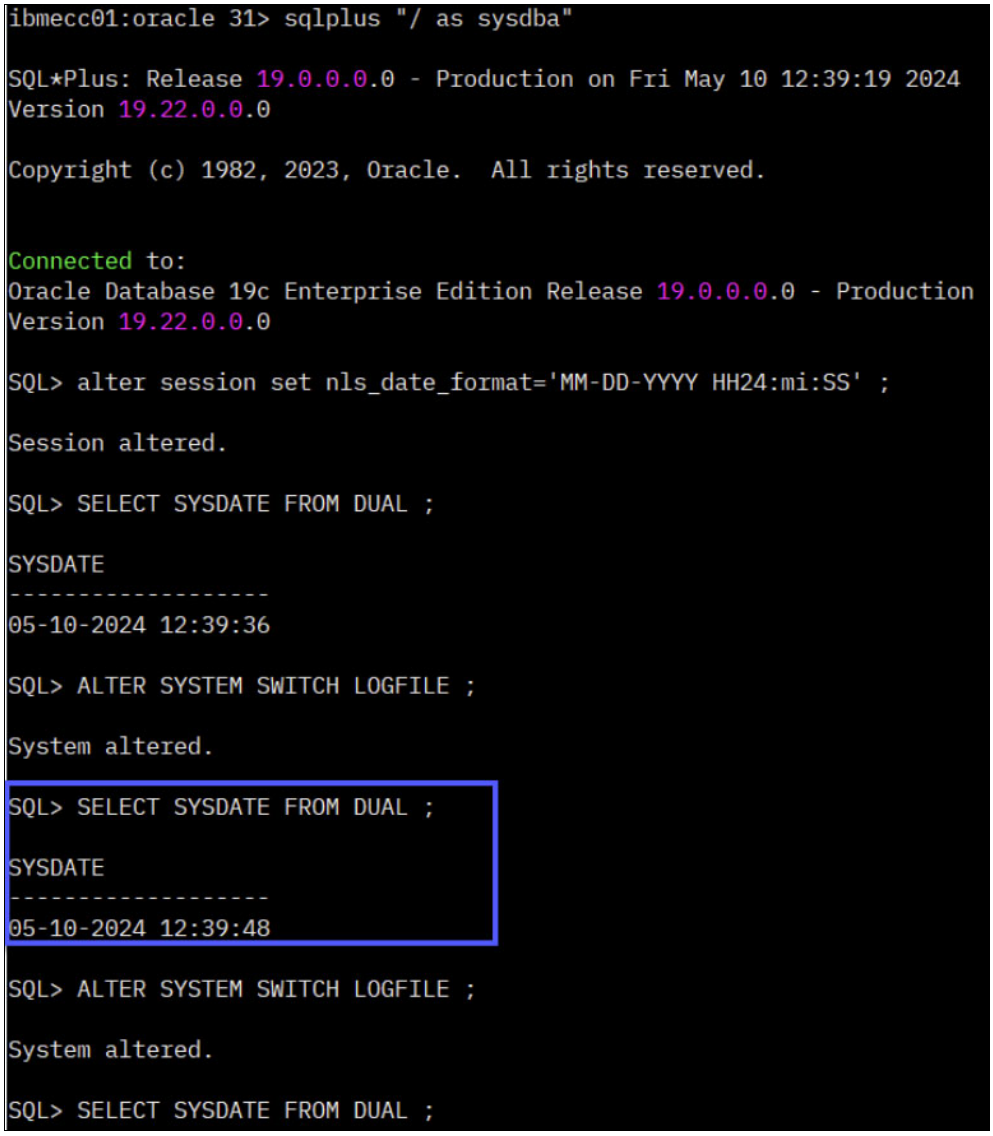

Figure 3-3 Incremental timestamp

5. Run the final incremental backup by using RMAN.

Validation and cross-check backups

When the backup finishes successfully, verify the integrity of the backup sets by using the RMAN commands that are shown in Example 3-14.

Example 3-14 RMAN command

```
RMAN> REPORT SCHEMA;
RMAN> LIST BACKUP SUMMARY;
RMAN> LIST BACKUPSET <TAG ID> ;
RMAN> BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;
RMAN> VALIDATE DATAFILE 10;
RMAN> VALIDATE BACKUPSET 3;
RMAN> CROSSCHECK BACKUP;
```

- **REPORT SCHEMA;**

  Lists database file and table spaces.
- **LIST BACKUP SUMMARY;**

  Lists all existing backups.
- **LIST BACKUP <TAG ID>;**

  List a specific backup by TAG.
- **BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVELOG ALL;**

  Validate backup contents.
- **VALIDATE DATAFILE 10;**

  Validates a specific data file.
- **VALIDATE BACKUPSET 3;**

  Validate a specific backup set.
- **CROSSCHECK BACKUP;**

  Synchronizes physical backups and copies with RMAN repository.

Creating a PFILE

Create a plain text database parameter file (PFILE) from the binary server parameter file (SPFILE). The database must remain mounted.

Run the following commands as user `oracle`:

```
setenv ORACLE_SID <SID>
sqlplus "/ as sysdba";
```

In SQL*Plus, create the parameter file:

```
SQL> create pfile='/backup/rman/init<SID>.ora' from spfile;
```

After migration, restore all RMAN parameters back to their original settings.

Shutdown and migration readiness

After completing the final incremental backup and validation, perform the following steps:

1. Shut down the source database to prepare for migration.
2. Stop using the source system.
3. For RMAN Option 2, create two backup sets (Level 0 and Level 1) to transfer to and restore on the target system.

3.3 Restoring the Oracle database on the target system

This section describes how to restore the Oracle database on the target system by using RMAN.

3.3.1 RMAN restore options

Choose one of these RMAN-based methods to restore the Oracle database:

- RMAN duplicate database
  - Uses the **`DUPLICATE`** command with a consistent backup set.
  - Automatically opens the database with the **`RESETLOGS`** command.

  Use this option only when a complete and consistent backup set is available on the target system.
- RMAN restore or recover database
  - Uses an initial Level 0 backup and applies incremental backups and archived redo logs.
  - Allows rolling forward to a specific point in time.

Both methods rely on the backup sets created in the previous RMAN backup procedures, as shown in Figure 3-4.

