# Databases for SAP applications on Amazon

# Databases for SAP applications on Amazon: ***Unable to locate subtitle***

Services or capabilities described in Amazon Web Services documentation might vary by Region. To see the differences applicable to the China Regions, see [Getting Started with Amazon Web Services in China](#) [(PDF)](#).

# Table of Contents

# Databases for SAP applications on Amazon

This section covers the following guides.

- [SAP on Amazon – IBM Db2 HADR with Pacemaker](#)
- [Databases for SAP on Amazon with Amazon FSx for NetApp ONTAP](#)
- [SAP ASE for SAP NetWeaver on Amazon: high availability configuration for SUSE Linux Enterprise Server (SLES) for SAP applications](#)
- [SAP ASE for SAP NetWeaver on Amazon: high availability configuration for Red Hat Enterprise Linux (RHEL) for SAP applications](#)

**Additional SAP on Amazon documentation**

- [General SAP guides](#)
- [SAP HANA on Amazon](#)
- [SAP NetWeaver on Amazon](#)
- [Amazon Launch Wizard for SAP](#)
- [Amazon Systems Manager for SAP](#)
- [Amazon SDK for SAP ABAP](#)
- [SAP BusinessObjects on Amazon](#)
- [Amazon Migration Hub Orchestrator](#)

# SAP on Amazon – IBM Db2 HADR with Pacemaker

SAP on Amazon – IBM Db2 HADR with Pacemaker

Deployment and Operations Guide

Last updated: March 2022

## About This Guide

This guide provides instructions on how to set up Amazon Web Services resources to deploy IBM Db2 High Availability Disaster Recovery (HADR) with Pacemaker for SAP NetWeaver on Amazon Elastic Compute Cloud (Amazon EC2) instances. This guide is for users who are responsible for planning, architecting and deploying IBM Db2 on Amazon for SAP NetWeaver-based applications.

## Overview

Instructions in this document are based on recommendations provided by SAP and IBM on Db2 deployment on Linux via the SAP notes and KB articles listed in Table 1.

> ⓘ **Note**
>
> When deploying IBM Db2 version 11.5 Mod Pack 6 (11.5.6) or higher, refer to the option recommended by IBM. For more information, see Integrated solution using Pacemaker.

*Table 1 - SAP NetWeaver on IBM Db2 OSS Notes*

| SAP OSS Note | Description |
| --- | --- |
| **1656099** | SAP Applications on Amazon: Supported DB/ OS and Amazon EC2 products |
| **1656250** | SAP on Amazon: Supported instance types |
| **1612105** | DB6: FAQ on Db2 High Availability Disaster Recovery (HADR) |

| SAP OSS Note | Description |
|---|---|
| 101809 | DB6: Supported Db2 Versions and Fix Pack Levels |
| 1168456 | SAP Db2 support info |
| 1600156 | SAP Db2 support on Amazon |

**What this guide doesn't do**

This document doesn't provide guidance on how to set up network and security constructs like Amazon Virtual Private Cloud (Amazon VPC), subnets, route tables, access control lists (ACLs), Network Address Translation (NAT) Gateway, Amazon Identity and Access Management (IAM) Roles, or Amazon Security Groups. It doesn't cover the high availability (HA) setup for the SAP Application Server Central Services/Enqueue Replication Server (ASCS/ERS), and focuses only on the database (DB) layer when covering the single points of failure (SPOF) for the SAP applications.

# Considerations

## Specialized Knowledge

To understand this document, you should have a good understanding of Amazon services, general networking concepts, Linux operating systems, and IBM Db2 administration.

Before you follow the instructions in this guide, we recommend that you become familiar with the following Amazon services. (If you are new to Amazon, see Getting Started with Amazon

- Amazon EC2
- Amazon EBS
- Amazon VPC
- Amazon CloudFormation
- Amazon Systems Manager
- Amazon Simple Storage Service
- Amazon Identity and Access Management (IAM)

# Technical Requirements

- Before you start the installation and configuration of IAM Db2 High Availability Disaster Recovery (HADR), ensure that you meet the following requirements:

- Your operating system is a supported Red Hat or SUSE version. Check SAP [product availability matrix](#) (PAM). Login required.

- Your database version is IBM Db2 10.5 or higher.

- Bring your own license (BYOL) for IBM Db2 and SAP application.

- Install [Amazon SAP Data Provider](#) on Amazon EC2 instances after installing IBM Db2 database.

- An Amazon account with permission to create resources.

- Access to SAP installation media for database and application.

- Amazon Business or Enterprise level support ([1656250 - SAP on Amazon: Support prerequisites](#)). Login required.

# Planning

## Architecture Options

SAP NetWeaver applications based on IBM Db2 can be installed in three different ways:

- **Standard system or single host installation**-- In this option, Advanced Business Application Programming (ABAP) Application Server Central Services/System Central Services (ASCS/SCS) and the database primary application server (PAS) of SAP NetWeaver run in a single Amazon EC2 instance. This option is suited for non-critical and non-production workloads.

- **Distributed system**-- In distributed systems, ASCS/SCS and the database PAS of SAP NetWeaver can run on separate Amazon EC2 instances. For example, you can choose to run ASCS and PAS on one Amazon EC2 instance, and the database on another Amazon EC2 instance, or other combinations. This option is suited for production and non-production workloads.

- **High availability system**-- For your SAP application to be highly available, you will need to protect the single point of failures. The database is one of the single points of failure in SAP applications.

Amazon recommends that you deploy primary and standby IBM Db2 databases in different Availability Zones (AZs) within an Amazon region. Figure 1 provides a high-level architecture for IBM Db2 high availability in Amazon. This option is suited for business-critical applications.
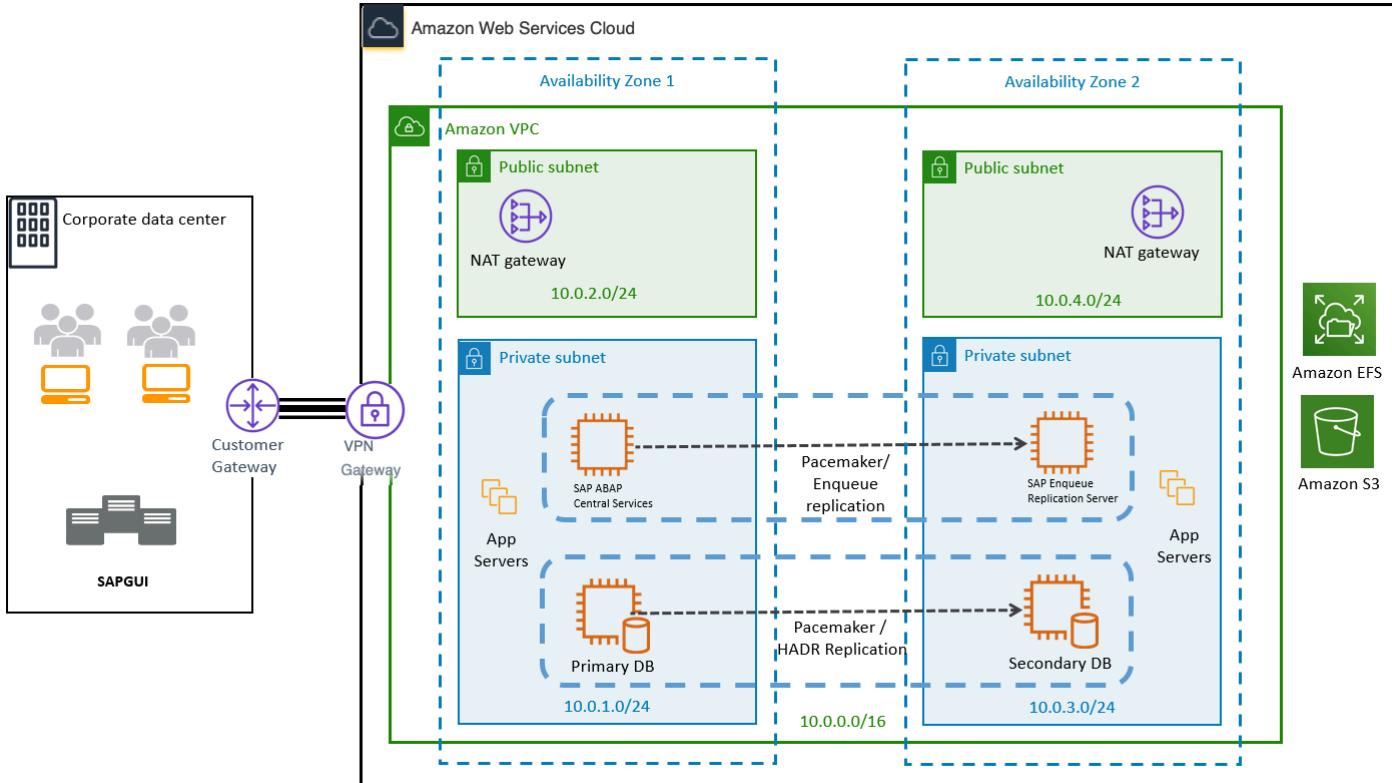


*Figure 1 – High-level HA architecture for SAP with IBM Db2 on Amazon*

# Security

Amazon provides security capabilities and services to securely run your SAP applications on the Amazon platform. In the context of IBM Db2 for SAP applications, you can use network services and features such as Amazon VPC, Amazon Virtual Private Network (Amazon VPN), Amazon Direct Connect, Amazon EC2 Security Groups, network access controls lists (NACLs), route tables, and more to restrict the access to your database.

## Network Security

The databases of SAP applications don't usually require direct user access. The end users access the application using SAP Graphical User Interface (GUI), SAP Web Dispatcher, or SAP Fiori. We recommend that you limit direct access to the EC2 instances to administrators only, for maintenance purpose.

IBM Db2 listens on TCP port 5912 by default. Depending on your VPC design, you should configure Amazon EC2 Security Groups, Network Access Control List (NaCls), and route tables to allow traffic to TCP Port 5912 from SAP primary application servers and additional application servers (PAS/AAS) and ABAP SAP Central Services/SAP Central Services (ASCS/SCS). To learn more about configuring the security group, see Security groups for your VPC.

## Encryption

Encryption is a security mechanism that converts plain text (readable data) into ciphertext. Amazon offers built-in encryption for Amazon EBS data volumes, boot volumes, and snapshots. The encryption process occurs automatically, and you don't need to manage encryption keys. This mechanism protects your EBS volumes at rest, and data in transit that passes between EC2 servers. This encryption level is offered at no additional cost.

You also can use the native IBM Db2 native database encryption feature if required.

## Sizing

SAP Quick Sizer is used to size SAP environment for new implementations. However, if you are migrating your existing SAP applications based on IBM Db2 to Amazon, consider using the following tools to right-size your SAP environment based on current utilization.

- **SAP Early Watch Alerts (EWA)**:--SAP EWA reports are provided by SAP regularly. These reports provide an overview of historical system utilization. Analyze these reports to see if your existing SAP system is over-utilized or under-utilized. Use this information to right-size your environment.
- **Linux native tools**:--Gather and analyze historical utilization data for CPU/Memory to right-size your environment. In case your source is IBM AIX, you can make use of nmon reports as well.
- **Amazon Services**-- Use services such as Amazon Migration Evaluator or Amazon Application Discovery Services that help with collecting usage and configuration data about your on-premises servers. Use this information to analyze and right-size your environment.

Because it's easy to scale up or scale down your Amazon EC2 instances on Amazon, consider the following while sizing your SAP environment on Amazon.

- You don't need to over-provision storage to meet future demand.
- SAP Quick Sizer tools provide sizing guidance based on assumptions that on 100% load (as per your inputs to tool), system utilization will not be more than 65%, so there is some buffer

built into SAP Quick Sizer recommendation. See SAP's [Quick Sizer guidance](#) for details. (Login required.)

# Operating System

You can deploy your SAP workload on SUSE Linux Enterprise Server (SLES) for SAP, Red Hat Enterprise Linux for SAP with High Availability and Update Services (RHEL for SAP with HA and US), or RHEL for SAP Solutions.

SLES for SAP and RHEL for SAP with HA and US are available in the [Amazon Marketplace](#) under an hourly or an annual subscription model.

## SLES

SLES for SAP provides additional benefits, including Extended Service Pack Overlap Support (ESPOS), configuration and tuning packages for SAP applications, and High Availability Extensions (HAE). See the [SUSE SLES for SAP product page](#). We strongly recommend using SLES for SAP instead of SLES for all your SAP workloads.

If you plan to use Bring Your Own Subscription (BYOS) images provided by SUSE, ensure that you have the registration code required to register your instance with SUSE to access repositories for software updates.

## RHEL

RHEL for SAP with HA and US provides access to Red Hat Pacemaker cluster software for High Availability, extended update support, and the libraries that are required to configure pacemaker HA. For details, see the [RHEL for SAP Offerings on Amazon FAQ](#) in the *Red Hat knowledgebase*.

If you plan to use the BYOS model with RHEL, either through the [Red Hat Cloud Access program](#) or another means, ensure that you have access to a RHEL for SAP Solutions subscription. For details, see [Overview of the Red Hat Enterprise Linux for SAP Solutions subscription](#) in the *Red Hat knowledgebase*.

The correct subscription is required to download the required packages for configuring the Pacemaker cluster.

## Compute

Amazon provides a wide array of SAP supported Amazon EC2 instances for your SAP workloads. See SAP Note 1656099 - SAP Applications on Amazon: Supported DB/OS and Amazon EC2 products for details. Based on the results of your sizing exercise, you can deploy your IBM Db2 on any of the SAP supported Amazon EC2 instances that meets your requirement.

## Storage

Amazon EBS volumes are designed to be highly available and durable. EBS volume data is replicated across multiple servers in an AZ to prevent the loss of data from the failure of any single component. Due to this built in protection, you don't have to configure RAID 1 for volumes containing database transaction log files and Db2 binaries.

We don't recommend RAID 5 for container files for data, index, or temporary tablespaces on Amazon for the following reasons:

- As mentioned previously, volumes are replicated within AZ by default.
- Parity write operations of RAID 5 consume some of the Input/Output Operations Per Second (IOPS) available to your volume and will reduce the overall Input/Output (IO) available for database operations by about 20-30% over RAID 0 configuration.

# Network

Ensure that you have your network constructs set up to deploy resources related to SAP NetWeaver. If you haven't already set up network components like Amazon VPC, subnets, and route tables, you can use Amazon Quick Start for VPC to easily deploy scalable VPC architecture. See the Amazon Quick Start for VPC reference deployment guide.

See the series of VPC Subnet Zoning Pattern blogs for VPC patterns that you should consider for SAP applications.

# Business Continuity

We recommend that you architect your business-critical applications to be fault tolerant. Depending on your availability requirements, there are different ways to achieve this. In this section we will discuss how you can set up highly available IBM Db2 for SAP applications.
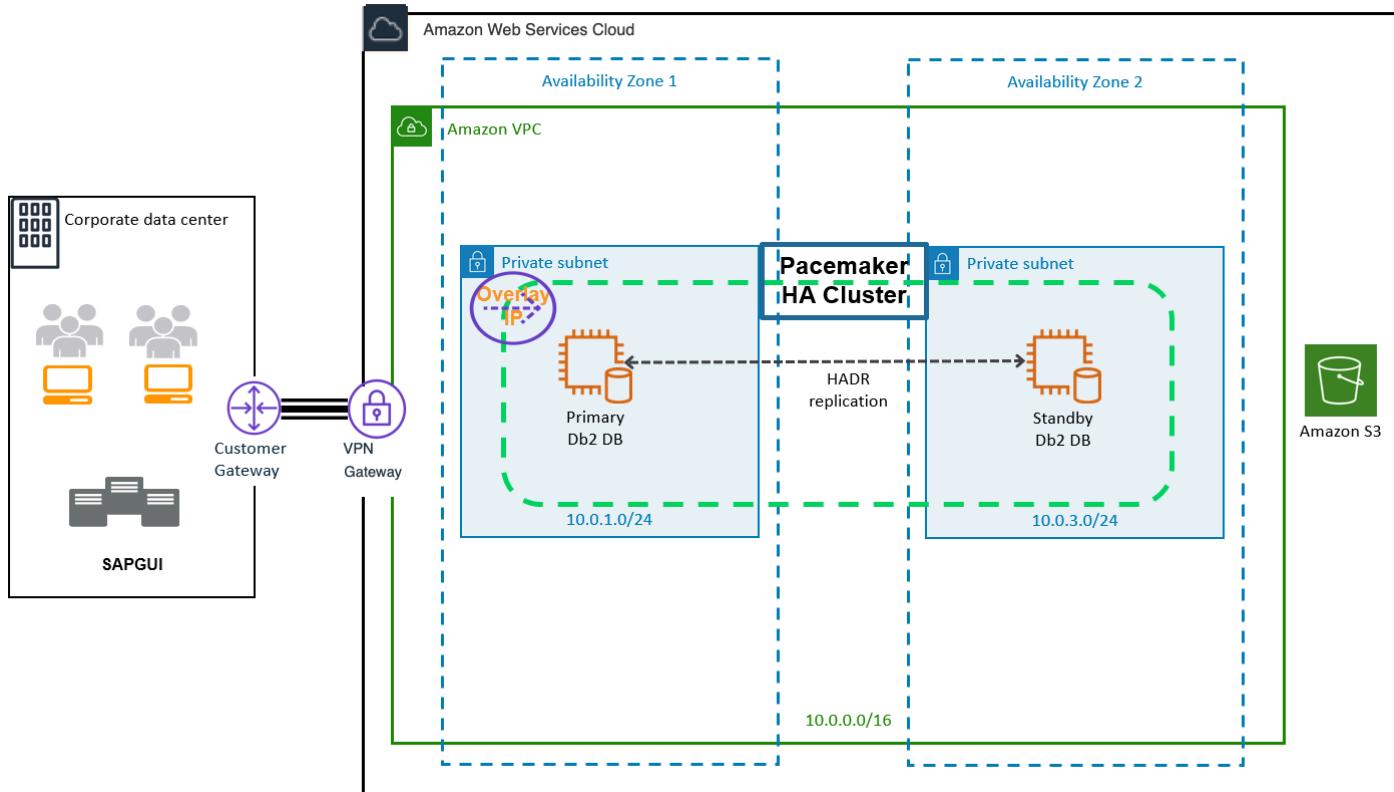
# High Availability

High availability for IBM Db2 database on Amazon can be configured with IBM HADR with Pacemaker:

One of the requirements for automated failover with IBM Db2 HADR on Amazon is Pacemaker. Implementing a Pacemaker cluster in Amazon is similar to deploying it in an on-premises setting. On Amazon, you need to deploy the cluster nodes in separate subnets, and we recommend that you have these subnets in different AZs.

Figure 2 provides an overview of architecture for IBM Db2 HADR with Pacemaker on Amazon. This includes the following components:

- A VPC configured with two private subnets across two AZs. This provides the network infrastructure for your IBM Db2 deployment.

- In private subnet, Linux servers are configured with Pacemaker to protect the IBM Db2 database.

- Overlay IP address (similar to a virtual IP address) that is relocatable between the primary and standby Db2 databases.

*Figure 2 - IBM Db2 HADR with Pacemaker*

# Deployment

This section discusses high level deployment process and steps. Table 2 lists the steps in the order they should be done, and each step's purpose.

*Table 2 – Steps to set up Amazon resources to deploy IBM Db2 HADR with Pacemaker for SAP NetWeaver on Amazon EC2 instances*

| Activity | Purpose |
|---|---|
| Step 1: Db2 Virtual Hostname | Decide on the virtual hostname for your Db2 database (for example, dbhadb2). |
| Step 2: Amazon Overlay IP | Decide on the Overlay IP for the dbhadb2 name (for example, `192.168.1.90` ). |
| Step 3: Provision Amazon Resources | Provision Amazon resources and configure security. |
| Step 4: SAP App and Db2 Install | Install SAP tier and Db2 primary and standby databases. |
| Step 5: Db2 HADR Setup | Configure Db2 HADR replication. |
| Step 6: Pacemaker Cluster Setup | Configure the Pacemaker cluster. |
| Step 7: Post Setup Configuration | Post tasks and manual configuration. |
| Step 8: Testing and Validation | Quality Assurance. |

## Step 1: Db2 Virtual Hostname

Decide on the virtual hostname for your Db2 database. For example, if your virtual hostname is dbhadb2, it would be configured in the SAP and `dbclient` profiles. See Step 7: Post Setup Configuration in this document for more information.

## Step 2: Amazon Overlay IP

Decide on IP address to use for your Overlay IP. An Overlay IP address is an Amazon-specific routing entry which can send network traffic to an instance, no matter which AZ the instance is located in.

One key requirement for the Overlay IP is that it should not be used elsewhere in your VPC or on-premises. It should be part of the private IP address range defined in [RFC1918](#). For example, if your VPC is configured in the range of `10.0.0.0/8` or `172.16.0.0/12`, you can use the Overlay IP from the range of `192.168.0.0/16`. Based on the number of HA setups you plan to have in your landscape, you can reserve the IP address by reserving a block from the private IP address to ensure there is no overlap.

Amazon worked on creating a resource agent, `aws-vpc-move-ip`, which is available along with the Linux Pacemaker. This agent updates the route table of the VPC where you have configured the cluster to always point to the primary DB.

All traffic within the VPC can reach the Overlay IP address via the route table. Traffic from outside the VPC, whether that is from another VPC or on-premises will require Amazon Transit Gateway (TGW) or Amazon NLB to reach the Overlay IP address. For more information on how to direct traffic to an Overlay IP address via Amazon TGW or Amazon NLB, see [SAP on Amazon High Availability with Overlay IP Address Routing](#).

## Step 3: Amazon Resources

Deciding the right storage layout is important to ensure you can meet required IO. EBS gp2 volumes balance price and performance for a wide variety of workloads, while io1 volumes provide the highest performance consistently for mission-critical applications. With these two options, you can choose a storage solution that meets your performance and cost requirements. For more information, see [Amazon EBS features](#) for more information.

## Step 4: SAP Netweaver and IBM Db2 Deployment

See the [SAP Standard Installation guide](#) based on your installation release to get the technical steps and prerequisites for SAP installation.

### Step 4a: Create EC2 Instances for Deploying SAP NetWeaver ASCS

See [SAP NetWeaver Environment Setup for Linux on Amazon](#) to learn how to set up an EC2 instance for SAP NetWeaver.

## Step 4b: Create EC2 Instances for IBM Db2 Primary and Standby Databases

Deploy two EC2 instances, one in each AZ, for your primary and standby databases.

## Step 4c: Disable Source/destination Check for the EC2 Instance Hosting the IBM Db2 Primary and Standby Databases

You need to disable source/destination check for your EC2 instance hosting primary and standby databases. This is required to route traffic via Overlay IP. See Changing the source or destination checking to learn more about how to disable source/destination check for your EC2 instance. You can use the following command line interface (CLI) to achieve this.

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

## Step 4d: Amazon IAM Role

For the Pacemaker setup, create two policies and attach them to the IAM role, which is attached to the Db2 primary and standby instance. This allows your EC2 instance to call the APIs which run during the failover process by Pacemaker.

- STONITH – Allows Ec2 instance to start, stop and reboot instances.

```
{
    "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:us-east-1:123456789012:instance/i-node1",
            "arn:aws:ec2:us-east-1:123456789012:instance/i-node2"
        ]
    }
  ]
}
```

- Overlay IP – Allows the Ec2 instance to update the route table in case of failover.

```
    {
  "Version":"2012-10-17",
  "Statement": [
      {
          "Sid": "VisualEditor0",
          "Effect": "Allow",
          "Action": "ec2:ReplaceRoute",
          "Resource": "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-XYZ"
      },
      {
          "Sid": "VisualEditor1",
          "Effect": "Allow",
          "Action": "ec2:DescribeRouteTables",
          "Resource": "*"
      }
  ]
}
```

Replace the following variables with the appropriate names:

- `region-name`: the name of the Amazon region.
- `account-id`: The name of the Amazon account in which the policy is used.
- `rtb-XYZ`: The identifier of the routing table which needs to be updated.
- `i-node1`: Instance ID for the Db2 primary instance.
- `i-node2`: Instance id for the Db2 standby instance.

## Step 4e: SAP Application and Db2 Software Install

**Prerequisites**:

- Before starting the installation, update the `/etc/hosts` files of database, ASCS, and application servers with the hostname and IP address of *your* database, ASCS and application servers. This ensures that all your instances can resolve each other's address during installation and configuration.

- You need to install the following packages in your instances: `tcsh.x86_64`, `ksh.x86_64`, `libaio.x86_64`, `libstdc++.x86_64`.

- Comment out the 5912 port entry in the `/etc/services` file (if it exists), as this port is used for the Db2 installation:

```
    #fis              5912/tcp           # Flight Information Services
    #fis              5912/udp           # Flight Information Services
    #fis              5912/sctp          # Flight Information Services
```

**SAP application and Db2 software installation (high-level instructions)**:

1. Install SAP ASCS using software provisioning manager (SWPM) on the Amazon EC2 instance. Choose the installation option depending on the scenario; for example, distributed or HA in case you plan to install ERS for app layer high availability.
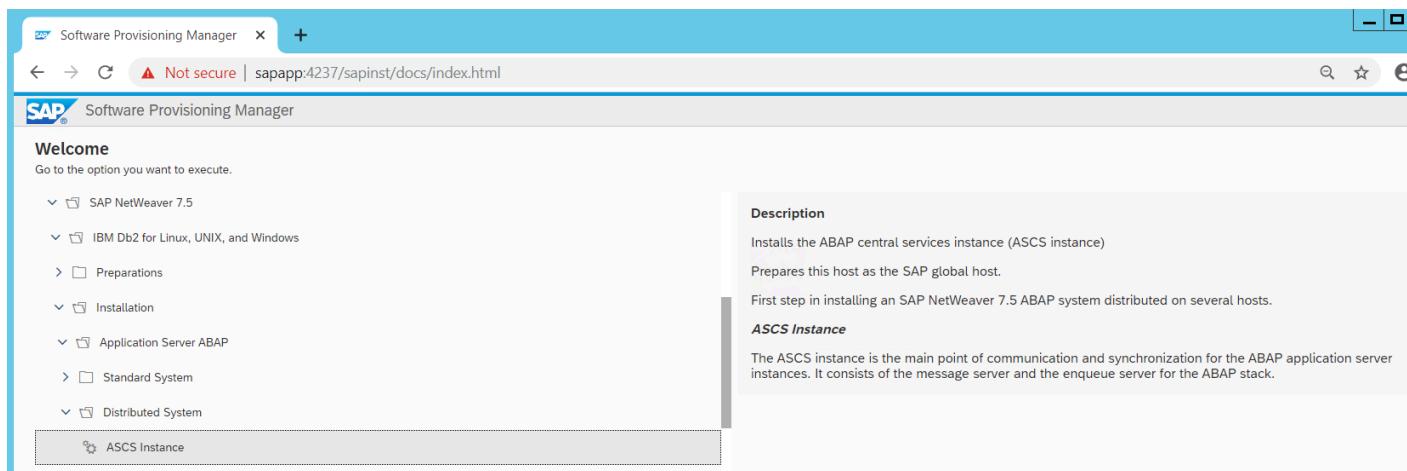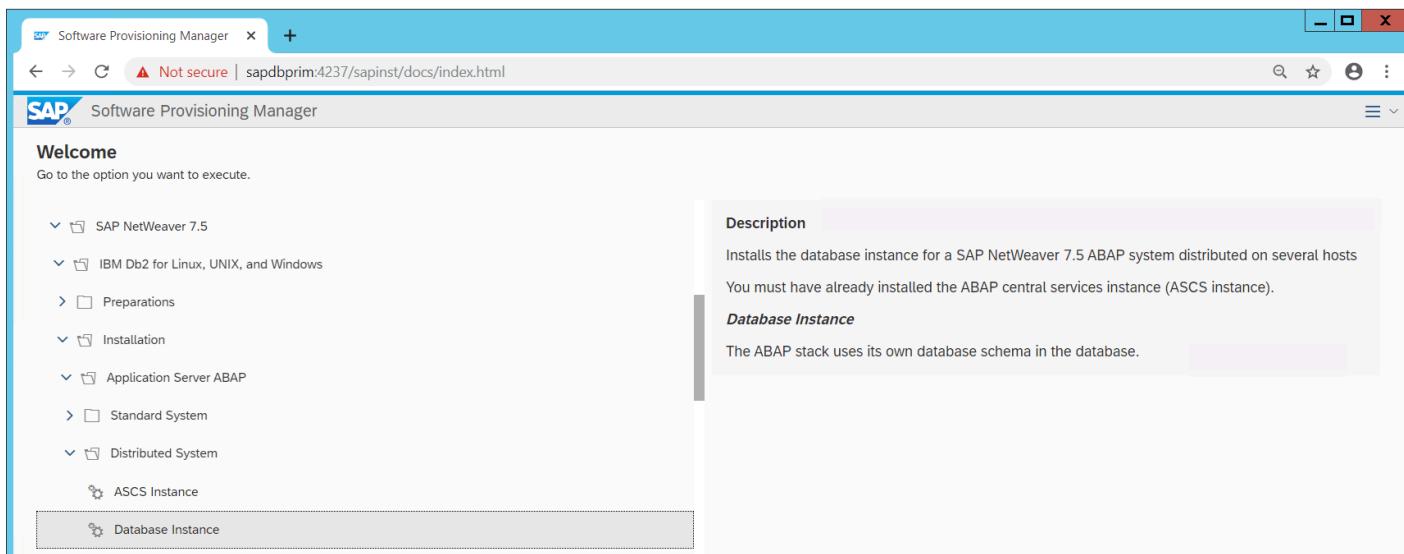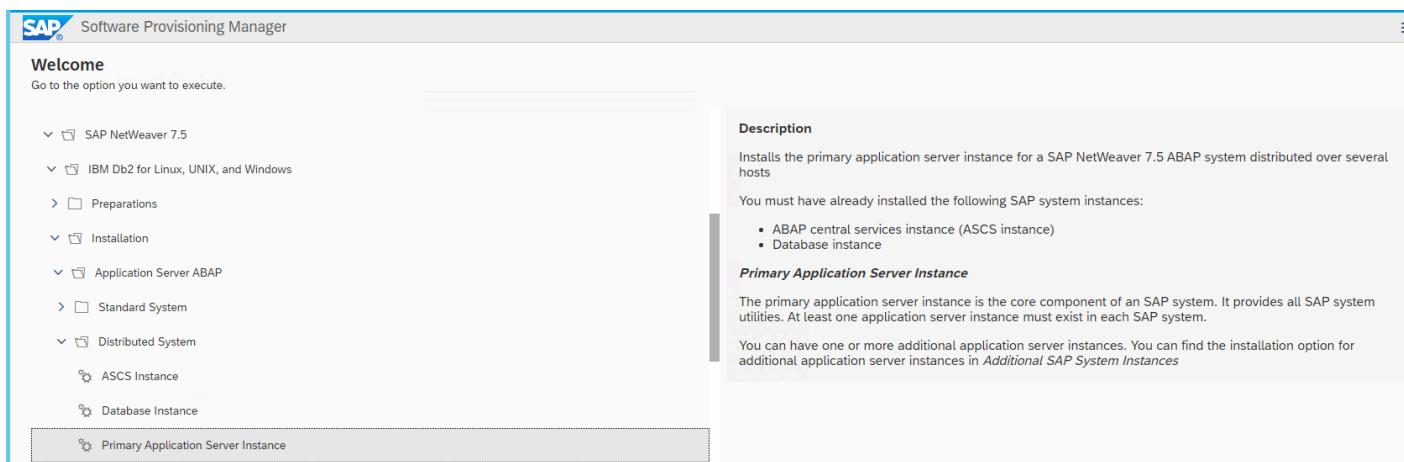


*Figure 3 – Install ASCS*

2. Install the primary database using SWPM on the Amazon EC2 instance hosted in AZ1.

*Figure 4 – Install the primary database*

3. Take a backup of the primary database.

4. Install the PAS instance. This can be the same EC2 instance used in step 1 if you want to install ASCS and PAS on one host.



*Figure 5 – Install the PAS instance*

5. For the standby DB installation:

a. Use the SAP homogeneous system copy procedure from SWPM with the option of **System copy** > **Target systems** > **Distributed** > **Database instance**.

b. In the SWPM parameter screen, when asked for system load type, choose Homogenous System and the backup/restore option.

c. When prompted by the SWPM, restore the backup you took during step 3 on the standby DB. You can exit the installer, because the subsequent installation is already completed on the primary database server.

You should now have your ASCS, primary Db2 database, and PAS (if running on a different host than ASCS) installed and running in AZ1 of your setup. In AZ2, you should have the standby Db2 database installed and running. You can optionally install an additional application server if required to support the workload. We recommend that you distribute your application servers in both AZs.

## Step 5: Db2 HADR Setup

The following steps explain how to set up HADR between the primary and standby databases. For additional references, see:

- IBM Db2 HADR documentation
- IBM Support Page

Table 3 details the steps for setup. Update the configuration commands according to your environment.

*Table 3 – Db2 HADR setup*

| System ID (SID) | STJ |
|---|---|
| Primary Db2 database hostname | dbprim00 |
| Standby Db2 database hostname | dbsec00 |
| Overlay IP | 192.168.1.81 |

**To set up Db2 HADR**:

1. Find the state of the primary database before HADR configuration by executing the following command:

```
> db2 get db cfg for STJ | grep HADR
```

```
        HADR database role                                   = STANDARD
        HADR local host name                 (HADR_LOCAL_HOST) =
        HADR local service name               (HADR_LOCAL_SVC) =
        HADR remote host name                (HADR_REMOTE_HOST) =
        HADR remote service name              (HADR_REMOTE_SVC) =
        HADR instance name of remote server  (HADR_REMOTE_INST) =
        HADR timeout value                       (HADR_TIMEOUT) = 120
        HADR target list                     (HADR_TARGET_LIST) =
        HADR log write synchronization mode     (HADR_SYNCMODE) = NEARSYNC
        HADR spool log data limit (4KB)      (HADR_SPOOL_LIMIT) = AUTOMATIC(0)
        HADR log replay delay (seconds)     (HADR_REPLAY_DELAY) = 0
        HADR peer window duration (seconds)  (HADR_PEER_WINDOW) = 0
        HADR SSL certificate label             (HADR_SSL_LABEL) =
```

2. The `HADR_LOCAL_SVC` and `HADR_REMOTE_SVC` parameters require an entry in your `/etc/services` file. If the entry is unavailable, update the `/etc/services` file to include the entry. Here is a sample `/etc/services` file entry. The SID is STJ.

```
    #  grep -i hadr /etc/services
    STJ_HADR_1   5950/tcp   # DB2 HADR log shipping
    STJ_HADR_2   5951/tcp   # DB2 HADR log shipping
```

3. Complete the following steps in your primary database (in this case, dbprim00) as the Db2 instance owner (in this case, `db2stj`):

```
    db2 update db cfg for STJ using HADR_LOCAL_HOST dbprim00
    db2 update db cfg for STJ using HADR_LOCAL_SVC STJ_HADR_1
    db2 update db cfg for STJ  using HADR_REMOTE_HOST dbsec00
    db2 update db cfg for STJ  using HADR_REMOTE_SVC STJ_HADR_2
    db2 update db cfg for STJ  using HADR_REMOTE_INST db2stj
    db2 update db cfg for STJ  using HADR_TIMEOUT 120
    db2 update db cfg for STJ  using HADR_SYNCMODE NEARSYNC
    db2 update db cfg for STJ  using HADR_PEER_WINDOW 300
    db2 update db cfg for STJ  using LOGINDEXBUILD ON
```

Here is the state after the configuration was updated:

```
    > db2 get db cfg for STJ | grep HADR
    HADR database role                                   = STANDARD
    HADR local host name                 (HADR_LOCAL_HOST) = dbprim00
    HADR local service name               (HADR_LOCAL_SVC) = STJ_HADR_1
    HADR remote host name                (HADR_REMOTE_HOST) = dbsec00
```

```
        HADR remote service name              (HADR_REMOTE_SVC) = STJ_HADR_2
        HADR instance name of remote server  (HADR_REMOTE_INST) = db2stj
        HADR timeout value                       (HADR_TIMEOUT) = 120
        HADR target list                     (HADR_TARGET_LIST) =
        HADR log write synchronization mode     (HADR_SYNCMODE) = NEARSYNC
        HADR spool log data limit (4KB)       (HADR_SPOOL_LIMIT) = AUTOMATIC(0)
        HADR log replay delay (seconds)      (HADR_REPLAY_DELAY) = 0
        HADR peer window duration (seconds)  (HADR_PEER_WINDOW) = 300
        HADR SSL certificate label               (HADR_SSL_LABEL) =
```

4. Run the following steps in your standby database (in this case `dbsec00`) as the Db2 instance owner (in this case, `db2stj`).

```
    db2 update db cfg for STJ using HADR_LOCAL_HOST dbsec00
    db2 update db cfg for STJ using HADR_LOCAL_SVC STJ_HADR_2
    db2 update db cfg for STJ  using HADR_REMOTE_HOST dbprim00
    db2 update db cfg for STJ  using HADR_REMOTE_SVC STJ_HADR_1
    db2 update db cfg for STJ  using HADR_REMOTE_INST db2stj
    db2 update db cfg for STJ  using HADR_TIMEOUT 120
    db2 update db cfg for STJ  using HADR_SYNCMODE NEARSYNC
    db2 update db cfg for STJ  using HADR_PEER_WINDOW 300
    db2 update db cfg for STJ  using LOGINDEXBUILD ON
```

Here's an example configuration:

```
    > db2 get db cfg for STJ | grep HADR
    HADR database role                                        = STANDBY
    HADR local host name                 (HADR_LOCAL_HOST) = dbsec00
    HADR local service name               (HADR_LOCAL_SVC) = STJ_HADR_2
    HADR remote host name                (HADR_REMOTE_HOST) = dbprim00
    HADR remote service name              (HADR_REMOTE_SVC) = STJ_HADR_1
    HADR instance name of remote server  (HADR_REMOTE_INST) = db2stj
    HADR timeout value                       (HADR_TIMEOUT) = 120
    HADR target list                     (HADR_TARGET_LIST) =
    HADR log write synchronization mode     (HADR_SYNCMODE) = NEARSYNC
    HADR spool log data limit (4KB)       (HADR_SPOOL_LIMIT) = AUTOMATIC(1200000)
    HADR log replay delay (seconds)      (HADR_REPLAY_DELAY) = 0
    HADR peer window duration (seconds)  (HADR_PEER_WINDOW) = 300
    HADR SSL certificate label               (HADR_SSL_LABEL) =
```

5. When using Linux pacemaker, use the following Db2 HADR parameters:

- `HADR peer window duration (seconds) (HADR_PEER_WINDOW) = 300`

- HADR timeout value (HADR_TIMEOUT) = 60

We recommend that you tune these parameters after testing the failover and takeover functionality. Because individual configuration can vary, the parameter might need adjustment.

6. After your primary and standby databases have been configured, start HADR on the standby server as the HADR standby.

```
db2 start hadr on database STJ as standby
```

7. Start HADR on the primary database.

```
db2 start hadr on database STJ as primary
DB20000I  The START HADR ON DATABASE command completed successfully.

db2pd -hadr -db STJ | head -20

Database Member 0 -- Database STJ -- Active --

                           HADR_ROLE = PRIMARY
REPLAY_TYPE = PHYSICAL
                      HADR_SYNCMODE = NEARSYNC
                          STANDBY_ID = 1
                      LOG_STREAM_ID = 0
                          HADR_STATE = PEER
                          HADR_FLAGS = TCP_PROTOCOL
               PRIMARY_MEMBER_HOST = dbprim00
                   PRIMARY_INSTANCE = db2stj
                                 ...
    HADR_CONNECT_STATUS = CONNECTED
```

## Step 6: Pacemaker Cluster Setup

In this section we'll discuss the cluster setup using Linux Pacemaker on both RHEL and SLES OS. The Linux Pacemaker works as a failover Orchestrator. It monitors both the primary and standby databases, and in the event of primary database server failure it initiates an automatic HADR takeover by the standby server. The resource agents this configuration uses are as following:

- STONITH resource agent for fencing.

- The db2 database resource, which is configured in a Primary/Standby configuration.

- The `Amazon-vpc-move-ip` resource, which is built by the Amazon team to handle the overlay IP switch from the Db2 primary instance to standby in the event of failure.

As mentioned in the [Operating System](#) section of this document, you need the correct subscription to download these resource agents.

**Important**: Change the shell environment for the db2<sid> user to /bin/ksh.

**To change the shell environment**:

1. Shut down both the database servers using `db2stop` while logged in as db2<sid>.
2. Install [Kornshell](#) (ksh) (if it's not already installed).
3. Run `sudo usermod -s /bin/ksh db2<sid>`.

# Step 6a. Setup on RHEL

This section focuses on setting up the Pacemaker cluster on the RHEL operating system.

**To set up the pacemaker cluster on RHEL**:

1. Basic cluster configuration – Install the required cluster packages using both database nodes.

    ```
    yum install –y pcs pacemaker fence-agents-aws
    yum install –y resource-agents
    ```

2. Start the cluster services.

    ```
    systemctl start pcsd.service
    systemctl enable pcsd.service
    ```

    **Note:** If you have subscribed to RHEL for SAP with HA and US products from Amazon Marketplace, run `mkdir -p /var/log/pcsd /var/log/cluster` before starting `pcsd.service`.

