

SAP HANA High Availability Cross-Zone Solution on Alibaba Cloud

With SUSE Linux Enterprise Server for SAP Applications

SUSE Linux Enterprise Server for SAP Applications 12 and later
Alibaba Cloud

Jinhui Li, Product Manager SAP Solutions on Alibaba Cloud (AliCloud Germany)
Bernd Schubert, SAP Solution Architect (SUSE)

SAP HANA High Availability Cross-Zone Solution on Alibaba Cloud

With SUSE Linux Enterprise Server for SAP Applications

Date: 2024-12-12

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications. This document explains how to deploy an SAP HANA High Availability solution cross different Zones on **Alibaba Cloud**. It is based on SUSE Linux Enterprise Server for SAP Applications 12 SP2. The concept however can also be used with SUSE Linux Enterprise Server for SAP Applications 12 SP3 or newer.

Disclaimer: Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Contents

- 1 Solution Overview 4
- 2 Infrastructure Preparation 6
- 3 Software Preparation 10
- 4 Configuring SAP HANA System Replication 12
- 5 Configuring High Availability Extension for SAP HANA 17
- 6 Example 25
- 7 Reference 28
- 8 Legal notice 29
- 9 GNU Free Documentation License 30

1 Solution Overview

1.1 SAP HANA System Replication

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

For details, refer to the SAP Help Portal, HANA system replication:

<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.03/en-US/b74e16a9e09541749a745f41246a065e.html> .

1.2 High Availability Extension Included with SUSE Linux Enterprise Server for SAP Applications

The SUSE High Availability Extension 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. This helps you to build your SAP HANA HA solution up more effectively.

For details, refer to the latest version of the document **SAP HANA SR Performance Optimized Scenario** at the SUSE documentation Web page:

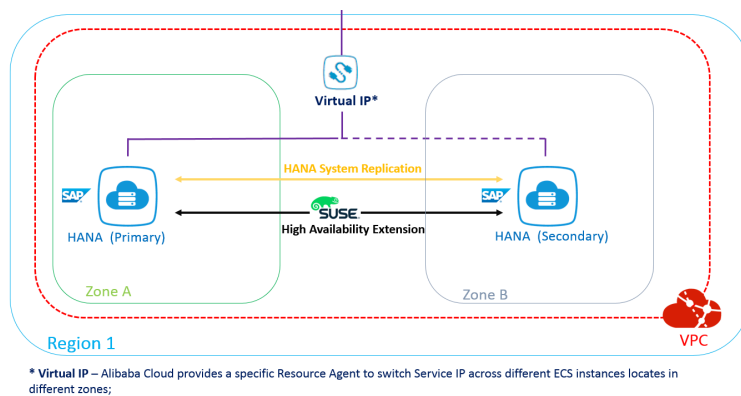
<https://documentation.suse.com/sbp/sap/> .

1.3 Architecture Overview

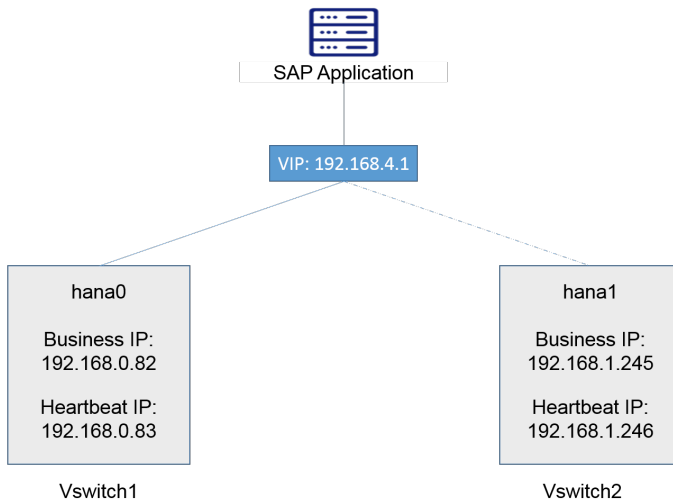
This document guides you on how to deploy an SAP HANA High Availability (HA) solution cross different Zones. The following list contains a brief architecture overview:

1. The High Availability Extension included with SUSE Linux Enterprise Server for SAP Applications is used to set up the HA Cluster;
2. SAP HANA System Replication is activated between the two HANA nodes;
3. Two HANA nodes are located in different Zones of the same Region;
4. The Alibaba Cloud Specific Virtual IP Resource Agent is used to allow Moving IP automatically switched to Active SAP HANA node; the Alibaba Cloud specific STONITH device is used for fencing.

