PUBLIC

SAP Adaptive Server Enterprise on AWS 16.0 SP04 PL01

Document Version: 1.0 – 2021-11-05

Quick Start Guide for Subscription Versions of SAP ASE on AWS, Public Cloud Edition, Premium Version

THE BEST RUN SAP

Content

1 Quick Start Guide for Subscription Versions of SAP ASE on AWS

This book describes the installation and administration of SAP ASE on the Amazon Web Services (AWS) platform.

2 Overview

Amazon Web Services (AWS) is a secure cloud platform that allows you to run, manage, and administer your SAP ASE server in the cloud.

Running SAP ASE on AWS enables you to quickly move your workloads from an on-premises server farm to the cloud, taking advantage of its lower cost of ownership and enterprise-level services, and faster deployment time. Once you are running on AWS, you pay only for the server time used, reducing the amount of hardware that remains idle because it is unnecessary. The start-up costs of running on AWS are also much lower than purchasing and maintaining your own hardware.

With a few exceptions, an SAP ASE running on AWS has the same functionality as an on-premises version, and you perform the same administrative tasks like backups, dumps, loads, and so on. Clients connect to the SAP ASE running on AWS in the same way as they connect to an on-premises SAP ASE.

SAP ASE version 16.0 SP04 PL01 running on AWS includes the Premium and the standard versions. These versions have the same functionality except the Premium version supports the MemScale option and HADR.

SAP offers a safe path for migrating your data from your on-premises data center to SAP ASE on AWS, regardless of the platform or dataserver you use. For more information, see *Migrating Data to SAP ASE on AWS*.

3 Requirements, Configuration, and Restrictions

SAP ASE starts with a default configuration and includes a number of restrictions. SAP ASE starts and stops when the Amazon Elastic Compute Cloud (EC2) instance on which it is running starts and stops.

Default Configuration

SAP ASE on AWS is configured with the following by default:

- 16 K page size.
- us_english default language, utf8 character set, and binary sort order.
- Client drivers (ODBC, JDBC) for Linux and Windows.
- Windows drivers in a compressed ZIP file for client development.
- SAP ASE cockpit, installed and enabled.
- Backup Server.
- Job Scheduler.
- SAP ASE optimized for each supported instance type (including CPU and memory), running on AWS resources optimized for SAP ASE.

SAP ASE on AWS includes the following options:

- Security and Directory Services
- Encryption
- Compression
- Partitioning
- ASE Workload Analyzer

Restrictions

SAP ASE on AWS does not support these items:

- Moving to an instance that is a different size
- Tivoli Storage Manager
- Simplified native access plans (only restricted in standard version)
- Non-essential 32-bit binaries
- Veritas Cluster Server (high availability)
- Active messaging
- XML Web Services
- Compatibility mode
- QPTune utility

- CTLIB and DBLIB
- Real-time data services
- EA Server
- Java Message Service (JMS)
- TIBCO Enterprise Message Service (EMS)
- IBM WebSphere MQ
- Web Services
- Process kernel mode
- Raw disk partitions
- Allpages locking and data-page locking
- SAP ASE Cluster Edition

Licensing

Your license ensures that all supported features are available for EC2 instances that are launched from SAP-hosted SAP ASE AMI (Amazon Machine Images). This license works if you create a private AMI from an EC2 instance that is launched from an SAP-hosted ASE AMI. However, it inherits the SAP product code and is metered by AWS.

These licenses do not restrict the number of cores you can use because they are metered by AWS. Instead, the license ensures it is valid when the EC2 is launched from an AMI that is metered by AWS, and restricts the features according to what is available, depending on the SAP ASE cloud edition.

3.1 SAP ASE on AWS Configuration

Keep configuration considerations in mind when running SAP ASE on AWS, including elastic IP addresses, storage, and the overall product architecture.

Elastic IP Addresses

If you don't specify an IP address at start-up, SAP ASE uses a global IP address. Because the IP address changes when you restart the instance, applications have difficulty using SAP ASE without making changes to the interfaces file they use or other connectivity mechanisms to update the address. For this reason, SAP ASE may associate the instance with an Elastic IP address, which provides a persistent address for the application. Elastic IP addresses also allow you to upgrade SAP ASE by instantiating a new instance and moving the data or data volumes, making the upgrade transparent to applications because you reassign the Elastic IP.

