

SAP

Best Practices

Best Practices

SAP HANA Intra-Availability Zone HA Deployment (Based on SLES HAE)

SAP HANA Intra-Availability Zone HA Deployment (Based on SLES HAE)

- Overview
- HANA HA architecture
- Preparations
 - Installation media
 - Access to installation media
- VPC planning
 - Network planning
 - Create a VPC
 - Create a service network
 - Create a heartbeat network
- Create HANA ECS instance
 - Create HANA primary node ECS
 - Create HANA backup node ECS
 - Configure shared storage
 - Configure ENI
 - Configure HaVip
- Configure HANA ECS
 - Modify the host name
 - Configure SSH password-free connection service
 - Configure the NTP service
 - Partition the HANA file system
 - Configure the heartbeat network
 - Configure the HaVip primary node
- Install the HANA database
- Install the HANA Studio

- ## Overview

HANA HA architecture

The diagram illustrates a multi-availability zone (AZ), multi-region database solution architecture on Alibaba Cloud. It is organized into three main layers:

- End User:** Located at the top, representing the external users of the system.
- Aliyun Region:** A dashed box containing the cloud resources, divided into two Availability Zones (AZ):
 - Left AZ:** Contains the **ECS - Studio** instance and the **ECS-DB (primary)** instance.
 - Right AZ:** Contains the **ECS - Bastion** instance and the **ECS-DB (backup)** instance.
- VPC:** A Virtual Private Cloud spanning both AZs, represented by a cloud icon.

Key Components and Connections:

- End User:** Connects to the **ECS - Bastion** instance in the right AZ.
- ECS - Studio:** Manages the database instances.
- ECS - Bastion:** Provides secure access to the database instances.
- HA (High Availability):** Indicated by a red double-headed arrow between the **ECS-DB (primary)** and **ECS-DB (backup)** instances.
- HANA system replication:** Indicated by a green double-headed arrow between the **ECS-DB (primary)** and **ECS-DB (backup)** instances.

Preparations

Installation media

Installation Package	File or Path	Description
----------------------	--------------	-------------

SUSE for SAP	SLE-12-SP2-SAP-x86_64-GM-DVD1.iso	The SUSE can be downloaded from the SuSE official website, and the software has a 60-day trial period.
SUSE for SAP	SLE-12-SP2-SAP-x86_64-GM-DVD2.iso	The SUSE can be downloaded from the SuSE official website, and the software has a 60-day trial period.
SAP HANA database installation package	HDB_SERVER_LINUX_X86_64	
SAP HANA client installation package	HDB_CLIENT_LINUX_X86_64	
SAP HANA Studio installation package	HDB_STUDIO_WINDOWS_X86_64	For Windows

Access to installation media

Access Method	Process	Remarks
Direct upload	Directly upload the package to the ECS	Upload through EIP or VPC
OSS + ossutil	Upload the package to OSS and then download it to ECS.	
OSS + ossfs	Upload the package to the OSS, and then use OSSFS to access the installation media in the OSS.	The source code is required for OSSFS installation on SUSE.

VPC planning

Network planning

Network	Location	Usage	Allocated Subnet
Service network	East China 2 zone A	For Business	192.168.10.0/24
Heartbeat network (redundant)	East China 2 zone A	For SR/HA	192.168.20.0/24


Host Name	Role	Heartbeat Address	Service Address	Virtual Address
hana01.poc.co	Hana primary	192.168.20.19	192.168.10.214	192.168.10.12

m	node			
hana02.poc.com	Hana backup node	192.168.20.20	192.168.10.215	192.168.10.12
HanaStudio	Hana Studio	None	192.168.10.210	None

Create a VPC

The Virtual Private Cloud (VPC) is an isolated network environment built on Alibaba Cloud. VPCs are logically isolated from each other. The VPC is your dedicated private network on the cloud. You have full control over your own VPC instance, including choosing the IP address range and configuring the route table and gateway. For more information and related documents, see the [product documentation](#).

Log on to the VPC console.



Create a service network

Create a service subnet as planned.

Create VSwitch ✕

• VPC

SAP_Network/vpc-gw8873skr23qym39a1fs6 ▼

Destination CIDR Block

192.168.0.0/16

• Name ?

SAP_Business_Network 20/128 ✓

• Zone ?

EU Central 1 Zone A ▼

Zone Resource ?

ECS ✓ RDS ✓ SLB ✓

• Destination CIDR Block

192

 -

168

 -

10

 -

0

 /

24

▼

⚠ The CIDR cannot be changed once the VPC is created.

Number of Available Private IPs

252

Description ?

OK

Cancel

Contact Us

Create a heartbeat network

Create a heartbeat subnet as planned.

Create VSwitch

×

VPC

SAP_Network/vpc-gw8873skr23qym39a1fs6

Destination CIDR Block

192.168.0.0/16

Name ?

SAP_Heartbeat_Network 21/128 ✓

Zone ?

EU Central 1 Zone A

Zone Resource ?

ECS ✓ RDS ✓ SLB ✓

Destination CIDR Block

192 - 168 - 20 - 0 / 24 ✓

ⓘ The CIDR cannot be changed once the VPC is created.

Number of Available Private IPs

252

Description ?

OK

Cancel

Contact Us

Create HANA ECS instance

Create HANA primary node ECS

ECS purchasing page

Access <https://www.aliyun.com/product/ecs> to open the purchasing page. Select an instance type under SAP HANA and click Buy.

Select a payment method

Select either of the following payment methods: Subscription and Pay-As-You-Go.

Select the region and zone.

Select the region and zone. By default, the zones are allocated randomly. You can select a zone according to your needs. For details about the region and zone selection, see [Region and Zone](#).

In this example, China East 2 zone A is selected.


Select instance specifications

Select an instance type certified by SAP HANA, namely, 56 vCPU 480GB (ecs.se1.14xlarge) in series III - Memory type se1 instance type family, or 80 vCPU 960GB (ecs.re4.20xlarge) in the enhanced memory type re4 instance type family. In this example, ecs.se1.14xlarge is selected.


Select image

You can select the public, custom, or shared image, or select an image from the market. The SUSE linux for SAP-12SP2 selected from the market is recommended for SAP HANA.

Note: Select the SUSE for SAP edition, instead of the SUSE 12 standard edition on the official website.



Click Marketplace Image to enter the image market. Enter the keyword sap for searching and select SUSE linux for SAP-12SP2.



Configure storage

System disk : Mandatory. Used to install the operating system. You need to specify the cloud disk type and capacity of the system disk.

Data disk : Optional. If you create a cloud disk as a data disk, you must specify the cloud disk type, capacity, quantity, and whether to encrypt. You can create an empty cloud disk, or use snapshot to create a cloud disk. A maximum of 16 cloud disks can be configured as the data disks.

The capacity of data disks needs to be adjusted according to the number of HANA instances.

Storage

- Disk specifications and performance

System Disk SSD Cloud Disk 60 GiB

SSD Cloud Disk 60 GiB 3600 IOPS

Guide to selecting SSD Cloud Disk/Ultra Cloud Disk/Basic Disk. [Learn More](#)

Data Disk 4/16

You have selected 4 Disks, you can still add 12 Disks

SSD Cloud Disk	1500	GiB	25000 IOPS	Quantity: 1	Device name	☐ Encrypted	Create from snapshot
SSD Cloud Disk	500	GiB	16800 IOPS	Quantity: 1	Device name	☐ Encrypted	Create from snapshot
SSD Cloud Disk	512	GiB	17160 IOPS	Quantity: 1	Device name	☐ Encrypted	Create from snapshot
SSD Cloud Disk	2000	GiB	25000 IOPS	Quantity: 1	Device name	☐ Encrypted	Create from snapshot

[+ Add Disk](#)

Select a network type

Click Next: Network and Security Group to configure the network and security group. :

1. Select a network type.

VPC , Select the VPC and switch. If you do not create the VPC or switch, you can retain the default VPC and switch.

2. Set the public network bandwidth.

If your instance does not need to access the public network or your VPC-type ECS instance uses an elastic public IP address (EIP) to access the public network, you do not need to assign a public IP address to your instance. The EIP can be unbind from the instance anytime.

Note: SAP HANA does not provide external services directly, so the instance does not need a public IP address.

Basic Configurations (Required) **2 Networking (Required)** System Configurations Grouping

Network * **VPC**

How to Select a Network

SAP_Business_Network switch1 Private IP Addresses Available: 252

If you need to create a new VPC, you can [Go to Console and Create](#)

VPC: SAP_Business_Network / vpc-gw8873skr23qym39a1f56	VSwitch: switch1 / vsw-gw80w6ss4ok1uzgnk6cp
VSwitch Zone: EU Central 1 Zone A	VSwitch CIDR Block: 192.168.0.0/24

Network Billing Method

☐ Assign public IP If you leave this box unchecked, your instance will not receive a public IP address and will not be accessible from the internet unless you bind an elastic IP (EIP). [Click bind an EIP to your instance.](#)

Bandwidth Pricing

Security Group * **Reselect Security Group** A security group is similar to a firewall, it is used to control connection requests. Go to the Security Groups page to see an overview. [Create Security Group](#)

- Security Group Limitations
- Security Group Scenarios

Select security group

Select a security group. If you do not create a security group, retain the default security group. For the rules of the default security group, see [Default security group rules](#).

ENI configuration


Note: The second ENI should be added after the ECS instance is successfully created.



Complete the system configuration, grouping, and ECS purchasing.

Create HANA backup node ECS

The creation of HANA backup node ECS is the same as that of HANA primary node ECS, except that the storage space allocation method is different. We recommend that the HANA backup node not be attached to an HANA backup volume as long as the HANA data storage space is sufficient.




Configure shared storage

ECS shared block storage refers to the data block-level storage device that allows multiple ECS instances to read and write data concurrently. It features high concurrency rate, high performance, and high reliability. A single block can be attached to a maximum of 16 ECS instances. For the operation procedure, watch the video [Attach a Shared Block to Multiple ECS Instances](#).

In this example, the shared block is used as the STNOITH of a HA cluster. Select the same zone as the ECS instance and attach the block to the ECS instance of the HA cluster.

Create shared block storage



Select at least 20 GB SSD for the STONITH.

Choose the Datacenter Region and Zone

EU Central 1 (Frankfurt) EU Central 1 Zone A	China North 1 (Qingdao) Random	China North 2 (Beijing) Random	China North 3 (Zhangjiakou) Random
China East 1 (Hangzhou) Random	China East 2 (Shanghai) Random	China South 1 (Shenzhen) Random	Hong Kong Random
Asia Pacific SE 1 (Singapore) Random	US West 1 (Silicon Valley) Random	US East 1 (Virginia) Random	Asia Pacific NE 1 (Tokyo) Random
Middle East 1 (Dubai) Random	Asia Pacific SE 2 (Sydney) Random		

Shared cloud disk can only be attached to ECS instances in the same zone. At present this product can be attached to up to 4 ECS instances and the region cannot be changed after you order it. Therefore choose carefully. Selection advice > [🔗](#)

Shared Cloud Disk

SSD Shared Cloud ... 20 GB 800 IOPS Encryption Create from snapshot

Instance Cost ¥ 0.00 per Hour

Add To Cart Buy Now

After the creation is successful, the following interface is displayed:

Attach shared block storage


Select the ECS instance to be attached to the HA cluster.

Configure ENI


ENI is a virtual network card that can be appended to an ECS instance in a VPC. With ENI, you can build highly available clusters, implement failover at a low cost, and achieve refined network management. All regions support the ENI. For more information, see [ENI](#).

Create an ENI

Log on to the **ECS console**, Select **Network and Security** > **ENI** from the left navigation bar. Select a region. Click **Create an ENI**.




An auxiliary ENI is successfully created.



Bind the HANA ECS instance.


Click Bind Instance for the auxiliary ENI to bind the HANA ECS instance.




Configure HaVip

Private High-Availability Virtual IP Address (HaVip) is a private IP resource which can be created and released independently. The uniqueness of HaVip is that you can broadcast the IP address on ECS using ARP. In this deployment, the HaVip is used as the virtual IP address of the cluster and is attached to each node in the cluster.

Create HaVip



The HaVip is used by the HANA instance to provide service, and is an IP address on the service subnet.



Bind HaVip

Click the ECS instance bound to the HA cluster. Ensure that each ECS instance in the cluster is bound.

HAVIP Details Refresh Delete

Information

ID	hvip-gw8hqu3xth9yucw38lvf	Status	Available
Region	EU Central 1 (Frankfurt)	Intranet IP	192.168.10.11
VPC ID	vpc-gw8873kr23qym39a1fs6	Created At	08/31/2018, 16:44:08
VSwitch	vsw-gw8xyqnf420m4j6gv1k	Description	Edit

Resources

Associate the HANA primary and backup nodes

Access the management page of the created HaVIP.

Click + to add the ECS instances to be associated with, and associate the HANA primary and backup nodes with the HaVip.