Alibaba Cloud Architecture - Overview:



Network Architecture - Overview:



1.4 Network Design

The following table contains information for the Network design.

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	

HANASTudio	HANA Studio	no	192.168.0.79	no
------------	-------------	----	--------------	----

2 Infrastructure Preparation

The next sections contain information about how to prepare the infrastructure.

2.1 Infrastructure List

To set up your infrastructure, the following components are required:

- 1 Virtual Private Cloud (VPC) network; [Virtual Private Cloud \(https://www.alibabacloud.com/product/vpc\)](https://www.alibabacloud.com/product/vpc)
- 2 Elastic Compute Service (ECS) instances in different zones of the same VPC; [Elastic Compute Service \(https://www.alibabacloud.com/product/ecs\)](https://www.alibabacloud.com/product/ecs)
- 2 Elastic Network Interfaces (ENI) - one for each ECS instance; [Elastic network interfaces \(https://www.alibabacloud.com/help/doc-detail/58496.htm\)](https://www.alibabacloud.com/help/doc-detail/58496.htm)
- Alibaba Cloud specific Virtual IP Resource Agent and STONITH device;
- NAT Gateway and SNAT entry; [NAT Gateway \(https://www.alibabacloud.com/help/product/44413.htm\)](https://www.alibabacloud.com/help/product/44413.htm)


2.2 Creating VPC

First, create a VPC via **Console→Virtual Private Cloud→VPCs→Create VPC**. In this example, a VPC named **suse_hana_ha** in the Region **EU Central 1 (Frankfurt)** has been created:



There should be at least two VSwitches (subnets) defined within the VPC network. Each VSwitch should be bound to a different Zone. In this example, the following two VSwitches (subnets) are defined:


- 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;



2.3 Creating ECS Instances

Two ECS instances are created in different Zones of the same VPC via **Console→Elastic Compute Service ECS→Instances→Create Instance**. Choose the "SUSE Linux Enterprise Server for SAP Applications" image from the Image Market place.

In this example, two ECS instances (hostname: *hana0* and *hana1*) are created in the Region **EU Central 1**, Zone A and Zone B, within the 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.



2.4 Creating ENIs and Binding to ECS Instances

Create two ENIs via **Console→ Elastic Compute Service ECS→Network and Security→ENI**, and attach one for each ECS instance, for HANA System Replication purposes. Configure the IP addresses of the ENIs to the subnet for HANA System Replication only.

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



Meanwhile, within the Guest OS, `/etc/hosts` should also be configured.

According to the example at hand, run the following two commands on both sites:

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

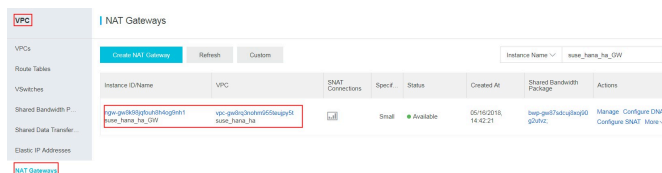
```
echo "192.168.1.245 hana1 hana1" >> /etc/hosts
```

The output looks as follows:

```
hana0:/tmp # cat /etc/hosts
#
# hosts        This file describes a number of hostname-to-address
#              mappings for the TCP/IP subsystem.  It is mostly
#              used at boot time, when no name servers are running.
#              On small systems, this file can be used instead of a
#              "named" name server.
# Syntax:
#
# IP-Address   Full-Qualified-Hostname  Short-Hostname
#
# special IPv6 addresses
::1           localhost ipv6-localhost ipv6-loopback
fe00::0       ipv6-localnet
ff00::0       ipv6-mcastprefix
ff02::1       ipv6-allnodes
ff02::2       ipv6-allrouters
ff02::3       ipv6-allhosts
127.0.0.1     localhost
10.251.192.226 iZ2zeigulbbi64490pjmdnZ
192.168.0.82  hana0 hana0
192.168.1.245 hana1 hana1
```

