

SAP NetWeaver High Availability Cluster 7.40 for the AWS Cloud

Setup Guide

SUSE Linux Enterprise Server for SAP Applications 12 SP3 and newer
SAP NetWeaver 7.40
Amazon Web Services

Fabian Herschel, Distinguished Architect SAP (SUSE)
Bernd Schubert, SAP Solution Architect (SUSE)
Stefan Schneider, Partner Solutions Architect (AWS)
Manas Srivastava, Senior Consultant SAP (AWS)

SAP NetWeaver High Availability Cluster 7.40 for the AWS Cloud

Setup Guide

Date: 2024-11-14

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP* applications.

This document explains how to deploy a high availability cluster solution using SUSE Linux Enterprise Server for SAP Applications 12 for the Enqueue Replication scenario on the AWS platform. The goal is to match the SAP NW- HA-CLU 7.40 certification specifications and goals. This document is based on SUSE Linux Enterprise Server for SAP Applications 12 SP3. The concept however can also be used with newer service packs of SUSE Linux Enterprise Server for SAP Applications.

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 About this Guide 4
- 2 Scope of This Document 5
- 3 Overview 6
- 4 Testing the AWS Agents 26
- 5 SAP Installation 28
- 6 Implementing the Cluster 36
- 7 Administration 47
- 8 AWS Specific Post Installation Tasks 51
- 9 Additional Implementation Scenarios 51
- 10 Appendix 69
- 11 Legal Notice 78
- 12 GNU Free Documentation License 79

1 About this Guide

1.1 Introduction

SUSE® Linux Enterprise Server for SAP Applications is the optimal platform to run SAP* applications with high availability. Together with a redundant layout of the technical infrastructure, single points of failure can be eliminated.

SAP* Business Suite is a sophisticated application platform for large enterprises and mid-size companies. Many critical business environments require the highest possible SAP* application availability.

The described cluster solution can be used for SAP* S/4 HANA and for SAP* SAP NetWeaver.

SAP NetWeaver is a common stack of middleware functionality used to support the SAP business applications. The SAP Enqueue Replication Server constitutes application level redundancy for one of the most crucial components of the SAP NetWeaver stack, the enqueue service. An optimal effect of the enqueue replication mechanism can be achieved when combining the application level redundancy with a high availability cluster solution, for example, as provided by SUSE Linux Enterprise Server for SAP Applications. The described concept has proven its maturity over several years of productive operations for customers of different sizes and branches.

1.2 Additional Documentation and Resources

Chapters in this manual contain links to additional documentation resources that are either available on the system or on the Internet.

For the latest documentation updates, see <https://documentation.suse.com>.

You can also find numerous whitepapers and other resources at the SUSE Linux Enterprise Server for SAP Applications resource library: <https://www.suse.com/products/sles-for-sap/resource-library/>.

This guide and other SAP specific best practices can be downloaded via <https://documentation.suse.com/sbp/sap/>. Here you can find guides for SAP HANA system replication automation and HA scenarios for SAP NetWeaver and SAP S/4 HANA.

1.3 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests

For services and support options available for your product, refer to <http://www.suse.com/support/>.

To report bugs for a product component, go to <https://scc.suse.com/support/> requests, log in, and select *Submit New SR* (Service Request).

Mail

For feedback on the documentation of this product, you can send a mail to doc-team@suse.com (<mailto:doc-team@suse.com>). Make sure to include the document title, the product version and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

2 Scope of This Document

This guide will detail how to:

- Plan a SUSE Linux Enterprise High Availability platform for SAP NetWeaver, including SAP Enqueue Replication Server.
- Set up a Linux high availability platform and perform a basic SAP NetWeaver installation including SAP Enqueue Replication Server on SUSE Linux Enterprise.
- Integrate the high availability cluster with the SAP control framework via `sap-suse-cluster-connector`, as certified by SAP.
- Install HA cluster solutions for the SAP HANA database on AWS as being described in SAP note 2309342 (SUSE Linux Enterprise High Availability Extension on AWS).

This guide focuses on the high availability (HA) of the central services. HA cluster solutions for the database and SAP NetWeaver instances are described in the best practice "Simple Stack" available on the SUSE Best Practices documentation Web page (see section "Additional documentation and resources"). For SAP HANA system replication follow the guides for the performance- or cost-optimized scenario.

3 Overview

This guide describes how to set up a pacemaker cluster using SUSE Linux Enterprise Server for SAP Applications 12 for the Enqueue Replication scenario on the AWS platform. This guide does not document how to install on premises pacemaker clusters. The goal is to match the SAP NW-HA-CLU 7.40 certification specifications and goals.

These goals include:

- Integration of the cluster with the SAP start framework *sapstartsrv* to ensure that maintenance procedures do not break the cluster stability
- Rolling Kernel Switch (RKS) awareness
- Standard SAP installation to improve support processes

The updated certification SAP NW-HA-CLU 7.40 has redefined some of the test procedures and described new expectations how the cluster should behave under special conditions. These changes allowed to improve the cluster architecture and to design it for easier usage and setup. Shared SAP resources are managed in AWS Elastic File Systems (EFS). The SAP instances themselves are installed on EFS file systems to allow switching over the file systems for proper functionality.

3.1 Using AWS Architectures in SUSE Linux Enterprise Server Pacemaker Clusters

SUSE Linux Enterprise Server pacemaker clusters will be installed in an AWS region. An AWS region consists of multiple availability zones. Availability zones are located in different data centers which are 10 to 50 kilometers apart. Availability zones have independent flood levels, electricity and network hookup. They are supposed to be independent. AWS recommends architectural patterns where redundant cluster nodes are being spread across availability zones (AZs) to allow a customer to overcome individual AZ failures.

An AWS Virtual Private Network (VPC) is spanning all AZs. We assume that a customer will:

- Have identified two availability zones to be used
- Have created subnets in the two AZs which can host the two nodes of a SUSE Linux Enterprise High Availability Extension cluster

- Use a routing table which is attached to the two subnets
- Optionally: host a Route53 private hosted naming zone to manage names in the VPC
- All components of the cluster should reside in the same Amazon Account. The use of networking components such as a route table in another account (shared VPC setup) is not supported by the cluster resource agent. If you do require a multi account landscape then we advise you to reach to your AWS representative to have a look at implementing a Transit GateWay for cross account/VPC access.

The AWS specific components can be installed in two configurations. Both configurations use the AWS Overlay IP address. An Overlay IP address is an AWS specific routing entry which can send network traffic to an instance, no matter which AZ the instance is located in.

The SUSE Linux Enterprise High Availability Extension cluster will update this routing entry as it is required. All SAP system components in the VPC can reach an AWS instance with an SAP system component inside a VPC through this Overlay IP address.

Overlay IP addresses have one disadvantage: they need to come from a CIDR range which is outside of the VPC. Otherwise they would be part of a subnet and a given availability zone.

On premises users like SAP GUIs cannot reach this IP address since the AWS Virtual Private Network (VPN) gateway will not route traffic to such an IP address. A customer has two options to overcome this limitation.

1. Use an SAP Router in the VPC. On premises users can reach it. The SAP router can relay traffic to the ASCS system.
2. Configure the additional Route 53 agent. Route 53 is the AWS specific name service. The cluster agent will change the IP address for a given name of the ASCS service. The on premises name server will need to delegate requests to the sub domain in the AWS VPC to this name service. On premises SAP GUI user will contact the ASCS through a name. Section TBD in the appendix explains how to integrate Route 53 with your local naming services.

3.2 Prerequisites for the AWS-Specific HA Installation

There are several prerequisites which need to be met before starting the installation:

- Have an AWS account
- Have an AWS user with admin rights, or at least rights to

- Create security groups
- Create EFS file systems
- Modify AWS routing tables
- Create policies and attach them to IAM roles
- Optional for Route53 agent installation
 - Create and modify A-records in a private hosted zone
- Understand your landscape
 - Know your region and its AWS name
 - Know your VPC and its AWS id
 - Know which availability zones you want to use in your VPC
 - Have a subnet in each of the availability zones
 - Have a routing table which is implicitly or explicitly attached to the two subnets
 - Have free IP addresses in the two subnets for your SAP installation and EFS mount points
 - Allow network traffic in between the two subnets Allow outgoing Internet access from the subnets
 - Optionally: Have a Route 53 private hosted zone which hosts a subdomain for instances in the two subnets
 - Have a resource record with a name and the IP address for the SAP central instance

Use the checklist in the appendix to make a note of all information needed before starting the installation.



Note

With regards to the EFS/NFS filesystems, with SLES for SAP 15, it is now possible to mount these via the Operating System and remove the NFS mount & unmount control to outside of cluster. This may simplify the ASCS / ERS environment. Refer to <https://documentation.suse.com/sbp/sap/html/SAP-S4HA10-setupguide-simplemount-sle15/index.html> for more detail.

3.2.1 Tagging the EC2 Instances

The EC2 instances will have host names which are automatically generated. Select host names which comply with SAP requirements, as detailed in SAP note 611361.

The cluster agents need to be able to identify the EC2 instances in the correct way. This happens through instance tags.

Tag the two EC2 instances through the console or the AWS Command Line Interface (CLI) with arbitrarily chosen tags like *cluster* and the host name as it will be shown in the command *uname*. Use the same tag (like *cluster*) and the individual host names for both instances. The AWS documentation (http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html) explains how to tag EC2 instances.



Note

Refrain from using non-ASCII characters in any tag assigned to cluster-managed resources.

3.2.2 Security Groups



Important

This section does not cover a discussion of SAP related ports in security groups, it only lists the ports which need to be available for the SUSE cluster only.

The following ports and protocols need to be configured to allow the two cluster nodes to communicate with each other:

- Port 5405 for inbound UDP: Used to configure the corosync communication layer. Port 5405 is being used in common examples. A different port may be used depending on the corosync configuration.
- Port 7630 for inbound TCP: Used by the SUSE "hawk" Web GUI.
- Enable ICMP: Used through a ping command in the AWS IP-move agent of the SUSE cluster.

We assume that there are no restrictions for outbound network communication.

3.2.3 Creating an AWS CLI Profile on Both EC2 Instances

The SUSE Linux Enterprise Server agents use the AWS Command Line Interface (CLI). This AWS CLI profile needs to be created for the root account *root* on both instances. The SUSE resources require a profile which creates output in text format. The name of the profile is arbitrary. The name chosen in this example is *cluster*. The region of the instance needs to be added as well. Replace the string *region-name* with your target region in the following example.

One way to create such a profile is to create a file */root/.aws/config* with the following content:

```
[default]
region = region-name
[profile cluster]
region = region-name
output = text
```

The other way is to use the *aws configure* CLI command in the following way:

```
# aws configure
AWS Access Key ID [None]:
AWS Secret Access Key [None]:
Default region name [None]: region-name
Default output format [None]:

# aws configure --profile cluster
AWS Access Key ID [None]:
AWS Secret Access Key [None]:
Default region name [None]: region-name
Default output format [None]: text
```

This command sequence generates a default profile and a cluster profile.

3.2.4 Configure HTTP Proxies

Configuring HTTP Proxies is not needed if the system has transparent access to the Internet. The resource agents execute the AWS CLI (Command Line Interface) commands. These commands send HTTP/HTTPS requests to an access point in the Internet. These access point are usually directly reachable.

However, systems which do not offer transparent Internet access will need to provide an HTTP/HTTPS proxy. The configuration of the proxy access is described in full detail in the AWS documentation.

Add the following environment variables to the root user's *.bashrc* file:

```
export HTTP_PROXY=http://a.b.c.d:n
export HTTPS_PROXY=http://a.b.c.d:m
```

If authentication is required, add the following environment variables instead of the ones above:

```
export HTTP_PROXY=http://username:password@a.b.c.d:n
export HTTPS_PROXY=http://username:password@a.b.c.d:m
```

The AWS Data Provider for SAP needs to reach the instance meta data service directly. Add the following environment variable to the root user's *.bashrc* file:

```
export NO_PROXY=169.254.169.254
```

3.2.5 Add a Second IP for each Cluster Instance

The cluster configuration will require two IP addresses for each cluster instance. Adding a second IP address on the instance will allow the SUSE cluster to implement a two ring corosync configuration. The two ring corosync configuration will allow the cluster nodes to communicate with each other using the secondary IP address in the event that there is an issue communicating with each other over the primary IP address.

Please refer to AWS documentation to understand how to assign a secondary IP address: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html#assignIP-existing>

After the secondary IP address is associated to the cluster instance in AWS, you will need to add the secondary IP address to the cluster instance. Update the file */etc/sysconfig/network/ifcfg-eth0*. Replace *XX.XX.XX.XX* with the new secondary IP address and replace 'XX' with the two digit subnet mask.

```
IPADDR_1='XX.XX.XX.XX/XX'
```

```
LABEL_1="1"
```

The system will read the file and add the secondary IP address after the cluster instance is rebooted. Additionally, executing the command below as root will add the IP address to the cluster instance network stack with out rebooting.

```
ip address add XX.XX.XX.XX/XX dev eth0
```

Replace XX.XX.XX.XX with the new secondary IP address and replace 'XX' with the two digit subnet mask.