Configure HANA ECS

Modify the host name

Configure domain name resolution on the two HANA servers of the HA cluster. Modify the host names as follows:

- Edit /etc/hostname.
- Set hostname
- Edit /etc/hosts and comment out the IPv6 part

Configure SSH password-free connection service

The SSH password-free connection service must be configured on the two HANA servers. The operation is as follows:

Configure the authentication public key

Run the following command on the HANA primary node:

```

hana01:~ # ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:jg4mJvQQUfE+qRcEEANL52xs4MJiRHGWT0mr+pu+g root@hana01
The key's randomart image is:
+----[RSA 2048]-----+
|+BX&=|
|oBo*|
|+o O.|
|+ * B|
| * X S|
|. = + o|
|. o.o..|
| o.o oo|
|.E ++o.|
+----[SHA256]-----+
hana01:~ # cd .ssh/
hana01:~/.ssh # ls
authorized_keys id_rsa id_rsa.pub
hana01:~/.ssh # ssh-copy-id -i id_rsa.pub hana02
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "id_rsa.pub"
The authenticity of host 'hana02 (192.168.10.215)' can't be established.
ECDSA key fingerprint is SHA256:RENeWTeZHFkuQghTUZ8W3cqGCXhbbpD8ZyJB8X7d8+8.
Are you sure you want to continue connecting (yes/no)? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys
Password:

Number of key(s) added: 1

Now try logging into the machine, with: "ssh 'hana02'"
and check to make sure that only the key(s) you wanted were added.
hana01:~/.ssh #

```

Run the following command on the HANA backup node:

```

hana02:~ # ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:9yhPCENstnBX36b/mo/6qebYvcIdNWlTA6lP2V0Bikw root@hana02
The key's randomart image is:
+----[RSA 2048]-----+
| . . .o+oo|
| . o . . oo|
| = . . O o|
| + E . = +|
| + S . = |
| + + = = .|
| + * * . o|
| = = . * |
| oo=+B+=|
+----[SHA256]-----+
hana02:~ # cd .ssh/
hana02:~/.ssh # ls
authorized_keys id_rsa id_rsa.pub
hana02:~/.ssh # ssh-copy-id -i id_rsa.pub hana01
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "id_rsa.pub"
The authenticity of host 'hana01 (192.168.10.214)' can't be established.
ECDSA key fingerprint is SHA256:hxxpg46rSSg6q6VI9i1SEaXc3j7EAxOeSTqwMgjGSHo.
Are you sure you want to continue connecting (yes/no)? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new keys
Password:

Number of key(s) added: 1

Now try logging into the machine, with: "ssh 'hana01'"
and check to make sure that only the key(s) you wanted were added.
hana02:~/.ssh #

```

Verify the configurations.

Verify the SSH password-free connection service: Log on to the nodes from each other through SSH. If both logon processes do not need a password, the service is successful.

Perform verification on the HANA primary node:

```

hana01:~/.ssh # ssh hana02
Last login: Sat Dec 30 18:29:05 2017 from 192.168.10.207

Welcome to Alibaba Cloud Elastic Compute Service !

hana02:~ #

```

Perform verification on the HANA backup node:

```
hana02:~/.ssh # ssh hana01
Last login: Sat Dec 30 18:26:32 2017 from 192.168.10.207

Welcome to Alibaba Cloud Elastic Compute Service !

hana01:~ #
```

Configure the NTP service

The nodes in the cluster need to synchronize time. In this example, the HANA primary node is configured as the NTP server, and the backup node is configured as the client.

HANA primary node:

```
# vim /etc/ntp.conf
server 127.127.1.0 # local clock (LCL)
fudge 127.127.1.0 stratum 10 # LCL is unsynchronized
# systemctl restart ntpd.service
# ntpq -p
```

```
hana01:~ # ntpq -p
      remote           refid      st t when poll reach  delay  offset  jitter
=====
LOCAL(0)        .LOCL.         10 l   8   64   1   0.000   0.000   0.000
hana01:~ #
```

HANA backup node:

```
# vim /etc/ntp.conf
server hana01 iburst
# systemctl restart ntpd.service
# ntpq -p
```

```
hana02:~ # ntpq -p
      remote           refid      st t when poll reach  delay  offset  jitter
=====
hana01          .INIT.         16 u   -   64   0   0.000   0.000   0.000
hana02:~ #
```

Note: If the offset between local time and NTP server time exceeds 1000 seconds, run the `systemctl stop ntpd.service` command to stop the ntpd service, and run `ntpdate 192.168.20.9` (replace 192.168.20.9 with the NTP server address) to synchronize the time manually. Then run `systemctl start ntpd.service` to restart the ntpd service.

Partition the HANA file system

File system partitioning differences between HANA primary and backup nodes:

HANA Node	File System Partition
HANA primary node	OS disk
HANA primary node	/hana/data
HANA primary node	/hana/log
HANA primary node	/hana/shared
HANA primary node	/hana/backup
HANA primary and backup nodes	Arbitration disk
HANA backup node	OS disk
HANA backup node	/hana/data
HANA backup node	/hana/log
HANA backup node	/hana/shared

Configure the heartbeat network


```
HANA01:~ # yast network
```

HANA primary node:

```
YaST2 - network @ hana01

Network Settings
-Global Options-Overview-Hostname/DNS-Routing-
Name      IP Address      Device  Note
Ethernet Card 0 192.168.10.214 eth0
Ethernet Card 1 192.168.20.19  eth1
```


HANA backup node::




Configure the HaVip primary node

After the HaVip is configured on Alibaba Cloud, the two ECS instances are in backup mode by default. The HaVip cannot be used for communication. It takes effect only after the HaVip primary node is configured. Therefore, you need to configure the HANA primary node to the HaVip primary node. Assign the HaVip to the ENI of the HANA primary node. This IP address is the additional address (or Linux subinterface) of the corresponding ENI.

```
HANA01:~ # yast network
```



After the configuration, the instance bound with HaVip turns into the primary state.



Install the HANA database

Note: HANA primary and backup nodes must have consistent system ID and instance ID. In this example, the system ID is HAN and instance ID is 00.

Check whether hdblcsm is an executable program. Install the HANA instances on the primary and backup nodes.

```
hana01:~/HDB_SERVER_LINUX_X86_64 # ./hdblcsm
```

```
SAP HANA Lifecycle Management - SAP HANA Database 2.00.020.00.1500920972
```

```
*****
```

```
Scanning software locations...
```

```
Detected components:
```

```
SAP HANA Database (2.00.020.00.1500920972) in /root/HDB_SERVER_LINUX_X86_64/server
```

```
Choose an action
```

```
Index | Action | Description
```

```
-----
```

```
1 | install | Install new system
2 | extract_components | Extract components
3 | Exit (do nothing) |
```

```
Enter selected action index [3]: 1
```

```
Enter Installation Path [/hana/shared]:
```

```
Enter Local Host Name [hana01]:
```

```
Do you want to add hosts to the system? (y/n) [n]:
```

```
Enter SAP HANA System ID: HAN
```

```
Enter Instance Number [00]:
```

```
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
Enter Location of Data Volumes [/hana/data/HAN]:
Enter Location of Log Volumes [/hana/log/HAN]:
Restrict maximum memory allocation? [n]:
Enter Certificate Host Name For Host 'hana01' [hana01]:
Enter SAP Host Agent User (sapadm) Password:
Confirm SAP Host Agent User (sapadm) Password:
Enter System Administrator (hanadm) Password:
Confirm System Administrator (hanadm) Password:
Enter System Administrator Home Directory [/usr/sap/HAN/home]:
Enter System Administrator Login Shell [/bin/sh]:
Enter System Administrator User ID [1000]:
Enter ID of User Group (sapsys) [79]:
Enter System Database User (SYSTEM) Password:
Confirm System Database User (SYSTEM) Password:
Restart system after machine reboot? [n]:

```

```
Summary before execution:
```

```
=====
```

```

SAP HANA Database System Installation
Installation Parameters
Remote Execution: ssh
Database Isolation: low
Installation Path: /hana/shared
Local Host Name: hana01
SAP HANA System ID: HAN
Instance Number: 00
Local Host Worker Group: default
System Usage: test
Location of Data Volumes: /hana/data/HAN
Location of Log Volumes: /hana/log/HAN
Certificate Host Names: hana01 -> hana01
System Administrator Home Directory: /usr/sap/HAN/home
System Administrator Login Shell: /bin/sh
System Administrator User ID: 1000
ID of User Group (sapsys): 79
Software Components
SAP HANA Database
Install version 2.00.020.00.1500920972
Location: /root/HDB_SERVER_LINUX_X86_64/server

```

```
Do you want to continue? (y/n): y
```

```

Installing components...
Installing SAP HANA Database...
Preparing package 'Saphostagent Setup'...
Preparing package 'Python Support'...
Preparing package 'Python Runtime'...
Preparing package 'Product Manifest'...
Preparing package 'Binaries'...
Preparing package 'Data Quality'...

```



```
Preparing package 'Krb5 Runtime'...
Preparing package 'Installer'...
Preparing package 'Ini Files'...
Preparing package 'HWCCT'...
Preparing package 'Documentation'...
Preparing package 'Delivery Units'...
Preparing package 'Offline Cockpit'...
Preparing package 'DAT Languages (EN, DE)'...
Preparing package 'DAT Languages (other)'...
Preparing package 'DAT Configfiles (EN, DE)'...
Preparing package 'DAT Configfiles (other)'...
Creating System...
Extracting software...
Installing package 'Saphostagent Setup'...
Installing package 'Python Support'...
Installing package 'Python Runtime'...
Installing package 'Product Manifest'...
Installing package 'Binaries'...
Installing package 'Data Quality'...
Installing package 'Krb5 Runtime'...
Installing package 'Installer'...
Installing package 'Ini Files'...
Installing package 'HWCCT'...
Installing package 'Documentation'...
Installing package 'Delivery Units'...
Installing package 'Offline Cockpit'...
Installing package 'DAT Languages (EN, DE)'...
Installing package 'DAT Languages (other)'...
Installing package 'DAT Configfiles (EN, DE)'...
Installing package 'DAT Configfiles (other)'...
Creating instance...
Installing SAP Host Agent version 7.21.26...
Starting SAP HANA Database system...
Starting 4 processes on host 'hana01' (worker):
Starting on 'hana01': hdbcompilesrv, hdbdaemon, hdbnamesrv, hdbpreprocessor
Starting 7 processes on host 'hana01' (worker):
Starting on 'hana01': hdbcompilesrv, hdbdaemon, hdbindexsrv, hdbnamesrv, hdbpreprocessor,
hdbwebdispatcher, hdbxsengine
Starting on 'hana01': hdbdaemon, hdbindexsrv, hdbwebdispatcher, hdbxsengine
Starting on 'hana01': hdbdaemon, hdbwebdispatcher, hdbxsengine
Starting on 'hana01': hdbdaemon, hdbwebdispatcher
All server processes started on host 'hana01' (worker).
Importing delivery units...
Importing delivery unit HCO_INA_SERVICE
Importing delivery unit HANA_DT_BASE
Importing delivery unit HANA_IDE_CORE
Importing delivery unit HANA_TA_CONFIG
Importing delivery unit HANA_UI_INTEGRATION_SVC
Importing delivery unit HANA_UI_INTEGRATION_CONTENT
Importing delivery unit HANA_XS_BASE
Importing delivery unit HANA_XS_DBUTILS
Importing delivery unit HANA_XS_EDITOR
Importing delivery unit HANA_XS_IDE
Importing delivery unit HANA_XS_LM
Importing delivery unit HDC_ADMIN
Importing delivery unit HDC_BACKUP
```

```

Importing delivery unit HDC_IDE_CORE
Importing delivery unit HDC_SEC_CP
Importing delivery unit HDC_SYS_ADMIN
Importing delivery unit HDC_XS_BASE
Importing delivery unit HDC_XS_LM
Importing delivery unit SAPUI5_1
Importing delivery unit SAP_WATT
Importing delivery unit HANA_SEC_CP
Importing delivery unit HANA_BACKUP
Importing delivery unit HANA_HDBLCM
Importing delivery unit HANA_SEC_BASE
Importing delivery unit HANA_SYS_ADMIN
Importing delivery unit HANA_ADMIN
Importing delivery unit HANA_WKLD_ANLZ
Installing Resident hdbldcm...
Updating SAP HANA Database Instance Integration on Local Host...
Regenerating SSL certificates...
Deploying SAP Host Agent configurations...
Creating Component List...
SAP HANA Database System installed
You can send feedback to SAP with this form: https://hana01:1129/lmsl/HDBLCM/HAN/feedback/feedback.html
Log file written to '/var/tmp/hdb_HAN_hdbldcm_install_2017-12-30_20.55.04/hdbldcm.log' on host 'hana01'.