2.5 Creating NAT Gateway and configure SNAT entry

Now create an NAT Gateway attached to the given VPC. In the example at hand, an NAT Gateway named **suse_hana_ha_GW** has been created:



Instance ID/Name	VPC	NAT Connections	Spec	Status	Created At	Shared Bandwidth Package	Actions
vpc-qvts3u0v050a9p5f suse_hana_ha_GW	vpc-qvts3u0v050a9p5f	1	Small	Available	05/16/2018, 14:42:21	bw-qvts3u0v050a9p5f	Manage Configure SNAT Configure DNAT

After having creating the NAT Gateway, you need to create a corresponding SNAT entry to allow ECS instances within the VPC to access public addresses on the Internet.



Note

An Alibaba Cloud specific STONITH device and Virtual IP Resource Agent are mandatory for the cluster. They need to access Alibaba Cloud OpenAPI through a public domain.

In the example at hand, two SNAT entries have been created, for ECS instances located in a different network range:

SNAT Entry ID	Source CIDR Block	VSwitch ID	Public IP	Status	Actions
snat-gslh3wsk3fatehggrl	192.168.0.0/24	vsw-gslh3wsk3fatehggrl	47.254.171.15	Available	Edit Remove
snat-gslh3wsk3fatehggrl	192.168.0.0/24	vsw-gslh3wsk3fatehggrl	47.254.171.16	Available	Edit Remove

2.6 Creating STONITH Device and Virtual IP Resource Agent

Download the STONITH fencing software with the following command:

```
wget http://sap-automation-cn-beijing.oss-cn-beijing.aliyuncs.com/software/aliyun-pacemaker_new.zip
```

For an HA solution, a fencing device is an essential requirement. Alibaba Cloud provides its own STONITH device, which allows the servers in the HA cluster to shut down the node that is not responsive. The STONITH device leverages Alibaba Cloud OpenAPI underneath the ECS instance, which is similar to a physical reset / shutdown in an on-premise environment.

```
hanal:/hana/tmp # ls
122.05 122.05.05 SAFHOSTAGENT 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/f63a3ba3357613caade16cf4fa69178eaaccd9cf2fc6692368024499083/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/28/045cf5cc1e80cc482c3bd1979287e62db73b3c42ef0e429f2d0d8dd2bfbef/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/24/84/29ce4167df5c4a320aad91e1178e5a1baf9cfe1c63

```

Configure Alibaba Cloud OpenAPI SDK and Client.

```
aliyuncli configure
```

Get your Access Key as shown below:



3 Software Preparation

The next sections contain information about the required software.

3.1 Software List

The following software must be available: - SUSE Linux Enterprise Server for SAP Applications 12 SP2 or newer - HANA Installation Media - SAP Host Agent Installation Media

3.2 High Availability Extension Installation

Both ECS instances are created with the SUSE Linux Enterprise Server for SAP Applications image. On both ECS instances, the High Availability Extension (with the major software components: Corosync and Pacemaker), and the package SAPHanaSR should be installed. To do so, you can use zypper.

First, install the pattern High Availability Extension on both nodes:

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

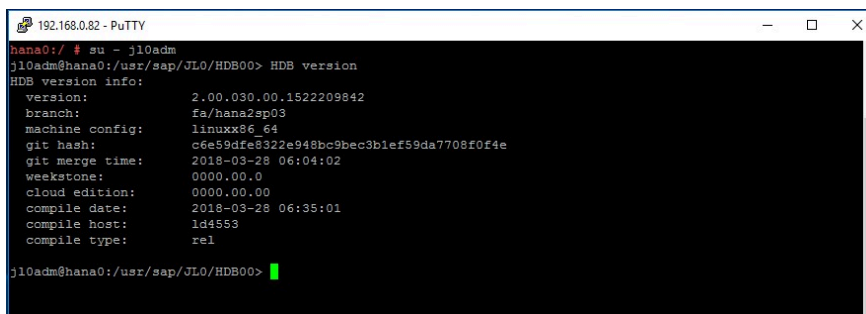
Now, install the Resource Agents for controlling the SAP HANA system replication on both cluster nodes:

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

3.3 SAP HANA Installation

Next, install the SAP HANA software on both ECS instances. Make sure the SAP HANA SID and Instance Number are the same (this is required by SAP HANA System Replication). It is recommended to use *hdblocm* to do the installation. For details refer to the SAP HANA Server Installation and Update Guide at <https://help.sap.com/viewer/2c1988d620e04368aa4103bf26f17727/2.0.03/en-US>.