3. Reset the password for user `hacluster` on both the DB nodes.

    ```
    passwd hacluster
    ```

4. Authorize the cluster. Make sure that both nodes are able to communicate with each other using the hostname.

```
[root@dbprim00 ~] pcs cluster auth dbprim00 dbsec00
Username: hacluster
Password:
dbprim00: Authorized
dbsec00: Authorized
[root@dbprim00 ~]
```

5. Create the cluster.

```
[root@dbprim00 ~] pcs cluster setup --name db2ha dbprim00 dbsec00
Destroying cluster on nodes: dbprim00, dbsec00...
dbsec00: Stopping Cluster (pacemaker)...
dbprim00: Stopping Cluster (pacemaker)...
dbprim00: Successfully destroyed cluster
dbsec00: Successfully destroyed cluster

Sending 'pacemaker_remote authkey' to 'dbprim00', 'dbsec00'
dbprim00: successful distribution of the file 'pacemaker_remote authkey'
dbsec00: successful distribution of the file 'pacemaker_remote authkey'
Sending cluster config files to the nodes...
dbprim00: Succeeded
dbsec00: Succeeded

Synchronizing pcsd certificates on nodes dbprim00, dbsec00...
dbprim00: Success
dbsec00: Success
Restarting pcsd on the nodes in order to reload the certificates...
dbprim00: Success
dbsec00: Success
[root@dbprim00 ~] pcs cluster enable --all
dbprim00: Cluster Enabled
dbsec00: Cluster Enabled
[root@dbprim00 ~] pcs cluster start --all
dbsec00: Starting Cluster...
dbprim00: Starting Cluster...
[root@dbprim00 ~]
```

**Note**: Adjust the `corosync` timeout.

6. Go to `/etc/corosync/corosync.conf` and add or modify the token value of totem to `30000`.

```
[root@dbprim00 corosync] more /etc/corosync/corosync.conf
totem {
      version: 2
      cluster_name: db2ha
      secauth: off
      transport: udpu
      token: 30000
}

nodelist {
      node {
            ring0_addr: dbprim00
            nodeid: 1
      }

      node {
            ring0_addr: dbsec00
            nodeid: 2
      }
}

quorum {
      provider: corosync_votequorum
      two_node: 1
}

logging {
      to_logfile: yes
      logfile: /var/log/cluster/corosync.log
      to_syslog: yes
}
```

7. Run `pcs cluster sync` to sync the changes on the standby database node.

```
[root@dbprim00 corosync] pcs cluster sync
dbprim00: Succeeded
dbsec00: Succeeded
```

8. Run `pcs cluster reload corosync` to make the changes effective.

```
[root@dbprim00 corosync] pcs cluster reload corosync
Corosync reloaded
```

9. To ensure of the changes are in place, run `corosync-cmapctl | grep totem.token`.

```
[root@dbprim00 corosync] corosync-cmapctl | grep totem.token
runtime.config.totem.token (u32) = 30000
runtime.config.totem.token_retransmit (u32) = 7142
runtime.config.totem.token_retransmits_before_loss_const (u32) = 4
totem.token (u32) = 30000
```

10 Before creating any resource, put the cluster in maintenance mode.

```
[root@dbprim00 ~] pcs property set maintenance-mode='true'
```

11 Create the STONITH resource. You will need the EC2 instance IDs for this operation. The default `pcmk` action is reboot. Replace the instance ID for `dbprim00` and `dbsec00` with the instance IDs of your setup.

If you want to have the instance remain in a stopped state until it has been investigated and then manually started, add `pcmk_reboot_action=off`. This setting is also required if you are running the Db2 on [Amazon EC2 Dedicated Hosts](#).

```
[root@dbprim00 ~] pcs stonith create clusterfence fence_aws
region=us-east-1 pcmk_host_map="dbprim00:i-09d1b1f105f71e5ed;dbsec00:i-
0c0d3444601b1d8c5" power_timeout=240 pcmk_reboot_timeout=480
pcmk_reboot_retries=4 op start timeout=300
op monitor timeout=60
```

12 Create the Db2 resource.

```
[root@dbprim00 ~] pcs resource create Db2_HADR_STJ db2
instance=db2stj dblist=STJ master meta notify=true resource-
stickiness=5000 op demote timeout=240 op promote timeout=240 op
start timeout=240 op stop timeout=240 op monitor interval=20s
timeout=120s op monitor interval=22s role=Master timeout=120s
```

**Note**: The timeout values here are default, which works for most deployments. We recommend that you test the timeouts in the QA setup extensively based on the test cases mentioned in the Appendix, and then tune it accordingly.

13 Create the Overlay IP resource agent. First, add the Overlay IP in the primary node.

```
[root@dbprim00 ~] ip address add 192.168.1.81/32 dev eth0
```

```
[root@dbprim00 ~] ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group
default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9001 qdisc mq state UP group
default qlen 1000
    link/ether 0e:10:e3:7b:6f:4f brd ff:ff:ff:ff:ff:ff
    inet 10.0.1.116/24 brd 10.0.1.255 scope global noprefixroute dynamic eth0
        valid_lft 2885sec preferred_lft 2885sec
    inet 192.168.1.81/32 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::c10:e3ff:fe7b:6f4f/64 scope link
        valid_lft forever preferred_lft forever
[root@dbprim00 ~]
```

14. Update the route table with the Overlay IP pointing to the Db2 primary instance:

```
aws ec2 create-route --route-table-id rtb-xxxxxxxx --destination-
cidr-block Overlay IP --instance-id i-xxxxxxxx
[root@dbprim00 ~] aws ec2 create-route --route-table-id rtb-
dbe0eba1 --destination-cidr-block 192.168.1.81/32 --instance-id
i-09d1b1f105f71e5ed
{
    "Return": true
}

[root@dbprim00 ~] pcs resource create db2-oip aws-vpc-move-ip
ip=192.168.1.81 interface=eth0 routing_table=rtb-dbe0eba1
```

**Note**: If you are using a different route table for both the subnets where you are deploying the primary and standby databases, you can specify them using a comma (,) in the resource creation command:

```
pcs resource create db2-oip aws-vpc-move-ip ip=192.168.1.81
interface=eth0 routing_table=rtb-xxxxx1,rtb-xxxxx2
```

15. Create a colocation constraint to bind the Overlay IP resource agent with the Db2 primary instance.

```
      [root@dbprim00 ~] pcs constraint colocation add db2-oip with master
Db2_HADR_STJ-master 2000
```

16.You can now remove the maintenance mode by using the following code:

```
      [root@dbprim00 ~] pcs property set maintenance-mode='false'
```

This is the final configuration of the cluster:

```
      [root@dbprim00 ~] pcs config show
      Cluster Name: db2ha
      Corosync Nodes:
        dbprim00 dbsec00
      Pacemaker Nodes:
        dbprim00 dbsec00

      Resources:
        Master: Db2_HADR_STJ-master
          Resource: Db2_HADR_STJ (class=ocf provider=heartbeat type=db2)
            Attributes: dblist=STJ instance=db2stj
            Meta Attrs: notify=true resource-stickiness=5000
            Operations: demote interval=0s timeout=120 (Db2_HADR_STJ-demote-
 interval-0s)
        monitor interval=20 timeout=60 (Db2_HADR_STJ-monitor-interval-20)
        monitor interval=22 role=Master timeout=60 (Db2_HADR_STJ-monitor-interval-22)
        notify interval=0s timeout=10 (Db2_HADR_STJ-notify-interval-0s)
        promote interval=0s timeout=120 (Db2_HADR_STJ-promote-interval-0s)
        start interval=0s timeout=120 (Db2_HADR_STJ-start-interval-0s)
        stop interval=0s timeout=120 (Db2_HADR_STJ-stop-interval-0s)
        Resource: db2-oip (class=ocf provider=heartbeat type=aws-vpc-move-ip)
        Attributes: interface=eth0 ip=192.168.1.81 routing_table=rtb-dbe0eba1
        Operations: monitor interval=60 timeout=30 (db2-oip-monitor-interval-60)
        start interval=0s timeout=180 (db2-oip-start-interval-0s)
        stop interval=0s timeout=180 (db2-oip-stop-interval-0s)

      Stonith Devices:
        Resource: clusterfence (class=stonith type=fence_aws)
          Attributes:
 pcmk_host_map=dbprim00:i-09d1b1f105f71e5ed;dbsec00:i-0c0d3444601b1d8c5
 pcmk_reboot_retries=4 pcmk_reboot_timeout=480 power_timeout=240 region=us-east-1
          Operations: monitor interval=60s (clusterfence-monitor-interval-60s)
      Fencing Levels:
```

```
        Location Constraints:
        Ordering Constraints:
        Colocation Constraints:
            db2-oip with Db2_HADR_STJ-master (score:2000) (rsc-role:Started) (with-rsc-
     role:Master)
          Ticket Constraints:

        Alerts:
          No alerts defined

        Resources Defaults:
          No defaults set
        Operations Defaults:
          No defaults set

        Cluster Properties:
          cluster-infrastructure: corosync
          cluster-name: db2ha
          dc-version: 1.1.18-11.el7_5.4-2b07d5c5a9
          have-watchdog: false

          Quorum:
          Options:
     [root@dbprim00 ~]
```

# Step 6b. Setup on SLES

This section focuses on setting up the Pacemaker cluster on the SLES operating system.

**Prerequisite**: You need to complete this on both the Db2 primary and standby instances.

**To create an Amazon CLI profile**:

The SLES operating system's resource agents use the [Amazon Command Line Interface](#) (CLI). You need to create the Amazon CLI profile for the root account on both instances: one with the default profile and the other with an arbitrary profile name (in this example, `cluster`) which creates output in text format. The region of the instance must be added as well.

1. Replace the string region-name with your target region in the following example.

```
     dbprim00:~  aws configure
```

```
        {aws} Access Key ID [None]:
        {aws} Secret Access Key [None]:
        Default region name [None]: us-east-1
        Default output format [None]:
        dbprim00:~  aws configure --profile cluster
        {aws} Access Key ID [None]:
        {aws} Secret Access Key [None]:
        Default region name [None]: us-east-1
        Default output format [None]: text
```

You don't need to provide the Access Key and Secret Access key, because access is controlled by the IAM role you created earlier in the setup.

2. Add a second private IP address.

3. You are required to add a second private IP address for each cluster instance. Adding a second IP address to the instance allows the SUSE cluster to implement a two-ring corosync configuration. The two-ring corosync configuration allows the cluster nodes to communicate with each other using the secondary IP address if there is an issue communicating with each other over the primary IP address.

   See [To assign a secondary private IPv4 address to a network interface](#).

4. Add a tag with an arbitrary "key" (in this case, `pacemaker`). The value of this tag is the hostname of the respective Db2 instance. This is required to enable Amazon CLI to use filters in the API calls.

5. Disable the source/destination check.

6. Ensure that the source/destination check is disabled, as described in [Step 4c](#).

7. Avoid deletion of cluster-managed IP addresses on the `eth0` interface.

8. Check if the package `cloud-netconfig-ec2` is installed with the following command:

```
        dbprim00:~  zypper info cloud-netconfig-ec2
```

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

```
        dbprim00:~  CLOUD_NETCONFIG_MANAGE='no'
```

10 Set up [NTP](#) (best with [YaST](#)). Use Amazon time service at 169.254.169.123, which is accessible from all EC2 instances. Enable ongoing synchronization.

11 Activate the public cloud module to get updates for the Amazon CLI:

```
dbprim00:~  SUSEConnect --list-extensions
dbprim00:~  SUSEConnect -p sle-module-public-cloud/12/x86_64
Registering system to registration proxy https://smt-ec2.susecloud.net
Updating system details on https://smt-ec2.susecloud.net ...
 Activating sle-module-public-cloud 12 x86_64 ...
-> Adding service to system ...
-> Installing release package ...
Successfully registered system
```

12.Update your packages on both the with the command:

```
dbprim00:~  zypper -n update
```

13.Install the resource agent pattern ha_sles.

```
dbprim00:~  zypper install -t pattern ha_sles
```

**To configure pacemaker: Configuration of the corosync.conf file**:

1. Use the following configuration in the /etc/corosync/corosync.conf file on both the Db2 primary and standby instances:

```
#Read the corosync.conf.5 manual page
totem {
    version: 2
    rrp_mode: passive
    token: 30000
    consensus: 36000
    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
     }
```

2. Replace the variables `ip-node-1` / `ip2-node-1` and `ip-node-2` / `ip2-node-2` with the IP addresses of your Db2 primary and standby instances, respectively. Replace `ip-local-node` with the IP address of the instance on which the file is being created.

   The chosen settings for `crypto_cipher` and `crypto_hash` are suitable for clusters in Amazon. They may be modified according to SUSE's documentation if you want strong encryption of cluster communication.

3. Start the cluster services and enable them on both the Db2 primary and standby instances.

```
        dbprim00:~ # systemctl start pacemaker
dbprim00:~ # systemctl enable pacemaker
Created symlink from /etc/systemd/system/multi-user.target.wants/pacemaker.service
  to /usr/lib/systemd/system/pacemaker.service.
```

4. Check the configuration with the following command:

```
        dbprim00:~  corosync-cfgtool -s
Printing ring status.
Local node ID 1
RING ID 0
        id      = 10.0.1.17
        status  = ring 0 active with no faults
RING ID 1
        id      = 10.0.1.62
        status  = ring 1 active with no faults
dbprim00:~

Cluster status:
dbprim00:~  crm_mon -1
Stack: corosync
Current DC: dbsec00 (version
 1.1.19+20181105.ccd6b5b10-3.13.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Fri Apr 17 14:09:56 2020
Last change: Fri Apr 17 13:38:59 2020 by hacluster via crmd on dbsec00

2 nodes configured
0 resources configured

Online: [ dbprim00 dbsec00 ]

No active resources
```

**To prepare the cluster for adding resources**:

1. To avoid cluster starting partially defined resources, set the cluster to maintenance mode. This deactivates all monitor actions.

```
    dbprim00:~  crm configure property maintenance-mode="true"
dbprim00:~  crm status
Stack: corosync
```

```
Current DC: dbprim00 (version
 1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Fri Apr 17 14:30:51 2020
Last change: Fri Apr 17 14:30:50 2020 by root via cibadmin on dbprim00

2 nodes configured
0 resources configured


              *** Resource management is DISABLED ***
   The cluster will not attempt to start, stop or recover services


Online: [ dbprim00 dbsec00 ]


No resources
```

2. Configuring Amazon-specific settings:

```
        dbprim00:/ha-files  vi crm-bs.txt
        dbprim00:/ha-files  more 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 off setting forces the agents to shut down the instance. You have the option of changing it to reboot if required.

3. Add the following configuration to the cluster:

```
        dbprim00:~  crm configure load update crm-bs.txt
```

**To configure the Amazon-specific STONITH resource**:

1. 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=180 timeout=60 \
        params tag=pacemaker profile=cluster
```

The EC2 tag `pacemaker` entry needs to match the tag chosen for the EC2 instances, and the name of the profile needs to match the previously configured Amazon profile as part of the prerequisite section.

2. Add the file to the configuration:

```
        dbprim00:/ha-files  vi aws-stonith.txt
 dbprim00:/ha-files  more aws-stonith.txt
 primitive res_{aws}_STONITH stonith:external/ec2 \
 op start interval=0 timeout=180 \
 op stop interval=0 timeout=180 \
 op monitor interval=120 timeout=60 \
 params tag=pacemaker profile=cluster

 dbprim00:/ha-files  crm configure load update aws-stonith.txt
```

3. Create the Db2 Primary/Standby resource.

4. Create a file with the following content. Change the value for SID, as per your configuration.

```
        primitive rsc_db2_db2stj_STJ db2 \
            params instance="db2stj" dblist="STJ" \
            op start interval="0" timeout="130" \
            op stop interval="0" timeout="120" \
            op promote interval="0" timeout="120" \
            op demote interval="0" timeout="120" \
            op monitor interval="30" timeout="60" \
            op monitor interval="31" role="Master" timeout="60"
 ms msl_db2_db2stj_STJ rsc_db2_db2stj_STJ \
            meta target-role="Started" notify="true"
```

5. Add the file to the configuration:

```
        dbprim00:/ha-files  vi db2res.txt
 dbprim00:/ha-files  more db2res.txt
 primitive rsc_db2_db2stj_STJ db2 \
        params instance="db2stj" dblist="STJ" \
        op start interval="0" timeout="130" \
```

```
        op stop interval="0" timeout="120" \
        op promote interval="0" timeout="120" \
        op demote interval="0" timeout="120" \
        op monitor interval="30" timeout="60" \
        op monitor interval="31" role="Master" timeout="60"
dbprim00:/ha-files  crm configure load update db2res.txt
```

6. Create the Overlay IP resource agent.

a. First, add the Overlay IP in the Db2 primary instance.

```
     dbprim00:~ ip address add 192.168.1.81/32 dev eth0
dbprim00:~ ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
 qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9001 qdisc mq state UP group default
 qlen 1000
    link/ether 0e:73:7f:b5:b2:95 brd ff:ff:ff:ff:ff:ff
    inet 10.0.1.17/24 brd 10.0.1.255 scope global eth0
       valid_lft forever preferred_lft forever
    inet 10.0.1.62/32 scope global eth0
       valid_lft forever preferred_lft forever
    inet 192.168.1.81/32 scope global eth0
       valid_lft forever preferred_lft forever
    inet6 fe80::c73:7fff:feb5:b295/64 scope link
       valid_lft forever preferred_lft forever

[root@dbprim00 ~]
```

b. Update the route table with the Overlay IP pointing to the Db2 primary instance.

```
     aws ec2 create-route --route-table-id rtb-xxxxxxxx --destination-cidr-block
  Overlay IP --instance-id i-xxxxxxxx
```

**Note**: If you are using different route table for both the subnets where you are deploying the primary and standby database, you can specify them with a comma (,) in the command preceding this note.

```
dbprim00:~  aws ec2 create-route --route-table-id rtb-dbe0eba1 --
destination-cidr-block 192.168.1.81/32 --instance-id i-05fc8801284585362
{
    "Return": true
}
```

The aws-vpc-move-ip resource agent call the Amazon command from the location /usr/bin, so ensure that there is a soft link pointing to the location where you have the awscli installed.

```
dbprim00:/usr/bin  which aws
/usr/local/bin/aws
dbprim00:/usr/bin  ls -ltr aws
lrwxrwxrwx 1 root root 18 Apr 18 17:44 aws -> /usr/local/bin/aws
```

c.  Create the file with the following content, and replace the Overlay IP and the route table ID based on your configuration. If you have multiple route tables associated with the subnet to which your instances belong, you can use a comma-separated list of routing tables.

   **Note**: Make sure you use the same profile name (which is cluster for this setup) that you used while configuring the Amazon CLI.

```
dbprim00:/ha-files  vi aws-move-ip.txt
dbprim00:/ha-files  more aws-move-ip.txt
primitive res_AWS_IP ocf:suse:aws-vpc-move-ip \
      params ip=192.168.1.81 routing_table=rtb-dbe0eba1 interface=eth0
profile=cluster \
      op start interval=0 timeout=180 \
      op stop interval=0 timeout=180 \
      op monitor interval=60 timeout=60
dbprim00:/ha-files  crm configure load update aws-move-ip.txt
```

d.  Create a colocation constraint to bind the Overlay IP resource agent with the Db2 primary instance.

```
dbprim00:/ha-files  more crm-cs.txt
colocation col_db2_db2stj_STJ 2000: res_AWS_IP:Started \
msl_db2_db2stj_STJ:Master
dbprim00:/ha-files  crm configure load update crm-cs.txt
dbprim00:/ha-files
```

e. Adjust the resource-stickiness and migration-threshold values.

```
dbprim00:~  crm configure rsc_defaults resource-stickiness=1000
dbprim00:~  crm configure rsc_defaults migration-threshold=5000
```

f. You can now remove maintenance-mode.

```
dbprim00:~  crm configure property maintenance-mode="false"
```

Final configuration of the cluster:

```
dbprim00:/ha-files  crm status
Stack: corosync
Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Sat Apr 18 18:45:53 2020
Last change: Sat Apr 18 16:01:26 2020 by root via cibadmin on dbprim00


2 nodes configured
4 resources configured


Online: [ dbprim00 dbsec00 ]


Full list of resources:


res_AWS_STONITH         (stonith:external/ec2): Started dbprim00
Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
Masters: [ dbprim00 ]
Slaves: [ dbsec00 ]
res_AWS_IP     (ocf::suse:aws-vpc-move-ip):    Started dbprim00


dbprim00:/ha-files  crm configure show
node 1: dbprim00
node 2: dbsec00
primitive res_AWS_IP ocf:suse:aws-vpc-move-ip \
params ip=192.168.1.81 routing_table=rtb-dbe0eba1 interface=eth0
profile=cluster \
op start interval=0 timeout=180 \
op stop interval=0 timeout=180 \
op monitor interval=60 timeout=60
primitive res_AWS_STONITH stonith:external/ec2 \
op start interval=0 timeout=180 \
```

```
        op stop interval=0 timeout=180 \
        op monitor interval=120 timeout=60 \
        params tag=pacemaker profile=cluster
        primitive rsc_db2_db2stj_STJ db2 \
        params instance=db2stj dblist=STJ \
        op start interval=0 timeout=130 \
        op stop interval=0 timeout=120 \
        op promote interval=0 timeout=120 \
        op demote interval=0 timeout=120 \
        op monitor interval=30 timeout=60 \
        op monitor interval=31 role=Master timeout=60
        ms msl_db2_db2stj_STJ rsc_db2_db2stj_STJ \
        meta target-role=Started notify=true
        colocation col_db2_db2stj_STJ 2000: res_AWS_IP:Started
    msl_db2_db2stj_STJ:Master
        property cib-bootstrap-options: \
        have-watchdog=false \
        dc-version="1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10" \
        cluster-infrastructure=corosync \
        maintenance-mode=false \
        stonith-enabled=true \
        stonith-action=off \
        stonith-timeout=600s
        rsc_defaults rsc-options: \
        resource-stickiness=1000 \
        migration-threshold=5000
        op_defaults op-options: \
        timeout=600 \
        record-pending=true
```

# Step 7: Post Setup Configuration

To enable SAP to connect to the Db2 virtual name, post-setup configuration tasks must be performed.

To perform post-setup configuration tasks:

1. Edit your SAP profile files:

```
    > vi DEFAULT.PFL

    SAPDBHOST = dbhadb2
```

```
        j2ee/dbhost = dbhadb2

        rsdb/reco_trials  = 10
        rsdb/reco_sleep_time  = 10
```

2. Update the two parameters (SAPDBHOST and `j2ee/dbhost`) to the virtual name you chose for your database server. You will have to update the `rsdb/reco*` parameters to greater than failover duration to avoid DB disconnect in case of failover. We recommend that you test these values in QA before setting it up in production.

3. Edit your Db2 client file:

```
    > cd /sapmnt/STJ/global/db6
    sappas01:stjadm 15> more db2cli.ini
    ; Comment lines start with a semi-colon.
    [STJ]
    Database=STJ
    Protocol=tcpip
    Hostname=dbhadb2
    Servicename=5912
    [COMMON]
    Diagpath=/usr/sap/STJ/SYS/global/db6/db2dump
```

Make sure the hostname parameter matches your Db2 virtual hostname.

4. After you change the entries and save your file, test your connection to the database server:

```
    sappas01:stjadm 17> R3trans -d
    This is R3trans version 6.26 (release 745 - 13.04.18 - 20:18:04).
    unicode enabled version
    R3trans finished (0000).
    sappas01:stjadm 18> startsap
    Checking db Database
    Database is running
    ---------------------------------------------
    Starting Startup Agent sapstartsrv
    OK
    Instance Service on host sappas01 started
    ---------------------------------------------
    starting SAP Instance D00
    Startup-Log is written to /home/stjadm/startsap_D00.log
    ---------------------------------------------
    /usr/sap/STJ/D00/exe/sapcontrol -prot NI_HTTP -nr 00 -function Start
```

```
    Instance on host sappas01 started
```

| Host data | | Database data | |
|---|---|---|---|
| Operating system | Linux | Database System | DB6 |
| Machine type | x86_64 | Release | 11.01.0303 |
| Server name | sappas01_STJ_00 | Name | STJ |
| Platform ID | 390 | Host | dbhadb2 |
| | | Owner | SAPSR3 |

*Figure 6 – SAP system status information*

You can check get the status/information of HADR in the transaction **DB02/dbacockpit** > **Configuration** > **Overview**.

| DB Name | STJ | DB Server | dbhadb2 | Started | | 06.02.2020 | 23:47:46 |
|---|---|---|---|---|---|---|---|
| | | DB Release | 11.01.0303 | | | | |

**Database Instance**

| Name | db2stj | Database Release | 0204010F |
|---|---|---|---|
| Partitionable | 0 | Service Level | DB2 v11.1.3.3 |
| Number of Partitions | 1 | Build Level | special_37682 |
| Address Space | 64 Bit | PTF | …04271300AMD64_37682 |
| | | Fix Pack | 3 |

**Operating System**

| Name | Linux | Host Name | dbprim00 |
|---|---|---|---|
| Version | 3 | Total CPUs | 4 |
| Release | 10 | Configured CPUs | 8 |
| | | Total Memory | 61.242 MB |

*Figure 7 – Transaction DB02 database instance information*

*Figure 8 – Transaction DB02 HADR information*

# Step 8: Testing and Validation

We recommend you define your failure scenarios and test them on your cluster. Unless otherwise specified, all tests are done with the primary node running on the primary server (`dbprim00`) and the standby node running on the standby server (`dbsec00`).

**Prerequisite**: Before running any tests, please ensure that:

- There is no error or failed action in the Pacemaker. This can be tested using pcs status. In case there is any failed action, check the cause in `/var/log/cluster/corosync.log` in the node on which it has failed, and then take the corrective action. You can clean the failed action using `pcs/crm` resource cleanup.

- There is no unintended location constraint set up. Using the `pcs/crm` resource, move the master from primary to standby to set a location constraint on the primary node which prevents any resource from starting on it. This can be identified using the `pcs/crm` constraint show. Note the ID of the location constraint, and then run `pcs/crm` constraint delete `<id>` to remove it.

- The Db2 HADR synchronization is working. This can be checked using db2pd `-hadr -db <DBSID>` and comparing the `LOG_FILE`, `PAGE`, and `POS` for primary and standby.

- Refer to Appendix 1 for detailed test cases on RHEL setup.

- Refer to Appendix 2 for detailed test cases on SLES Setup

# Operations

In this section we will cover some of the native Amazon services that help you with day-to-day operations of your IBM Db2 database for SAP applications.

## Monitoring

Amazon provides multiple native services to monitor and manage your infrastructure and applications on Amazon. Services like Amazon CloudWatch and Amazon CloudTrail can be leveraged to monitor your underlying infrastructure and APIs, respectively.

CloudWatch provides ready-to-use key performance indicators (KPIs) that you can use to monitor CPU utilization and disk utilization.

You can also create custom metrics for monitoring IBM Db2.

With Amazon CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your Amazon infrastructure. Amazon CloudTrail is enabled on all Amazon accounts, and records your account activity upon account creation. You can view and download the last 90 days of your account activity for create, modify, and delete operations of supported services without the need to manually set up CloudTrail.

## Backup and Recovery

You need to regularly back up your operating system and database to recover them in case of failure. Amazon provides various services and tools that you can use to back up your IBM Db2 database of SAP applications.

### Amazon Backup

Amazon Backup is a fully managed backup service centralizes and automates the backup of data across Amazon services. Using Amazon Backup, you can centrally configure backup policies and monitor backup activity for Amazon resources, such as EBS volumes, Amazon EC2 instances, and Amazon Elastic File System (Amazon EFS). Amazon Backup automates and consolidates backup tasks previously performed service-by-service, removing the need to create custom scripts and manual processes. Amazon Backup provides a fully managed, policy-based backup solution, simplifying your backup management and enabling you to meet your business and regulatory backup compliance requirements.

## AMI

You can use the [Amazon Management Console](#) or the Amazon CLI to create a new [Amazon Machine Image](#) (Amazon AMI) of your existing SAP system. This can be used to recover your existing SAP system or create a clone.

The Amazon CLI create image command creates a new AMI based on an existing Amazon EC2 instance. The new AMI contains a complete copy of the operating system and its configuration, software configurations, and optionally all EBS volumes that are attached to the instance.

A simple command to create an AMI with reboot (if running) of your EC2 instance (with instance ID `i-0b09a25c58929de26` as example) including all attached EBS volumes:

```
aws ec2 create-image --instance-id i-0b09a25c58929de26 --name "My server"
```

A simple command to create AMI without reboot (if running[underline]) of your EC2 instance (with instance ID `i-0b09a25c58929de26` as example) including all attached EBS volumes:

```
aws ec2 create-image --instance-id i-0b09a25c58929de26 --name "My server" --no-reboot
```

## Amazon EBS Snapshots

You can back up your Amazon EBS volumes to Amazon S3 by taking point-in-time [snapshots](#). Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved.

Snapshots are suited to backup SAP file systems like `/usr/sap/*` , `/sapmnt/*`. We do not recommend using snapshots to back up your volumes containing data and log files. If you decide to take snapshots for your database volume snapshot, keep in mind that for consistency you should use Microsoft's [Volume Shadow Copy Service](#) and use the [run command](#) to back up or shut down your database before Snapshots is triggered.

A simple command to create a snapshot of volume (with volume `id vol-1234567890abcdef0` as example):

```
aws ec2 create-snapshot --volume-id vol-1234567890abcdef0 --description "This
is my volume snapshot."
```

## Database Backups

One of following methods can be used for IBM Db2 database backup:

- **With native tools to take backup on disk**-- Backup requires high throughput compared to Input/Output Operations Per Second (IOPS). We recommend using st1 disk, which provides maximum throughput of 500MB/s per volume. Once the backup completes on disk it can be moved to an Amazon S3 bucket via scripts.
- **With third party backint tools**-- There are many third-party tools from partners like Commvault and Veritas that use SAP backint interface and store backups directly in Amazon S3 buckets.

## Storage

The storage services we use across this guide are:

### Amazon EBS

Amazon EBS provides persistent storage for SAP applications and databases. EBS volume size can be increased or their type can be changed (for example, gp2 to io1) without downtime requirements. For more information, see Modifying Amazon EBS volume.

Once you have extended the volume, you need to extend the drive with your Linux volume manager software.

### Amazon S3

Amazon S3 does not need you to explicitly provision storage at all – you simply pay for what you use.

# Operating System Maintenance

Operating system maintenance across large estates of EC2 instances can be managed by:

- Tools specific to each operating system such as SUSE Manager and Red Hat Smart Management.
- 3rd party products such as those available on the Amazon Marketplace.
- Using Amazon Systems Manager.

Here are some key operating system maintenance tasks that can help with:

## Patching

Follow SAP recommended patching processes to update your landscape on Amazon. For operating system patching, with Amazon Systems Manager Patch Manager you can roll out OS patches as per your corporate policies. There are multiple key features such as:

- Scheduling based on tags
- Auto-approving patches with lists of approved and rejected patches
- Defining patch baselines

Amazon Systems Manager Patch Manager integrates with IAM, Amazon CloudTrail, and Amazon CloudWatch Events to provide a secure patching experience that includes event notifications and the ability to audit usage. For details about the process, see How Patch Manager Operations Work. If Amazon Systems Manager Patch Manager does not fulfil your requirements, there are third-party products available as well. Some of these are available via the Amazon Marketplace.

## Maintenance Window

Amazon Systems Manager Maintenance Windows enables you to define a schedule for when to perform potentially disruptive actions on your instances, such as patching an operating system, updating drivers, or installing software or patches.

## Automation Using Documents

Amazon Systems Manager Automation simplifies common maintenance and deployment tasks of Amazon EC2 instances and other Amazon resources. Automation enables you to do the following:

- Build automation workflows to configure and manage instances and Amazon resources.
- Create custom workflows or use pre-defined workflows maintained by Amazon.
- Receive notifications about Automation tasks and workflows by using Amazon CloudWatch Events.
- Monitor automation progress and execution details by using the Amazon EC2 or the Amazon Systems Manager console.

## Business Continuity

Amazon recommends periodically scheduling business continuity process validations by executing disaster recovery tests. This planned activity helps to flush out any potential unknowns, and helps

the organization deal with any real disaster in a streamlined manner. Depending on your disaster recovery architecture it may include:

- Backup/Recovery of databases from S3.
- Creation of systems from AMI and point-in-time recovery via snapshots.
- Changing EC2 instance size of pilot light systems.
- Validation of integration (AD/DNS, email, 3$^{rd}$ party, and more)

## Support

SAP requires customers to have a minimum [Amazon Business Support](#) plan with Amazon. This ensures that any critical issues raised with SAP are also handled by Amazon on priority. Amazon business support provides a less than one-hour response time for production-down scenarios. You can also choose to have an Amazon enterprise support plan, which provides a less than 15-minute response time for business-critical systems, along with other benefits. See [Amazon Enterprise Support](#).

For any SAP application issues, Amazon suggests raising an incident with SAP via the SAP support portal. After the first level of investigation, SAP can redirect the incident to Amazon support if they find an infrastructure related issue which needs to be managed by Amazon. However, if you choose to raise support issues for SAP applications with Amazon support, we cannot redirect the tickets to SAP. For any infrastructure related issues, you can raise the issue directly with Amazon support.

## Cost Optimization

Resources (CPU, memory, additional application servers, system copies for different tests/validations and more) required the SAP landscape change over time. Amazon recommends monitoring system utilization, and the need for existing systems, on a regular basis to take actions that will reduce cost. In cases of databases like IBM Db2 as we cannot scale out only opportunity to right size database server is by scaling up/down or shutting it down if not required. A few suggestions to consider:

- Consider reserved instances or savings plans over on-demand instances if your requirement is to run 24-7, 365 days a year. Reserved instances provide up to 75% discount over on-demand instances. See [Amazon EC2 pricing](#).
- Consider running occasionally required systems like training and sandbox on-demand for the duration required.

- Monitor CPU and memory utilization overtime for other non-production systems like Dev/QA, and right-size them when possible.

# Appendix 1: Testing on RHEL Setup

## Test Case 1: Manual Failover

**Procedure**: Use the command `pcs resource move <Db2 master resource name>`.

```
[root@dbprim00 profile] pcs resource move Db2_HADR_STJ-master
Warning: Creating location constraint cli-ban-Db2_HADR_STJ-master-on-dbprim00 with
a score of -INFINITY for resource
Db2_HADR_STJ-master on-dbprim00 with a score of -INFINITY for resource
Db2_HADR_STJ-
master on node dbprim00.
This will prevent Db2_HADR_STJ-master from running on dbprim00
until the constraint is removed. This will be the case even if
dbprim00 is the last node in the cluster.
[root@dbprim00 profile]
```

**Expected result**: The Db2 primary node is moved from primary node to standby node.

```
[root@dbprim00 profile] pcs status
Cluster name: db2ha
Stack: corosync
Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
Last updated: Sat Feb  8 08:54:04 2020
Last change: Sat Feb  8 08:53:02 2020 by root via crm_resource on dbprim00

2 nodes configured
4 resources configured

Online: [ dbprim00 dbsec00 ]

Full list of resources:

clusterfence    (stonith:fence_aws):    Started dbprim00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
    Masters: [ dbsec00 ]
Stopped: [ dbprim00 ]
db2-oip         (ocf::heartbeat:aws-vpc-move-ip):     Started dbsec00
```

```
    Daemon Status:
    corosync: active/enabled
    pacemaker: active/enabled
    pcsd: active/enabled
    [root@dbprim00 profile]
```

**Followup actions**: Remove the location constraint.

When using a manual command for moving the resource, there is location constraint created on the node (in this case, the primary node) that prevents running the Db2 resource in standby mode.

**To remove the location constraint:**

1. Use the following command to remove the location constraint:

```
     pcs config show
    Location Constraints:
    Resource: Db2_HADR_STJ-master
    Disabled on: dbprim00 (score:-INFINITY) (role: Started) (id:cli-ban-
  Db2_HADR_STJ-master-on-dbprim00)

    [root@dbprim00 profile] pcs constraint delete cli-ban-Db2_HADR_STJ-master-on-
  dbprim00
```

2. Start the Db2 instance as standby on the new standby node, logged in as db2<sid>. Next, clean up the error logged in as root.

```
    db2stj> db2start
    02/08/2020 09:11:29      0    0    SQL1063N   DB2START processing was successful.
    SQL1063N   DB2START processing was successful.

    db2stj> db2 start hadr on database STJ as standby
    DB20000I   The START HADR ON DATABASE command completed successfully.

    [root@dbprim00 ~] pcs resource cleanup
    Cleaned up all resources on all nodes
    [root@dbprim00 ~] pcs status
    Cluster name: db2ha
    Stack: corosync
    Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with
  quorum
    Last updated: Sat Feb  8 09:13:17 2020
```

```
    Last change: Sat Feb  8 09:12:26 2020 by hacluster via crmd on dbprim00


    2 nodes configured
    4 resources configured


    Online: [ dbprim00 dbsec00 ]


    Full list of resources:


    clusterfence   (stonith:fence_aws):    Started dbprim00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
        Masters: [ dbsec00 ]
        Slaves: [ dbprim00 ]
    db2-oip        (ocf::heartbeat:aws-vpc-move-ip):     Started dbsec00


    Daemon Status:
        corosync: active/enabled
        pacemaker: active/enabled
        pcsd: active/enabled
    [root@dbprim00 ~]
```

## Test Case 2: Shut Down the Primary EC2 Instance

**Procedure**: Using Amazon Console or CLI to stop the EC2 instance and simulate EC2 failure.

**Expected result**: The Db2 primary node is moved to the standby server.

```
    [root@dbsec00 db2stj] pcs status
    Cluster name: db2ha
    Stack: corosync
    Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
    Last updated: Sat Feb  8 09:44:16 2020
    Last change: Sat Feb  8 09:31:39 2020 by hacluster via crmd on dbsec00


    2 nodes configured
    4 resources configured


    Online: [ dbsec00 ]
    OFFLINE: [ dbprim00 ]


    Full list of resources:
```

```
    clusterfence    (stonith:fence_aws):      Started dbsec00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
        Masters: [ dbsec00 ]
        Stopped: [ dbprim00 ]
    db2-oip         (ocf::heartbeat:aws-vpc-move-ip):       Started dbsec00

    Daemon Status:
        corosync: active/enabled
        pacemaker: active/enabled
        pcsd: active/enabled
```

**Followup action**: Start the EC2 instance and then start Db2 as standby on the standby instance as you did in Test Case 1. Do not include location constraint removal this time.

## Test Case 3: Stop the Db2 Instance on the Primary Instance

**Procedure**: Log in to the Db2 primary instance as db2<sid> (db2stj) and run db2stop force.

```
    db2stj> db2stop force
    02/12/2020 12:40:03      0   0    SQL1064N  DB2STOP processing was successful.
    SQL1064N  DB2STOP processing was successful.
```

**Expected result**: The Db2 primary node is failed over to standby server. The standby node continues to be on the old primary in a stopped state. There is a failed monitoring action.

```
    [root@dbsec00 db2stj] pcs status
    Cluster name: db2ha
    Stack: corosync
    Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
    Last updated: Wed Feb 12 16:55:56 2020
    Last change: Wed Feb 12 13:58:11 2020 by hacluster via crmd on dbsec00

    2 nodes configured
    4 resources configured

    Online: [ dbprim00 dbsec00 ]

    Full list of resources:

    clusterfence    (stonith:fence_aws):      Started dbsec00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
        Masters: [ dbsec00 ]
```

```
       Stopped: [ dbprim00 ]
    db2-oip          (ocf::heartbeat:aws-vpc-move-ip):        Started dbsec00


    Failed Actions:
    * Db2_HADR_STJ_start_0 on dbprim00 'unknown error' (1): call=34, status=complete,
  exitreason='',
     last-rc-change='Wed Feb 12 16:55:32 2020', queued=1ms, exec=6749ms



    Daemon Status:
       corosync: active/enabled
       pacemaker: active/enabled
    pcsd: active/enabled
    [root@dbsec00 db2stj]
```

**Followup action**: Start the EC2 instance, then start Db2 as standby on the standby instance as you did in Test Case 2. Clear the failed monitoring error.

# Test Case 4: End the Db2 Process (db2sysc) on the Node that Runs the Primary Database

**Procedure**: Log in to the Db2 primary instance as root and then run `ps -ef|grep db2sysc`. Note the process ID (PID) and then end it.

```
    [root@dbprim00 ~] ps -ef|grep db2sysc
    root       5809 30644  0 18:54 pts/1     00:00:00 grep --color=auto
    db2sysc
    db2stj   26982 26980  0 17:12 pts/0     00:00:28 db2sysc 0
    [root@dbprim00 ~] kill -9 26982
```

**Expected result**: The Db2 primary node is failed over to the standby server. The standby node is in the old primary in a stopped state.

```
    [root@dbprim00 ~] pcs status
    Cluster name: db2ha
    Stack: corosync
    Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
    Last updated: Wed Feb 12 18:54:50 2020
    Last change: Wed Feb 12 18:53:12 2020 by hacluster via crmd on dbsec00

    2 nodes configured
    4 resources configured
```

```
    Online: [ dbprim00 dbsec00 ]

    Full list of resources:

    clusterfence    (stonith:fence_aws):     Started dbsec00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
        Masters: [ dbsec00 ]
        Stopped: [ dbprim00 ]
    db2-oip         (ocf::heartbeat:aws-vpc-move-ip):       Started dbsec00

    Failed Actions:
    * Db2_HADR_STJ_start_0 on dbprim00 'unknown error' (1): call=57, status=complete,
 exitreason='',
    last-rc-change='Wed Feb 12 18:54:37 2020', queued=0ms, exec=6777ms


    Daemon Status:
        corosync: active/enabled
        pacemaker: active/enabled
        pcsd: active/enabled
```

**Followup action**: Start the EC2 instance and start Db2 as standby on the standby instance, as you did in Test Case 2. Clear the failed monitoring alert.

# Test Case 5: End the Db2 Process (db2sysc) on the Node that Runs the Standby Database

**Procedure**: Log in to the Db2 standby instance as root and run `ps -ef|grep db2sysc`. Note the PID and then end it.

```
    [root@dbsec00 db2stj] ps -ef|grep db2sysc
    db2stj   24194 24192  1 11:55 pts/1     00:00:01 db2sysc 0
    root     26153  4461  0 11:57 pts/0     00:00:00 grep --color=auto
    db2sysc
    [root@dbsec00 db2stj] kill -9 24194
```

**Expected result**: The db2sysc process is restarted on the Db2 standby instance. There is a monitoring failure event record in the cluster.

```
    [root@dbprim00 ~] pcs status
```

```
    Cluster name: db2ha
    Stack: corosync
    Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
    Last updated: Fri Feb 14 11:59:22 2020
    Last change: Fri Feb 14 11:55:54 2020 by hacluster via crmd on dbsec00


    2 nodes configured
    4 resources configured


    Online: [ dbprim00 dbsec00 ]


    Full list of resources:


    clusterfence    (stonith:fence_aws):     Started dbsec00
    Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
        Masters: [ dbprim00 ]
        Slaves: [ dbsec00 ]
    db2-oip         (ocf::heartbeat:aws-vpc-move-ip):        Started dbprim00


    Failed Actions:
    * Db2_HADR_STJ_monitor_20000 on dbsec00 'not running' (7): call=345,
 status=complete, exitreason='',
    last-rc-change='Fri Feb 14 11:57:57 2020', queued=0ms, exec=0ms



    Daemon Status:
        corosync: active/enabled
        pacemaker: active/enabled
        pcsd: active/enabled



    [root@dbsec00 db2stj] ps -ef|grep db2sysc
    db2stj   26631 26629  1 11:57 ?         00:00:01 db2sysc 0
    root     27811  4461  0 11:58 pts/0     00:00:00 grep --color=auto db2sysc
```

**Follow-up action**: Clear the monitoring error.

# Test Case 6: Simulating a Crash of the Node that Runs the Primary Db2

**Procedure**: Log in to the Db2 primary instance as root and run echo  'c' > /proc/sysrq-trigger.

```
    [root@dbprim00 ~] echo 'c' > /proc/sysrq-trigger
```

```
####################################################################################

    Session stopped
        - Press <return> to exit tab
        - Press R to restart session
        - Press S to save terminal output to file

    Network error: Software caused connection abort
```

**Expected result**: The primary Db2 should failover to standby node. The standby is in a stopped state on the previous primary.

```
[root@dbsec00 ~] pcs status
Cluster name: db2ha
Stack: corosync
Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
Last updated: Fri Feb 21 15:38:43 2020
Last change: Fri Feb 21 15:33:17 2020 by hacluster via crmd on dbsec00

2 nodes configured
4 resources configured

Online: [ dbprim00 dbsec00 ]

Full list of resources:

 clusterfence    (stonith:fence_aws):    Started dbsec00
 Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
     Masters: [ dbsec00 ]
     Stopped: [ dbprim00 ]
 db2-oip         (ocf::heartbeat:aws-vpc-move-ip):       Started dbsec00

Failed Actions:
* Db2_HADR_STJ_start_0 on dbprim00 'unknown error' (1): call=15, status=complete,
exitreason='',
    last-rc-change='Fri Feb 21 15:38:31 2020', queued=0ms, exec=7666ms


Daemon Status:
    corosync: active/enabled
    pacemaker: active/enabled
    pcsd: active/enabled
```

**Followup action**: Start the EC2 instance and then start Db2 as standby on the standby instance as you did in Test Case 2. Clear the monitoring error.

