

**SAP HANA**

## **Overview**

**Issue** 05

**Date** 2018-11-15

**Copyright © Huawei Technologies Co., Ltd. 2019. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

---

# Contents

---

<b>1 Introduction to SAP HANA.....</b>	<b>1</b>
<b>2 Advantages of Deploying SAP HANA on the Public Cloud.....</b>	<b>3</b>
<b>3 Required Resources.....</b>	<b>4</b>
3.1 Licenses.....	4
3.2 Public Cloud Services.....	4
<b>4 Application Scenarios.....</b>	<b>6</b>
4.1 Description.....	6
4.2 Single-Node Deployment Where HA Is Not Required.....	8
4.3 Single-Node Deployment Where HA Is Required.....	14
<b>5 Service Continuity.....</b>	<b>17</b>
<b>6 Support.....</b>	<b>19</b>
6.1 Contact Information.....	19
6.2 HUAWEI CLOUD Technical Support.....	19
6.3 SAP Support.....	19
<b>A Change History.....</b>	<b>20</b>

---

# 1 Introduction to SAP HANA

---

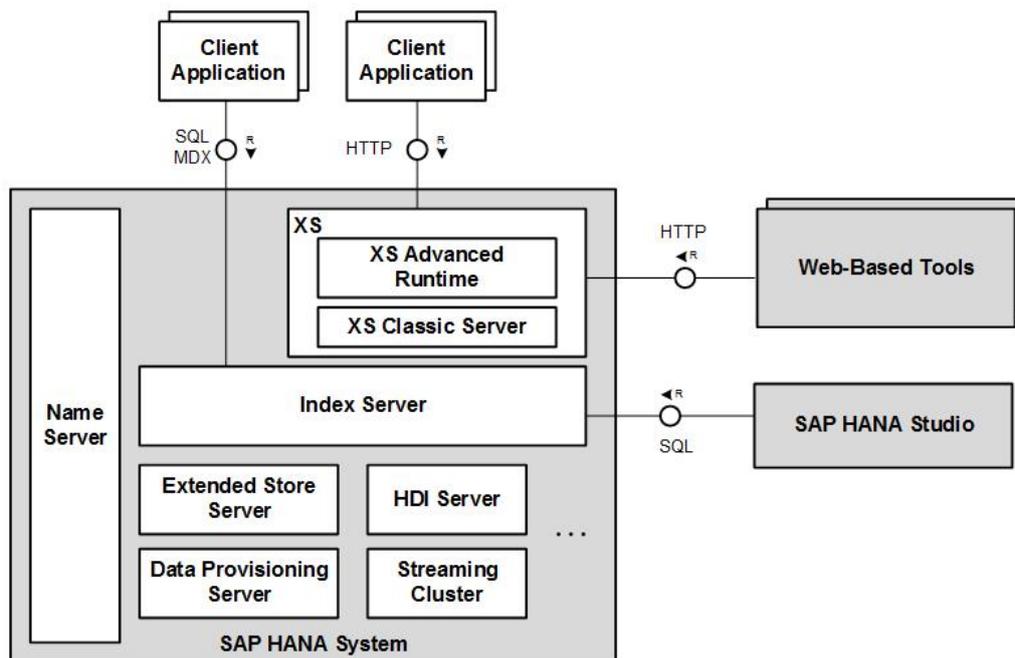
SAP High-Performance Analytic Appliance (HANA) is a high-performance real-time data computing platform launched by SAP in June 2011. The platform leverages the memory computing technology, which enables users to directly query and analyze a large quantity of real-time business data.