Storage on AWS

SAP ASE uses these types of AWS storage:

- Amazon Elastic Block Store (EBS) – persistent storage that provides sufficient I/O bandwidth and reliability. SAP ASE creates all the storage for $SYBASE, data, and other files on EBS volumes (aside from the error log and backup files).
- Instance store – stores all databases on Amazon EBS.
- Amazon Simple Storage Service (Amazon S3) – space to which you can manually move files.

Amazon EBS provides volume types that differ in performance, characteristics, and price. SAP ASE uses the General Purpose SSD (gp2) EBS volume type. gp2 volumes offer cost-effective storage for a wide variety of workloads. Baseline performance scales linearly at 3 IOPS per gigabyte of volume size.

This table shows the throughput and IOPS for different volume types:

	Solid-state drives (SSDs)		**Hard disk drives (HDDs)**	
Volume type	(Default) General Purpose SSD	Provisioned IOPS SSD	Throughput Optimized HDD	Cold HDD
API name	gb2	io1	st1	sc1
Volume size	1 GiB – 16 TiB	4 GiB – 16 TiB	500 GiB – 16 TiB	500 GiB – 16 TiB
Maximum IOPS per volume	10,000	32,000	500	250
Maximum throughput per volume	160 MiB per second	500 MiB per second	500 MiB per second	250 MiB per second
Maximum IOPS per instance	80,000	80,000	80,000	80,000
Maximum throughput per instance	1,750 MiB per second	1,750 MiB per second	1,750 MiB per second	1,750 MiB per second
Dominant performance attribute	IOPS	IOPS	MiBs per second	MiBs per second

For more information about EBS volume types, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#w146aac23c29c17b9c29b4c14 .

You can change the volume type SAP ASE uses according to the needs of your application (for example, from gp2 to io1). For additional information about changing the size, IOPS, or type of an EBS volume, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html?icmpid=docs_ec2_console .

Architecture of SAP ASE on AWS

This image describes the architecture of the components and networking requirements of SAP ASE running on AWS, which includes SAP ASE and its storage and IP addresses inside the virtual private cloud and availability zone, and outside of this is additional S3 and Glacier storage:

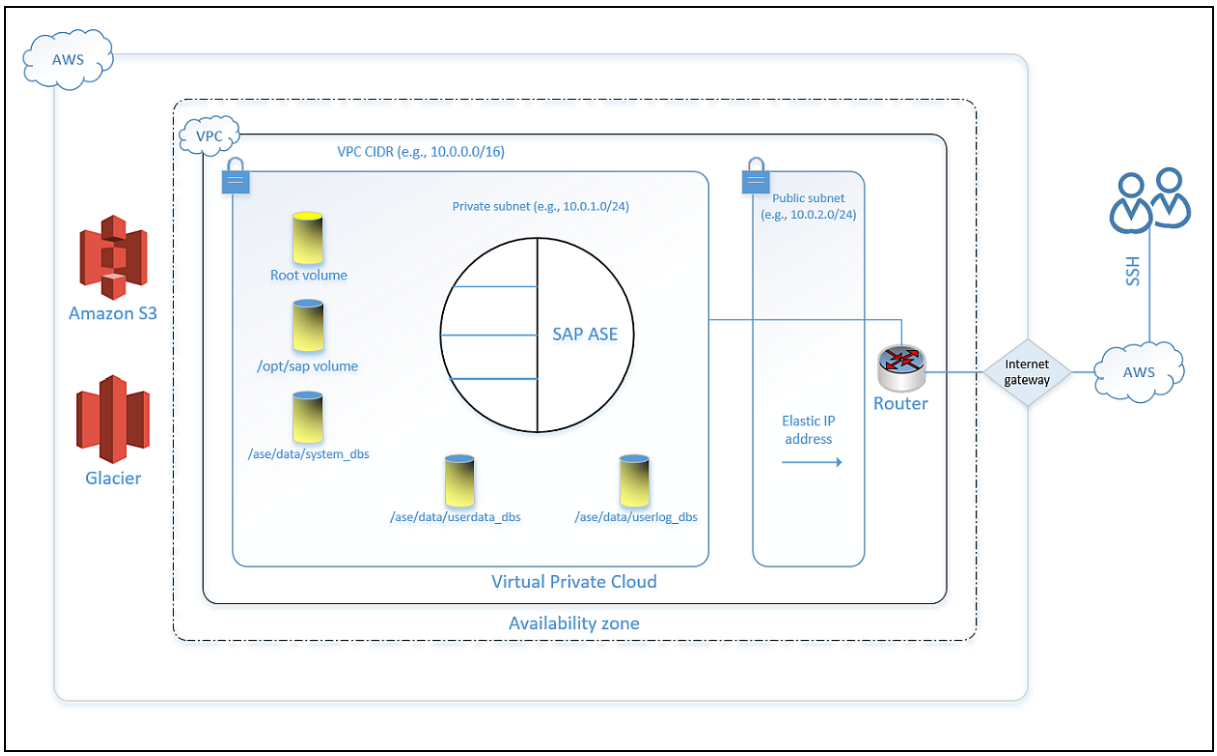

3.1.1 Security

SAP ASE running on AWS provides many levels of security to protect your data.

To prevent unnecessary exposure, ports on EC2 instances are restricted to those required for accessing SAP ASE and related products. You should restrict `ssh` access to the EC2 host by using an IP-address whitelist.

SAP ASE supports column- and database-level encryption, which can provide an additional layer of security. SSL is enabled on AWS by default, providing security for the data between the client and the server. All files containing data (for example, EBS) are located on encrypted devices. You can use AWS-generated or your own keys for encryption.

You can also manually enable SSL on SAP ASE to ensure a secure connection. See SSL Overview.

3.1.2 MemScale Option

SAP ASE version 16.0 SP03 PL06 and later on AWS supports the MemScale option.

The MemScale option increases transactional throughput, minimizes latency, and improves the efficiency of SAP ASE in high-core count machines and leverages newer memory architectures, resulting in faster query execution and response time. It also allows faster storage performance.

The MemScale option is disabled by default. You must enable the MemScale option during your SAP ASE installation on AWS.

The features that comprise the MemScale option are described below.

Latch-Free Indexes

Latch-free indexes allow you to create indexes that do not require latches during traversals or scans, which can reduce contention. See Using Latch-Free Indexes.

Simplified Native Access Plans (SNAP) on Linux

SNAP compiles lava execution plans into native code, which is then invoked directly, allowing for faster execution of extreme online transaction processing (XOLTP) queries. See Simplified Native Access Plans.

Transactional Memory

Transactional memory reduces spinlock contention within the lock manager and the buffer manager, which can impact SAP ASE throughput and response time. See Transactional Memory.

In Memory Database option

In-memory databases store transactionally active data, providing a storage repository for hot data and extends the table's standard, page-based storage. The in-memory database option includes:

- In-memory database (IMDB) – a page-oriented, zero-disk footprint in-memory database that uses the SAP ASE buffer cache to provide in-memory storage for the entire database. See In-Memory Databases.
- In-memory row storage (IMRS) – an in-memory, row-oriented storage cache that stores frequently accessed data from "hot" tables or partitions. Older, in-frequently accessed data are stored in the page-based buffer cache. Distinct from an IMDB, only part of a disk-resident database's data is stored in-memory in an IMRS, while large amounts of the database are stored on disk, accessed through the buffer cache. See In-Memory Row Storage.
  IMRS cache supports these features:

- Data row caching of hot data rows. See Data Row Caching.
- Indexes with hash caching. See Hash-Cache BTree Indexes.
- In-memory and on-disk multiversion concurrency control (MVCC) – snapshot isolation in the IMRS. See Multiversion Concurrency Control.

4 Creating an SAP ASE Instance

Configure SAP ASE to create an instance by providing information to the AWS wizard.

4.1 Preparing for the Installation

Before installing SAP ASE on AWS, sign up for AWS and create the IAM user, key pair, virtual private cloud, and Elastic IP address.

Create Your AWS Account

Create an AWS account before installing SAP ASE (AWS only charges you for the services you use):

1. Go to https://aws.amazon.com/ .
2. Click *Create a Free Account*.
3. Enter your information:
   - AWS account name (note this name for future reference).
   - Email address.
   - Password, then confirm.

AWS automatically prompts for your phone number so it can call and verify your identity.

Create the IAM User

Users who are granted AWS Identity and Access Management (IAM) privileges can access and use AWS resources and services. Only users with root privileges can grant the IAM user role to another user. There are no additional charges for adding the IAM privileges to your users.