3.2.6 Disable the Source/Destination Check for the Cluster Instances

The source/destination check can be disabled through scripts using the AWS command line interface (AWS-CLI). The following command needs to be executed one time for both EC2 instances, which are supposed to receive traffic from the Overlay IP address:

```
# aws ec2 modify-instance-attribute --profile cluster --instance-id EC2-instance --no-source-dest-check
```

The system on which this command gets executed needs temporarily a role with the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1424870324000",
      "Effect": "Allow",
      "Action": [ "ec2:ModifyInstanceAttribute" ],
      "Resource": [
        "arn:aws:ec2:region-name:account-id:instance/instance-a",
        "arn:aws:ec2:region-name:account-id:instance/instance-b"
      ]
    }
  ]
}
```

Replace the individual parameter for the region, the account identifier and the two identifiers for the EC2 instances with appropriate values.

The source/destination check can be disabled as well from the AWS console. It takes the execution of the following drop-down box in the console for both EC2 instances (see below).

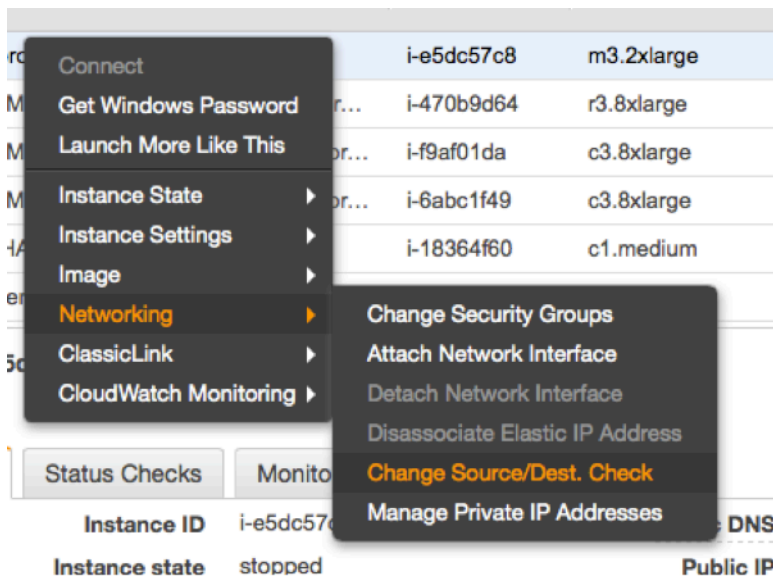


FIGURE 1: DISABLE SOURCE/DESTINATION CHECK AT CONSOLE

3.2.7 Avoid Deletion of Cluster-Managed IP Address on the eth0 Interface

SUSE Linux Enterprise Server 12 SP3 is the first version which ships the cloud-netconfig package. This package will remove any secondary IP address which is managed by the cluster agents from the eth0 interface. This can cause service interruptions for users of the HA service. Perform the following task on all cluster nodes:

Check whether the package cloud-netconfig-ec2 is installed with the command

```
# zypper info cloud-netconfig-ec2
```

Update the file `/etc/sysconfig/network/ifcfg-eth0` if this package is installed. Change the following line to a „no“ setting or add the line if the package is not yet installed:

```
CLOUD_NETCONFIG_MANAGE='no'
```

3.2.8 AWS Roles and Policies

The SAP ASCS and ESR will run the SUSE Linux Enterprise Server Pacemaker software and the agents. This software needs several AWS IAM privileges to operate the cluster. Create a new role for every ASCS/ESR cluster and associate this role to the two instances. Attach the policies detailed below to this role.

3.2.8.1 AWS Data Provider Policy

Every cluster node will operate an SAP system. SAP systems on AWS require an installation of the “AWS Data Provider for SAP”. The data provider needs a policy to access AWS resources. Use the policy as described in the “AWS Data Provider for SAP Installation and Operations Guide“, section “IAM Roles” and attach it to the role of the instance. This policy can be used by all SAP systems. It takes only one policy in an AWS account. The policy does not contain any instance-specific privileges.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "EC2:DescribeInstances",
        "EC2:DescribeVolumes"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "cloudwatch:GetMetricStatistics",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::aws-data-provider/config.properties"
    }
  ]
}
```

3.2.8.2 STONITH Policy

The instances of the SUSE cluster need the privilege to start and stop the other nodes in the cluster. Create a policy with a name like *stonith-policy* with the following content and attach it to the cluster role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1424870324000",
      "Effect": "Allow",
```

```

    "Action": [
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeTags"
    ],
    "Resource": "*"
},
{
    "Sid": "Stmt1424870324001",
    "Effect": "Allow",
    "Action": [
        "ec2:ModifyInstanceAttribute",
        "ec2:RebootInstances",
        "ec2:StartInstances",
        "ec2:StopInstances"
    ],
    "Resource": [
        "arn:aws:ec2:region-name:aws-account:instance/i-node1",
        "arn:aws:ec2:region-name:aws-account:instance/i-node2"
    ]
}
]
}

```

Replace the variable *aws-account* with the appropriate AWS account identifier. Replace the variables *i-node1* and *i-node2* with the AWS instance-ids of your two cluster nodes of (hacert01) and (hacert02). Replace the variable *region-name* with the name of your AWS region (Example: us-east-1 for the N. Virginia region). This policy is dependent of the instances of your cluster. You need to create a separate policy for every cluster!

3.2.9 Overlay IP Agent Policy

The Overlay IP agent changes a routing entry in an AWS routing table. Create a policy with a name like *Manage-Overlay-IP-Policy* and attach it to the role of the cluster instances:

```

{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": "ec2:ReplaceRoute",
            "Resource": "arn:aws:ec2:region-name:account-id:route-table/rtb-XYZ"
        },
        {

```

```

        "Sid": "VisualEditor1",
        "Effect": "Allow",
        "Action": "ec2:DescribeRouteTables",
        "Resource": "*"
    }
]
}

```

This policy allows the agent to update the routing tables which get used. Replace the following variables with the appropriate names:

- `region-name` : the name of the AWS region
- `account-id` : The name of the AWS account in which the policy is getting used
- `rtb-XYZ` : The identifier of the routing table which needs to be updated

3.2.10 Route 53 Updates

Optionally, you can install the Route 53 agent in the cluster. The following policy is needed only when the Route 53 agent is used. Create a policy with the name *Route53-Update* and attach it to the role of the two cluster nodes:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1471878724000",
      "Effect": "Allow",
      "Action": "route53:GetChange",
      "Resource": "arn:aws:route53::change/*"
    },
    {
      "Sid": "Stmt1471878724001",
      "Effect": "Allow",
      "Action": "route53:ChangeResourceRecordSets",
      "Resource": "arn:aws:route53::hostedzone/hosted zone ID/full name"
    },
    {
      "Sid": "Stmt1471878724002",
      "Effect": "Allow",
      "Action": [
        "route53:ListResourceRecordSets",
        "route53:ChangeResourceRecordSets"
      ]
    }
  ]
}

```



```
    "Resource": "arn:aws:route53:::hostedzone/hosted zone ID"
  }
]
}
```

This policy is specific to the hosted zone and the resource record set. Replace the variables *hosted zone ID* with the AWS ID of the hosted zone. Replace *full name* with the name of the entry. An individual Route 53 policy needs to be created for every cluster.

3.3 Add Overlay IP Addresses to Routing Table

Manually add two routing entries to the routing table which is assigned to the two subnets. The IP addresses need to be outside of the CIDR range of the VPC. Use the AWS console and search for “VPC”.

- Select VPC
- Click “Route Tables” in the left column
- Select route table used for SAP ASCS subnets
- Click tabulator “Routes”
- Click “Edit”
- Scroll to the end of the list and click “Add another route”

3.3.1 Add the Service IP Address for Your ASCS Service

Add the service IP address of the ASCS service (node hacert01). Use as filter /32 (example: 192.168.10.1/32). Add the Elastic Network Interface (ENI) name of your instance which is initially serving as ASCS. Save your changes by clicking “Save”.

This is the service IP address with the name sapha1as.

3.3.2 Add the Service IP Address for Your ERS Service

Add the service IP address of the ERS service (node hacert02). Use as filter /32 (example: 192.168.10.2/32). Add the Elastic Network Interface (ENI) name of your instance which is initially serving as ERS. Save your changes by clicking “Save”.

This is the IP address with the name sapha1er.

3.4 EFS File System

The cluster requires an NFS file system provided by AWS Elastic File System (EFS). The file system will manage:

- /usr/sap/HA1 data for ASCS00, ERS10, D02, DVEBMGS01 and the other application servers and the SYS directory
- /sapmnt

You need the identifier of your VPC and the subnet identifiers of the subnets in which you plan to operate the two cluster nodes. It is also feasible to pick other subnets. These subnets need to be reachable by the two cluster nodes and they need to be in the same availability zone (AZ) for high availability reasons. The option “General Purpose” will be sufficient.

Note down the DNS name of your specific EFS server. AWS name services will resolve it internally to your VPC to an IP address which is in your availability zone. We refer to this name as "efs-name" when we need to mount the file systems.

We use one file system for two future mount points (/usr/sap/HA1/ASCS00, /usr/SAP/HA1/ESR10). This keeps the administration level low and provides more throughput. AWS throughput in EFS is based on the total size of the file system.

Log in to one of the two cluster nodes and create several directories in the EFS file system through a temporary mount. As user *root*, execute the following commands:

```
# mount efs-name: /mnt
# mkdir -p /mnt/ASCS00 /mnt/ERS10 /mnt/D01 /mnt/DVEBMGS01 /mnt/D02 /mnt/SYS /mnt/sapmnt /
mnt/sapcd
# umount /mnt
```

Create additional directories for other application servers. This NFS file system will be used for all of them.

Mount the two mount points in the cluster nodes. Execute the following command on both cluster nodes as root:

```
# mkdir -p /sapmnt /usr/sap/HA1/SYS
```

Add the following two lines to the file /etc/fstab on the two instances which will run the SAP ASCS and the ERS service.

```
efs-name:SYS /usr/sap/HA1/SYS nfs4
    rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2 0 0
```

```
efs-name:sapmnt /sapmnt nfs4
rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2 0 0
```

Replace *efs-name* with the appropriate DNS name.

Mount the file system as root with the command

```
# mount /usr/sap/HA1/SYS
# mount /sapmnt
```

3.4.1 Enable Cluster Instances to Use the Overlay IP Address

The two cluster instances need the Overlay IP address to be configured as secondary IP address on their standard interface *eth0*. This can be achieved by the command:

```
# ip address add OVERLAY-IP dev eth0
```

Execute this command with root privileges on both instances. Add the ASCS IP address on the ASCS node *hacert01*. Add the Enqueue Replication address on the ERS node *hacert02*.

3.5 Differences to Previous Cluster Architecture

The concept is different to the old stack with the master-slave architecture. With the new certification we switch to a more simple model with primitives. This means we have on one machine the ASCS with its own resources and on the other machine the ERS with its own resources.

3.6 Five Systems for ASCS, ERS, Database and Additional SAP Instances

This guide describes the installation of a distributed SAP system on the five systems. In this setup only two systems are in the cluster. The database and SAP dialog instances could also be added to the cluster by either adding the three nodes to the cluster or by installing the database on either of the nodes. However we recommend to install the database on a separate cluster.



Note

The cluster in this guide only manages the SAP instances ASCS and ERS, because of the focus of the SAP NW-HA-CLU 7.40 certification.

If your database is SAP HANA, we recommend to set up the performance optimized system replication scenario using our automation solution SAPHanaSR. The SAPHanaSR automation should be set up in an own two node cluster. The setup is described in a separate best practice available at our best practice page. <https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices/>

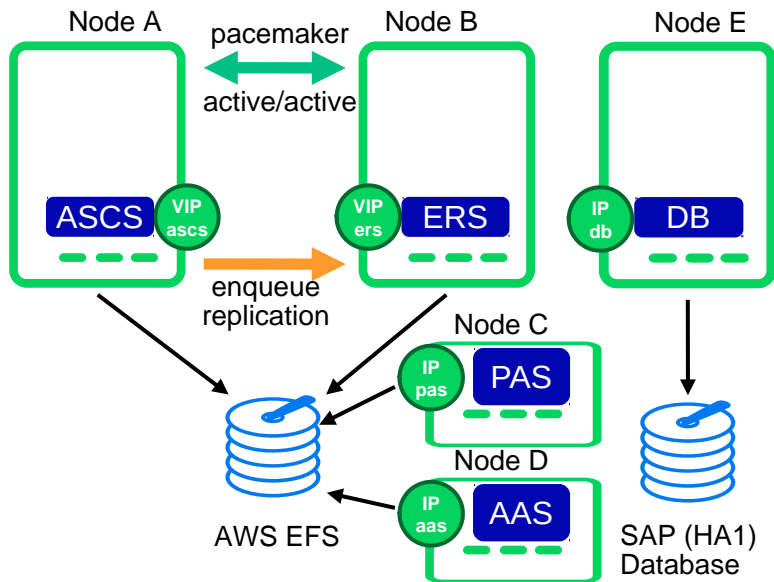


FIGURE 2: FIVE SYSTEMS FOR THE CERTIFICATION SETUP

CLUSTERED MACHINES

- One machine (hacert01) for ASCS; Hostname: sapha1as
- One machine (hacert02) for ERS; Hostname: sapha1er

NON-CLUSTERED MACHINE

- One machine for DB; Hostname: sapha1db
- One machine for PAS; Hostname: sapha1ci
- One machine for AAS; Hostname: sapha1d2

3.7 High Availability for the Database

Depending on your needs you could also increase the availability of the database, if your database is not already high available by design.

3.7.1 SAP HANA System Replication

A perfect enhancement of the five node scenario described in this document is to implement an SAP HANA system replication (SR) automation.

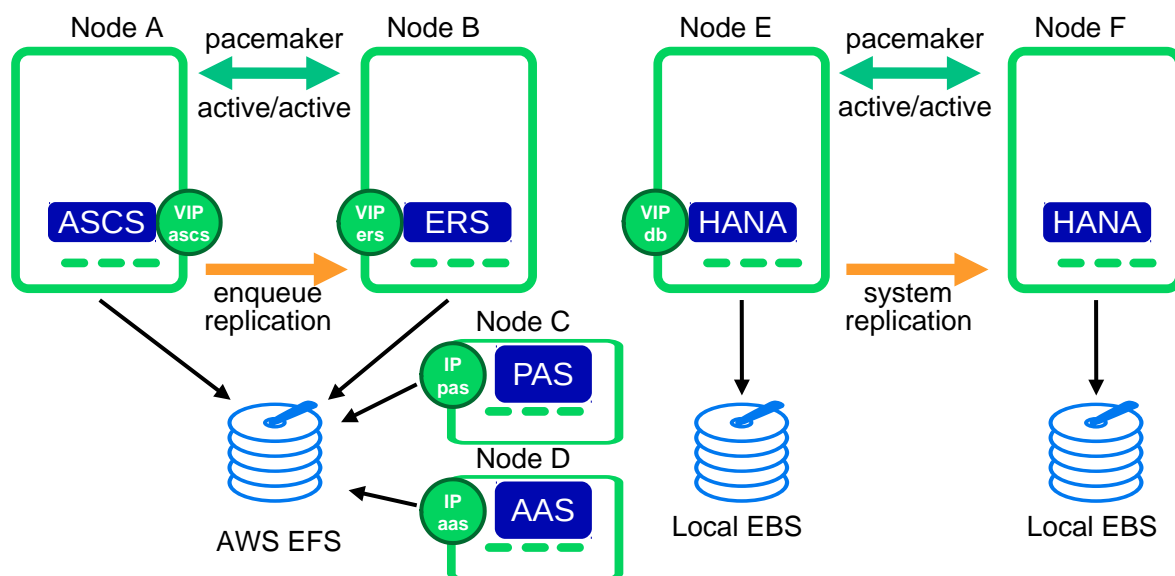


FIGURE 3: ONE CLUSTER FOR CENTRAL SERVICES, ONE FOR SAP HANA SR

The following Databases are supported in combination with this scenario:

- SAP HANA DATABASE 1.0
- SAP HANA DATABASE 2.0

3.8 Integration of SAP NetWeaver Into the Cluster Using the Cluster Connector

The integration of the HA cluster through the SAP control framework using the sap-suse-cluster-connector is of special interest. The sapstartsrv controls SAP instances since SAP Kernel versions 6.40. One classic problem of running SAP instances in a highly available environment is as follows: If an SAP administrator changes the status (start/stop) of an SAP instance without using the interfaces provided by the cluster software, the cluster framework will detect that as an error status and will bring the SAP instance into the old status by either starting or stopping the SAP instance. This can result in very dangerous situations, if the cluster changes the status of an SAP instance during some SAP maintenance tasks. This new updated solution enables the central component sapstartsrv to report state changes to the

cluster software, and therefore avoids the previously described dangerous situations. (See also the article "Using sap_vendor_cluster_connector for interaction between cluster framework and sapstartsrv" at <https://blogs.sap.com/2014/05/08/using-sapvendorclusterconnector-for-interaction-between-cluster-framework-and-sapstartsrv/comment-page-1/>).