# Appendix 2: Testing on SLES Setup

## Test Case 1: Manual Failover

**Procedure**: Use the command `crm resource move <Db2 primary resource name> force` to move the primary Db2 instance to standby node.

```
dbprim00:  crm resource move msl_db2_db2stj_STJ force
INFO: Move constraint created for rsc_db2_db2stj_STJ
```

**Expected result**: The Db2 primary node is moved from the primary node (`dbprim00`) to the standby node (`dbsec00`).

```
dbprim00:~  crm status
Stack: corosync
Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Sat Apr 25 19:03:20 2020
Last change: Sat Apr 25 19:02:26 2020 by root via crm_resource on dbprim00


2 nodes configured
4 resources configured


Online: [ dbprim00 dbsec00 ]


Full list of resources:

res_AWS_STONITH         (stonith:external/ec2): Started dbsec00
Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
    Masters: [ dbsec00 ]
    Stopped: [ dbprim00 ]
res_AWS_IP     (ocf::suse:aws-vpc-move-ip):    Started dbsec00
```

**Follow-up actions**: Start the Db2 instance as standby on the new standby node, logged in as db2<sid>. Clean up the error logged in as root.

```
db2stj> db2start
```

```
04/25/2020 19:05:27     0   0   SQL1063N  DB2START processing was successful.
SQL1063N  DB2START processing was successful.


db2stj> db2 start hadr on database STJ as standby
DB20000I  The START HADR ON DATABASE command completed successfully.
```

**Remove location constraint**: When using a manual command to move the resource, there is a location constraint created on the node (in this case primary node) which is run, preventing the Db2 resource from running in standby mode.

Use the following command to remove the location constraint.

```
dbprim00:  crm resource clear msl_db2_db2stj_STJ
```

Once the constraint is removed, the standby instance starts automatically.

```
    dbprim00:  crm status
Stack: corosync
Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Sat Apr 25 19:05:29 2020
Last change: Sat Apr 25 19:05:18 2020 by root via crm_resource on dbprim00

2 nodes configured
4 resources configured

Online: [ dbprim00 dbsec00 ]

Full list of resources:

res_AWS_STONITH         (stonith:external/ec2): Started dbsec00
Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
    Masters: [ dbsec00 ]
    Slaves: [ dbprim00 ]
res_AWS_IP     (ocf::suse:aws-vpc-move-ip):    Started dbsec00
```

# Test Case 2: Shut Down the Primary EC2 Instance

**Procedure**: Using Amazon console or CLI, stop the EC2 instance to simulate EC2 failure.

Expected Result: The Db2 primary node is moved to a standby server (dbsec00).

```
dbsec00:~  crm status
Stack: corosync
Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Sat Apr 25 19:19:32 2020
Last change: Sat Apr 25 19:18:16 2020 by root via crm_resource on dbprim00


2 nodes configured
4 resources configured


Online: [ dbsec00 ]
OFFLINE: [ dbprim00 ]


Full list of resources:


res_AWS_STONITH       (stonith:external/ec2): Started dbsec00
Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
    Masters: [ dbsec00 ]
    Stopped: [ dbprim00 ]
res_AWS_IP     (ocf::suse:aws-vpc-move-ip):   Started dbsec00
```

**Follow-up action**: Start the EC2 instance and the standby node should start on `dbprim00`.

# Test Case 3: Stop the Db2 Instance on the Primary Instance

**Procedure**: Log in to the Db2 primary instance (`dbprim00`) as db2<sid> (`db2stj`) and run `db2stop force`.

```
db2stj> db2stop force
02/12/2020 12:40:03     0   0   SQL1064N  DB2STOP processing was successful.
SQL1064N  DB2STOP processing was successful.
```

**Expected result**: The Db2 primary node will failover on primary instance. The standby remains on the standby instance. There is a failed resource alert.

```
dbsec00:~  crm status
Stack: corosync
Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
Last updated: Sat Apr 25 19:29:38 2020
Last change: Sat Apr 25 19:23:04 2020 by root via crm_resource on dbprim00
```

```
    2 nodes configured
    4 resources configured


    Online: [ dbprim00 dbsec00 ]


    Full list of resources:


    res_AWS_STONITH        (stonith:external/ec2): Started dbprim00
    Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
        Masters: [ dbprim00 ]
        Slaves: [ dbsec00 ]
    res_AWS_IP     (ocf::suse:aws-vpc-move-ip):    Started dbprim00


    Failed Resource Actions:
    * rsc_db2_db2stj_STJ_demote_0 on dbprim00 'unknown error' (1): call=74,
 status=complete, exitreason='',
    last-rc-change='Sat Apr 25 19:27:21 2020', queued=0ms, exec=175ms
```

**Followup action**: Clear the failed cluster action.

```
    dbsec00:~  crm resource cleanup
    Waiting for 1 reply from the CRMd. OK
```

# Test Case 4: End the Db2 Process (db2sysc) on the Node that Runs the Primary Database

**Procedure**: Log in to the Db2 primary instance as root and run `ps -ef|grep db2sysc`. Note the PID and then end it.

```
    dbprim00:~  ps -ef|grep db2sysc
    db2stj   11690 11688  0 19:27 ?        00:00:02 db2sysc 0
    root     15814  4907  0 19:31 pts/0    00:00:00 grep --color=auto db2sysc
    [root@dbprim00 ~] kill -9 11690
```

**Expected result**: The Db2 primary node is restarted on the primary instance. The standby node remains on the standby instance. There is a failed resource alert.

```
    dbsec00:~  crm status
    Stack: corosync
    Current DC: dbsec00 (version
 1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
```

```
    Last updated: Sat Apr 25 19:29:38 2020
    Last change: Sat Apr 25 19:23:04 2020 by root via crm_resource on dbprim00


    2 nodes configured
    4 resources configured


    Online: [ dbprim00 dbsec00 ]


    Full list of resources:


    res_AWS_STONITH        (stonith:external/ec2): Started dbprim00
    Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
        Masters: [ dbprim00 ]
        Slaves: [ dbsec00 ]
    res_AWS_IP      (ocf::suse:aws-vpc-move-ip):    Started dbprim00


    Failed Resource Actions:
    * rsc_db2_db2stj_STJ_demote_0 on dbprim00 'unknown error' (1): call=74,
 status=complete, exitreason='',
    last-rc-change='Sat Apr 25 19:27:21 2020', queued=0ms, exec=175ms
```

**Followup action**: Clear the failed cluster action.

```
    dbsec00:~  crm resource cleanup
    Waiting for 1 reply from the CRMd. OK
```

# Test Case 5: End the Db2 Process (db2sysc) on the Node that Runs the Standby Database

**Procedure**: Log in to the standby DB instance (dbsec00) as root, then run ps -ef|grep db2sysc. Note the PID and then end it.

```
    dbsec00:~  ps -ef| grep db2sysc
    db2stj   16245 16243  0 19:23 ?        00:00:04 db2sysc 0
    root     28729 28657  0 19:38 pts/0    00:00:00 grep --color=auto db2sysc
    dbsec00:~  kill -9 16245
```

**Expected result**: The db2sysc process is restarted on the standby DB instance. There is a monitoring failure event recorded in the cluster.

```
    dbsec00:~  crm status
```

```
    Stack: corosync
    Current DC: dbsec00 (version
1.1.19+20181105.ccd6b5b10-3.16.1-1.1.19+20181105.ccd6b5b10) - partition with quorum
    Last updated: Sat Apr 25 19:40:23 2020
    Last change: Sat Apr 25 19:23:04 2020 by root via crm_resource on dbprim00


    2 nodes configured
    4 resources configured


    Online: [ dbprim00 dbsec00 ]


    Full list of resources:


    res_AWS_STONITH        (stonith:external/ec2): Started dbprim00
    Master/Slave Set: msl_db2_db2stj_STJ [rsc_db2_db2stj_STJ]
        Masters: [ dbprim00 ]
        Slaves: [ dbsec00 ]
    res_AWS_IP     (ocf::suse:aws-vpc-move-ip):    Started dbprim00


    Failed Resource Actions:
    * rsc_db2_db2stj_STJ_monitor_30000 on dbsec00 'not running' (7): call=387,
status=complete, exitreason='',
    last-rc-change='Sat Apr 25 19:39:24 2020', queued=0ms, exec=0ms
```

**Followup action**: Clear the monitoring error.

```
    dbsec00:~  crm resource cleanup
    Waiting for 1 reply from the CRMd. OK
```

# Test Case 6: Simulating a Crash of the Node that Runs the Primary Db2

**Procedure**: Log in to the Db2 primary instance as root, then run echo 'c' > /proc/sysrq-trigger.

```
    dbprim00:~  echo 'c' > /proc/sysrq-trigger
    Session stopped
        - Press <return> to exit tab
        - Press R to restart session
        - Press S to save terminal output to file


    Network error: Software caused connection abort
```

**Expected result**: The primary Db2 should failover to standby node.vThe standby is in a stopped state on the previous primary (`dbprim00`).

```
[root@dbsec00 ~] crm status
Cluster name: db2ha
Stack: corosync
Current DC: dbsec00 (version 1.1.18-11.el7_5.4-2b07d5c5a9) - partition with quorum
Last updated: Fri Feb 21 15:38:43 2020
Last change: Fri Feb 21 15:33:17 2020 by hacluster via crmd on dbsec00

2 nodes configured
4 resources configured

Online: [ dbprim00 dbsec00 ]

Full list of resources:

clusterfence    (stonith:fence_aws):    Started dbsec00
Master/Slave Set: Db2_HADR_STJ-master [Db2_HADR_STJ]
     Masters: [ dbsec00 ]
     Stopped: [ dbprim00 ]
db2-oip        (ocf::heartbeat:aws-vpc-move-ip):      Started dbsec00

Failed Actions:
* Db2_HADR_STJ_start_0 on dbprim00 'unknown error' (1): call=15, status=complete,
exitreason='',
  last-rc-change='Fri Feb 21 15:38:31 2020', queued=0ms, exec=7666ms


Daemon Status:
    corosync: active/enabled
    pacemaker: active/enabled
    pcsd: active/enabled
```

**Followup action**: Start the EC2 instance and then start Db2 as standby on the standby instance as you did in [Test Case 2](#).

# FAQ

**Question:** Can I use Database Migration Service to migrate and deploy SAP NetWeaver on IBM Db2 based applications?

**Answer:** No, Amazon DMS supports IBM Db2 as a source, but it is not certified by SAP for SAP NetWeaver-based applications.

# Document Revisions

| Date | Change |
|---|---|
| **December 2020** | First publication |

# Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current Amazon product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from Amazon and its affiliates, suppliers or licensors. Amazon products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of Amazon to its customers are controlled by Amazon agreements, and this document is not part of, nor does it modify, any agreement between Amazon and its customers.

# Databases for SAP on Amazon with Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. You can now deploy and operate SAP applications using IBM Db2, SAP MaxDB, SAP ASE, Oracle, and MSSQL on Amazon with FSx for ONTAP. For more information, see Amazon FSx for NetApp ONTAP.

If you are a first-time user, see How Amazon FSx for NetApp ONTAP works.

**Topics**

- Instances and sizing
- Supported instance types
- Sizing
- Create storage virtual machines (SVM)
- Volume configuration and layout
- File system setup
- Architecture diagrams for databases with Amazon FSx for NetApp ONTAP
- Host setup for databases with Amazon FSx for NetApp ONTAP
- Installing the databases

## Instances and sizing

The following rules and limitations are applicable for deploying SAP applications on Amazon using IBM Db2, SAP MaxDB, SAP ASE, Oracle, and MSSQL with FSx for ONTAP.

- Amazon EC2 instance where you plan to deploy the database and FSx for ONTAP must be in the same subnet.
- Use separate storage virtual machines (SVM) for binarie, data, and log volumes of the database. This ensures that your I/O traffic flows through different IP addresses and TCP sessions.
- The following file systems must have their separate FSx for ONTAP volumes. See the following tabs.

**Example**

IBM Db2

- /db2/<DBSID>/sapdata1
- /db2/<DBSID>/sapdata1
- /db2/<DBSID>/log_dir
- /db2/<DBSID>/saptmp<x>
- /db2/<DBSID>/log_arch
- /usr/sap

SAP MaxDB

- /sapdb/<SID>/sapdata
- /sapdb/<SID>/saplog
- /sapdb/<SID>/backup
- /usr/sap

SAP ASE

- /sybase/<SID>/sapdata_1
- /sybase/<SID>/sapdata_n
- /sybase/<SID>/saplog_1
- /sybase/<SID>/saptmp
- /sybase/<SID>/sapdiag
- /sybasebackup
- /usr/sap

Oracle

- /oracle/<DBSID>
- /oracle/<DBSID>/oraarch
- /oracle/<DBSID>/saparch
- /oracle/<DBSID>/sapreorg
- /oracle/<DBSID>/sapbackup
- /oracle/<DBSID>/saptrace
- /oracle/<DBSID>/sapdata<x>
- /oracle/<DBSID>/origlogA

- /oracle/<DBSID>>/origlogB

- /oracle/<DBSID>/mirrlogA

- /oracle/<DBSID>/mirrlogB

MSSQL

- MSSQL data

- MSSQL log

- MSSQL tempdb

- SQL binaries (including SQL Server system databases)

- Backup directories

- FSx for ONTAP creates endpoints for accessing your file system in a Amazon VPC route table. FSx for ONTAP uses your Amazon VPC main route table. We recommend configuring your file system to use all of your Amazon VPC route tables that are associated with the subnets in which your clients are located. Alternatively, you can specify one or more route tables for Amazon FSx to use when you create your file system.

- *Oracle* – SAP with Oracle Database is supported on Amazon EC2 instances with Amazon FSx for NetApp ONTAP only using Oracle Direct NFS (dNFS) with NFSv3, NFSv4 and NFSv4.1. Kernel NFS mounted filesystems are not supported by SAP for any Oracle files, such as database files, redo logs, and control files. For more information, see [SAP Note 2358420 - Oracle Database Support for Amazon Web Services EC2](link) (requires SAP portal access).

# Supported instance types

See the following tabs for details.

**Example**

IBM Db2

    FSx for ONTAP is supported for SAP applications using IBM Db2 database in a Single or Multi-Availability Zone deployment. You can use FSx for ONTAP as the primary storage solution for IBM Db2 database devices and backup volumes with supported Amazon EC2 instances.

SAP MaxDB

    FSx for ONTAP is supported for SAP MaxDB version 7.9.10 or higher (including liveCache) with SAP applications in a Single or Multi-Availability Zone deployment. You can use FSx for ONTAP

as the primary storage solution for SAP MaxDB `data`, `log`, `binary`, and `backup` volumes with supported Amazon EC2 instances.

SAP ASE

FSx for ONTAP is supported for SAP ASE version 16.0 for Business Suite and higher with SAP NetWeaver based applications in a Single or Multi-Availability Zone deployment. You can use FSx for ONTAP as the primary storage solution for SAP ASE `data`, `log`, `sapdiag`, `saptmp`, and `backup` volumes with supported Amazon EC2 instances.

Oracle

FSx for ONTAP is supported for Oracle 19c (minimum 19.5.0) for Business Suite and higher with SAP NetWeaver based applications in a Single or Multi-Availability Zone deployment. You can use FSx for ONTAP as the primary storage solution for Oracle `data`, `log`, `binaries`, `archive`, `backup`, and other volumes volumes with supported Amazon EC2 instances.

You can use the Quick Sizer tool from SAP to calculate the compute requirement in SAPS for a greenfield (new) Oracle deployment. It enables you to choose an instance type that meets your business requirements and is cost-efficient. You can use source system utilisation and workload patterns, such as SAP EarlyWatch alert reports, source system specification: CPU+ memory and source system SAPS rating for migrations.

MSSQL

FSx for ONTAP is supported for SAP applications using MSSQL database in a Single or Multi-Availability Zone deployment. You can use FSx for ONTAP as the primary storage solution for MSSQL `data`, `log`, `tempdb`, `SQL binaries`, and `backup` volumes with supported Amazon EC2 instances.

For a complete list of supported Amazon EC2 instances for the databases, see [SAP Note 1656099 - SAP Applications on Amazon EC2: Supported DB/OS and Amazon EC2 products](#) (requires SAP portal access).

# Sizing

A single FSx for ONTAP file system can provide a maximum output of 80,000 IOPS. Select multiple FSx for ONTAP file systems if your I/O requirement exceeds 80,000 IOPS.

You can configure the throughput capacity of FSx for ONTAP when you create a new file system by scaling up to 4 GB/s of read throughput and 1000 MB/s of write throughput in a single Availability

Zone deployment. In a multi-Availability Zone deployment, you can create a file system by scaling up to 4 GB/s of read throughput and 1800 MB/s of write throughput. For more information, see [Performance details](#).

### Oracle

You can calculate the IOPS required by your database by querying the system tables over a period of time and selecting the highest value. You can get this information from the GV$SYSSTAT dynamic performance view from Oracle. This view is continuously updated while the database is open and in use. IOPS requirements can also be validated through Oracle Enterprise Manager and Automatic Workload Repository reports which use these views to gather data. For more information, see [Estimating IOPS for an existing database](#).

### MSSQL

For migrating your existing SAP applications based on SQL Server to Amazon, use Windows Performance Monitor to get information about IOPS and throughput required for FSx for ONTAP file system. To open Windows Performance Monitor, run `perfmon` at command prompt. IOPS and throughput data is provided by the following performance counters:

- Peak Disk reads/sec + Peak disk writes/sec = IOPS
- Peak Disk read bytes/sec + Peak disk write bytes/sec = throughput

We also recommend that you get the IOPS and throughput data over a typical workload cycle to get a good estimate for your requirements. Ensure that the FSx for ONTAP file system you provision for SQL server supports these I/O and throughput requirements.

For sizing your SAP on SQL environment, you can use:

- SAP Quick Sizer (for new implementations)
- SAP Early Watch Alerts reports

Running MSSQL with FSx for ONTAP file system separates compute and storage. Therefore, SQL server running on smaller EC2 instances connected to FSx for ONTAP can perform the same as SQL server running on a much larger EC2 instance. A smaller instance can also lower the total cost of ownership for your SAP landscape on Amazon. For standard database workloads, you can run memory optimized instance classes, such as `r6i`, `r6in`, `r6a`, `r7i`, and `r7a`. For a complete list of supported instances, see [Amazon EC2 instance types for SAP on Amazon](#).

# Create storage virtual machines (SVM)

You get one SVM per FSx for ONTAP file system by default. You can create additional SVMs at any time. We recommend a separate SVM for each volume. You don't need to join your file system to Active Directory for the database. For more information, see [Managing FSx for ONTAP storage virtual machines](#).

# Volume configuration and layout

Before you create an FSx for ONTAP file system, determine the total storage space you need for your system. You can increase the storage size later. To decrease the storage size, you must create a new file system.

The storage capacity of your file system should align with the needs of your database's system volumes. You must also consider the capacity required for snapshots, if applicable.

The following lists the example volume layout for databases.

**IBM Db2**

| Volume name | Junction name | Linux mount points |
| --- | --- | --- |
| <SID>-usrsap | /<SID>-usrsap | /usr/sap |
| <SID>-sapmnt | /<SID>-sapmnt | /sapmnt |
| <DBSID>-db2dbsid | /<DBSID>-db2dbsid | /db2/db2<db2sid> |
| <DBSID>-dbpath | /<DBSID>-dbpath | /db2/<DBSID> |
| <DBSID>-sapdata1 | /<DBSID>-sapdata1 | /db2/<DBSID>/sapdata1 |
| <DBSID>-sapdata<x> | /<DBSID>-sapdata<x> | /db2/<DBSID>/sapdata<x> |
| <DBSID>-saptmp1 | /<DBSID>-saptmp1 | /db2/<DBSID>/saptmp1 |
| <DBSID>-saptmp<x> | /<DBSID>-saptmp<x> | /db2/<DBSID>/saptmp<x> |
| <DBSID>-logdir | /<DBSID>-logdir | /db2/<DBSID>/log_dir |
| <DBSID>-logarch | /<DBSID>-logarch | /db2/<DBSID>/log_arch |

| <DBSID>-db2dump | /<DBSID>-db2dump | /db2/<DBSID>/db2dump |
|---|---|---|
| <DBSID>-backup | /<DBSID>-backup | /db2backup |

## SAP MaxDB

| Volume name | Junction name | Linux mount points |
|---|---|---|
| <DBSID>-sapmnt | <DBSID>-sapmnt | /sapmnt/ |
| <DBSID>-usrsap | <DBSID>-usrsap | /usr/sap |
| <DBSID>-sapdata | <DBSID>-sapdata | /sapdb/<DBSID>/sapdata |
| <DBSID>-saplog | <DBSID>-saplog | /sapdb/<DBSID>/saplog |
| <DBSID>-backup | <DBSID>-backup | /sapdb/<DBSID>/backup |
| <DBSID>-sapdb | <DBSID>-sapdb | /sapdb |

## SAP ASE

| Volume name | Junction name | Linux mount points |
|---|---|---|
| <SID>-sapmnt | <SID>-sapmnt | /sapmnt/ |
| <SID>-usrsap | <SID>-usrsap | /usr/sap |
| <SID>-sysbase | <SID>-sysbase | /sysbase |
| <SID>-sapdata_1 | <SID>-sapdata_1 | /sysbase/<SID>/sapdata_1 |
| <SID>-saplog_1 | <SID>-saplog_1 | /sysbase/<SID>/saplog_1 |
| <SID>-sapdiag | <SID>-sapdiag | /sysbase/<SID>/sapdiag |
| <SID>-saptmp | <SID>-saptmp | /sysbase/<SID>/saptmp |
| <SID>-backup | <SID>-backup | /sysbasebackup |

## Oracle

| Volume name | Junction name | Linux mount points |
|---|---|---|
| <DBSID>-oracle | /<DBSID>-oracle | /oracle |
| <DBSID>-oracle-<DBSID> | /<DBSID>-oracle-<DBSID> | /oracle/<DBSID> |
| <DBSID>-oraarch | /<DBSID>-oraarch | /oracle/<DBSID>/oraarch |
| <DBSID>-saparch | /<DBSID>-saparch | /oracle/<DBSID>/saparch |
| <DBSID>-sapreorg | /<DBSID>-sapreorg | /oracle/<DBSID>/sapreorg |
| <DBSID>-sapbackup | /<DBSID>-sapbackup | /oracle/<DBSID>/sapbackup |
| <DBSID>-saptrace | /<DBSID>-saptrace | /oracle/<DBSID>/saptrace |
| <DBSID>-sapdata1 | /<DBSID>-sapdata1 | /oracle/<DBSID>/sapdata1 |
| <DBSID>-sapdata<x> | /<DBSID>-sapdata<x> | /oracle/<DBSID>/sapdata<x> |
| <DBSID>-origlogA | /<DBSID>-origlogA | /oracle/<DBSID>/origlogA |
| <DBSID>-origlogB | /<DBSID>-origlogB | /oracle/<DBSID>/origlogB |
| <DBSID>-mirrlogB | /<DBSID>-mirrlogB | /oracle/<DBSID>/mirrlogB |
| <DBSID>-mirrlogA | /<DBSID>-mirrlogA | /oracle/<DBSID>/mirrlogA |

## MSSQL

| Volume name | Directory | Description |
|---|---|---|
| <SID>-sapdata | <drive>:\<SAPSID>DATA0<br><br><drive>:\<SAPSID>DATA1<br><br><drive>:\<SAPSID>DATA<N> | Directory for SAP database data files |

| <SID>-saplog | <drive>:\<SAPSID>log<N> | Directory for SAP database transaction log |
| <SID>-tempdb | <drive>:\Tempdb | Directory for temporary database data files |
| <SID>-sqldbexe | <drive>:\Program Files\Microsoft SQL Server | Directory for SQL server program files, and `master`, `msdb`, and `model` data files |
| <SID>-backup | <drive>:\Backup | SQL server data and log backup directory |

It is recommended to switch on Storage Efficiency, but to keep inactive data compression disabled.

# File system setup

To create a FSx for ONTAP file system, see Step 1: Create an Amazon FSx for NetApp ONTAP file system. For more information, see Managing FSx for ONTAP file systems.

After creating a FSx for ONTAP file system, you must complete additional file system setup.

**Topics**

- Set administrative password
- Sign in to the management endpoint via SSH
- Set TCP max transfer size
- Disable snapshots
- Configuration settings for dynamically allocating storage – MSSQL

## Set administrative password

If you did not create an administrative password during FSx for ONTAP file system creation, you must set an ONTAP administrative password for `fsxadmin` user.

The administrative password enables you to access the file system via SSH, the ONTAP CLI, and REST API. To use tools like NetApp SnapCenter, you must have an administrative password.

# Sign in to the management endpoint via SSH

Get the DNS name of the management endpoint from Amazon console. Sign in to the management endpoint via SSH, using the `fsxadmin` user and administrative password.

```
ssh fsxadmin@management.<file-system-id>.fsx.<aws-region>.amazonaws.com Password:
```

# Set TCP max transfer size

We recommend a TCP max transfer size of 262,144 for your database systems. Elevate the privilege level to *advanced* and use the following command on each SVM.

```
set advanced
nfs modify -vserver <svm> -tcp-max-xfer-size 262144
set admin
```

# Disable snapshots

FSx for ONTAP automatically enables a snapshot policy for volumes that take hourly snapshots. The default policy offers limited value due to missing application awareness. We recommend disabling the automatic snapshots by setting the policy to none.

```
volume modify -vserver <vserver-name> -volume <volume-name> -snapshot-policy none
```

You can use SnapCenter, a backup management software offered by NetApp to automate backup and restore of your workloads. SnapCenter provides a plug-in for Oracle and MSSQL databases.

# Configuration settings for dynamically allocating storage – MSSQL

This section only applies to MSSQL database.

NetApp recommends that storage space be dynamically allocated to each volume or logical unit number as data is written, instead of allocating space upfront. The following table provides the configuration settings for dynamically allocating storage.

| Setting | Configuration |
| --- | --- |
| Volume guarantee | None (set by default) |

| Setting | Configuration |
|---|---|
| LUN reservation | Enabled |
| fractional_reserve | 0% (set by default) |
| snap reserve | 0% |
| Autodelete | volume/oldest_first |
| Autosize | On |
| try_first | Autogrow |
| Volume tiering policy | Snapshot only |
| Snapshot policy | None |
| space allocation | Enabled |
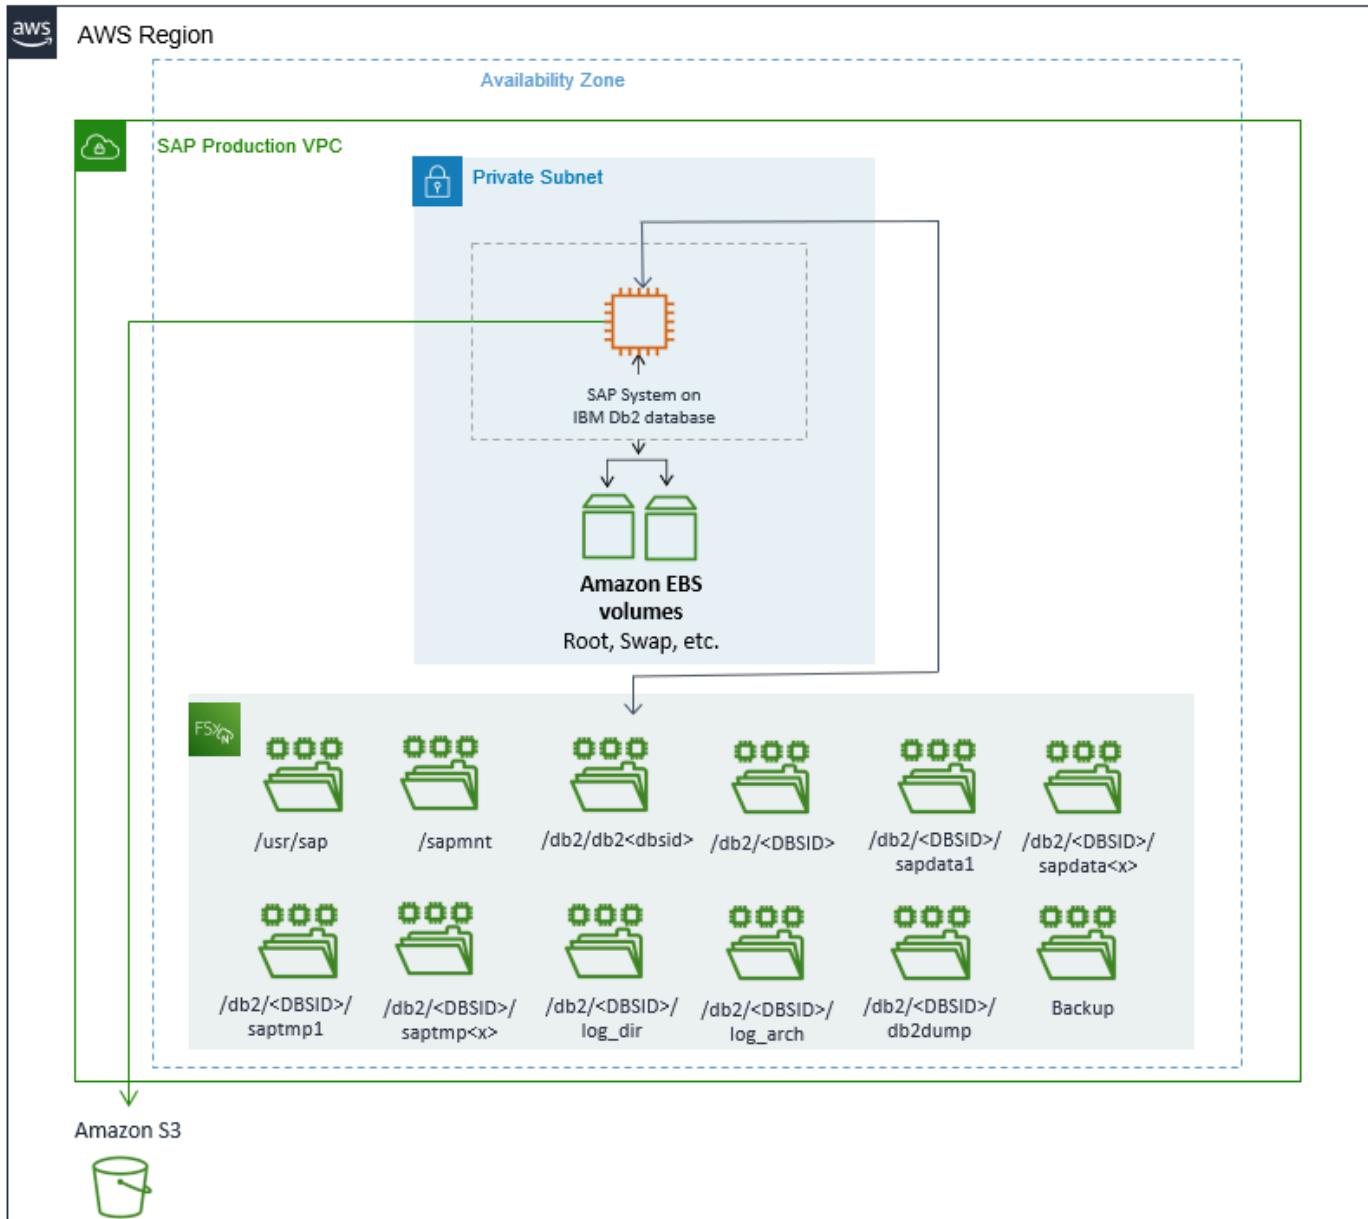
# Architecture diagrams for databases with Amazon FSx for NetApp ONTAP

See the following tabs for the architecture diagram of each database.

**Example**
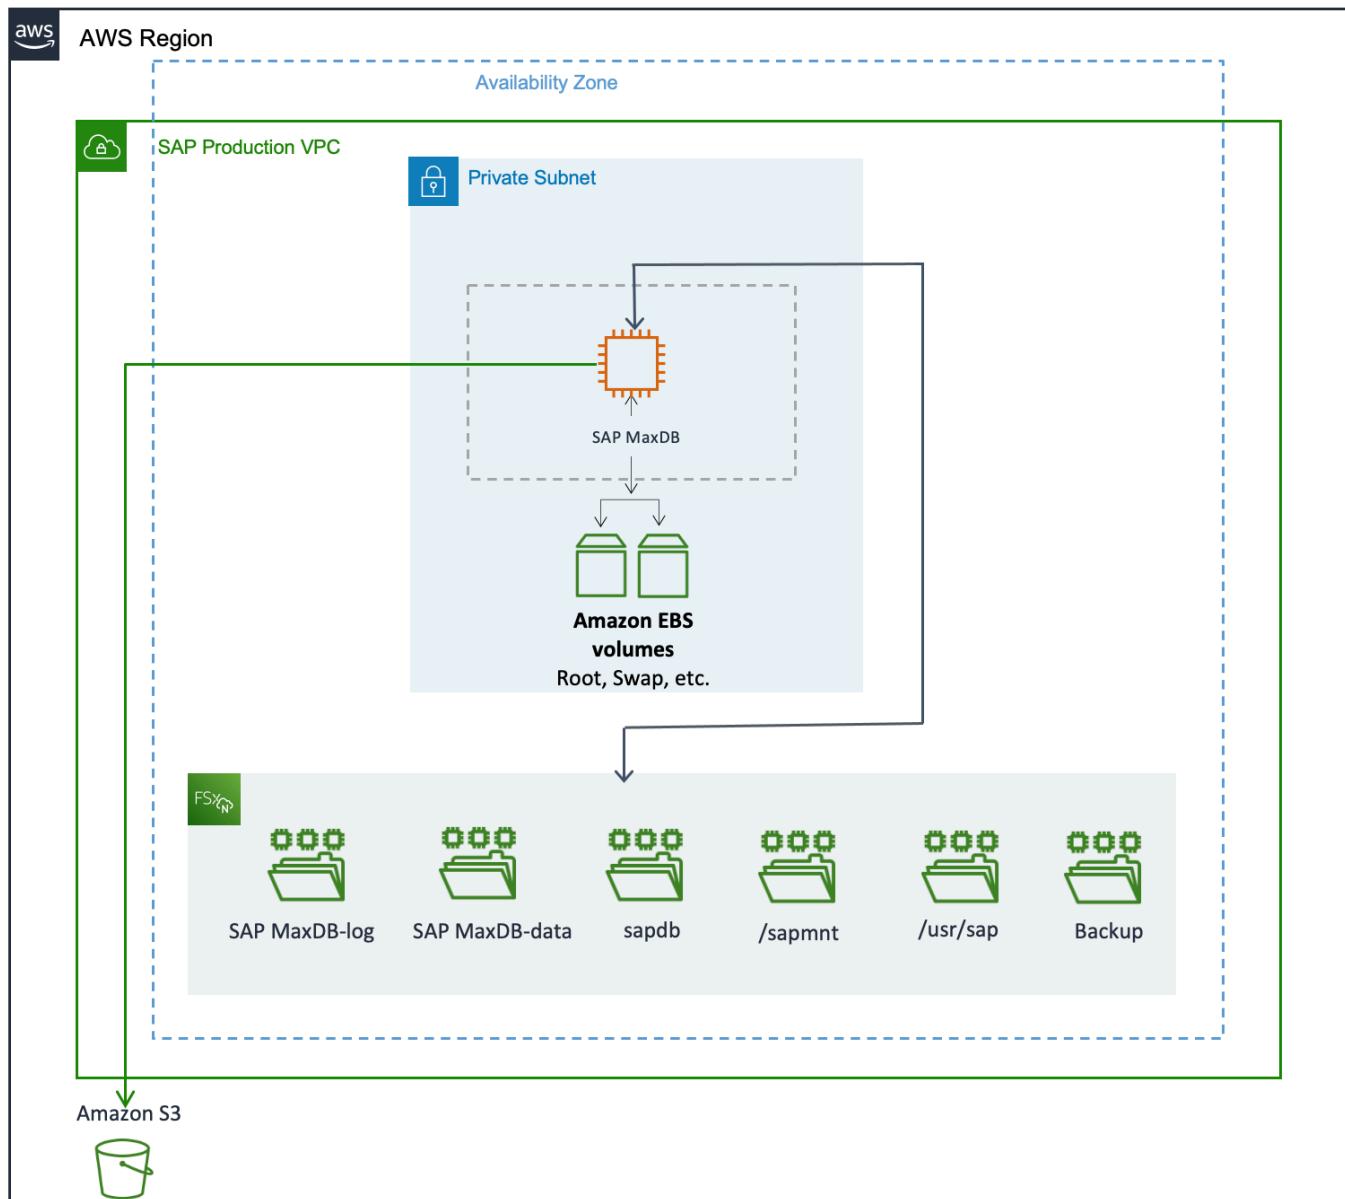
IBM Db2

   The following diagram presents the setup for IBM Db2 system with FSx for ONTAP.
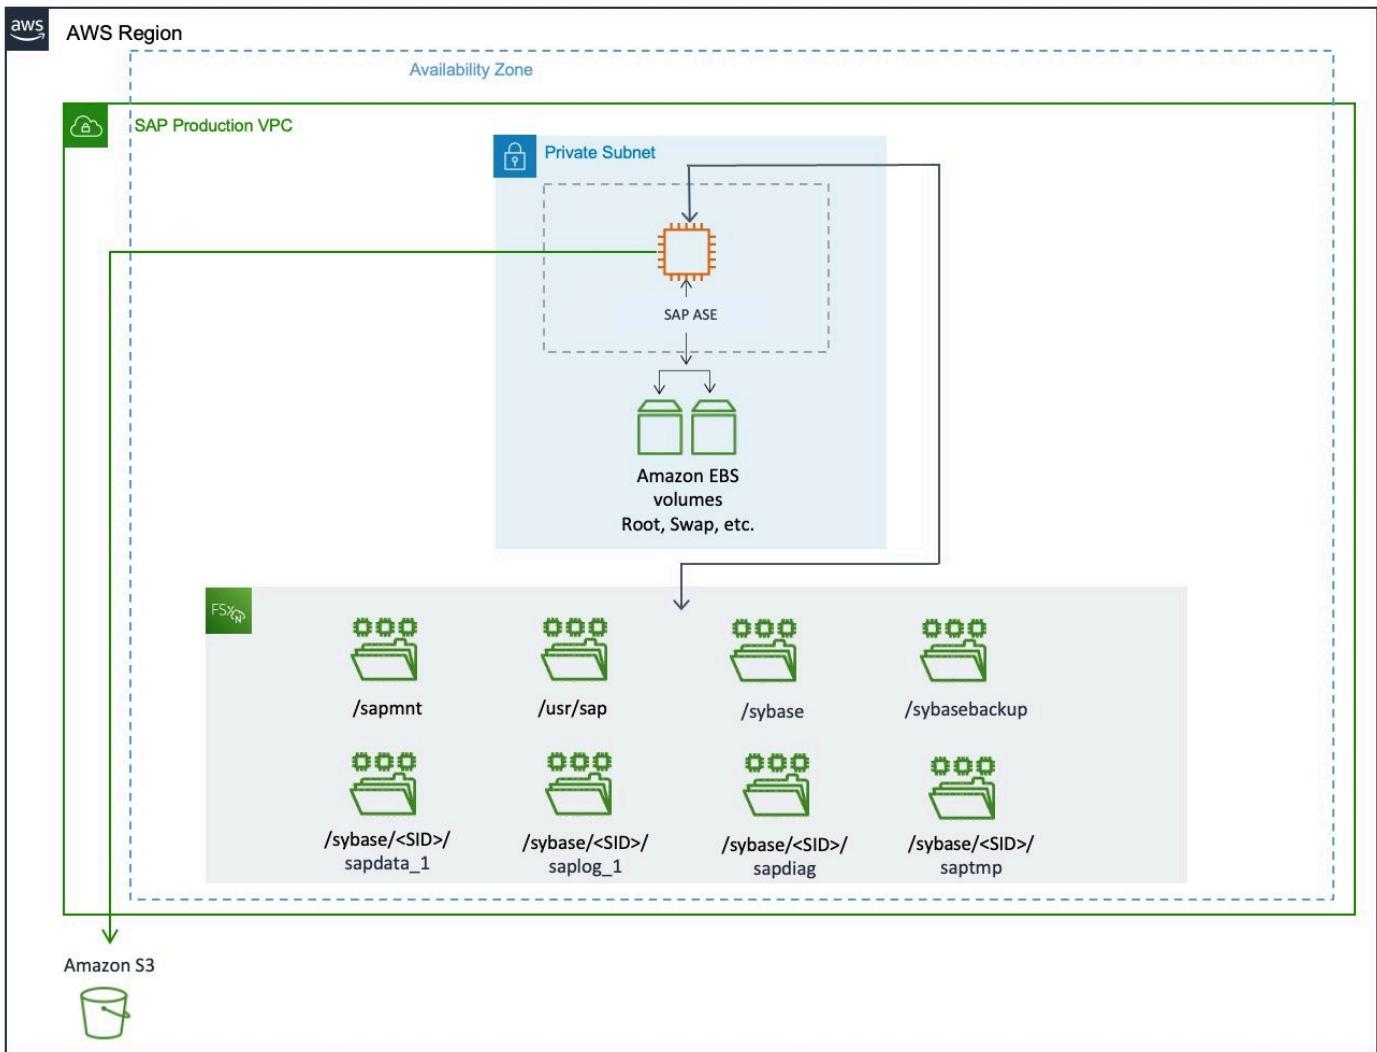
## SAP MaxDB