Add IAM privileges:

1. Go to https://console.aws.amazon.com/iam/home#/home .
2. Select *Users* from the dashboard.
3. Click *Add User*.
4. Follow the online instructions to add the IAM privileges to the user.

For more information about IAM, see http://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html .

Create a Key Pair

AWS requires that you specify a key pair when you create your SAP ASE instance. Key pairs are region-specific. You can either import your own key pair or create one from your AWS instance.

Create a key pair from AWS:

1. From the EC2 console *Dashboard*, select *Network* and *Security* > *Key Pairs*.
2. Select *Create Key Pair* or *Import Key Pair*.
3. Follow the online instructions.

For more information about AWS key pairs, see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html .

Select a Virtual Private Cloud (VPC)

The Amazon Virtual Private Cloud (VPC) lets you set up a virtual private network that you define for your resources. When you create your instance, you specify the VPC in which your instance runs. Generally, a VPC contains an instance launched within it. These instances are isolated from the outside world; however, they can share information and connect to each other.

You can communicate directly with your instance running on this VPC from your data center using SSH.

You may use either the default Amazon VPC or configure your own when you deploy an instance. Users often select the default VPC because it offers faster deployment. In this case, users are not concerned with their instance's visibility to other instances deployed in the default VPC.

1. To determine the default VPC and its subnet:
    1. From the *Services* page, select *VPC* (under the *Networking & Content Delivery* heading).
    2. Select *Your VPCs* from the *VPC Dashboard*.
    3. Identify the default VPC by locating the value of *Yes* in the *Default VPC* column, and note its VPC ID value:

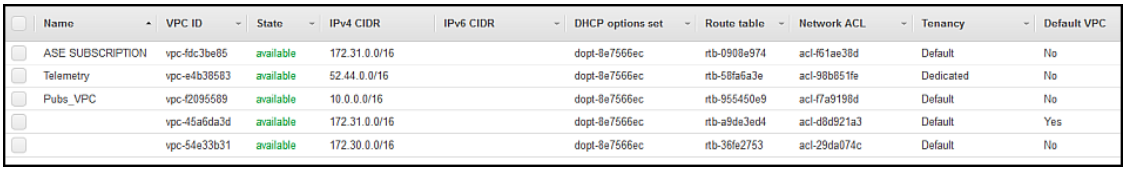

Name	VPC ID	State	IPv4 CIDR	IPv6 CIDR	DHCP options set	Route table	Network ACL	Tenancy	Default VPC
ASE SUBSCRIPTION	vpc-fdc3be85	available	172.31.0.0/16		dopt-8e7566ec	rtb-0908e974	acl-f61ae38d	Default	No
Telemetry	vpc-e4b38583	available	52.44.0.0/16		dopt-8e7566ec	rtb-58fa6a3e	acl-98b851fe	Dedicated	No
Pubs_VPC	vpc-f2095589	available	10.0.0.0/16		dopt-8e7566ec	rtb-955450e9	acl-f7a9198d	Default	No
	vpc-45a6da3d	available	172.31.0.0/16		dopt-8e7566ec	rtb-a9de3ed4	acl-d8d921a3	Default	Yes
	vpc-54e33b31	available	172.30.0.0/16		dopt-8e7566ec	rtb-36fe2753	acl-29da074c	Default	No

You will use the VPC ID when you configure your instance. In this example, the VPC ID is vpc-45a6da3d.

4. Verify the *DNS hostnames* line in the *Description* tab is set to *Enabled* for this VPC.

If it is not:

1. Select *Actions* > *Edit DNS hostnames*.
2. Check the box for *enable*.
3. Click *Save*.
4. Click *Close*.

See *DNS Support in Your VPC* on https://docs.aws.amazon.com/vpc/latest/userguide/vpc-dns.html#vpc-dns-hostnames for more information about enabling and disabling DNS hostnames.