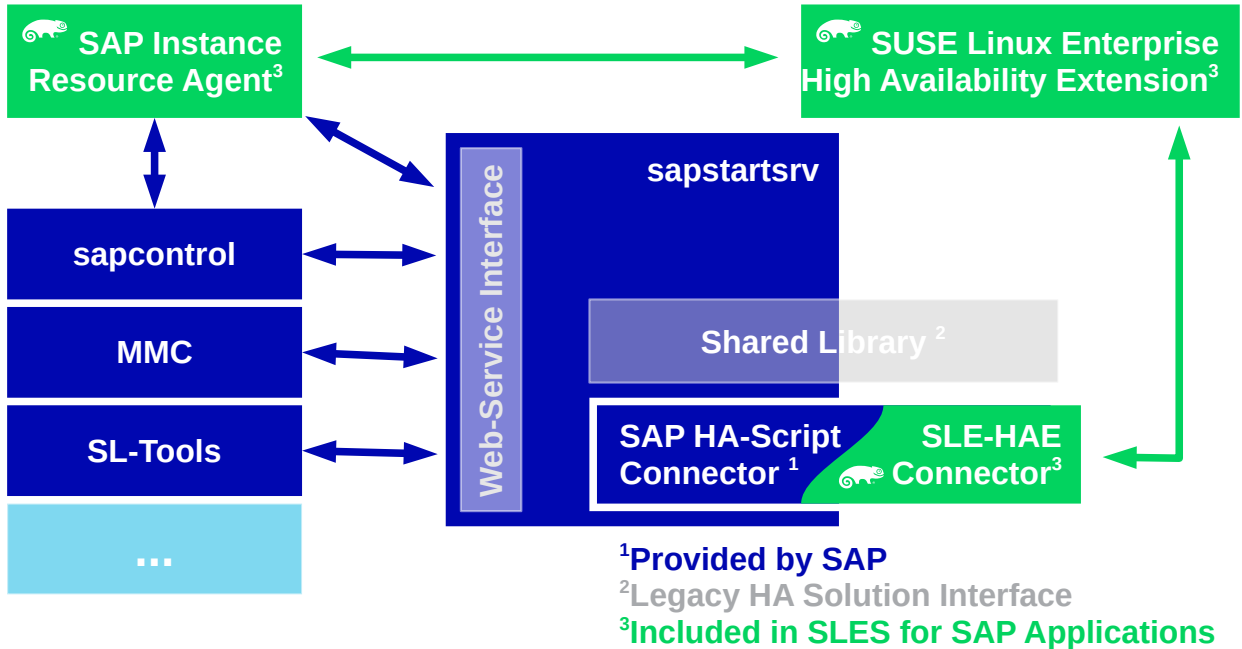


FIGURE 4: CLUSTER CONNECTOR TO INTEGRATE THE CLUSTER WITH THE SAP START FRAMEWORK



Note

For this scenario we are using an updated version of the sap-suse-cluster-connector which implements the API version 3 for the communication between the cluster framework and the sapstartsrv.

The new version of the sap-suse-cluster-connector now allows to start, stop and 'migrate' an SAP instance. The integration between the cluster software and the sapstartsrv also implements to run checks of the HA setup using either the command line tool sapcontrol or even the SAP management consoles (SAP MMC or SAP MC).

3.9 Disks and Partitions

For all SAP file systems beside the EFS file systems we are using XFS.

3.9.1 EFS File Systems for Cluster ASCS and ERS

Create the following sub directories on both cluster nodes as root:

```
# mkdir -p /usr/sap/HA1/ASCS00 /usr/sap/HA1/ERS10
```

The file systems for the ASCS and ERS instances need to be shared and assigned to the cluster nodes hacert01 and hacert02. Create an EFS file system.

During the SAP installation we need the filesystems `/usr/sap/HA1/ASCS00` to be mounted on hacert01 and `/usr/sap/HA1/ERS10` to be mounted on hacert02.

```
hacert01: efs-name:/ASCS00 /usr/sap/HA1/ASCS00
hacert02: efs-name:/ERS10 /usr/sap/HA1/ERS10
```

Replace the variable `efs-name` with the appropriate DNS name of the EFS file system.



Note

hacert01 and hacert02 operate in different availability zones. They need mount points named "efs-name" which are individual to the availability zone. Use the DNS name provided by AWS. The DNS name will point to the files system mount point local to a given Availability Zone. During the SAP installation we need `/usr/sap/HA1/ASCS00` to be mounted on hacert01 and `/usr/sap/HA1/ERS10` to be mounted on hacert02.

3.10 IP Addresses and Virtual Names

Check, if the file `/etc/hosts` contains at least the following address resolutions. Add those entries, if they are missing. The 10.0.0.0 addresses in the example below are primary IP addresses within the VPC CIDR block. The 192.168.201.0 addresses are the Overlay IP addresses for the virtual services. The listing below lists a virtual IP address for the database server. An SAP system installation against a virtual database server address will allow to upgrade the database server to be a protected cluster service in a later step.

```
10.0.0.111 hacert01
10.0.0.112 hacert02
10.0.0.113 hacert03
10.0.0.114 saphalci
10.0.0.115 saphald2
192.168.201.116 saphalas
```

```
192.168.201.117 saphaler
192.168.201.118 saphaldb
```

3.11 Mount Points and NFS Shares

In our setup the directory `/usr/sap` is part of the root file system. You could of course also create a dedicated file system for that area and mount `/usr/sap` during the system boot. As `/usr/sap` also contains the SAP control file `sapservices` and the `saphostagent`, the directory should not be placed on a shared file system between the cluster nodes.

We need to create the directory structure on all nodes which might be able to run the SAP resource. The `SYS` directory will be on a NFS share for all nodes.

- Creating mount points and mounting NFS share at all nodes
- Replace *efs-name* with the appropriate DNS name.

EXAMPLE 1: SAP NETWEAVER 7.4

```
# mkdir -p /sapmnt
# mkdir -p /usr/sap/HA1/{ASCS00,D02,DVEBMGS01,ERS10,SYS}
# mount -t nfs efs-name:/sapmnt /sapmnt
# mount -t nfs efs-name:/SYS /usr/sap/HA1/SYS
# mount -t nfs efs-name:/sapcd /sapcd
```

EXAMPLE 2: SAP NETWEAVER 7.5

```
# mkdir -p /sapmnt
# mkdir -p /usr/sap/HA1/{ASCS00,D01,D02,ERS10,SYS}
# mount -t nfs efs-name:/sapmnt /sapmnt
# mount -t nfs efs-name:/SYS /usr/sap/HA1/SYS
# mount -t nfs efs-name:/sapcd /sapcd
```

- Only HANA: creating mount points for database at `hacert03`:

```
# mkdir -p /hana/{shared,data,log}
```





FIGURE 5: FILE SYSTEM LAYOUT INCLUDING NFS SHARES

We prepare the three servers for the distributed SAP installation.

- Server 1 (hacert01) will be used to install the ASCS SAP instance.
- Server 2 (hacert02) will be used to install the ERS SAP instance
- Server 3 (hacert03) will be used to install the database.
- Server 4 (sapha1ci) will be used to install the PAS SAP instance.
- Server 5 (sapha1d2) will be used to install the AAS SAP instance.
- Mounting the instance and database file systems at one specific node
- As a result the directory `/usr/sap/HA1/` should now look like:

```
# ls -l /usr/sap/HA1/
total 0
drwxr-xr-x 1 haladm sapsys 70 28. Mär 17:26 ./
drwxr-xr-x 1 root sapsys 58 28. Mär 16:49 ../
drwxr-xr-x 7 haladm sapsys 58 28. Mär 16:49 ASCS00/
drwxr-xr-x 1 haladm sapsys 0 28. Mär 15:59 D02/
drwxr-xr-x 1 haladm sapsys 0 28. Mär 15:59 DVEBMGS01/
drwxr-xr-x 1 haladm sapsys 0 28. Mär 15:59 ERS10/
drwxr-xr-x 5 haladm sapsys 87 28. Mär 17:21 SYS/
```



Note

The owner of the directory and files are changed during the SAP installation. By default all of them are owned by root.

4 Testing the AWS Agents

Test the AWS agents before you start up the cluster. The tests will show whether the AWS role and the policies are being configured correctly. The tests should execute without AWS CLI errors

4.1 Testing the Overlay IP Agents

Replace the following variables in the commands

- *ip_address*: The service IP addresses of the ASCS and the ERS system
- *rtb-table* : The name of AWS routing table of the Overlay IP address
- *cluster* : replace the AWS CLI profile name if needed

The variables will need to match the variables in the OCF primitives later on!

Run the following commands as root on both systems:

```
OCF_RESKEY_address=ip_address \  
OCF_RESKEY_routing_table=rtb-table \  
OCF_RESKEY_interface=eth0 OCF_ROOT=/usr/lib/ocf OCF_RESKEY_profile=cluster \  
/usr/lib/ocf/resource.d/suse/aws-vpc-move-ip start  
  
OCF_RESKEY_address=ip_address \  
OCF_RESKEY_routing_table=rtb-table \  
OCF_RESKEY_interface=eth0 OCF_ROOT=/usr/lib/ocf OCF_RESKEY_profile=cluster \  
/usr/lib/ocf/resource.d/suse/aws-vpc-move-ip monitor  
  
OCF_RESKEY_address=ip_address \  
OCF_RESKEY_routing_table=rtb-table \  
OCF_RESKEY_interface=eth0 OCF_ROOT=/usr/lib/ocf OCF_RESKEY_profile=cluster \  
/usr/lib/ocf/resource.d/suse/aws-vpc-move-ip stop
```

Check for AWS CLI access issues and fix the AWS Policy. Use the AWS console to check whether the IP address got added after *start*. Use the AWS console to check whether the IP got removed after *stop*. Use the other cluster node to execute some access commands (ping, SSH etc.) Recheck and fix all network related settings if this does not work.

4.2 Testing Agents Mounting EFS File Systems

Test the monitoring function first. Replace the following variable in the commands:

- *efs-name*: The name of the EFS file system

Run the following commands as root on both cluster nodes:

```
OCF_RESKEY_device="efs-name:ASCS00" \  
OCF_RESKEY_directory="/usr/sap/HA1/ASCS00" OCF_RESKEY_fstype=nfs4 \  
OCF_RESKEY_options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
OCF_ROOT=/usr/lib/ocf /usr/lib/ocf/resource.d/heartbeat/Filesystem start  
  
OCF_RESKEY_device="efs-name:ERS10" \  
OCF_RESKEY_directory="/usr/sap/HA1/ERS10" OCF_RESKEY_fstype=nfs4 \  
OCF_RESKEY_options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
OCF_ROOT=/usr/lib/ocf /usr/lib/ocf/resource.d/heartbeat/Filesystem start  
  
df -k  
  
OCF_RESKEY_device="efs-name:ASCS00" \  
OCF_RESKEY_directory="/usr/sap/HA1/ASCS00" OCF_RESKEY_fstype=nfs4 \  
OCF_RESKEY_options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
OCF_ROOT=/usr/lib/ocf /usr/lib/ocf/resource.d/heartbeat/Filesystem stop  
  
OCF_RESKEY_device="efs-name:ERS10" \  
OCF_RESKEY_directory="/usr/sap/HA1/ERS10" OCF_RESKEY_fstype=nfs4 \  
OCF_RESKEY_options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
OCF_ROOT=/usr/lib/ocf /usr/lib/ocf/resource.d/heartbeat/Filesystem stop  
  
df -k
```

Check with the command *df -k* whether the file systems got mounted and unmounted. Potential problems arise with an incorrect *efs-name* or with missing subdirectories.

4.3 Optional: Testing Route 53 Agents

This test should be conducted if the Route 53 agent will be used. Test the monitoring function first. Replace the following variables in the commands:

- *hosted zone id*: The ID of the hosted private zones
- *fullname* : The full name of the service name including sub domain and a trailing dot.
- *cluster* : replace the AWS CLI profile name if needed

The variables will need to match the variables in the OCF primitives later on!

Run the following commands as root on both systems:

```
OCF_RESKEY_hostedzoneid=hosted zone id OCF_RESKEY_ttl=10 \  
  OCF_RESKEY_fullname=fullname OCF_ROOT=/usr/lib/ocf \  
  OCF_RESKEY_profile=cluster \  
  /usr/lib/ocf/resource.d/heartbeat/aws-vpc-route53 monitor  
  
OCF_RESKEY_hostedzoneid=hosted zone id OCF_RESKEY_ttl=10 \  
  OCF_RESKEY_fullname=fullname OCF_ROOT=/usr/lib/ocf \  
  OCF_RESKEY_profile=cluster \  
  /usr/lib/ocf/resource.d/heartbeat/aws-vpc-route53 start  
  
OCF_RESKEY_hostedzoneid=hosted zone id OCF_RESKEY_ttl=10 \  
  OCF_RESKEY_fullname=fullname OCF_ROOT=/usr/lib/ocf \  
  OCF_RESKEY_profile=cluster \  
  /usr/lib/ocf/resource.d/heartbeat/aws-vpc-route53 stop
```

Fix any problems in monitoring first. Try a start as second test and a stop as last test.

5 SAP Installation

The overall procedure to install the distributed SAP is:

- Installing the ASCS instance for the central services
- Installing the ERS to get a replicated enqueue scenario
- Preparing the ASCS and ERS installations for the cluster take-over
- Installing the Database

- Installing the primary application server instance (PAS)
- Installing additional application server instances (AAS)

The result will be a distributed SAP installation as illustrated here:

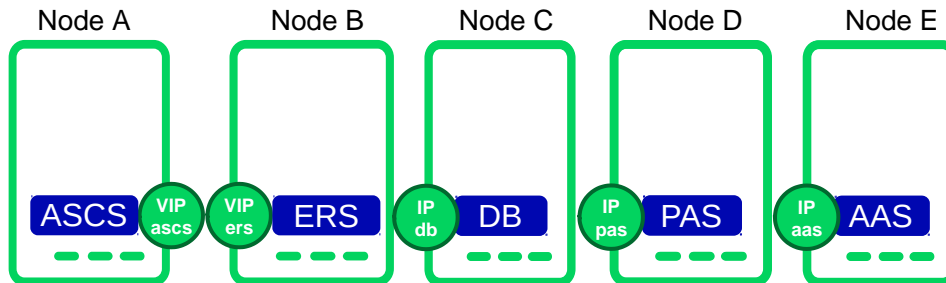


FIGURE 6: DISTRIBUTED INSTALLATION OF THE SAP SYSTEM

5.1 Linux User and Group Number Scheme

Whenever asked by the SAP software provisioning manager (SWPM) which Linux User IDs or Group IDs to use, refer to the following table as an example.