The following diagram presents the setup for SAP MaxDB system with FSx for ONTAP.
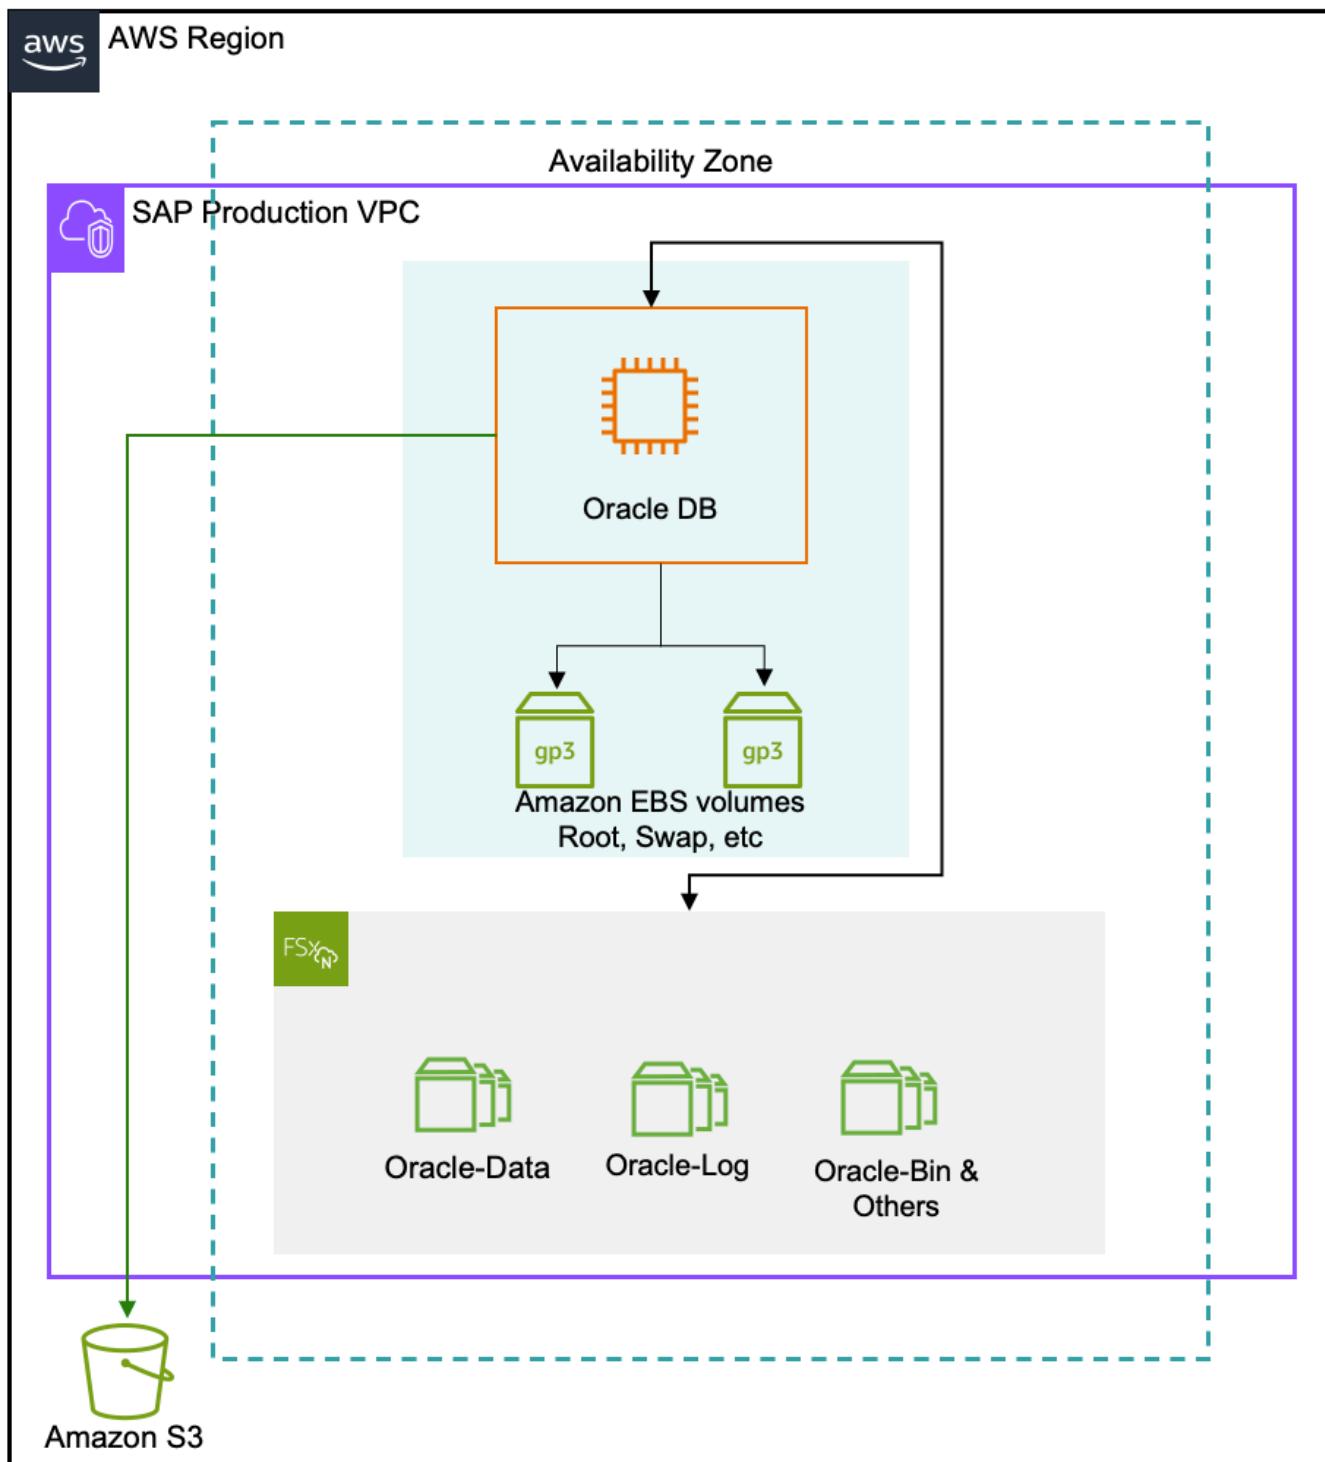
## SAP ASE

The following diagram presents the setup for SAP ASE system with FSx for ONTAP.
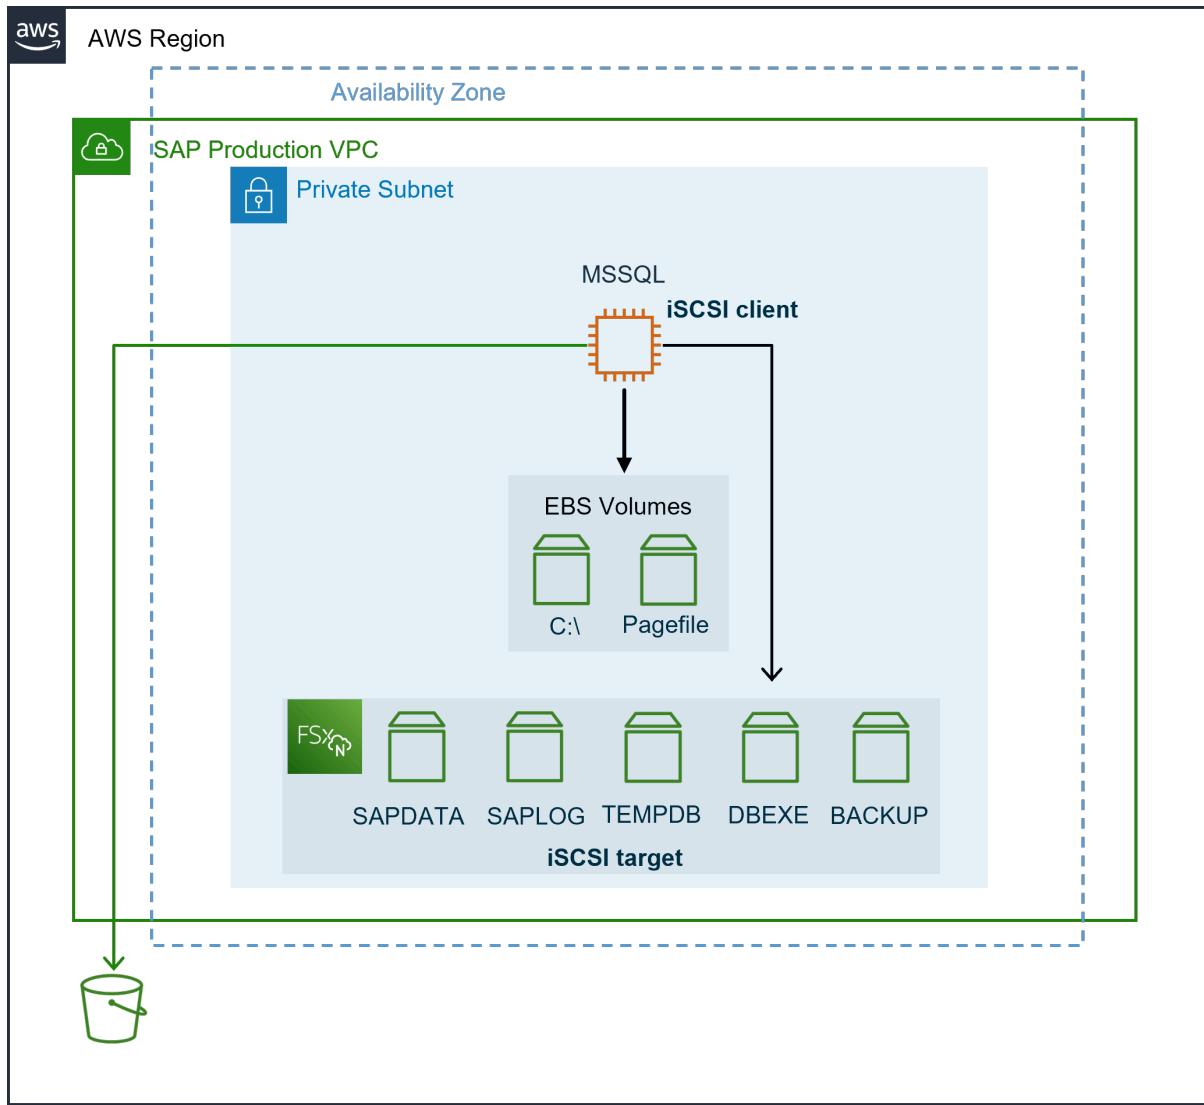
## Oracle

The following diagram presents the setup for Oracle with FSx for ONTAP.
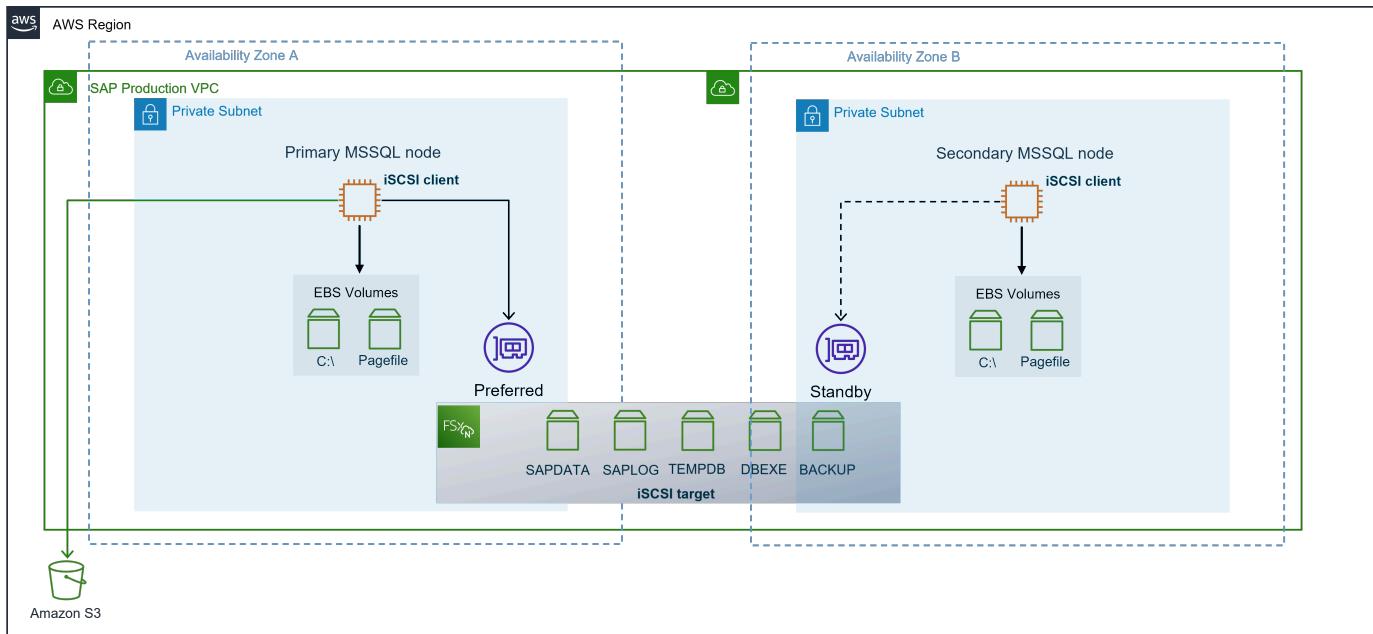
MSSQL

The following diagram presents the single Availability Zone setup for MSSQL database with FSx for ONTAP.

The following diagram presents a high availability setup for MSSQL database with FSx for ONTAP.

FSx for ONTAP supports both, iSCSI and SMB protocol to be used for SQL server deployments on Amazon. The iSCSI protocol works at the block level, and is expected to drive higher performance for SQL server with OLTP-type workloads than the same system configured with SMB. We recommend configuring your MSSQL on FSx for ONTAP using the iSCSI protocol.

FSx for ONTAP file systems are set up redundant by default. Each file system has a preferred (active) and a standby (passive) file server. FSx for ONTAP file systems provide management and protocol specific endpoints for each file server either within an Availability Zone (single-AZ) or across Availability Zones (Multi-AZ). For more information, see Availability, durability, and deployment options.

The SQL server FCI nodes access your FSx for ONTAP file system through elastic network interfaces (ENI). These network interfaces reside in Amazon VPC that you associate with your file system. Clients access the FSx for ONTAP file system via these ENIs (Preferred and Standby).

As a good practice, the active SQL server FCI node should be in the same subnet as the FSx for ONTAP file system preferred subnet. This enables best throughput and low latency, avoiding unnecessary inter-Availability Zone network traffic.

# Host setup for databases with Amazon FSx for NetApp ONTAP

This section walks you through an example host setup for deploying a database system on Amazon using Amazon FSx for NetApp ONTAP as the primary storage solution.

**Topics**

- [Linux kernel parameters](#)
- [Network File System (NFS)](#)
- [Create mount points](#)
- [Mount file systems](#)
- [Host setup for MSSQL](#)

# Linux kernel parameters

See the following tabs for details.

**Example**

IBM Db2

Use the following steps to setup the Linux kernel parameters.

1. Create a file named `91-NetApp-DB2.conf` with the following configurations in the `/etc/sysctl.d` directory.

   ```
   net.core.rmem_max = 16777216
   net.core.wmem_max = 16777216
   net.ipv4.tcp_rmem = 4096 131072 16777216
   net.ipv4.tcp_wmem = 4096 16384 16777216
   net.core.netdev_max_backlog = 300000
   net.ipv4.tcp_slow_start_after_idle=0
   net.ipv4.tcp_no_metrics_save = 1
   net.ipv4.tcp_moderate_rcvbuf = 1
   net.ipv4.tcp_window_scaling = 1
   net.ipv4.tcp_timestamps = 1
   net.ipv4.tcp_sack = 1
   ```

2. Increase the max sessions slots for NFSv4 to 180.

   ```
   echo options nfs max_session_slots=180 > /etc/modprobe.d/nfsclient.conf
   ```

3. Set `sunrpc.tcp_slot_table_entries = 128` in `/etc/sysctl.conf`.

4. Reboot your instance for the kernel parameters and NFS settings to take effect.

SAP MaxDB

Use the following steps to setup the Linux kernel parameters.

1. Create a file named `91-NetApp-MaxDB.conf` with the following configurations in the `/etc/sysctl.d` directory.

```
net.core.rmem_max = 16777216
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384  16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

2. Increase the max sessions slots for NFSv4 to 180.

```
echo options nfs max_session_slots=180 > /etc/modprobe.d/nfsclient.conf
```

3. Reboot your instance for the kernel parameters and NFS settings to take effect.

SAP ASE

Use the following steps to setup the Linux kernel parameters.

1. Create a file named `91-NetApp-ASE.conf` with the following configurations in the `/etc/sysctl.d` directory.

```
net.core.rmem_max = 16777216
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384  16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
```

```
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
```

2. Increase the max sessions slots for NFSv4 to 180.

```
echo options nfs max_session_slots=180 > /etc/modprobe.d/nfsclient.conf
```

3. Reboot your instance for the kernel parameters and NFS settings to take effect.

Oracle

Use the following steps to setup the Linux kernel parameters.

1. Take a backup of `/etc/sysctl.conf` file by using the following command.

```
sudo cp /etc/sysctl.conf /etc/sysctl.conf.bak
```

2. Open the `/etc/sysctl.conf` file in a text editor, add the following entries to it, and save the file.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.core.optmem_max = 16777216
net.core.somaxconn = 4096
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_slow_start_after_idle=0
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384  16777216
net.ipv4.tcp_max_syn_backlog = 16348
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_dsack = 1
net.ipv4.tcp_sack = 1
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_syn_retries = 8
```

3. Increase the max sessions slots for NFSv4 to 180.

```
echo options nfs max_session_slots=180 > /etc/modprobe.d/nfsclient.conf
```

4. Reboot your instance for the kernel parameters and NFS settings to take effect.

# Network File System (NFS)

Network File System (NFS) version 4 and higher requires user authentication. You can authenticate with Lightweight Directory Access Protocol (LDAP) server or local user accounts.

If you are using local user accounts, the NFSv4 domain must be set to the same value on all Linux servers and SVMs. Set the following parameters in the `/etc/idmapd.conf` file on the Linux hosts.

```
Domain = <domain name>
Nobody-User = root
Nobody-Group = root
```

> ⓘ **Note**
>
> You must restart the `nfs-idmapd.service` service after making changes to the domain.

**Oracle Direct Network File System (dNFS)**

SAP only supports deployment of Oracle database on Amazon FSx for NetApp ONTAP using Direct NFS with NFSv3, NFSv4 and NFSv4.1.

You must setup the `oranfstab` file before enabling dNFS client control of NFS. For more details, refer [Deploying Oracle Direct NFS](). On configuration completion, query the V$DNFS_SERVERS view to verify that the server name, NFS mount points, and NFS version shown in the output match the configuration in `oranfstab` file.

# Create mount points

See the following for details.

**Example**

IBM Db2

Create the following mount points on your Amazon EC2 instance.

```
mkdir -p /sapmnt
mkdir -p /usr/sap
mkdir -p /db2/db2<dbsid>
mkdir -p /db2/<DBSID>
mkdir -p /db2/<DBSID>/sapdata1
mkdir -p /db2/<DBSID>/sapdata<x>
mkdir -p /db2/<DBSID>/saptmp1
mkdir -p /db2/<DBSID>/saptmp<x>
mkdir -p /db2/<DBSID>/db2dump
mkdir -p /db2/<DBSID>/log_dir
mkdir -p /db2/<DBSID>/log_arch
mkdir -p /db2backup
```

## SAP MaxDB

Create the following mount points on your Amazon EC2 instance.

```
mkdir -p /sapmnt
mkdir -p /usr/sap
mkdir -p /sapdb/<DBSID>/sapdata
mkdir -p /sapdb/<DBSID>/saplog
mkdir -p /backup
mkdir -p /sapdb
```

## SAP ASE

Create the following mount points on your Amazon EC2 instance.

```
mkdir -p /sapmnt
mkdir -p /usr/sap
mkdir -p /sybase
mkdir -p /sybase/<SID>/sapdata_1
mkdir -p /sybase/<SID>/saplog_1
mkdir -p /sybase/<SID>/sapdiag
mkdir -p /sybase/<SID>/saptmp
mkdir -p /sybasebackup
```

## Oracle

Create the following mount points on your Amazon EC2 instance.

```
mkdir -p /oracle
```

```
mkdir -p /oracle/<DBSID>
mkdir -p /oracle/<DBSID>/sapdata1
mkdir -p /oracle/<DBSID>/sapdata2
mkdir -p /oracle/<DBSID>/sapdata3
mkdir -p /oracle/<DBSID>/sapdata4
mkdir -p /oracle/<DBSID>/origlogA
mkdir -p /oracle/<DBSID>/origlogB
mkdir -p /oracle/<DBSID>/mirrlogA
mkdir -p /oracle/<DBSID>/mirrlogB
mkdir -p /oracle/<DBSID>/sapreorg
mkdir -p /oracle/<DBSID>/saptrace
mkdir -p /oracle/<DBSID>/saparch
mkdir -p /oracle/<DBSID>/sapcheck
mkdir -p /oracle/<DBSID>/oraarch
mkdir -p /oracle/<DBSID>/sapbackup
```

# Mount file systems

The created file systems must be mounted as NFS file systems on Amazon EC2. See the following for details.

The following table is an example recommendation of NFS options for different IBM Db2 file systems.

**Example**

IBM Db2

| File systems | NFS mount options | | | |
|---|---|---|---|---|
| | **Common** | **NFS version** | **NFS transfer size** | **nconnect** |
| Db2 data | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=8 |
| Db2 log | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |

| Backup | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |
|---|---|---|---|---|
| Db2 binary | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |

**Example**

Add the following lines to `/etc/fstab` to preserve mounted file systems during an instance reboot. You can then run `mount -a` to mount the NFS file systems.

```
<SVM NFSIP>:/<SID>-sapmnt /sapmnt nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<SID>-usrsap /usr/sap nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-db2dbsid /db2/db2<dbsid> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-dbpath /db2/<DBSID> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-sapdata1 /db2/<DBSID>/sapdata1 nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-sapdata<x> /db2/<DBSID>/sapdata<x> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-saptmp1 /db2/<DBSID>/saptmp1 nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-saptmp<x> /db2/<DBSID>/saptmp<x> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-logdir /db2/<DBSID>/log_dir nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-logarch /db2/<DBSID>/log_arch nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-db2dump /db2/<DBSID>/db2dump nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFSIP>:/<DBSID>-backup /db2backup nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

The following table is an example recommendation of NFS options for different SAP MaxDB file systems.

SAP MaxDB

| File systems | NFS mount options | | | |
|---|---|---|---|---|
| | **Common** | **NFS version** | **NFS transfer size** | **nconnect** |
| SAP MaxDB data | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=8 |
| SAP MaxDB log | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |
| Backup | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |
| SAP MaxDB binary | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |

**Example**

Add the following lines to `/etc/fstab` to preserve mounted file systems during an instance reboot. You can then run `mount  -a` to mount the NFS file systems.

```
<BackupSVM NFSIP>:/<DBSID>_sapmnt /sapmnt nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<BackupSVM NFSIP>:/<DBSID>_usrsap /usr/sap nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<BackupSVM NFSIP>:/<DBSID>_backup /backup nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SAPLOGSVM NFSIP>:/<DBSID>_sapdb /sapdb nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SAPLOGSVM NFSIP>:/<DBSID>_log /sapdb/<DBSID>/saplog nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

```
<SAPLOGSVM NFSIP>:/<DBSID>_sapdata /sapdb/<DBSID>/sapdata nfs
  rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

The following table is an example recommendation of NFS options for different SAP ASE file systems.

SAP ASE

| File systems | NFS mount options | | | |
|---|---|---|---|---|
| | **Common** | **NFS version** | **NFS transfer size** | **nconnect** |
| SAP ASE data | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=8 |
| SAP ASE log | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |
| Backup | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |
| SAP ASE binary | rw,bg,har d,timeo=6 00,noatime, | vers=4,mi norversio n=1,lock, | rsize=262 144,wsize =262144, | nconnect=2 |

**Example**

Add the following lines to `/etc/fstab` to preserve mounted file systems during an instance reboot. You can then run `mount -a` to mount the NFS file systems.

```
<SVM-ASEDBData NFSIP>:/<SID>-sapdata_1 /sybase/<SID>/sapdata_1 nfs
  rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBLog NFSIP>:/<SID>-saplog_1 /sybase/<SID>/saplog_1 nfs
  rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBShared NFSIP>:/<SID>-sapmnt /sapmnt nfs
  rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

```
<SVM-ASEDBShared NFSIP>:/<SID>-usrsap /usr/sap nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBShared NFSIP>:/<SID>-sybase /sybase nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBShared NFSIP>:/<SID>-sapdiag /sybase/<SID>/sapdiag nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBShared NFSIP>:/<SID>-saptmp /sybase/<SID>/saptmp nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM-ASEDBBackup NFSIP>:/<SID>-backup /sybasebackup nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

You must assign a virtual hostname to each SVM NFS IP address in `/etc/hosts` file. These hostnames are used in the `oranfstab` file for values of the NFS server (`server: `) during Oracle dNFS configuration.

The following table is an example recommendation of NFS options for different SAP Oracle file systems.

Oracle

| File systems | NFS mount options | | | |
|---|---|---|---|---|
| | **Common** | **NFS version** | **NFS transfer size** | **nconnect** |
| Oracle data | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=8 |
| Oracle log | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |
| Oracle backup | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |
| Oracle binary | rw,bg,hard,timeo=600,noatime, | vers=4,minorversion=1,lock, | rsize=262144,wsize=262144, | nconnect=2 |

**Example**

Add the following lines to `/etc/fstab` to preserve mounted file systems during an instance reboot. You can then run `mount -a` to mount the NFS file systems.

```
<SVM NFS Host>:/<DBSID>-oracle /oracle nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-oracle<DBSID> /oracle/<DBSID> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-oraarch /oracle/<DBSID>/oraarch nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-saparch /oracle/<DBSID>/saparch nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-sapreorg /oracle/<DBSID>/sapreorg nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-sapbackup /oracle/<DBSID>/sapbackup nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-saptrace /oracle/<DBSID>/saptrace nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-sapdata1 /oracle/<DBSID>/sapdata1 nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-sapdata2 /oracle/<DBSID>/sapdata<x> nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-origlogA /oracle/<DBSID>/origlogA nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-origlogB /oracle/<DBSID>/origlogB nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-mirrlogB /oracle/<DBSID>/mirrlogB nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
<SVM NFS Host>:/<DBSID>-mirrlogA /oracle/<DBSID>/mirrlogA nfs
 rw,bg,hard,timeo=600,noatime,vers=4,minorversion=1,lock,rsize=262144,wsize=262144,nconnect=
```

<SVM NFS Host> are the values that you have assigned in the `/etc/hosts` file for each SVM NFS IP address.

- Changes to the `nconnect` parameter take effect only if the NFS file system is unmounted and mounted again.
- Client systems must have unique host names when accessing FSx for ONTAP. If there are systems with the same name, the second system may not be able to access FSx for ONTAP.
- For RHEL operating system, the `nconnect` parameter is supported only on RHEL 8.3 or higher.

# Host setup for MSSQL

Use the following procedure to create volumes and LUNs for MSSQL server.

1. Connect to your FSx for ONTAP system using SSH. For more information, see Managing FSx for ONTAP resources using NetApp applications.

2. Use the following SSH command to create the required volumes:

```
vol create -vserver <SvmName> -volume <VolumeName> -aggregate aggr1 -size
 <VolumeSize> -state online -tiering-policy snapshot-only -percent-snapshot-space 0 -
 autosize-mode grow -snapshot-policy none -security-style ntfs
volume modify -vserver <SvmName> -volume <VolumeName> -fractional-reserve 0
volume modify -vserver <SvmName> -volume <VolumeName> -space-mgmt-try-first vol_grow
volume snapshot autodelete modify -vserver <SvmName> -volume <VolumeName> -delete-
order oldest_first
```

3. Start the `iSCSI` service with PowerShell using elevated privileges in Windows servers, and set the startup type to *Automatic*.

```
Start-service -Name msiscsi
Set-Service -Name msiscsi -StartupType Automatic
```

4. Install Multipath-IO with PowerShell using elevated privileges in Windows servers.

```
Install-WindowsFeature -name Multipath-IO -Restart
```

5. Find the Windows initiator name with PowerShell using elevated privileges in Windows servers.

```
Get-InitiatorPort | select NodeAddress
```

6. Connect to Storage virtual machines (SVM) using putty, and create an `iGroup`.

```
igroup create -igroup <iGroupName> -protocol iscsi -ostype windows -initiator
 <InitiatorName>
```

7. Use the following SSH command to create LUNs inside each volume. To achieve I/O alignment with the operating system partitioning scheme, use `windows_2008` as the recommended LUN type.

```
lun create -path /vol/<VolumeName>/<LUNName> -size <LUNSize> -ostype windows_2008 -
space-allocation enabled
```

8. Use the following SSH command to map the `igroup` to the LUNs that you just created:

```
lun show
lun map -path <LUNPath> -igroup <iGroupName>
```

9. Get the iSCSI endpoint IP addresses (preferred subnet and standby subnet) from the FSx for ONTAP console. Choose your SVM on the Storage Virtual Machines page.

   The iSCSI endpoint IP addresses are used as iSCSI targets in the next step. For more information, see Provisioning iSCSI for Windows. .. On the Windows server, go to iSCSI initiator settings, and connect to your iSCSI targets (FSx for ONTAP iSCSI endpoints). Go to **Discovery** > **Discover Portal**. Enter the iSCSI IP address from previous step, and select **Advanced**. From Local Adapter, select **Microsoft iSCSI Initiator**. From Initiator IP, select **IP of the server**. .. From iSCSI Initiator Settings, select **Targets** and choose **Connect** and **Enable muti-path**. .. For best performance, add more sessions. NetApp recommends creating five iSCSI sessions. Select **Properties** > **Add session** > **Advanced**, and repeat the previous step. For further details, see SQL Server on Amazon EC2 using Amazon FSx for NetApp ONTAP.

10 Initialize disks with the following PowerShell command:

```
$disks = Get-Disk | where PartitionStyle -eq raw
foreach ($disk in $disks) {Initialize-Disk $disk.Number}
```

11 Run the partition and format commands with PowerShell.

```
New-Partition -DiskNumber 1 -DriveLetter F -UseMaximumSize
Format-Volume -DriveLetter F -FileSystem NTFS -AllocationUnitSize 65536
```


LUNs can also be created using SnapCenter.

You can also automate volume and LUN creation using the PowerShell script provided in the Appendix B of SQL Server on Amazon EC2 using Amazon FSx for NetApp ONTAP.

# Installing the databases

See the following tabs for more information.

**Example**

IBM Db2

You must install the IBM Db2 system as per the instructions provided by SAP. For more information, see Installation of SAP Systems Based on the Application Server ABAP of SAP NetWeaver 7.3 EHP1 to 7.52 on UNIX: IBM Db2 for Linux, UNIX, and Windows.

SAP MaxDB

You must install the SAP MaxDB as per the instructions provided by SAP. For more information, see Installation of SAP Systems Based on the Application Server ABAP of SAP NetWeaver 7.3 EHP1 to 7.52 on UNIX: SAP MaxDB.

**Backup and restore**

Backup and restore operations are supported by standard SAP MaxDB tools. For more information, see SAP Note 1928060 - Data backup and recovery with file system backup(requires access to the SAP portal).

SAP ASE

You must install the SAP ASE as per the instructions provided by SAP. You can select the relevant guide from the Guide Finder for SAP NetWeaver and ABAP Platform on SAP website.

**Backup and restore**

FSx for ONTAP snapshot is a read-only image of an FSx for ONTAP volume at a point-in-time. Snapshots offer protection against accidental deletion or modification of files in your volumes. Your users can easily view and/or restore individual files or folders from an earlier snapshot. For more information, see Working with snapshots.

Backup and restore operations are also supported by standard SAP ASE tools. You can check the following SAP Notes (requires SAP portal access) to learn more.

- SAP Note 1585981 - SYB: Ensuring Recoverability for SAP ASE
- SAP Note 1588316 - SYB: Configure automatic database and log backups
- SAP Note 1618817 - SYB: How to restore an SAP ASE database server (UNIX)
- SAP Note 1887068 - SYB: Using external backup and restore with SAP ASE

## Oracle

You must install the Oracle database as per the instructions provided by SAP. You can select the relevant guide from the Guide Finder for SAP NetWeaver and ABAP Platform on SAP website.

**Backup and restore**

FSx for ONTAP snapshot is a read-only image of an FSx for ONTAP volume at a point-in-time. Snapshots offer protection against accidental deletion or modification of files in your volumes. Your users can easily view and/or restore individual files or folders from an earlier snapshot. For more information, see Working with snapshots.

You can also use the plug-in for Oracle database offered by NetApp SnapCenter. The plug-in takes application-consistent backups using NetApp snapshots and Oracle Recovery Manager.

Backup and restore operations are also supported by standard Oracle database for SAP tools, such as BRTools. You can check the following resources from SAP and Oracle to learn more.

- SAP Database Guide: Oracle
- Oracle Database Backup To Cloud: Amazon Simple Storage Service (S3)
- SAP Note 2358420 - Oracle Database Support for Amazon Web Services EC2 (requires SAP portal access)
- SAP Note 1656250 - SAP on Amazon: Support prerequisites (requires SAP portal access)

## MSSQL

You must install the MSSQL database as per the instructions provided by SAP. You can select the relevant guide from the Guide Finder for SAP NetWeaver and ABAP Platform on SAP website.

**Backup and restore**

FSx for ONTAP snapshot is a read-only image of an FSx for ONTAP volume at a point-in-time. Snapshots offer protection against accidental deletion or modification of files in your volumes. Your users can easily view and/or restore individual files or folders from an earlier snapshot. For more information, see Working with snapshots.

Backup and restore operations are also supported by standard SQL server tools. For further details, see Back Up and Restore of SQL Server Databases.

For point-in-time resilient restores and immutable backups, we storing 3 days of snapshots on a local disk, and replicating older backups via SnapVault. Replicate older backups to a secondary (different Availability Zone) FSx for ONTAP filesystem with capacity pool enabled. For more information, see Managing storage capacity.

# SAP ASE for SAP NetWeaver on Amazon: high availability configuration for SUSE Linux Enterprise Server (SLES) for SAP applications

This topic applies to SUSE Linux Enterprise Server (SLES) operating system for SAP NetWeaver running SAP Adaptive Server Enterprise (ASE) database on Amazon cloud. It covers the instructions for configuration of a pacemaker cluster for SAP ASE database when deployed on Amazon EC2 instances in two different Availability Zones within an Amazon Region and FSx for ONTAP as the storage layer.

This topic covers the implementation of high availability using the cold standby method. For more information, see SAP Note 1650511 – SYB: High Availability Offerings with SAP Adaptive Server Enterprise (requires SAP portal access).

**Topics**

- Planning
- Architecture diagram
- Deployment
- Operations

# Planning

This section covers the following topics.

**Topics**

- Prerequisites
- Reliability
- SAP and SUSE references
- Concepts

## Prerequisites

You must meet the following prerequisites before commencing setup.

**Topics**

- [Deployed cluster infrastructure](#)
- [Supported operating system](#)
- [Required access for setup](#)

## Deployed cluster infrastructure

Ensure that your Amazon networking requirements and Amazon EC2 instances where SAP workloads are installed, are correctly configured for SAP. For more information, see [SAP NetWeaver Environment Setup for Linux on Amazon](#).

See the following SAP ASE pacemaker cluster specific requirements.

- Two cluster nodes created in private subnets in separate Availability Zones within the same Amazon VPC and Amazon Region

- Access to the route table(s) that are associated with the chosen subnets

  For more information, see [Overlay IP](#).

- Targeted Amazon EC2 instances must have connectivity to the Amazon EC2 endpoint via internet or a Amazon VPC endpoint.

## Supported operating system

Protecting SAP ASE database with a pacemaker cluster requires packages from SUSE, including targeted cluster resource agents for SAP and Amazon that may not be available in standard repositories.

For deploying SAP applications on SUSE, SAP and SUSE recommend using SUSE Linux Enterprise Server for SAP applications (SLES for SAP). SLES for SAP provides additional benefits, including Extended Service Pack Overlap Support (ESPOS), configuration and tuning packages for SAP applications, and High Availability Extensions (HAE). For more details, see SUSE website at [SUSE Linux Enterprise Server for SAP Applications](#).

SLES for SAP is available at [Amazon Marketplace](#) with an hourly or annual subscription. You can also use the bring your own subscription (BYOS) model.

## Required access for setup

The following access is required for setting up the cluster.

- An IAM user with the following privileges.

  - modify Amazon VPC route tables

  - modify Amazon EC2 instance properties

  - create IAM policies and roles

  - create Amazon EFS file systems

- Root access to the operating system of both cluster nodes

- SAP administrative user access – `<syb>adm`

  In case of a new install, this user is created by the install process.

# Reliability

The SAP Lens of the Well-Architected framework, in particular the Reliability pillar, can be used to understand the reliability requirements for your SAP workload.

SAP ASE is a single point of failure in a highly available SAP architecture. The impact of an outage of this component must be evaluated against factors, such as, recovery point objective (RPO), recovery time objective (RTO), cost and operation complexity. For more information, see Reliability in SAP Lens - Amazon Well-Architected Framework.

## SAP and SUSE references

In addition to this guide, see the following references for more details.

- SAP Note: 1650511 - SYB: High Availability Offerings with SAP Adaptive Server Enterprise
- SAP Note: 1656099 - SAP Applications on Amazon: Supported DB/OS and Amazon EC2 products
- SAP Note: 1984787 - SUSE Linux Enterprise Server 12: Installation Notes
- SAP Note: 2578899 - SUSE Linux Enterprise Server 15: Installation Notes
- SAP Note: 1275776 - Linux: Preparing SLES for SAP environments

You must have SAP portal access for reading all SAP Notes.

# Concepts

This section covers Amazon concepts.

**Topics**

- [Availability Zones](#)
- [Overlay IP](#)
- [Shared VPC](#)
- [Amazon FSx for NetApp ONTAP](#)
- [Pacemaker - STONITH fencing agent](#)

## Availability Zones

Availability Zone is one or more discreet data centers with redundant power, networking, and connectivity in an Amazon Region. For more information, see [Regions and Availability Zones](#).

For mission critical deployments of SAP on Amazon where the goal is to minimise the recovery time objective (RTO), we suggest distributing single points of failure across Availability Zones. Compared with single instance or single Availability Zone deployments, this increases resilience and isolation against a broad range of failure scenarios and issues, including natural disasters.

Each Availability Zone is physically separated by a meaningful distance (many kilometers) from another Availability Zone. All Availability Zones in an Amazon Region re interconnected with high-bandwidth, low-latency network, over fully redundant, dedicated metro fiber. This enables synchronous replication. All traffic between Availability Zones is encrypted.

## Overlay IP

Overlay IP enables a connection to the application, regardless of which Availability Zone (and subnet) contains the active primary node.

When deploying an Amazon EC2 instance in Amazon, IP addresses are allocated from the CIDR range of the allocated subnet. The subnet cannot span across multiple Availability Zones, and therefore the subnet IP addresses may be unavailable after faults, including network connectivity or hardware issues which require a failover to the replication target in a different Availability Zone.

To address this, we suggest that you configure an overlay IP, and use this in the connection parameters for the application. This IP address is a non-overlapping RFC1918 private IP address

from outside of VPC CIDR block and is configured as an entry in the route table or tables. The route directs the connection to the active node and is updated during a failover by the cluster software.

You can select any one of the following RFC1918 private IP addresses for your overlay IP address.

- 10.0.0.0 – 10.255.255.255 (10/8 prefix)
- 172.16.0.0 – 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 – 192.168.255.255 (192.168/16 prefix)

If, for example, you use the 10/8 prefix in your SAP VPC, selecting a 172 or a 192 IP address may help to differentiate the overlay IP. Consider the use of an IP Address Management (IPAM) tool such as Amazon VPC IP Address Manager to plan, track, and monitor IP addresses for your Amazon workloads. For more information, see What is IPAM?

The overlay IP agent in the cluster can also be configured to update multiple route tables which contain the Overlay IP entry if your subnet association or connectivity requires it.

**Access to overlay IP**

The overlay IP is outside of the range of the VPC, and therefore cannot be reached from locations that are not associated with the route table, including on-premises and other VPCs.

Use Amazon Transit Gateway as a central hub to facilitate the network connection to an overlay IP address from multiple locations, including Amazon VPCs, other Amazon Regions, and on-premises using Amazon Direct Connect or Amazon Client VPN.

If you do not have Amazon Transit Gateway set up as a network transit hub or if it is not available in your preferred Amazon Region, you can use a Network Load Balancer to enable network access to an overlay IP.

For more information, see SAP on Amazon High Availability with Overlay IP Address Routing.

## Shared VPC

An enterprise landing zone setup or security requirements may require the use of a separate cluster account to restrict the route table access required for the Overlay IP to an isolated account. For more information, see Share your VPC with other accounts.

Evaluate the operational impact against your security posture before setting up shared VPC. To set up, see Shared VPC – optional.

## Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed Amazon service.

FSx for ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx for ONTAP supports popular data security and anti-virus applications. For more information, see [What is Amazon FSx for NetApp ONTAP?](#)

## Pacemaker - STONITH fencing agent

In a two-node cluster setup for a primary resource and its replication pair, it is important that there is only one node in the primary role with the ability to modify your data. In the event of a failure scenario where a node is unresponsive or incommunicable, ensuring data consistency requires that the faulty node is isolated by powering it down before the cluster commences other actions, such as promoting a new primary. This arbitration is the role of the fencing agent.

Since a two-node cluster introduces the possibility of a fence race in which a dual shoot out can occur with communication failures resulting in both nodes simultaneously claiming, "I can't see you, so I am going to power you off". The fencing agent is designed to minimise this risk by providing an external witness.

SLES supports several fencing agents, including the one recommended for use with Amazon EC2 Instances (external/ec2). This resource uses API commands to check its own instance status - "Is my instance state anything other than running?" before proceeding to power off its pair. If it is already in a stopping or stopped state it will admit defeat and leave the surviving node untouched.

# Architecture diagram

The following diagram shows the cold standby SAP ASE cluster setup with FSx for ONTAP.

SLES for SAP – Pacemaker Cluster for SAP ASE

# Deployment

This section covers the following topics.

**Topics**

- [Settings and prerequisites](#)
- [SAP ASE and cluster setup](#)
- [Cluster configuration](#)

# Settings and prerequisites

The cluster setup uses parameters, including `DBSID` that is unique to your setup. It is useful to predetermine the values with the following examples and guidance.

**Topics**

- [Define reference parameters for setup](#)
- [Amazon EC2 instance settings](#)

- [Operating system prerequisites](#)
- [IP and hostname resolution prerequisites](#)
- [FSx for ONTAP prerequisites](#)
- [Shared VPC – optional](#)

## Define reference parameters for setup

The cluster setup relies on the following parameters.

**Topics**

- [Global Amazon parameters](#)
- [Amazon EC2 instance parameters](#)
- [SAP and Pacemaker resource parameters](#)
- [SLES cluster parameters](#)

### Global Amazon parameters

| Name | Parameter | Example |
|------|-----------|---------|
| Amazon account ID | `<account_id>` | `123456789100` |
| Amazon Region | `<region_id>` | `us-east-1` |

- Amazon account – For more details, see [Your Amazon account ID and its alias](#).
- Amazon Region – For more details, see [Describe your Regions](#).

### Amazon EC2 instance parameters

| Name | Parameter | Primary example | Secondary example |
|------|-----------|-----------------|-------------------|
| Amazon EC2 instance ID | `<instance_id>` | `i-xxxxins tidforhost1` | `i- xxxxinsti dforhost2` |
| Hostname | `<hostname>` | `slxdbhost01` | `slxdbhost02` |

| Name | Parameter | Primary example | Secondary example |
|------|-----------|-----------------|-------------------|
| Host IP | <host_ip> | 10.1.10.1 | 10.1.20.1 |
| Host additional IP | <host_add itional_ip> | 10.1.10.2 | 10.1.20.2 |
| Configured subnet | <subnet_id> | subnet-xx xxxxxxxxs ubnet1 | subnet-xx xxxxxxxxs ubnet2 |

- Hostname – Hostnames must comply with SAP requirements outlined in SAP Note 611361 - Hostnames of SAP ABAP Platform servers (requires SAP portal access).

  Run the following command on your instances to retrieve the hostname.