SAP HANA has the following features:

- Rapid data processing  
SAP HANA data is stored in the in-memory database (IMDB), which features rapid data input/output (I/O). Success stories show that the data processing rate provided by SAP HANA is 1,000 or even 10,000 times faster than that provided by customers' original database systems.
- Support for online analytical processing (OLAP) and online transaction processing (OLTP)  
OLAP is a main application in data warehouse (DWH) systems. It supports complex analysis, focuses on decision support, and provides query results that are easy to understand.  
OLTP is a main application of traditional relational databases. It processes basic and routine transactions, such as banking transactions.
- Support for both structured and unstructured data  
Structured data is a type of information represented using a unified structure, such as digits and symbols. The other data that cannot be represented using a unified structure is called unstructured data, such as texts, images, audio, and HTML.
- Various peripheral application software  
The application software includes data replication (DR) and business intelligence (BI).
- Structured query language (SQL), multidimensional expressions (MDX), and BI consumer services (BICS) interfaces for the application layer
- Support for multiple data importing modes

**Figure 1-1** shows the SAP HANA architecture.

Figure 1-1 SAP HANA architecture



## Application Scenarios

SAP HANA applies to following scenarios:

- Accelerator for other databases
- Data source for report analysis
- OLAP DWH
- OLTP database
- Unified SAP database platform

# 2 Advantages of Deploying SAP HANA on the Public Cloud

---

Deploying SAP HANA on the public cloud makes full use of the advantages of cloud services, which is easy to expand and use, reducing the purchase cost and improving the deployment efficiency.

- **Low cost**  
You only need to purchase the SAP HANA resources required in such scenarios as DEV, TST, and TRN. This significantly reduces purchasing costs.
- **High deployment efficiency**  
The time required for deploying the SAP HANA system on the public cloud has been shortened from several days to several hours. This significantly improves deployment efficiency and reduces manpower costs.
- **Simple expansion**  
On the public cloud, you can flexibly expand the capacity of a single node or expand resources from a single node to a cluster.
- **Flexible configuration**  
If your SAP HANA system is required to apply in such scenarios as DEV, TST, and TRN, you can purchase different SAP HANA deployment modes to suit the application requirements.

# 3 Required Resources

## 3.1 Licenses

SAP HANA is authorized in Bring Your Own License (BYOL) mode. In this mode, you must log in to the SAP technical support website and apply for a license.

In addition to applying for a license, you must purchase public-cloud-related resources.

## 3.2 Public Cloud Services

The cloud services required for deploying SAP HANA on the public cloud are as follows:

- **Elastic Cloud Server (ECS)**

SAP HANA software is deployed on both ECSs and HANA ECSs. Specifically, SAP HANA is deployed on HANA ECSs. NFS servers, SAP HANA Studio, and NAT servers are deployed on ECSs.

 **NOTE**

- In the SAP HANA cluster deployment mode, the NFS server shares its shared and backup volumes with all SAP HANA nodes.
- A HANA ECS is a type of high-performance ECS.
- For more information about SAP HANA Studio and NAT servers, see *SAP HANA User Guide (Single-Node Deployment)*.
- **Elastic Volume Service (EVS)**  
All ECSs used in SAP HANA scenarios have EVS disks attached.
- **Virtual Private Cloud (VPC)**  
All ECSs in SAP HANA scenarios locate in the same VPC. They are isolated using VPC subnets and security groups for network security.
- **Image Management Service (IMS)**  
Images are required for creating ECSs and HANA ECSs.
- **Volume Backup Service (VBS)**  
VBS backs up the volumes attached to SAP HANA nodes.
- **Scalable File Service (SFS)**

In some deployment scenarios, an SFS file system is created to provide the backup volumes and the shared path to SAP HANA nodes.

---

# 4 Application Scenarios

---

## 4.1 Description

This document describes how to deploy SAP HANA systems within an AZ.

For details about deploying SAP HANA systems in HA and DR scenarios across AZs or regions, see the *SAP HANA HA and DR Guide*.

### Scale-up and Scale-out

SAP HANA nodes can be expanded in the scale-up or scale-out mode.

- Scale-up mode

Also called single-node system mode, in which an SAP HANA system contains only one valid node. If high availability (HA) is required, construct such an architecture using two single nodes through system replication. In scale-up mode, when the system requires expansion, add CPUs, memory capacity, and hard disks to the node.

The public cloud does not support scale-up expansion for a node on which SAP HANA is running.

- Scale-out mode

Also called cluster system mode, in which an SAP HANA system contains multiple nodes. In scale-out mode, when the system requires expansion, add more nodes to the system.