Group	sapinst	1000
Group	sapsys	1001
Group	sapadm	3000
Group	sdba	3002
User	haladm	3000
User	sdb	3002
User	sqdha1	3003
User	sapadm	3004
User	h04adm	4001

5.2 Installing ASCS on hacert01

Temporarily we need to set the service IP address which is used later in the cluster as local IP, because the installer wants to resolve or use it. Make sure to use the correct virtual host name for each installation step. Take care of the file systems like `efs-name:/ASCS00` and `/sapcd/` which might also need to be mounted.

```
# ip address add 192.168.201.116/32 dev eth0
# mount efs-name:/ASCS00 /usr/sap/HA1/ASCS00
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1as
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → ASCS Instance
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → ASCS Instance
- SID id HA1
- Use instance number 00
- Deselect using FQDN
- All passwords: use SuSE1234
- Double-check during the parameter review, if virtual name **sapha1as** is used

5.3 Installing ERS on hacert02

Temporarily we need to set the service IP address which is used later in the cluster as local IP, because the installer wants to resolve or use it. Make sure to use the correct virtual host name for each installation step.

```
# ip address add 192.168.201.117/32 dev eth0
# mount efs-name:/ERS10 /usr/sap/HA1/ERS10
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1er
```

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Enqueue Replication Server Instance
- Use instance number 10

- Deselect using FQDN
- Double-check during the parameter review, if virtual name **sapha1er** is used
- If you get an error during the installation about permissions, change the ownership of the ERS directory

```
# chown -R haladm:sapsys /usr/sap/HA1/ERS10
```

- If you get a prompt to manually stop/start the ASCS instance, log in to hacert01 as user haladm and call sapcontrol.

```
# sapcontrol -nr 00 -function Stop      # to stop the ASCS
# sapcontrol -nr 00 -function Start    # to start the ASCS
```

5.4 Poststeps for ASCS and ERS

5.4.1 Stopping ASCS and ERS

On *hacert01*

```
# su - haladm
# sapcontrol -nr 00 -function Stop
# sapcontrol -nr 00 -function StopService
```

On *hacert02*

```
# su - haladm
# sapcontrol -nr 10 -function Stop
# sapcontrol -nr 10 -function StopService
```

5.4.2 Maintaining *sapservices*

Ensure that the file */usr/sap/sapservices* holds both entries (ASCS + ERS) on both cluster nodes. Modify the file by copying the missing entry from both hosts across. Alternatively add the missing command string with *sapstartsrv*.

Example steps for ASCS profile on the ERS host:

```
# cd /usr/sap/hostctrl/exe
# export LD_LIBRARY_PATH=.
```

```
# ./sapstartsrv pf=/usr/sap/HA1/SYS/profile/HA1_ASCS00_saphalas -reg
```

This allows the sapstartsrv clients to start the service like:

As user ha1adm

```
# sapcontrol -nr 10 -function StartService HA1
```

The `/usr/sap/sapservices` looks like (typically one line per instance):

```
#!/bin/sh
LD_LIBRARY_PATH=/usr/sap/HA1/ASC00/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/HA1/ASC00/exe/sapstartsrv pf=/usr/sap/HA1/SYS/profile/HA1_ASC00_saphalas -D -u
ha1adm
LD_LIBRARY_PATH=/usr/sap/HA1/ERS10/exe:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH; /usr/
sap/HA1/ERS10/exe/sapstartsrv pf=/usr/sap/HA1/ERS10/profile/HA1_ERS10_saphaler -D -u
ha1adm
```

5.4.3 Integrating the Cluster Framework Using the sap-suse-cluster-connector Package

Install the package **sap-suse-cluster-connector** version > 3.0.0 from our repositories:

```
# zypper install sap-suse-cluster-connector
```



Note

The package `sap-suse-cluster-connector` with version 3.0.x implements the SUSE SAP API version 3. New features like SAP Rolling Kernel Switch (RKS) and migration of ASCS are only supported with this new version.

For the ERS and ASCS instance edit the instance profile files `HA1_ASCS00_saphalas` and `HA1_ERS10_saphaler` in the profile directory `/usr/sap/HA1/SYS/profile/`.

You need to tell the sapstartsrv service to load the HA script connector library and to use the `sap-suse-cluster-connector`.

```
service/halib = $(DIR_EXECUTABLE)/saphascriptco.so
service/halib_cluster_connector = /usr/bin/sap_suse_cluster_connector
```

Add the user `ha1adm` to the unix user group `haclient`.

```
# usermod -a -G haclient ha1adm
```


5.4.4 Adapting SAP Profiles to Match the SAP NW-HA-CLU 7.40 Certification

For the ASCS, change the start command from *Restart_Program_xx* to *Start_Program_xx* for the enqueue server (enserver). This change tells the SAP start framework **not** to self-restart the enqueue process. Such a restart would lead in loss of the locks.

EXAMPLE 3: FILE /USR/SAP/HA1/SYS/PROFILE/HA1_ASCS00_SAPHA1AS

```
Start_Program_01 = local $_EN pf=$_PF)
```

Optionally you could limit the number of restarts of services (in the case of ASCS this limits the restart of the message server).

For the ERS change instance the start command from *Restart_Program_xx* to *Start_Program_xx* for the enqueue replication server (enrepsrver).

EXAMPLE 4: FILE /USR/SAP/HA1/SYS/PROFILE/HA1_ERS10_SAPHA1ER

```
Start_Program_00 = local $_ER pf=$_PFL NR=$(SCSID)
```

5.4.5 Starting ASCS and ERS

On hacert01

```
# su - haladm
# sapcontrol -nr 00 -function StartService HA1
# sapcontrol -nr 00 -function Start
```

On hacert02

```
# su - haladm
# sapcontrol -nr 10 -function StartService HA1
# sapcontrol -nr 10 -function Start
```

5.5 Installing DB on hacert03 (Example SAP HANA)

The HANA DB has very strict HW requirements. The storage sizing depends on many indicators. Check the supported configurations at [SAP HANA Hardware Directory \(https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630\)](https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630) and [SAP HANA TDI \(https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html\)](https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html).

Install the HANA file systems as being described in the section "Planning the Deployment" of the [AWS SAP HANA on the AWS Cloud: Quick Start Reference Deployment](http://docs.aws.amazon.com/quickstart/latest/sap-hana/planning.html) (<http://docs.aws.amazon.com/quickstart/latest/sap-hana/planning.html>) ↗

Consider to install the database against an Overlay IP address which acts like a service IP address. This will allow to upgrade the database to run in a SUSE Linux Enterprise High Availability Extension for SAP cluster:

```
# ip address add 192.168.201.118/32 dev eth0
# mount /dev/sdc1 /hana/shared
# mount /dev/sdc2 /hana/log
# mount /dev/sdc3 /hana/data
# cd /sapcd/SWPM/
# ./sapinst SAPINST_USE_HOSTNAME=sapha1db
```

- We are installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Database Instance
- Profile directory /sapmnt/HA1/profile
- Deselect using FQDN
- Database parameters enter DBSID is H04; Database Host is sapha1db; Instance Number is 00
- Database System ID enter Instance Number is 00; SAP Mount Directory is /hana/shared
- Account parameters change them in case of custom values needed
- Cleanup select **Yes**, remove operating system users from group'sapinst'....
- Double-check during the parameter review, if virtual name **sapha1db** is used

5.6 Installing the Primary Application Server (PAS) on sapha1ci

Add the following mount points to the */etc/fstab* file on host sapha1ci. Replace the string "efs_fs_local_az" with the IP address of your EFS service in your availability zone.

```
efs-name:sapcd /sapcd nfs4 rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2 0 0
efs-name:sapmnt /sapmnt nfs4 rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2 0 0
efs-name:DVEBMGS01 /usr/sap/HA1/DVEBMGS01 nfs4 rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2 0 0
```

```
efs-name:SYS      /usr/sap/HA1/SYS      nfs4
rsize=1048576,ws=1048576,hard,timeo=600,retrans=2 0 0
```

Replace the variable *efs-name* with the appropriate DNS name.

Create mount directories and mount the file systems

```
# mkdir -p /sapcd /sapmnt /usr/sap/HA1/SYS /usr/sap/HA1/DVEBMGS01
# mount -a
```

Install the PAS server with the *sapinst* tool.

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Primary Application Server Instance (PAS)
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Primary Application Server Instance (PAS)
- Use instance number 01
- Deselect using FQDN
- For our hands-on setup use a default secure store key
- Do not install Diagnostic Agent
- No SLD
- Double-check during the parameter review, if virtual name **sapha1ci** is used

5.7 Installing an Additional Application Server (AAS) on sapha1d2

Add the following mount points to the */etc/fstab* file on host *sapha1d2*. Replace the string "efs_fs_local_az" with the IP address of your EFS service in your availability zone.

```
efs-name:sapcd    /sapcd            nfs4
rsize=1048576,ws=1048576,hard,timeo=600,retrans=2 0 0
efs-name:sapmnt   /sapmnt           nfs4
rsize=1048576,ws=1048576,hard,timeo=600,retrans=2 0 0
efs-name:D02      /usr/sap/HA1/D02  nfs4
rsize=1048576,ws=1048576,hard,timeo=600,retrans=2 0 0
efs-name:SYS      /usr/sap/HA1/SYS  nfs4
rsize=1048576,ws=1048576,hard,timeo=600,retrans=2 0 0
```

Replace the variable *efs-name* with the appropriate DNS name. Create mount directories and mount the file systems

```
# mkdir -p /sapcd /sapmnt /usr/sap/HA1/SYS /usr/sap/HA1/D02
# mount -a
```

Install the AAS server with the *sapinst* tool.

- SWPM option depends on SAP NetWeaver version and architecture
 - Installing SAP NetWeaver 7.40 SR2 → MaxDB → SAP-Systems → Application Server ABAP → High-Availability System → Additional Application Server Instance (AAS)
 - Installing SAP NetWeaver 7.5 → SAP HANA Database → Installation → Application Server ABAP → High-Availability System → Additional Application Server Instance (AAS)
- Use instance number 02
- Deselect using FQDN
- Do not install Diagnostic Agent
- Double-check during the parameter review, if name **sapha1d2** is used

6 Implementing the Cluster

The main procedure to implement the cluster is:

- Install the cluster software, if not already done during the installation of the operating system
- Configure the cluster communication framework corosync
- Configure the cluster resource manager
- Configure the cluster resources



Note

Before we continue to set up the cluster we first stop all SAP instances, remove the (manually added) IP addresses on the cluster nodes and unmount the file systems which will be controlled by the cluster later.

TASKS

1. Set up NTP (best with YaST). Use AWS time service at 169.254.169.123 which is accessible from all EC2 instances. Enable ongoing synchronization.
2. Install pattern *ha_sles* on both cluster nodes

```
# zypper install -t pattern ha_sles
```

Activate the public cloud module to get updates for the AWS CLI (Command Line Interface):

```
# SUSEConnect --list-extensions  
# SUSEConnect -p sle-module-public-cloud/12/x86_64
```

Update your packages with the command:

```
# zypper update
```

6.1 Configuring the Cluster Base

TASKS

- Install and configure the cluster stack at the first machine

6.1.1 Configuration of System Logging

SUSE recommends to use rsyslogd for logging with the SUSE cluster. This is a default configuration. Some AWS AMIs however use syslogd logging. Perform the following commands as root on all cluster nodes:

```
# zypper install rsyslog
```

Use option 1 (deinstallation of competing software, syslogd). Reboot both nodes.

6.1.2 Corosync Configuration

6.1.2.1 Configuration of the *corosync.conf* File

The configuration will have an IP address for node node-1 which is supposed to be ip-node-1 and ip2-node-1. Node node-2 has an ip address to which we refer as ip-node-2 and ip2-node-2.

All cluster nodes are required to have a local configuration file `/etc/corosync/corosync.conf` which will be structured as follows.

The relevant information is being located in the two sections describing interface and nodelist. The other entries can be configured as needed for a specific implementation.



Note

AWS requires a specific manual corosync configuration.

Use the following configuration in the `/etc/corosync/corosync.conf` file on both cluster nodes:

```
# Read the corosync.conf.5 manual page
totem {
  version: 2
  rrp_mode: passive
  token: 30000
  consensus: 32000
  token_retransmits_before_loss_const: 10
  max_messages: 20
  crypto_cipher: none
  crypto_hash: none
  clear_node_high_bit: yes
  interface {
    ringnumber: 0
    bindnetaddr: <ip-local-node>
    mcastport: 5405
    ttl: 1
  }
  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>
ring1_addr: <ip2-node-1>
nodeid: 1
}
node {
ring0_addr: <ip-node-2>
ring1_addr: <ip2-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
}

```

Replace the variables *ip-node-1* / *ip2-node-1* and *ip-node-2* / *ip2-node-2* with the IP addresses of your two cluster instances. Replace *ip-local-node* with the IP address of the server the file is being created.

The chosen settings for *crypto_cipher* and *crypto_hash* are suitable for clusters in AWS. They may be modified according to SUSE's documentation if strong encryption of cluster communication is desired.

6.1.3 Starting the Cluster

The next step is to start the cluster with the command on both nodes:

```
# systemctl start pacemaker
```

6.1.4 Checking the Configuration

The configuration can be checked with the command:

```
# corosync-cfgtool -s
```

It will create a result like the following one for a cluster node with the IP address 10.0.0.111:

```

Printing ring status.
Local node ID 1
RING ID 0
id = 10.0.0.111

```

```
status = ring 0 active with no faults
```

The cluster in question has been using ring 0, the node had the ID 1.

- The `crm_mon -1` output should look like this:

```
Stack: corosync
Current DC: hacert01 (version 1.1.15-19.15-e174ec8) - partition with quorum
Last updated: Wed Dec 6 16:02:42 2017
Last change: Wed Dec 6 15:44:45 2017 by hacluster via crmd on hacert01

2 nodes configured
0 resources configured

Online: [ hacert01 hacert02 ]