  ```
  hostname
  ```

- Amazon EC2 instance ID – run the following command (IMDSv2 compatible) on your instances to retrieve instance metadata.

  ```
  /usr/bin/curl --noproxy '*' -w "\n" -s -H "X-aws-ec2-metadata-token: $(curl --noproxy
    '*' -s -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-
  token-ttl-seconds: 21600")" http://169.254.169.254/latest/meta-data/instance-id
  ```

  For more details, see Retrieve instance metadata and Instance identity documents.

**SAP and Pacemaker resource parameters**

| Name | Parameter | Example |
|------|-----------|---------|
| DBSID | <DBSID> or <dbsid> | ASD |
| Virtual hostname | <db_virt_hostname> | slxvdb |
| Database Overlay IP | <ase_db_oip> | 172.16.0.29 |
| VPC Route Tables | <rtb_id> | rtb-xxxxxroutetable1 |

| Name | Parameter | Example |
|------|-----------|---------|
| FSx for ONTAP mount points | `<ase_db_fs>` | `svm-xxx.fs-xxx.fsx.us-east-1.amazonaws.com` |

- SAP details – SAP parameters must follow the guidance and limitations of SAP and Software Provisioning Manager. Refer to [SAP Note 1979280 - Reserved SAP System Identifiers (SAPSID) with Software Provisioning Manager](#) for more details.

  Post-installation, use the following command to find the details of the instances running on a host.

```
sudo /usr/sap/hostctrl/exe/saphostctrl -function ListDatabases
```

- Overlay IP – This value is defined by you. For more information, see [Overlay IP](#).
- FSx for ONTAP mount points – This value is defined by you. Consider the required mount points specified in [SAP ASE on Amazon with Amazon FSx for NetApp ONTAP](#).

### SLES cluster parameters

| Name | Parameter | Example |
|------|-----------|---------|
| Cluster user | `cluster_user` | `hacluster` |
| Cluster password | `cluster_password` | |
| Cluster tag | `cluster_tag` | `pacemaker` |
| Amazon CLI cluster profile | `aws_cli_cluster_profile` | `cluster` |

## Amazon EC2 instance settings

Amazon EC2 instance settings can be applied using Infrastructure as Code or manually using Amazon Command Line Interface or Amazon Management Console. We recommend Infrastructure as Code automation to reduce manual steps, and ensure consistency.

**Topics**

- [Create IAM roles and policies](#)

- [Amazon Overlay IP policy](#)

- [Assign IAM role](#)

- [Modify security groups for cluster communication](#)

- [Disable source/destination check](#)

- [Review automatic recovery and stop protection](#)

- [Create Amazon EC2 resource tags used by Amazon EC2 STONITH agent](#)


## Create IAM roles and policies

In addition to the permissions required for standard SAP operations, two IAM policies are required for the cluster to control Amazon resources on ASCS. These policies must be assigned to your Amazon EC2 instance using an IAM role. This enables Amazon EC2 instance, and therefore the cluster to call Amazon services.

Create these policies with least-privilege permissions, granting access to only the specific resources that are required within the cluster. For multiple clusters, you need to create multiple policies.

For more information, see {https---docs-aws-amazon-com-AmazonEC2-latest-UserGuide-iam-roles-for-amazon-ec2-html-ec2-instance-profile}[IAM roles for Amazon EC2].

### STONITH policy

The SLES STONITH resource agent (`external/ec2`) requires permission to start and stop both the nodes of the cluster. Create a policy as shown in the following example. Attach this policy to the IAM role assigned to both Amazon EC2 instances in the cluster.

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeInstances",
        "ec2:DescribeTags"
      ],
```

```
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances"
      ],
      "Resource": [
        "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
 east-1:123456789012:instance/i-1234567890abcdef0",
        "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
 east-1:123456789012:instance/i-1234567890abcdef0"
      ]
    }
  ]
}
```

**Amazon Overlay IP policy**

The SLES Overlay IP resource agent (`aws-vpc-move-ip`) requires permission to modify a routing entry in route tables. Create a policy as shown in the following example. Attach this policy to the IAM role assigned to both Amazon EC2 instances in the cluster.

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ec2:ReplaceRoute",
            "Resource": [
                "arn:aws:ec2:us-east-1:123456789012:route-table/
 rtb-0123456789abcdef0",
                "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0"
                    ]
        },
        {
            "Effect": "Allow",
            "Action": "ec2:DescribeRouteTables",
            "Resource": "*"
        }
    ]
}
```

> ⓘ **Note**
>
> If you are using a Shared VPC, see [Shared VPC cluster resources](#).

**Assign IAM role**

The two cluster resource IAM policies must be assigned to an IAM role associated with your Amazon EC2 instance. If an IAM role is not associated to your instance, create a new IAM role for cluster operations. To assign the role, go to https://console.aws.amazon.com/ec2/, select your instance(s), and then choose **Actions** > **Security** > **Modify IAM role**.

**Modify security groups for cluster communication**

A security group controls the traffic that is allowed to reach and leave the resources that it is associated with. For more information, see [Control traffic to your Amazon resources using security groups](#).

In addition to the standard ports required to access SAP and administrative functions, the following rules must be applied to the security groups assigned to both Amazon EC2 instances in the cluster.

**Inbound**

| Source | Protocol | Port range | Description |
|--------|----------|------------|-------------|
| The security group ID (its own resource ID) | **UDP** | 5405 | Allows UDP traffic between cluster resources for corosync communication |
| Bastion host security group or CIDR range for administration | **TCP** | 7630 | *optional* Used for SLES Hawk2 Interface for monitoring and administration using a Web Interface<br><br>For more information, see [Configuri](#) |

| Source | Protocol | Port range | Description |
|--------|----------|------------|-------------|
|        |          |            | ng and Managing Cluster Resources with Hawk2 in the SUSE documentation. |

> **ⓘ Note**
>
> Note the use of the UDP protocol.

If you are running a local firewall, such as `iptables`, ensure that communication on the preceding ports is allowed between two Amazon EC2 instances.

**Disable source/destination check**

Amazon EC2 instances perform source/destination checks by default, requiring that an instance is either the source or the destination of any traffic it sends or receives.

In the pacemaker cluster, source/destination check must be disabled on both instances receiving traffic from the Overlay IP. You can disable check using Amazon CLI or Amazon Management Console.

Amazon CLI

- Use the [modify-instance-attribute](#) command to disable source/destination check.

  Run the following commands for both instances in the cluster.

  **

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost1> --no-source-dest-check
```

**

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost2> --no-source-dest-check
```

Amazon Management Console

- Ensure that the **Stop** option is checked in https://console.aws.amazon.com/ec2/.



**Review automatic recovery and stop protection**

After a failure, cluster-controlled operations must be resumed in a coordinated way. This helps ensure that the cause of failure is known and addressed, and the status of the cluster is as expected. For example, verifying that there are no pending fencing actions.

This can be achieved by not enabling pacemaker to run as a service at the operating system level or by avoiding auto restarts for hardware failure.

If you want to control the restarts resulting from hardware failure, disable simplified automatic recovery and do not configure Amazon CloudWatch action-based recovery for Amazon EC2 instances that are part of a pacemaker cluster. Use the following commands on both Amazon EC2 instances in the pacemaker cluster, to disable simplified automatic recovery via Amazon CLI. If making the change via Amazon CLI, run the command for both Amazon EC2 instances in the cluster.

> ⓘ **Note**
>
> Modifying instance maintenance options will require admin privileges not covered by the IAM instance roles defined for operations of the cluster.

```
aws ec2 modify-instance-maintenance-options --instance-id <i-xxxxinstidforhost1> --
auto-recovery disabled
```

```
aws ec2 modify-instance-maintenance-options --instance-id <i-xxxxinstidforhost2> --
auto-recovery disabled
```

To ensure that STONITH actions can be executed, you must ensure that stop protection is disabled for Amazon EC2 instances that are part of a pacemaker cluster. If the default settings have been modified, use the following commands for both instances to disable stop protection via Amazon CLI.

> ⓘ **Note**
>
> Modifying instance attributes will require admin privileges not covered by the IAM instance roles defined for operations of the cluster.
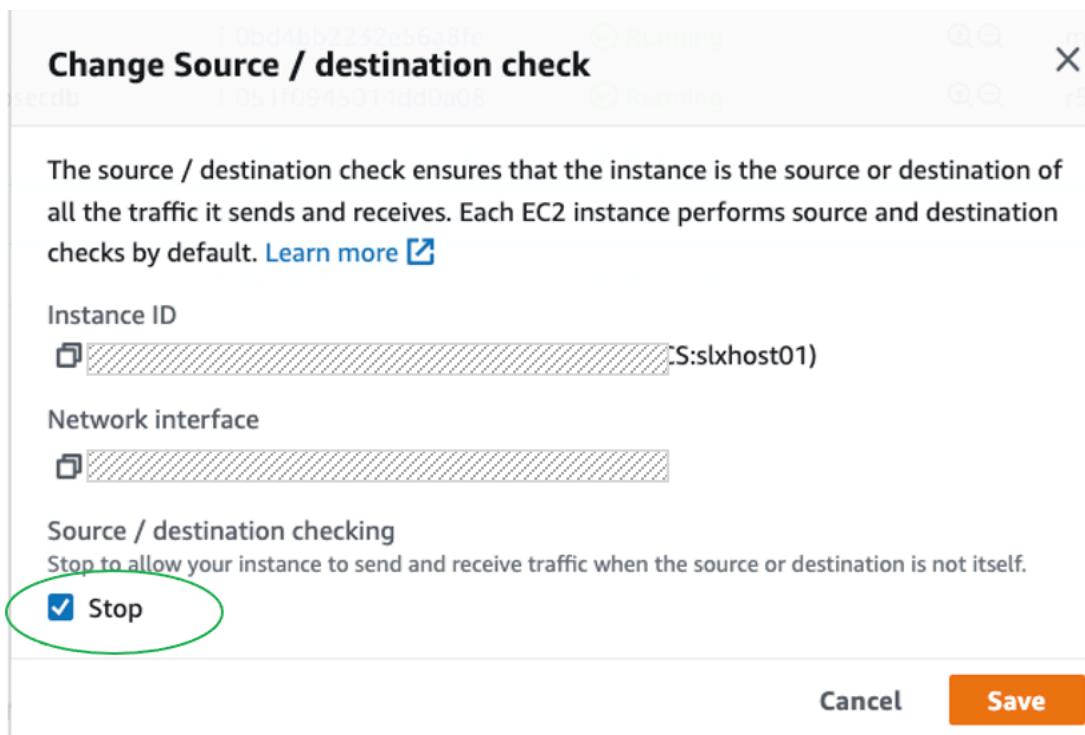
```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost1> --no-disable-
api-stop
```

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost2> --no-disable-
api-stop
```

**Create Amazon EC2 resource tags used by Amazon EC2 STONITH agent**

Amazon EC2 STONITH agent uses Amazon resource tags to identify Amazon EC2 instances. Create tag for the primary and secondary Amazon EC2 instances via Amazon Management Console or Amazon CLI. For more information, see Tagging your Amazon resources.

Use the same tag key and the local hostname returned using the command `hostname` across instances. For example, a configuration with the values defined in Global Amazon parameters would require the tags shown in the following table.

| Amazon EC2 | Key example | Value example |
|---|---|---|
| Instance 1 | pacemaker | slxdbhost1 |
| Instance 2 | pacemaker | slxdbhost2 |

You can run the following command locally to validate the tag values and IAM permissions to describe the tags.

```
aws ec2 describe-tags --filters "Name=resource-id,Values=<instance_id>"
 "Name=key,Values=<pacemaker_tag>" --region=<region> --output=text | cut -f5
```

## Operating system prerequisites

This section covers the following topics.

**Topics**

- [Root access](#)
- [Install missing operating system packages](#)
- [Update and check operating system versions](#)
- [System logging](#)
- [Time synchronization services](#)
- [Amazon CLI profile](#)
- [Pacemaker proxy settings](#)

### Root access

Verify root access on both cluster nodes. The majority of the setup commands in this document are performed with the root user. Assume that commands should be run as root unless there is an explicit call out to choose otherwise.

### Install missing operating system packages

This is applicable to both cluster nodes. You must install any missing operating system packages.

The following packages and their dependencies are required for the pacemaker setup. Depending on your baseline image, for example, SLES for SAP, these packages may already be installed.

```
aws-cli
chrony
cluster-glue
corosync
crmsh
dstat
fence-agents
ha-cluster-bootstrap
iotop
pacemaker
patterns-ha-ha_sles
resource-agents
rsyslog
sap-suse-cluster-connectorsapstartsrv-resource-agents
```

We highly recommend installing the following additional packages for troubleshooting.

```
zypper-lifecycle-plugin
supportutils
yast2-support
supportutils-plugin-suse-public-cloud
supportutils-plugin-ha-sap
```

> ⚠️ **Important**
>
> Ensure that you have installed the newer version `sap-suse-cluster-connector`
> (**dashes**), and not the older version `sap_suse_cluster_connector` that uses
> underscores.

Use the following command to check packages and versions.

```
for package in aws-cli chrony cluster-glue corosync crmsh dstat fence-agents ha-
cluster-bootstrap iotop pacemaker patterns-ha-ha_sles resource-agents rsyslog sap-suse-
cluster-connector sapstartsrv-resource-agents zypper-lifecycle-plugin supportutils
 yast2-support supportutils-plugin-suse-public-cloud supportutils-plugin-ha-sap; do
echo "Checking if ${package} is installed..."
RPM_RC=$(rpm -q ${package} --quiet; echo $?)
if [ ${RPM_RC} -ne 0 ];then
echo "   ${package} is missing and needs to be installed"
fi
```

```
done
```

If a package is not installed, and you are unable to install it using `zypper`, it may be because SUSE Linux Enterprise High Availability extension is not available as a repository in your chosen image. You can verify the availability of the extension using the following command.

```
zypper repos
```

To install or update a package or packages with confirmation, use the following command.

```
zypper install <package_name(s)>
```

**Update and check operating system versions**

You must update and confirm versions across nodes. Apply all the latest patches to your operating system versions. This ensures that bugs are addresses and new features are available.

You can update the patches individually or use the `zypper` update. A clean reboot is recommended prior to setting up a cluster.

```
zypper update
reboot
```

Compare the operating system package versions on the two cluster nodes and ensure that the versions match on both nodes.

**System logging**

This is applicable to both cluster nodes. We recommend using the `rsyslogd` daemon for logging. It is the default configuration in the cluster. Verify that the `rsyslog` package is installed on both cluster nodes.

`logd` is a subsystem to log additional information coming from the STONITH agent.

```
systemctl enable --now logd
systemctl status logd
```

**Time synchronization services**

This is applicable to both cluster nodes. Time synchronization is important for cluster operation. Ensure that `chrony rpm` is installed, and configure appropriate time servers in the configuration file.

You can use Amazon Time Sync Service that is available on any instance running in a VPC. It does not require internet access. To ensure consistency in the handling of leap seconds, don't mix Amazon Time Sync Service with any other `ntp` time sync servers or pools.

Create or check the `/etc/chrony.d/ec2.conf` file to define the server.

```
#Amazon EC2 time source config
server 169.254.169.123 prefer iburst minpoll 4 maxpoll 4
```

Start the `chronyd.service`, using the following command.

```
systemctl enable --now chronyd.service
systemctl status chronyd
```

For more information, see https://docs.aws.amazon.com/AmazonEC2/latest/UserGuide/set-time.html[Set the time for your Linux instance].

**Amazon CLI profile**

This is applicable to both cluster nodes. The cluster resource agents use Amazon Command Line Interface (Amazon CLI). You need to create an Amazon CLI profile for the root account on both instances.

You can either edit the config file at `/root/.aws` manually or by using [aws configure](#) Amazon CLI command.

You can skip providing the information for the access and secret access keys. The permissions are provided through IAM roles attached to Amazon EC2 instances.

```
aws configure
{aws} Access Key ID [None]:
{aws} Secret Access Key [None]:
Default region name [None]: <region_id>
```

```
Default output format [None]:
```

The profile name is configurable. The name chosen in this example is **cluster** – it is used in [Create Amazon EC2 resource tags used by Amazon EC2 STONITH agent](). The Amazon Region must be the default Amazon Region of the instance.

```
aws configure --profile <cluster>
{aws} Access Key ID [None]:
{aws} Secret Access Key [None]:
Default region name [None]: <region_id>
Default output format [None]:
```

**Pacemaker proxy settings**

This is applicable to both cluster nodes. If your Amazon EC2 instance has been configured to access the internet and/or Amazon Cloud through proxy servers, then you need to replicate the settings in the pacemaker configuration. For more information, see [Use an HTTP proxy]().

Add the following lines to `/etc/sysconfig/pacemaker`.

```
http_proxy=http://<proxyhost>:<proxyport>
https_proxy= http://<proxyhost>:<proxyport>
no_proxy=127.0.0.1,localhost,169.254.169.254,fd00:ec2::254
```

Modify `proxyhost` and `proxyport` to match your settings. Ensure that you exempt the address used to access the instance metadata. Configure `no_proxy` to include the IP address of the instance metadata service – **169.254.169.254** (IPV4) and **fd00:ec2::254** (IPV6). This address does not vary.

## IP and hostname resolution prerequisites

This section covers the following topics.

**Topics**

- [Primary and secondary IP addresses]()
- [Add initial VPC route table entries for overlay IPs]()
- [Add overlay IPs to host IP configuration]()
- [Hostname resolution]()

**Primary and secondary IP addresses**

This is applicable to both cluster nodes. We recommend defining a redundant communication channel (a second ring) in `corosync` for SUSE clusters. The cluster nodes can use the second ring to communicate in case of underlying network disruptions.

Create a redundant communication channel by adding a secondary IP address on both nodes.

Add a secondary IP address on both nodes. These IPs are only used in cluster configurations. They provide the same fault tolerance as a secondary Elastic Network Interface (ENI). For more information, see Assign a secondary private IPv4 address.

On correct configuration, the following command returns two IPs from the same subnet on both, the primary and secondary node.

```
ip -o -f inet addr show eth0 | awk -F " |/" '{print $7}'
```

These IP addresses are required for `ring0_addr` and `ring1_addr` in `corosync.conf`.

**Add initial VPC route table entries for overlay IPs**

You need to add initial route table entries for overlay IPs. For more information on overlay IP, see Overlay IP.

Add entries to the VPC route table or tables associated with the subnets of your Amazon EC2 instance for the cluster. The entries for destination (overlay IP CIDR) and target (Amazon EC2 instance or ENI) must be added manually for SAP ASE database. This ensures that the cluster resource has a route to modify. It also supports the install of SAP using the virtual names associated with the overlay IP before the configuration of the cluster.

**Modify or add a route to a route table using Amazon Management Console**

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Route Tables**, and select the route table associated with the subnets where your instances have been deployed.
3. Choose **Actions**, **Edit routes**.
4. To add a route, choose **Add route**.
   a. Add your chosen overlay IP address CIDR and the instance ID of your primary instance for SAP ASE database. See the following table for an **example**.

| Destination | 172.16.0.29/32 |
| Target | i-xxxxinstidforhost1 |

5. Choose **Save changes**.

The preceding steps can also be performed programmatically. We suggest performing the steps using administrative privileges, instead of instance-based privileges to preserve least privilege. CreateRoute API isn't necessary for ongoing operations.

Run the following command as a dry run on both nodes to confirm that the instances have the necessary permissions.

```
aws ec2 replace-route --route-table-id <rtb-xxxxxroutetable1> --destination-cidr-
block <172.16.0.29/32> --instance-id <i-xxxxinstidforhost1> --dry-run --profile
 <aws_cli_cluster_profile>
```

**Add overlay IPs to host IP configuration**

You must configure the overlay IP as an additional IP address on the standard interface to enable SAP install. This action is managed by the cluster IP resource. However, to install SAP using the correct IP addresses prior to having the cluster configuration in place, you need to add these entries manually.

If you need to reboot the instance during setup, the assignment is lost, and must be re-added.

See the following **example**. You must update the command with your chosen IP addresses.

On EC2 instance 1, where you are installing SAP ASE database, add the overlay IP.

```
ip addr add <172.16.0.29/32> dev eth0
```

**Hostname resolution**

This is applicable to both cluster nodes. You must ensure that both instances can resolve all hostnames in use. Add the hostnames for cluster nodes to `/etc/hosts` file on both cluster nodes. This ensures that hostnames for cluster nodes can be resolved even in case of DNS issues. See the following example.

```
cat /etc/hosts
<10.1.10.1 slxdbhost01.example.com slxdbhost01>
<10.1.20.1 slxdbhost02.example.com slxdbhost02>
<172.16.0.29 slxvdb.example.com slxvdb>
```

In this example, the secondary IPs used for the second cluster ring are not mentioned. They are only used in the cluster configuration. You can allocate virtual hostnames for administration and identification purposes.

> ⚠️ **Important**
>
> The overlay IP is out of VPC range, and cannot be reached from locations not associated with the route table, including on-premises.

## FSx for ONTAP prerequisites

This section covers the following topics.

**Topics**

- Shared file systems
- Create volumes and file systems

**Shared file systems**

Amazon FSx for NetApp ONTAP is supported for SAP ASE database file systems.

FSx for ONTAP provides fully managed shared storage in Amazon Cloud with data access and management capabilities of ONTAP. For more information, see Create an Amazon FSx for NetApp ONTAP file system.

Select a file system based on your business requirements, evaluating the resilience, performance, and cost of your choice.

The SVM's DNS name is your simplest mounting option. The file system DNS name automatically resolves to the mount target's IP address on the Availability ZOne of the connecting Amazon EC2 instance.

```
svm-id.fs-id.fsx.aws-region.amazonaws.com
```

> ⓘ **Note**
>
> Review the `enableDnsHostnames` and `enableDnsSupport` DNS attributes for your VPC. For more information, see [View and update DNS attributes for your VPC](#).

**Create volumes and file systems**

You can review the following resources to understand the FSx for ONTAP mount points for SAP ASE database.

- [Host setup for SAP ASE](#)
- SAP – [Setup of Database Layout](#) (ABAP)
- SAP – [Setup of Database Layout](#) (JAVA)

The following are the FSx for ONTAP mount points covered in this topic.

| Unique NFS Location (example) | File system location |
|---|---|
| SVM-xxx:/sybase | /sybase |
| SVM-xxx:/asedata | /sybase/<DBSID>/sapdata_1 |
| SVM-xxx:/aselog | /sybase/<DBSID>/saplog_1 |
| SVM-xxx:/sapdiag | /sybase/<DBSID>/sapdiag |
| SVM-xxx:/saptmp | /sybase/<DBSID>/saptmp |
| SVM-xxx:/backup | /sybasebackup |
| SVM-xxx:/usrsap | /usr/sap |

Ensure that you have properly mounted the file systems, and the necessary adjustments for host setup have been performed. See [Host setup for SAP ASE](#). You can temporarily add the entries to `/etc/fstab` to not lose them during a reboot. The entries must be removed prior to configuring the cluster. The cluster resource manages the mounting of the NFS.

You need to perform this step only on the primary Amazon EC2 instance for the initial installation.

Review the mount options to ensure that they match with your operating system, NFS file system type, and SAP's latest recommendations.

Use the following command to check that the required file systems are available.

```
df -h
```

## Shared VPC – *optional*

Amazon VPC sharing enables you to share subnets with other Amazon accounts within the same Amazon Organizations. Amazon EC2 instances can be deployed using the subnets of the shared Amazon VPC.

In the pacemaker cluster, the `aws-vpc-move-ip` resource agent has been enhanced to support a shared VPC setup while maintaining backward compatibility with previous existing features.

The following checks and changes are required. We refer to the Amazon account that owns Amazon VPC as the sharing VPC account, and to the consumer account where the cluster nodes are going to be deployed as the cluster account.

This section covers the following topics.

**Topics**

- [Minimum version requirements](#)
- [IAM roles and policies](#)
- [Shared VPC cluster resources](#)

**Minimum version requirements**

The latest version of the `aws-vpc-move-ip` agent shipped with SLES15 SP3 supports the shared VPC setup by default. The following are the minimum version required to support a shared VPC Setup:

- SLES 12 SP5 - resource-agents-4.3.018.a7fb5035-3.79.1.x86_64
- SLES 15 SP2 - resource-agents-4.4.0+git57.70549516-3.30.1.x86_64
- SLES 15 SP3 - resource-agents-4.8.0+git30.d0077df0-8.5.1

## IAM roles and policies

Using the overlay IP agent with a shared Amazon VPC requires a different set of IAM permissions to be granted on both Amazon accounts (sharing VPC account and cluster account).

### Sharing VPC account

In sharing VPC account, create an IAM role to delegate permissions to the EC2 instances that will be part of the cluster. During the IAM Role creation, select "Another Amazon account" as the type of trusted entity, and enter the Amazon account ID where the EC2 instances will be deployed/running from.

After the IAM role has been created, create the following IAM policy on the sharing VPC account, and attach it to an IAM role. Add or remove route table entries as needed.

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": "ec2:ReplaceRoute",
      "Resource": [
        "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0",
        "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": "ec2:DescribeRouteTables",
      "Resource": "*"
    }
  ]
}
```

Next, edit move to the "Trust relationships" tab in the IAM role, and ensure that the Amazon account you entered while creating the role has been correctly added.

### Cluster account

In cluster account, create the following IAM policy, and attach it to an IAM role. This is the IAM Role that is going to be attached to the EC2 instances.

## Amazon STS policy

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::123456789012:role/sharing-vpc-account-cluster-role"
    }
  ]
}
```

## STONITH policy

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances"
      ],
      "Resource": [
        "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0",
        "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0"
      ]
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": "ec2:DescribeInstances",
      "Resource": "*"
    }
  ]
}
```

**Shared VPC cluster resources**

The cluster resource agent `aws-vpc-move-ip` also uses a different configuration syntax. When configuring the `aws-vpc-move-ip` resource agent, the following new parameters must be used:

- lookup_type=NetworkInterfaceId
- routing_table_role="arn:aws:iam::<account_id>:role/<VPC-Account-Cluster-Role>"

The following IP Resource for SAP ASE database needs to be created.

```
crm configure primitive rsc_ip_SD_ASEDB ocf:heartbeat:aws-vpc-move-ip params
 ip=172.16.0.29
routing_table=rtb-xxxxxroutetable1 interface=eth0 profile=cluster
 lookup_type=NetworkInterfaceId
routing_table_role="arn:aws:iam::<sharing_vpc_account_id>:role/
<sharing_vpc_account_cluster_role>"
op start interval=0 timeout=180s op stop interval=0 timeout=180s op monitor
 interval=20s
timeout=40s
```

# SAP ASE and cluster setup

This section covers the following topics.

**Topics**

- Install SAP ASE database
- Cluster prerequisites
- Create cluster and node associations

## Install SAP ASE database

The following topics provide information about installing SAP ASE database on Amazon Cloud in a highly available cluster. Review SAP Documentation for more details.

**Topics**

- Use SWPM
- Install SAP database instance
- Check SAP host agent version

**Use SWPM**

Before running SAP Software Provisioning Manager (SWPM), ensure that the following prerequisites are met.

- If the operating system groups for SAP are pre-defined, ensure that the user identifier (UID) and group identifier values for `sapadm`, `<syb>adm`, and `sapsys` are consistent across both instances.
- You have downloaded the most recent version of Software Provisioning Manager for your SAP version. For more information, see SAP Documentation [Software Provisioning Manager](#).
- Ensure that routes, overlay IPs, and virtual host names are mapped to the instance where the installation is run. This is to ensure that the virtual hostname for SAP ASE database is available on the primary instance. For more information, see [IP and hostname resolution prerequisites](#).
- Ensure that FSx for ONTAP mount points are available, either in `/etc/fstab` or using the mount command. For more information, see [File system prerequisites](#). If you are adding the entries in `/etc/fstab`, ensure that they are removed before configuring the cluster.

**Install SAP database instance**

The commands in this section use the example values provided in [Define reference parameters for setup](#).

Install SAP ASE database on `slxdbhost01` with virtual hostname `slxvdb`, using the high availability option of Software Provisioning Manager (SWPM) tool. You can use the `SAPINST_USE_HOSTNAME` parameter to install SAP using a virtual hostname.

```
<swpm location>/sapinst SAPINST_USE_HOSTNAME=<slxvdb>
```

> **ⓘ Note**
>
> Before installing SAP ASE database, ASCS and ERS must be installed, and the `/sapmnt` directory must be available on the database server.

**Check SAP host agent version**

The SAP host agent is used for ASE database instance control and monitoring. This agent is used by SAP cluster resource agents and hooks. It is recommended that you have the latest version

installed on both instances. For more details, see [SAP Note 2219592 – Upgrade Strategy of SAP Host Agent](#).

Use the following command to check the version of the host agent.

```
/usr/sap/hostctrl/exe/saphostexec -version
```

## Cluster prerequisites

This section covers the following topics.

**Topics**

- [Update the hacluster password](#)
- [Setup passwordless authentication between nodes](#)
- [Create an authentication key for corosync](#)

### Update the `hacluster` password

This is applicable to both cluster nodes. Change the password of the operating system user `hacluster` using the following command.

```
passwd hacluster
```

### Setup passwordless authentication between nodes

For a more comprehensive and easily consumable view of cluster activity, SUSE provides additional reporting tools. Many of these tools require access to both nodes without entering a password. SUSE recommends performing this setup for root user. For more details, see *Configuration to collect cluster report as root with root SSH access between cluster nodes* section in SUSE Documentation https://www.suse.com/support/kb/doc/?id=000017501#:~:text=The %20hb_report%20utility%20(on%20newer,an%20incident%20to%20be%20investigated.[Usage of hb_report for SLES HAE].

### Create an authentication key for `corosync`

If you want to configure `corosync` to use cryptographic techniques for ensuring authenticity and privacy of the messages, you need to generate a private key. The executable `corosync-keygen` creates this key and writes it to `/etc/corosync/authkey`.

Use the following command on Node 1 as root.

```
corosync-keygen
```

Use `scp` or a temporary shared NFS location to copy an identical file on the second node at the same location. For example, on `slxdbhost01`.

```
scp -p /etc/corosync/authkey root@<slxdbhost02>:/etc/corosync
```

## Create cluster and node associations

This section covers the following topics.

**Topics**

- [Stop services for initial configuration](#)
- [File modifications and key values](#)
- [Sample corosync.conf file](#)

**Stop services for initial configuration**

This is applicable to both cluster nodes. The cluster service `pacemaker` must be in a stopped state when performing cluster configuration.

Run the following command to check if `pacemaker` is running.

```
systemctl status pacemaker
```

Run the following command to stop `pacemaker`.

```
systemctl stop pacemaker
```

**File modifications and key values**

`corosync.conf` is the configuration file for the `corosync` executable. Copy the contents of the [Sample corosync.conf file](#) to `/etc/corosync/corosync.conf` on both nodes.

Ensure the following when copying the file.

- Ensure that the node list IP addresses match the primary and secondary IPs on each host (not the overlay IP)

- Ensure that the file is same on both nodes, with the exception of `bindnetaddr` that should match the relevant local primary IP address on each node.

- Ensure that the token value is set to 30000. This timeout specifies the time taken in milliseconds until a token loss is declared after not receiving a token. This is important for the stability of the cluster.

**Sample `corosync.conf` file**

The following is a sample `corosync.conf` file.

Ensure that the file is same on both nodes, with the exception of `bindnetaddr` that should match the relevant local primary IP address on each node.

```
#Read the corosync.conf.5 manual page
totem {
  version: 2
  rrp_mode: passive
  token: 30000
  consensus: 36000
  token_retransmits_before_loss_const: 10
  max_messages: 20
  crypto_cipher: aes256
  crypto_hash: sha1
  clear_node_high_bit: yes
  interface {
    ringnumber: 0
    bindnetaddr: <local_ip>
    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: <primary_host_ip>
   ring1_addr: <primary_host_additional_ip>
   nodeid: 1
   }
   node {
   ring0_addr: <secondary_host_ip>
   ring1_addr: <secondary_host_additional_ip>
   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
}
```

The following table displays example substitutions for IP addresses using the sample IP addresses provided in this document. The <local_ip> configuration differs between hosts.

| IP address type | Primary host | Secondary host |
| --- | --- | --- |
| <local_ip> | **10.1.10.1** | **10.1.20.1** |
| <primary_host_ip> | 10.1.10.1 | 10.1.10.1 |
| <primary_host_additional_ip> | 10.1.10.2 | 10.1.10.2 |
| <secondary_host_ip> | 10.1.20.1 | 10.1.20.1 |
| <secondary_host_additional_ ip> | 10.1.20.2 | 10.1.20.2 |

# Cluster configuration

This section covers the following topics.

**Topics**

- [Cluster resources](#)
- [Sample configuration (crm config)](#)

## Cluster resources

This section covers the following topics.

**Topics**

- [Enable and start the cluster](#)
- [Check cluster status](#)
- [Prepare for resource creation](#)
- [Reset configuration – optional](#)
- [Cluster bootstrap](#)
- [Create Amazon EC2 STONITH resource](#)
- [Create file system resources](#)
- [Create overlay IP resources](#)
- [Create SAP ASE database resource](#)
- [Activate cluster](#)

### Enable and start the cluster

This is applicable to both cluster nodes. Run the following command to enable and start the `pacemaker` cluster service on both nodes.

```
systemctl enable --now pacemaker
```

or

```
systemctl start pacemaker
```

By enabling the `pacemaker` service, the server automatically joins the cluster after a reboot. This ensures that your system is protected. Alternatively, you can start the `pacemaker` service manually on boot. You can then investigate the cause of failure. However, this is generally not required for SAP NetWeaver ASCS cluster.

Run the following command to check the status of the `pacemaker` service.

```
systemctl status pacemaker
# pacemaker.service - Pacemaker High Availability Cluster Manager
     Loaded: loaded (/usr/lib/systemd/system/pacemaker.service; <enabled>; vendor
 preset: disabled)
     Active: <active (running)> since Tue XXXX-XX-XX XX:XX:XX XXX; XXh ago
       Docs: man:pacemakerd
             https://clusterlabs.org/pacemaker/doc/en-US/Pacemaker/2.0/html-single/
Pacemaker_Explained/index.html
   Main PID: 1899 (pacemakerd)
Enable cluster service (optional)
```

**Check cluster status**

Once the cluster service `pacemaker` is started, check the cluster status with `crm mon` command, as shown in the following example.

```
crm_mon -1
Cluster Summary:
  * Stack: corosync
  * Current DC: <slxdbhost01> (version 2.0.xxxxxxxxxx) - partition with quorum
  * Last updated:
  * Last change:  by hacluster via crmd on slxdbhost01
  * 2 nodes configured
  * 0 resource instances configured

Node List:
  * Online: [ <slxdbhost0> <slxdbhost02> ]

Active Resources:
  * No active resources
```

The primary (`slxdbhost01`) and secondary (`slxdbhost02`) must show up as online.

You can find the ring status and the associated IP address of the cluster with `corosync-cfgtool` command, as shown in the following example.

```
corosync-cfgtool -s
Printing ring status.
Local node ID 1
RING ID 0
        id      = <10.1.10.1>
        status  = ring 0 active with no faults
RING ID 1
        id      = <10.1.10.2>
        status  = ring 1 active with no faults
```

**Prepare for resource creation**

To ensure that the cluster does not perform any unexpected actions during setup of resources and configuration, set the maintenance mode to true.

Run the following command to put the cluster in maintenance mode.

```
crm maintenance on
```

**Reset configuration – *optional***

> ⓘ **Note**
>
> The following instructions help you reset the complete configuration. Run these commands only if you want to start setup from the beginning. You can make minor changes with the `crm edit` command.

Run the following command to back up the current configuration for reference.

```
crm config show > /tmp/crmconfig_backup.txt
```

Run the following command to clear the current configuration.

```
crm configure erase
```

> ⚠ **Important**
>
> Once the preceding erase command is executed, it removes all of the cluster resources from Cluster Information Base (CIB), and disconnects the communication from `corosync` to

the cluster. Before starting the resource configuration run `crm cluster restart`, so that cluster reestablishes communication with `corosync`, and retrieves the configuration. The restart of cluster removes *maintenance mode*. Reapply before commencing additional configuration and resource setup.

### Cluster bootstrap

Configure the cluster bootstrap parameters by running the following commands.

```
crm configure rsc_defaults resource-stickiness=1
crm configure rsc_defaults migration-threshold=3
crm configure property stonith-enabled="true"
crm configure property stonith-action="off"
crm configure property stonith-timeout="300s"
crm configure op_defaults timeout="300s"
crm configure op_defaults record-pending="true"
```

### Create Amazon EC2 STONITH resource

Modify the following command to match your configuration values.

```
crm configure primitive res_{aws}_STONITH stonith:external/ec2 op start interval=0
  timeout=180s op stop interval=0 timeout=180s op monitor interval=180s timeout=60s
  params tag=pacemaker profile=cluster pcmk_delay_max=30
```

**profile** – this refers to the Amazon CLI profile crated during setup. In the preceding command, *cluster* is the profile name.

### Create file system resources

Mounting and unmounting file system resources to align with the location of SAP ASE database is done using cluster resources.

Modify and run the following commands to create these file system resources.

### /sybase

```
crm configure primitive rsc_fs_<DBSID>_sybase ocf:heartbeat:Filesystem params
  device="<nfs.fqdn>:/sybase" directory="/sybase" fstype="nfs4" options="
```

```
rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,timeo=600,r
op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
timeout=40s
```

## /sybase/<DBSID>/sapdata_1

```
crm configure primitive rsc_fs_<DBSID>_data ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/asedata" directory="/sybase/<DBSID>/sapdata_1" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=8,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/saplog_1

```
crm configure primitive rsc_fs_<DBSID>_log ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/aselog" directory="/sybase/<DBSID>/saplog_1" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/sapdiag

```
crm configure primitive rsc_fs_<DBSID>_diag ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/sapdiag" directory="/sybase/<DBSID>/sapdiag" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/saptmp

```
crm configure primitive rsc_fs_<DBSID>_tmp ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/saptmp" directory="/sybase/<DBSID>/saptmp" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybasebackup

```
crm configure primitive rsc_fs_<DBSID>_bkp ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/sybasebackup" directory="/backup" fstype="nfs4"
```

```
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /usr/sap

```
crm configure primitive rsc_fs_<DBSID>_sap ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/usrsap" directory="/usr/sap" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

### Notes

- Review the mount options to ensure that they match with your operating system, NFS file system type, and the latest recommendations from SAP and Amazon.

- `<nfs.fqdn>` must be the alias of the FSx for ONTAP resource. For example, `svm-xxxxx.fs-xxxxx.<region>.amazonaws.com`.

- Your file system structure can vary – it can have multiple data file systems. The preceding examples must be adapted to your environment.

### Create overlay IP resources

The IP resource provides the details necessary to update the route table entry for overlay IP.

Modify and run the following command to create IP resources.

```
crm configure primitive rsc_ip_<DBSID>_ASEDB ocf:heartbeat:aws-vpc-move-ip params
 ip=172.16.0.29 routing_table=rtb-xxxxxroutetable1 interface=eth0 profile=cluster op
 start interval=0 timeout=180s op stop interval=0 timeout=180s op monitor interval=20s
 timeout=40s
```

### Notes

- If more than one route table is required for connectivity or because of subnet associations, the `routing_table` parameter can have multiple values separated by a comma. For example, `routing_table=rtb-xxxxxroutetable1, rtb-xxxxxroutetable2`.

- Additional parameters – `lookup_type` and `routing_table_role` are required for shared VPC. For more information, see Shared VPC – optional.

**Create SAP ASE database resource**

SAP ASE database is started and stopped using cluster resources.

Modify and run the following command to create the SAPDatabase resource.

```
crm configure primitive rsc_ase_<DBSID>_ASEDB ocf:heartbeat:SAPDatabase SID=<DBSID>
 DBTYPE=SYB STRICT_MONITORING=TRUE op start timeout=300 op stop timeout=300
```

Create the cluster resource group, and the resources together in the order in which the services will be started and stopped.