In accordance with the node expansion modes, SAP HANA can be deployed in the following scenarios:

- Single-node deployment: applies in OLTP scenarios. You can choose HA configuration as required. SAP HANA receives and processes data quickly, provides processing results in a short period of time, and rapidly responds to user operations.
- Cluster deployment: applies in OLAP scenarios. As a data warehouse, SAP HANA offers support for decision-makers and senior managers. It can quickly and flexibly process complex queries on a large amount of data based on the analysts' requests. It can also provide decision makers with the query results intuitively and clearly. Using SAP HANA, decision makers can obtain accurate information about the enterprise operating status, learn object demands, and make correct decisions.

## Application Scenarios and Deployment Modes

Recommended deployment modes vary depending on systems.

- **Production (PRD):** indicates the production system where SAP HANA is formally used.
- **Quality Assure (QAS):** indicates the quality assurance system where SAP HANA functions, performance, and reliability are fully verified.
- **Development (DEV):** indicates the development system where the development engineers configure and verify the compatibility between application software and SAP HANA and continuously optimize the application software.
- **Training (TRN):** indicates the training and demonstration system where you provide a training or demonstration after deploying SAP HANA.
- **Test (TST):** indicates the test scenario where the development engineers test application software and SAP HANA to verify the functions of application software after application software development is complete.

### NOTE

In Suite on HANA (SoH) scenario, SAP HANA works with SAP business suites, such as Enterprise Resource Planning (ERP) or Customer Requirement Management (CRM). In this scenario, SAP HANA provides OLTP functions. The SAP HANA process latency is the key concern.

In Business Warehouse on HANA (BWoH) scenario, SAP HANA works with SAP Business Warehouse. In this scenario, SAP HANA provides OLAP functions and supports rapid computing and analyzing on massive data. The SAP HANA processing performance and the network bandwidths between SAP HANA nodes are the key concern.

**Table 4-1** Systems and deployment schemes

System	SoH	BWoH (Single-Node)	BWoH (Cluster)
PRD	Single-node scenario where HA is required	Single-node scenario where HA is required	Cluster scenario where two NFS servers are deployed
QAS	Single-node, regardless of whether HA is required or not	Single-node, regardless of whether HA is required or not	Single-node scenario where HA is not required, cluster scenario where two NFS servers are required, or cluster scenario where single NFS server is required
DEV	Single-node scenario where HA is not required	Single-node scenario where HA is not required	Single-node scenario where HA is not required
TRN	Single-node scenario where HA is not required	Single-node scenario where HA is not required	Single-node scenario where HA is not required
TST	Single-node scenario where HA is not required	Single-node scenario where HA is not required	Single-node scenario where HA is not required