Full list of resources:
```

6.2 Configuring Cluster Resources

We need a changed SAPInstance resource agent for SAP NetWeaver to **not** use the master-slave construct anymore and to move to a more cluster-like construct. This allows to start and stop the ASCS and the ERS itself and **not** only the complete master-slave.

For this there is a new functionality for the ASCS needed to follow the ERS. The ASCS needs to mount the shared memory table of the ERS to avoid the loss of locks.

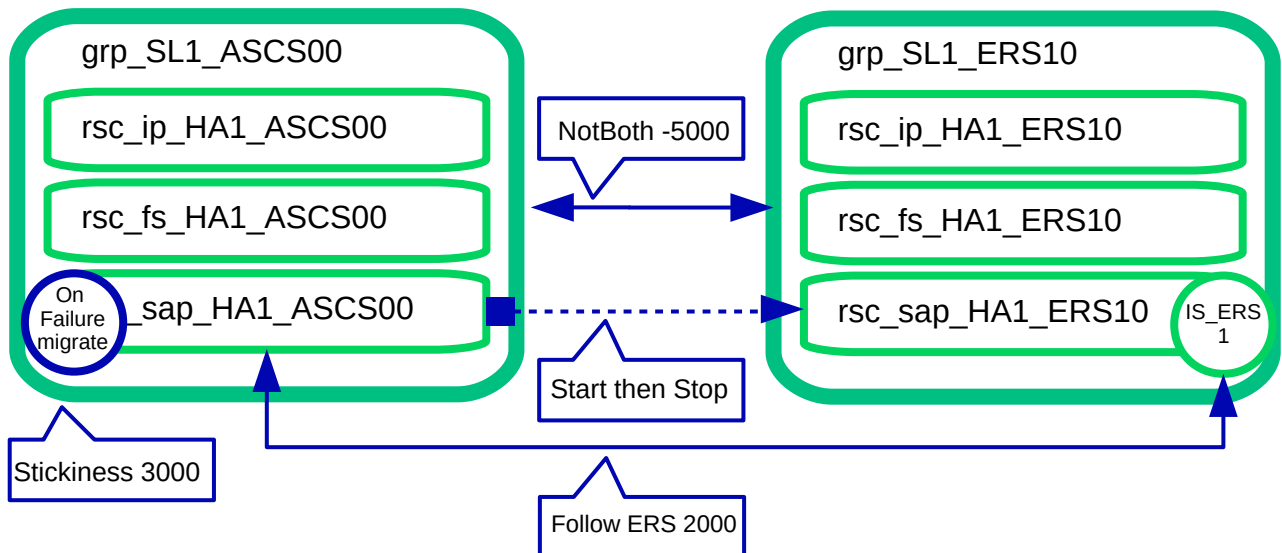


FIGURE 7: RESOURCES AND CONSTRAINTS

The implementation is done using a new flag "runs_ers_\$SID" within the RA, enabled with the help of the resource parameter "IS_ERS=TRUE".

There is the option to add a Route 53 agent. The architecture will then look as follows:

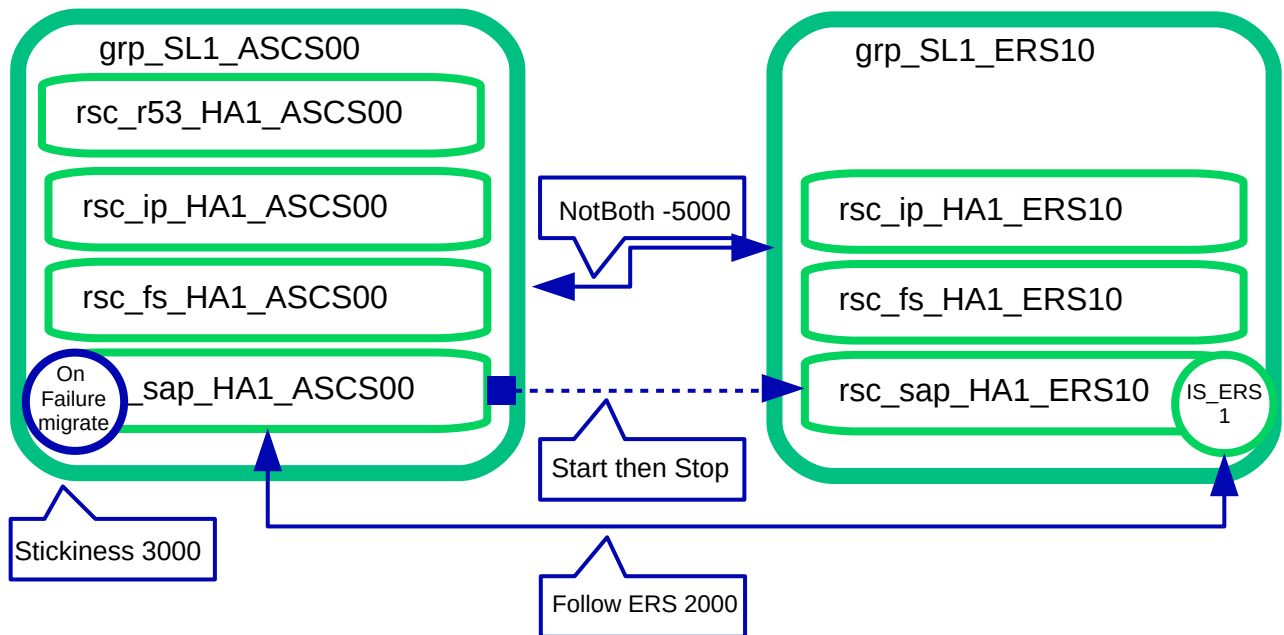


FIGURE 8: RESOURCES AND CONSTRAINTS

6.2.1 Preparing the Cluster for Adding the Resources

To avoid that the cluster starts partially defined resources we set the cluster to the maintenance mode. This deactivates all monitor actions.

As user root

```
# crm configure property maintenance-mode="true"
```

6.2.2 Configuring AWS-Specific Settings

Execute the following commands on one of the two cluster nodes:

```
# vi crm-bs.txt
```

Enter the following information to the file `crm-bs.txt`:

```
property cib-bootstrap-options: \
```

```

stonith-enabled="true" \
stonith-action="off" \
stonith-timeout="600s"
rsc_defaults rsc-options: \
resource-stickiness=1 \
migration-threshold=3
op_defaults op-options: \
timeout=600 \
record-pending=true

```

The setting *off* forces the agents to shut down the instance. This is desirable to avoid split brain scenarios on AWS.

Add the configuration to the cluster:

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

6.2.3 Configuration of AWS specific Stonith Resource

Create a file with the following content:

```

primitive res_AWS_STONITH stonith:external/ec2 \
op start interval=0 timeout=180 \
op stop interval=0 timeout=180 \
op monitor interval=300 timeout=60 \
params tag=pacemaker profile=cluster

```

The EC2 tag *pacemaker* entry needs to match the tag chosen for the EC2 instances. The value for this tag will contain the host name. The name of the profile (*cluster* in this example) will need to match the previously configured AWS profile.

Name this file for example *aws-stonith.txt* and add this file to the configuration. The following command needs to be issued as root. It uses the file name *aws-stonith.txt*:

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

6.2.4 Configuring the Resources for the ASCS

First we configure the resources for the file system, IP address and the SAP instance. Of course you need to adapt the parameters to your environment.

Create a file with your editor of choice with a name *aws-ascv.txt*. Add the ASCS primitive and the ASCS group to it. Do not forget to save your changes.

EXAMPLE 5: ASCS PRIMITIVE

```
primitive rsc_fs_HA1_ASCS00 Filesystem \  
  params device="efs-name:/ASCS00" directory="/usr/sap/HA1/ASCS00" \  
         fstype="nfs4" \  
         options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=200s timeout=40s  
primitive rsc_ip_HA1_ASCS00 ocf:suse:aws-vpc-move-ip \  
  params address=192.168.201.116 routing_table=rtb-table \  
         interface=eth0 profile=cluster \  
  op start interval=0 timeout=180 \  
  op stop interval=0 timeout=180 \  
  op monitor interval=60 timeout=60  
primitive rsc_sap_HA1_ASCS00 SAPIInstance \  
  operations $id=rsc_sap_HA1_ASCS00-operations \  
  op monitor interval=120 timeout=60 on-fail=restart \  
  params InstanceName=HA1_ASCS00_saphalas \  
         START_PROFILE="/sapmnt/HA1/profile/HA1_ASCS00_saphalas" \  
         AUTOMATIC_RECOVER=false \  
  meta resource-stickiness=5000 failure-timeout=60 \  
         migration-threshold=1 priority=10
```

Replace the variable *efs-name* with the name of your EFS server.

Replace the variable *rtb-table* with the identifier of the appropriate AWS routing table for the subnets. The name of the AWS CLI profile (*cluster* in this example) will need to match the previously configured AWS profile.

EXAMPLE 6: ASCS GROUP

```
group grp_HA1_ASCS00 \  
  rsc_ip_HA1_ASCS00 rsc_fs_HA1_ASCS00 rsc_sap_HA1_ASCS00 \  
  meta resource-stickiness=3000
```

Create a txt file *aws_ascs.txt* with your preferred text editor, enter both examples (primitives and group) to that file and load the configuration to the cluster manager configuration.

```
# crm configure load update aws_ascs.txt
```

6.2.5 Optional: Including Route 53

Name this file for example *aws-route53.txt* and add this file to the configuration. The following command needs to be issued as root. It uses the file name *aws-route53.txt*:

Enter the following primitive before or after the existing primitives in the editor:

EXAMPLE 7: ROUTE53 PRIMITIVE

```
primitive rsc_r53_HA1_ASCS00 ocf:heartbeat:aws-vpc-route53 \  
  params hostedzoneid=route-53-name ttl=10 fullname=name-full. profile=cluster \  
  op start interval=0 timeout=180 \  
  op stop interval=0 timeout=180 \  
  op monitor interval=300 timeout=180
```

Replace the variable *route-53-name* with the name of the associated private hosted Route 53 zone.

Replace the variable *name-full.* with the fully qualified host name with matches the private hosted Route 53 zone.

The agent uses a time-to-live (ttl) of 10 seconds in this example. Change this parameter if needed.

Insert the *rsc_r53_HA1_ASCS00* after the *rsc_ip_HA1_ASCS00*. This will force the group to update then Route 53 as second item after the Overlay IP address.

EXAMPLE 8: ROUTE53 GROUP

```
group grp_HA1_ASCS00 \  
  rsc_ip_HA1_ASCS00 rsc_r53_HA1_ASCS00 \  
  rsc_fs_HA1_ASCS00 rsc_sap_HA1_ASCS00 \  
  meta resource-stickiness=3000
```

Create a txt file *aws-route53.txt* with your preferred text editor, enter both examples (primitives and group) to that file and load the configuration to the cluster manager configuration. Use the following command as root and modify ASCS group in the editor.

```
# crm configure load update aws-route53.txt
```



Note

Version 1.0.2 of the Route 53 agent will not work if the EC2 metadata contains a string like "local-ipv4" in the userdata section!

6.2.6 Configuring the Resources for the ERS

Second we configure the resources for the file system, IP address and the SAP instance. Of course you need to adapt the parameters to your environment.

Replace *efs-name* with the name of your EFS server.

The specific parameter *IS_ERS = true* should only be set for the ERS instance.

EXAMPLE 9: ERS PRIMITIVE

```
primitive rsc_fs_HA1_ERS10 Filesystem \  
  params device="efs-name:/ERS10" directory="/usr/sap/HA1/ERS10" fstype=nfs4 \  
  options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \  
  op start timeout=60s interval=0 \  
  op stop timeout=60s interval=0 \  
  op monitor interval=200s timeout=40s  
primitive rsc_ip_HA1_ERS10 ocf:suse:aws-vpc-move-ip \  
  params address=192.168.201.117 routing_table=rtb-table \  
  interface=eth0 profile=cluster \  
  op start interval=0 timeout=180 \  
  op stop interval=0 timeout=180 \  
  op monitor interval=60 timeout=60  
primitive rsc_sap_HA1_ERS10 SAPIInstance \  
  operations $id=rsc_sap_HA1_ERS10-operations \  
  op monitor interval=120 timeout=60 on-fail=restart \  
  params InstanceName=HA1_ERS10_saphaler \  
    START_PROFILE="/sapmnt/HA1/profile/HA1_ERS10_saphaler" \  
    AUTOMATIC_RECOVER=false IS_ERS=true \  
  meta priority=1000
```

Replace the variable *rtb-table* with the identifier of the appropriate AWS routing table for the subnets. The name of the AWS CLI profile (*cluster* in this example) will need to match the previously configured AWS profile.

EXAMPLE 10: ERS GROUP

```
group grp_HA1_ERS10 \  
  rsc_ip_HA1_ERS10 rsc_fs_HA1_ERS10 rsc_sap_HA1_ERS10
```

Create a txt file (like *aws_crm_ers.txt*) with your preferred text editor, enter both examples (primitives and group) to that file and load the configuration to the cluster manager configuration.

As user root

```
# crm configure load update aws_crm_ers.txt
```

6.2.7 Configuring the Colocation Constraints Between ASCS and ERS

The constraints between the ASCS and ERS instance are needed to define that the ASCS instance should start up exactly on the cluster node running the ERS instance after a failure (`loc_sap_HA1_failover_to_ers`). This constraint is needed to ensure that the locks are not lost after an ASCS instance (or node) failure.

If the ASCS instance has been started by the cluster the ERS instance should be moved to an "other" cluster node (`col_sap_HA1_no_both`). This constraint is needed to ensure that the ERS will synchronize the locks again and the cluster is ready for an additional take-over.

EXAMPLE 11: LOCATION CONSTRAINT

```
colocation col_sap_HA1_no_both -5000: grp_HA1_ERS10 grp_HA1_ASCS00
location loc_sap_HA1_failover_to_ers rsc_sap_HA1_ASCS00 \
    rule 2000: runs_ers_HA1 eq 1
order ord_sap_HA1_first_start_ascs Optional: rsc_sap_HA1_ASCS00:start \
    rsc_sap_HA1_ERS10:stop symmetrical=false
```

Create a text file (like `crm_col.txt`) with your preferred text editor, enter all three constraints to that file and load the configuration to the cluster manager configuration.

Issue the following command as root:

```
# crm configure load update crm_col.txt
```

6.2.8 Activating the Cluster

Now the last step is to end the cluster maintenance mode and to allow the cluster to detect already running resources.

Issue the following command as root:

```
# crm configure property maintenance-mode="false"
```

The cluster will now start the ASCS and the ERS system. This can take a few minutes. Check progress with the command:

```
# crm status
```

7 Administration

7.1 Dos and Don'ts

7.1.1 Never Stop the ASCS Instance

For normal operation **do not stop** the ASCS SAP instance with any tool such as cluster tools or SAP tools. The stop of the ASCS instance might lead to a loss of enqueue locks. Because following the new SAP NW-HA-CLU 7.40 certification the cluster must allow local restarts of the ASCS. This feature is needed to allow Rolling Kernel Switch (RKS) updates without reconfiguring the cluster.



Warning

Stopping the ASCS instance might lead into the loss of SAP enqueue locks during the start of the ASCS on the same node.

7.1.2 How to Migrate ASCS

To **migrate** the ASCS SAP instance you should use the SAP tools such as the SAP management console. This will trigger `sapstartsrv` to use the `sap-suse-cluster-connector` to migrate the ASCS instance. As user `ha1adm` you might call the following command to migrate-away the ASCS. The migrate-away will always migrate the ASCS to the ERS side which will keep the SAP enqueue locks.

As ha1adm

```
# sapcontrol -nr 00 -function HAFailoverToNode ""
```

7.1.3 Never Block Resources

With SAP NW-HA-CLU 7.40 it is **not longer allowed to block resources** from being controlled manually. This using the variable `BLOCK_RESOURCES` in `/etc/sysconfig/sap_suse_cluster_connector` is not allowed anymore.

7.1.4 Always Use Unique Instance Numbers

Currently all SAP **instance numbers controlled by the cluster must be unique**. If you need to have multiple dialog instances such as D00 running on different systems they should be not controlled by the cluster.

7.1.5 How to Set the Cluster in Maintenance Mode

Procedure to set the cluster into maintenance mode can be done as *root* or *sidadm*.

As user root

```
# crm configure properties maintenance-mode="true"
```

As user ha1adm (the full path is needed)

```
# /usr/sbin/crm configure properties maintenance-mode="true"
```

7.1.6 Procedure to End the Cluster Maintenance

As user root

```
# crm configure properties maintenance-mode="false"
```

7.1.7 Cleanup Resources

How to **cleanup resource failures**? Failures of the ASCS will be automatically deleted to allow a failback after the configured period of time. For all other resources you can cleanup the status including the failures:

As user root