```

Verify the HANA installation on the primary and backup nodes by checking the HANA process status.

```


hana01:~$ sudo SERVER_LINUX_V06_04 # su hanada
sh-4.3# cd /usr/sap/HAN/
HDB000/ # ps/ hdbldcm_uploads/ home/
sh-4.3# cd /usr/sap/HAN/HDB000
sh-4.3# cd: Not a directory
sh-4.3# cd /usr/sap/HAN/HDB000
sh-4.3# ls
HDB HDBMain.sh HDBSettings.sh HDBSettings.sh backup exe hana01 hdbenv.csh hdbenv.sh work xteras
sh-4.3# ./HDB info
USER      PID  PPID  %CPU  VSZ    RSS COMMAND
hanada    11920 11919 0.0   21068 3972 sh
hanada    12032 11920 12.5  12132 3152 /bin/sh ./HDB info
hanada    12007 12032 0.0   63944 3946 \_ ps fx -u hanada -o user,pid,ppid,pcpu,vsz,rsz,args
hanada    10426 1 0.0   21628 2780 sapstart pf=/hana/shared/HAN/profile/HAN_HDB000_hana01
hanada    10453 10426 0.0 320956 5708 \_ /usr/sap/HAN/HDB000/hana01/trace/hdb.sapHAN_HDB000 -d -nv -f /usr/sap/HAN/HDB000/hana01/daemon.ini pf=/usr/sap/HAN/SIS/profile/HAN_HDB000_hana01
hanada    10469 10453 70.6 9275380 6082952 \_ hdbnameserver
hanada    10444 10453 3.6 2605460 425980 \_ hdbcompleserver
hanada    10646 10453 88.3 5220120 2364604 \_ hdbpreprocessor
hanada    10635 10453 70.1 9470700 6750200 \_ hdbindexserver -port 30003
hanada    10687 10453 3.5 4219736 1281488 \_ hdbvsengine -port 30007
hanada    11038 10453 2.9 2904840 464484 \_ hdbwebdispatcher
hanada    10323 1 0.4 502580 25900 /usr/sap/HAN/HDB000/exe/sapstartsv pf=/hana/shared/HAN/profile/HAN_HDB000_hana01 -D -u hanada
sh-4.3#

```


Install the HANA Studio

Configure a Windows ECS


Double-click the hdbsetup executable file in the Studio installation package.




Click "Next" .



Click "Install" .



Complete the installation and close the program.




Configure HANA system replication


Back up the database

Connect the HANA Studio to the HANA primary node to back up the database.

System-level database backup




Tenant-level database backup




Enable HANA system replication on the primary node

Activate HANA system replication on the primary node.



Maintain the logic names on the primary node.



Register the backup node to the primary node

Copy the PKI SSFS file on the primary node to the corresponding location on the backup node :


```
/usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
/usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
```

Location of the PKI SSFS file on the primary node:


```
/usr/sap/HAN/SYS/global/security/rsecssfs/data
/usr/sap/HAN/SYS/global/security/rsecssfs/key
```

Note: When copying the file, do not delete the original file owner; otherwise, some operations may be failed due to insufficient rights.

Register the backup node on the HANA Studio console.



Maintain the backup node information using the system replication method.



Check the HANA system replication status.

Install and configure SLES 12 Cluster HA

Install the SUSE HAE software


Add the local source to the primary and backup nodes.

```
#zypper addrepo iso:/?iso=/root/SLE-12-SP2-HA-DVD-x86_64-GM-CD1.iso SAP1
#zypper addrepo iso:/?iso=/root/SLE-12-SP2-HA-DVD-x86_64-GM-CD2.iso SAP2
```

Note: The ISO path needs to be adjusted.

```
hana01#yast
```

Select all software packages on the right.



Select the dependent package and click Accept.


```

yast2 - pw_single @ HAN02

Performing Installation
Actions performed:
  Downloading crash (download size 884.5 KiB)
  Downloading release-notes-ha (download size 235.2 KiB)
  Downloading ldirectord (download size 112.3 KiB)
  Downloading hawk2 (download size 20.14 MiB)
  Downloading ha-cluster-bootstrap (download size 31.1 KiB)
  Downloading patterns-ha-ha_sles (download size 8 KiB)
  Downloading yacl2-nagha (download size 60.8 KiB)
  Installing crash-2.2.0.git.1475924149.Bob212-17.6.noarch.rpm (installed size 2.30 MiB)
  Installing release-notes-ha-12.2.20161007-11.13.noarch.rpm (installed size 498.1 KiB)
  Installing ldirectord-3.9.7.git.146192870a.c07c2ba-10.3.x86_64.rpm (installed size 184.7 KiB)
  Installing hawk2-2.0.0.git.1468429505.0129e30-24.32.x86_64.rpm (installed size 52.09 MiB)
Installing Packages... (Remaining: 53.21 MiB, 4 packages)
[help] [Abort]

```

Configure the cluster

Generate the cluster configuration file

Generate the corosync.conf file on the HANA primary node.

```
hana01:~ # yast cluster
```

The configuration is as follows. Other configuration options retain the default values.

```

yast2 - cluster @ hana01

Cluster - Communication Channels
Transport:
  Channel:
    Basic Network Address: 192.168.20.0 (192.168.20.0)
    Multicast Address: 192.168.20.0
    Multicast Port: 4444
    Header Address:
      IP: (Redundant IP (Node ID)
      192.168.20.10 192.168.20.14
      192.168.20.20 192.168.20.25
    Cluster Name: cluster
    Expected Votes: 2
    crp mode: passive
    [x] Auto Generate Node ID
    [Add][Del][Edit]
    [Cancel] [Finish]

Cluster - Configure corosyncd
Corosyncd is a daemon which helps to duplicate firewall status between cluster nodes.
yast2 can help to configure some basic aspects of corosyncd.
You need to start it with the ocf/heartbeat-corosyncd.
Indicated Interface: IP: 192.168.20.24
Multicast Address:
Group Number:
Generate /etc/corosyncd/corosyncd.conf
[Cancel] [Finish]

```

Copy the corosync.conf file to the HANA backup node.

```
hana01# scp /etc/corosync/corosync.conf hana02:/etc/corosync/corosync.conf
```

Start the cluster

Run the following commands on both nodes:

```
# rcpacemaker start
```

View the cluster status.

```
# crm_mon -1
Stack: corosync
Current DC: hana01 (version 1.1.16-4.8-77ea74d) - partition with quorum
Last updated: Tue Nov 7 23:13:06 2017
Last change: Tue Nov 7 23:13:05 2017 by hacluster via crmd on hana01

2 nodes configured
0 resources configured

Online: [ hana01 hana02 ] # Both nodes should be in online state.

No active resources
```

Close the STONITH (which will be configured later).

```
# crm
crm(live)# configure
crm(live)configure# property stonith-enabled=false
crm(live)configure# commit
```


Enable web-based configuration.

(1) Set the HA cluster user password on hana01 to hacluster.

```
# passwd hacluster
New password:
BAD PASSWORD: it is based on a dictionary word
Retype new password:
passwd: password updated successfully

# systemctl restart hawk.service
```

(2) Access <https://192.168.10.214:7630/> (through HANA Studio ECS) with the user name and password hacluster.



Configure the SBD arbitration disk.

Disable the cluster on hana01 and hana02.

```
# rcpacemaker stop
```

View disk information.

```
Hana01:~ # lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sr0 11:0 1 1024M 0 rom
vda 253:0 0 60G 0 disk
└─vda1 253:1 0 60G 0 part /
vdb 253:16 0 1024G 0 disk
└─vdb1 253:17 0 1024G 0 part /hana/data
vdc 253:32 0 512G 0 disk
└─vdc1 253:33 0 512G 0 part /hana/log
vdd 253:48 0 512G 0 disk
└─vdd1 253:49 0 512G 0 part /hana/shared
vde 253:64 0 20G 0 disk
```

The cloud disk vde uses shared block storage..

Configure watchdog on hana01 and hana02.

```
# echo softdog > /etc/modules-load.d/watchdog.conf
# systemctl restart systemd-modules-load
# systemctl status systemd-modules-load
● systemd-modules-load.service - Load Kernel Modules
Loaded: loaded (/usr/lib/systemd/system/systemd-modules-load.service; static; vendor preset: disabled)
Active: active (exited) since Mon 2018-01-01 21:18:41 CST; 8s ago
Docs: man:systemd-modules-load.service(8)
man:modules-load.d(5)
Process: 2300 ExecStart=/usr/lib/systemd/systemd-modules-load (code=exited, status=0/SUCCESS)
Main PID: 2300 (code=exited, status=0/SUCCESS)

Jan 01 21:18:41 s4hsvra systemd[1]: Starting Load Kernel Modules...
Jan 01 21:18:41 s4hsvra systemd[1]: Started Load Kernel Modules.
```

```
# lsmod | grep dog
softdog 16384 0

# vim /etc/init.d/boot.local
modprobe softdog

# vim /etc/sysconfig/sbd
SBD_DEVICE="/dev/vde"
SBD_OPTS="-W"
```

Create the SBD partition on hana01.

```
sbd -d /dev/vde -4 30 -1 15 create

# sbd -d /dev/vde dump
==Dumping header on disk /dev/vde
Header version : 2.1
UUID : 94d700ee-837b-46c7-95cc-27f3d1ffc9f
Number of slots : 255
Sector size : 512
Timeout (watchdog) : 15
Timeout (allocate) : 2
Timeout (loop) : 1
Timeout (msgwait) : 30
==Header on disk /dev/vde is dumped
```

Description of msgwait timeout and watchdog timeout:

-4 indicates the msgwait timeout. In the preceding example, the msgwait timeout interval is 30s.

-1 indicates the watchdog timeout. In the preceding example, the watchdog timeout interval is 15s.

The minimum value for simulated package detection is 15s.

If SBD stays in the multi-path group, the timeout interval required by SBD needs to be modified, because the MPIO detection along the path is time consuming. If msgwait times out, it is assumed that the message has been transmitted to the target node. For multi-path, the delay is the time for switching to the next path when MPIO detects a path failure. You may need to test this function in your system environment. If the SBD on the node does not reset the package detection timer in time, the node is automatically stopped. The watchdog timeout interval must be shorter than the msgwait timeout interval. The former should be a half of the later.

The following formula expresses the relationships between the three values:

Timeout (msgwait) = (Timeout (watchdog) * 2)

stonith-timeout = Timeout (msgwait) + 20%

For more information, run the man sbd command.

Configure SBD on hana01 and hana02.

```
# vim /etc/sysconfig/sbd
SBD_DEVICE="/dev/vde"
SBD_OPTS="-W"
```

The SBD program automatically starts on hana01 and hana02 when the system is started.

```
# systemctl enable sbd
```

Enable the cluster on hana01 and hana02.

```
# rcpacemaker start
```

Modify the cluster SBD parameters on hana01.

```
# crm configure
crm(live)configure# primitive stonith_sbd stonith:external/sbd params pcmk_delay_max=30
crm(live)configure# commit
crm(live)configure# exit
```

Enable STONITH.

```
# crm configure
crm(live)configure# property stonith-enabled="true"
crm(live)configure# property stonith-timeout="40s"
crm(live)configure# property no-quorum-policy="ignore"
crm(live)configure# property default-resource-stickiness="1000"
crm(live)configure# commit
crm(live)configure# exit
```

Note: We recommend that you set stonith-timeout to 40s (calculated based on the previous formula).

View the SBD process and service on hana01 and hana02.

```
# ps -ef | grep sbd
root 5946 1 0 22:44 ? 00:00:01 sbd: inquisitor
root 5947 5946 0 22:44 ? 00:00:00 sbd: watcher: /dev/vde - slot: 0 - uuid: 94d700ee-837b-46c7-95cc-27f3d1ffc9f
root 5948 5946 0 22:44 ? 00:00:01 sbd: watcher: Pacemaker
root 5949 5946 0 22:44 ? 00:00:00 sbd: watcher: Cluster
root 6915 2540 0 23:25 pts/0 00:00:00 grep --color=auto sbd

# systemctl status sbd
● sbd.service - Shared-storage based fencing daemon
Loaded: loaded (/usr/lib/systemd/system/sbd.service; enabled; vendor preset: disabled)
Active: active (running) since Tue 2017-12-26 22:44:51 CST; 41min ago
Process: 5934 ExecStart=/usr/sbin/sbd $SBD_OPTS -p /var/run/sbd.pid watch (code=exited, status=0/SUCCESS)
Main PID: 5946 (sbd)
Tasks: 4 (limit: 512)
CGroup: /system.slice/sbd.service
├─5946 sbd: inquisitor
├─5947 sbd: watcher: /dev/vde - slot: 0 - uuid: 94d700ee-837b-46c7-95cc-27f3d1ffc9f
└─5948 sbd: watcher: Pacemaker
```

```
└─5949 sbd: watcher: Cluster
```

```
Dec 26 22:44:50 node001 systemd[1]: Starting Shared-storage based fencing daemon...
```

```
Dec 26 22:44:51 node001 systemd[1]: Started Shared-storage based fencing daemon.
```

Verify the SBD configuration.

Note: Ensure that important processes on hana02 have been closed.

```
Hana01# sbd -d /dev/vde message hana02 reset
```

If hana2 is restarted properly, the SBD disk is successfully configured.

Integrate SAP HANA with SUSE HA

Add SAP HANA resources

Open the SUSE Hawk management interface, click Wizards, and select the HSR options to maintain HANA information.

The screenshot shows the SUSE Hawk management interface. On the left is a dark sidebar with a menu. The main area displays the 'SAPHanaSR' configuration wizard. The wizard has a progress indicator at the top with a blue circle containing the number '1'. Below this, a light gray box contains the text 'Create a single primitive resource of type IPAddr2.' The configuration fields are as follows:

Field	Value	Status
sid	HAN	✓
SAP SID		
ino	00	✓
SAP Instance Number		
ip	192.168.10.12	✓
Virtual IP Address		

At the bottom of the form are three buttons: 'Cancel', 'Back', and 'Verify'.

The script after successful configuration is as follows:

Changes applied successfully.