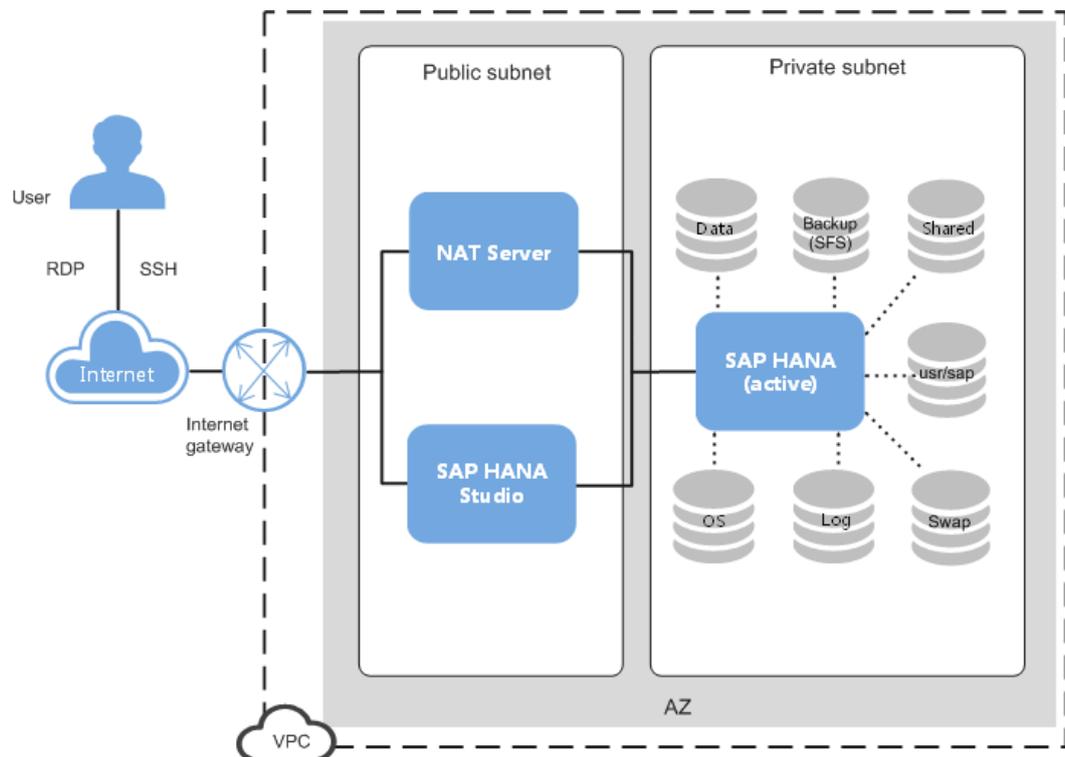
## 4.2 Single-Node Deployment Where HA Is Not Required

Single-node deployment is usually applied in OLTP scenarios. You can choose the HA configuration as required.

### System Architecture

**Figure 4-1** shows the system architecture in single-node scenario where HA is not required.

**Figure 4-1** Single-node scenario where HA is not required



The concepts involved in the preceding figure are as follows:

- VPC network: All SAP HANA nodes are deployed within a VPC network, and all nodes in the HANA system belong to the same AZ to ensure network security.
- Public subnet:
  - Network Address Translation (NAT) instance ECS: allows you to access SAP HANA nodes using SSH.
  - SAP HANA Studio ECS: runs the SAP HANA Studio software. You can use RDP or SSH to access the SAP HANA Studio ECS and manage the SAP HANA system.
- Private subnet:
  - SAP HANA node: used for deploying the SAP HANA software. An SAP HANA server has the following disks attached:
    - OS disk: provides the directory for installing the OS.

- Data volume: periodically stores the data transmitted from the SAP HANA IMDB (a database running in high-performance memory). The period is 5 minutes by default.
- Log volume: stores the data triggered by an event. When an event, for example, a record or a batch of records are updated, is triggered for the server IMDB, the system will write the latest IMDB data into the log volume.
- Shared volume: stores the SAP HANA installation software and SAP HANA database log files.
- Backup volume: stores SAP HANA database backup files.
- Usr/sap: used to mount to the **/usr/sap** directory.
- Swap volume: Linux swap space.

## Requirements on SAP HANA Nodes

SAP HANA node specifications used in SoH scenario must meet the following requirements:

**Table 4-2** and **Table 4-5** list the specifications of HANA ECSs certified by SAP on HUAWEI CLOUD. **Table 4-3** and **Table 4-4** list the ECSs provided by HUAWEI CLOUD for non-production systems.

**Table 4-2** Large-memory E1 ECS specifications

ECS Type	pCPUs	Memory (GB)	Flavor
Large-memory	16	470	e1.4xlarge
	32	940	e1.8xlarge

**Table 4-3** Large-memory E2 ECS specifications

ECS Type	pCPUs	Memory (GB)	Flavor
Large-memory	12	256	e2.3xlarge

**Table 4-4** Large-memory ET2 ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor
Large-memory	8	128	et2.2xlarge.16
	18	256	et2.4xlarge.14
	36	512	et2.8xlarge.14

**Table 4-5** Large-memory E3 ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor
Large-memory	28	348	e3.7xlarge.12
	56	696	e3.14xlarge.12
	104	1466	e3.26xlarge.14
	208	2932	e3.52xlarge.14
	208	4096	e3.52xlarge.20

SAP HANA node specifications used in the BWoH scenario must meet the following requirements:

**Table 4-6** and **Table 4-8** list the specifications of HANA ECSs certified by SAP on HUAWEI CLOUD. **Table 4-7** lists the ECSs provided by HUAWEI CLOUD for non-production systems.