```
# crm resource cleanup RESOURCE-NAME
```



Warning

You should not cleanup the complete group of the ASCS resource as this might lead into an unwanted cluster action to take over the complete group onto the node where ERS instance is running.

7.2 Test the Cluster

We strongly recommend that you at least process the following tests before you plan going into production with your cluster.

7.2.1 Check Product Names with HAGetFailoverConfig

Check if the name of the SUSE cluster solution is shown in the output of sapcontrol or SAP management console. This test checks the status of the SAP NetWeaver cluster integration.

As user ha1adm

```
# sapcontrol -nr 00 -function HAGetFailoverConfig
```

7.2.2 Start SAP Checks Using HACheckConfig and HACheckFailoverConfig

Check if the HA configuration tests are showing no errors.

As user ha1adm

```
# sapcontrol -nr 00 -function HACheckConfig  
# sapcontrol -nr 00 -function HACheckFailoverConfig
```

7.2.3 Manually Migrate ASCS

Check if manually migrating the ASCS using HA tools works properly.

As user root

```
# crm resource migrate rsc_sap_HA1_ASCS00 force  
## wait till the ASCS is been migrated to the ERS host  
# crm resource unmigrate rsc_sap_HA1_ASCS00
```

7.2.4 Migrate ASCS Using HAFailoverToNode

Check if moving the ASCS instance using SAP tools like sapcontrol does work properly.

As user ha1adm

```
# sapcontrol -nr 00 -function HAFailoverToNode ""
```

7.2.5 Test ASCS Migration After Failure

Check if the ASCS instance moves correctly after a node failure.

As user root

```
## on the ASCS host
# echo b >/proc/sysrq-trigger
```

7.2.6 Inplace Restart of ASCS Using Stop and Start

Check if the inplace re-start of the SAP resources have been processed correctly. The SAP instance should not failover to an other node, it must start on the same node where it has been stopped.



Warning

This test will force the SAP system to **lose** the enqueue locks. **This test should not be processed during production.**

As user ha1adm

```
## example for ASCS
# sapcontrol -nr 00 -function Stop
## wait till the ASCS is completely down
# sapcontrol -nr 00 -function Start
```

7.2.7 Additional Tests to Perform

- Automated restart of the ASCS (simulating RKS)
- Check the recoverable and non-recoverable outage of the message server process
- Check the non-recoverable outage of the SAP enqueue server process
- Check the outage of the SAP Enqueue Replication Server
- Check the outage and restart of sapstartsrv
- Check the Rolling Kernel Switch procedure (RKS), if possible
- Check the simulation of an upgrade
- Check the simulation of cluster resource failures

8 AWS Specific Post Installation Tasks

The optional installation of the Route 53 agent will update the DNS name of the central instance as needed. The Route 53 naming tables may need to be made visible for on premises users like SAP GUI users. This happens through updating the on-premises name servers to delegate name resolution to Route 53. This forwarding of name resolution requests acquires an extra configuration in the AWS VPC.

Active directory users need to configure an Active Directory Connector as described in [<https://aws.amazon.com/de/blogs/security/how-to-set-up-dns-resolution-between-on-premises-networks-and-aws-using-aws-directory-service-and-amazon-route-53/>].

DNS server users need to implement bind forwarding EC2 instances as described in [<http://www.scalingbits.com/aws/dnsfailover/backpropagation>].

9 Additional Implementation Scenarios

9.1 Adaptive Server Enterprise Replication Failover Automation Integration

9.1.1 FM Integration with SUSE Linux Enterprise High Availability Extension Cluster

Standard SAP on AWS for an HA setup is Multi-AZ deployment with ASCS, Primary DB running in one AZ and their counterpart ERS and Secondary DB running in the second AZ of the same region. The Primary Application Server & Additional Application servers based on the load can be distributed in both AZ's as well to provide resiliency. Considering a scenario where SAP NetWeaver or Business Suite system is running on SAP Sybase ASE. The completely automated HA for the ABAP Stack (ASCS) is provided by the SUSE Linux Enterprise High Availability Extension Cluster. For the Sybase ASE DB the HA feature is provided with the Always On configuration and the failover orchestration is done by the Fault Manager (FM) utility which traditionally was installed on a third host (other than the Primary & Secondary DB). In an SAP world the FM utility comes along with an SAP DB dependent Kernel and gets installed in the ASCS Work

directory `/usr/sap/<SID>/ASCS<instnr>/exe/`. The failover of the ASCS instance along with the associated directories (provided they are installed on a shared file system using either Amazon EFS or NFS) is taken care by the SUSE Linux Enterprise High Availability Extension Cluster.

9.1.2 Sybase ASE Always On

SAP Sybase ASE comes with an Always On feature which provides native HA & DR capability. The always-on option is a high availability and disaster recovery (HADR) system that consists of two SAP ASE servers: One is designated as the primary server, on which all transaction processing takes place. The other acts as a warm standby (called "standby server" in DR mode, and as a "companion" in HA mode) for the primary server, and contains copies of designated databases from the primary server. The failover orchestration is carried out by ASE provided utility called Fault Manager. The Fault Manager monitors the various components of the HADR environment – Replication Management Agent (RMA), ASE, Replication Server, applications, databases, and the operating system. Its primary job is to ensure the high-availability (zero data loss during failover) of the ASE cluster by initiating automatic failover with minimal manual intervention. In an SAP Stack, the fault manager utility (`sybdbfm`) comes as part of the DB (Sybase ASE) dependent SAP Kernel. Refer to the SAP Standard ASE HA-DR guide (<https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.6/en-US/a6645e28bc2b1014b54b8815a64b87ba.html>) for setting up the Sybase ASE DB in HA mode.

Important

In the following section we use sometimes examples and sometimes general examples. In the general are terms like `<SID>`; `<instance nr>`. They must be adapted to your environment. As an example, `su - <sid>adm` means `su - ha7adm` or in capital letters `cd /usr/sap/<SID>/ASCS<instance nr>/work` means `cd /usr/sap/HA7/ASCS00/work`

9.1.3 Database Host Preparation

This guide does not duplicate the official HADR documentation. The following procedure describes the key points which you need to take care of.

EXAMPLE 12: INSTALLATION 32-BIT ENVIRONMENT

```
# zypper install glibc-32bit libgcc_s1-32bit
```

For the example this software stack is used:

- SL TOOLSET 1.0 — SWPM → 1.0 SP25 for NW higher than 7.0x
- saphostagent → 7.21 patch 41
- SAP Kernel → 7.53 PL421
- SAP Installation Export → (51051806_1)
- Sybase RDBMS → ASE 16.0.03.06 RDBMS (51053561_1)



Note

Very useful is that short table of installation information which helps to be prepared for the next steps. SAP Adaptive Server Enterprise - Installation Worksheet <https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/16.0.3.6/en-US/3fe35550f3814b2bb411d5494976e25a.html>



Important

The Fault Manager is enhanced to work in this setup. The minimal version's which support this scenario are * SAP Kernel 749 PL632 * SAP Kernel 753 PL421

9.1.4 Database Installation for Replication Scenario

The installation can be done with the SWPM which is provided by SAP.

INSTALLING THE PRIMARY DATABASE WITH SWPM:

- SWPM option depends on SAP NetWeaver version and architecture
 - Software Provisioning Manager 1.0 SP 25 → SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Installation → Application Server ABAP → High-Availability System → Database Instance

The following information is requested from the wizard:

- Master Password <secure >
- SAP System Code Page: Unicode (default)

- Uncheck: → Set FQDN for SAP system
- Sybase database Administrator UID: 2003
- In our test setup we uncheck → Use separate devices for sybmgmtldb database

After the basis installation is finished the primary database must be prepared for the replication. First the user **sa** must be unlocked.

```
# su - syb<sid>
# isql -Usapss0 -P <secure password> -S<SID> -X -w1900
# 1> go
# 1> exec sp_locklogin sa, 'unlock'
# 2> go
# Account unlocked.
# (return status = 0)
# 1> quit
```

In the next step, install the SRS software with a response file and enter the following command as user *syb* <sid>:

```
# /sapcd/ase-16.0.03.06/BD_SYBASE_ASE_16.0.03.06_RDBMS_for_BS_/SYBASE_LINUX_X86_64/
setup.bin -f /sybase/HA7/srs-setup.txt -i silent
```

Activate HADR on primary node with a response file and enter the following command as user *syb* <sid>:

```
# setuphadr /sybase/HA7/HA7_primary_lin.rs.txt
```



Note

If the installation stops with an error message as displayed here, perform the steps explained below:

```
Clean up environment.
Environment cleaned up.
Error: Fail to connect to "PRIM" site SAP ASE at "<hostname>:4901".
```

Check if the host name and port number are correct and the database server is up and running. If everything is correct and network connection should be available, it might help to modify the *interface* file. Try to add a new line in the */sybase/<SID>/interfaces* file for the <SID> section with the IP address of the corresponding host name.

```
# vi /sybase/<SID>/interfaces
...
master tcp ether <hostname> 4901
master tcp ether 172.17.1.21 4901
...
```

Create a secure store key entry for the database:

```
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -function
LiveDatabaseUpdate -dbname <SID> -dbtype syb -dbuser DR_admin -dbpass <Secure password>
-updatemethod Execute -updateoption TASK=SET_USER_PASSWORD -updateoption USER=DR_ADMIN
```

INSTALLING THE COMPANION DATABASE WITH SWPM:

- SWPM option depends on SAP NetWeaver version and architecture
 - Software Provisioning Manager 1.0 SP 25 → SAP NetWeaver AS for ABAP 7.52 → SAP ASE → Database Replication → Setup of Replication Environment

The following information is requested from the wizard:

- Replication System Parameters → SID, Master Password, check Set up a secondary database instance
- Primary Database server → host name or virt. name
- Primary Database server port → default is 4901, depends on the setup of your primary server

After the basis installation is finished the companion database must be prepared for the replication. First the user **sa** must be unlocked.

```
# su - syb<sid>
# isql -Usapss0 -P <secure password> -S<SID> -X -w1900
# 1> go
# 1> exec sp_locklogin sa, 'unlock'
# 2> go
# Account unlocked.
# (return status = 0)
# 1> quit
```

Next step installing the SRS software with a response file on the companion site and enter the following command as user **syb<sid>**:

```
# /sapcd/ase-16.0.03.06/BD_SYBASE_ASE_16.0.03.06_RDBMS_for_BS_/SYBASE_LINUX_X86_64/
setup.bin -f /sybase/HA7/srs-setup.txt -i silent
```

Activate HADR on companion node with a response file and enter the following command as user syb <sid>:

```
# setuphadr /sybase/HA7/HA7_companion_lin.rs.txt
```



Note

In certain circumstances the installation is not successful. It could help to set up the primary system again and install the companion afterward.



Note

If the system is reinstalled and the companion system reports **Missing read/write permissions** for this directory */tmp/.SQLAnywhere*, check the permission on both node. In case the ownership must be changed run the setup again on both nodes. Start with the **Primary**.

Creating a secure store key entry for the database:

```
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -function  
LiveDatabaseUpdate -dbname <SID> -dbtype syb -dbuser DR_admin -dbpass <Secure password>  
-updatemethod Execute -updateoption TASK=SET_USER_PASSWORD -updateoption USER=DR_ADMIN
```

9.1.5 Fault Manager Installation

The Fault Manager is configured on the ASCS host. The benefit from this setup is that the syb-dbfm service can be monitored and tracked with the existing pacemaker for the ASCS / ERS replication.

```
# su - <sid>adm  
# cd /usr/sap/<SID>/ASCS00/exe/  
# sybdbfm install
```

EXAMPLE 13: FAULT MANAGER INSTALLATION:

```
replication manager agent user DR_admin and password set in Secure Store.  
Keep existing values (yes/no)? (yes)  
SAPHostAgent connect user sapadm and password set in Secure Store.
```



```

Keep existing values (yes/no)? (yes)
Enter value for primary database host: (suse7cl1)
suse7db1
Enter value for primary database name: (HA7)
Enter value for primary database port: (4901)
Enter value for primary site name: (Site1)
Enter value for primary database heart beat port: (13777)
Enter value for standby database host: (suse7cl1)
suse7db2
Enter value for standby database name: (HA7)
Enter value for standby database port: (4901)
Enter value for standby site name : (Site2)
Enter value for standby database heart beat port: (13787)
Enter value for fault manager host: (s7as-service)
Enter value for heart beat to heart beat port: (13797)
Enter value for support for floating database ip: (no)
Enter value for use SAP ASE Cockpit if it is installed and running: (no)

```

Update the values as per your environment for the Primary DB & companion DB host name, SID & Site Name. Make sure to use the virtual host name for the ASCS host. When the Fault Manager is installed, profile for it will be created in the `/sapmnt/<SID>/profile` by the name `SYBHA.PFL` and will have the configuration details. Restart the ASCS Instance which will also start the Fault Manager that has been added to the start profile as below:



Note

In case of a reinstallation it might be better to overwrite the existing user name and password in the secure store for the `sapadm` and `DR_admin` if the old values are not 100% known.

```

# cat /sapmnt/<SID>/profile/<SID>_ASCS00_<virt. ASCS hostname>
....
#-----
# copy sybdbfm and dependent
#-----
_CP_SYBDBFM_ARG1 = list:$(DIR_CT_RUN)/instancedb.lst
Execute_06 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $_CP_SYBDBFM_ARG1)
_CP_SYBDBFM_ARG2 = list:$(DIR_GLOBAL)/syb/linuxx86_64/cpe_sybodbc.lst
_CP_SYBDBFM_ARG3 = source:$(DIR_GLOBAL)/syb/linuxx86_64/sybodbc
Execute_07 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $_CP_SYBDBFM_ARG2)
$_CP_SYBDBFM_ARG3)
#-----
# Start sybha

```

```
#-----
_SYBHAD = sybdbfm.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
_SYBHA_PF = $(DIR_PROFILE)/SYBHA.PFL
Execute_08 = local rm -f $_SYBHAD
Execute_09 = local ln -s -f $(DIR_EXECUTABLE)/sybdbfm$(FT_EXE) $_SYBHAD
Restart_Program_02 = local $_SYBHAD hadm pf=$_SYBHA_PF
#-----
....
```

The status of the FM can be checked as below. Navigate to the ASCS work directory and then run *sybdbfm.sap.<SID>_ASCS<instance number> status* :

As user ha1adm

```
# cd /usr/sap/<SID>/ASCS<instance number>/work
# ./sybdbfm.sap<SID>_ASCS<instance number> status

fault manager running, pid = 23234, fault manager overall status = OK, currently
executing in mode PAUSING
sanity check report (208).
node 1: server <DB server1>, site <site name one>.
db host status: OK.
db status OK hadr status PRIMARY.
node 2: server <DB server2>, site <site name two>.
db host status: OK.
db status OK hadr status STANDBY.
replication status: SYNC_OK.
```

Checking the log file is also a suitable measure to validate the status.

As user ha1adm

```
# cd /usr/sap/<SID>/ASCS<instance number>/work
# tail -f dev_sybdbfm
# ...