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

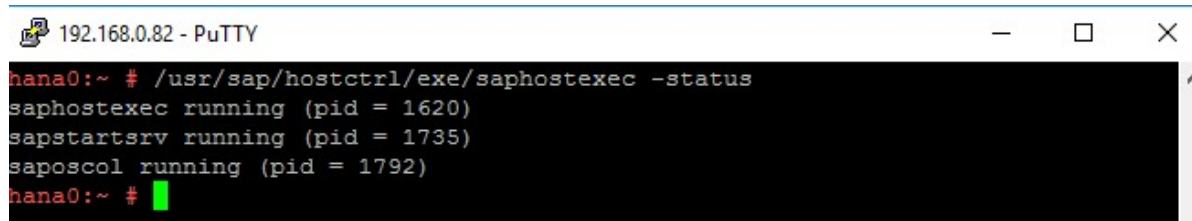
A terminal window titled '192.168.0.82 - PuTTY' showing a user 'j10adm' at 'hana0'. The user runs 'su - j10adm' and then 'hdb version'. The output displays version information for SAP HANA, including version, branch, machine config, git hash, git merge time, weekstone, cloud edition, compile date, compile host, and compile type.

```
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>
```

3.4 SAP Host Agent Installation

When you have finished the HANA installation with *hdblcm* as mentioned above, the SAP Host Agent should already be installed on your server. To install it manually, refer to the article [Installing SAP Host Agent Manually: https://help.sap.com/saphelp_nw73ehp1/helpdata/en/8b/92b1cf6d5f4a7eac40700295ea687f/content.htm?no_cache=true](https://help.sap.com/saphelp_nw73ehp1/helpdata/en/8b/92b1cf6d5f4a7eac40700295ea687f/content.htm?no_cache=true).

Check the SAP Host Agent status after you have installed SAP HANA with *hdblcm* on *hana0* and *hana1*:



```
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:~ #
```

4 Configuring SAP HANA System Replication

The following sections detail how to configure SAP HANA System Replication.

4.1 Backing up SAP HANA on Primary ECS Instance

To do a backup on SAP 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 back up systemDB and all tenantDB as shown below in the example):

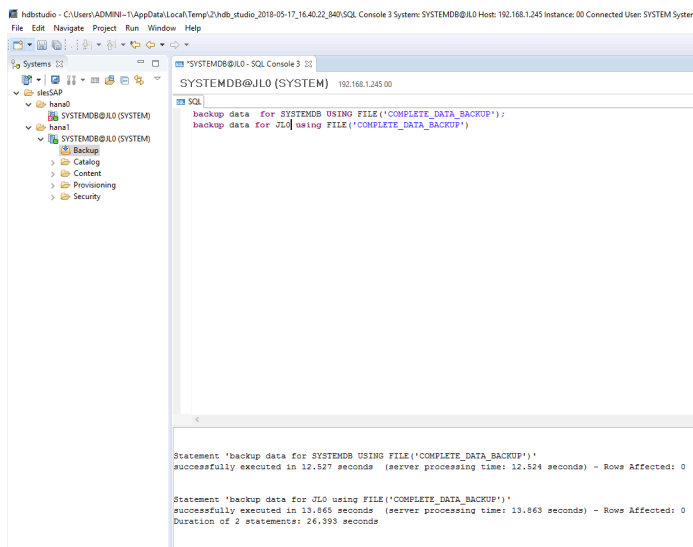
```
BACKUP DATA for <DATABASE> using FILE('COMPLETE_DATA_BACKUP')
```

Command line example:

```
BACKUP DATA for SYSTEMDB using FILE('COMPLETE_DATA_BACKUP')
```

```
BACKUP DATA for JL0 using FILE('COMPLETE_DATA_BACKUP')
```

In this example for HDB Studio, the SAP HANA database backup is executed on both ECS instances as shown below:



4.2 Configuring SAP HANA System Replication on Primary Node

Log on to 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 the following file content:

```
vi /hana/shared/<SID>/global/hdb/custom/config/global.ini
```