**Table 4-6** Large-memory E2 ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor
Large-memory	18	445	e2.4xlarge
	36	890	e2.9xlarge

**Table 4-7** Large-memory ET2 ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor
Large-memory	8	128	et2.2xlarge.16
	18	256	et2.4xlarge.14
	36	512	et2.8xlarge.14

**Table 4-8** Large-memory E3 ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor
Large-memory	28	348	e3.7xlarge.12
	56	696	e3.14xlarge.12
	104	1466	e3.26xlarge.14
	208	2932	e3.52xlarge.14
	208	4096	e3.52xlarge.20

SAP HANA OSs and disks must meet the following requirements:

 **NOTE**

- SAP HANA ECS volumes include log volumes, data volumes, shared volumes, backup volumes, and /usr/sap volumes.
- A shared disk can be attached to multiple ECSs, while a non-shared disk can only be attached to one ECS.

**Table 4-9** Requirements on HANA ECS OS in single-node deployment mode

Item	Specifications
OS	SUSE Linux Enterprise Server 12 SP1 for SAP or later

**Table 4-10** Requirements on E1 and E2 HANA ECS disks in single-node deployment scenarios

Disk	Type	Sharing Mode	Size
OS volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	N/A
Log volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	<ul style="list-style-type: none"> <li>● When the memory size is less than or equal to 512 GB, the log volume capacity is half of the memory size and rounded up for decimal places.</li> <li>● When the memory size is greater than 512 GB, the log volume capacity is 512 GB.</li> </ul>
Data volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	The recommended size is at least 1.2 times that of the memory size.
Shared volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	The recommended size is at least 1.2 times that of the memory size.
Backup volume	SFS	N/A	The recommended capacity is three times or more of the memory size.
SBD volume	<b>Ultra-high I/O (Latency optimized)</b>	Shared disk (SCSI)	10 GB
/usr/sap volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	50 GB

Disk	Type	Sharing Mode	Size
Swap volume	<b>Ultra-high I/O (Latency optimized)</b>	Non-shared disk	10 GB

**Table 4-11** Requirements on ET2 HANA ECS disks in single-node deployment scenarios

Disk	Type	Sharing Mode	Size
OS volume	<b>High I/O</b>	Non-shared disk	-
Log volume	<b>Ultra-high I/O</b>	Non-shared disk	For details, see <a href="#">Table 4-12</a> .
Data volume	<b>Ultra-high I/O</b>	Non-shared disk	Create an EVS disk. Use LVM to create soft partitions and logically divide the disk into data volumes. For details, see <a href="#">Table 4-12</a> .
Shared volume	<b>High I/O</b>	Non-shared disk	The recommended size is at least 1.2 times that of the memory size.
Backup volume	SFS	N/A	The recommended size is at least three times that of the memory size.
SBD volume	<b>High I/O</b>	Shared disk (SCSI)	10 GB
/usr/sap volume	<b>High I/O</b>	Non-shared disk	50 GB
Swap volume	<b>High I/O</b>	Non-shared disk	10 GB

**Table 4-12** Recommended log and data volume specifications for ET2 ECSs

Flavor	Log Volume Size (GB)	Data Volume Size
et2.2xlarge.16	100	2 x 200 GB EVS disks
et2.4xlarge.14	200	2 x 250 GB EVS disks
et2.8xlarge.14	512	2 x 400 GB EVS disks

