

A low-angle, upward-looking photograph of a jet engine's nozzle section. The image shows several large, circular, copper-colored nozzle inlets arranged in a circular pattern. Each inlet has a red center cap. The surrounding structure is metallic and complex, with various pipes and components visible. The background is a clear blue sky.

Bert Vanstechelman

The SAP® HANA Deployment Guide

- ▶ SAP HANA sizing, capacity planning guidelines, and data tiering
- ▶ Backup and recovery options and procedures
- ▶ Deployment options and data provisioning scenarios
- ▶ Software and hardware virtualization in SAP HANA

Table of Contents

Introduction	7
Content	7
Acknowledgments	9
1 Architecture	11
1.1 Technology	11
1.2 Editions	11
1.3 Use Cases	14
1.4 On-Premise	19
1.5 Capacity Planning	24
1.6 Data Tiering	41
2 Landscape Layout	47
2.1 Deployment	47
2.2 SAP HANA and the SAP Application Server on one system	62
2.3 Data Provisioning	64
3 Scalability	75
3.1 Scale-Out	76
3.2 Scale-Out and High Availability	77
3.3 Scale-Out and Data Partitioning	78
4 Disaster Recovery	85
4.1 What is Disaster Recovery?	86
4.2 Recovery Point Objective versus Recovery Time Objective	87
4.3 Concepts	88
4.4 Backup Options	93

5	Replication	99
5.1	Storage Replication	99
5.2	System Replication	102
5.3	Log Shipping	109
5.4	Storage Replication versus System Replication versus Log Shipping	111
6	High Availability	115
6.1	High Availability versus Disaster Recovery?	116
6.2	High Availability Clusters	116
7	Virtualization	125
7.1	Virtualization on VMware	126
7.2	Virtualization on IBM Power VM	132
8	Conclusion and Summary	135
8.1	Cluster Software versus Virtualization	135
8.2	Storage Replication versus System Replication	140
8.3	System Replication versus Storage Replication versus Virtualization	144
8.4	Understand your Performance Load!	145
9	Appendix: References	147
9.1	SAP Support Notes	147
9.2	Articles and Other Publications	148
A	The Author	153
B	Index	154
C	Disclaimer	157

2 Landscape Layout

Options, options, options. There are many options for the deployment of SAP HANA systems. There is the standard or standalone SAP HANA system (SCOS), multitenant database containers (MDC), multiple components on one database (MCOD), multiple components on one system (MCOS) and finally virtualization for SAP HANA.

There are various aspects influencing the choice between the different deployment options. The impact on high-availability and disaster recovery needs to be considered and the required software change management landscape with its development, testing, quality assurance and production systems has to be mapped to SAP HANA hardware infrastructure.

2.1 Deployment

This section discusses the various different types of technical deployment options (see Figure 2.1), such as:

- ▶ single application on one SAP HANA system (SCOS)
- ▶ multitenant database containers (MDC)
- ▶ multiple applications on one SAP HANA system (MCOD)
- ▶ multiple SAP HANA systems on one host (MCOS)
- ▶ SAP HANA with virtualization




Figure 2.1: SAP HANA deployment options

2.1.1 Single Application on One SAP HANA System (SCOS)

The standard SAP HANA deployment is a single SAP HANA application running in a single database schema in a single SAP HANA database as part of an SAP HANA system or, as SAP calls it, a single application on one SAP HANA system (SCOS). This is a simple, straightforward scenario that is supported for all scenarios without restriction.

For example, two SAP HANA appliances are sufficient for a two-system SAP BW landscape (development and production). There is no failover for the production system in this setup. This might be acceptable for customers where reporting is not considered business critical.