1. Configure cluster resources

```
primitive rsc_ip_HAN_HDB000 ocf:heartbeat:IPaddr2
params
ip="192.168.10.216"
op start timeout="20" op stop timeout="20"
op monitor interval="10" timeout="20"

primitive rsc_SAPHanaTopology_HAN_HDB000 ocf:suse:SAPHanaTopology
params
SID="HAN"
InstanceNumber="00"
op monitor interval="10" timeout="600"
op start interval="0" timeout="600"
op stop interval="0" timeout="300"

primitive rsc_SAPHana_HAN_HDB000 ocf:suse:SAPHana
params
SID="HAN"
InstanceNumber="00"
PREFER_SITE_TAKEOVER="yes"
AUTOMATED_REGISTER="no"
DUPLICATE_PRIMARY_TIMEOUT="7200"
op start interval="0" timeout="3600"
op stop interval="0" timeout="3600"
op promote interval="0" timeout="3600"
op monitor interval="60" role="Master" timeout="700"
op monitor interval="61" role="Slave" timeout="700"

ms msl_SAPHana_HAN_HDB000 rsc_SAPHana_HAN_HDB000
meta clone-max="2" clone-node-max="1" interleave="true"

clone cln_SAPHanaTopology_HAN_HDB000 rsc_SAPHanaTopology_HAN_HDB000
meta is-managed="true" clone-node-max="1" interleave="true"

colocation col_saphana_ip_HAN_HDB000 2000: rsc_ip_HAN_HDB000:Started msl_SAPHana_HAN_HDB000

```

Verify the cluster status

Check the resource status in the cluster.


SUSE Hawk

MANAGE
Status
Dashboard
History
CONFIGURATION
Add Resource
Add Constraint
Wizards
Edit Configuration
Cluster Configuration
Command Log
ACCESS CONTROL
Roles
Targets



Resources (7)		Nodes (2)	
Status	Name	Location	
+	cln_SAPHanaTopology_HAN_HDB000	hana01, hana02	
+	msl_SAPHana_HAN_HDB000	hana01, hana02	
+	rsc_ip_HAN_HDB000	hana01	
+	sap_ip	hana01	
+	stonith_sbd	hana02	

Check the node status in the cluster.




Test the SAP HANA HA failover


Test the HANA primary node failure

Ensure that HA works normally before the test.


Check the HANA node status.




Check the HANA system replication status.



Check the cluster node column information.



Check the service status in the cluster.



Test the primary node failover.

Forcibly close hana01 on the ECS console..

Stop

?

Operation will be executed on the selected **1 Instances** . Are you sure you want to proceed?

Stopped By:

☐ Stop
 ☒ Force Stop

If you use force stop mode, all unsaved data on the specified instance will be lost.

☒ I confirm to forcibly Stop the selected instances.

Stopped By:

☐ Keep Stopped Instances and Continue Billing

Operation will be executed on the selected **1 Instances** . The instances will not be billed after being stopped.

Stopping Subscription based instances does not change their expiration date.

If you need to stop an instance for system disk replacement, disk reinitialization, instance upgrade, or private IP address modification, we recommend that you select Keep Stopped Instances and Continue Billing to avoid startup failure.

OK

Cancel

VIP(sap_ip) floats to hana02.

SUSE Hawk

Batch hcluster Help ? Logout

MANAGE

Status

Dashboard

History

CONFIGURATION

Add Resource

Add Constraint

Wizards

Edit Configuration

Cluster Configuration

Command Log

Dashboard

+ Add Cluster

Local Status

sap_ip			
stonith_sbd			
res_ip_HANA_HDB00			
msl_SAPHana_HANA_HDB00			
cln_SAPHanaTopology_HANA_HDB00			

Slave node, taking over

HANA status after failover.

Systems

HANA

HANA (SYSTEM) Hana VIP

Backup

Catalog

Content

Provisioning

Security

HANA (SYSTEM) Standby Node

Backup

Catalog

Content

Provisioning

Security

Old primary node cannot be connected

Old backup node can be connected

HANA (SYSTEM) Hana VIP

192.168.10.12 00

Last Update: 2018-1-7 10:00:56

Interval: 60 Seconds

Overview

Logspace

Alerts

Performance

Volumes

Configuration

System Information

Diagnosis Files

Trace Configuration

Enter your filter

Visible rows: 0/0

Filters

Save as File

HOST	SECONDARY_HOST	REPLICATION_MODE	REPLICATION_STATUS	REPLICATION_STATUS_DETAILS	PORT	VOLUME_ID	SIZE_TB
No more system replication status							

Properties

Error Log

Logspace Log

Type filter text


Message

Recover the HANA primary node


Test the recovery of the HANA primary node.

Enable the ECS of hana01 and start the cluster software pacemaker.


```
hana01:~ # rcpacemaker start
```




Configure HSR on the console..




Register hana01 as the backup node..




Set the synchronization mode to **syncmem**.




Check the HANA node status.



Check the HSR copy status.



Check the HAE cluster status and clean up the nodes reporting errors. After cleanup, the cluster is recovered.



The HA cluster recovers, and the HANA backup node starts to provide services.



SAP HANA HA Cross-Zone with SLES HAE

SAP HANA HA Cross-Zone solution on SUSE Linux Enterprise Server for SAP Applications

- Solution Overview
 - SAP HANA System Replication
 - HAE of SUSE Linux Enterprise Server for SAP Applications
 - Architecture Overview
 - Network Design
- Infrastructure Preparation
 - Infrastructure List
 - Creating VPC
 - Creating ECS Instances
 - Creating ENIs and binding to ECS instances
 - Creating NAT Gateway and configure SNAT entry
 - Creating STONITH device and Virtual IP Resource Agent
- Software Preparation
 - Software List
 - HAE Installation
 - SAP HANA Installation
 - SAP Host Agent Installation
- Configuring SAP HANA System Replication
- Configuring HAE for SAP HANA
 - Configuration of Corosync
 - Configuration of Pacemaker
 - Verify the HA takeover
- Example
 - Example Cluster Configuration
 - Example for `/etc/corosync/corosync.conf`
- Reference

Version Control:

Version	Revision Date	Types Of Changes	Effective Date
1.0			2018/3/7
1.1	2018/7/04	Add corosync and cluster configuration example	2018/7/04

Solution Overview

SAP HANA System Replication

SAP HANA provides a feature called System Replication which is available in every SAP HANA installation offering an inherent disaster recovery support.

For details, please refer to SAP Help Portal HANA system replication.

HAE of SUSE Linux Enterprise Server for SAP Applications


SUSE High Availability Extension (HAE) is a high availability solution based on Corosync and Pacemaker. With SUSE Linux Enterprise Server for SAP Applications, SUSE provides SAP specific Resource Agents (SAPHana, SAPHanaTopology etc.) used by Pacemaker to help users to buildup SAP HANA HA solution more effectively.

For details, please refer to latest version of SAP HANA SR Performance Optimized Scenario at SUSE documentation center.

Architecture Overview

This document guides you on how to deploy a SAP HANA HA solution cross different Zones. Following is a brief architecture:

- HAE of SUSE Linux Enterprise Server for SAP Applications is used to setup the HA Cluster;
- SAP HANA System Replication is activated between the two HANA nodes;
- Two HANA nodes locates in different Zones of the same Region;
- Alibaba Cloud Specific Virtual IP Resource Agent is used to allow Moving IP automatically switched to Active SAP HANA node; Alibaba Cloud specific STONITH device is used for fencing;



* Virtual IP – Alibaba Cloud provides a specific Resource Agent to switch Service IP across different ECS instances locates in different zones;

Network Desian

network	location	usage	subnet
---------	----------	-------	--------

business	eu-central-1 zone A	For Business	192.168.0.82/24
heartbeat	eu-central-1 zone B	For SR/HA	192.168.1.245/24

hostname	role	heartbeat IP	business IP	virtual IP
hana0	Hana primary node	192.168.0.83	192.168.0.82	192.168.4.1
hana1	Hana secondary node	192.168.1.246	192.168.1.245	192.168.4.1
HanaStudio	Hana Studio	no	192.168.0.79	no

Infrastructure Preparation


Infrastructure List

- 1 VPC network;
- 2 ECS instances in different zones of the same VPC;
- 2 Elastic Network Interface (ENI), one for each ECS instance;
- Alibaba Cloud specific Virtual IP Resource Agent and STONITH device;
- NAT Gateway and SNAT entry;

Creating VPC


First of all, a VPC should be created.

In this example, we create a VPC named `suse_hana_ha` in EU Central 1 (frankfurt) Region as follow:



There should be at least 2 VSwitches(subnets) defined within the VPC network, each VSwitch bound to a different Zone. In this example, we have following 2 VSwitches(subnets):


- Switch1 192.168.0.0/24 Zone A, for SAP HANA Primary Node
- Switch2 192.168.1.0/24 Zone B, for SAP HANA Secondary Node



Creating ECS Instances

Two ECS instances are created in different Zones of the same VPC. Choose the "SUSE Linux Enterprise Server for SAP Applications" image from the Image Market place.

In this example, 2 ECS instances (hostname: hana0 and hana1) are created in eu-central-1 Region, Zone A and Zone B, within VPC: suse_hana_ha, with SUSE Linux Enterprise Server for SAP Applications 12 SP2 image from the Image Market Place. Host hana0 is the primary SAP HANA database node, and hana1 is the secondary SAP HANA database node.



Creating ENIs and binding to ECS instances

Create two ENIs, and attach one for each ECS instance, for HANA System Replication purpose. Configure the IP addresses of the ENIs to the subnet for HANA System Replication only.


In this example, the ENIs are attached to ECS instances hana0 and hana1, and IP addresses are configured as 192.168.0.83 and 192.168.1.246 within the same VSwitches of hana0 and hana1, and put in the VPC: suse_hana_ha

ENI ID	VPC	Status	Destination CIDR Block	Default VSwitch	Zone	Number of Available Private IPs	Actions
eni-gw8q3m4f8a9q9h155	vpc-gw8q3m4f8a9q9h155	Available	192.168.1.0/24	No	EU Central 1 Zone B	248	Manage Delete Purchase
eni-gw8q3m4f8a9q9h156	vpc-gw8q3m4f8a9q9h155	Available	192.168.1.0/24	No	EU Central 1 Zone A	247	Manage Delete Purchase

Meanwhile, within the Guest OS, /etc/hosts should also be configured as well. In this example, please run following two commands on both sites:


```
echo "192.168.0.82 hana0 hana0" >> /etc/hosts
```

```
echo "192.168.1.246 hana1 hana1" >> /etc/hosts
```




Creating NAT Gateway and configure SNAT entry

First of all, create a NAT Gateway attached to the given VPC; In our example, we create a NAT Gateway named suse_hana_ha_GW as follows:



After creating NAT Gateway, you need to create corresponding SNAT entry to allow ECS instances within the VPC can access public address on Internet. (Caution: Alibaba Cloud specific STONITH device and Virtual IP Resource Agent are mandatory for cluster and they need to access Alibaba Cloud OpenAPI through a public domain);

In our example, we create to two SNAT entries, for ECS instances locates in different network range as follows:



Creating STONITH device and Virtual IP Resource Agent

Download software from with following command:

`wget http://repository-iso.oss-cn-beijing.aliyuncs.com/ha/aliyun-ecs-pacemaker.tar.gz`

```
hanal:/hana/tmp # ls
122.05  122.05.rpm  SAPHOSTAGENT  sap-host-agent  aliyun  corosync.conf
hanal:/hana/tmp # wget http://repository-iso.oss-cn-beijing.aliyuncs.com/ha/aliyun-ecs-pacemaker.tar.gz
--2018-03-05 16:12:13-- http://repository-iso.oss-cn-beijing.aliyuncs.com/ha/aliyun-ecs-pacemaker.tar.gz
Resolving repository-iso.oss-cn-beijing.aliyuncs.com (repository-iso.oss-cn-beijing.aliyuncs.com)... 59.110.190.40
Connecting to repository-iso.oss-cn-beijing.aliyuncs.com (repository-iso.oss-cn-beijing.aliyuncs.com) [59.110.190.40]:80
... connected.
HTTP request sent, awaiting response... 200 OK
Length: 4125 (4.0K) [application/x-gzip]
Saving to: 'aliyun-ecs-pacemaker.tar.gz'

100%[=====] 4,125  --.-K/s  in 0s

2018-03-05 16:12:13 (503 MB/s) - 'aliyun-ecs-pacemaker.tar.gz' saved [4125/4125]

hanal:/hana/tmp #
```

Extract the package and install the software

`tar -xvf aliyun-ecs-pacemaker.tar.gz`

`./install`

```

hana1:/hana/tmp # ls -l
total 3102992
drwxrwxr-x 8 root root      4096 Mar 30  2017 122.05
-rw-r--r-- 1 root root 3177440026 Dec 24 18:14 122.05.tgz
drwxr-xr-x 2 root root      4096 Feb 28 12:14 SAPHOSTAGENT
drwxrwxrwx 4 root root      131 Feb 28 12:14 SAP_HANA_CLIENT
drwxr-xr-x 3 root root        69 Mar  5 15:31 aliyun
-rw-r--r-- 1 root root      4125 Mar  2 18:24 aliyun-ecs-pacemaker.tar.gz
-rw-r--r-- 1 root root     3957 Mar  1 10:29 corosync.conf
hana1:/hana/tmp # tar -xvf aliyun-ecs-pacemaker.tar.gz
aliyun-ecs-pacemaker/ecs-pacemaker/
aliyun-ecs-pacemaker/ecs-pacemaker/fence_aliyun.py
aliyun-ecs-pacemaker/ecs-pacemaker/vpc-move-ip
aliyun-ecs-pacemaker/install.sh
hana1:/hana/tmp # ./install