```
crm configure group grp_<DBSID>_ASEDB rsc_fs_<DBSID>_sybase rsc_fs_<DBSID>_data
 rsc_fs_<DBSID>_log rsc_fs_<DBSID>_diag rsc_fs_<DBSID>_tmp rsc_fs_<DBSID>_bkp
 rsc_fs_<DBSID>_sap rsc_aws_stonith_<DBSID> rsc_ase_<DBSID>_ASEDB
```

**Activate cluster**

Use `crm config show` and `crm config edit` commands to review that all the values have been entered correctly.

On confirmation of correct values, set the maintenance mode to false using the following command. This enables the cluster to take control of the resources.

```
crm maintenance off
```

See the [Sample configuration](#).

# Sample configuration (crm config)

```
node 1: slxdbhost01
node 2: slxdbhost02
primitive rsc_ase_ASD_ASEDB SAPDatabase \
        params SID=ASD DBTYPE=SYB STRICT_MONITORING=TRUE \
        op start timeout=300 interval=0s \
        op stop timeout=300 interval=0s \
        op monitor timeout=60s interval=120s \
        meta target-role=Started
primitive rsc_aws_stonith_ASD stonith:external/ec2 \
        params tag=pacemaker profile=cluster pcmk_delay_max=30 \
        op start interval=0 timeout=180s \
        op stop interval=0 timeout=180s \
        op monitor interval=180s timeout=60s
```

```
primitive rsc_fs_ASD_bkp Filesystem \
        params device="svm-091efa9986c8e93c7.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/backup" directory="/sybasebackup" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_data Filesystem \
        params device="svm-0e6e2738a9ca391ce.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/asedata" directory="/sybase/ASD/sapdata_1" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=8,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_diag Filesystem \
        params device="svm-091efa9986c8e93c7.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/sapdiag" directory="/sybase/ASD/sapdiag" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_log Filesystem \
        params device="svm-0895fe73884c12f83.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/aselog" directory="/sybase/ASD/saplog_1" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_sap Filesystem \
        params device="svm-091efa9986c8e93c7.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/usrsap" directory="/usr/sap" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_sybase Filesystem \
        params device="svm-091efa9986c8e93c7.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/sybase" directory="/sybase" fstype=nfs4
```

```
        options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_fs_ASD_tmp Filesystem \
        params device="svm-091efa9986c8e93c7.fs-0c3a4a5162a325aea.fsx.us-
east-1.amazonaws.com:/saptmp" directory="/sybase/ASD/saptmp" fstype=nfs4
 options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,ti
 \
        op start timeout=60s interval=0 \
        op stop timeout=60s interval=0 \
        op monitor interval=20s timeout=40s
primitive rsc_ip_SD_ASEDB aws-vpc-move-ip \
        params ip=172.16.0.29 routing_table=rtb-0b3f1d6196f45300d interface=eth0
 profile=cluster \
        op start interval=0 timeout=180s \
        op stop interval=0 timeout=180s \
        op monitor interval=20s timeout=40s
group grp_ASD_ASEDB rsc_fs_ASD_sybase rsc_fs_ASD_data rsc_fs_ASD_log rsc_fs_ASD_diag
 rsc_fs_ASD_tmp rsc_fs_ASD_bkp rsc_fs_ASD_sap rsc_ip_SD_ASEDB rsc_ase_ASD_ASEDB
property cib-bootstrap-options: \
        maintenance-mode=false \
        stonith-enabled=true \
        stonith-action=off \
        stonith-timeout=300s \
        last-lrm-refresh=1686941627 \
        have-watchdog=false \
        dc-version="2.1.2+20211124.ada5c3b36-150400.4.9.2-2.1.2+20211124.ada5c3b36" \
        cluster-infrastructure=corosync
rsc_defaults rsc-options: \
        resource-stickiness=1 \
        migration-threshold=1
op_defaults op-options: \
        timeout=300s \
        record-pending=true
```

# Operations

This section covers the following topics.

**Topics**

- [Analysis and maintenance](#)

- [Testing](#)

# Analysis and maintenance

This section covers the following topics.

**Topics**

- [Viewing the cluster state](#)

- [Performing planned maintenance](#)

- [Post-failure analysis and reset](#)

- [Alerting and monitoring](#)

## Viewing the cluster state

You can view the state of the cluster in two ways - based on your operating system or with a web based console provided by SUSE.

**Topics**

- [Operating system based](#)

- [SUSE Hawk2](#)

### Operating system based

There are multiple operating system commands that can be run as root or as a user with appropriate permissions. The commands enable you to get an overview of the status of the cluster and its services. See the following commands for more details.

```
crm status
```

Sample output:

```
slxdbhost01:~  crm status
Cluster Summary:
 * Stack: corosync
```

```
* Current DC: slxdbhost02 (version 2.1.2+20211124.ada5c3b36-150400.4.9.2-
2.1.2+20211124.ada5c3b36) - partition with quorum
* Last updated: Sat Jun 17 01:16:10 2023
* Last change: Sat Jun 17 01:15:31 2023 by root via crm_resource on slxdbhost01
* 2 nodes configured
* 10 resource instances configured
Node List:
* Online: [ slxdbhost01 slxdbhost02 ]
Full List of Resources:
* rsc_aws_stonith_ASD (stonith:external/ec2): Started slxdbhost02
* Resource Group: grp_ASD_ASEDB:
* rsc_fs_ASD_sybase (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_data (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_log (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_diag (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_tmp (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_bkp (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_fs_ASD_sap (ocf::heartbeat:Filesystem): Started slxdbhost01
* rsc_ip_SD_ASEDB (ocf::heartbeat:aws-vpc-move-ip): Started slxdbhost01
* rsc_ase_ASD_ASEDB (ocf::heartbeat:SAPDatabase): Started slxdbhost01
```

The following table provides a list of useful commands.

| Command | Description |
|---------|-------------|
| crm_mon | Display cluster status on the console with updates as they occur |
| crm_mon -1 | Display cluster status on the console just once, and exit |
| crm_mon -Arnf | -A Display node attributes

-n Group resources by node

-r Display inactive resources

-f Display resource fail counts |
| crm help | View more options |
| crm_mon --help-all | View more options |

**SUSE Hawk2**

Hawk2 is a web-based graphical user interface for managing and monitoring pacemaker highly availability clusters. It must be enabled on every node in the cluster, to point your web browser on any node for accessing it. Use the following command to enable Hawk2.

```
systemctl enable --now hawk
systemctl status hawk
```

Use the following URL to check security groups for access on port 7630 from your administrative host.

```
https://your-server:7630/

e.g https://slxdbhost01:7630
```

For more information, see [Configuring and Managing Cluster Resources with Hawk2](#) in the SUSE Documentation.

## Performing planned maintenance

The cluster connector is designed to integrate the cluster with SAP start framework (`sapstartsrv`), including the rolling kernel switch (RKS) awareness. Stopping and starting the SAP system using `sapcontrol` should not result in any cluster remediation activities as these actions are not interpreted as failures. Validate this scenario when testing your cluster.

There are different options to perform planned maintenance on nodes, resources, and the cluster.

**Topics**

- [Maintenance mode](#)
- [Placing a node in standby mode](#)
- [Moving a resource (not recommended)](#)

**Maintenance mode**

Use maintenance mode if you want to make any changes to the configuration or take control of the resources and nodes in the cluster. In most cases, this is the safest option for administrative tasks.

On

- Use one of the following commands to turn on maintenance mode.

```
crm maintenance on
```

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

Off

- Use one of the following commands to turn off maintenance mode.

```
crm maintenance off
```

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

**Placing a node in standby mode**

To perform maintenance on the cluster without system outage, the recommended method for moving active resources is to place the node you want to remove from the cluster in standby mode.

```
crm node standby <slxdbhost01>
```

The cluster will cleanly relocate resources, and you can perform activities, including reboots on the node in standby mode. When maintenance activities are complete, you can re-introduce the node with the following command.

```
crm node online <slxdbhost01>
```

**Moving a resource (not recommended)**

Moving individual resources is not recommended because of the migration or move constraints that are created to lock the resource in its new location. These can be cleared as described in the info messages, but this introduces an additional setup.

```
slxdbhost01:~  crm resource move grp_ASD_ASEDB force
INFO: Move constraint created for grp_ASD_ASEDB
```

```
INFO: Use `crm resource clear grp_ASD_ASEDB` to remove this constraint
```

Use the following command once the resources have relocated to their target location.

```
<slxdbhost01>:~  crm resource clear grp_ASD_ASEDB
```

## Post-failure analysis and reset

A review must be conducted after each failure to understand the source of failure as well the reaction of the cluster. In most scenarios, the cluster prevents an application outage. However, a manual action is often required to reset the cluster to a protective state for any subsequent failures.

**Topics**

- Checking the logs
- Cleanup crm status
- Restart failed nodes or pacemaker
- Further analysis

### Checking the logs

Start your troubleshooting by checking the operating system log `/var/log/messages`. You can find additional information in the cluster and pacemaker logs.

- **Cluster logs** – updated in the `corosync.conf` file located at `/etc/corosync/corosync.conf`.
- **Pacemaker logs** – updated in the `pacemaker.log` file located at `/var/log/pacemaker`.
- **Resource agents** – `/var/log/messages`

Application based failures can be investigated in the SAP work directory.

### Cleanup `crm status`

If failed actions are reported using the `crm status` command, and if they have already been investigated, then you can clear the reports with the following command.

```
crm resource cleanup <resource> <hostname>
```

**Restart failed nodes or `pacemaker`**

It is recommended that failed (or fenced) nodes are not automatically restarted. It gives operators a chance to investigate the failure, and ensure that the cluster doesn't make assumptions about the state of resources.

You need to restart the instance or the pacemaker service based on your approach.

**Further analysis**

The following commands consolidate information from both nodes, highlighting key events and differentiating between originating node to make the analysis clear.

```
crm history events

crm history log
```

If further analysis from SUSE is required, an `hb_report` may be requested. For more information, see SUSE Documentation – [Usage of hb_report for SLES HAE](#).

> **ⓘ Note**
>
> `crm history events` and `hb_report` rely on passwordless ssh being set up between the nodes.

# Alerting and monitoring

**Using the cluster alert agents**

Within the cluster configuration, you can call an external program (an alert agent) to handle alerts. This is a *push* notification. It passes information about the event via environment variables.

The agents can then be configured to send emails, log to a file, update a monitoring system, etc. For example, the following script can be used to access Amazon SNS.

```
#!/bin/sh

#alert_sns.sh
#modified from /usr/share/pacemaker/alerts/alert_smtp.sh.sample

##############################################################################
```

```
#SETUP
* Create an SNS Topic and subscribe email or chatbot
* Note down the ARN for the SNS topic
* Give the IAM Role attached to both Instances permission to publish to the SNS Topic
* Ensure the aws cli is installed
* Copy this file to /usr/share/pacemaker/alerts/alert_sns.sh or other location on BOTH
 nodes
* Ensure the permissions allow for hacluster and root to execute the script
* Run the following as root (modify file location if necessary and replace SNS ARN):

#SLES:
crm configure alert aws_sns_alert /usr/share/pacemaker/alerts/alert_sns.sh meta
 timeout=30s timestamp-format="%Y-%m-%d_%H:%M:%S" to { arn:aws:sns:region:account-
id:myPacemakerAlerts  }
#RHEL:
pcs alert create id=aws_sns_alert path=/usr/share/pacemaker/alerts/alert_sns.sh meta
 timeout=30s timestamp-format="%Y-%m-%d_%H:%M:%S"
pcs alert recipient add aws_sns_alert value=<arn:aws:sns:region:account-
id:myPacemakerAlerts>


#Additional information to send with the alerts.

node_name=`uname -n`
sns_body=`env | grep CRM_alert_`
#Required for SNS
TOKEN=$(/usr/bin/curl --noproxy '*' -s -X PUT "http://169.254.169.254/latest/api/token"
 -H "X-aws-ec2-metadata-token-ttl-seconds: 21600")

#Get metadata
REGION=$(/usr/bin/curl --noproxy '*' -w "\n" -s -H "X-aws-ec2-metadata-token: $TOKEN"
 http://169.254.169.254/latest/dynamic/instance-identity/document | grep region | awk -
F\" '{print $4}')


sns_subscription_arn=${CRM_alert_recipient}

#Format depending on alert type
case ${CRM_alert_kind} in
    node)
      sns_subject="${CRM_alert_timestamp} ${cluster_name}: Node '${CRM_alert_node}' is
 now '${CRM_alert_desc}'"
    ;;
    fencing)
      sns_subject="${CRM_alert_timestamp} ${cluster_name}: Fencing ${CRM_alert_desc}"
    ;;
```

```
    resource)
      if [ ${CRM_alert_interval} = "0" ]; then
          CRM_alert_interval=""
      else
          CRM_alert_interval=" (${CRM_alert_interval})"
      fi
      if [ ${CRM_alert_target_rc} = "0" ]; then
          CRM_alert_target_rc=""
      else
          CRM_alert_target_rc=" (target: ${CRM_alert_target_rc})"
      fi
      case ${CRM_alert_desc} in
          Cancelled)
             ;;
          *)
             sns_subject="${CRM_alert_timestamp}: Resource operation
 '${CRM_alert_task}${CRM_alert_interval}' for '${CRM_alert_rsc}' on
 '${CRM_alert_node}': ${CRM_alert_desc}${CRM_alert_target_rc}"
             ;;
      esac
      ;;
    attribute)
      sns_subject="${CRM_alert_timestamp}: The '${CRM_alert_attribute_name}' attribute
 of the '${CRM_alert_node}' node was updated in '${CRM_alert_attribute_value}'"
      ;;
    *)
      sns_subject="${CRM_alert_timestamp}: Unhandled $CRM_alert_kind alert"
      ;;
 esac

 #Use this information to send the email.
 aws sns publish --topic-arn "${sns_subscription_arn}" --subject "${sns_subject}" --
 message "${sns_body}" --region ${REGION}
```

## Testing

We recommend scheduling regular fault scenario recovery testing at least annually, and as part of the operating system or SAP kernel updates that may impact operations. For more details on best practices for regular testing, see SAP Lens – [Best Practice 4.3 – Regularly test business continuity plans and fault recovery](#).

The tests described here simulate failures. These can help you understand the behavior and operational requirements of your cluster.

In addition to checking the state of cluster resources, ensure that the service you are trying to protect is in the required state. Can you still connect to SAP? Are locks still available in SM12?

Define the recovery time to ensure that it aligns with your business objectives. Record recovery actions in runbooks.

**Topics**

- [Test 1: Stop SAP ASE database using sapcontrol](#)
- [Test 2: Unmount FSx for ONTAP file system on primary host](#)
- [Test 3: Kill the database processes on the primary host](#)
- [Test 4: Simulate hardware failure of an individual node, and repeat for other node](#)
- [Test 5: Simulate a network failure](#)
- [Test 6: Simulate an NFS failure](#)
- [Test 7: Accidental shutdown](#)

## Test 1: Stop SAP ASE database using `sapcontrol`

**Simulate failure** – On primary host as root:

```
/usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname ASD -dbtybe syb -force
```

**Expected behavior** – SAP ASE database is stopped, and the `SAPDatabase` resource agent enters a failed state. The cluster will failover the database to the secondary instance.

**Recovery action** – No action required.

## Test 2: Unmount FSx for ONTAP file system on primary host

**Simulate failure** – On primary host as root:

```
umount -l /sybase/ASD/sapdata_1
```

**Expected behavior** – The `rsc_fs` resource enters a failed state. The cluster stops the SAP ASE database, and will failover to the secondary instance.

**Recovery action** – No action required.

## Test 3: Kill the database processes on the primary host

**Simulate failure** – On primary host as `syb<sid>`:

```
ps -ef |grep -i sybaasd
kill -9 <PID>
```

**Expected behavior** – SAP ASE database fails, and the `SAPDatabase ` resource enters a failed state. The cluster will failover the database to the secondary instance.

**Recovery action** – No action required.

## Test 4: Simulate hardware failure of an individual node, and repeat for other node

**Notes** – To simulate a system crash, you must first ensure that `/proc/sys/kernel/sysrq` is set to 1.

**Simulate failure** – On the primary host as root:

```
echo 'c' > /proc/sysrq-trigger
```

**Expected behavior** – The node which has been killed fails. The cluster moves the resource (SAP ASE database) that was running on the failed node to the surviving node.

**Recovery action** – Start the EC2 node.

## Test 5: Simulate a network failure

**Notes** – See the following list.

- Iptables must be installed.
- Use a subnet in this command because of the secondary ring.
- Check for any existing iptables rules as iptables -F will flush all rules.
- Review pcmk_delay and priority parameters if you see neither node survives the fence race.

**Simulate failure** – On the primary host as root:

```
iptables -A INPUT -s <CIDR_of_other_subnet> -j DROP; iptables -A OUTPUT -d
  <CIDR_of_other_subnet> -j DROP
```

**Expected behavior** – The cluster detects the network failure, and fences one of the nodes to avoid a split-brain situation.

**Recovery action** – If the node where the command was run survives, execute iptables -F to clear the network failure. Start the EC2 node.

## Test 6: Simulate an NFS failure

**Notes** – See the following list.

- Iptables must be installed.

- Check for any existing iptables rules as iptables -F will flush all rules.

- Although rare, this is an important scenario to test. Depending on the activity it may take some time (10 min +) to notice that I/O to EFS is not occurring and fail either the Filesystem or SAP resources.

**Simulate failure** – On the primary host as root:

```
iptables -A OUTPUT -p tcp --dport 2049 -m state --state NEW,ESTABLISHED,RELATED -j
  DROP; iptables -A INPUT -p tcp --sport 2049 -m state --state ESTABLISHED -j DROP
```

**Expected behavior** – The cluster detects that NFS is not available, and the SAPDatabase resource agent fails, and moves to the FAILED state.

**Recovery action** – If the node where the command was run survives, execute iptables -F to clear the network failure. Start the EC2 node.

## Test 7: Accidental shutdown

**Notes** – See the following list.

- Avoid shutdowns without cluster awareness.

- We recommend the use of systemd to ensure predictable behaviour.

- Ensure the resource dependencies are in place.

**Simulate failure** – Login to Amazon Management Console, and stop the instance or issue a shutdown command.

**Expected behavior** – The node which has been shut down fails. The cluster moves the resource (SAP ASE database) that was running on the failed node to the surviving node. If systemd and resource dependencies are not configured, you may notice that while the EC2 instance is shutting down gracefully, the cluster will detect an unclean stop of cluster services on the node and will fence the EC2 instance being shut down.

**Recovery action** – Start the EC2 node and pacemaker service.

# SAP ASE for SAP NetWeaver on Amazon: high availability configuration for Red Hat Enterprise Linux (RHEL) for SAP applications

This topic applies to Red Hat Enterprise Linux (RHEL) for SAP with high availability and update services operating system for SAP NetWeaver running on SAP ASE database on Amazon cloud. It covers the instructions for configuration of a pacemaker cluster for SAP ASE database when deployed on Amazon EC2 instances in two different Availability Zones within an Amazon Region.

This topic covers the implementation of high availability using the cold standby method. For more information, see SAP Note 1650511 – SYB: High Availability Offerings with SAP Adaptive Server Enterprise (requires SAP portal access).

**Topics**

- Planning
- Architecture diagram
- Deployment
- Operations

# Planning

This section covers the following topics.

**Topics**

- Prerequisites
- Reliability
- SAP and Red Hat references
- Concepts

## Prerequisites

You must meet the following prerequisites before commencing setup.

**Topics**

- [Deployed cluster infrastructure](#)
- [Supported operating system](#)
- [Required access for setup](#)

## Deployed cluster infrastructure

Ensure that your Amazon networking requirements and Amazon EC2 instances where SAP workloads are installed, are correctly configured for SAP. For more information, see [SAP NetWeaver Environment Setup for Linux on Amazon](#).

See the following SAP ASE cluster specific requirements.

- Two cluster nodes created in private subnets in separate Availability Zones within the same Amazon VPC and Amazon Region
- Access to the route table(s) that are associated with the chosen subnets

  For more information, see [Amazon – Overlay IP](#).
- Targeted Amazon EC2 instances must have connectivity to the Amazon EC2 endpoint via internet or a Amazon VPC endpoint.

## Supported operating system

Protecting SAP ASE database with a pacemaker cluster requires packages from Red Hat, including targeted cluster resource agents for SAP and Amazon that may not be available in standard repositories.

SAP and Red Hat recommend the use of Red Hat Enterprise Linux for SAP. Starting with Red Hat Enterprise Linux 8 (RHEL 8), either RHEL for SAP Solutions or RHEL for SAP Applications are required for running SAP applications in production environments. See [SAP Note 1656099 - SAP Applications on Amazon: Supported DB/OS and Amazon EC2 products](#) (requires SAP portal access).

Built on the Red Hat Enterprise Linux operating system, Red Hat Enterprise Linux for SAP expands existing capabilities, lso you can get the most out of SAP's powerful analytics and data management portfolio. See [Red Hat Enterprise Solutions for SAP](#) product page from Red Hat.

Red Hat Enterprise Linux High Availability (HA) provides all the necessary packages for configuring pacemaker-based clusters. Extended Update Support (E4S) provides support on specific minor releases for 4 years from general availability.

Red Hat Enterprise Linux for SAP with HA and US is available on [Amazon Marketplace](#) under an hourly or an annual subscription model or can be accessed using a BYOS subscription model.

## Required access for setup

The following access is required for setting up the cluster.

- An IAM user with the following privileges.
  - modify Amazon VPC route tables
  - modify Amazon EC2 instance properties
  - create IAM policies and roles
  - create Amazon EFS file systems
- Root access to the operating system of both cluster nodes
- SAP administrative user access – `<syb>adm`

  In case of a new install, this user is created by the install process.

# Reliability

SAP ASE database is a single point of failure in a highly available SAP architecture. We recommend evaluating the impact of design decisions on cost, operation, and reliability. For more information, see [Reliability](#) in SAP Lens - Amazon Well-Architected Framework.

# SAP and Red Hat references

In addition to this guide, see the following references for more details.

- Red Hat: [Is there a High Availability resource agent for SAP (Sybase) ASE database, and how can I configure it in a Red Hat Enterprise Linux HA Cluster?](#)
- Red Hat: [Red Hat Enterprise Linux for SAP offerings on Amazon Web Services FAQ](#)
- [SAP Note: 1656099 - SAP Applications on Amazon: Supported DB/OS and Amazon EC2 products](#)
- [SAP Note: 1618572 - Linux: Support Statement for RHEL on Amazon Web Services](#)
- [SAP Note: 2002167 - Red Hat Enterprise Linux 7.x: Installation and Upgrade](#)
- [SAP Note: 2772999 - Red Hat Enterprise Linux 8.x: Installation and Upgrade](#)

You must have SAP portal access for reading all SAP Notes.

# Concepts

This section covers Amazon, SAP, and Red Hat concepts.

**Topics**

- Amazon – Availability Zones
- Amazon – Overlay IP
- Amazon – Shared VPC
- Amazon FSx for NetApp ONTAP
- Pacemaker - STONITH fencing agent

## Amazon – Availability Zones

Availability Zone is one or more data centers with redundant power, networking, and connectivity in an Amazon Region. For more information, see Regions and Availability Zones.

For mission critical deployments of SAP on Amazon where the goal is to minimise the recovery time objective (RTO), we suggest distributing single points of failure across Availability Zones. Compared with single instance or single Availability Zone deployments, this increases resilience and isolation against a broad range of failure scenarios and issues, including natural disasters.

Each Availability Zone is physically separated by a meaningful distance (many kilometers) from another Availability Zone. All Availability Zones in an Amazon Region re interconnected with high-bandwidth, low-latency network, over fully redundant, dedicated metro fiber. This enables synchronous replication. All traffic between Availability Zones is encrypted.

## Amazon – Overlay IP

Overlay IP enables a connection to the application, regardless of which Availability Zone (and subnet) contains the active primary node.

When deploying instances in Amazon, it is necessary to allocate an IP from a pre-existing subnet. Subnets have a classless inter-domain routing (CIDR) IP assignment from the VPC which resides entirely within one Availability Zone. This CIDR IP assignment cannot span multiple Availability Zones or be reassigned to an instance in a different Availability Zone after faults, including network connectivity or hardware issues which require a failover to the replication target.

To address this, we suggest that you configure an overlay IP, and use this in the connection parameters for the application. This IP address is a non-overlapping RFC1918 private IP address

from outside of VPC CIDR block and is configured as an entry in the route table or tables. The route directs the connection to the active node and is updated during a failover by the cluster software.

You can select any one of the following RFC1918 private IP addresses for your overlay IP address.

- 10.0.0.0 – 10.255.255.255 (10/8 prefix)
- 172.16.0.0 – 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 – 192.168.255.255 (192.168/16 prefix)

If you use the 10/8 prefix in your SAP VPC, selecting a 172 or a 192 IP address may help to differentiate the overlay IP. Consider the use of an IP Address Management (IPAM) tool such as Amazon VPC IP Address Manager to plan, track, and monitor IP addresses for your Amazon workloads. For more information, see What is IPAM?

The overlay IP agent in the cluster can also be configured to update multiple route tables which contain the Overlay IP entry if your subnet association or connectivity requires it.

**Access to overlay IP**

The overlay IP is outside of the range of the VPC, and therefore cannot be reached from locations that are not associated with the route table, including on-premises and other VPCs.

Use Amazon Transit Gateway as a central hub to facilitate the network connection to an overlay IP address from multiple locations, including Amazon VPCs, other Amazon Regions, and on-premises using Amazon Direct Connect or Amazon Client VPN.

If you do not have Amazon Transit Gateway set up as a network transit hub or if it is not available in your preferred Amazon Region, you can use a Network Load Balancer to enable network access to an overlay IP.

For more information, see SAP on Amazon High Availability with Overlay IP Address Routing.

## Amazon – Shared VPC

An enterprise landing zone setup or security requirements may require the use of a separate cluster account to restrict the route table access required for the Overlay IP to an isolated account. For more information, see Share your VPC with other accounts.

Evaluate the operational impact against your security posture before setting up shared VPC. To set up, see Shared VPC – optional.

## Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed Amazon service.

FSx for ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx for ONTAP supports popular data security and anti-virus applications. For more information, see [What is Amazon FSx for NetApp ONTAP?](#)

## Pacemaker - STONITH fencing agent

In a two-node cluster setup for a primary resource and its replication pair, it is important that there is only one node in the primary role with the ability to modify your data. In the event of a failure scenario where a node is unresponsive or incommunicable, ensuring data consistency that can require you to isolate the faulty node by powering it down before the cluster commences other actions, such as promoting a new primary. This arbitration is the role of the fencing agent.

Since a two-node cluster introduces the possibility of a fence race in which a dual shoot out can occur with communication failures resulting in both nodes simultaneously claiming, "I can't see you, so I am going to power you off". The fencing agent is designed to minimise this risk by providing an external witness.

Red Hat supports several fencing agents, including the one recommended for use with Amazon EC2 Instances (`fence_aws`). This resource uses API commands to check its own instance status - "Is my instance state anything other than running?" before proceeding to power off its pair. If it is already in a stopping or stopped state it will admit defeat and leave the surviving node untouched.

# Architecture diagram

The following diagram shows the cold standby SAP ASE cluster setup with FSx for ONTAP.

RHEL for SAP – Pacemaker Cluster for SAP ASE

# Deployment

This section covers the following topics.

**Topics**

- [Settings and prerequisites](#)
- [SAP and cluster setup](#)
- [Cluster configuration](#)

## Settings and prerequisites

The cluster setup uses parameters, including `DBSID` that is unique to your setup. It is useful to predetermine the values with the following examples and guidance.

**Topics**

- [Define reference parameters for setup](#)
- [Amazon EC2 instance settings](#)

- [Operating system prerequisites](#)

- [IP and hostname resolution prerequisites](#)

- [FSx for ONTAP prerequisites](#)

- [Shared VPC – optional](#)

## Define reference parameters for setup

The cluster setup relies on the following parameters.

**Topics**

- [Global Amazon parameters](#)

- [Amazon EC2 instance parameters](#)

- [SAP and Pacemaker resource parameters](#)

- [RHEL cluster parameters](#)

### Global Amazon parameters

| Name | Parameter | Example |
| --- | --- | --- |
| Amazon account ID | `<account_id>` | `123456789100` |
| Amazon Region | `<region_id>` | `us-east-1` |

- Amazon account – For more details, see [Your Amazon account ID and its alias](#).

- Amazon Region – For more details, see [Describe your Regions](#).

### Amazon EC2 instance parameters

| Name | Parameter | Primary example | Secondary example |
| --- | --- | --- | --- |
| Amazon EC2 instance ID | `<instance_id>` | `i-xxxxins tidforhost1` | `i- xxxxinsti dforhost2` |
| Hostname | `<hostname>` | `rhxdbhost01` | `rhxdbhost02` |

| Name | Parameter | Primary example | Secondary example |
|------|-----------|-----------------|-------------------|
| Host IP | `<host_ip>` | `10.1.10.1` | `10.1.20.1` |
| Host additional IP | `<host_add itional_ip>` | `10.1.10.2` | `10.1.20.2` |
| Configured subnet | `<subnet_id>` | `subnet-xx xxxxxxxxs ubnet1` | `subnet-xx xxxxxxxxs ubnet2` |

- Hostname – Hostnames must comply with SAP requirements outlined in SAP Note 611361 - Hostnames of SAP ABAP Platform servers (requires SAP portal access).

  Run the following command on your instances to retrieve the hostname.

  ```
  hostname
  ```

- Amazon EC2 instance ID – run the following command (IMDSv2 compatible) on your instances to retrieve instance metadata.

  ```
  /usr/bin/curl --noproxy '*' -w "\n" -s -H "X-aws-ec2-metadata-token: $(curl --noproxy
   '*' -s -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-
  token-ttl-seconds: 21600")" http://169.254.169.254/latest/meta-data/instance-id
  ```

  For more details, see Retrieve instance metadata and Instance identity documents.

**SAP and Pacemaker resource parameters**

| Name | Parameter | Example |
|------|-----------|---------|
| DBSID | `<DBSID>` or `<dbsid>` | ARD |
| Virtual hostname | `<db_virt_hostname>` | rhxvdb |
| Database Overlay IP | `<ase_db_oip>` | 172.16.0.23 |
| VPC Route Tables | `<rtb_id>` | rtb-xxxxxroutetable1 |

| Name | Parameter | Example |
|------|-----------|---------|
| FSx for ONTAP mount points | `<ase_db_fs>` | `svm-xxx.fs-xxx.fsx .us- east-1.am azonaws.com` |

- SAP details – SAP parameters, including SID and instance number must follow the guidance and limitations of SAP and Software Provisioning Manager. Refer to [SAP Note 1979280 - Reserved SAP System Identifiers (SAPSID) with Software Provisioning Manager](#) for more details.

  Post-installation, use the following command to find the details of the instances running on a host.

  ```
  sudo /usr/sap/hostctrl/exe/saphostctrl -function ListDatabases
  ```

- Overlay IP – This value is defined by you. For more information, see [Overlay IP](#).

- FSx for ONTAP mount points – This value is defined by you. Consider the required mount points specified in [SAP ASE on Amazon with Amazon FSx for NetApp ONTAP](#).

**RHEL cluster parameters**

| Name | Parameter | Example |
|------|-----------|---------|
| Cluster name | `cluster_name` | `rhelha` |
| Cluster user | `cluster_user` | `hacluster` |
| Cluster password | `cluster_password` | |

# Amazon EC2 instance settings

Amazon EC2 instance settings can be applied using Infrastructure as Code or manually using Amazon Command Line Interface or Amazon Management Console. We recommend Infrastructure as Code automation to reduce manual steps, and ensure consistency.

**Topics**

- [Create IAM roles and policies](#)

- [Amazon Overlay IP policy](#)

- [Assign IAM role](#)

- [Modify security groups for cluster communication](#)

- [Disable source/destination check](#)

- [Review automatic recovery and stop protection](#)

## Create IAM roles and policies

In addition to the permissions required for standard SAP operations, two IAM policies are required for the cluster to control Amazon resources on ASCS. These policies must be assigned to your Amazon EC2 instance using an IAM role. This enables Amazon EC2 instance, and therefore the cluster to call Amazon services.

Create these policies with least-privilege permissions, granting access to only the specific resources that are required within the cluster. For multiple clusters, you need to create multiple policies.

For more information, see [IAM roles for Amazon EC2](#).

### STONITH policy

The RHEL STONITH agent requires permission to start and stop both the nodes of the cluster. Create a policy as shown in the following example. Attach this policy to the IAM role assigned to both Amazon EC2 instances in the cluster.

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeInstances",
                "ec2:DescribeTags"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:StartInstances",
```

```
            "ec2:StopInstances"
          ],
          "Resource": [
            "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0",
            "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0"
          ]
        }
    ]
}
```

**Amazon Overlay IP policy**

The RHEL Overlay IP resource agent (`aws-vpc-move-ip`) requires permission to modify a routing entry in route tables. Create a policy as shown in the following example. Attach this policy to the IAM role assigned to both Amazon EC2 instances in the cluster.

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ec2:ReplaceRoute",
            "Resource": [
                "arn:aws:ec2:us-east-1:123456789012:route-table/
rtb-0123456789abcdef0",
                "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0"
                ]
        },
        {
            "Effect": "Allow",
            "Action": "ec2:DescribeRouteTables",
            "Resource": "*"
        }
    ]
}
```

> (i) **Note**
>
>   If you are using a Shared VPCs, see Shared VPC – optional.

---

## Assign IAM role

The two cluster resource IAM policies must be assigned to an IAM role associated with your Amazon EC2 instance. If an IAM role is not associated to your instance, create a new IAM role for cluster operations. To assign the role, go to https://console.aws.amazon.com/ec2/, select each or both instance(s), and then choose **Actions** > **Security** > **Modify IAM role**.

## Modify security groups for cluster communication

A security group controls the traffic that is allowed to reach and leave the resources that it is associated with. For more information, see [Control traffic to your Amazon resources using security groups](#).

In addition to the standard ports required to access SAP and administrative functions, the following rules must be applied to the security groups assigned to both Amazon EC2 instances in the cluster.

**Inbound**

| Source | Protocol | Port range | Description |
|---|---|---|---|
| The security group ID (its own resource ID) | **UDP** | 5405 | Allows UDP traffic between cluster resources for corosync communication |

> ⓘ **Note**
>
> Note the use of the UDP protocol.

If you are running a local firewall, such as `iptables`, ensure that communication on the preceding ports is allowed between two Amazon EC2 instances.

## Disable source/destination check

Amazon EC2 instances perform source/destination checks by default, requiring that an instance is either the source or the destination of any traffic it sends or receives.

In the pacemaker cluster, source/destination check must be disabled on both instances receiving traffic from the Overlay IP. You can disable check using Amazon CLI or Amazon Management Console.

**Example**

Amazon CLI

Use the [modify-instance-attribute](#) command to disable source/destination check.

Run the following commands on both instances in the cluster.

- Primary example –

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost1> --no-
source-dest-check
```

- Secondary example –

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost2> --no-
source-dest-check
```

Amazon Management Console

Ensure that the **Stop** option is checked in https://console.aws.amazon.com/ec2/.

**Review automatic recovery and stop protection**

After a failure, cluster-controlled operations must be resumed in a coordinated way. This helps ensure that the cause of failure is known and addressed, and the status of the cluster is as expected. For example, verifying that there are no pending fencing actions.

This can be achieved by not enabling pacemaker to run as a service at the operating system level or by avoiding auto restarts for hardware failure.

If you want to control the restarts resulting from hardware failure, disable simplified automatic recovery and do not configure Amazon CloudWatch action-based recovery for Amazon EC2 instances that are part of a pacemaker cluster. Use the following commands on both Amazon EC2 instances in the pacemaker cluster, to disable simplified automatic recovery via Amazon CLI.

If making the change via Amazon CLI, run the command for both Amazon EC2 instances in the cluster.

> **ⓘ Note**
>
> Modifying instance maintenance options will require admin privileges not covered by the IAM instance roles defined for operations of the cluster.

```
aws ec2 modify-instance-maintenance-options --instance-id <i-xxxxinstidforhost1> --
auto-recovery disabled
```

```
aws ec2 modify-instance-maintenance-options --instance-id <i-xxxxinstidforhost2> --
auto-recovery disabled
```

To ensure that STONITH actions can be executed, you must ensure that stop protection is disabled for Amazon EC2 instances that are part of a pacemaker cluster. If the default settings have been modified, use the following commands for both instances to disable stop protection via Amazon CLI.

> **ⓘ Note**
>
> Modifying instance attributes will require admin privileges not covered by the IAM instance roles defined for operations of the cluster.

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost> --no-disable-api-
stop
```

```
aws ec2 modify-instance-attribute --instance-id <i-xxxxinstidforhost2> --no-disable-
api-stop
```

## Operating system prerequisites

This section covers the following topics.

**Topics**

- [Root access](#)

- [Install missing operating system packages](#)

- [Update and check operating system versions](#)

- [Stop and disable nm-cloud-setup](#)

- [Time synchronization services](#)

- [Amazon CLI profile](#)

- [Pacemaker proxy settings](#)

**Root access**

Verify root access on both cluster nodes. The majority of the setup commands in this document are performed with the root user. Assume that commands should be run as root unless there is an explicit call out to choose otherwise.

**Install missing operating system packages**

This is applicable to both cluster nodes. You must install any missing operating system packages.

The following packages and their dependencies are required for the pacemaker setup. Depending on your baseline image, for example, RHEL for SAP, these packages may already be installed.

```
awscli
chrony
corosync
pcs
pacemaker
fence-agents-aws
resource-agents-sap (Version resource-agents-sap-3.9.5-124.el7.x86_64 or higher)
sap-cluster-connector
```

We highly recommend installing the following additional packages for troubleshooting.

```
sysstat
pcp-system-tools
sos
```

See Red Hat documentation [What are all the Performance Co-Pilot (PCP) RPM packages in RHEL?](#)

> ⓘ **Note**
>
> The preceding list of packages is not a complete list required for running SAP applications. For the complete list, see [SAP and Red Hat references](#).

Use the following command to check packages and versions.

```
for package in awscli chrony corosync pcs pacemaker fence-agents-aws resource-agents-
sap sap-cluster-connector sysstat pcp-system-tools sos; do
echo "Checking if ${package} is installed..."
RPM_RC=$(rpm -q ${package} --quiet; echo $?)
if [ ${RPM_RC} -ne 0 ];then
echo "   ${package} is missing and needs to be installed"
fi
done
```

If a package is not installed, and you are unable to install it using `yum`, it may be because Red Hat Enterprise Linux for SAP extension is not available as a repository in your chosen image. You can verify the availability of the extension using the following command.

```
yum repolist
```

To install or update a package or packages with confirmation, use the following command.

```
yum install <package_name(s)>
```

**Update and check operating system versions**

You must update and confirm versions across nodes. Apply all the latest patches to your operating system versions. This ensures that bugs are addresses and new features are available.

You can update the patches individually or use the `yum update`. A clean reboot is recommended prior to setting up a cluster.

```
yum update
reboot
```

Compare the operating system package versions on the two cluster nodes and ensure that the versions match on both nodes.

## Stop and disable `nm-cloud-setup`

This is applicable on both cluster nodes. If you are using Red Hat 8.6 or later, the following services must be stopped and disabled on both the cluster nodes. This prevents the NetworkManager from removing the overlay IP address from the network interface.

```
systemctl disable nm-cloud-setup.timer
systemctl stop nm-cloud-setup.timer
systemctl disable nm-cloud-setup
systemctl stop nm-cloud-setup
```

## Time synchronization services

This is applicable to both cluster nodes. Time synchronization is important for cluster operation. Ensure that `chrony rpm` is installed, and configure appropriate time servers in the configuration file.

You can use Amazon Time Sync Service that is available on any instance running in a VPC. It does not require internet access. To ensure consistency in the handling of leap seconds, don't mix Amazon Time Sync Service with any other `ntp` time sync servers or pools.

Create or check the `/etc/chrony.d/ec2.conf` file to define the server.

```
# Amazon EC2 time source config
server 169.254.169.123 prefer iburst minpoll 4 maxpoll 4
```

Start the `chronyd.service`, using the following command.

```
systemctl enable --now chronyd.service
systemctl status chronyd
```

For more information, see [Set the time for your Linux instance](#).

## Amazon CLI profile

This is applicable to both cluster nodes. The cluster resource agents use Amazon Command Line Interface (Amazon CLI). You need to create an Amazon CLI profile for the root account on both instances.

You can either edit the config file at `/root/.aws` manually or by using [aws configure](#) Amazon CLI command.

You can skip providing the information for the access and secret access keys. The permissions are provided through IAM roles attached to Amazon EC2 instances.

```
aws configure
{aws} Access Key ID [None]:
{aws} Secret Access Key [None]:
Default region name [None]: <region_id>
Default output format [None]:
```

**Pacemaker proxy settings**

This is applicable to both cluster nodes. If your Amazon EC2 instance has been configured to access the internet and/or Amazon Cloud through proxy servers, then you need to replicate the settings in the pacemaker configuration. For more information, see Use an HTTP proxy.

Add the following lines to /etc/sysconfig/pacemaker.

```
http_proxy=http://<proxyhost>:<proxyport>
https_proxy= http://<proxyhost>:<proxyport>
no_proxy=127.0.0.1,localhost,169.254.169.254,fd00:ec2::254
```

Modify `proxyhost` and `proxyport` to match your settings. Ensure that you exempt the address used to access the instance metadata. Configure `no_proxy` to include the IP address of the instance metadata service – **169.254.169.254** (IPV4) and **fd00:ec2::254** (IPV6). This address does not vary.

# IP and hostname resolution prerequisites

This section covers the following topics.

**Topics**

- Add initial VPC route table entries for overlay IPs
- Add overlay IPs to host IP configuration
- Hostname resolution

**Add initial VPC route table entries for overlay IPs**

You need to add initial route table entries for overlay IPs. For more information on overlay IP, see Overlay IP.

Add entries to the VPC route table or tables associated with the subnets of your Amazon EC2 instance for the cluster. The entries for destination (overlay IP CIDR) and target (Amazon EC2 instance or ENI) must be added manually for ASCS and ERS. This ensures that the cluster resource has a route to modify. It also supports the install of SAP using the virtual names associated with the overlay IP before the configuration of the cluster.

**Modify or add a route to a route table using Amazon Management Console**

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.

2. In the navigation pane, choose **Route Tables**, and select the route table associated with the subnets where your instances have been deployed.

3. Choose **Actions**, **Edit routes**.

4. To add a route, choose **Add route**.

5. Add your chosen overlay IP address CIDR and the instance ID of your primary instance for SAP ASE database. See the following table for an **example**.

| Destination | 172.16.0.23/32 |
| --- | --- |
| Target | i-xxxxinstidforhost1 |

6. Choose **Save changes**.

The preceding steps can also be performed programmatically. We suggest performing the steps using administrative privileges, instead of instance-based privileges to preserve least privilege. CreateRoute API isn't necessary for ongoing operations.