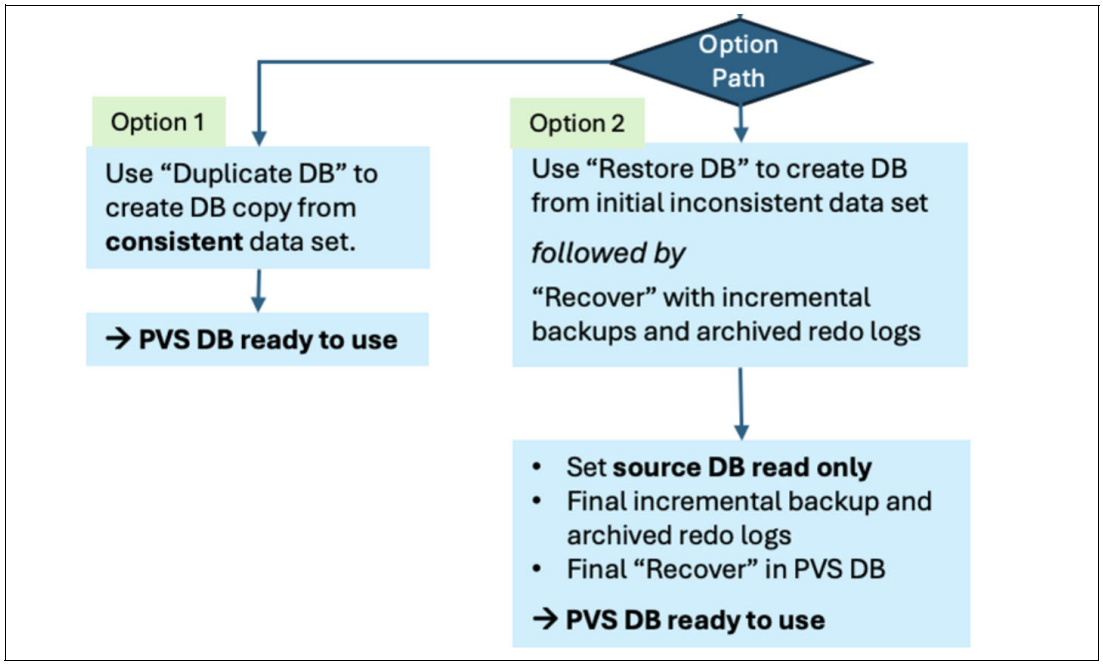

Figure 3-4 Restore options

3.3.2 Assumptions

Before starting, ensure that these conditions are met:

- Ensure that an IBM Cloud account and a PowerVS workspace exist.
- The target AIX LPAR matches the source system in compute, memory, and storage.

- Oracle Grid Infrastructure (if you use ASM) and Oracle RDBMS homes are installed on the target system with the same version and fix level as the source.
- All required backup files are accessible to RMAN through local storage, NFS, or IBM Cloud Object Storage.

3.3.3 Connectivity and preparation: Oracle on IBM PowerVS

This section explains how to prepare connectivity for Oracle on IBM PowerVS before performing a restore.

From the target instance CLI, follow these steps to perform a restore:

1. Log in as the Oracle user to the PowerVS instance.
2. Confirm these items:
   - The Oracle software version and fix level match the source environment.
   - Device names for Oracle data files and table spaces must match the source database. If they do not match, use SET NEWNAME directives (not covered in this book).
   - ASM disk groups (if used) exist and have sufficient capacity.
   - Backup files are available to RMAN.

3.3.4 Option 1: Creating an RMAN duplicate database

This option uses a consistent Level 0 backup to create a duplicate database.

Log in as the Oracle user, set `ORACLE_SID <SID>`, start SQL*Plus, and connect to the database as a system database administrator:

```
sqlplus "/ as sysdba"
```

Starting the target database in NOMOUNT mode

Set the environment, start the instance in `NOMOUNT` mode, and create an SPFILE, as shown in Example 3-15.

*Example 3-15 SQL*Plus command*

```
SQL> startup nomount pfile='/backup/rman/ora/initEC6.ora';
SQL> create spfile='/oracle/EC6/19.0.0/dbs/spfileEC6.ora' from
pfile='/backup/rman/ora/initEC6.ora';
SQL> shutdown immediate;
SQL> startup nomount;
SQL> show parameter spfile;
```

Verify the SPFILE location, as shown in Example 3-16.

Example 3-16 Parameter

```
NAME                 TYPE            VALUE
---------          -------------    ---------------------------------
spfile               string          /oracle/EC6/19/dbs/spfileEC6.ora
```

Creating the RMAN duplicate database: Option 1

Create an RMAN duplicate database script that references the appropriate backup location. Place the file (for example, `rman_duplicate.cmd`) in the Oracle user's home directory. Ideally, create the number of restore channels that match the number of channels that are used for the backup. If 60 channels are used during the backup, match that number in the script.

Note: This script shows the first 60 ALLOCATE AUXILIARY CHANNEL entries, one per line.

Create a script (for example, **`rman_duplicate.cmd`**) with the content that is shown in Example 3-17.

Example 3-17 RMAN duplicate database script: rman_duplicate.cmd

```
set encryption on identified by passw0rd;
set decryption identified by passw0rd;
run {
    allocate auxiliary channel ch1 device type disk;
    allocate auxiliary channel ch2 device type disk;
    allocate auxiliary channel ch3 device type disk;
    allocate auxiliary channel ch4 device type disk;
    allocate auxiliary channel ch5 device type disk;
    allocate auxiliary channel ch6 device type disk;
    allocate auxiliary channel ch7 device type disk;
    allocate auxiliary channel ch8 device type disk;
    allocate auxiliary channel ch9 device type disk;
    allocate auxiliary channel ch10 device type disk;
    allocate auxiliary channel ch11 device type disk;
    allocate auxiliary channel ch12 device type disk;
    allocate auxiliary channel ch13 device type disk;
    allocate auxiliary channel ch14 device type disk;
    allocate auxiliary channel ch15 device type disk;
    allocate auxiliary channel ch16 device type disk;
    allocate auxiliary channel ch17 device type disk;
    allocate auxiliary channel ch18 device type disk;
    allocate auxiliary channel ch19 device type disk;
    allocate auxiliary channel ch20 device type disk;
    allocate auxiliary channel ch21 device type disk;
    allocate auxiliary channel ch22 device type disk;
    allocate auxiliary channel ch23 device type disk;
    allocate auxiliary channel ch24 device type disk;
    allocate auxiliary channel ch25 device type disk;
    allocate auxiliary channel ch26 device type disk;
    allocate auxiliary channel ch27 device type disk;
    allocate auxiliary channel ch28 device type disk;
    allocate auxiliary channel ch29 device type disk;
    allocate auxiliary channel ch30 device type disk;
    allocate auxiliary channel ch31 device type disk;
    allocate auxiliary channel ch32 device type disk;
    allocate auxiliary channel ch33 device type disk;
    allocate auxiliary channel ch34 device type disk;
    allocate auxiliary channel ch35 device type disk;
    allocate auxiliary channel ch36 device type disk;
    allocate auxiliary channel ch37 device type disk;
    allocate auxiliary channel ch38 device type disk;
```

```
    allocate auxiliary channel ch39 device type disk;
    allocate auxiliary channel ch40 device type disk;
    allocate auxiliary channel ch41 device type disk;
    allocate auxiliary channel ch42 device type disk;
    allocate auxiliary channel ch43 device type disk;
    allocate auxiliary channel ch44 device type disk;
    allocate auxiliary channel ch45 device type disk;
    allocate auxiliary channel ch46 device type disk;
    allocate auxiliary channel ch47 device type disk;
    allocate auxiliary channel ch48 device type disk;
    allocate auxiliary channel ch49 device type disk;
    allocate auxiliary channel ch50 device type disk;
    allocate auxiliary channel ch51 device type disk;
    allocate auxiliary channel ch52 device type disk;
    allocate auxiliary channel ch53 device type disk;
    allocate auxiliary channel ch54 device type disk;
    allocate auxiliary channel ch55 device type disk;
    allocate auxiliary channel ch56 device type disk;
    allocate auxiliary channel ch57 device type disk;
    allocate auxiliary channel ch58 device type disk;
    allocate auxiliary channel ch59 device type disk;
    allocate auxiliary channel ch60 device type disk;
    duplicate database to <DBSID> backup location '/backup/rman/ec6_option1'
nofilenamecheck noredo;
}
```