```

Install Alibaba Cloud OpenAPI SDK

pip install aliyun-python-sdk-ecs aliyun-python-sdk-vpc aliyuncli

```

hana1:~ # pip install aliyun-python-sdk-ecs aliyun-python-sdk-vpc aliyuncli
Collecting aliyun-python-sdk-ecs
  Downloading http://mirrors.aliyun.com/pypi/packages/02/f6/f63a3ba3357613caade16cf4fa69178eaaccc9cf2fc6692368024499083/aliyun-python-sdk-ecs-4.6.4.tar.gz (41kB)
    100% |#####| 51kB 31.4MB/s
Collecting aliyun-python-sdk-vpc
  Downloading http://mirrors.aliyun.com/pypi/packages/39/eb/ec999243088e6d59710f9e634ca4b43c8ff895f256aa5e3a5ca9b2303982/aliyun-python-sdk-vpc-3.0.2.tar.gz
Collecting aliyuncli
  Downloading http://mirrors.aliyun.com/pypi/packages/f1/6c/ecb0313299cc8f886b75d4a117c6ead329ec89562fb1b82e1c6f083d3c5a/aliyuncli-2.1.9-py2-none-any.whl (69kB)
    100% |#####| 71kB 41.4MB/s
Collecting aliyun-python-sdk-core>=2.0.2 (from aliyun-python-sdk-ecs)
  Downloading http://mirrors.aliyun.com/pypi/packages/fa/28/045cf5cc1e80cc482c3bd1979287e62db73b3c42ef0e429f2d0d8dd2bfbfe/aliyun-python-sdk-core-2.8.1.tar.gz
Collecting colorama<0.3.3,>=0.2.5 (from aliyuncli)
  Downloading http://mirrors.aliyun.com/pypi/packages/24/84/29ce4167d1f5c4a320aaad91e1178e5a1baf9cfe1c63e6a75f4d4c323cf/aliyun-python-sdk-core-2.8.1.tar.gz

```

Configure Alibaba Cloud OpenAPI SDK and Client


aliyuncli configure

```

hana1:~ #
hana1:~ # aliyuncli configure
Aliyun Access Key ID [None]: LTAIPYT3XsfmFFNs
Aliyun Access Key Secret [None]: wjM285NLWOS8dx32bewKwPIjTipHUf
Default Region Id [None]: cn-hangzhou
Default output format [None]: table
hana1:~ #

```

You can get your Access Key from following:



Software Preparation

Software List

- SUSE Linux Enterprise Server for SAP Applications 12 SP2
- HANA Installation Media
- SAP Host Agent Installation Media

HAE installation

Both ECS instances are created with the SUSE Linux Enterprise Server for SAP Applications image. Both ECS instances should install the HAE component, as well as package SAPHanaSR. In this example, we install HAE (major software component: Corosync and Pacemaker), and SAPHanaSR on both ECS instances as follows:

Install the pattern High Availability on both nodes. To do so, for example, use zypper :

```
zypper in -t pattern ha_sles
```


Now the Resource Agents for controlling the SAP HANA system replication needs to be installed at both cluster nodes.

```
zypper in SAPHanaSR SAPHanaSR-doc
```

SAP HANA Installation

Install SAP HANA software on both ECS instances, and make sure the SAP HANA SID and Instance Number are the same (requirement by SAP HANA System Replication). It is recommended to use hdblcm to do the installation. For details please refer to [SAP HANA Server Installation and Update Guide](#).

In this example, both node are installed with SAP HANA (Rev. 2.00.030.00), and SID: **JL0**, Instance Number: **00**.



The image displays two terminal windows side-by-side, both running PuTTY. The top window is titled '192.168.0.82 - PuTTY' and shows a command prompt on 'hana0'. The user 'j10adm' has executed 'su - j10adm' and then 'HDB version'. The output shows the HDB version information for the system.

```
hana0:/ # su - j10adm
j10adm@hana0:/usr/sap/JL0/HDB00> HDB version
HDB version info:
version:          2.00.030.00.1522209842
branch:           fa/hana2sp03
machine config:   linuxx86_64
git hash:         c6e59dfe8322e948bc9bec3b1ef59da7708f0f4e
git merge time:   2018-03-28 06:04:02
weekstone:        0000.00.0
cloud edition:    0000.00.00
compile date:     2018-03-28 06:35:01
compile host:     1d4553
compile type:     rel

j10adm@hana0:/usr/sap/JL0/HDB00>
```

The bottom window is titled '192.168.1.245 - PuTTY' and shows a command prompt on 'hana1'. The user 'j10adm' has executed 'su - j10adm' and then 'HDB version'. The output shows the HDB version information for the system.


```
hana1:/ # su - j10adm
j10adm@hana1:/usr/sap/JL0/HDB00> HDB version
HDB version info:
version:          2.00.030.00.1522209842
branch:           fa/hana2sp03
machine config:   linuxx86_64
git hash:         c6e59dfe8322e948bc9bec3b1ef59da7708f0f4e
git merge time:   2018-03-28 06:04:02
weekstone:        0000.00.0
cloud edition:    0000.00.00
compile date:     2018-03-28 06:35:01
compile host:     1d4553
compile type:     rel

j10adm@hana1:/usr/sap/JL0/HDB00>
```

SAP Host Agent Installation

When you have finished hana installation with hdblcsm as mentioned above, the SAP Host Agent should already be installed on your server. In case you want to install it manually, please kindly refer to [Installing SAP Host Agent Manually](#).

In this example, you can check SAP Host Agent status after you have SAP HANA with hdblcsm installed on hana0 and hana1 as follows:



```
192.168.0.82 - PuTTY
hana0:~ # /usr/sap/hostctrl/exe/saphostexec -status
saphostexec running (pid = 1620)
sapstartsrv running (pid = 1735)
saposcol running (pid = 1792)
hana0:~ #
```

Configuring SAP HANA System Replication

Backup HANA on primary ECS instance

To do backup on HANA, you can either use SAP HANA studio or hdbsql as the client command tool.

The backup command is:


For HANA 1 single container mode:

```
BACKUP DATA USING FILE('COMPLETE_DATA_BACKUP');
```

For HANA 2 with multitenant as default mode (You should backup systemDB and also all tenantDB as shown below in our example):

```
BACKUP DATA for <DATABASE> using FILE('COMPLETE_DATA_BACKUP')
```

In this example, we execute SAP HANA database backup on both ECS instances as follows:



Configuring SAP HANA System Replication on primary node

- Log onto the primary node with: `su - <sid>adm;`
[sidadm] should be replaced by your SAP HANA database SID. In our example it is `su - jl0adm;`
- Stop HANA with: `HDB stop;`
- Change following file content as user root:
`/hana/shared/<SID>/global/hdb/custom/config/global.ini>/global/hdb/custom/config/global.ini``

Add following content:

```

[system_replication_hostname_resolution]
<IP> = <HOSTNAME>
  
```

[IP]] should be address of the ENI (heartbeat IP address for HANA system replication) attached to the Secondary node;

[HOSTNAME] should be hostname of the Secondary node;

In this example, we have following configuration:

```

[system_replication_hostname_resolution]
192.168.1.246 = hana1
  
```

Configuring SAP HANA System Replication on secondary node

Same as above for primary, but use IP and hostname of primary node

In this example, we have following configuration:

```

[system_replication_hostname_resolution]
  
```

192.168.0.83 = hana0

Enable SAP HANA System Replication on primary node

a) Log onto the primary node with: su - <sid>adm;

b) Start HANA with: HDB start;

c) Enable System Replication with:

hdbnsutil -sr_enable --name= [primary location name]

[primary location name] should be replaced by location of your primary HANA node.

In this example, we use following command:

hdbnsutil -sr_enable --name=hana0

CAUTION: all above operations are done on primary node.

Register the Secondary node to the Primary HANA node

a) Log onto the secondary node with: su - <sid>adm;

b) Stop HANA with: HDB stop;

c) Register the Secondary HANA node to the Primary HANA node by running following command:

hdbnsutil -sr_register --remoteHost=[location of primary Node] --remoteInstance=[instance number of primary node] --replicationMode=sync --name=[location of the secondary node] --operationMode=logreplay

In this example, we use following command:

hdbnsutil -sr_register --name=hana1 --remoteHost=hana0 --remoteInstance=00 --replicationMode=sync --operationMode=logreplay

d) Start HANA with: HDB start;

e) Verify the System Replication Status with:

hdbnsutil -sr_state

In this example, we have following status on secondary HANA node hana1:


```
j10adm@hana1:/usr/sap/JL0/HDB00> hdbnsutil -sr_state

System Replication State
~~~~~

online: true

mode: sync
operation mode: logreplay
site id: 2
site name: hana1

is source system: false
is secondary/consumer system: true
has secondaries/consumers attached: false
is a takeover active: false
active primary site: 1

primary masters: hana0

Host Mappings:
~~~~~

hana1 -> [hana1] hana1
hana1 -> [hana0] hana0

Site Mappings:
~~~~~
hana0 (primary/primary)
    |--hana1 (sync/logreplay)

Tier of hana0: 1
Tier of hana1: 2

Replication mode of hana0: primary
Replication mode of hana1: sync

Operation mode of hana0: primary
Operation mode of hana1: logreplay

Mapping: hana0 -> hana1
done.
```

CAUTION: all above operations are done on secondary node.

Configuring HAE for SAP HANA

STONITH: fence_aliyun

For a HA solution, a fencing device is a must. Alibaba Cloud provides its own STONITH device, which

allows the servers in the HA cluster to shut down the other which is not responsible. The STONITH device leverage Alibaba Cloud OpenAPI underneath the ECS instance, which is similar to a physical reset / shutdown on a on-premise environment.


Configuration of Corosync

It is desirable that, you add more redundancy for messaging (Heartbeat) by using separate ENIs attached to the ECS instances with separate network range. On Alibaba Cloud, it is strongly suggested that, only using Unicast for the transport setting in Corosync. Follow the following steps to configure Corosync:

Create Keys

Run `corosync-keygen` on primary HANA node. The generated key will be located in the file: `/etc/corosync/authkey`.

In our example, we execute the command on hana1:



```
192.168.0.82 - PuTTY
hana0:/etc/corosync # ls
authkey  corosync.conf  corosync.conf.final  uidgid.d
hana0:/etc/corosync #
```

Configure `/etc/corosync/corosync.conf` with following content as root on primary HANA node:

```
totem {
  version: 2
  token: 5000
  token_retransmits_before_loss_const: 6
  secauth: on
  crypto_hash: sha1
  crypto_cipher: aes256
  clear_node_high_bit: yes
  interface {
    ringnumber: 0
    bindnetaddr: **IP-address-for-heart-beating-for-the-current-server**
    mcastport: 5405
    ttl: 1
  }
  # On Alibaba Cloud, transport should be set to udpu, means: unicast
  transport: udpu
}
```

```

}
logging {
  fileline: off
  to_logfile: yes
  to_syslog: yes
  logfile: /var/log/cluster/corosync.log
  debug: off
  timestamp: on
  logger_subsys {
    subsys: QUORUM
    debug: off
  }
}
nodelist {
  node {
    ring0_addr: **ip-node-1**
    nodeid: 1
  }
  node {
    ring0_addr: **ip-node-2**
    nodeid: 2
  }
}
quorum {
  # Enable and configure quorum subsystem (default: off)
  # see also corosync.conf.5 and votequorum.5
  provider: corosync_votequorum
  expected_votes: 2
  two_node: 1
}

```

IP-address-for-heart-beating-for-the-current-server should be replaced by the IP address of the current server, used for messaging (heartbeat) or HANA System Replication. In our example, we use IP address of ENI of the current node (192.168.0.83 for hana0 and 192.168.1.246 for hana1); **Caution:** this value will be different on primary and secondary node. **nodelist** directive is used to list all nodes in the cluster.

ip-node-1 and **ip-node-2** should be replaced by the IP addresses of the ENIs attached to ECS instances for Heartbeat Purpose or HANA System Replication Purpose (in this example it should be 192.168.0.83 for hana0 and 192.168.1.246 for hana1).

After completing edit of `/etc/corosync/corosync.conf` on primary HANA node, copy the `/etc/corosync/authkey` and `/etc/corosync/corosync.conf` to `/etc/corosync` on the secondary HANA node with following command:

```

scp /etc/corosync/authkey root@hostnameOfSecondaryNode:/etc/corosync
scp /etc/corosync/corosync.conf root@hostnameOfSecondaryNode:/etc/corosync

```

In our example, we execute following command:

```

hana1:/ # scp /etc/corosync/authkey root@hana2:/etc/corosync
hana1:/ # scp /etc/corosync/corosync.conf root@hana2:/etc/corosync

```

After copy the `corosync.conf` to the secondary node, please kindly configure the `bindnetaddr` as

above to the local heart beating IP address.

Configuration of pacemaker

For SAP HANA HA solution, we need to configure 7 Resource Agents and corresponding constraints in Pacemaker.

CAUTION the following pacemaker configuration only need to be done on one node (normally primary node).

1. Cluster bootstrap and more

Add configuration of bootstrap and default setting of resource and operations to the cluster; Save following scripts in a file: crm-bs.txt