2019 09/11 15:34:30.523 (23234) ----- Log messages -----

2019 09/11 15:34:30.523 (23234) Info: saphostcontrol: Executing LiveDatabaseUpdate

2019 09/11 15:34:30.523 (23234) Info: saphostcontrol: LiveDatabaseUpdate successfully
executed

2019 09/11 15:34:30.524 (23234) call is running.
2019 09/11 15:34:30.534 (23234) call exited (exit code 0).
2019 09/11 15:34:30.534 (23234) db status is:
DB_OK.
2019 09/11 15:34:42.561 (23234) *** sanity check report (136)***.
```

```

2019 09/11 15:34:42.562 (23234) node 1: server <DB server1>, site <site name one>.
2019 09/11 15:34:42.562 (23234) db host status: OK.
2019 09/11 15:34:42.562 (23234) db status OK hadr status PRIMARY.
2019 09/11 15:34:42.562 (23234) node 2: server <DB server2>, site <site name two>.
2019 09/11 15:34:42.562 (23234) db host status: OK.
2019 09/11 15:34:42.562 (23234) db status OK hadr status STANDBY.
2019 09/11 15:34:42.562 (23234) replication status: SYNC_OK.
2019 09/11 15:34:57.688 (23234) *** sanity check report (137)***.
2019 09/11 15:34:57.688 (23234) node 1: server <DB server1>, site <site name one>.
2019 09/11 15:34:57.688 (23234) db host status: OK.
2019 09/11 15:34:57.688 (23234) db status OK hadr status PRIMARY.
2019 09/11 15:34:57.688 (23234) node 2: server <DB server2>, site <site name two>.
2019 09/11 15:34:57.688 (23234) db host status: OK.
2019 09/11 15:34:57.688 (23234) db status OK hadr status STANDBY.
2019 09/11 15:34:57.688 (23234) replication status: SYNC_OK.
2019 09/11 15:35:12.827 (23234) *** sanity check report (138)***.
2019 09/11 15:35:12.827 (23234) node 1: server <DB server1>, site <site name one>.
2019 09/11 15:35:12.827 (23234) db host status: OK.
2019 09/11 15:35:12.827 (23234) db status OK hadr status PRIMARY.
2019 09/11 15:35:12.827 (23234) node 2: server <DB server2>, site <site name two>.
2019 09/11 15:35:12.827 (23234) db host status: OK.
2019 09/11 15:35:12.827 (23234) db status OK hadr status STANDBY.
2019 09/11 15:35:12.827 (23234) replication status: SYNC_OK.
# ...

```

Below are a few parameters that need to be updated in the *SYBHA.PFL* to make the failover working.

```

ha/syb/support_cluster = 1
ha/syb/failover_if_unresponsive = 1
ha/syb/allow_restart_companion = 1
ha/syb/set_standby_available_after_failover = 1
ha/syb/chk_restart_repserver = 1
ha/syb/cluster_fmhost1 = Hostname for Node 1 of the ASCS HA Setup
ha/syb/cluster_fmhost2 = Hostname for Node 2 of the ASCS HA Setup
ha/syb/use_boot_file_always = 1
ha/syb/dbfmhost = virtual hostname of ASCS instance

```

Details of all the FM parameters can be found in the **SAP ASE HA DR User Guide**. Those highlighted in bold are of interest for the setup. Since the FM is installed with the ASCS which can failover from Node 1 to Node 2, the parameters *ha/syb/cluster_fmhost1* and *ha/syb/cluster_fmhost2* provide the physical host names of both nodes where the FM can potentially run.

In a scenario where the complete Availability Zone (AZ1) goes down and the ASCS and Primary database are running there, the DB failover is not triggered until the ASCS failover is complete and the FM is up and running in the 2nd Availability Zone (AZ2). The FM then needs to read the

boot file to get the prior state of the ASE DB. This is mandatory to ensure that FM can trigger the failover correctly. The parameter `ha/syb/use_boot_file_always=1` makes sure that the FM always reads from the boot file which is part of the work directory (the same for ASCS and FM) and failover along with FM.

9.2 Cluster Integration of Fault Manager

We implement the FM in the pacemaker environment as part of the ASCS instance.

EXAMPLE 14: FM IS INTEGRATED AS INCLUDED SERVICE ALONG WITH THE ASCS.

The cluster configuration for the *primitive rsc_sap_<SID>_ASCS <instance number>* have to be modified. In the example we use:

- `<SID> ⇒ HA1`
- `<instance number> ⇒ 00`
- `virtual hostname ⇒ sapha1as`

```
# crm configure edit rsc_sap_HA1_ASCS00
```

```
primitive rsc_sap_HA1_ASCS00 SAPIInstance \  
    operations $id=rsc_sap_HA1_ASCS00-operations \  
    op monitor interval=11 timeout=60 on-fail=restart \  
    params InstanceName=HA1_ASCS00_sapha1as START_PROFILE="/sapmnt/HA1/profile/  
HA1_ASCS00_sapha1as" \  
    AUTOMATIC_RECOVER=false MONITOR_SERVICES="sybdbfm|msg_server|enserver" \  
    meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1  
    priority=10
```

The FM service is not part of the default observed SAP instance services. If we specify the **MONITOR_SERVICES** all default settings are overwritten by the named services. That means we have to count all services which are shown as a result of the `sapcontrol -nr 00 -function GetProcessList` command. The example above is for an ENSA1 configuration.



Note

The cluster configuration is different for ENSA1 and ENS2 installation. The names for the **MONITOR_SERVICES** differ between this two versions.

9.3 Operating a Pacemaker-Controlled and FM-Monitored ASE Replication Setup

To implement an FM controlled ASE replication setup with pacemaker integration for FM some special rules need to be followed. First of all it needs to be detailed how to check the status of the replication and FM itself. This chapter will also give guidelines how to improve the takeover time and how to control such an environment.

EXAMPLE 15: CHECKING THE STATUS OF THE DATABASE SITUATION

Check the status and locate the primary site. *As user ha1adm on the ASCS host*

```
# cd /usr/sap/<SID>/ASCS<instance nr>/work
# ./sybdbfm.sap<SID>_ASCS<instance nr> status
```

And check the log file *dev_sybdbfm*

```
2019 09/13 09:58:52.200 (3290) *** sanity check report (2)***.
2019 09/13 09:58:52.200 (3290) node 1: server sapdb1, site LeonRot.
2019 09/13 09:58:52.200 (3290) db host status: OK.
2019 09/13 09:58:52.200 (3290) db status OK hadr status STANDBY.
2019 09/13 09:58:52.200 (3290) node 2: server sapdb2, site Orlando.
2019 09/13 09:58:52.201 (3290) db host status: OK.
2019 09/13 09:58:52.201 (3290) db status OK hadr status PRIMARY.
2019 09/13 09:58:52.201 (3290) replication status: SYNC_OK.
```

As user root on the database host

```
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname <SID> -
dbtype syb -function GetDatabaseSystemStatus
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname <SID> -
dbtype syb -function GetDatabaseStatus
# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname <SID>
-dbtype syb -function LiveDatabaseUpdate -updatemethod Check -updateoption
TASK=REPLICATION_STATUS
```

As user syb <sid> on the database host

```
# isql -UDR_admin -P <secure password> -S<db host>:4909 -X -w 1000
1> sap_status active_path
2> go
```

The application server (PAS and AAS) environment must be adapted for the DB failover situation (takeover). On each host which is providing a dialog server (PAS; AAS) the *.dbenv.sh* and or *.dbenv.csh* file needs to be extended.

EXAMPLE 16: MODIFY THE DB ENVIRONMENT SETTINGS ON THE DIALOG SERVER

Add the missing value and extend the settings as below shown on each host who runs a dialog application server.

As user ha1adm

```
# vi .dbenv.csh
...
setenv dbs_syb_server <server1:server2>
setenv dbs_syb_ha 1
...
```

As user ha1adm

```
# vi .dbenv.sh
...
dbs_syb_server=<server1:server2>
export dbs_syb_server
dbs_syb_ha=1
export dbs_syb_ha
...
```



Important

The instance must be restarted to activate the changes.

EXAMPLE 17: OS SETTINGS FOR FASTER REACTION TIME AFTER PRIMARY DB HOST IS DOWN

The default `tcp_retries` value is too high and causes a very long takeover time. With ASE16 PL7 the behavior is modified. Up to this patch the change below improves the takeover time.

As user root

```
# echo 3 >/proc/sys/net/ipv4/tcp_retries2
## makes the changes online
# vi /etc/sysctl.conf
...
net.ipv4.tcp_retries2 = 3
...
## makes the changes reboot persistent
```

EXAMPLE 18: STARTING AND STOPPING THE SAP SYSTEM AND DATABASES IN REPLICATION MODE

If Fault Manager is monitoring the Primary and Companion database and Fault Manager is monitored by Pacemaker there is a special procedure necessary to start and stop the system.

IN GENERAL THESE STEPS ARE IMPORTANT TO START THE SYSTEM:

- Start companion database + replication server
- Start primary database + replication server
- Change cluster maintenance mode to false
 - Start ASCS with FM (automatic)
 - Start ERS (automatic)
- Start PAS and AAS instances
- Optional: release cluster maintenance mode, if the SAP system was started manually
 - File system must be mounted and IP must be set manually
 - As user *<sid> adm* with *sapcontrol -nr <instance number> -function StartSystem*

As user *root* on companion database host

```
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID>_REP -dbtype syb
```

As user *root* on primary database host

```
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID> -dbtype syb
# /usr/sap/hostctrl/exe/saphostctrl -function StartDatabase -dbname <SID>_REP -dbtype syb
```

As user *root* on one of the Pacemaker host for ASCS and ERS

```
# crm configure property maintenance-mode=false
```

As user *<sid> adm* on the host for PAS or AAS

```
# sapcontrol -nr <instance number> -function StartSystem
```



Note

If the system should start one by one, use the command *sapcontrol -nr <instance number> -function StartSystem*. The direction must be: ASCS; ERS; PAS; AAS.

IN GENERAL THESE STEPS ARE IMPORTANT TO STOP THE SYSTEM:

- Set cluster maintenance mode to *true*
- Stop PAS and AAS instances
- Stop ASCS with FM
- Stop ERS
- Stop primary database + replication server
- Stop companion database + replication server

As user root

```
# crm configure property maintenance-mode=true  
# crm status
```

As user <sid> adm on one of the Pacemaker host for ASCS and ERS or PAS / AAS

```
# sapcontrol -nr <instance number> -function StopSystem
```



Note

If the system should stop one by one, use the command `sapcontrol -nr <instance number> -function Stop` on each instance host. The direction must be: AAS; PAS; ASCS; ERS.

As user root on primary database host

```
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID> -dbtype syb  
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID>_REP -dbtype syb
```

As user root on companion database host

```
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID> -dbtype syb  
# /usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname <SID>_REP -dbtype syb
```



Important

The Pacemaker-controlled server must be stopped in a proper way, too. Depending on the stonith method which is implemented different procedures are available.

As user root on one cluster node

```
# crm cluster run "crm cluster stop"
```

As user root on each node

```
# reboot  
## or  
# poweroff
```

9.3.1 Testing the Replication and FM Cluster Integration

Important for each high availability solution is an extensive testing procedure. That ensures that the solution is working as expected in case of a failure.

EXAMPLE 19: TRIGGERING A DATABASE FAILOVER AND MONITORING IF FM IS WORKING

Check the status and locate the primary site. As user *ha1adm* on the ASCS host

```
# cd /usr/sap/<SID>/ASCS<instance nr>/work  
# ./sybdbfm.sap<SID>_ASCS<instance nr> status
```

And check the log file *dev_sybdbfm*

```
2019 09/13 09:58:52.200 (3290) *** sanity check report (2)***.  
2019 09/13 09:58:52.200 (3290) node 1: server sapdb1, site LeonRot.  
2019 09/13 09:58:52.200 (3290) db host status: OK.  
2019 09/13 09:58:52.200 (3290) db status OK hadr status STANDBY.  
2019 09/13 09:58:52.200 (3290) node 2: server sapdb2, site Orlando.  
2019 09/13 09:58:52.201 (3290) db host status: OK.  
2019 09/13 09:58:52.201 (3290) db status OK hadr status  
PRIMARY.  
2019 09/13 09:58:52.201 (3290) replication status: SYNC_OK.
```

- Now destroy the primary database server.
- Monitor the takeover process with the FM on the ASCS host.

As user *ha1adm* on the ASCS host

```
# cd /usr/sap/<SID>/ASCS<instance nr>/work  
# tail -f dev_sybdbfm
```

EXAMPLE 20: SELECTED OUTPUT FROM THE TAKEOVER PROCESS.

```
...
2019 09/13 11:08:38.301 (3290) * sanity check report (270)* .
2019 09/13 11:08:38.301 (3290) node 1: server sapdb1, site LeonRot.
2019 09/13 11:08:38.301 (3290) db host status: OK.
2019 09/13 11:08:38.301 (3290) db status OK hadr status STANDBY.
2019 09/13 11:08:38.301 (3290) node 2: server sapdb2, site Orlando.
2019 09/13 11:08:38.301 (3290) db host status: OK.
2019 09/13 11:08:38.301 (3290) db status OK hadr status PRIMARY.
2019 09/13 11:08:38.301 (3290) replication status: SYNC_OK.
2019 09/13 11:08:50.416 (3290) ERROR in function SimpleFetch (1832)
(SQLExecDirect failed): (30046) [08S01] [SAP][ASE ODBC Driver]Connection to the
server has been lost. Unresponsive Connection was disconnected during command
timeout. Check the server to determine the status of any open transactions.
2019 09/13 11:08:50.416 (3290) ERROR in function SimpleFetch (1832)
(SQLExecDirect failed): (30149) [HYT00] [SAP][ASE ODBC Driver]The command has timed
out.
2019 09/13 11:08:50.416 (3290) execution of statement master..sp_hadr_admin
get_request, '1' failed.
2019 09/13 11:08:50.416 (3290) ERROR in function SimpleFetch (1824)
(SQLAllocStmt failed): (30102) [HY010] [SAP][ASE ODBC Driver]Function sequence
error
2019 09/13 11:08:50.416 (3290) execution of statement select top 1
convert( varchar(10), @@hadr_mode ) || ' ' || convert( varchar(10), @@hadr_state )
from sysobjects failed.
2019 09/13 11:08:50.416 (3290) disconnect connection
2019 09/13 11:09:22.505 (3290) ERROR in function SQLConnectWithRetry (1341)
(SQLConnectWithRetry failed): (30293) [HY000] [SAP][ASE ODBC Driver]The socket
failed to connect within the timeout specified.
2019 09/13 11:09:22.505 (3290) ERROR in function SQLConnectWithRetry (1341)
(SQLConnectWithRetry failed): (30012) [08001] [SAP][ASE ODBC Driver]Client unable
to establish a connection
2019 09/13 11:09:22.505 (3290) connected with warnings (555E69805100)
2019 09/13 11:09:22.505 (3290) ERROR in function SimpleFetch (1824)
(SQLAllocStmt failed): (30293) [HY000] [SAP][ASE ODBC Driver]The socket failed to
connect within the timeout specified.
2019 09/13 11:09:22.505 (3290) ERROR in function SimpleFetch (1824)
(SQLAllocStmt failed): (30012) [08001] [SAP][ASE ODBC Driver]Client unable to
establish a connection
2019 09/13 11:09:22.505 (3290) execution of statement select top 1
convert( varchar(10), @@hadr_mode ) || ' ' || convert( varchar(10), @@hadr_state )
from sysobjects failed.
2019 09/13 11:09:22.505 (3290) disconnect connection
2019 09/13 11:09:22.505 (3290) primary site unusable.
...
2019 09/13 11:09:22.984 (3290) primary site unusable.
```