Run the following script:

```
rman auxiliary / cmdfile=rman_duplicate.cmd
```

A typical script appears in Example 3-18.

Example 3-18 Running the script

```
sql clone "alter system set  db_name =
 ''EC6'' comment=
 ''Modified by RMAN duplicate'' scope=spfile";
   sql clone "alter system set  db_unique_name =
 ''EC6'' comment=
 ''Modified by RMAN duplicate'' scope=spfile";
   shutdown clone immediate;
   startup clone force nomount
   restore clone primary controlfile from
'/backup/rman/ec6_option1/option1_cf_c-2254911489-20240516-00';
   alter clone database mount;
```

RMAN clones the database successfully and opens it with **RESETLOGS**. The database is now open, and RMAN confirms that the duplication finished successfully.

Important:

- Add enough **ALLOCATE AUXILIARY CHANNEL** statements to match the number of channels that are used during backup.
- Use **NOFILENAMECHECK** to skip directory validation when source and target file paths are identical.
- Use **NOREDO** to indicate that redo logs are not applied because this backup is consistent.

3.3.5 Checking the archive log status

This section explains how to verify the archive log status after restoring the database.

Complete this step before continuing with the recovery process:

```
sqlplus "/ as sysdba"
SQL> archive log list;
```

The expected output appears in Example 3-19.

Example 3-19 Expected output

```
Database log mode              No Archive Mode
Automatic archival             Disabled
Archive destination            /oracle/EC6/oraarch/EC6arch
Oldest online log sequence     1
Current log sequence           1
```

3.3.6 Restarting the database in archive mode

This section explains how to restart the database and enable ARCHIVELOG mode on the target database.

Complete the following steps before continuing with the backup procedure:

```
SQL> SHUTDOWN IMMEDIATE;
SQL> STARTUP MOUNT;
SQL> ALTER DATABASE ARCHIVELOG;
SQL> ALTER DATABASE OPEN;
SQL> ARCHIVE LOG LIST;
```

The following options control RMAN behavior during restore operations:

- **NOFILENAMECHECK** prevents RMAN from validating directory paths when source and target file paths are identical.
- **NOREDO** tells RMAN not to check for redo logs because a full offline backup is being restored.

Finally, create a database backup on the target system that includes the archive files, as described in Back up the Source Oracle Database by using RMAN.

3.3.7 Option 2: RMAN restore and recover database

This option restores the database from an initial Level 0 backup and applies incremental backups and archived redo logs to roll forward to a specific point in time.

Starting the target database in NOMOUNT mode

Start the environment by running the **setenv ORACLE_SID <SID>** command, and start the Oracle instance in NOMOUNT mode directly, specifying the PFILE.

Note: In this example, Oracle runs in an SAP environment, so ORACLE_HOME was /oracle/EC9/19.0.0. The environment configuration might vary. Pay close attention when copying and pasting commands, and replace the Oracle <SID> with the system SID in all commands that are shown in Example 3-20.

Example 3-20 SQL commands

```
SQL> startup nomount pfile='/backup/rman/ora/initEC6.ora';
SQL> create spfile='/oracle/EC6/19.0.0/dbs/spfileEC6.ora' from
pfile='/backup/rman/ora/initEC6.ora';
SQL> shutdown immediate;
SQL> startup nomount;
SQL> show parameter spfile;
```

The output in Example 3-21 shows the parameter.

Example 3-21 SQL command output

```
NAME                 TYPE            VALUE
---------           -------------    --------------------------------
spfile               string          /oracle/EC6/19/dbs/spfileEC6.ora
```

Determining the control file path in the backup file set

The control file (CF) was included in the full backup (Level 0). Determine the full path of the control file in the Level 0 backup set shown in Example 3-22.

Example 3-22 Backup command

```
find /backup/rman/ec6_option2 -name "lev0_cf_*"
```

The find command prints all file names, including the path that match the name pattern:

```
/backup/rman/ec6_option2/lev0_cf_c-2252531432-20240529-01
```

When output is missing, see the backup that is created with the Option 2 Level 0 procedure for the control file location and verify that the backup file set is on the target server.

Creating the RMAN restore script

Create and run the RMAN restore script (for example o**ption2_restore_level0.rman**), as shown in Example 3-23.

Example 3-23 Restore script

```
connect target /
connect target /
set ENCRYPTION ALGORITHM 'AES256' ;
SET DECRYPTION IDENTIFIED BY passw0rd;
run {
  restore controlfile from
'/backup/rman/ec6_option2/lev0_cf_c-2252531432-20240529-01';
  alter database mount;
  restore database from tag IBMECCO2_EC6_LEV0;
```

```
  recover database;
}
```

Run the **option2_restore_level0.rman** script.

Explanation of the restore script (before running the script)

Review these considerations before running the script:

- Use the backup TAG to specify which backup to restore.
- Start the database in NOMOUNT mode before running the script.
- Adjust the control file location to match the backup directory
- If the backup files are encrypted, RMAN requires the encryption password to restore the database files.

What the script does

This section describes the actions that the script performs:

- Restores the control file that contains the RMAN catalog.
- Runs **RESTORE DATABASE FROM TAG** command to restore the database from the specified backup tag.
- Runs the **RECOVER DATABASE** command to apply archived redo logs included in the backup.

Restoring the incremental backups (Level 1)

After the Level 0 backup is restored successfully, restore the incremental Level 1 backup sets.