```
property $id='cib-bootstrap-options' \
stonith-enabled="true" \
stonith-action="off" \
stonith-timeout="150s"
rsc_defaults $id="rsc-options" \
resource-stickiness="1000" \
migration-threshold="5000"
op_defaults $id="op-options" \
timeout="600"
```

Execute command to add setting to the cluster:

crm configure load update crm-bs.txt

2. STONITH device

This part defines Aliyun STONITH devices in the cluster;

Save following scripts in a file: crm-stonith.txt

```
primitive res_ALIYUN_STONITH_1 stonith:fence_aliyun \
op monitor interval=120 timeout=60 \
params pcmk_host_list=<primary node hostname> port=<primary node instance id> \
access_key=<access key> secret_key=<secret key> \
region=<region> \
meta target-role=Started
primitive res_ALIYUN_STONITH_2 stonith:fence_aliyun \
op monitor interval=120 timeout=60 \
params pcmk_host_list=<secondary node hostname> port=<secondary node instance id> \
access_key=<access key> secret_key=<secret key> \
region=<region> \
meta target-role=Started
location loc_<primary node hostname>_stonith_not_on_<primary node hostname>
res_ALIYUN_STONITH_1 -inf: <primary node hostname>
#Stonith 1 should not run on primary node because it is controlling primary node
location loc_<secondary node hostname>_stonith_not_on_<secondary node hostname>
res_ALIYUN_STONITH_2 -inf: <secondary node hostname>
#Stonith 2 should not run on secondary node because it is controlling secondary node
```

[secondary node hostname] / **[primary node hostname]** should be replaced by the real hostname of your secondary node;

[secondary node instance id] / **[secondary node instance id]** should be replaced by the real instance-id of your secondary node; you can get this from the console;

[access key] should be replaced with real access key;

[secret key] should be replaced with real secret key;

[region] should be replaced with real region name where the node locates;

Execute command to add the resource to the cluster:

crm configure load update crm-stonith.txt

3. SAPHanaTopology

This part defines a SAPHanaTopology RA, and a clone of SAPHanaTopology on both nodes in the cluster. Save following scripts in a file: crm-saphanatop.txt

```
primitive rsc_SAPHanaTopology_<SID>_HDB<instance number> ocf:suse:SAPHanaTopology \
operations $id="rsc_SAPHanaTopology_<SID>_HDB<instance number>-operations" \
op monitor interval="10" timeout="600" \
op start interval="0" timeout="600" \
op stop interval="0" timeout="300" \
params SID="<SID>" InstanceNumber="<instance number>"
clone cln_SAPHanaTopology_<SID>_HDB<instance number>
rsc_SAPHanaTopology_<SID>_HDB<instance number> \
meta clone-node-max="1" interleave="true"
```

[SID] should be replaced by the real SAP HANA SID;

[instance number] should be replaced by the real SAP HANA Instance Number;

Execute command to add resources to the cluster:

crm configure load update crm-saphanatop.txt

4. SAPHana

This part defines a SAPHana RA, and a Multi-state resource of SAPHana on both nodes in the cluster. Save following scripts in a file: crm-saphana.txt

```
primitive rsc_SAPHana_<SID>_HDB<instance number> ocf:suse:SAPHana \
operatoins $id="rsc_sap_<SID>_HDB<instance number>-operations" \
op start interval="0" timeout="3600" \
op stop interval="0" timeout="3600" \
op promote interval="0" timeout="3600" \
op monitor interval="60" role="Master" timeout="700" \
op monitor interval="61" role="Slave" timeout="700" \
params SID="<SID>" InstanceNumber="<instance number>" PREFER_SITE_TAKEOVER="true" \
DUPLICATE_PRIMARY_TIMEOUT="7200" AUTOMATED_REGISTER="false"
ms msl_SAPHana_<SID>_HDB<instance number> rsc_SAPHana_<SID>_HDB<instance number> \
meta clone-max="2" clone-node-max="1" interleave="true"
```

[SID] should be replaced by the real SAP HANA SID;

[instance number] should be replaced by the real SAP HANA Instance Number;

Execute command to add resources to the cluster:

crm configure load update crm-saphana.txt

5. Virtual IP

This part defines a Virtual IP RA in the cluster. Save following scripts in a file: crm-vip.txt.

```
primitive res_vip_<SID>_HDB<instance number> ocf:aliyun:vpc-move-ip \
op monitor interval=60 \
meta target-role=Started \
params address=<virtual_IPv4_address> routing_table=<route_table_ID> interface=eth0
```

[virtual_IP4_address] should be replaced by the real IP address you prefer to provide service;

[route_table_ID] should be replaced by the route table ID of your VPC;

[SID] should be replaced by the real SAP HANA SID;

[instance number] should be replaced by the real SAP HANA Instance Number;

Execute command to add the resource to the cluster:

crm configure load update crm-vip.txt

6. Constraints

Two constraints are organizing the correct placement of the virtual IP address for the client database access and the start order between the two resource agents SAPHana and SAPHanaTopology. Save following scripts in a file: crm-constraint.txt

```
colocation col_SAPHana_vip_<SID>_HDB<instance number> 2000: rsc_vip_<SID>_HDB<instance number>:started \
msl_SAPHana_<SID>_HDB<instance number>:Master
order ord_SAPHana_<SID>_HDB<instance number> Optional:
cln_SAPHanaTopology_<SID>_HDB<instance number> \
msl_SAPHana_<SID>_HDB<instance number>
```

[SID] should be replaced by the real SAP HANA SID;

[instance number] should be replaced by the real SAP HANA Instance Number;

Execute command to add the resource to the cluster:

crm configure load update crm-constraint.txt

7. check cluster status

a) Start HANA HA Cluster on both nodes

Execute command: systemctl start pacemaker

b) Monitor the HANA HA Cluster

Execute command: systemctl status pacemaker

Execute command: crm_mon -r

In our example we have following result:

```

192.168.0.82 - PuTTY
hana0:/ # systemctl status pacemaker
● pacemaker.service - Pacemaker High Availability Cluster Manager
   Loaded: loaded (/usr/lib/systemd/system/pacemaker.service; disabled; vendor preset: disabled)
   Active: active (running) since Wed 2018-06-20 14:04:48 UTC; 58min ago
     Docs: man:pacemakerd
           http://clusterlabs.org/doc/en-US/Pacemaker/1.1-pcs/html/Pacemaker_Explained/index.html
   Main PID: 32044 (pacemakerd)
     Tasks: 7 (limit: 512)
    CGroup: /system.slice/pacemaker.service
            └─32044 /usr/sbin/pacemakerd -f
              └─32047 /usr/lib64/pacemaker/cib
                └─32048 /usr/lib64/pacemaker/stonithd
                  └─32049 /usr/lib64/pacemaker/lrmd
                    └─32050 /usr/lib64/pacemaker/attrd
                      └─32051 /usr/lib64/pacemaker/pengine
                        └─32052 /usr/lib64/pacemaker/crmd
^C
hana0:/ # crm status
Stack: corosync
Current DC: hana0 (version 1.1.15-21.1-e174ec8) - partition with quorum
Last updated: Wed Jun 20 15:03:39 2018
Last change: Wed Jun 20 15:03:36 2018 by root via crm_attribute on hana0

2 nodes configured
7 resources configured

Online: [ hana0 hana1 ]

Full list of resources:

res_ALIYUN_STONITH_0 (stonith:fence_aliyun): Started hana1
res_ALIYUN_STONITH_1 (stonith:fence_aliyun): Started hana0
rsc_vip_JLO_HDB00 (ocf::aliyun:vpc-move-ip): Started hana0
Master/Slave Set: msl_SAPHana_JLO_HDB00 [rsc_SAPHana_JLO_HDB00]
   Masters: [ hana0 ]
   Slaves: [ hana1 ]
Clone Set: cln_SAPHanaTopology_JLO_HDB00 [rsc_SAPHanaTopology_JLO_HDB00]
   Started: [ hana0 hana1 ]

hana0:/ #

```

Meanwhile, please kindly check, if a new entry [virtual_IP4_address] is added into the route table of VPC.

In our example, we have following:

Route Table Details					
Route Table ID		vpc-gw81g1d3cp14tzyub		VPC ID	
Name		-		Route Table Type	
Created At		05/16/2018, 10:37:20		System	
				Description	
				-	
Route Entry List					
Add Route Entry Refresh					
Destination CIDR Block	Status	Next Hop	Type	Actions	
192.168.4.1/32	Available	Instance ID: i-gw81g1d3cp14tzyub Instance Type: ECS Instance	Custom	Delete	
0.0.0.0/0	Available	Instance ID: ngw-gw81g1d3cp14tzyub Instance Type: NAT Gateway	Custom	Delete	
192.168.0.0/24	Available	-	System		
192.168.1.0/24	Available	-	System		

Verify the HA takeover

- Shutdown the primary node;

Check the status of Pacemaker as follows:

```

192.168.1.245 - PuTTY
hanal:~ # crm status
Stack: corosync
Current DC: hanal (version 1.1.15-21.1-e174ec8) - partition with quorum
Last updated: Wed Jun 20 15:19:59 2018
Last change: Wed Jun 20 15:19:44 2018 by root via crm_attribute on hanal

2 nodes configured
7 resources configured

Online: [ hanal ]
OFFLINE: [ hana0 ]

Full list of resources:

res_ALIYUN_STONITH_0 (stonith:fence_aliyun): Started hanal
res_ALIYUN_STONITH_1 (stonith:fence_aliyun): Stopped
rsc_vip_JL0_HDB00 (ocf::aliyun:vpc-move-ip): Started hanal
Master/Slave Set: msl_SAPHana_JL0_HDB00 [rsc_SAPHana_JL0_HDB00]
Masters: [ hanal ]
Stopped: [ hana0 ]
Clone Set: cln_SAPHanaTopology_JL0_HDB00 [rsc_SAPHanaTopology_JL0_HDB00]
Started: [ hanal ]
Stopped: [ hana0 ]

hanal:~ # █

```

Compare the entry of route table in VPC as follows:

Destination CIDR Block	Status	Next Hop	Type	Actions
192.168.4.1/32	Available	Instance ID: i-gw8byf3m4f6a8os6rke9 Instance Type: EC2 Instance	Custom	Delete
0.0.0.0/0	Available	Instance ID: ngw-gw89g9ouh8m0g0rh1 Instance Type: NAT Gateway	Custom	Delete
192.168.0.0/24	Available	-	System	
192.168.1.0/24	Available	-	System	

Example

Example Cluster Configuration

In our example, the cluster configuration (you can check it via command "crm configure show") should have content as below:

```

node 1: hana0 \
attributes hana_jl0_vhost=hana0 hana_jl0_srmode=sync hana_jl0_remoteHost=hana1 hana_jl0_site=hana0
lpa_jl0_lpt=10 hana_jl0_op_mode=logreplay
node 2: hana1 \
attributes lpa_jl0_lpt=1529509236 hana_jl0_op_mode=logreplay hana_jl0_vhost=hana1 hana_jl0_site=hana1
hana_jl0_srmode=sync hana_jl0_remoteHost=hana0
primitive res_ALIYUN_STONITH_0 stonith:fence_aliyun \

```



```

op monitor interval=120 timeout=60 \
params pcmk_host_list=hana0 port=i-gw8byf3m4f9a8os6rke8 access_key=<access key> secret_key=<secret key>
region=eu-central-1 \
meta target-role=Started
primitive res_ALIYUN_STONITH_1 stonith:fence_aliyun \
op monitor interval=120 timeout=60 \
params pcmk_host_list=hana1 port=i-gw8byf3m4f9a8os6rke9 access_key=<access key> secret_key=<secret key>
region=eu-central-1 \
meta target-role=Started
primitive rsc_SAPHanaTopology_JL0_HDB00 ocf:suse:SAPHanaTopology \
operations $id=rsc_SAPHanaTopology_JL0_HDB00-operations \
op monitor interval=10 timeout=600 \
op start interval=0 timeout=600 \
op stop interval=0 timeout=300 \
params SID=JL0 InstanceNumber=00
primitive rsc_SAPHana_JL0_HDB00 ocf:suse:SAPHana \
operations $id=rsc_SAPHana_JL0_HDB00-operations \
op start interval=0 timeout=3600 \
op stop interval=0 timeout=3600 \
op promote interval=0 timeout=3600 \
op monitor interval=60 role=Master timeout=700 \
op monitor interval=61 role=Slave timeout=700 \
params SID=JL0 InstanceNumber=00 PREFER_SITE_TAKEOVER=true DUPLICATE_PRIMARY_TIMEOUT=7200
AUTOMATED_REGISTER=false
primitive rsc_vip_JL0_HDB00 ocf:aliyun:vpc-move-ip \
op monitor interval=60 \
meta target-role=Started \
params address=192.168.4.1 routing_table=vtb-gw8fii1g1d8cp14tzynub interface=eth0
ms msl_SAPHana_JL0_HDB00 rsc_SAPHana_JL0_HDB00 \
meta clone-max=2 clone-node-max=1 interleave=true target-role=Started
clone cln_SAPHanaTopology_JL0_HDB00 rsc_SAPHanaTopology_JL0_HDB00 \
meta clone-node-max=1 interleave=true
colocation col_SAPHana_vip_JL0_HDB00 2000: rsc_vip_JL0_HDB00:Started msl_SAPHana_JL0_HDB00:Master
location loc_hana0_stonith_not_on_hana0 res_ALIYUN_STONITH_0 -inf: hana0
location loc_hana1_stonith_not_on_hana1 res_ALIYUN_STONITH_1 -inf: hana1
order ord_SAPHana_JL0_HDB00 Optional: cln_SAPHanaTopology_JL0_HDB00 msl_SAPHana_JL0_HDB00
property cib-bootstrap-options: \
have-watchdog=false \
dc-version=1.1.15-21.1-e174ec8 \
cluster-infrastructure=corosync \
stonith-action=off \
stonith-enabled=true \
stonith-timeout=150s \
last-lrm-refresh=1529503606 \
maintenance-mode=false
rsc_defaults rsc-options: \
resource-stickiness=1000 \
migration-threshold=5000
op_defaults op-options: \
timeout=600

```

Example for /etc/corosync/corosync.conf

In our example, the corosync.conf should on hana1 should have content as below:

```
totem{
version: 2
token: 5000
token_retransmits_before_loss_const: 6
secauth: on
crypto_hash: sha1
crypto_cipher: aes256
clear_node_high_bit: yes
interface {
ringnumber: 0
bindnetaddr: 192.168.0.83
mcastport: 5405
ttl: 1
}
# On Alibaba Cloud, transport should be set to udpu, means: unicast
transport: udpu
}
logging {
fileline: off
to_logfile: yes
to_syslog: yes
logfile: /var/log/cluster/corosync.log
debug: off
timestamp: on
logger_subsys {
subsys: QUORUM
debug: off
}
}
nodelist {
node {
ring0_addr: 192.168.0.83
nodeid: 1
}
node {
ring0_addr: 192.168.1.246
nodeid: 2
}
}
quorum {
# Enable and configure quorum subsystem (default: off)
# see also corosync.conf.5 and votequorum.5
provider: corosync_votequorum
expected_votes: 2
two_node: 1
}
```

Reference

- Pacemaker 1.1 Configuration Explained
- SAP HANA SR Performance Optimized Scenario
- SAP HANA system replication - SAP Help Portal

Microsoft SQL Server on Alibaba Cloud

Microsoft SQL Server on Alibaba Cloud

- Getting Started
- Prerequisites
 - Alibaba Cloud account and RAM
 - SQL Server Images and Version
 - ECS instance
 - VPC
 - Cloud Disk
 - OSS
 - Shared block storage
- Installation
 - Create SQL Server instance
 - Windows Settings
- High Availability
- Backup
- Performance
 - Separate Cloud Disk
 - Table Compression
 - Enable Buffer Pool Extension(BPE)
 - Max degree of parallelism setting
- Monitor
 - Monitoring Capability
 - Alarm Capability
- Management Studio
 - Remote SQL Server Management Studio
 - Default SQL Server Management Studio

Getting Started

Make use of Alibaba Cloud reliable and flexible cloud computing infrastructure and platform services can help you to run Microsoft SQL Server more stably and smoothly. Microsoft SQL server on Alibaba cloud Elastic Compute Service(ECS) just like when you installed on-premises, you are responsible for installation, administering the database, including backups and recovery, patching the operating system and the database, tuning of the operating system and database parameters, managing

security, and configuring high availability or replication. It gives you complete control over every setting. To learn what is ECS, please refer to the [detail](#). This article provides you with the best practice for Microsoft SQL Server instance running on ECS instance. Please be aware this article used **SQL server 2016 Enterprise Edition** as the example, so not every option is appropriate for every version. It also presents a simplistic procedure to guide you for running your Microsoft SQL server. If you want to choose Alibaba RDS for SQL server, please refer to [Quick Start for SQL Server](#).

Prerequisites

Alibaba Cloud account and RAM

You must have registered to an Alibaba Cloud account. We recommend you to enable RAM to manage your account. Resource Access Management (RAM) is a cloud service that helps you manage user identities and control resources access. Using RAM, you can create and manage user accounts, and control the operation permissions that these user accounts possess for resources under your account, for example, employees, systems, and applications. For the detail information please refer to [Quick Start](#) and [RAM best practices](#).

SQL Server Images and Version

Alibaba cloud support Bring Your Own License (BYOL) and Images. SQL Server Enterprise, Standard, and Express Editions are licensed for production use. For Enterprise and Standard Editions, contact your software vendor for the installation media. You can find purchasing information and a directory of Microsoft partners on Microsoft official purchasing website. Free editions you can find at Microsoft official website: [SQL Server Downloads](#).

ECS instance

Elastic Compute Service (ECS) is a type of computing service that features elastic processing capabilities. ECS has a simpler and more efficient management mode than physical servers. You can create instances, change the operating system, and add or release any number of ECS instances at any time to fit your business needs. An ECS instance is a virtual computing environment that includes CPU, memory, and other basic computing components. An instance is the core component of ECS and is the actual operating entity offered by Alibaba Cloud. Other resources, such as disks, images, and snapshots, can only be used in conjunction with an ECS instance. Before create SQL server instances you have to create ECS instances first using the [ECS console](#), about the detail information, please refer to [Create ECS instances](#).

VPC

Virtual Private Cloud (VPC) creates an isolated network environment for you SQL Server environment. You can select an IP address range, divide networks, and configure the routing list and gateway. The interflow of VPC intranet and between VPC and physical IDC machine rooms can be realized among

regions or users. About how to create VPC please refer to the detail.

Cloud Disk

Ultra Cloud Disk: When you create ECS instance, Ultra Cloud Disk as the system disk provides a high-performance location for operating system and windows page file.

SSD Cloud Disk: When you create ECS instance we recommend you choose SSD cloud disk store the database files, tempdb, log file separately. Separate SSD cloud disks provide high performance and high reliability.

- High performance: A single SSD cloud disk provides a maximum of 20,000 random reading/writing IOPS and 300 MBps throughput of storage performance.
- $IOPS = \min\{1200 + 30 * disk_size, 20000\}$. The base is 1200 IOPS, and each GB provides 30 random IOPS up to a maximum of 20,000.
- $Throughput = \min\{80 + 0.5 * disk_size, 300\}$ MBps. The base is 80 MBps, and each GB adds an additional 0.5 MBps up to a maximum of 300 MBps throughput performance.
- Reliability: SSD cloud disks use Alibaba Cloud's Apsara distributed storage technology, based on three distributed copies, which can guarantee 99.9999999% data reliability.

For how to create a cloud disk, please refer to [create a cloud disk](#).

OSS

Alibaba Cloud Object Storage Service (OSS) is a network-based data access service. OSS enables you to store and retrieve unstructured data including text files, images, audios, and videos. We recommend you backup your SQL Server database into OSS. For how to use OSS please see [Get started with Object Storage Service](#)

Shared block storage

Shared Block Storage is designed for the high availability architecture of enterprise-class applications and provide shared access to block storage devices in a Share-everything architecture, such as the SQL Server always on with WSFC node architecture, which is common among government departments, enterprises, and financial customers, and the high availability server cluster architecture. For about shared block storage detail, please see [Shared block storage FAQ](#)

Installation

This section provides general information about how to create a SQL server instance on Elastic Compute Service (ECS). The tutorial includes the following tasks :

Create SQL Server instance

We recommend you to close the windows update setting before you create your SQL Server instance as below:

1. Click "run" button and input gpedit.msc
2. You will open the "Local Computer Policy" :

- Select "Administrative template"
- Select " Windows Components"
- Select " Windows Update"


3. Change the "Configure Automatic Updates" option to disable status.

After the SQL server instance creating you can decide whether need to enable the update setting. It is the same method with on premise to create your SQL Server instance on ECS instance. You can choose installation wizard, command-line or using a configuration file. For how to create or install please refer to Microsoft website.


Windows Settings

This section provides you the settings about how to configure windows settings to better optimize your SQL Server instance. We recommend you to setting the windows server firewall rule to specify the IP addresses for your client computer. It is very important security policy when you create your SQL Server instance to be able to connect to the database from other client machines. Configure the firewall to allow incoming traffic:

- Open your windows firewall with advanced security.




- Click right button to create a new inbound rule.




- Select your program path.


Input: %ProgramFiles%\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\Binn\sqlservr.exe




- Allow the connection.



- Name your rule, for example: mysqlserver.



- Set your remote policy.



The system default network settings are usually sufficient. Alibaba cloud offers you the high capacity and performance network. When you create ECS instances you can choose the network bandwidth from 1M up to 100M, about the bandwidth, Please refer to ECS Bandwidth FAQs. Windows requires anti-virus software to be installed. Install enterprise level anti-virus software and enable virus library updating and real-time protection, however, if the antivirus software is not configured correctly, it can negatively impact your database performance. Microsoft provides advice about how to choose antivirus software.


High Availability

We recommend you to use Windows Server Failover Clustering and SQL Server AlwaysOn Availability Groups as your SQL Server high availability solution on ECS instances.


The Always On feature must be enabled for the server instance 'sqlserver' before you can create an availability group on this instance. To enable Always On:

- Open the SQL Server Configuration Manager.
- Select SQL Server Services.
- Right-click the SQL Server instance name.
- Select Properties, and use the Always On High Availability tab of the SQL Server Properties dialog. (ObjectExplorer).

You should create 3 ECS instances, one is for the DC(domain controller)and DNS, the other two are the cluster nodes. You need to create a shared cloud disk as the shared block storage as below:



and then choose the region and size, please pay attention you can only buy 20GB at least as the shared block storage.



For how to plan, configure your WSFC and Always on group please see Microsoft official website.

Backup

The SQL Server backup and restore component provides an essential safeguard for protecting critical data stored in your SQL Server databases. We strongly recommend place the databases, logs, backups on separate SSD Cloud disk. Placing the data and backups on SSD Cloud disk also enhances the I/O performance for both writing backups and the production use of the database. A backup and restore strategy contains a backup portion and a restore portion. Designing an effective backup and restore strategy requires careful planning, implementation, and testing. There is no difference

between doing SQL Server database backup and restore from Alibaba cloud and on premise version. [Backing Up and Restoring How-to Topics \(SQL Server Management Studio\)](#) and [Backing Up and Restoring How-to Topics \(Transact-SQL\)](#)

provides best practice for how to implement a solid backup and maintenance action.

Use the Cloud SSD disk to store your backups and then copy them into OSS bucket or you can use the windows task scheduler copy them as the regular task.

Performance

This section provides you with how to tuning your SQL server instance performance on ECS instances. Running SQL Server on ECS environment we recommend you continue using the same database performance tuning options that are applicable to SQL Server in on-premises server environment SQL Server Enterprise Edition has a long list of added capabilities over Standard Edition. If you are migrating an existing license to ECS, there are some performance options that you might consider implementing.

Separate Cloud Disk

We recommend you to place the databases, logs, backups on separate SSD Cloud disk during you create your SQL Server instances

Table Compression

Generally, data compression reduces the space occupied by the data. It can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk. Data compression can be performed for a table, clustered index, non-clustered index. We recommend you to enable table and index compression. It might seem counter-intuitive that compressing tables could make your system perform faster, but in most cases, that's what happens. The tradeoff is using a small amount of CPU cycles to compress the data and eliminate the extra disk IO required to read and write the bigger blocks. Generally, the less disk IO your system uses, the better its performance will be. Instructions for estimating and enabling table and index compression please refer to [Microsoft website](#)

Enable buffer pool extension (BPE)

We recommend you to use the buffer pool extension to speed data access. The buffer pool extension feature enables you to push clean pages to the SSD Cloud disk, instead of dropping them. This works along the same lines as virtual memory, which is to say by swapping, and gives you access to the clean pages on the SSD Cloud disk, which is faster than you would get by going to the regular disk to fetch the data. This technique is not nearly as fast as having enough memory, but it can give you a modest increase in throughput when your available memory is low. For how to enable BPE and the technology detail please refer to [Microsoft website](#)

Max degree of parallelism setting

We recommend you to configure the max degree of parallelism option to 8. When your SQL Server instance runs on ECS that has more than one processor, it detects the best degree of parallelism, that is, the number of processors employed to run a single statement, for each parallel plan execution. You can use the max degree of parallelism option to limit the number of processors to use in parallel plan execution.


This value is set using `sp_configure` system procedure or you use SQL Server Management Studio. The default value is 0 which means there is no upper limit and SQL Server can use all available processors. If you set Max Degree of Parallelism to 1 then all queries will execute serially. This setting is ignored on servers with a single processor. Occasionally you might find that parallelism actually hinders performance of some queries. In this case the cost of initializing and synchronizing parallel plans might exceed the benefit of running portions of the query on multiple threads. If you feel that serial execution of a particular query can provide better performance you can override this setting using MAXDOP option within an individual query. For how to configure please refer to [Microsoft website](#)

Monitor

We recommend you to use **CloudMonitor** to monitor your ECS instances. Make sure that your ECS monitoring agents are functional to collect metric data. Otherwise, you must install the agent manually. For more information, see [How to install CloudMonitor agent](#).

More metrics are available.

More than 20 metrics are supported, such as `cpu.user`, `cpu.system`, `cpu.iowait`, `netout.packages`, `netout.errorpackage`. For OS metrics, the collection granularity is 15 seconds. [Which metrics are supported in the latest version?](#)



Monitoring capability

CloudMonitor allows more than 30 metrics covering CPU, memory, disk, and network to meet the basic monitoring and O&M requirements of the servers. Click [here](#) to view the full list of metrics the switch.

Alarm capability

CloudMonitor provides alarm service for all metrics, allowing you to set alarm rules for individual servers, application groups, and all the other resources. You can use the alarm service as per your business requirements. CloudMonitor provides Host monitoring metrics to set alarm rules for individual servers, application groups, and all the other resources. You can use the alarm services as per your business requirements. You can use the alarm service directly in the host monitoring list, or use it in your application group once you add servers to the group. You can add the alarm rules directly in the host monitoring list, or use it in your application group once you add servers to the group. For how to create an alarm service, please see [here](#)

Management Studio

You can use SQL Server Management Studio to perform most administrative tasks. This section provides you with how to manage your SQL server instance on ECS instances.

Remote SQL Server Management Studio

Microsoft offered SQL Server Management Studio to configure SQL Server databases. You can download and installed it on your desktop, connect to database remotely.

Default SQL Server Management Studio

You also can use the default SQL Server Management Studio which running on the instance itself. With this method you should connect to your SQL Server instance through RDP. SQL Server 2012 and SQL Server 2014 both include the SQL Server Management Studio by default. For SQL Server 2016, you must download the SQL Server Management Studio from the Microsoft website and install it on the instance. By default, SQL Server uses Windows Authentication mode to control remote access to SQL Server itself. If you need to use SQL Server Authentication mode, change the authentication mode.

ECS Metrics Collector for SAP Deployment Guide

- Overview
- Create a RAM role
- Configure the RAM role
- Install the SAP Host Agent
- Install and uninstall the ECS Metrics Collector
 - Based on Linux
 - Based on Windows

Version Control

Version	Revision Date	Types Of Changes	Effective Date
1.0			2018/5/31
1.1	2018/6/11	Description of installation preparations is updated.	2018/6/12


Overview

The ECS Metrics Collector is a monitor agent that used by the SAP system on the cloud platform to collect required VM configurations and physical resource usage information.

When the SAP system runs on ECS, the SAP host agent obtains SAP system monitoring information (such as operating system, network, storage, and SAP architecture information) through the metadata service and open APIs, and provides the information to SAP application for event analysis and system performance analysis.

You need to install Metrics Collector for SAP for each ECS instance (database or application) running the SAP system.

The following figure is the overall architecture of Metrics Collector.



Configure the RAM role


To monitor the ECS Metrics Collector, you need to configure the certain RAM role and access rights.

Note: The configuration of RAM role is permanently valid for your account.

For more information about RAM role configuration, see Use the instance RAM role on the console.


1. Log on to the ECS console.
2. Click RAM in the left navigation bar.

If RAM is unavailable in the navigation bar, click Monitoring and Management and select RAM.




You need to activate the RAM function if you use this function first time.

3. Open the RAM console, select **Role Management**, and click **Create a Role**.



4.In the “Select the role type” step, select **Service Role**.



5.In the “Enter the type information” step, select **ECS Elasic Compute Service**.

Create Role

1: Select Role Type

2: Enter Type

3: Configure Basic

4: Role created

Select the trusted services that can use this role to access your cloud resources.

ApiGateway Service

When configure the FunctionCompute as API Gateway backend service , you must create an API Gateway trusted service role. The API Gateway will play the role to call your FunctionCompute.

ECS Elastic Compute Service

Elastic Compute Service will play the role to access your other resources.

RDS ApsaraDB for RDS

RDS will play the role to access your other resources.

Previous

6.In the “Configure basic information” step, set a role name. For example, you can set the role name to **ecs-metrics-collector**. Click **Create**.

Create Role

1: Select Role Type

2: Enter Type

3: Configure Basic

4: Role created

* Role Name:

ecs-metrics-collector

Names must be 1-64 characters long. They may only contain letters, numbers, and hyphens.


Description:

ecs-metrics-collector

Previous


Create

7.After the role is created, click **Authorize** to go to the next step.




8.Click **Edit Authorization Policy**. Select the policy names **AliyunECSReadOnlyAccess** and **AliyunCloudMonitorReadOnlyAccess**, and assign the policies to your RAM role.

The two policies authorize the Metrics Collector to only read ECS information and only read CloudMonitor resources.



9.Click **OK** to complete policy.




Configure RAM role


Attach the created RAM role to your SAP ECS instance.

Open the ECS console and click **Instance** to find out your ECS instance.

Click **More** and select **Bind/Unbind RAM Role** from the drop-down list.



Select the RAM role you have created.



Click OK to attach the role.

For more information about RAM binding and unbinding, see [Use the instance RAM role on the console](#)

Install the SAP Host Agent

You need an SAP market place account for logon.

For the suggested version of SAP Host Agent, see [1031096 - Installing Package SAPHOSTAGENT](#)

Install and uninstall the ECS Metrics Collector

The installation of ECS Metrics Collector and the data reading from open APIs require a connection to the Internet. If your ECS instance does not have an Internet IP address, apply for an EIP and bind it to the ECS instance. For details, see [Apply for an EIP address](#) and [Bind an EIP](#).

Based on Linux

Install the ECS Metrics Collector



1. Log on to the SAP ECS instance as a root user.

To use the root rights, you need sudo, so you must belong to the sudo group.

2. Install the ECS Metrics Collector by using the Cloud Assistant as follows:

The Metrics Collector needs the support of python. If you have not installed pip for the ECS instance, install it first.