Add the following content:

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

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

[HOSTNAME] should be the hostname of the Secondary node;

In the example at hand, the configuration is as follows:

```
[system_replication_hostname_resolution]

192.168.1.246 = hana1
```

4.3 Configuring SAP HANA System Replication on Secondary Node

Perform the same steps as outlined above for the Primary node on the Secondary node. However, do not forget to use here the IP and hostname of the Primary node instead of the Secondary node. In the example at hand, the configuration is as follows:

```
[system_replication_hostname_resolution]

192.168.0.83 = hana0
```

4.4 Enable SAP HANA System Replication on Primary Node

Log on to the primary node with: `su - <sid>adm;`

Start HANA with: `HDB start;`

Enable System Replication with:

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

[primary location name] should be replaced by the location of your primary SAP HANA node. For the example at hand, the following command is used:

```
hdbnsutil -sr_enable --name=hana0
```



Note

All of the above operations are done on the Primary node.

4.5 Register the Secondary Node to the Primary SAP HANA Node

Log on to the secondary node with: `su - <sid>adm;`

Stop SAP HANA with: `HDB stop;`

Register the Secondary SAP HANA node to the Primary SAP HANA node by running the 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
```

For the example at hand, the following command is used:

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

Start SAP HANA with: HDB start;

Verify the System Replication Status with the following command:

```
hdbnsutil -sr_state
```

```

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.

```




Note

All of the above operations are done on the Secondary node.

5 Configuring High Availability Extension for SAP HANA

5.1 Configuration of Corosync

It is recommended that you add more redundancy for messaging (Heartbeat) by using separate ENIs attached to the ECS instances with a separate network range.

On Alibaba Cloud, it is strongly suggested to only use Unicast for the transport setting in Corosync.

Follow the steps below to configure Corosync:

- Create Keys

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

In the example at hand, the command is executed 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` as root on the Primary SAP HANA node with the following content:

```
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 SAP HANA System Replication. In the example at hand, the IP address of ENI of the current node (192.168.0.83 for *hana0* and 192.168.1.246 for *hana1*) is used.



Note

This value will be different on Primary and Secondary node. The **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 the example at hand, it should be 192.168.0.83 for *hana0* and 192.168.1.246 for *hana1*.

After completing the editing of */etc/corosync/corosync.conf* on the Primary HANA node, copy the files */etc/corosync/authkey* and */etc/corosync/corosync.conf* to */etc/corosync* on the Secondary SAP HANA node with the following command:

```
scp /etc/corosync/authkey root@**hostnameOfSecondaryNode**:/etc/corosync
```

In the example at hand, the following command is executed:

```
hana1:/ # scp /etc/corosync/authkey root@hana2:/etc/corosync
```

After you have copied the *corosync.conf* to the Secondary node, configure the **bindnetaddr** as above to the local heart beating IP address.

5.2 Configuration of Pacemaker

For the SAP HANA High Availability solution, you need to configure seven resources and the corresponding constraints in Pacemaker.



Note

The following Pacemaker configuration only needs to be done on one node. It is usually done on the Primary node.

5.2.1 Cluster Bootstrap and More

Add the configuration of the bootstrap and default setting of the resource and operations to the cluster. Save the 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 the command below to add the setting to the cluster:

```
crm configure load update crm-bs.txt
```

5.2.2 STONITH Device

This part defines the Aliyun STONITH devices in the cluster.

Save the following scripts in a file for **SLES12 SP2 and older releases**: *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
```

Save the following scripts in a file for **SLES12 SP3 and newer releases**: *crm-stonith.txt*.

```
primitive res_ALIYUN_STONITH_1 stonith:fence_aliyun \
    op monitor interval=120 timeout=60 \
        params plug=<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 plug=<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
```

Be sure to implement the following changes:

- **[primary node hostname]** / **[secondary node hostname]** should be replaced by the real hostname of your primary and secondary hostname, in our case is hana0 and hana1.
- **[primary node instance id]** / **[secondary node instance id]** should be replaced by the real instance-id of your ECS instance, you can get this from the Alibaba Cloud console. For SLES 12 SP3 or newer releases, please use *plug* instead of *port*.
- **[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 is located.

Execute the command below to add the resource to the cluster:

```
crm configure load update crm-stonith.txt
```

5.2.3 SAPHanaTopology

This part defines an SAPHanaTopology RA, and a clone of the SAPHanaTopology on both nodes in the cluster. Save the 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"

```

Be sure to implement the following changes:

- **[SID]** should be replaced by the real SAP HANA SID.
- **[instance number]** should be replaced by the real SAP HANA Instance Number.