Restoring the incremental backups (Level 1) except for the last one

Catalog and apply incremental backups (**RECOVER DATABASE**), in the target directory. If multiple incremental backups are provided, repeat for all incremental backups except the last one. For the final incremental backup, use point-in-time recovery:

```
RMAN> alter session set nls_date_format='DD-MM-YYYY HH24:MI:SS';
RMAN> recover database until time '13-01-2024 11:52:12';
```

Cataloging the final incremental backup (Level 1)

The final incremental backup requires recovery only to a specific point in time to allow the database to open.

This command assumes that all incremental backups are stored in directories under `/backup/rman`. From the RMAN prompt, run:

```
RMAN> catalog start with '/backup/rman';
```

Important: The target database can remain in MOUNT state and have incremental Level 1 backups and archived `redo` logs continuously, as needed. Catalog new backup pieces and archived logs, and recover until ready for final cutover.

Example 3-24 shows the backup files from the first incremental Level 1 backup.

Example 3-24 Backup catalog output

```
searching for all files that match the pattern /backup/rman
List of Files Unknown to the Database
=====================================
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dk2s0rg6_1460_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dl2s0rgc_1461_1_1
```

```
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dl2s0rgc_1461_2_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dm2s0rgc_1462_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dn2s0rgc_1463_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_do2s0rgc_1464_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dp2s0rgc_1465_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dq2s0rge_1466_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_dr2s0rgh_1467_1_1
File Name: /backup/rman/ec6_option2_inc1/inc1_EC6_ds2s0rgn_1468_1_1
       ...Lines Omitted ...
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_eo2s0rlq_1496_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_ep2s0rlu_1497_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_eq2s0rm3_1498_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_er2s0rm6_1499_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_es2s0rma_1500_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_EC6_eu2s0rmi_1502_1_1
File Name: /backup/rman/ibmecc02/ec6_option2_inc1/inc1_cf_c-2252531432-20240529-02

Do you really want to catalog the above files (enter YES or NO)? YES
cataloging files...
cataloging done
```

Restoring the final incremental backup (Level 1) to a point in time

For all incremental backups except the final one, run the following command in RMAN:

```
RMAN> RECOVER DATABASE;
```

Because the backup was created with encryption, configure RMAN to decrypt the backup by using the password; otherwise errors such as `ORA-19913: unable to decrypt backup` and `ORA-28365: wallet is not open` appear, as shown in Example 3-25.

Example 3-25 Backup encryption

```
RMAN> SET ENCRYPTION ALGORITHM 'AES256' ;
RMAN> SET DECRYPTION IDENTIFIED BY passw0rd;
RMAN> recover database ;
```

After cataloging the final incremental backup pieces, recover the database to the point in time identified during the backup process. After this final recovery action, do not apply further changes from a recovery perspective.

Set the time and date format:

```
RMAN> alter session set nls_date_format='DD-MM-YYYY HH24:mi:SS' ;
```

Recover the database to a specific time:

```
RMAN> recover database until time '13-01-2024 11:52:12';
```

Opening the database with RESETLOGS

After applying all backups and archived logs, run the following **`sqlplus`** commands as user `oracle`:

```
sqlplus "/ as sysdba"
SQL> alter database open RESETLOGS;
```

Checking the database mode

Verify that ARCHIVELOG mode is enabled by running the following steps:

1. Log in as user `oracle` and run the **`sqlplus`** command:

   ```
   sqlplus "/ as sysdba"
   ```

2. Check the archive log list:

   ```
   SQL> archive log list;
   ```

3. Create a plain text parameter file for reference:

   ```
   SQL> create pfile='/backup/rman/init<SID>.ora' from spfile;
   ```

Final note

After completing the restore and recovery process, shut down the source database permanently to help ensure consistency. This step concludes the restore recover approach to database migration to PowerVS.

4

Migrating SAP ERP 6 on IBM Db2 to IBM Power Virtual Server

This scenario focuses on migrating SAP ERP 6 that runs on IBM Db2 to IBM Power Virtual Server (PowerVS).

The following migration options are available:

- A simple backup and restore method
- A more advanced approach that uses IBM Db2 High Availability and Disaster Recovery (HADR)

This chapter describes both methods and provides resource links to supporting resources.

This chapter describes the following topics:

- “Overview”
- “Migration option 1: Backup and restore”
- “Migration option 2: IBM Db2 HADR”

© Copyright IBM Corp. 2025.

4.1 Overview

This book outlines the process for migrating an SAP ERP 6.0 system running on IBM Db2 to an IBM Power Virtual Server (PowerVS) environment. Two primary options are available for migrating IBM Db2 databases to the target system. Each option offers distinct advantages and considerations.

Migration options include the following approaches:

- Backup and restore
  - Uses standard administrative procedures such as stopping the system, creating a database backup, and restoring it on the target environment.
  - Suitable for non-production or less time-sensitive systems.
  - Requires full system downtime during the migration.
- IBM Db2 HADR
  - Minimizes downtime, making it ideal for migrating SAP production systems.
  - Uses IBM Db2 HADR technology to synchronize the source and target databases in real time and combines this process with an online backup.
  - Requires more advanced configuration and setup than the backup and restore method.

4.2 Migration option 1: Backup and restore

Back up and restore is a common method for migrating development or test SAP systems. This option uses a backup of the database from the source system and restores it on a preinstalled target system.

Finish the migration by using these steps:

1. Finish the migration by using the steps described in 4.2.1, “Preparing for migration” on page 69.
2. Shut down the SAP application and deactivate the database on the source system.
3. Use the IBM Db2 offline backup on the source system.
4. Transfer the backup files.
5. Restore the database on the target system.
6. Start the SAP target system.

Review these characteristics of the backup and restore approach before starting the migration:

- Downtime

  The SAP system is unavailable during migration, so downtime must be planned.
- Fallback

  The source SAP system remains unchanged. If needed, the migration can be stopped and the source system restarted.
- Complexity

  This method uses standard administrative tasks and is the least complex option.

4.2.1 Preparing for migration

This section explains how to prepare the source system for migration.

To prepare for migration, complete the following steps:

1. Start a C shell session as the `root` user on the source system:

   ```
   $ csh
   ```

   The syntax for defining environment variables depends on the shell type. The default shell type for Db2 and SAP administrator users `db2<sid>` and `<sid>adm`) is the C shell. When using the same shell, copy the command examples and paste them into the session.

2. Define the hostname of the target migration system. Replace the example value with the system-specific value:

   ```
   set TARGETSERVER=cdb6ecc1
   ```

   Fast path: Save these environment variable commands to a temporary text file and reuse them on the target server.