**Table 4-13** Requirements on E3 ECS disks in single-node deployment scenarios

Disk	Type	Sharing Mode	Size
OS volume	<b>High I/O</b>	Non-shared disk	-
Log volume	<b>Ultra-high I/O</b>	Non-shared disk	For details, see <a href="#">Table 4-14</a> .
Data volume	<b>Ultra-high I/O</b>	Non-shared disk	Create an EVS disk. Use LVM to create soft partitions and logically divide the disk into data volumes. For details, see <a href="#">Table 4-14</a> .
Shared volume	<b>High I/O</b>	Non-shared disk	The recommended size is at least 1.2 times that of the memory size.
Backup volume	SFS	N/A	The recommended size is at least three times that of the memory size.
SBD volume	<b>High I/O</b>	Shared disk (SCSI)	10 GB
/usr/sap volume	<b>High I/O</b>	Non-shared disk	100 GB
Swap volume	<b>High I/O</b>	Non-shared disk	10 GB

**Table 4-14** Recommended log and data volume specifications for E3 ECSs

Flavor	Log Volume Size (GB)	Data Volume Size
e3.7xlarge.12	200	2 x 250 GB EVS disks
e3.14xlarge.12	512	2 x 450 GB EVS disks
e3.26xlarge.14	512	3 x 600 GB EVS disks
e3.52xlarge.14	512	5 x 750 GB EVS disks
e3.52xlarge.20	512	5 x 950 GB EVS disks

## Requirements on Other Nodes

Other nodes include the NAT server and SAP HANA Studio nodes. [Table 4-15](#) lists the requirements on these nodes.

**Table 4-15** Data planning for other nodes

Node	Specifications
SAP HANA Studio	<ul style="list-style-type: none"><li>● OS: <b>NOTE</b> Based on service requirements, use a Windows or Linux ECS to deploy SAP HANA Studio.</li><li>- Windows: Windows Server 2012 R2 or Windows Server 2008 R2</li><li>- Linux: SUSE Linux Enterprise Server 12 SP1 or later</li><li>● Flavor: s1.xlarge (4 vCPUs and 16 GB memory capacity)</li><li>● Disk: 80 GB system disk</li></ul>
NAT server	<ul style="list-style-type: none"><li>● OS: SUSE Linux Enterprise Server 12 SP1 or later</li><li>● Flavor: s1.medium (1 vCPU and 4 GB memory capacity) or higher</li><li>● Disk: 40 GB system disk</li></ul>

## Deployment

For instructions about how to create ECSs or HANA ECSs on the management console and install the SAP HANA software on the ECS or HANA ECS, see the *SAP HANA User Guide (Single-Node Deployment)*.

## 4.3 Single-Node Deployment Where HA Is Required

Single-node deployment is usually applied in OLTP scenarios. You can choose the HA configuration as required.