```
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
pip -V #Check the pip version
```

```
aliyun_installer -i ecs-metrics-collector
```

Enter the ID of the latest version for package_id. In this example, the latest version 8 is used.

```
hana01:/usr/local/share/aliyun-assist # aliyun_installer -i ecs-metrics-collector
package_id  name      version arch  publisher
5      ecs-metrics-collector  0.2    x64    aliyun-inc
6      ecs-metrics-collector  0.3    x64    aliyun-inc
8      ecs-metrics-collector  0.4    x64    aliyun-inc
Please input the package_id you want to install.
8
The package_id you input is 8.
The package you want to install is ecs-metrics-collector.
Downloading...
Check MD5
Unzip
Installing...
see details in installation log file: /var/log/ecs_metrics_collector/install.log
```

```
[DEBUG] prepare_env
[DEBUG] account_check
[DEBUG] distro_check
[DEBUG] init_system_check
[DEBUG] system_restraint_check
[DEBUG] package_check
[DEBUG] python_env_check
[DEBUG] install_collector
[DEBUG] install_watchmen
[DEBUG] install_service
```

Installation success !

Verify the installation of Metrics Collector

3. Verify the ECS Metrics Collector service.

```
systemctl status ecs_metrics_collector
```

```
hana01:/usr/local/share/aliyun-assist # systemctl status ecs_metrics_collector
● ecs_metrics_collector.service - ECS Metrics Collector for SAP
   Loaded: loaded (/etc/systemd/system/ecs_metrics_collector.service; enabled; vendor preset: disabled)
   Active: active (running) since Thu 2018-05-24 19:56:39 CST; 4min 45s ago
     Process: 5412 ExecStartPost=/usr/local/sbin/ecs_metrics_collector_watchmen --set (code=exited, status=0/SUCCESS)
     Process: 5405 ExecStart=/usr/local/sbin/ecs_metrics_collector -d (code=exited, status=0/SUCCESS)
    Main PID: 5410 (python)
       Tasks: 2 (limit: 512)
      CGroup: /system.slice/ecs_metrics_collector.service
              └─5410 python /usr/local/sbin/ecs_metrics_collector -d

May 24 19:56:39 hana01 systemd[1]: Starting ECS Metrics Collector for SAP...
May 24 19:56:39 hana01 systemd[1]: ecs_metrics_collector.service: Supervising process 5410 which is not our child. We'll most lik...t exists.
May 24 19:56:39 hana01 systemd[1]: Started ECS Metrics Collector for SAP.
Hint: Some lines were ellipsized, use -l to show in full.
```

4. Check whether the watching task and automatic update task are configured in crontab and whether the running log is recorded.

```
cat /etc/cron.d/ecs_metrics_collector
cat /var/log/ecs_metrics_collector/watchmen.log
```

```
hana01:/usr/local/share/aliyun-assist # cat /etc/cron.d/ecs_metrics_collector
*/1 * * * * root /usr/local/sbin/ecs_metrics_collector_watchmen --watching >& /dev/null
*/5 * * * * root /usr/local/sbin/ecs_metrics_collector_watchmen --update >& /dev/null
hana01:/usr/local/share/aliyun-assist # cat /var/log/ecs_metrics_collector/watchmen.log
[Thu May 24 19:56:39 CST 2018] [watchmen] [INFO] register watchmen
[Thu May 24 19:57:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 19:58:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 19:59:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:00:01 CST 2018] [watchmen] [INFO] check_update
[Thu May 24 20:00:02 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:01:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:02:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:03:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:04:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:05:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
[Thu May 24 20:05:02 CST 2018] [watchmen] [INFO] check_update
[Thu May 24 20:06:01 CST 2018] [watchmen] [INFO] detect collector running, pid 5410
hana01:/usr/local/share/aliyun-assist #
```

The two tasks are automatically added to crontab during the installation of ECS Metrics Collector. The automatic update task checks for available updates of ECS Metrics Collector on the Cloud Assistant server every hour. If an update of the latest version is detected, an automatic upgrade is performed.

Check the collected metric data

5. Verify the collected data.

```
curl localhost:8888 | vim -
```

```
<?xml version="1.0" encoding="UTF-8"?>
<metrics>
  <metric context="vm" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Data Provider Version</name>
    <value>0.4</value>
  </metric>
  <metric context="host" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Cloud Provider</name>
    <value>Alibaba Cloud</value>
  </metric>
  <metric context="vm" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Instance Type</name>
    <value>ecs.se1.4xlarge</value>
  </metric>
  <metric context="host" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Virtualization Solution</name>
    <value>ECS Virt</value>
  </metric>
  <metric context="host" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Virtualization Solution Version</name>
    <value>2.0</value>
  </metric>
  <metric context="host" category="config" type="string" unit="none" last-refresh="1527134681" refresh-interval="60">
    <name>Hardware Manufacturer</name>
    <value>Alibaba Cloud</value>
  </metric>
</metrics>
```

Uninstall the ECS Metrics Collector

Log on to your SAP ECS instance as a root user.

Run the

aliyun_installer -u ecs-metrics-collector command to uninstall the ECS Metrics Collector.


```
hanasr01:~ # aliyun_installer -u ecs-metrics-collector
Uninstalling...
/etc/cron.d/ecs_metrics_collector not exist
/etc/init.d/ecs_metrics_collector not exist
/etc/systemd/system/ecs_metrics_collector not exist

Uninstallation success !
```

More commands

For more commands related to Metrics Collector, run the `aliyun_installer --h` command.

Based on Windows



1. Log on to your SAP ECS instance as an administrator.

2. Install the ECS Metrics Collector by using the Cloud Assistant. Right-click the Start menu, right-click CMD, and choose Run as Administrator to open the CMD window. Switch to the directory of Cloud Assistant to perform the installation.

The latest version of Cloud Assistant is 1.0.0.107. Modify the directory name according to the latest version.

```
cd "C:\ProgramData\aliyun\assist\[1.0.0.107]"
aliyun_installer.exe -l
aliyun_installer.exe -i ecs_metrics_collector
```

Select the ID of the latest version for `package_id`.

```

选定 管理员: C:\Windows\system32\cmd.exe

C:\ProgramData\aliyun\assist\1.0.0.107>aliyun_installer.exe -l
package_id    name    version arch    publisher
?    ecs_metrics_collector    1.0.0.51    x64    aliyun-inc

C:\ProgramData\aliyun\assist\1.0.0.107>aliyun_installer.exe -i ecs_metrics_collector
The latest version of ecs_metrics_collector in the software repository is 1.0.0.51
Downloading...
Check MD5
Unzip
Installing...
Installation success.

```

Verify the installation of ECS Metrics Collector

3. Verify the ECS Metrics Collector service.

Right-click the taskbar and choose Task Manager > Service. Check and ensure that the Ecs metrics collector service status is **running**.



4. Right-click the Start menu, choose Run, enter `taskschd.msc`, and press Enter. Open the task schedule program and click Task Scheduler Library to check that the update detection task has been scheduled.




A task is created during the installation of ECS Metrics Collector to check for version updates of ECS Metrics Collector on the Cloud Assistant server every hour. If an update of the latest version is detected, an automatic upgrade is performed.

Check the collected metric data

5. Verify the collected data.


Open the browser and enter `http://localhost:8888A` page similar to the following is displayed. If all parameters have values, the configuration is successful. If a value is abnormal, check the RAM role

configuration.



Uninstall the ECS Metrics Collector

1. Log on to your SAP ECS instance as an administrator. Switch to the installation directory and double-click `uninst.exe`.



More commands

For more commands related to Metrics Collector, run the `aliyun_installer.exe --h` command.