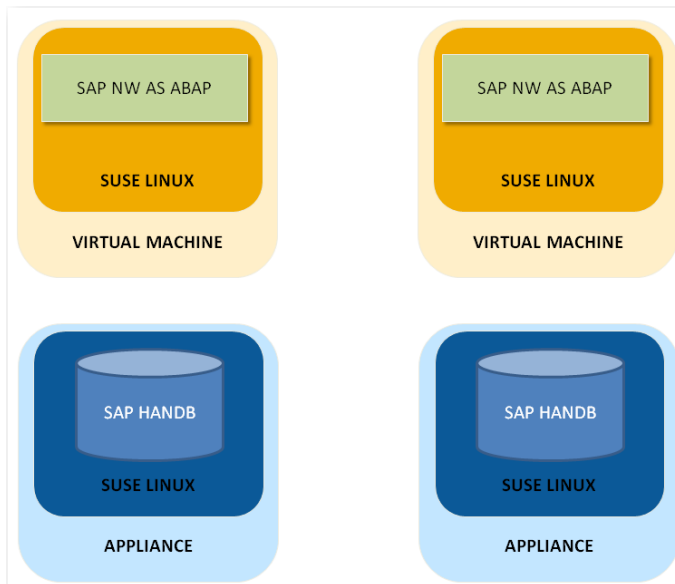


Figure 2.2: SAP HANA standard deployment option

The standard SAP HANA deployment option (see Figure 2.2) is a system layout which was often used in the early days of SAP HANA and is still used by customers who are only deploying SAP HANA for a limited use case such as SAP BW.

Note that in the example above, only the SAP HANA database is installed on the appliance. The SAP NetWeaver application server is deployed on a different, in this case, virtual machine running on SUSE Enterprise Linux.

2.1.2 Multitenant Database Containers (MDC)

SAP HANA supports multiple isolated databases in a single SAP HANA system. These are referred to as multitenant database containers. The multitenant database container setup of SAP HANA is comparable to the SQL-Server or Sybase instance layout. There is an SAP HANA instance, a system database and several tenant databases. The system database is used for central system administration. It is also the database from which recoveries of the tenant databases are initiated.

An SAP HANA system installed in multiple-container mode is identified by a single system ID (SID). An SID and a database name identify databases. From an administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA studio, connect to the system or the tenant databases.

All the databases in a multiple-container system share the same installation of database system software, the same computing resources, and the same system administration. As a result, software upgrades or system maintenance impact all databases. In addition, system replication applies to the whole SAP HANA instance; that is, for all tenant databases including the system database.

Newly created tenant databases are automatically integrated into the replication process after they are backed up.

However, each database is fully isolated when it comes to:

- ▶ Security—Each tenant DB has its own users and authorizations which are completely separate from the other tenant databases. The database catalog and repository are also isolated in each tenant DB.
- ▶ Backups—With SAP HANA multitenant database containers, if each application is deployed on its own tenant DB, then each can be backed up and recovered independently.
- ▶ Moving and copying tenant DB's—Tenant databases can be moved or copied using the backup and restore capabilities. This only needs downtime for the tenant database affected. The other tenant databases can stay online. Simply perform a backup and then either create a new tenant database and restore the backup into this tenant database, or restore the copy into an existing tenant database.

- ▶ Traces and logs—Each tenant database has his own set of trace and log files.

In general, all applications that are supported to run on a single database SAP HANA system are also supported to run on an MDC system. Tools exist to convert a single-container system to a multiple-container system.

Many customers use MDC to consolidate several SAP HANA databases into one SAP HANA system. This setup minimizes the number of appliances and reduces total cost of ownership (TCO).

Consider the following example:

The customer has SAP ERP, SAP PO and SAP CRM landscapes. Every landscape consists of a development, acceptance and production system. In addition, system replication is a requirement for all production systems.

The following SAP HANA landscape has been designed (see Figure 2.3):

- ▶ There are two appliances; one for the production and another for the non-production SAP HANA systems.
- ▶ The appliance for production hosts an MDC installation for the production SAP-HANA systems. The MDC consists of a system database and three production databases, one each for SAP ERP, SAP PO and SAP CRM.
- ▶ The appliance for the non-production systems hosts two MDC installations. One MDC installation for the development and another for acceptance systems. Each MDC installation has one system database and three non-production databases, one each for SAP ERP, SAP PO and SAP CRM. The two MDC installations have their own SID and software installation and are actually MCOS (multiple SAP HANA installations on one system). MCOS is explained in detail in Section 2.1.4.
- ▶ On the appliance for the non-production systems, there is an MDC installation for the failover of the production MDC system. This SAP HANA system has the same layout as the SAP HANA system on the production system and system replication is set up between both.
- ▶ Only SAP HANA is installed on the appliances. The SAP application servers for SAP ERP, SAP PO and SAP CRM are installed on two ESX servers running VMware. High availability for the production application servers is guaranteed by VMware HA (see Section 7.1.1).