Execute the command below to add the resources to the cluster:

```
crm configure load update crm-saphanatop.txt
```

5.2.4 SAPHana

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

```

primitive rsc_SAPHana_<SID>_HDB<instance number> ocf:suse:SAPHana \
    operations $id="rsc_SAPHana_<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"

```

Be sure to implement the following changes:

- **[SID]** should be replaced by the real SAP HANA SID.
- **[instance number]** should be replaced by the real SAP HANA Instance Number.

Execute the command below to add the resources to the cluster:

```
crm configure load update crm-saphana.txt
```

5.2.5 Virtual IP

This part defines a Virtual IP RA in the cluster. Save the 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
```

Be sure to implement the following changes:

- **[virtual_IP4_address]** should be replaced by the real IP address you prefer to provide the 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 the command below to add the resource to the cluster:

```
crm configure load update crm-vip.txt
```

5.2.6 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 the 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>
```

Be sure to implement the following changes:

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

Execute the command below to add the resource to the cluster:

```
crm configure load update crm-constraint.txt
```

5.3 Check the Cluster Status

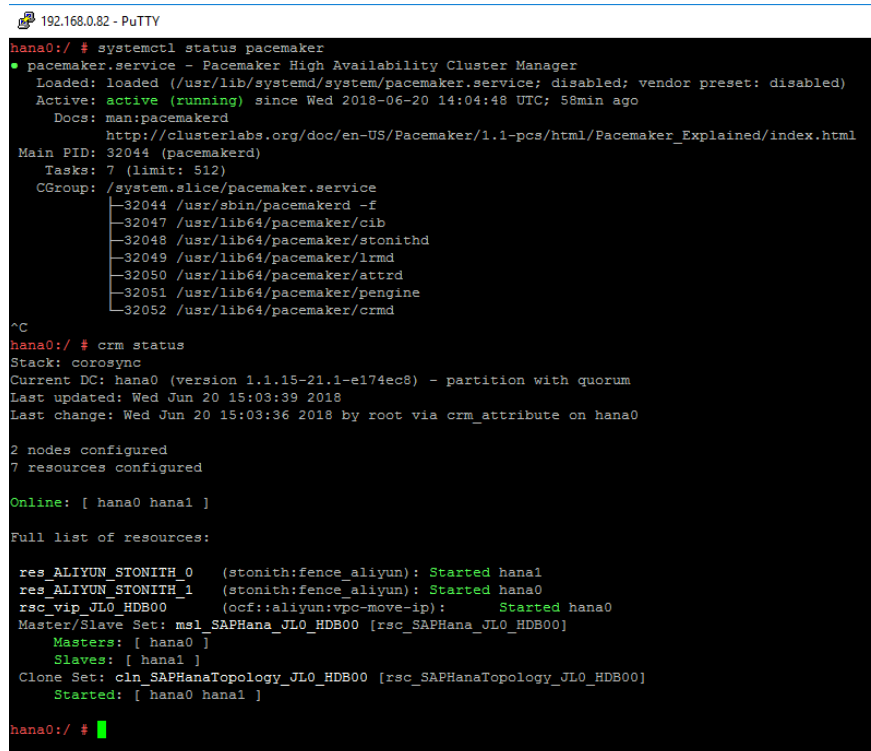
Start the SAP HANA High Availability Cluster on both nodes:

```
systemctl start Pacemaker
```

Monitor the SAP HANA High Availability Cluster with the following commands:

```
systemctl status pacemaker
crm_mon -r
```