3. Define the SAP instance administrator. Replace the example with the system-specific value:

   ```
   set SIDADM=th1adm
   ```

4. Define the IBM Db2 database name. Replace the example value with the system-specific value:

   ```
   set DBNAME=th1
   ```

5. Define the IBM Db2 database administrator:

   ```
   set DB2ADM=db2th1
   ```

6. Define a directory to store the backup files:

   ```
   set BACKUPDIR=/db2/backup
   ```

7. Create the backup directory and assign ownership:

   ```
   mkdir -p $BACKUPDIR
   chown $DB2ADM $BACKUPDIR
   ```

Requirement: Ensure that the source and target SAP systems run the same versions of SAP ERP software and IBM Db2 database.

Tip: Use **`GET_DBSIZE_INFO`** to determine the required backup space. For more information, see the GET_DBSIZE_INFO procedure.

4.2.2 Shutting down the SAP application and deactivating the database

This section describes how to shut down the SAP application and deactivate the database on the source server. These steps are typically performed during system maintenance or migration procedures.

Stopping the SAP application

To stop the SAP application and keep the database active:

1. Log in as the SAP instance administrator (`$SIDADM`) by running the following command:

   ```
   su - $SIDADM
   ```

2. Stop the SAP system:

   ```
   stopsap r3
   ```

   a. Optionally, check whether SAP is in the required state by running:

   ```
   stopsap check
   ```

   The expected output appears in Example 4-1.

   Example 4-1 Output from the stopsap check command

   ```
   Checking db6 Database
   Database is running
   -----------------------------------------
   Checking SAP TH1 Instance ASCS01
   -----------------------------------------
   Instance ASCS01 is not running
   Checking SAP TH1 Instance D00
   -----------------------------------------
   Instance D00 is not running
   ```

   b. Exit the `$SIDADM` user SAP session:

   ```
   exit
   ```

Deactivating the database

To deactivate the database, complete the following steps:

1. Switch to the Db2 (`$DB2ADM`) user by running the following command:

   ```
   su - $DB2ADM
   ```

2. Optionally, check for active applications by running the following command:

   ```
   db2 list applications
   ```

   The expected output appears in the following example:

   ```
   SQL1611W  No data was returned by Database System Monitor.
   ```

3. Set the environment variables for the Db2 (`$DB2ADM`) user by running the following command:

   ```
   set DBNAME=th1
   set BACKUPDIR=/db2/backup
   ```

4. Deactivate the database by running the following command:

   ```
   db2 deactivate database $DBNAME
   ```

   The expected output appears in the following example:

   ```
   DB20000I The DEACTIVATE DATABASE command completed successfully.
   ```

Tip: If applications are still listed, stop the applications and check again. If external applications do not stop, disconnect them from the database server by running the following command:

```
db2 force applications all
```

4.2.3 Using the IBM Db2 offline backup on the source system

This section describes how to create an offline backup of the IBM Db2 database on the source system. The backup directory was created in 4.2.1, “Preparing for migration” on page 69.

Creating the offline backup

1. Verify that the backup directory has enough space to store the backup files:

   ```
   df -m $BACKUPDIR
   ```

2. Start the IBM Db2 backup:

   ```
   db2 backup database $DBNAME to $BACKUPDIR compress
   ```

   This command can take a long time, depending on the database size.

   The expected output appears in the following example:

   ```
   Backup successful. The timestamp for this backup image is : 20240730170709
   ```

3. Store the timestamp for use in subsequent steps. Save it in an environment variable:

   ```
   set TIMESTAMP=<your time stamp>
   ```

   Note: This timestamp appears at the end of the backup command output. It is required for restoring the database on the target system.

4. Change to the backup folder by running the following command:

   ```
   cd $BACKUPDIR
   ```

5. Check for the backup files that contain the timestamp, as shown in Example 4-2:

   ```
   ls -l *$TIMESTAMP*
   ```

   Example 4-2 Sample backup file name that is generated by Db2

   ```
   TH1.0.db2th1.DBPART000.20240730170709.001
   ```

   Files that are listed in the output are transferred to the target system in the next step.

6. (Optional) To track the backup progress, run the following command:

   ```
   db2 list utilities show detail
   ```

   Fast path: The IBM Db2 backup time `TIMESTAMP` has the format `YYYYMMDDHHMMSS` (Year-Month-Day-Hour-Minute-Second), for example `20240730170709`. Append this `TIMESTAMP` definition to the list of saved commands from 4.2.1, “Preparing for migration” on page 69.

4.2.4 Transferring the backup files

This section explains how to transfer backup files from the source SAP system to the target system.

Perform this step before starting the restore process:

```
scp ${BACKUPDIR}/*${TIMESTAMP}*
${TARGETSERVER}:${BACKUPDIR}
```

This example uses secure copy protocol (SCP), which is slower than other data transfer options. Transfer backup files directly to an IBM AIX LPAR in PowerVS or to IBM Cloud Object Storage.

If SCP or SFTP is used with IBM Cloud Object Storage, ensure that an IBM FileManage Gateway service is available or install and configure an SFTP server within or next to the target PowerVS environment to receive the transfer.

A faster option is the IBM high-performance Aspera product for data transfer. In many situations, IBM Aspera can transfer data several times faster than traditional TCP-based protocols. For more information, see IBM Aspera Technologies. This reference also contains the Accelerated network transfer migration guide.

4.2.5 Restoring the database on the target system

After transferring the backup files, restore the database on the target system to finish the migration process. This procedure helps ensure that the target environment is properly configured and ready for use.

Perform the following steps to restore a database on the target system:

1. Log in to the target server by running the following command:

   ```
   ssh $DB2ADM@$TARGETSERVER
   ```

   Requirement: The target server does not recognize the environment variables that are defined on the source system. Define `$DBNAME`, `$TIMESTAMP`, and `$SIDADM` on the target system. To define these variables, run the list of commands that were saved in 4.2.1, “Preparing for migration” on page 69 and 4.2.3, “Using the IBM Db2 offline backup on the source system” on page 71, or define the environment variables again.

2. Verify that the system does not have active application connections:

   ```
   db2 list applications
   ```

   The expected output appears in the following example:

   ```
   SQL1611W  No data was returned by Database System Monitor.
   ```

   Tip: Ensure that the IBM Db2 database is deactivated before starting the full restore procedure.

3. Deactivate the IBM Db2 database:

   ```
   db2 deactivate database $DBNAME
   ```

4. Restore the database by using the backup image (see Example 4-3).

   Example 4-3 Database restore command

   ```
   db2 restore database $DBNAME
            from $BACKUPDIR
            taken at $TIMESTAMP
   ```

   Tip: Append **`replace existing`** to the **`restore`** command to avoid the overwrite prompt.

5. When the restore command runs, Db2 prompts for confirmation to overwrite the existing database because the target database has the same name as the database in the backup image but contains different data. The expected output appears in Example 4-4.