B Index

A

Appliance 20

B

Backup

- Backup to file 93
- Backup using Backint 94
- Backup using Storage Snapshot 95

Business Continuity 86

C

Configuration Files 92

D

Data Backup 89

Data Tiering, Near-Line Storage
42

Delivery Options, On-premise
19

Delta Backup 90

Delta Merge 79

Deployment

- MCOD 52
- MCOS 54
- MDC 50
- SCOS 49
- Virtualization 55

Differential backups 90

Direct Extractor Connection
(DXC) 72

Disaster Recovery 86

Distributed SAP HANA, Host
auto-failover 77

Dynamic Tiering 43

E

EMC, RecoverPoint 111

Extended Application Services
18

Extraction-Transformation-Load
(ETL) 69

H

Hadoop 45

Hardware Configuration Check
Tool 21

High availability 116

High Availability

- Active & active 118
- Active & passive 118
- Cluster 116
- Logical host 117
- Split-brain 117

I

IBM POWER 132

Live Partition Mobility 133

Incremental backups 90

L

Log Backup 91

P

Partitioning

- Database 78
- Tables 78

R

Recovery Point Objective 87

Recovery Time Objective 87

Red Hat, High Availability Add-on 123

S

SAP Central Services 119

 SAP Enqueue Service 120

 SAP Message Service 120

SAP Certified Technology

 Specialist (Edition 2015) 21

SAP Components 119

 Application Servers 120

 Database 119

 Primary Application Server 120

 Replicated Enqueue 120

 SAP Central Services 119

SAP Data Services 69

SAP HANA

 Application Development 18

 auto-restart 128

 Data Mart 16

 SAP NetWeaver 14

SAP HANA Cloud Integration 12

SAP HANA Cloud Platform

 SAP HANA Enterprise Cloud (HEC) 12

 SAP HANA One 12

SAP HANA Cloud Platform 12

SAP HANA Live 16

SAP HANA Vora 46

SAP IQ 42

 Multiplex 43

 Simplex 43

SAP Landscape Transformation Replication Server (SAP LT) 65

SAP Mobilink 71

 UltraLite 71

SAP NetWeaver

 Dual-stack 15

SAP NetWeaver 14

ABAP 14

JAVA 14

SAP Replication Server (SRS) 69

Scale-out 76

Scale-up 75

Sizing

 Appliance 35

 Quick Sizer 24

 Report for Business Suite Systems 25

 Report for SAP BW 27

 TDI 35

 T-shirt 33

Smart Data Access 73

Smart Data Streaming 70

Snapshot

 Database aware 96

 Database snapshot 96

 Database Unaware 96

 Storage snapshots 95

Storage Replication 99

 Asynchronous 100

 Point-in-time 101

 Semi-synchronous 101

 Synchronous 100

SUSE Linux Enterprise

 High Availability 123

Symantec ApplicationHA 129

System Replication 102

 Asynchronous 105

 Log shipping 109

 Synchronous 104

 Synchronous in-memory 104

 Synchronous with full sync option 104

T

Table Partitioning

 Hash 81

 Range partitioning 81

 Round-robin 81

Tailored Data Center Integration

20

TDI

Enterprise Network 22

Enterprise Storage 21

Transaction Log-Based Data

Replication 69

V

Virtualization

Fujitsu Physical Partitioning 59

Hardware 58

Hewlett Packard nPartitions

58

Hitachi LPAR 2.0 57

Huawei FusionSphere 58

Hypervisor 56

IBM PowerVM 57

KVM 58

Lenovo FlexNode 59

Software 56

VMware vSphere 57

XEN 58

VMware

Fault Tolerance (FT) 128

High Availability (HA) 127

Site Recovery Manager 102

vMotion 131

vSphere 126