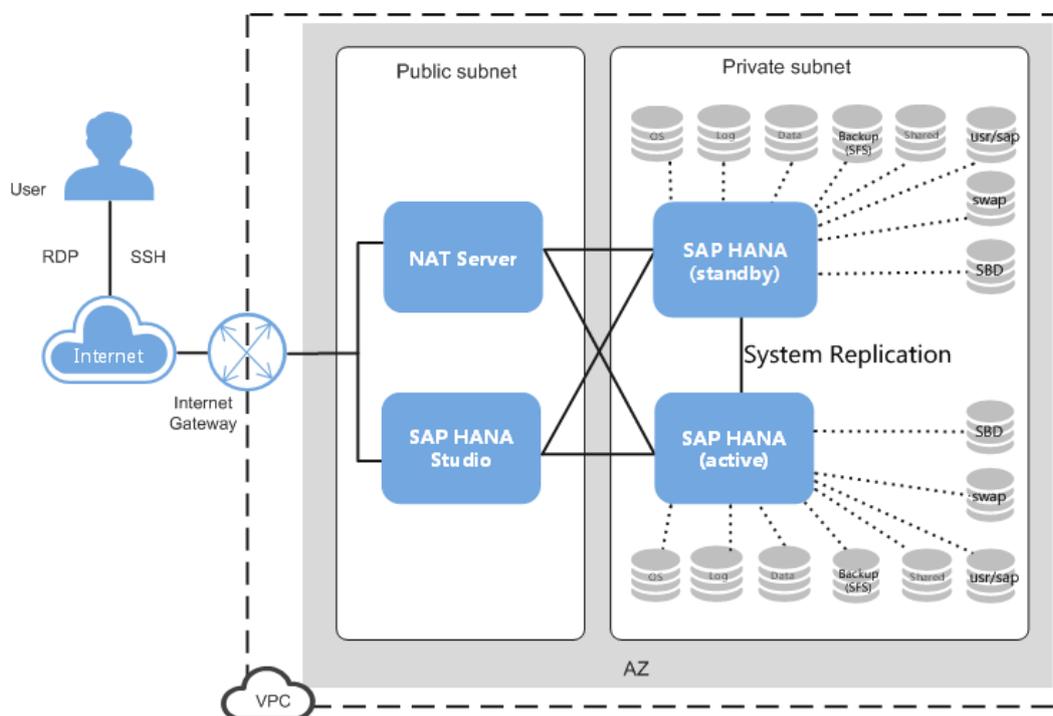
### System Architecture

[Figure 4-2](#) and [Figure 4-3](#) show the system architecture in single-node scenario where HA is required.

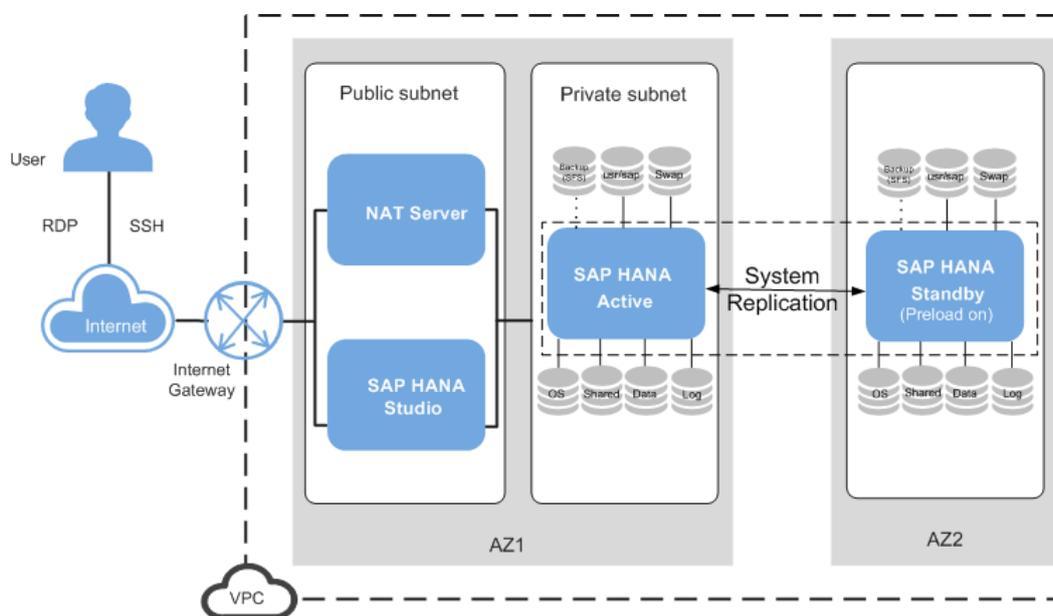
#### NOTE

In the single-node scenario where HA is required, active/standby switchovers can be manually performed, or automatically performed using scripts.

**Figure 4-2** Single-node HA deployment within an AZ



**Figure 4-3** Single-node HA deployment across AZs



The concepts involved in this scenario are the same as those involved in [Single-Node Deployment Where HA Is Not Required](#). The differences between the two deployment modes are as follows: Two independent SAP HANA nodes in the same AZ or different AZs are connected to the SAP HANA Studio for management. The two SAP HANA nodes synchronize data and implement HA using System Replication.

Compared with the mode in which HA is not required in the single-node scenario, the mode in which HA is required has the following requirements:

- You must create two nodes and synchronize time, format disks, and install the SAP HANA software on them.
- The two SAP HANA nodes must connect to SAP HANA Studio.
- You must manually configure HA on the two SAP HANA nodes using System Replication.
- You must configure scripts for the two SAP HANA nodes for active/standby switchovers.  
This only applies to SAP HANA nodes running the OS SUSE Linux Enterprise Server 12 SP1 for SAP or later.
- In the single-node scenario where SAP HANA nodes are deployed across AZs for HA, the backup volume must be provided by Scalable File Service (SFS) and three ECSs are required. Each ECS is bound to a SCSI disk and iSCSI configuration is required for SBD.