Run the following command as a dry run on both nodes to confirm that the instances have the necessary permissions.

```
aws ec2 replace-route --route-table-id <rtb-xxxxxroutetable1> --destination-cidr-
block <172.16.0.23/32> --instance-id <i-xxxxinstidforhost1> --dry-run --profile
 <aws_cli_cluster_profile>
```

**Add overlay IPs to host IP configuration**

You must configure the overlay IP as an additional IP address on the standard interface to enable SAP install. This action is managed by the cluster IP resource. However, to install SAP using the

correct IP addresses prior to having the cluster configuration in place, you need to add these entries manually.

If you need to reboot the instance during setup, the assignment is lost, and must be re-added.

See the following **examples**. You must update the commands with your chosen IP addresses.

On EC2 instance 1, where you are installing SAP ASE database, add the overlay IP.

```
ip addr add <172.16.0.23/32> dev eth0
```

**Hostname resolution**

This is applicable to both cluster nodes. You must ensure that both instances can resolve all hostnames in use. Add the hostnames for cluster nodes to `/etc/hosts` file on both cluster nodes. This ensures that hostnames for cluster nodes can be resolved even in case of DNS issues. See the following example.

```
cat /etc/hosts
<10.1.10.1 rhxdbhost01.example.com rhxdbhost01>
<10.1.20.1 rhxdbhost02.example.com rhxdbhost02>
<172.16.0.23 rhxvdb.example.com rhxvdb>
```

> ⚠️ **Important**
>
> The overlay IP is out of VPC range, and cannot be reached from locations not associated with the route table, including on-premises.

## FSx for ONTAP prerequisites

This section covers the following topics.

**Topics**

- [Shared file systems](#)
- [Create volumes and file systems](#)

**Shared file systems**

Amazon FSx for NetApp ONTAP is supported for SAP ASE database file systems.

FSx for ONTAP provides fully managed shared storage in Amazon Cloud with data access and management capabilities of ONTAP. For more information, see Create an Amazon FSx for NetApp ONTAP file system.

Select a file system based on your business requirements, evaluating the resilience, performance, and cost of your choice.

The SVM's DNS name is your simplest mounting option. The file system DNS name automatically resolves to the mount target's IP address on the Availability ZOne of the connecting Amazon EC2 instance.

`svm-id.fs-id.fsx.aws-region.amazonaws.com`

> ⓘ **Note**
>
> Review the `enableDnsHostnames` and `enableDnsSupport` DNS attributes for your VPC. For more information, see View and update DNS attributes for your VPC.

## Create volumes and file systems

You can review the following resources to understand the FSx for ONTAP mount points for SAP ASE database.

- Host setup for SAP ASE
- SAP – Setup of Database Layout (ABAP)
- SAP – Setup of Database Layout (JAVA)

The following are the FSx for ONTAP mount points covered in this topic.

| Unique NFS Location (example) | File system location |
| --- | --- |
| SVM-xxx:/sybase | /sybase |
| SVM-xxx:/asedata | /sybase/<DBSID>/sapdata_1 |
| SVM-xxx:/aselog | /sybase/<DBSID>/saplog_1 |
| SVM-xxx:/sapdiag | /sybase/<DBSID>/sapdiag |

| Unique NFS Location (example) | File system location |
|---|---|
| SVM-xxx:/saptmp | /sybase/<DBSID>/saptmp |
| SVM-xxx:/backup | /sybasebackup |
| SVM-xxx:/usrsap | /usr/sap |

Ensure that you have properly mounted the file systems, and the necessary adjustments for host setup have been performed. See [Host setup for SAP ASE](#). You can temporarily add the entries to `/etc/fstab` to not lose them during a reboot. The entries must be removed prior to configuring the cluster. The cluster resource manages the mounting of the NFS.

You need to perform this step only on the primary Amazon EC2 instance for the initial installation.

Review the mount options to ensure that they match with your operating system, NFS file system type, and SAP's latest recommendations.

Use the following command to check that the required file systems are available.

```
df -h
```

## Shared VPC – *optional*

Amazon VPC sharing enables you to share subnets with other Amazon accounts within the same Amazon Organizations. Amazon EC2 instances can be deployed using the subnets of the shared Amazon VPC.

In the pacemaker cluster, the `aws-vpc-move-ip` resource agent has been enhanced to support a shared VPC setup while maintaining backward compatibility with previous existing features.

The following checks and changes are required. We refer to the Amazon account that owns Amazon VPC as the sharing VPC account, and to the consumer account where the cluster nodes are going to be deployed as the cluster account.

This section covers the following topics.

**Topics**

- [Minimum version requirements](#)

- [IAM roles and policies](#)

- [Shared VPC cluster resources](#)

**Minimum version requirements**

The latest version of the `aws-vpc-move-ip` agent shipped with Red Hat 8.2 supports the shared VPC setup by default. The following are the minimum version required to support a shared VPC Setup:

- Red Hat 7.9 - resource-agents-4.1.1-61.10

- Red Hat 8.1 - resource-agents-4.1.1-33.10

- Red Hat 8.2 - resource-agents-4.1.1-44.12

**IAM roles and policies**

Using the overlay IP agent with a shared Amazon VPC requires a different set of IAM permissions to be granted on both Amazon accounts (sharing VPC account and cluster account).

**Sharing VPC account**

In sharing VPC account, create an IAM role to delegate permissions to the EC2 instances that will be part of the cluster. During the IAM Role creation, select "Another Amazon account" as the type of trusted entity, and enter the Amazon account ID where the EC2 instances will be deployed/running from.

After the IAM role has been created, create the following IAM policy on the sharing VPC account, and attach it to an IAM role. Add or remove route table entries as needed.

```
{
    "Version":"2012-10-17",
    "Statement": [
      {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": "ec2:ReplaceRoute",
        "Resource": [
            "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0",
            "arn:aws:ec2:us-east-1:123456789012:route-table/rtb-0123456789abcdef0"
        ]
      },
```

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

Next, edit move to the "Trust relationships" tab in the IAM role, and ensure that the Amazon account you entered while creating the role has been correctly added.

## Cluster account

In cluster account, create the following IAM policy, and attach it to an IAM role. This is the IAM Role that is going to be attached to the EC2 instances.

### Amazon STS policy

```
{
    "Version":"2012-10-17",
    "Statement": [
      {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": "sts:AssumeRole",
        "Resource": "arn:aws:iam::123456789012:role/sharing-vpc-account-cluster-role"
      }
    ]
}
```

### STONITH policy

```
{
    "Version":"2012-10-17",
    "Statement": [
      {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": [
            "ec2:StartInstances",
            "ec2:StopInstances"
        ],
```

```
        "Resource": [
            "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0",
            "arn:aws:ec2:us-east-1:123456789012:instance/arn:aws:ec2:us-
east-1:123456789012:instance/i-1234567890abcdef0"
        ]
    },
    {
        "Sid": "VisualEditor1",
        "Effect": "Allow",
        "Action": "ec2:DescribeInstances",
        "Resource": "*"
    }
    ]
}
```

**Shared VPC cluster resources**

The cluster resource agent `aws-vpc-move-ip` also uses a different configuration syntax. When configuring the `aws-vpc-move-ip` resource agent, the following new parameters must be used:

- lookup_type=NetworkInterfaceId
- routing_table_role="arn:aws:iam::<account_id>:role/<VPC-Account-Cluster-Role>"

The following IP Resource for SAP ASE database needs to be created.

```
pcs resource create rsc_ip_ARD_ASEDB ocf:heartbeat:aws-vpc-move-ip ip=172.16.0.23
 interface=eth0  routing_table=rtb-xxxxxroutetable1 lookup_type=NetworkInterfaceId
  routing_table_role="arn:aws:iam::<sharing_vpc_account_id>:role/
<sharing_vpc_account_cluster_role>" op monitor interval=20s timeout=40s --group
 rsc_asedb_group
```

# SAP and cluster setup

This section covers the following topics.

**Topics**

- [Install SAP](#)
- [Cluster prerequisites](#)
- [Create cluster and node associations](#)

# Install SAP

The following topics provide information about installing SAP ASE database on Amazon Cloud in a highly available cluster. Review SAP Documentation for more details.

**Topics**

- [Use SWPM with high availability](#)
- [Install SAP database instance](#)
- [Check SAP host agent version](#)

## Use SWPM with high availability

Before running SAP Software Provisioning Manager (SWPM), ensure that the following prerequisites are met.

- If the operating system groups for SAP are pre-defined, ensure that the user identifier (UID) and group identifier (GID) values for `<syb>adm`, `sapadm`, and `sapsys` are consistent across both instances.
- You have downloaded the most recent version of Software Provisioning Manager for your SAP version. For more information, see SAP Documentation [Software Provisioning Manager](#).
- Ensure that routes, overlay IPs, and virtual host names are mapped to the instance where the installation will run. This is to ensure that the virtual hostname for SAP ASE database is available on the primary instance. For more information, see [IP and hostname resolution prerequisites](#).
- Ensure that FSx for ONTAP mount points are available, either in `/etc/fstab` or using the mount command. For more information, see [File system prerequisites](#). If you are adding the entries in `/etc/fstab`, ensure that they are removed before configuring the cluster.

## Install SAP database instance

The commands in this section use the example values provided in [Define reference parameters for setup](#).

Install SAP ASE database on `<rhxdbhost01>` with virtual hostname `rhxvdb`, using the high availability option of Software Provisioning Manager (SWPM) tool. You can use the `SAPINST_USE_HOSTNAME` parameter to install SAP using a virtual hostname.

```
<swpm location>/sapinst SAPINST_USE_HOSTNAME=<rhxvdb>
```

> ⓘ **Note**
>
> Before installing SAP ASE database, ASCS and ERS must be installed, and the `/sapmnt` directory must be available on the database server.

**Check SAP host agent version**

The SAP host agent is used for ASE database instance control and monitoring. This agent is used by SAP cluster resource agents and hooks. It is recommended that you have the latest version installed on both instances. For more details, see [SAP Note 2219592 – Upgrade Strategy of SAP Host Agent](#).

Use the following command to check the version of the host agent.

```
/usr/sap/hostctrl/exe/saphostexec -version
```

# Cluster prerequisites

This section covers the following topics.

**Topics**

- [Update the hacluster password](#)
- [Setup passwordless authentication between nodes](#)

**Update the `hacluster` password**

This is applicable to both cluster nodes. Change the password of the operating system user `hacluster` using the following command.

```
passwd hacluster
```

**Setup passwordless authentication between nodes**

For a more comprehensive and easily consumable view of cluster activity, Red Hat provides additional reporting tools. Many of these tools require access to both nodes without entering a password. Red Hat recommends performing this setup for root user.

For more details, see Red Hat documentation [How to setup SSH Key passwordless login in Red Hat Enterprise Linux?](#)

## Create cluster and node associations

This section covers the following topics.

**Topics**

- [Start pcsd service](#)
- [Reset configuration – optional](#)
- [Authenticate pcs with user hacluster](#)
- [Setup node configuration](#)

### Start `pcsd` service

This is applicable on both cluster nodes. Run the following command to enable and start the cluster service `pcsd` (pacemaker/corosync configuration system daemon) on both, the primary and secondary node.

```
systemctl start pcsd.service
systemctl enable pcsd.service
```

Run the following command to check the status of cluster service.

```
systemctl status pcsd.service
# pcsd.service - PCS GUI and remote configuration interface
   Loaded: loaded (/usr/lib/systemd/system/pcsd.service; enabled; vendor preset:
 disabled)
   Active: active (running) since Fri 2023-01-13 14:15:32 IST; 7min ago
     Docs: man:pcsd(8)
           man:pcs(8)
 Main PID: 1445 (pcsd)
    Tasks: 1 (limit: 47675)
   Memory: 27.1M
   CGroup: /system.slice/pcsd.service
           ##1445 /usr/libexec/platform-python -Es /usr/sbin/pcsd
```

## Reset configuration – *optional*

> **ⓘ Note**
>
> The following instructions help you reset the complete configuration. Run these commands
> only if you want to start setup from the beginning. You can make minor changes with the
> `crm edit` command.

Run the following command to back up the current configuration for reference.

```
pcs config show > /tmp/pcsconfig_backup.txt
```

Run the following command to clear the current configuration.

```
pcs cluster destroy
```

### Authenticate `pcs` with user `hacluster`

The following command authenticates `pcs` to the `pcs` daemon on cluster nodes. It should be run
on only one of the cluster nodes. The username and password for the `pcs` user must be the same,
and the username should be <hacluster>.

**RHEL 7.x**

```
pcs cluster auth <rhxdbhost01> <rhxdbhost02>
Username: <hacluster>
Password:
rhxhost02: Authorized
rhxhost01: Authorized
```

**RHEL 8.x**

```
pcs host auth <rhxdbhost01> <rhxdbhost02>
Username: <hacluster>
Password:
rhxhost02: Authorized
rhxhost01: Authorized
```

### Setup node configuration

The following command configures the `cluster configuration` file, and syncs the configuration on both nodes. It should be run on only one of the cluster nodes.

**RHEL 7.x**

```
pcs cluster setup --name <rhelha> <rhxdbhost01> <rhxdbhost02>
Destroying cluster on nodes: <rhxdbhost01>, <rhxdbhost02>...
<rhxdbhost02>: Stopping Cluster (pacemaker)...
<rhxdbhost01>: Stopping Cluster (pacemaker)...
<rhxdbhost02>: Successfully destroyed cluster
<rhxdbhost01>: Successfully destroyed cluster

Sending 'pacemaker_remote authkey' to '<rhxdbhost01>', '<rhxdbhost02>'
<rhxdbhost01>: successful distribution of the file 'pacemaker_remote authkey'
<rhxdbhost02>: successful distribution of the file 'pacemaker_remote authkey'
Sending cluster config files to the nodes...
<rhxdbhost01>: Succeeded
<rhxdbhost02>: Succeeded

Synchronizing pcsd certificates on nodes <rhxdbhost01>, <rhxdbhost02>...
<rhxdbhost01>: Success
<rhxdbhost02>: Success
Restarting pcsd on the nodes in order to reload the certificates...
<rhxdbhost01>: Success
<rhxdbhost02>: Success.
```

**RHEL 8.x**

```
pcs cluster setup <rhelha> <rhxdbhost01> <rhxdbhost02>
        No addresses specified for host '<rhxdbhost01>', using '<rhxdbhost01>'
        No addresses specified for host '<rhxdbhost02>', using '<rhxdbhost02>'
        Destroying cluster on hosts: '<rhxdbhost01>', '<rhxdbhost02>'...
        <rhxdbhost01>: Successfully destroyed cluster
        <rhxdbhost02>: Successfully destroyed cluster
        Requesting remove 'pcsd settings' from '<rhxdbhost01>', '<rhxdbhost02>'
        <rhxdbhost01>: successful removal of the file 'pcsd settings'
        <rhxdbhost02>: successful removal of the file 'pcsd settings'
        Sending 'corosync authkey', 'pacemaker authkey' to '<rhxdbhost01>',
 '<rhxdbhost02>'
        <rhxdbhost01>: successful distribution of the file 'corosync authkey'
        <rhxdbhost01>: successful distribution of the file 'pacemaker authkey'
```

```
<rhxdbhost02>: successful distribution of the file 'corosync authkey'
<rhxdbhost02>: successful distribution of the file 'pacemaker authkey'
Sending 'corosync.conf' to '<rhxdbhost01>', '<rhxdbhost02>'
<rhxdbhost01>: successful distribution of the file 'corosync.conf'
<rhxdbhost02>: successful distribution of the file 'corosync.conf'
Cluster has been successfully set up.
```

# Cluster configuration

This section covers the following topics.

**Topics**

- Cluster resources
- Sample configuration (pcs config show)

## Cluster resources

This section covers the following topics.

**Topics**

- Enable and start the cluster
- Increase corosync totem timeout
- Check cluster status
- Prepare for resource creation
- Cluster bootstrap
- Create fence_aws STONITH resource
- Create file system resources
- Create overlay IP resources
- Create SAP ASE database resource
- Activate cluster

### Enable and start the cluster

This is applicable to both cluster nodes. Run the following command to enable and start the
`pacemaker` cluster service on both nodes.

```
pcs cluster enable --all
rhxdbhost01: Cluster Enabled
rhxdbhost02: Cluster Enabled


pcs cluster start --all
rhxdbhost01: Starting Cluster...
rhxdbhost02: Starting Cluster...
```

By enabling the `pacemaker` service, the server automatically joins the cluster after a reboot. This ensures that your system is protected. Alternatively, you can start the `pacemaker` service manually on boot. You can then investigate the cause of failure. However, this is generally not required for SAP NetWeaver ASCS cluster.

**Increase corosync totem timeout**

**RHEL 7.x**

1. Edit the `/etc/corosync/corosync.conf` file in all cluster nodes to increase the token value or to add a value if it is not present.

```
totem {
    version: 2
    secauth: off
    cluster_name: my-rhel-sap-cluster
    transport: udpu
    rrp_mode: passive
    token: 29000   <------ Value to be set
}
```

2. Reload the corosync with the following command, on any one of the cluster nodes. This does not cause any downtime.

```
pcs cluster reload corosync
```

3. Use the following command to confirm if the changes are active.

```
corosync-cmapctl | grep totem.token
Runtime.config.totem.token (u32) = 29000
```

## RHEL 8.x

Use the following command to increase the token value or to add a value if it is not present.

```
pcs cluster config update totem token=29000
```

## Check cluster status

Once the cluster service `pacemaker` is started, check the cluster status with `pcs status` command, as shown in the following example. Both the primary (`rhxdbhost01`) and secondary (`rhxdbhost02`) servers should be seen as online.

```
pcs status
Cluster name: rhelha

WARNINGS:
No stonith devices and stonith-enabled is not false

Cluster Summary:
  * Stack: corosync
  * Current DC: rhxdbhost01 (version 2.0.3-5.el8_2.5-4b1f869f0f) - partition with
 quorum
  * Last updated: Tue Jan 10 21:32:15 2023
  * Last change:  Tue Jan 10 19:46:50 2023 by hacluster via crmd on rhxdbhost01
  * 2 nodes configured
  * 0 resource instances configured

Node List:
  * Online: [ rhxdbhost01 rhxdbhost02 ]

Full List of Resources:
  * No resources

Daemon Status:
  corosync: active/enabled
  pacemaker: active/enabled
  pcsd: active/enabled
```

## Prepare for resource creation

To ensure that the cluster does not perform any unexpected actions during setup of resources and configuration, set the maintenance mode to true.

Run the following command to put the cluster in maintenance mode.

```
pcs property set maintenance-mode=true
```

## Cluster bootstrap

Configure the cluster bootstrap parameters by running the following commands.

```
pcs resource defaults update resource-stickiness=1
pcs resource defaults update migration-threshold=3
```

## Create `fence_aws` STONITH resource

Modify the commands in the following table to match your configuration values.

```
pcs stonith create rsc_aws_stonith_<DBSID> fence_aws region=us-east-1
  pcmk_host_map="rhxdbhost01:i-xxxxinstidforhost1;rhxdbhost02:i-xxxxinstidforhost2"
  power_timeout=240 pcmk_reboot_timeout=300 pcmk_reboot_retries=2 pcmk_delay_max=30
  pcmk_reboot_action=reboot op start timeout=180 op stop timeout=180 op monitor
  interval=180 timeout=60
```

> ⓘ **Note**
>
> The default `pcmk` action is reboot. If you want to have the instance remain in a
> stopped state until it has been investigated, and then manually started again,
> add `pcmk_reboot_action=off`. Any High Memory (u-**tb1.** ) instance or metal
> instance running on a dedicated host won't support reboot, and will require
> `pcmk_reboot_action=off`.

## Create file system resources

Mounting and unmounting file system resources to align with the location of SAP ASE database is
done using cluster resources.

Modify and run the following commands to create these file system resources.

### /sybase

```
pcs resource create rsc_fs_<DBSID>_sybase ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/sybase" directory="/sybase" fstype="nfs4" options="
 rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,timeo=600,r
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/sapdata_1

```
pcs resource create rsc_fs_<DBSID>_data ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/asedata" directory="/sybase/<DBSID>/sapdata_1" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=8,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/saplog_1

```
pcs resource create rsc_fs_<DBSID>_log ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/aselog" directory="/sybase/<DBSID>/saplog_1" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/sapdiag

```
pcs resource create rsc_fs_<DBSID>_diag ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/sapdiag" directory="/sybase/<DBSID>/sapdiag" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybase/<DBSID>/saptmp

```
pcs resource create rsc_fs_<DBSID>_tmp ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/saptmp" directory="/sybase/<DBSID>/saptmp" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

## /sybasebackup

```
pcs resource create rsc_fs_<DBSID>_bkp ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/sybasebackup" directory="/backup" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

### /usr/sap

```
pcs resource create rsc_fs_<DBSID>_sap ocf:heartbeat:Filesystem params
 device="<nfs.fqdn>:/usrsap" directory="/usr/sap" fstype="nfs4"
options="rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
 op start timeout=60s interval=0 op stop timeout=60s interval=0 op monitor interval=20s
 timeout=40s
```

### Notes

- Review the mount options to ensure that they match with your operating system, NFS file system type, and the latest recommendations from SAP and Amazon.
- <nfs.fqdn> must be the alias of the FSx for ONTAP resource. For example, `fs-xxxxxx.efs.xxxxxx.amazonaws.com`.
- It is important to create the resources in the proper mount order/sequence.

### Create overlay IP resources

The IP resource provides the details necessary to update the route table entry for overlay IP.

Use the following command to create an SAP ASE database IP resource.

```
pcs resource create rsc_ip_<DBSID>_ASEDB ocf:heartbeat:aws-vpc-move-ip ip=172.16.0.23
 interface=eth0 routing_table=rtb-xxxxxroutetable1 op monitor interval=20s timeout=40s
 --group rsc_asedb_group
```

### Notes

- If more than one route table is required for connectivity or because of subnet associations, the `routing_table` parameter can have multiple values separated by a comma. For example, `routing_table=rtb-xxxxxroutetable1, rtb-xxxxxroutetable2`.
- Additional parameters – `lookup_type` and `routing_table_role` are required for shared VPC. For more information, see Shared VPC – optional.

**Create SAP ASE database resource**

SAP ASE database is started and stopped using cluster resources.

Modify and run the following command to create the SAPDatabase resource.

```
pcs resource create rsc_ase_<DBSID>_ASEDB ocf:heartbeat:SAPDatabase SID=<DBSID>
 DBTYPE=SYB STRICT_MONITORING=TRUE op start timeout=300 op stop timeout=300 --group
 grp_<DBSID>_ASEDB
```

Use the following command for more details on the resource.

```
pcs resource describe ocf:heartbeat:SAPDatabase
```

**Activate cluster**

Use `pcs config show` to review that all the values have been entered correctly.

On confirmation of correct values, set the maintenance mode to false using the following
command. This enables the cluster to take control of the resources.

```
pcs property set maintenance-mode=false
```

See the Sample configuration.

# Sample configuration (pcs config show)

The following sample configuration is based on ENSA2.

```
pcs config show
Cluster Name: rhelha
Corosync Nodes:
 rhxdbhost01 rhxdbhost02
Pacemaker Nodes:
 rhxdbhost01 rhxdbhost02

Resources:
 Group: rsc_asedb_group
  Meta Attrs: resource-stickiness=5000
  Resource: rsc_vip_asedb (class=ocf provider=heartbeat type=aws-vpc-move-ip)
   Attributes: interface=eth0 ip=172.16.0.23 routing_table=rtb-0b3f1d6196f45300d
   Operations: monitor interval=60s timeout=30s (rsc_vip_asedb-monitor-interval-60s)
               start interval=0s timeout=180s (rsc_vip_asedb-start-interval-0s)
```

```
                 stop interval=0s timeout=180s (rsc_vip_asedb-stop-interval-0s)
  Resource: rsc_fs_sybase (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-09794aeece44cc025.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/sybase directory=/sybase force_unmount=safe fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
   Operations: monitor interval=20s timeout=40s (rsc_fs_sybase-monitor-interval-20s)
              start interval=0s timeout=60s (rsc_fs_sybase-start-interval-0s)
              stop interval=0s timeout=60s (rsc_fs_sybase-stop-interval-0s)
  Resource: rsc_fs_data (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-01c02d046ae5a24a2.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/asedata directory=/sybase/ARD/sapdata_1 force_unmount=safe
 fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=8,tim
   Operations: monitor interval=20s timeout=40s (rsc_fs_data-monitor-interval-20s)
              start interval=0s timeout=60s (rsc_fs_data-start-interval-0s)
              stop interval=0s timeout=60s (rsc_fs_data-stop-interval-0s)
  Resource: rsc_fs_log (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-04cd525dbd0b354d2.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/aselog directory=/sybase/ARD/saplog_1 force_unmount=safe
 fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
   Operations: monitor interval=20s timeout=40s (rsc_fs_log-monitor-interval-20s)
              start interval=0s timeout=60s (rsc_fs_log-start-interval-0s)
              stop interval=0s timeout=60s (rsc_fs_log-stop-interval-0s)
  Resource: rsc_fs_sapdiag (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-09794aeece44cc025.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/sapdiag directory=/sybase/ARD/sapdiag force_unmount=safe
 fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
   Operations: monitor interval=20s timeout=40s (rsc_fs_sapdiag-monitor-interval-20s)
              start interval=0s timeout=60s (rsc_fs_sapdiag-start-interval-0s)
              stop interval=0s timeout=60s (rsc_fs_sapdiag-stop-interval-0s)
  Resource: rsc_fs_saptmp (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-09794aeece44cc025.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/saptmp directory=/sybase/ARD/saptmp force_unmount=safe
 fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
   Operations: monitor interval=20s timeout=40s (rsc_fs_saptmp-monitor-interval-20s)
              start interval=0s timeout=60s (rsc_fs_saptmp-start-interval-0s)
              stop interval=0s timeout=60s (rsc_fs_saptmp-stop-interval-0s)
  Resource: rsc_fs_backup (class=ocf provider=heartbeat type=Filesystem)
   Attributes: device=svm-09794aeece44cc025.fs-04af26e8311974f41.fsx.us-
east-1.amazonaws.com:/backup directory=/sybasebackup force_unmount=safe fstype=nfs4
 options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
```

```
     Operations: monitor interval=20s timeout=40s (rsc_fs_backup-monitor-interval-20s)
                 start interval=0s timeout=60s (rsc_fs_backup-start-interval-0s)
                 stop interval=0s timeout=60s (rsc_fs_backup-stop-interval-0s)
   Resource: rsc_fs_usrsap (class=ocf provider=heartbeat type=Filesystem)
     Attributes: device=svm-09794aeece44cc025.fs-04af26e8311974f41.fsx.us-
 east-1.amazonaws.com:/usrsap directory=/usr/sap force_unmount=safe fstype=nfs4
  options=rw,noatime,vers=4.1,rsize=262144,wsize=262144,namlen=255,hard,proto=tcp,nconnect=2,tim
     Operations: monitor interval=20s timeout=40s (rsc_fs_usrsap-monitor-interval-20s)
                 start interval=0s timeout=60s (rsc_fs_usrsap-start-interval-0s)
                 stop interval=0s timeout=60s (rsc_fs_usrsap-stop-interval-0s)
   Resource: sybaseARD (class=ocf provider=heartbeat type=SAPDatabase)
     Attributes: DBTYPE=SYB SID=ARD STRICT_MONITORING=TRUE
     Operations: methods interval=0s timeout=5s (sybaseARD-methods-interval-0s)
                 monitor interval=120s timeout=60s (sybaseARD-monitor-interval-120s)
                 start interval=0s timeout=300 (sybaseARD-start-interval-0s)
                 stop interval=0s timeout=300 (sybaseARD-stop-interval-0s)
Stonith Devices:
 Resource: clusterfence (class=stonith type=fence_aws)
  Attributes: pcmk_delay_max=45
 pcmk_host_map=rhxdbhost01:i-03939ad3f07e14e3f;rhxdbhost02:i-09f138e3a1290bfde
 pcmk_reboot_action=off pcmk_reboot_retries=4 pcmk_reboot_timeout=600 power_timeout=240
 region=us-east-1
  Operations: monitor interval=300 timeout=60 (clusterfence-monitor-interval-300)
              start interval=0s timeout=600 (clusterfence-start-interval-0s)
Fencing Levels:
Location Constraints:
Ordering Constraints:
Colocation Constraints:
Ticket Constraints:
Alerts:
 No alerts defined
Resources Defaults:
  Meta Attrs: rsc_defaults-meta_attributes
    migration-threshold=1
Operations Defaults:
  No defaults set
Cluster Properties:
 cluster-infrastructure: corosync
 cluster-name: rhelha
 dc-version: 2.1.2-4.el8_6.7-ada5c3b36e2
 have-watchdog: false
 last-lrm-refresh: 1693394303
 maintenance-mode: false
Tags:
```

```
  No tags defined
Quorum:
  Options:
```

# Operations

This section covers the following topics.

**Topics**

- [Analysis and maintenance](#)

- [Testing](#)

# Analysis and maintenance

This section covers the following topics.

**Topics**

- [Viewing the cluster state](#)

- [Performing planned maintenance](#)

- [Post-failure analysis and reset](#)

- [Alerting and monitoring](#)

## Viewing the cluster state

You can view the state of the cluster based on your operating system.

**Operating system based**

There are multiple operating system commands that can be run as root or as a user with appropriate permissions. The commands enable you to get an overview of the status of the cluster and its services. See the following commands for more details.

```
pcs status
```

Sample output:

```
pcs status
Cluster name: rhelha
Cluster Summary:
* Stack: corosync
* Current DC: <rhxdbhost01> (version 2.1.2-4.el8_6.5-ada5c3b36e2) - partition with
  quorum
* Last updated: Wed Apr 12 19:38:46 2023
* Last change: Mon Apr 10 14:55:08 2023 by root via crm_resource on <rhxdbhost01>
* 2 nodes configured
* 10 resource instances configured
Node List:
* Online: [ awnulaeddb awnulaeddbha ]
Full List of Resources:
* rsc_aws_stonith_ARD (stonith:fence_aws): Started awnulaeddb
* Resource Group: grp_ARD_ASEDB:
* rsc_ip_ARD_ASEDB (ocf::heartbeat:aws-vpc-move-ip): Started <rhxdbhost01>
* rsc_fs_ARD_sybase (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_fs_ARD_data (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_fs_ARD_log (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_fs_ARD_sapdiag (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
SAP NetWeaver on {aws} SAP NetWeaver Guides
Analysis and maintenance
140
* rsc_fs_ARD_saptmp (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_fs_ARD_backup (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_fs_ARD_usrsap (ocf::heartbeat:Filesystem): Started <rhxdbhost01>
* rsc_ase_ARD_ASEDB (ocf::heartbeat:SAPDatabase): Started <rhxdbhost01>
Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

The following table provides a list of useful commands.

| Command | Description |
| --- | --- |
| crm_mon | Display cluster status on the console with updates as they occur |
| crm_mon -1 | Display cluster status on the console just once, and exit |

| Command | Description |
|---|---|
| crm_mon -Arnf | -A Display node attributes |
| | -n Group resources by node |
| | -r Display inactive resources |
| | -f Display resource fail counts |
| pcs help | View more options |
| crm_mon --help-all | View more options |

## Performing planned maintenance

The cluster connector is designed to integrate the cluster with SAP start framework (`sapstartsrv`), including the rolling kernel switch (RKS) awareness. Stopping and starting the SAP system using `sapcontrol` should not result in any cluster remediation activities as these actions are not interpreted as failures. Validate this scenario when testing your cluster.

There are different options to perform planned maintenance on nodes, resources, and the cluster.

**Topics**

- [Maintenance mode](#)
- [Placing a node in standby mode](#)
- [Moving a resource (not recommended)](#)

### Maintenance mode

Use maintenance mode if you want to make any changes to the configuration or take control of the resources and nodes in the cluster. In most cases, this is the safest option for administrative tasks.

**Example**

On

Use the following command to turn on maintenance mode.

```
pcs property set maintenance-mode="true"
```

Off

- Use the following command to turn off maintenance mode.

```
pcs property set maintenance-mode="false"
```

## Placing a node in standby mode

To perform maintenance on the cluster without system outage, the recommended method for moving active resources is to place the node you want to remove from the cluster in standby mode.

```
pcs node standby <rhxdbhost01>
```

The cluster will cleanly relocate resources, and you can perform activities, including reboots on the node in standby mode. When maintenance activities are complete, you can re-introduce the node with the following command.

```
pcs node unstandby <rhxdbhost01>
```

## Moving a resource (not recommended)

Moving individual resources is not recommended because of the migration or move constraints that are created to lock the resource in its new location. These can be cleared as described in the info messages, but this introduces an additional setup.

```
<rhxdbhost01>:~ # pcs resource move <grp_ARD_ASEDB>

Note: Move constraint created for <grp_ARD_ASEDB> to <rhxdbhost02>
Note: Use "pcs constraint location remove cli-prefer-grp_ARD_ASEDB" to remove this
  constraint.
```

## Post-failure analysis and reset

A review must be conducted after each failure to understand the source of failure as well the reaction of the cluster. In most scenarios, the cluster prevents an application outage. However, a manual action is often required to reset the cluster to a protective state for any subsequent failures.

**Topics**

- [Checking the logs](#)
- [Cleanup pcs status](#)
- [Restart failed nodes or pacemaker](#)
- [Further analysis](#)

## Checking the logs

Start your troubleshooting by checking the operating system log `/var/log/messages`. You can find additional information in the cluster and pacemaker logs.

- **Cluster logs** – updated in the `corosync.conf` file located at `/etc/corosync/corosync.conf`.
- **Pacemaker logs** – updated in the `pacemaker.log` file located at `/var/log/pacemaker`.
- **Resource agents** – `/var/log/messages`

Application based failures can be investigated in the SAP work directory.

## Cleanup `pcs status`

If failed actions are reported using the `crm status` command, and if they have already been investigated, then you can clear the reports with the following command.

```
pcs resource cleanup <resource> <hostname>
```

```
pcs stonith cleanup
```

> **ⓘ Note**
>
> Use the help command to understand the impact of these commands.

## Restart failed nodes or `pacemaker`

It is recommended that failed (or fenced) nodes are not automatically restarted. It gives operators a chance to investigate the failure, and ensure that the cluster doesn't make assumptions about the state of resources.

You need to restart the instance or the pacemaker service based on your approach.

**Further analysis**

If further analysis from Red Hat is required, they may request an sos report, or logs of the cluster from `crm_report` or `pcs cluster report`.

**sos report** – The sos report command is a tool that collects configuration details, system information, and diagnostic information from a Red Hat Enterprise Linux system. For instance, the running kernel version, loaded modules, and system and service configuration files. The command also runs external programs to collect further information, and stores this output in the resulting archive. For more information, see Red Hat documentation [What is an sos report and is it different from an sosreport?](#)

**crm report** – collects the cluster logs/information from the node where the command is being run. For more information, see Red Hat documentation [How do I generate a crm_report from a RHEL 6 or 7 High Availability cluster node using pacemaker?](#)

```
crm_report
```

**pcs cluster report** – command collects the cluster logs/information from all the nodes involved in the cluster.

```
pcs cluster report <destination_path>
```

> ⓘ **Note**
>
> The `pcs cluster report` command relies on passwordless ssh being set up between the nodes.

## Alerting and monitoring

**Using the cluster alert agents**

Within the cluster configuration, you can call an external program (an alert agent) to handle alerts. This is a *push* notification. It passes information about the event via environment variables.

The agents can then be configured to send emails, log to a file, update a monitoring system, etc. For example, the following script can be used to access Amazon SNS.

```
#!/bin/sh

#alert_sns.sh
#modified from /usr/share/pacemaker/alerts/alert_smtp.sh.sample

##############################################################################
#SETUP
* Create an SNS Topic and subscribe email or chatbot
* Note down the ARN for the SNS topic
* Give the IAM Role attached to both Instances permission to publish to the SNS Topic
* Ensure the aws cli is installed
* Copy this file to /usr/share/pacemaker/alerts/alert_sns.sh or other location on BOTH
 nodes
* Ensure the permissions allow for hacluster and root to execute the script
* Run the following as root (modify file location if necessary and replace SNS ARN):

#SLES:
crm configure alert aws_sns_alert /usr/share/pacemaker/alerts/alert_sns.sh meta
 timeout=30s timestamp-format="%Y-%m-%d_%H:%M:%S" to <{ arn:aws:sns:region:account-
id:myPacemakerAlerts  }>

#RHEL:
pcs alert create id=aws_sns_alert path=/usr/share/pacemaker/alerts/alert_sns.sh meta
 timeout=30s timestamp-format="%Y-%m-%d_%H:%M:%S"
pcs alert recipient add aws_sns_alert value=<arn:aws:sns:region:account-
id:myPacemakerAlerts>

#Additional information to send with the alerts.
node_name=`uname -n`
sns_body=`env | grep CRM_alert_`

#Required for SNS
TOKEN=$(/usr/bin/curl --noproxy '*' -s -X PUT "http://169.254.169.254/latest/api/token"
 -H "X-aws-ec2-metadata-token-ttl-seconds: 21600")

#Get metadata
REGION=$(/usr/bin/curl --noproxy '*' -w "\n" -s -H "X-aws-ec2-metadata-token: $TOKEN"
 http://169.254.169.254/latest/dynamic/instance-identity/document | grep region | awk -
F\" '{print $4}')

sns_subscription_arn=${CRM_alert_recipient}

#Format depending on alert type
```

```
case ${CRM_alert_kind} in
   node)
      sns_subject="${CRM_alert_timestamp} ${cluster_name}: Node '${CRM_alert_node}' is
 now '${CRM_alert_desc}'"
   ;;
   fencing)
      sns_subject="${CRM_alert_timestamp} ${cluster_name}: Fencing ${CRM_alert_desc}"
   ;;
   resource)
      if [ ${CRM_alert_interval} = "0" ]; then
         CRM_alert_interval=""
      else
         CRM_alert_interval=" (${CRM_alert_interval})"
      fi
      if [ ${CRM_alert_target_rc} = "0" ]; then
         CRM_alert_target_rc=""
      else
         CRM_alert_target_rc=" (target: ${CRM_alert_target_rc})"
      fi
      case ${CRM_alert_desc} in
         Cancelled)
            ;;
         *)
            sns_subject="${CRM_alert_timestamp}: Resource operation
 '${CRM_alert_task}${CRM_alert_interval}' for '${CRM_alert_rsc}' on
 '${CRM_alert_node}': ${CRM_alert_desc}${CRM_alert_target_rc}"
            ;;
      esac
   ;;
   attribute)
      sns_subject="${CRM_alert_timestamp}: The '${CRM_alert_attribute_name}' attribute
 of the '${CRM_alert_node}' node was updated in '${CRM_alert_attribute_value}'"
   ;;
   *)
      sns_subject="${CRM_alert_timestamp}: Unhandled $CRM_alert_kind alert"
   ;;
esac

#Use this information to send the email.
aws sns publish --topic-arn "${sns_subscription_arn}" --subject "${sns_subject}" --
message "${sns_body}" --region ${REGION}
```

# Testing

We recommend scheduling regular fault scenario recovery testing at least annually, and as part of the operating system or SAP kernel updates that may impact operations. For more details on best practices for regular testing, see SAP Lens – Best Practice 4.3 – Regularly test business continuity plans and fault recovery.

The tests described here simulate failures. These can help you understand the behavior and operational requirements of your cluster.

In addition to checking the state of cluster resources, ensure that the service you are trying to protect is in the required state. Can you still connect to SAP? Are locks still available in SM12?

Define the recovery time to ensure that it aligns with your business objectives. Record recovery actions in runbooks.

**Topics**

- Test 1: Stop SAP ASE database using sapcontrol
- Test 2: Unmount FSx for ONTAP file system on primary host
- Test 3: Kill the database processes on the primary host
- Test 4: Simulate hardware failure of an individual node
- Test 5: Simulate a network failure
- Test 6: Simulate an NFS failure
- Test 7: Accidental shutdown

## Test 1: Stop SAP ASE database using `sapcontrol`

**Simulate failure** – On rhxdbhost01 as root:

```
/usr/sap/hostctrl/exe/saphostctrl -function StopDatabase -dbname ARD -dbtybe syb -force
```

**Expected behavior** – SAP ASE database is stopped, and the `SAPDatabase` resource agent enters a failed state. The cluster will failover the database to the secondary instance.

**Recovery action** – No action required.

## Test 2: Unmount FSx for ONTAP file system on primary host

**Simulate failure** – On rhxdbhost01 as root:

```
umount -l /sybase/ARD/sapdata_1
```

**Expected behavior** – The `rsc_fs` resource enters a failed state. The cluster stops the SAP ASE database, and will failover to the secondary instance.

**Recovery action** – No action required.

## Test 3: Kill the database processes on the primary host

**Simulate failure** – On rhxdbhost01 as root:

```
ps -ef |grep -i sybaard
kill -9 <PID>
```

**Expected behavior** – SAP ASE database fails, and the `SAPDatabase` resource enters a failed state. The cluster will failover the database to the secondary instance.

**Recovery action** – No action required.

## Test 4: Simulate hardware failure of an individual node

**Notes** – To simulate a system crash, you must first ensure that `/proc/sys/kernel/sysrq` is set to 1.

**Simulate failure** – On the primary host as root:

```
echo 'c' > /proc/sysrq-trigger
```

**Expected behavior** – The node which has been killed fails. The cluster moves the resource (SAP ASE database) that was running on the failed node to the surviving node.

**Recovery action** – Start the EC2 node.

## Test 5: Simulate a network failure

**Notes** – See the following list.

- Iptables must be installed.

- Use a subnet in this command because of the secondary ring.

- Check for any existing iptables rules as iptables -F will flush all rules.

- Review pcmk_delay and priority parameters if you see neither node survives the fence race.

**Simulate failure** – On either node as root:

```
iptables -A INPUT -s <CIDR_of_other_subnet> -j DROP; iptables -A OUTPUT -d
  <CIDR_of_other_subnet> -j DROP
```

**Expected behavior** – The cluster detects the network failure, and fences one of the nodes to avoid a split-brain situation.

**Recovery action** – If the node where the command was run survives, execute iptables -F to clear the network failure. Start the EC2 node.

## Test 6: Simulate an NFS failure

**Notes** – See the following list.

- Iptables must be installed.

- Check for any existing iptables rules as iptables -F will flush all rules.

- Although rare, this is an important scenario to test. Depending on the activity it may take some time (10 min +) to notice that I/O to EFS is not occurring and fail either the Filesystem or SAP resources.

**Simulate failure** – On the primary host as root:

```
iptables -A OUTPUT -p tcp --dport 2049 -m state --state NEW,ESTABLISHED,RELATED -j
  DROP; iptables -A INPUT -p tcp --sport 2049 -m state --state ESTABLISHED -j DROP
```

**Expected behavior** – The cluster detects that NFS is not available, and the `SAPDatabase` resource agent fails, and moves to the FAILED state.

**Recovery action** – If the node where the command was run survives, execute iptables -F to clear the network failure. Start the EC2 node.

## Test 7: Accidental shutdown

**Notes** – See the following list.

- Avoid shutdowns without cluster awareness.

- We recommend the use of systemd to ensure predictable behaviour.

- Ensure the resource dependencies are in place.

**Simulate failure** – Login to Amazon Management Console, and stop the instance or issue a shutdown command.

**Expected behavior** – The node which has been shut down fails. The cluster moves the resource (SAP ASE database) that was running on the failed node to the surviving node. If systemd and resource dependencies are not configured, you may notice that while the EC2 instance is shutting down gracefully, the cluster will detect an unclean stop of cluster services on the node and will fence the EC2 instance being shut down.

**Recovery action** – Start the EC2 node and pacemaker service.

# Amazon Backint Agent for SAP ASE

Amazon Backint Agent for SAP ASE is a backup and restore application for SAP ASE relational database management systems (RDBMS) running on Amazon EC2 instances. Amazon Backint Agent runs as a standalone application that integrates with your existing workflows to back up your SAP ASE database to Amazon S3.

Amazon Backint Agent runs on the SAP ASE database instance, where backups and transaction logs are transferred from SAP ASE to Amazon Backint agent. Based on the configurations in your agent file, Amazon Backint agent stores your files in Amazon S3. To restore your SAP ASE database, specify the backup location Amazon S3.

**Topics**

- [How Amazon Backint Agent for SAP ASE works](#)
- [Architecture Components](#)
- [Workflow](#)
- [Pricing](#)
- [Supported Database Versions](#)
- [Supported Operating Systems](#)
- [Supported Amazon Regions](#)
- [Prerequisites](#)
- [Installation](#)
- [Backup and Restore](#)
- [Troubleshooting](#)
- [Version history](#)

# How Amazon Backint Agent for SAP ASE works

Amazon Backint Agent for SAP ASE is built on the same foundation as the Amazon Backint Agent for SAP HANA and provides a similar experience and capabilities.

Amazon Backint Agent implements the SAP ASE Backup Server Archive API. DUMP and LOAD commands can be initiated via, for example, the SAP ASE command line interface.

# Architecture Components

Architecture components include the SAP ASE database, SAP ASE backup server, Amazon Backint Agent (for ASE), and an implementation of the ASE backup server archive API for Amazon Backint Agent.

## SAP ASE Backup Server

SAP ASE backup server is a component of the SAP ASE database system that manages database backup and restore operations. It's a separate process that runs alongside the main database server and is specifically designed to handle the movement of data during backup and recovery operations.

## Backup Server Archive API

The SAP ASE Backup Server API provides an interface for third-party applications to integrate with ASE's backup and recovery operations. This API enables communication between the ASE database engine and external storage systems, allowing for efficient backup and restore operations without requiring intermediate storage.
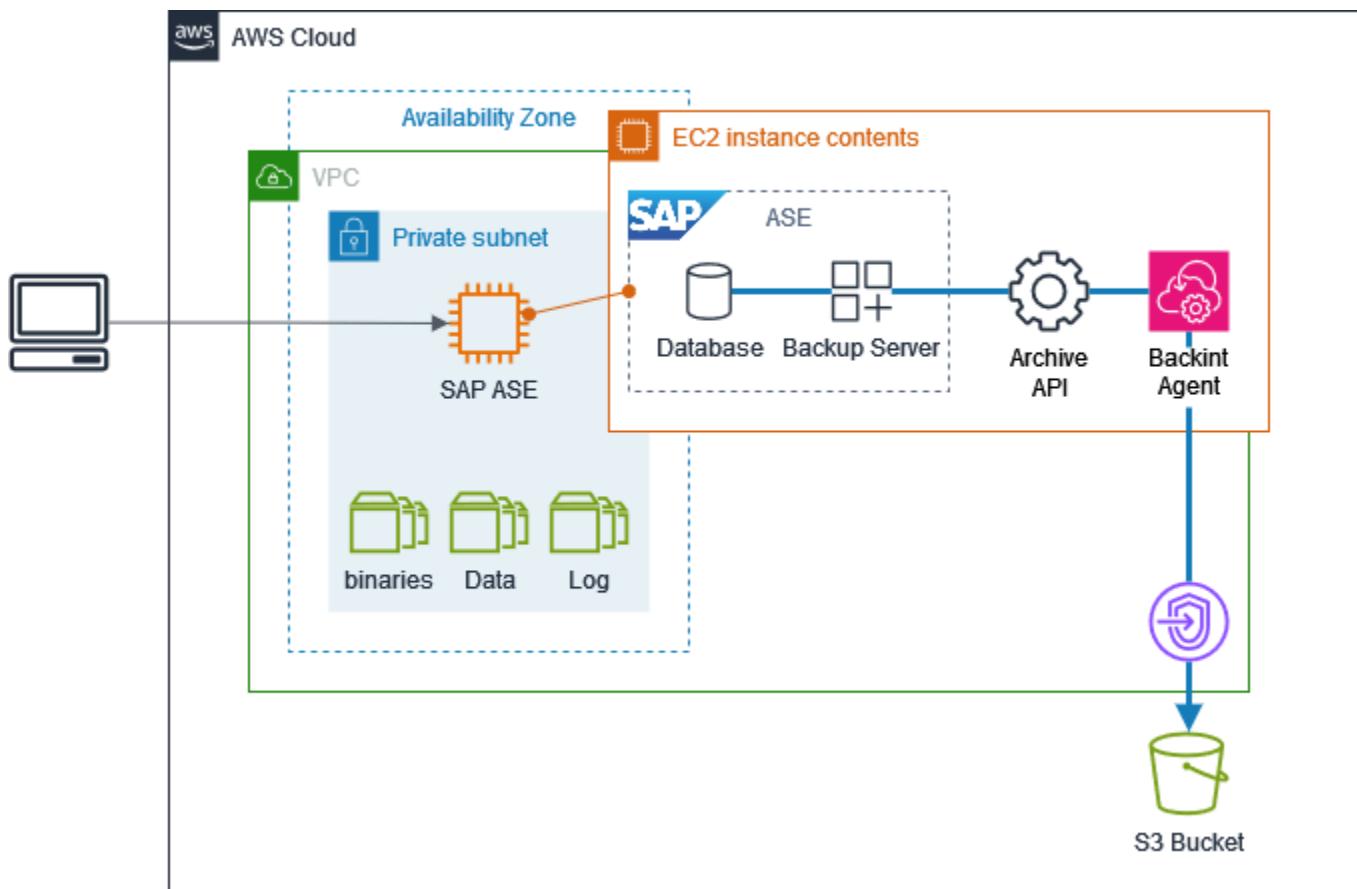
## Amazon Backint Agent

Backint Agent is responsible for streaming backup-related data to and from the destination in Amazon S3. The agent Backint Agent is largely independent of the database management system.

# Workflow

SAP ASE database, SAP ASE backup server, Amazon Backint Agent for ASE and Amazon Backint Agent API are installed on a single EC2 instance.

1. Administrator initiates backup or restore via methods, which are native to SAP ASE. For example, a database statement via isql.
2. The **SAP ASE database** reads data pages from database devices and transfers these pages to the Backup Server through the IPC channel.
3. **Backup Server** opens byte-streams (one per stripe) and invokes Backint-specific implementation of the Archive API (Backint Archive API).
4. **Archive API** handles communication between Backup Server and Amazon Backint Agent.
5. **Backint Agent** streams data to (backup) and from (restore) **Amazon S3**.

# Pricing

Amazon Backint agent is free of charge. You pay only for the underlying Amazon services that you use, for example Amazon S3. See Amazon S3 pricing for more information.

# Supported Database Versions

Amazon Backint agent for SAP ASE supports the following SAP ASE database versions.

- SAP ASE 16.0
- SAP ASE 16.1

Please note that only SAP ASE databases on Amazon EC2 are supported.

# Supported Operating Systems

Amazon Backint agent is supported on the following operating systems.

- SUSE Linux Enterprise Server

- SUSE Linux Enterprise Server for SAP

- Red Hat Enterprise Linux for SAP

All major and minor versions that are currently supported for SAP ASE are supported. Please refer to the SAP Product Availability Matrix (PAM) at support.sap.com/pam for more information.

# Supported Amazon Regions

Amazon Backint agent is available in all commercial Regions. China (Beijing), China (Ningxia), and GovCloud are not supported.

# Prerequisites

This section provides information on mandatory prerequisites for Amazon Backint Agent for SAP ASE.

**Topics**

- Amazon Identity and Access Management (IAM)
- Amazon S3 Bucket

## Amazon Identity and Access Management (IAM)

To enable S3 bucket access for your Amazon EC2 instance, create or update an inline IAM policy with the following permissions and attach it to your EC2 service role. Replace the resource names, such as the S3 bucket name, to match your setup. You must provide the Amazon Region and Amazon S3 bucket owner account ID along with the Amazon S3 bucket name and KMS Key for Encryption.

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketPolicyStatus",
```

```
                "s3:GetBucketLocation",
                "s3:ListBucket",
                "s3:GetBucketAcl",
                "s3:GetBucketPolicy"
            ],
            "Resource": [
                "arn:aws:s3:::your-bucket-name/*",
                "arn:aws:s3:::your-bucket-name"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:GenerateDataKey"
            ],
            "Resource": "arn:aws:kms:us-east-1:123456789012:key/your-kms-key-id"
        },
          {
            "Effect": "Allow",
            "Action": [
                "s3:PutObjectTagging",
                "s3:PutObject",
                "s3:GetObject",
                "s3:DeleteObject"
            ],
            "Resource": "arn:aws:s3:::your-bucket-name/your-folder-name/*"
        }
    ]
}
```

> **ⓘ Note**
>
> If you want to allow cross-account backup and restore, you must add your account details under a principal element in your policy. For more information about principal policies, see [Amazon JSON Policy Elements: Principal](#) in the *Amazon Identity and Access Management User Guide*. In addition, you must ensure that the S3 bucket policies allow your account to perform the actions specified in the IAM policy example above. For more information, see the example for [Bucket owner granting cross-account bucket permissions](#) in the *Amazon S3 Developer Guide*.

For more information about managed and inline policies, see the [IAM User Guide](#).

# Amazon S3 Bucket

In preparation for installation of the Amazon Backint agent, identify or create an S3 bucket in the target Region where SAP ASE backups will be stored. The bucket must have been created after May 2019 in order to be compatible with Amazon Backint agent and must have public access blocked, as backups will fail if public access is enabled.

Amazon Backint agent supports backing up to Amazon S3 with VPC endpoints. Using an Amazon S3 gateway endpoint can improve performance, help prevent timeouts, enhance security, and reduce costs. For more information, see [VPC Endpoints](#).

## S3 storage classes

Amazon Backint agent supports storing SAP ASE database backups in S3 Standard, S3 Standard-IA, S3 One Zone-IA, and S3 Intelligent-Tiering storage classes. By default, backups use the S3 Standard storage class, but this can be changed through the Amazon Backint agent configuration file, [S3 Lifecycle rules](#), or directly via APIs.

Amazon Backint Agent for SAP ASE does not support the following storage classes: . Reduced Redundancy . Deep Archive . Glacier

While S3 Intelligent-Tiering can automatically move objects to archival tiers, Amazon Backint agent requires objects to be in standard access tiers for recovery or deletion operations. For more information, see [Amazon S3 Storage Classes](#) in the *Amazon S3 Developer Guide*.

## Encryption

Amazon Backint agent supports encrypting your SAP ASE backup files while storing them in Amazon S3, using server-side encryption with Amazon Key Management Service (KMS). You can encrypt your backups with an Amazon managed key (aws/s3) or you can use your own customer managed key stored in Amazon KMS. To encrypt your backup files with Amazon KMS keys (Amazon managed or customer managed), you must provide the KMS key ARN during installation or update the Amazon Backint agent configuration file later. To learn more about encrypting your S3 objects using Amazon KMS, see [How Amazon S3 uses Amazon KMS](#) in the  *Amazon Key Management Service Developer Guide*. Alternatively, you can enable default encryption for your Amazon S3 bucket using either Amazon KMS keys or keys managed by Amazon S3 (SSE-S3). To learn more about enabling default encryption for your bucket, see [How do I enable default encryption for an Amazon S3 bucket?](#) in the *Amazon S3 Console User Guide*.

## Object locking

You can store objects using a *write-once-read-many* (WORM) model with S3 Object Lock. Use S3 Object Lock if you want to prevent your SAP ASE backup files from being accidentally deleted or overwritten for a specific time period or indefinitely. If S3 Object Lock is enabled, you can't delete your SAP ASE backups stored in Amazon S3 using SAP ASE Cockpit, SAP ASE Studio, or SQL commands until the retention period expires. To learn about S3 Object Lock, see Locking objects using S3 Object Lock in the *Amazon S3 Developer Guide*.

## Object tagging

By default, Amazon Backint agent adds a tag called `AmazonBackintAgentVersion` when it stores your SAP ASE backup files in your S3 bucket. This tag helps to identify the Amazon Backint version and the SAP ASE version used when backing up your SAP ASE database. You can list the value of the tags from S3 console or using APIs. To disable default tagging, modify the Amazon Backint agent configuration file. See Additional Parameters for more information.

## Data perimeter

Amazon Backint Agent must be installed on your Amazon EC2 instance. To download the installation binary, your EC2 instance needs access to the Amazon-managed S3 buckets where the installer packages are hosted. If your organization uses data perimeter policies to control access to Amazon S3 in your environment, you might need to explicitly allow these service-owned buckets so that the EC2 instance can retrieve the required installer. The following policy shows an example Service control policy to permit access to service-owned resources through the resource perimeter, relevant service owned buckets are listed in the NotResource element of the policy.

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "EnforceResourcePerimeterAWSResources",
      "Effect": "Deny",
      "Action": "*",
      "NotResource": [
        "arn:aws:s3:::awssap-backint-agent-ase",
        "arn:aws:s3:::awssap-backint-agent-ase/*"
      ],
      "Condition": {
        "StringNotEqualsIfExists": {
```

```
            "aws:ResourceOrgID": "<organization id>",
            "aws:PrincipalTag/dp:exclude:resource": "true"
        }
    }
  }
 ]
}
```

The following policy shows an example VPC endpoint policy allowing access to specific service-owned resources through a VPC endpoint. Relevant service owned buckets are listed in the Resource element of the statement.

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Sid": "AllowRequestsToAWSOwnedResources",
            "Effect": "Allow",
            "Principal": "*",
            "Action": [
                "s3:GetObject",
                "s3:ListBucket"
            ],
            "Resource": [
                "arn:aws:s3:::awssap-backint-agent-ase",
                "arn:aws:s3:::awssap-backint-agent-ase/*"
            ]
        }
    ]
}
```

### Amazon Command Line Interface (CLI)

Amazon Backint agent installation leverages the Amazon CLI to validate S3 bucket properties. To install or update to the Amazon CLI, see Install or update to the latest version of the Amazon CLI.

## Installation

This section provides information to help you install the Amazon Backint agent using the Amazon Backint for SAP ASE installer. It also provides information to help you configure the agent, view logs, and get the current agent version.

**Topics**

- [Backint for SAP ASE Installer](#)
- [Verify Authenticity before Installation](#)
- [Configuration](#)

# Backint for SAP ASE Installer

Amazon Backint Installer (`install-aws-backint-agent-ase`) is the recommended method for installing Amazon Backint Agent for ASE on an Amazon EC2 instance running SAP ASE.

1. Download `install-aws-backint-agent-ase` from Amazon S3.

   ```
   wget https://s3.amazonaws.com/awssap-backint-agent-ase/binary/latest/install-aws-
   backint-agent-ase
   ```

2. Verify authenticity of installer. (See section [the section called "Verify Authenticity"](#))
3. Start installation by running installer as `root`.

The Installer installs Amazon Backint Agent in a central directory (`/sybase/shared/aws-backint-agent-ase`). In addition, symbolic links are created at a database-specific location (e.g., `/sybase/Amazon/lib`). This allows multiple instances of SAP ASE running on the same Amazon EC2 instance to use Amazon Backint Agent.

**Installation logs** are written to `install_aws_backint_YYYYMMDD_HHMMSS.log` in the current working directory.

## Interactive Mode

Interactive mode provides a guided installation experience. You are asked to provide each configuration value before the installation starts.

```
./install-aws-backint-agent-ase
```

## Silent Mode

Silent mode allows for unattended installation without user interaction. It requires a response file - either created manually, or by running the installer in interactive mode and selecting "Create Response File".

```
./install-aws-backint-agent-ase \
    --mode silent \
    --config-file install-aws-backint-agent-ase-response.rsp
```

**Example: Response File**

```
[DEFAULT]
s3_bucket = my-s3-bucket
s3_bucket_folder = my-backup-folder
s3_sse_kms_arn = arn:aws:kms:us-east-1:111122223333:key/1abcd9b9-
ab12-1a2a-1abc-12345abc12a3
s3_sse_kms_enabled = True
aws_account_id = 111122223333
aws_region = us-east-1
sid = AWS
installation_directory = /sybase/shared
binary_location = aws-backint-agent-ase.rpm
verify_agent_binary = True
create_response_file = True
install_action = install
```

## Upgrade

The installer supports upgrading an existing installation of Amazon Backint Agent for SAP ASE. Follow the same steps as in the previous section. The installer will detect that the currently installed version is lower and perform the upgrade.

In case of issues with the upgrade process, uninstall the existing installation first.

## Uninstall

In case you want to uninstall Amazon Backint Agent for SAP ASE.

1. Remove the RPM package using the RPM package manager.

```
rpm -e aws-backint-agent-ase
```

2. Remove the Installation directory.

```
rm -r /sybase/shared/aws-backint-agent-ase
```

3. Remove stale symbolic links in database installation directory.

```
find <install dir>/lib -type l -xtype l -delete
```

# Verify Authenticity before Installation

Files for installing Amazon Backint Agent for SAP ASE - the installer (`install-aws-backint-agent-ase`) and the RPM package (`aws-backint-agent-ase.rpm`) - support signature verification. Use the Amazon Backint public key to verify that the downloaded files are original and unmodified.

## Verify Installer

1. Download public key

```
> wget https://s3.amazonaws.com/awssap-backint-agent/binary/public-key/aws-backint-agent.gpg
```

2. Import public key into your GPG keyring

```
> gpg --import aws-backint-agent.gpg

gpg: key 1E65925B: public key "AWS Backint Agent" imported
gpg: Total number processed: 1
gpg: imported: 1 (RSA: 1)
```

Make a note of the key value, as you will need it in the next step. In the preceding example, the key value is 1E65925B.

3. Verify fingerprint

```
> gpg --fingerprint 1E65925B

pub 2048R/1E65925B 2020-03-18
Key fingerprint = BD35 7A5F 1AE9 38A0 213A 82A8 80D8 5C5E 1E65 925B
uid [ unknown] AWS Backint Agent
```

The fingerprint should be equal to the following:

```
BD35 7A5F 1AE9 38A0 213A 82A8 80D8 5C5E 1E65 925B
```

> **ⓘ Note**
>
> If the fingerprint string doesn't match, don't install the agent. Contact Amazon Web Services.

4. Download the signature file `install-aws-backint-agent-ase.sig` of the installer.

```
> wget https://s3.amazonaws.com/awssap-backint-agent-ase/binary/latest/install-aws-
backint-agent-ase.sig
```

5. Verify installer (`install-aws-backint-agent-ase`) and signature (`install-aws-backint-agent-ase.sig`).

```
> gpg --verify \
    install-aws-backint-agent-ase.sig \
    install-aws-backint-agent-ase
```

> **ⓘ Note**
>
> If the output includes the phrase BAD signature, check whether you performed the procedure correctly. If you continue to get this response, contact Amazon Web Services and avoid using the downloaded files.

## Verify RPM package

By default, the installer verifies the signature of the RPM package automatically before installation. Installation fails if the authenticity of the RPM package - downloaded automatically or manually - fails.

Follow the instructions below to verify the signature manually.

1. Import GPG key into RPM database.

```
> rpm --import aws-backint-agent.gpg
```

2. Verify the signature.

```
> rpm -Kv aws-backint-agent-ase.rpm
```

```
aws-backint-agent-ase.rpm:
    Header V3 RSA/SHA1 Signature, key ID 1e65925b: OK
    Header SHA1 digest: OK
    V3 RSA/SHA1 Signature, key ID 1e65925b: OK
    MD5 digest: OK
```

> ⓘ **Note**
>
> If the command fails (return code != 0) or the key ID does not match the key ID of the
> GPG public key, do not install the RPM package. Contact Amazon Web Services.

# Configuration

Installing Amazon Backint Agent for SAP ASE with the installer creates a configuration file
automatically. The configuration file is located at `/sybase/shared/aws-backint-agent-ase/aws-backint-agent-config.yaml`. The following tables summarize the configuration
parameters added as part of the Amazon Backint agent installation process.

## Default Parameters

| Parameter | Description |
| --- | --- |
| S3BucketAwsRegion | Amazon Region of your Amazon S3 bucket. For example, `us-east-1`` . |
| S3BucketName | Name of the Amazon S3 bucket, where you want to store your SAP ASE backup files. For example, `amzon-s3-demo-bucket` . |
| S3BucketFolder | Name of the folder in the Amazon S3 bucket, where you want to store your SAP ASE backup files. For example, `my-folder` . |
| S3BucketOwnerAccountID | 12-digit Amazon account ID of the Amazon S3 bucket owner. For example, `111122223333` |
| S3SseEnabled | Specifies whether Amazon KMS encryption is enabled. |

| Parameter | Description |
|-----------|-------------|
| S3SseKmsArn | ARN of the kms-key that Amazon Backint agent can use to encrypt the backup files stored in Amazon S3. For example, `arn:aws:kms:<us-east-1>: 111122223333:key/5 bfbc9b9-ab12-ab12-a123-11111xxx22xx` . |
| LogFile | Location of the Amazon Backint agent log file. |

## Additional Parameters

You can add the following additional parameters to the configuration file.

| Parameter | Description | Default value | Allowed values |
|-----------|-------------|---------------|----------------|
| BackupObjectTags | Additional S3 object tags. | None | JSON string with the following syntax: `[{Key=string,Value=string}, {Key=string,Value=string},…]` |
| EnableTagging | Enables or disables default object tagging for backups files stored in S3. Tagging helps to identify the Amazon Backint version and SAP HANA version used during the backup. | `true` | `true, false` |
| LogLevel | Logging level for agent logs. | `info` | `info, debug` |

| Parameter | Description | Default value | Allowed values |
|-----------|-------------|---------------|----------------|
| `LogRotationFrequency` | Log file rotation frequency. | `never` | `minute`, `hour`, `day`, `never` |
| `S3StorageClass` | S3 storage class that is used for storing backups. | `STANDARD` | `STANDARD`, `STANDARD_IA`, `ONEZONE_IA`, `INTELLIGENT_TIERING` |
| `UploadConcurrency` | Number of Amazon S3 threads that can work in parallel during backup. | `100` | 1 to 200 |
| `UploadChannelSize` | Number of files that can be uploaded in parallel to the S3 bucket during the backups. | `10` | 1 to 32 |
| `DownloadConcurrency` | Number of Amazon S3 threads that can work in parallel during restore. | `100` | 1 to 200 |

# Backup and Restore

This guide shows you how to backup and restore SAP ASE databases using Amazon Backint Agent for SAP ASE. You'll learn how to install and configure the agent, create backups, restore databases, and manage logs.

**Topics**

- Interactive SQL
- Placeholders
- Backup Location
- Backup Striping
- Create Backup Configuration
- Database Backup
- Transaction Log Backup

- [Database Restore](#)
- [Transaction Log Restore](#)
- [Database Restore to Different Target Database](#)

## Interactive SQL

Interactive SQL (`isql`) is a command line interactive SQL parser to the SAP ASE server. Refer to [Using interactive isql from the Command Line](#) for more information.

## Placeholders

The commands in this document use placeholders for illustration. Their values need to be set according to your environment configuration.

| Placeholder | Description | Example Value |
|---|---|---|
| `<DB_NAME>` | Name of the database | `ASE` |
| `<MY_PREFIX>` | Custom prefix for backup identification | `my_custom_prefix` |
| `<MY_S3_LOCATION>` | Complete S3 path for restore operations | `sapaws_ASE_COMPLETE_my_custom_prefix_2/01-06-2025_1159/` |
| `<SID>` | SAP ASE System ID | `ASE` |
| `<HOSTNAME>` | Database instance hostname | `sapaws` |
| `<STRIPES>` | Number of stripes | `1` |
| `<TIMESTAMP>` | UTC timestamp of backup with format DD-MM-YYYY_hhmm | `01-06-2025_1159` |

## Backup Location

Your backup location depends on:

1. Agent configuration (S3 bucket and folder)

2. SAP ASE System ID (SID)

3. Database instance hostname

4. Your custom prefix

5. Number of stripes

6. UTC timestamp

If you're performing a backup (load), Backint agent constructs the S3 location automatically. If you're restoring a particular backup (load) created with Amazon Backint Agent, you need to specify the S3 location manually.

Database backups:

```
<FOLDER>/<SID>/<HOSTNAME>_<DB_NAME>_COMPLETE_<MY_PREFIX>_<STRIPES>/<TIMESTAMP>/
```

Transaction log backups:

```
<FOLDER>/<SID>/<HOSTNAME>_<DB_NAME>_TRANSACTION_<MY_PREFIX>_<STRIPES>/<TIMESTAMP>/
```

## Finding backups via the Amazon CLI

The Amazon CLI command below allows you to list available backups in Amazon S3.

Database backups:

```
aws s3 ls s3://<BUCKET>/<FOLDER>/<SID>/
<HOSTNAME>_<DB_NAME>_COMPLETE_<MY_PREFIX>_<STRIPES>/
```

Transaction log backups:

```
aws s3 ls s3://<BUCKET>/<FOLDER>/<SID>/
<HOSTNAME>_<DB_NAME>_TRANSACTION_<MY_PREFIX>_<STRIPES>/
```

**Example:**

```
> aws s3 ls s3://my_s3_bucket/my_s3_folder/ABC/myhost_ABC_COMPLETE_my_daily_backup_1/

                         PRE 07-01-2025_120000/
```

```
PRE 07-02-2025_120000/
PRE 07-03-2025_120000/
PRE 07-04-2025_120000/
```

# Backup Striping

Backup Striping allows adjusting the number of simultaneous backup and restore streams to and from Amazon S3. The ideal number of stripes depends mainly on the size of your database. Striping is defined by appending the following statement to your dump or load command. Striping is supported for DUMP and LOAD commands for both, database and transaction log backups.

Add the following statement (n-1) times for n stripes. The number of stripes in a load statement need to match the number of stripes when creating the backup. The maximum number of stripes is 32.

```
stripe on "awsbackint::<MY_PREFIX/MY_S3_LOCATION>"
```

# Create Backup Configuration

SAP ASE provides several methods to configure backup operations. The stored procedure sp_config_dump offers a straightforward way to set backup parameters and specify Amazon Backint Agent as the backup target.

> **ⓘ Note**
>
> Using a backup configuration will not allow you to specify a custom prefix. Instead, SAP ASE will generate a prefix with format <DB_NAME>.DB.YYYYMMDD.hhmmss.<COUNTER>.000 automatically. Creating a dump configuration is optional and not required in case you prefer using the statements in the following section.

Amazon Backint Agent for SAP ASE supports the following parameters:

| Parameter | Description | Required |
|---|---|---|
| @ext_api='awsbackint' | Specifies Amazon Backint Agent as the external API. | Yes |

| Parameter | Description | Required |
|---|---|---|
| `@num_stripes` | Number of stripe devices used during the dump operation. The default is 1. | No |
| `@blocksize` | Block size for the dump device, overriding the default block size for the device. The value must be at least 1 database page, and an exact multiple of the database page size. For optimal performance, specify blocksize as a power of 2 (such as 65,536, 131,072, or 262,144). | No |
| `@compression` | Compression level for compressed dumps. By default, compression is disabled. | No |

## Example

Create database dump configuration `my_custom_config`, using the external API for Amazon Backint Agent for SAP ASE, three stripes, a block size of 16,384 bytes, and compression level of 100.

```
sp_config_dump
@config_name='my_custom_config',
@ext_api='awsbackint',
@num_stripes='3',
@blocksize='16384',
@compression='100',
go
```

Once you have created a dump configuration, initiate a dump to Amazon S3 with the following statement.

```
dump database ASE using config='my_custom_config'
go
```

# Database Backup

Use the following statement to create a full database dump of your database in Amazon S3.

```
dump database <DB_NAME> to "awsbackint::<MY_PREFIX>"
```

## Example

Dump database ASE to S3 with prefix `my_custom_prefix`, using two stripes.

```
dump database ASE to "awsbackint::my_custom_prefix"
stripe on "awsbackint::my_custom_prefix"
go
```

# Transaction Log Backup

Use the following statement to create a transaction log dump of your database in Amazon S3.

```
dump transaction <DB_NAME> to "awsbackint::<MY_PREFIX>"
```

## Example

Dump transaction log of database ASE to S3 with prefix `my_custom_prefix`, using one stripes.

```
dump transaction ASE to "awsbackint::my_custom_prefix"
go
```

# Database Restore

Use the following statement to restore your database from Amazon S3.

```
load database <DB_NAME> from "awsbackint::<MY_S3_LOCATION>"
```

## Example

Load database ASE from S3, using two stripes.

```
load database ASE from
  "awsbackint::sapaws_ASE_COMPLETE_my_custom_prefix_2/01-06-2025_1159/"
stripe on "awsbackint::sapaws_ASE_COMPLETE_my_custom_prefix_2/01-06-2025_1159/"
go
```

## Transaction Log Restore

Use the following statement to restore your transaction logs from Amazon S3.

```
load transaction <DB_NAME> from "awsbackint::<MY_S3_LOCATION>"
```

### Example

Load transaction logs for database ASE from S3, using one stripes.

```
load database ASE from
  "awsbackint::sapaws_ASE_COMPLETE_my_custom_prefix_2/01-06-2025_1159/"
go
```

## Database Restore to Different Target Database

```
load database <TARGET_DB_NAME> from "awsbackint::<MY_S3_LOCATION>::<SOURCE_DB_NAME>"
stripe on "awsbackint::<MY_S3_LOCATION>::<SOURCE_DB_NAME>"
```

# Troubleshooting

This section helps you diagnose and resolve common issues with Amazon Backint Agent for SAP ASE. You'll find information about logging locations, performance optimization techniques, and solutions to frequently encountered problems.

**Topics**

- [Prerequisites](#)
- [Logging](#)
- [Performance](#)
- [Common Issues and Solutions](#)
- [Best Practices](#)

- [Getting Help](#)

## Prerequisites

Before troubleshooting, ensure you have:

- Amazon Backint Agent for ASE properly installed and configured
- Valid Amazon credentials with appropriate permissions
- Access to Amazon CloudTrail logs for permission debugging
- SAP ASE database administrator privileges

## Logging

Log files are located at the following locations. To get help with issues, submit all log files to Amazon Support through the Amazon Support Center.

| Log Type | Location |
|---|---|
| Installation logs | `install_aws_backint_YYYYMMDD_HHMMSS.log` in the current working directory |
| Backint Agent logs | `/sybase/shared/aws-backint-agent-ase/logs/aws-backint-agent.log` |
| Archive API logs | `/sybase/shared/aws-backint-agent-ase/logs/aws-backint-agent-archive-api.log` |
| Backup Server log | `<install_dir>/ASE-16_<minor_version>/install/<server_name>_BS.log` |
| Database log | `<install_dir>/ASE-16_<minor_version>/install/<server_name>.log` |

## Enable Detailed Logging

1. Configure the Backint Agent for verbose logging

2. Monitor Amazon CloudTrail for API call failures

3. Check SAP ASE error logs for database-specific issues

## Key Amazon Log Locations

- Amazon CloudTrail: Amazon Management Console > CloudTrail

# Performance

This section contains information on common measures for optimizing performance of Amazon Backint Agent for SAP ASE.

## Striping

Striping allows adjusting the number of simultaneous backup and restore streams to and from Amazon S3. The optimal number of stripes depends on a variety of factors, such as database size, network, and Amazon EC2 instance type. Three (3) to four (4) stripes are a reasonable choice for most database sizes. For more information on striping, refer to Backup Striping.

## Main Memory Consumption During LOAD Operations

Main memory consumption of Amazon Backint Agent for SAP ASE typically ranges between 10 and 20 GiB during LOAD operations. In case you experience significantly higher values, consider setting the following parameters in `/sybase/shared/aws-backint-agent-ase/aws-backint-agent-config.yaml`.

Lower values reduce memory consumption but may increase LOAD operation time. The default values are shown below:

| Parameter | Value |
|---|---|
| downloadConcurrencyDefault | 10 |
| downloadChannelSizeDefault | 5 |

# Common Issues and Solutions

This section covers common issues you may encounter while using Amazon Backint Agent for SAP ASE.

## Install custom version

By default, the installer automatically downloads the latest available version of Amazon Backint Agent for SAP ASE. In case you need to install a specific version of the RPM, you can download the RPM package of any previous version from S3 and provide the RPM package to the installer.

```
# Download installer
wget https://s3.amazonaws.com/awssap-backint-agent-ase/binary/<version>/install-aws-
backint-agent-ase
# Download RPM package
wget https://s3.amazonaws.com/awssap-backint-agent-ase/binary/<version>/aws-backint-
agent-ase.rpm
# Add execute permission to installer
chmod +x install-aws-backint-agent-ase
# Run installer
./install-aws-backint-agent-ase --location aws-backint-agent-ase.rpm
```

## S3 Permission Issues

### S3 Access Denied During Backup Operations

**Problem**: Backup fails with `AccessDenied` error when attempting to write to S3 bucket.

**Symptoms**:

```
Error creating upload id for bucket:<bucket-name> key:<object-key>
error:AccessDenied: User: arn:aws:sts::<account-id>:assumed-role/<role-name>/<instance-
id>
is not authorized to perform: s3:PutObject on resource: "arn:aws:s3:::<bucket-name>/
<object-path>"
because no identity-based policy allows the s3:PutObject action
```

**Root Cause**: The IAM role or user lacks the required S3 permissions for the specified bucket path.

**Solution**:

1. Verify the S3 bucket policy includes the correct resource path:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:PutObjectTagging",
    "s3:PutObject",
    "s3:GetObject",
    "s3:DeleteObject"
  ],
  "Resource": [
    "arn:aws:s3:::your-bucket-name/your-folder-name/*"
  ]
}
```

2. Ensure the resource path in your IAM policy matches the actual S3 prefix used by the Backint Agent.

3. Test permissions using Amazon CLI:

```
aws s3 cp test-file.txt s3://your-bucket-name/your-folder-name/
```

**Time to Resolution**: Typically 2-5 minutes once the correct path is identified.

**S3 ListBucket Permission Denied During Restore**

**Problem**: Database restore fails when the Backint Agent cannot list S3 bucket contents.

**Symptoms**:

```
ListObjects error: User: arn:aws:sts::<account-id>:assumed-role/<role-name>/<instance-
id>
is not authorized to perform: s3:ListBucket on resource: "arn:aws:s3:::<bucket-name>"
because no identity-based policy allows the s3:ListBucket action
```

**Root Cause**: Missing s3:ListBucket permission in the IAM policy.

**Solution**:

1. Add the s3:ListBucket permission to your IAM policy:

```
{
```

```
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::your-bucket-name"
      ]
    }
```

### S3 GetObject Permission Denied During Restore

**Problem**: Restore operation fails when attempting to download backup files from S3.

**Symptoms**:

```
Key metadata fetch failed for bucket:<bucket-name> key:<object-key>
error:Forbidden: Forbidden
status code: 403
```

**Root Cause**: The IAM role lacks `s3:GetObject` permission for the backup files.

**Solution**:

1. Ensure your IAM policy includes `s3:GetObject` permission:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::your-bucket-name/your-folder-name/*"
  ]
}
```

2. Verify the object exists and the path is correct using Amazon CLI:

```
aws s3 ls s3://your-bucket-name/your-folder-name/
```

# Backup and Restore Configuration Issues

## Stripe Count Mismatch During Restore

**Problem**: Restore fails due to mismatch between backup and restore stripe configuration.

**Symptoms**:

```
Error: Mismatch stripe count, load requested with: 2 stripes instead of 1
```

**Root Cause**: The restore command specifies a different number of stripes than used during backup.

**Solution**:

1. Check the original backup configuration to determine the correct stripe count.
2. Use the same stripe count for restore as was used for backup:

```
-- For single stripe backup
LOAD DATABASE your_db FROM "awsbackint::backup_prefix"

-- For multi-stripe backup
LOAD DATABASE your_db FROM "awsbackint::backup_prefix"
STRIPE ON "awsbackint::backup_prefix"
```

3. If unsure about the original stripe count, check the S3 bucket structure to count the stripe files.

## Named Pipe Creation Failure

**Problem**: Backup fails when attempting to reuse the same prefix after a previous failed attempt.

**Symptoms**:

```
Error creating named pipe for stripe 0: File exists
```

**Root Cause**: Previous backup attempt left temporary files that weren't cleaned up.

**Solution**:

1. **Immediate workaround**: Use a different backup prefix:

```
DUMP DATABASE your_db TO "awsbackint::new_unique_prefix"
```

2. **Long-term solution**: Clean up temporary files from the previous attempt:

```
# Remove temporary pipes and files from the Backint Agent working directory
rm -f /sybase/shared/aws-backint-ase/pipes/<database>/<prefix>
```

3. **Best practice**: Always use unique prefixes for each backup operation.


## Configuration File Issues

### Configuration File Permission Errors

**Problem**: Backint Agent fails to start due to incorrect file permissions on the configuration file.

**Symptoms**:

```
Error reading config file Failed to validate config file permission,
other users outside of group cannot have execution permission.
```

**Root Cause**: The configuration file has overly permissive file permissions.

**Solution**:

1. Set correct permissions on the configuration file:

```
chmod 640 /sybase/shared/aws-backint-agent-ase/aws-backint-agent-config.yaml
```

2. Ensure the file is owned by the correct user and group:

```
chown sap_user:sap_group /sybase/shared/aws-backint-agent-ase/aws-backint-agent-config.yaml
```

3. Verify permissions:

```
ls -la /sybase/shared/aws-backint-agent-ase/aws-backint-agent-config.yaml
```


### Missing S3 Bucket Configuration

**Problem**: Backint Agent fails to start due to missing or incorrect S3 bucket configuration.

**Symptoms**:

```
Error: S3 Bucket Name not found in aws-backint-agent-config.yaml
```

**Root Cause**: The configuration file is missing the S3 bucket name or has incorrect permissions.

**Solution**:

1. Verify the S3 bucket name is correctly specified in the configuration file.

2. Ensure the configuration file has proper read permissions.

3. Check that the IAM role has access to the specified S3 bucket.

## KMS Encryption Issues

### KMS Key Type Mismatch

**Problem**: Backup fails when using an asymmetric KMS key for S3 encryption.

**Symptoms**:

```
KMS.InvalidKeyUsageException: You cannot generate a data key with an asymmetric CMK
```

**Root Cause**: S3 server-side encryption requires a symmetric KMS key, but an asymmetric key was specified.

**Solution**:

1. Use a symmetric KMS key for S3 encryption.

2. Create a new symmetric key if needed:

   ```
   aws kms create-key --description "S3 encryption key for SAP backups" --key-usage
     ENCRYPT_DECRYPT
   ```

3. Update your S3 bucket configuration or Backint Agent configuration to use the symmetric key.

### KMS Permission Denied

**Problem**: Backup fails due to insufficient KMS permissions.

**Symptoms**:

```
AccessDenied: User: arn:aws:sts::<account-id>:assumed-role/<role-name>/<instance-id>
 is not authorized to perform: kms:GenerateDataKey on resource:
  arn:aws:kms:<region>:<account-id>:key/<key-id>
with an explicit deny in an identity-based policy
```

**Root Cause**: The IAM role lacks the required KMS permissions.

**Solution**:

1. Add KMS permissions to your IAM policy:

```
{
  "Effect": "Allow",
  "Action": [
    "kms:GenerateDataKey",
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:region:account-id:key/key-id"
  ]
}
```

2. Ensure there are no explicit deny policies blocking KMS access.

3. Test KMS access:

```
aws kms generate-data-key --key-id your-key-id --key-spec AES_256
```

# Best Practices

## Backup Naming Convention

- Use consistent and descriptive prefixes
- Include timestamp information for easy identification
- Avoid special characters that might cause parsing issues

## S3 Path Structure

The Backint Agent creates the following S3 path structure:

```
<bucket-name>/<path-prefix>/<SID>/
sapci_<Database_Name>_COMPLETE_<custom_prefix>_<stripe_number>/<date>_<time>/
```

## Restore Command Format

When restoring, use the exact path format:

```
LOAD DATABASE <DB_NAME> FROM "awsbackint::<backup_location_without_database_prefix>"
```

> ⓘ **Note**
>
> Do not include the database name as a prefix in the restore path.

## IAM Policy Template

Use this comprehensive IAM policy template for the Backint Agent:

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketPolicyStatus",
                "s3:GetBucketLocation",
                "s3:ListBucket",
                "s3:GetBucketAcl",
                "s3:GetBucketPolicy"
            ],
            "Resource": [
                "arn:aws:s3:::your-bucket-name/*",
                "arn:aws:s3:::your-bucket-name"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:GenerateDataKey"
            ],
```

```
            "Resource": "arn:aws:kms:us-east-1:123456789012:key/your-kms-key-id"
        },
          {
            "Effect": "Allow",
            "Action": [
                "s3:PutObjectTagging",
                "s3:PutObject",
                "s3:GetObject",
                "s3:DeleteObject"
            ],
            "Resource": "arn:aws:s3:::your-bucket-name/your-folder-name/*"
        }
    ]
}
```

# Getting Help

If you continue to experience issues:

1. Check Amazon Service Health Dashboard for S3 or KMS service issues

2. Review Amazon CloudTrail logs for detailed error information

3. Contact Amazon Support with specific error messages and request IDs

4. Consult the Amazon Backint Agent for SAP documentation

# Version history

The following table summarizes the changes for each release of Amazon Backint Agent for SAP ASE.

| Version | Details | Release date |
|---------|---------|--------------|
| 1.0.0 | Initial release | July 4th, 2025 |