Example 4-4 Command overwrite exit db

```
SQL2523W  Warning!  Restoring to an existing database that is different from
the database on the backup image, but have matching names. The target database
is overwritten by the backup version. The rollforward recovery logs associated
with the target database is deleted.
Do you want to continue ? (y/n)
```

Enter `Y` and press Enter.

The expected output appears in the following example:

```
DB20000I The RESTORE DATABASE command completed successfully.
```

6. Complete the restore by running this command:

```
db2 rollforward db $DBNAME to end of backup and stop
```

Tip: If the restore fails, drop the database on the target system in advance by using the following command:

```
db2 drop database $DBNAME
```

Then repeat the **restore** command.

The target database now contains the source system data.

4.2.6 Starting the SAP target system

After restoring the database, start the SAP target system to finish the migration process.

Perform the following steps:

1. Switch to the SAP administrator by running the following command:

```
login $SIDADM
```

2. Start SAP by running the following command:

```
startsap
```

The expected output is similar to Example 4-5 and includes messages that confirm the database is running and that SAP instances are starting.

Example 4-5 The startsap command output

```
Checking db6 Database
Database is running
-----------------------------------------
Starting Startup Agent sapstartsrv
OK
Instance Service on host <TARGET> started
-----------------------------------------
starting SAP Instance ASCS01
Startup-Log is written to /home/th1adm/startsap_ASCS01.log
-----------------------------------------
/usr/sap/TH1/ASCS01/exe/sapcontrol -prot NI_HTTP -nr 01 -function Start
Instance on host <TARGET> started
```

```
Starting Startup Agent sapstartsrv
OK
Instance Service on host <TARGET> started
-----------------------------------------
starting SAP Instance D00
Startup-Log is written to /home/th1adm/startsap_D00.log
-----------------------------------------
/usr/sap/TH1/D00/exe/sapcontrol -prot NI_HTTP -nr 00 -function Start
Instance on host <TARGET> started
```

Important: Adjust the SAP system DNS record to point to the target SAP server to help ensure that client systems connect to the target server.

The migration is complete:

- The SAP system on the source server is down.
- The SAP system on the target server is up and running.

4.3 Migration option 2: IBM Db2 HADR

Migration option 2 reduces downtime by combining backup, recovery, and database synchronization with IBM Db2 HADR. After the databases are synchronized and the environment is ready for migration, switch the core migration to the preinstalled target SAP server.

If the source system uses an IBM Db2 cluster or Db2 HADR is enabled, extend the existing IBM Db2 HADR configuration. For more information, see 4.3.1, “Special case: Migrating a source IBM Db2 cluster” on page 75.

If HADR is not enabled on the source system, complete the following steps:

1. Prepare for migration.
2. Ensure that IBM Db2 archive logging is enabled.
3. Create an online backup.
4. Transfer backup files to the target system.
5. Restore the target system.
6. Define HADR local and remote service ports on the target and source systems.
7. Configure HADR on the target and source systems.
8. Start HADR and verify data synchronizing.
9. Run the core migration.

Here are the characteristics of migration option 2: Db2 HADR:

- Downtime (general)

  A short downtime occurs when switching from the source to the target server. Before the switch, the source SAP system remains operational. After the switch, the target SAP system takes over.

- Downtime (circular logging)

  HADR requires IBM Db2 to run in archive logging mode. If the IBM Db2 database uses circular logging, migrate it to archive logging mode before enabling HADR. This step requires a one-time database downtime.

- Fallback

  The source SAP system remains operational until the switch. If steps before or after the switch fail, reverse them to bring the source SAP system back online.

4.3.1 Special case: Migrating a source IBM Db2 cluster

Extend the HADR configuration to include the target server. Add the target system as an auxiliary node to the existing IBM Db2 HADR configuration.

For more information, see Steps for adding a new Auxiliary Standby to an existing Db2 HADR Pair.

4.3.2 Preparing for migration

Before starting the migration process, configure the SAP system and supporting components to help ensure a smooth migration. These steps help establish network connectivity, security, and environment settings required for HADR and database restoration.

Preparing the SAP system for migration

Complete the following steps to prepare the SAP system for migration:

1. Define two TCP port numbers for the setup. HADR requires two open network ports:
   - A local port for outgoing traffic
   - A remote port for incoming traffic
2. Configure both ports on each IBM Db2 systems, but in reverse order. This example uses ports `5920/tcp` and `5921/tcp`, but different ports can be selected if needed.
3. Ensure that firewalls allow TCP traffic between the source and target systems.
4. Configure the required firewall rules on all involved firewalls on both systems.

Setting the environment variables

Set the environment variables according to the configuration by completing the following steps.

Fast path: Save these commands in a temporary text file and reuse them on the source SAP server. Use the same commands on the target server, as described in Restoring the database on the target system.

Requirement: Ensure that the target server runs the same version of SAP ERP software and IBM Db2 database as the source system.

1. Define the hostname for the target migration system. Replace the example value with the target hostname:

   ```
   set TARGETSERVER=cdb6ecc1
   ```

2. Define the SAP instance administrator. Replace the example value with the system value:

   ```
   set SIDADM=th1adm
   ```

3. Define the IBM Db2 database name with the target hostname. The value `th1` is an example. Change it to match the system:

   ```
   set DBNAME=th1
   ```

4. Define the IBM Db2 database administrator:

   ```
   set DB2ADM=db2th1
   ```

5. Define a directory to store backup files:

   ```
   set BACKUPDIR=/db2/backup
   ```

Tip: This backup directory needs enough space to store the compressed backup files. Determine the current IBM Db2 database size by calling the `GET_DBSIZE_INFO` procedure. For more information, see the GET_DBSIZE_INFO procedure.

6. Create the directory if it does not exist by running the following command:

   ```
   mkdir -p $BACKUPDIR
   ```

7. Transfer the ownership to the Db2 administrator:

   ```
   chown $DB2ADM $BACKUPDIR
   ```

4.3.3 Ensuring that archive logging is enabled

HADR requires archive logging to be enabled. By default, IBM Db2 uses circular logging after installing HADR. Most SAP systems on IBM Db2 already have archive logging enabled.

Checking the archive logging method

Perform these steps to verify the logging configuration:

1. Switch to the IBM Db2 database administrator by running the following command:

   ```
   su - $DB2ADM
   ```

2. Retrieve both log archive method configuration options by running the following command:

   ```
   db2 get db cfg for th2 | grep LOGARCHMET
   ```

3. Continue to the next section if at least one log archive method is set to `ON`. Ensure that both are not set to `OFF`. The expected output for enabled archive logging appears in the following example:

```
First log archive method                (LOGARCHMETH1) = DISK:/db2/logarch/
Second log archive method               (LOGARCHMETH2) = OFF
```

4. If both log archive methods are `OFF`, configure archive logging as shown in the following lines:

```
First log archive method                (LOGARCHMETH1) = OFF
Second log archive method               (LOGARCHMETH2) = OFF
```

Configuring the archive logging

HADR requires archive logging. If both **LOGARCHMETH1** and **LOGARCHMETH2** are set to `OFF`, complete the following steps to enable archive logging.