## Requirements on SAP HANA Nodes

The requirements for each node are the same as those described in [Requirements on SAP HANA Nodes](#).

## Requirements on Other Nodes

Other nodes include the NAT server and SAP HANA Studio nodes. For details, see [Requirements on Other Nodes](#).

## Deployment

For instructions about how to create ECSs or HANA ECSs on the management console and install the SAP HANA software on the ECS or HANA ECS, see the *SAP HANA User Guide (Single-Node Deployment)*.

# 5 Service Continuity

Service continuity is ensured from the aspect of automatic service restarting, HA, and backup based on public cloud and SAP HANA capabilities.

## Automatic Service Restarting

SAP HANA provides the reliability mechanism. Its Daemon process monitors other SAP HANA processes, including the NameServer, IndexServer, Preprocessor, Complieserver, Statisticsserver, and Xsengine processes. When an SAP HANA process becomes faulty, the system automatically restarts the process.

## HA

In single-node deployment:

- When the physical server where an SAP HANA node is deployed becomes faulty, the system automatically creates an SAP HANA node on another physical server, attaches the log and data volumes to the new node, and switches services from the faulty node to the new node to ensure service continuity.
- In HA mode, the active and standby nodes synchronize data using System Replication. When the active SAP HANA node becomes faulty, you can manually switch services to the standby node.

In cluster deployment:

SAP HANA provides host auto failover (HAF). When a worker node in the cluster becomes faulty, the system automatically switches services from the faulty node to the standby node for cluster HA.

## Backup

The following backup modes are supported:

- Based on shared files  
This backup mode adapts to the SAP HANA backup mechanism. Two SAP HANA nodes (NFS servers) constitute an HA NAS system. SAP HANA data is backed up to the backup volume. Then, the data can be backed up to OBS buckets.
- Based on volumes

This backup mode is provided based on volume backup on the public cloud platform. In this mode, you must back up and restore data by yourself and ensure data consistency in SAP HANA data and log volumes.

- Data can only be backed up within an AZ. To do so, create a snapshot for volumes. Then, the snapshot can be used for rolling back data or creating new volumes. The newly created volumes are within the AZ where the original volumes locate.
- Data can be backed up in OBS buckets. To do so, create a snapshot for volumes and back up data based on the snapshot. Then, data can be restored both within and across the AZ.

# 6 Support

---

## 6.1 Contact Information

If you encounter any SAP problems, submit tickets to SAP and set the value of SAP component to **BC-OP-LNX-HUAWEI**. Your problem will be handled by the joint support team of SAP and Huawei.

## 6.2 HUAWEI CLOUD Technical Support

Professional HUAWEI CLOUD engineer teams provide you with 24/7 services.

Technical support center of HUAWEI CLOUD:

<https://www.huaweicloud.com/en-us/about/protection.html>

HUAWEI CLOUD hotline: 4000-955-988

## 6.3 SAP Support

Visit [SAP Support Portal Home](#), set the value of SAP component to **BC-OP-LNX-HUAWEI**, and submit your problem.

# A Change History

What's New	Released On
This release incorporates the following change: Added the E3 ECS specifications.	2018-11-15
This release incorporates the following changes: Supported the cross-AZ deployment in the single-node mode. Two independent SAP HANA nodes in different AZs are connected to the SAP HANA Studio for management. The two SAP HANA nodes synchronize data and implement HA using System Replication. For details, see <a href="#">Single-Node Deployment Where HA Is Required</a> .	2018-08-30
This release incorporates the following change: Added SUSE Linux Enterprise Server 12 SP2 for SAP.	2018-02-12
This release incorporates the following changes: <ul style="list-style-type: none"><li>● Updated deployment schemes and ECS specifications.</li><li>● Added the ET2 ECS specifications.</li></ul>	2018-01-10
This issue is the first official release.	2017-11-10