```

2019 09/13 11:09:22.984 (3290) * sanity check report (271)* .
2019 09/13 11:09:22.984 (3290) node 1: server sapdb1, site LeonRot.
2019 09/13 11:09:22.984 (3290) db host status: OK.
2019 09/13 11:09:22.984 (3290) db status OK hadr status STANDBY.
2019 09/13 11:09:22.984 (3290) node 2: server sapdb2, site Orlando.
2019 09/13 11:09:22.984 (3290) db host status: UNUSABLE.
2019 09/13 11:09:22.984 (3290) db status DB INDOUBT hadr status UNREACHABLE.
2019 09/13 11:09:22.984 (3290) replication status: SYNC_OK.
2019 09/13 11:09:23.047 (3290) doAction: Primary database is declared dead or
unusable.
2019 09/13 11:09:23.047 (3290) disconnect connection
2019 09/13 11:09:23.047 (3290) database host cannot be reached.
2019 09/13 11:09:23.047 (3290) doAction: failover.
...
2019 09/13 11:11:55.497 (3290) * sanity check report (273)* .
2019 09/13 11:11:55.497 (3290) node 1: server sapdb1, site LeonRot.
2019 09/13 11:11:55.497 (3290) db host status: OK.
2019 09/13 11:11:55.497 (3290) db status OK hadr status PRIMARY.
2019 09/13 11:11:55.497 (3290) node 2: server sapdb2, site Orlando.
2019 09/13 11:11:55.497 (3290) db host status: UNUSABLE.
2019 09/13 11:11:55.498 (3290) db status DB INDOUBT hadr status UNREACHABLE.
2019 09/13 11:11:55.498 (3290) replication status: UNKNOWN.
2019 09/13 11:11:55.555 (3290) doAction: Standby database is declared dead or
unusable.
2019 09/13 11:11:55.555 (3290) disconnect connection
2019 09/13 11:11:55.555 (3290) doAction: Companion db host is declared unusable.
2019 09/13 11:11:55.555 (3290) doAction: no action defined.
2019 09/13 11:11:58.568 (3290) Error: NIECONN_REFUSED (No route to host),
NiRawConnect failed in plugin_fopen()
...
host is coming back online ##
2019 09/13 11:18:45.579 (3290) call is running.
2019 09/13 11:18:45.589 (3290) call exited (exit code 0).
2019 09/13 11:18:45.589 (3290) db status is: DB_OK.
2019 09/13 11:18:45.589 (3290) doAction: Standby database is declared dead or
unusable.
2019 09/13 11:18:45.589 (3290) disconnect connection
2019 09/13 11:18:45.589 (3290) doAction: Companion db host is declared ok.
2019 09/13 11:18:45.589 (3290) doAction: restart database.
2019 09/13 11:18:45.805 (3290) Webmethod returned successfully
...
2019 09/13 11:22:43.677 (3290) * sanity check report (286)* .
2019 09/13 11:22:43.677 (3290) node 1: server sapdb1, site LeonRot.
2019 09/13 11:22:43.677 (3290) db host status: OK.
2019 09/13 11:22:43.677 (3290) db status OK hadr status PRIMARY.
2019 09/13 11:22:43.677 (3290) node 2: server sapdb2, site Orlando.
2019 09/13 11:22:43.677 (3290) db host status: OK.

```

```

2019 09/13 11:22:43.677 (3290) db status OK hadr status STANDBY.
2019 09/13 11:22:43.677 (3290) replication status: SYNC_OK.
...

```

As user root

```

# /usr/sap/hostctrl/exe/saphostctrl -user sapadm <secure password> -dbname <SID>
-dbtype syb -function LiveDatabaseUpdate -updatemethod Check -updateoption
TASK=REPLICATION_STATUS

Webmethod returned successfully
Operation ID: 5254001F87CB1EE9B5C34755C99DDDFa

----- Response data -----
TASK_NAME=REPLICATION_STATUS
REPLICATION_STATUS=active
PRIMARY_SITE=<site1>
STANDBY_SITE=<site2>
REPLICATION_MODE=sync
ASE transaction log backlog (MB)=0
Replication queue backlog (MB)=0
TASK_STATUS=OK
----- Log messages -----
Info: saphostcontrol: Executing LiveDatabaseUpdate
Info: saphostcontrol: LiveDatabaseUpdate successfully executed

```

EXAMPLE 21: TRIGGERING AN FM FAILURE

Killing the FM process more than five times will bring pacemaker in action. Up to five times the saphostagent will take care of the SAP process. If this fail-count is reached in a specific time window, the service will not be restarted.

As user ha1adm

```

# pkill -9 sybdbfm
## check that the PID has changed
# sapcontrol -nr 00 -function GetProcessList
# pkill -9 sybdbfm
...
# sapcontrol -nr 00 -function GetProcessList
...
sybdbfm, , GRAY, Stopped, , , 11154
...

```

Now pacemaker will restart the FM instance locally first. *As user root*

```

# crm_mon -lrfn
...
Migration Summary:

```

```
* Node <hostname>:
rsc_sap_WAS_ASCS00: migration-threshold=3 fail-count=1 last-failure='Fri Sep 13
13:46:39 2019
...
```



Note

If the **fail-count** reaches the defined threshold, the ASCS is moved away from that host.

10 Appendix

10.1 CRM Configuration

The complete crm configuration for SAP system HA1

```
#
# nodes
#
node 1084753931: hacert01
node 1084753932: hacert02
#
# primitives for ASCS and ERS
#
primitive res_AWS_STONITH stonith:external/ec2 \
    op start interval=0 timeout=180 \
    op stop interval=0 timeout=180 \
    op monitor interval=300 timeout=60 \
    params tag=pacemaker profile=cluster
primitive rsc_fs_HA1_ASCS00 Filesystem \
    params device="efs-name:/ASC00" \
    directory="/usr/sap/HA1/ASC00" fstype=nfs4 \
    options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \
    op start timeout=60s interval=0 \
    op stop timeout=60s interval=0 \
    op monitor interval=200s timeout=40s
primitive rsc_fs_HA1_ERS10 Filesystem \
    params device="efs-name:/ERS10" \
    directory="/usr/sap/HA1/ERS10" fstype=nfs4 \
    options="rsize=1048576,wsiz=1048576,hard,timeo=600,retrans=2" \
    op start timeout=60s interval=0 \
    op stop timeout=60s interval=0 \
```

```

op monitor interval=20s timeout=40s
primitive rsc_ip_HA1_ASCS00 ocf:suse:aws-vpc-move-ip \
    params address=192.168.201.116 routing_table=rtb-table-name \
    interface=eth0 profile=cluster \
    op start interval=0 timeout=180 \
    op stop interval=0 timeout=180 \
    op monitor interval=120 timeout=60
primitive rsc_ip_HA1_ERS10 ocf:suse:aws-vpc-move-ip \
    params address=192.168.201.117 routing_table=rtb-table-name \
    interface=eth0 profile=cluster \
    op start interval=0 timeout=180 \
    op stop interval=0 timeout=180 \
    op monitor interval=120 timeout=60
primitive rsc_r53_HA1_ASCS00 aws-vpc-route53 \
    params hostedzoneid=hosted-zone-id ttl=10 \
    fullname=full-name profile=cluster \
    op start interval=0 timeout=180 \
    op stop interval=0 timeout=180 \
    op monitor interval=300 timeout=180
primitive rsc_sap_HA1_ASCS00 SAPIInstance \
    operations $id=rsc_sap_HA1_ASCS00-operations \
    op monitor interval=120 timeout=60 on-fail=restart \
    params InstanceName=HA1_ASCS00_saphalas \
        START_PROFILE="/sapmnt/HA1/profile/HA1_ASCS00_saphalas" \
        AUTOMATIC_RECOVER=false \
    meta resource-stickiness=5000 failure-timeout=60 migration-threshold=1 \
        priority=10
primitive rsc_sap_HA1_ERS10 SAPIInstance \
    operations $id=rsc_sap_HA1_ERS10-operations \
    op monitor interval=120 timeout=60 on-fail=restart \
    params InstanceName=HA1_ERS10_saphaler \
        START_PROFILE="/sapmnt/HA1/profile/HA1_ERS10_saphaler" \
        AUTOMATIC_RECOVER=false IS_ERS=true \
    meta priority=1000
#
# group definitions for ASCS and ERS
#
group grp_HA1_ASCS00 rsc_ip_HA1_ASCS00 \
    rsc_r53_HA1_ASCS00 rsc_fs_HA1_ASCS00 \
    rsc_sap_HA1_ASCS00 \
    meta resource-stickiness=3000
group grp_HA1_ERS10 rsc_ip_HA1_ERS10 \
    rsc_fs_HA1_ERS10 rsc_sap_HA1_ERS10

#
# constraints between ASCS and ERS
#

```

```

colocation col_sap_HA1_not_both -5000: grp_HA1_ERS10 grp_HA1_ASCS00
location loc_sap_HA1_failover_to_ers rsc_sap_HA1_ASCS00 \
  rule 2000: runs_ers_HA1 eq 1
order ord_sap_HA1_first_ascs Optional: rsc_sap_HA1_ASCS00:start rsc_sap_HA1_ERS10:stop
symmetrical=false
#
# crm properties and more
#
property cib-bootstrap-options: \
  have-watchdog=false \
  dc-version=1.1.15-21.1-e174ec8 \
  cluster-infrastructure=corosync \
  stonith-enabled=true \
  stonith-action=off \
  stonith-timeout=600s \
  last-lrm-refresh=1513844735
rsc_defaults rsc-options: \
  resource-stickiness=1 \
  migration-threshold=3
op_defaults op-options: \
  timeout=600 \
  record-pending=true

```

10.2 Checklist AWS Installation

Check your AWS configuration upfront and gather the following AWS items before you start the installation:

Checklist AWS installation	
Item	Status/Value
<i>SLES subscription and update status</i>	
All systems have a SLES for SAP subscription	
All systems have a public cloud channel	
All system have been updated to use the latest patch level	

Checklist AWS installation	
<i>AWS User Privileges for the installing person</i>	
Creation of EC2 instances and EBS volumes	
Creation security groups	
Creation EFS file systems	
Modification of AWS routing tables	
Creation policies and attach them to IAM roles	
Optional for Route53 agent installation	
Create and modify A-records in a private hosted zone	
Potentially needed :Creation of subnets and routing tables	
<i>VPC and Network</i>	
VPC Id	
CIDR range of VPC	
Subnet id A for systems in first AZ	
Subnet id B for systems in second AZ	
Routing table id for subnet A and B	
Is this routing table associated with both subnets?	

Checklist AWS installation	
Alternative: Is it associated to VPC? Subnets do not have their own ones	
<i>Optional: Route 53 configuration</i>	
Name of private hosted Route 53 zone	
Name of DHCP option set (Verify options!)	
Is option set associated to VPC?	
<i>AWS Policies Creation</i>	
Name of data provider policy	
Name of STONITH policy	
Name of Move IP (Overlay IP) policy	
Optionally: Name of Route53 policy	
<i>First cluster node (ASCS and ERS)</i>	
instance id	
ENI id	
IP address	
host name	
instance is associated to subnet A?	
instance has all 3 or 4 policies attached?	
EC2 tag <i>pacemaker</i> set with host name?	

Checklist AWS installation	
AWS CLI profile <i>cluster</i> created and set to <i>text</i> ?	
source/destination check disabled?	
<i>Second cluster node (ASCS and ERS)</i>	
instance id	
ENI id	
IP address	
host name	
instance is associated to subnet B?	
instance has all 3 or 4 policies attached?	
EC2 tag <i>pacemaker</i> set with host name?	
AWS CLI profile <i>cluster</i> created and set to <i>text</i> ?	
source/destination check disabled?	
<i>PAS system</i>	
IP address	
host name	
instance is associated to subnet A or B?	
instance has data provider policy attached?	

Checklist AWS installation	
<i>AAS system</i>	
IP address	
host name	
instance is associated to subnet A or B	
instance has data provider policy attached?	
<i>DB system (is potentially node 1 of a database failover cluster)</i>	
instance id	
ENI id	
IP address	
host name	
instance is associated to subnet A?	
instance has data provider policy attached? A cluster node has 2 to 3 more policies attached	
<i>Overlay IP address: service ASCS</i>	
IP address	
Has it been added to routing table?	
Does it point to the ENI of first node?	
<i>Overlay IP address: service ERS</i>	
IP address	
Has it been added to routing table?	

Checklist AWS installation	
Does it point to the ENI of the second node?	
<i>Optional: Overlay IP address DB server</i>	
IP address	
Has it been added to routing table?	
Does it point to the ENI of the DB server?	
<i>Optional: Route 53 configuration</i>	
The Route 53 private hosted zone has an A record with the name of the ASCS system the IP address of the first cluster node	
<i>Creation of EFS file system</i>	
DNS name of EFS file system	
<i>Internet access</i>	
All instances have Internet access ? Check routing tables	
Alternative: Add HTTP proxies for data providers and cluster software	

10.3 Related SAP Notes

- 953653 - Rolling Kernel Switch (<https://launchpad.support.sap.com/#/notes/953653/E>)
- 1153713 - Problems with SAP Management Console (Java) (<https://launchpad.support.sap.com/#/notes/1153713/E>)

- 1588667 - SAP on AWS: Overview of related SAP Notes and Web-Links (<https://launchpad.support.sap.com/#/notes/1588667/E>)
- 1656099 - SAP Applications on AWS: Supported DB/OS and AWS EC2 products (<https://launchpad.support.sap.com/#/notes/1656099/E>)
- 1656250 - SAP on AWS: Support prerequisites (<https://launchpad.support.sap.com/#/notes/1656250/E>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP Applications (<https://launchpad.support.sap.com/#/notes/1763512/E>)
- 1984787 - SUSE LINUX Enterprise Server 12: Installation notes (<https://launchpad.support.sap.com/#/notes/1984787/E>)
- 2077934 - Rolling kernel switch in HA environments (<https://launchpad.support.sap.com/#/notes/2077934/E>)
- 2235581 - SAP HANA: Supported Operating Systems (<https://launchpad.support.sap.com/#/notes/2235581/E>)
- 2254173 - Linux: Rolling Kernel Switch in Pacemaker based NetWeaver HA environments (<https://launchpad.support.sap.com/#/notes/2254173/E>)
- 2309342 - SUSE Linux Enterprise High Availability Extension on AWS for SAP HANA (<https://launchpad.support.sap.com/#/notes/2309342/E>)
- 2369910 - SAP Software on Linux: General information (<https://launchpad.support.sap.com/#/notes/2369910/E>)
- 2777438 - SYB: Database Fault Management (<https://launchpad.support.sap.com/#/notes/2777438/E>)

11 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.

12 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.