In the example at hand, the result looks as follows:



```
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-el74ec8) - 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 ]

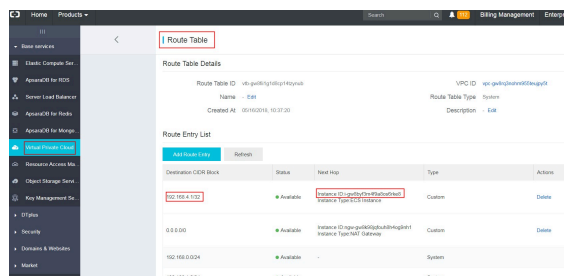
Full list of resources:

res_ALIYUN_STONITH_0 (stonith:fence_aliyun): Started hana1
res_ALIYUN_STONITH_1 (stonith:fence_aliyun): Started hana0
res_vip_JL0_HDB00 (ocf::aliyun:vpc-move-ip): Started hana0
Master/Slave Set: msl_SAPHana_JL0_HDB00 [rsc_SAPHana_JL0_HDB00]
Masters: [ hana0 ]
Slaves: [ hana1 ]
Clone Set: cln_SAPHanaTopology_JL0_HDB00 [rsc_SAPHanaTopology_JL0_HDB00]
Started: [ hana0 hana1 ]

hana0:/ #
```

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

In the example at hand, the following result is shown:



Route Table Details					
Route Table ID: vpc-gzkyg1n0y1k0y01			VPC ID: vpc-gzkyg1n0y1k0y01		
Name: -			Route Table Type: System		
Created At: 2018-06-20 10:37:20			Description: -		
Route Entry List					
Destination CIDR Block	Status	Next Hop	Type	Actions	
192.168.4.1/32	Available	Instance-Local Forwarding	Custom	Delete	
0.0.0.0/0	Available	Instance-Local Forwarding	Custom	Delete	
192.168.0.0/24	Available	-	System		
192.168.1.0/24	Available	-	System		

5.4 Verify the High Availability Takeover

Shut down the primary node.

Check the status of Pacemaker as shown below:

```
192.168.1.245 - PuTTY
hana1:~ # crm status
Stack: corosync
Current DC: hana1 (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 hana1

2 nodes configured
7 resources configured

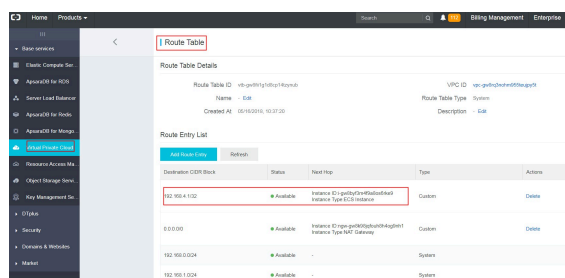
Online: [ hana1 ]
Offline: [ hana0 ]

Full list of resources:

res_ALIYUN_STONITH_0 (stonith:fence_aliyun): Started hana1
res_ALIYUN_STONITH_1 (stonith:fence_aliyun): Stopped
rsc_vip_JL0_HDB00 (ocf::aliyun:vpc-move-ip): Started hana1
Master/Slave Set: msl_SAPHana_JL0_HDB00 [rsc_SAPHana_JL0_HDB00]
Masters: [ hana1 ]
Stopped: [ hana0 ]
Clone Set: cln_SAPHanaTopology_JL0_HDB00 [rsc_SAPHanaTopology_JL0_HDB00]
Started: [ hana1 ]
Stopped: [ hana0 ]

hana1:~ #
```

Compare the entry of the route table in the VPC as shown below:



6 Example

6.1 Example of Cluster Configuration

You can check the cluster configuration via the command `crm configure show`. For the example at hand, the cluster configuration should display the following content:

```
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-gw8fiilgld8cp14tzynub 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

```

6.2 Example of /etc/corosync/corosync.conf

For the example at hand, the *corosync.conf* on *hana1* should display the following content:

```

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
}

```


7 Reference

- Pacemaker 1.1 Configuration Explained https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/1.1/html/Pacemaker_Explained/ .
- SAP HANA SR Performance Optimized Scenario <https://documentation.suse.com/sbp/all/html/SLES4SAP-hana-sr-guide-PerfOpt-12/index.html> .
- SAP HANA system replication - SAP Help Portal <https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.03/en-US/b74e16a9e09541749a745f41246a065e.html> .
- SAP Applications on Alibaba Cloud: Supported Products and IaaS VM Types <https://launchpad.support.sap.com/#/notes/2552731> .

8 Legal notice

Copyright © 2006–2024 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> .

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Below we draw your attention to the license under which the articles are published.

9 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.