Schedule downtime to complete this change, then perform the following steps:

1. Log in as the SAP administrator by running the following command:

   ```
   $ su - $SIDADM
   ```

2. Stop the SAP system, but keep the database running by running the following command:

   ```
   $ stopsap r3
   ```

3. Exit the SAP administrative session by running the following command:

   ```
   $ exit
   ```

4. Switch to the IBM Db2 administrator account by running the following command:

   ```
   $ su - $DB2ADM
   ```

5. Ensure that environment variables such as **DBNAME** and **BACKUPDIR** are set.
6. Define a directory for log archive files by running the following command:

   ```
   $ set LOGARCHDIR=/db2/log_archive
   ```

7. Create the log archive directory by running the following command:

   ```
   $ mkdir $LOGARCHDIR
   ```

8. Verify the directory and permissions by running the following command:

   ```
   $ ls -ld $LOGARCHDIR
   ```

   Ensure that the Db2 administrator owns the directory and has full permissions (`rwx`).

   The sample output appears in the following example:

   ```
   drwxr-xr-x 7 db2th1 dbth1adm 256 Aug 30 11:24 /db2/log_archive
   ```

9. Configure Db2 to use this directory for archive logging by running the following command:

   ```
   db2 "update db cfg for $DBNAME using LOGARCHMETH1 DISK:$LOGARCHDIR"
   ```

   The expected output appears in the following example:

   ```
   DB20000I The UPDATE DATABASE CONFIGURATION command completed successfully.
   ```

10. After this change, perform a backup by running the following command:

    ```
    db2 backup database $DBNAME to $BACKUPDIR compress
    ```

    The time that it takes the backup to finish depends on the database size.

11. Restart the SAP system by running the following command:

    ```
    su - $SIDADM
    startsap
    ```

4.3.4 Creating an online backup

This section explains how to create an online Db2 database backup to minimize SAP system downtime.

Important: If possible, create the log archive directory on a separate partition. The archive log partition depends on the LPAR disk setup, which is outside the scope of this document. If the source system is an IBM Db2 cluster, ensure that the log archive directory is a shared directory between nodes.

Complete the following steps before starting the backup process:

1. Switch to the Db2 admin user by running the following command:

   ```
   su - $DB2ADM
   ```

2. Start the online backup, including log files by running the following command:

   ```
   db2 backup db $DBNAME online to $BACKUPDIR compress include logs
   ```

   The backup takes time to finish. When the backup finishes, Db2 prints a timestamp:

   ```
   Backup successful. The timestamp for this backup image is : 20240913091058
   ```

3. Record the timestamp. It is required in the next two steps. Store it as an environment variable by running the following command:

   ```
   set TIMESTAMP=<your time stamp>
   ```

Attention: The IBM Db2 backup timestamp uses the format `YYYYMMDDHHMMSS` (Year-Month-Day-Hour-Minute-Second) and looks like `20240913091058`. This timestamp is needed on the target SAP server again. Append this `TIMESTAMP` definition command to the list of saved commands from 4.2.1, “Preparing for migration” on page 69.

4.3.5 Transferring backup files to the target system

To transfer backup files from the source SAP system to the target system, use a secure copy protocol (SCP) command. This method ensures encrypted file transfer over SSH.

Run the following SCP command:

```
scp ${BACKUPDIR}/*${TIMESTAMP}*
${TARGETSERVER}:${BACKUPDIR}
```

This command uses SCP. Ensure that the target system is accessible and that the required permissions are set.

Tip: For faster transfers, consider using IBM Aspera, which provides high-performance data transfer capabilities compared to traditional TCP-based protocols.

4.3.6 Restoring the target system

To restore the target system, complete the following steps:

1. Log in to the target server by running the following command:

   ```
   ssh $DB2ADM@$TARGETSERVER
   ```

2. Restore the database by using the online backup (see Example 4-6).

   Example 4-6 db2 restore command

   ```
   db2 restore database $DBNAME from $BACKUPDIR taken at $TIMESTAMP
   ```

   Db2 prompts for confirmation before it overwrites the existing database.

   Attention: The IBM Db2 backup timestamp uses the format `YYYYMMDDHHMMSS` (Year-Month-Day-Hour-Minute-Second), for example, `20240913091058`. Append this `TIMESTAMP` definition to the list of saved commands in 4.2.1, “Preparing for migration” on page 69.

3. Enter `Y` and press Enter to confirm. The expected output appears in the following example:

   ```
   SQL2523W  Warning!  Restoring to an existing database that is different from
   the database on the backup image, but have matching names. The target database
   will be overwritten by the backup version. The rollforward recovery logs
   associated with the target database will be deleted. Do you want to continue?
   (y/n)
   ```

   In contrast to option 1, do not roll forward archive logs. HADR takes this state to synchronize changed data in the next step. Because the source system is operational, source data changes. To transfer the most recent changes, configure IBM Db2 HADR database synchronization.

Restriction: The target server does not recognize the environment variables that are defined on the source system. Define `$DBNAME`, `$TIMESTAMP`, and `$SIDADM` on the target system. To do this task, run the saved commands from 4.2.1, “Preparing for migration” on page 69 and 4.3.4, “Creating an online backup” on page 78, or define the environment variables again.

If the backup completes successfully, you should see the following messages:

```
Backup status: The backup finished successfully. The timestamp for this backup
image is 20240913091058:
DB20000I  The RESTORE DATABASE command completed successfully.
```

Tip: If the restore fails, drop the database on the target system by using the following command:

```
db2 drop database $DBNAME
```

Then, repeat the **`restore`** command.

4.3.7 Defining HADR local and remote service ports

This section explains how to define high availability disaster recovery (HADR) local and remote service ports on the target and source systems.

Complete the following steps before configuring HADR:

1. Verify that the `/etc/services` file does not contain these TCP port numbers on either system. Then, configure the port numbers in this file on both servers.

   In the example, ports `5920/tcp` and `5921/tcp` are used. However, different ports can be defined if needed.

   Tip: Configuring local and remote HADR ports in `/etc/services` improves readability of the IBM Db2 configuration because service names can be used instead of port numbers.

   Table 4-1 shows the TCP port assignments in `/etc/services` on both servers.

Table 4-1 Table 4-1 TCP port assignments in /etc/services on both servers

Server name	Value on source server	Value on target server	Comment
db2th1ha_l	5921/tcp	5920/tcp	Local port
db2th1ha_r	5920/tcp	5921/tcp	Remote port

2. Use a preferred text editor to update `/etc/services` on both servers.
3. Verify the configuration by running the following command:

   ```
   grep -e '592[01]/' /etc/services
   ```

   The expected output on the source server appears in the following example:

   ```
   db2th1ha_l        5921/tcp        # Db2 HADR local port
   db2th1ha_r        5920/tcp        # Db2 HADR remote port
   ```

   The expected output on the target server appears in the following example:

   ```
   db2th1ha_l        5920/tcp        # Db2 HADR local port
   db2th1ha_r        5921/tcp        # Db2 HADR remote port
   ```

4.3.8 Configuring HADR on the target and source system

Configure the HADR parameters in Table 4-2 on both systems. Use identical local and remote service entries (`HADR_LOCAL_SVC` and `HADR_REMOTE_SVC`) because the switch is configured in `/etc/services`.

Table 4-2 HADR parameter overview on both servers

Db2 HADR parameter	Value on source server	Value on target server	Comment
HADR_LOCAL_HOST	<source host name>	<target host name>	Hostname of the system where the command is run.
HADR_LOCAL_SVC	db2th1ha_l	db2th1ha_l	Service names that are defined in /etc/services.
HADR_REMOTE_HOST	<target host name>	<source host name>	Hostname of the other system.

Db2 HADR parameter	Value on source server	Value on target server	Comment
HADR_REMOTE_SVC	db2th1ha_r	db2th1ha_r	Remote port as defined in /etc/services.
HADR_REMOTE_INST	<db2 instance name>	<db2 instance name>	Db2 instance name on the other system (not the database name).
LOGINDEXBUILD	ON	ON	Set to ON on for both systems.
HADR_SYNCMODE	<a valid sync mode>	<a valid sync mode>	For more information, see HADR Synchronization Mode.

Configure the local and remote hostnames (HADR_LOCAL_HOST and HADR_REMOTE_HOST) on both systems.

Note: High availability disaster recovery (HADR) synchronization mode explains different IBM Db2 HADR synchronization options and their benefits.

Use these commands to configure HADR parameters on both systems:

1. Set the local host by running the following command:

   ```
   db2 "update db cfg for $DBNAME using HADR_LOCAL_HOST `hostname`"
   ```

2. Set the local service by running the following command:

   ```
   db2 "update db cfg for $DBNAME using HADR_LOCAL_SVC db2th1ha_l"
   ```

3. Set the remote host. Verify that $TARGETSERVER points to the other server:

   ```
   db2 "update db cfg for $DBNAME using HADR_REMOTE_HOST $TARGETSERVER"
   ```

4. Set the remote service by running the following command:

   ```
   db2 "update db cfg for $DBNAME using HADR_REMOTE_SVC db2th1ha_r"
   ```

5. Configure the remote IBM Db2 instance by running the following command:

   ```
   db2 "update db cfg for $DBNAME using HADR_REMOTE_INST $DBINST"
   ```

6. Enable the log index build on both systems by running the following command:

   ```
   db2 "update db cfg for $DBNAME using LOGINDEXBUILD ON"
   ```

7. Define the HADR synchronization mode by running the following command. The example uses async.

   ```
   db2 "update db cfg for $DBNAME using HADR_SYNCMODE async"
   ```

   For more information about available modes and their benefits, see HADR Synchronization Mode.

4.3.9 Starting HADR and checking the synchronized data

This section describes how to start IBM Db2 HADR on both servers and verify that data synchronization is working correctly.

Complete the following steps to start HADR and verify data synchronization:

1. Log in as `$DB2ADM` on the target server and start HADR in standby mode by running the following command:

   ```
   db2 start hadr on db $DBNAME as standby
   ```

   The expected output appears in Example 4-7.

   Example 4-7 Command output

   ```
   DB20000I  The START HADR ON DATABASE command completed successfully.
   ```

2. Log in as `$DB2ADM` on the source server and start HADR in primary mode by running the following command:

   ```
   db2 start hadr on db $DBNAME as primary
   ```

   The expected output appears in Example 4-8.

   Example 4-8 Command output

   ```
   DB20000I  The START HADR ON DATABASE command completed successfully.
   ```

3. Verify the HADR state by running the following command:

   ```
   db2pd -d $DBNAME -hadr
   ```

   Check the output for the values that are shown in Table 4-3.

Table 4-3 HADR status values for both servers

Field	Source server	Target server
HADR_ROLE	PRIMARY	STANDBY
HADR_STATE	PEER	PEER
HADR_CONNECT_STATUS	CONNECTED	CONNECTED

4.3.10 Running the core migration

This section describes how to finish the migration process by stopping the SAP system on the source server and initiating a takeover on the target server.

Complete the migration from the source server to the target server by completing the following steps:

1. Stop the SAP system on the source server. Log in as `$SIDADM` and run the following command:

   ```
   $ stopsap
   ```

   Wait for the command to finish.

2. Initiate a takeover on the target serve. Log in as `$DB2ADM` and run the following command:

   ```
   $ db2 takeover hadr on database $DBNAME
   ```

Tip: The **takeover** command includes an option by force that can help if the source system was not shut down cleanly. For more information, see the TAKEOVER HADR command.

3. Verify that the HADR role changed from `standby` to `primary` on the target server by running the following command:

   ```
   db2pd -d th2 -hadr | grep ROLE
   ```

4. If the IBM Db2 HADR role is `primary` on the target server, start the SAP system on the target server. Log in as $SIDADM and run the following command:

   ```
   $ startsap
   ```

Reviewing the final state

Verify the system configuration after migration to help ensure readiness for future operations or rollback scenarios.

- The IBM Db2 database remains configured for synchronization, but is inactive at this stage. Keep the HADR configuration in place if a reverse migration is planned (for example, in a disaster recovery scenario).
- If the source system is removed after migration, remove the HADR configuration, as described in How to remove existing HADR configuration.

Here are the migration statuses:

- `The migration finished successfully.`
- `The SAP system on the source server is offline.`
- `The SAP system on the target server is up and running.`

Important: Update the SAP system DNS record to point to the target SAP server. This adjustment helps ensure that systems that run the SAP logon GUI can connect to the target server.

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513
- *SAP HANA on IBM Power Systems Architectural Summary*, REDP-5569
- *SAP HANA on IBM Power Systems Virtual Servers: Hybrid Cloud Solution*, REDP-5693

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

© Copyright IBM Corp. 2025.

Redbooks

Migrating SAP Workloads to IBM Power Virtual Server: A Practical Guide for Hybrid Cloud Transformation

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8577-00

ISBN 0738462349

Printed in U.S.A.

Get connected