5. Select *Subnets* from the *VPC Dashboard*.
6. Identify the VPC ID listed on the *VPC* column that is associated with the default VPC, and make a note of the subnet associated with this VPC (there may be many of them):

Name	Subnet ID	State	VPC	IPv4 CIDR	Available IPv4	IPv6 CIDR	Availability Zone	Route Table
	subnet-33607a75	available	vpc-54e33b31	172.30.2.0/24	251		us-east-1c	rtb-36fe2753
Telemetry Subnet	subnet-2269ad0f	available	vpc-e4b38583 | Telemetry	52.44.0.0/16	65531		us-east-1a	rtb-58fa6a3e
	subnet-6daa521a	available	vpc-54e33b31	172.30.1.0/24	251		us-east-1b	rtb-36fe2753
	subnet-be16e5e3	available	vpc-45a6da3d	172.31.32.0/20	4086		us-east-1c	rtb-a9de3ed4
Pubs_Subnet	subnet-d95aa3be	available	vpc-f2095589 | Pubs_VPC	10.0.0.0/24	250		us-east-1d	rtb-6656521a | My...
ASE Subnet	subnet-8b0878ef	available	vpc-fdc3be85 | ASE SUBSCRIPTI...	172.31.0.0/16	65530		us-east-1d	rtb-0908e974
	subnet-b41c269c	available	vpc-54e33b31	172.30.0.0/24	250		us-east-1a	rtb-36fe2753
	subnet-f94d50f5	available	vpc-45a6da3d	172.31.48.0/20	4086		us-east-1f	rtb-a9de3ed4
	subnet-43463527	available	vpc-45a6da3d	172.31.0.0/20	4090		us-east-1d	rtb-a9de3ed4
	subnet-ab11ef84	available	vpc-45a6da3d	172.31.80.0/20	4052		us-east-1a	rtb-a9de3ed4
	subnet-47d5bf0c	available	vpc-45a6da3d	172.31.16.0/20	4079		us-east-1b	rtb-a9de3ed4
	subnet-68810957	available	vpc-45a6da3d	172.31.64.0/20	4089		us-east-1e	rtb-a9de3ed4

You will use this subnet when you configure your instance.

> **i Note**
>
> Subnets have internal IP addresses with the subnet's CIDR. Make a note of this IP address because they are used to limit the number of instances deployed in each subnet. For example, the number of instance deployable within the IPv4 CIDR for IP address 172.31.16.0/20 is:
>
> $[2^{12} - (4 \text{ reserved instances}) = 4,092]$

2. Create a VPC and configure its rules for connecting from within and outside AWS:
   1. Go to https://console.aws.amazon.com/vpc/ .
   2. Click *Start VPC Wizard*.
   3. Select the table for the VPC you need to create with one of the following configurations:
      - A single public subnet.
      - Public and private subnets.
      - Public and private subnets and hardware access.
      - Private subnet with hardware VPN access.

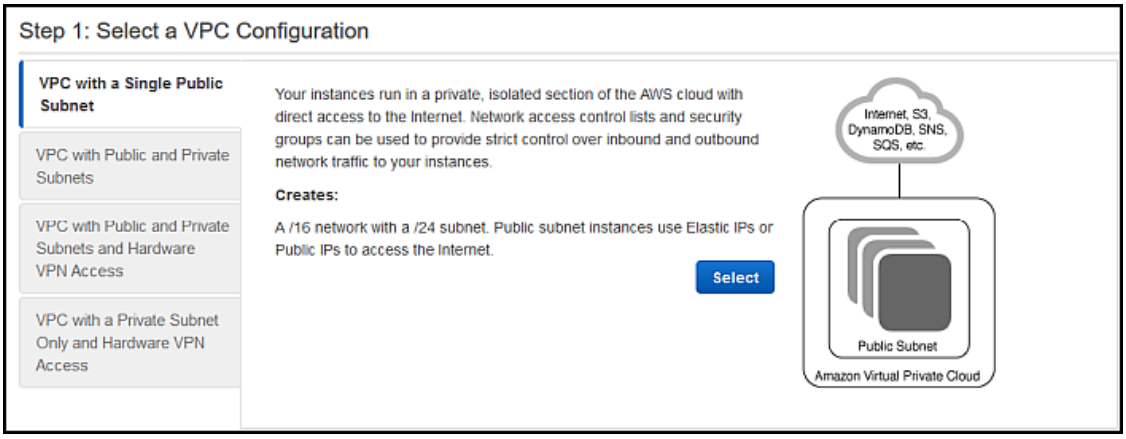

1. Depending on your selection, the wizard asks for the appropriate networking information. Use the AWS defaults if you do not have a preference for the networking information, or if you are going to customize the network selections later.
   The image below shows the information required for a VPC with public and private subnets:

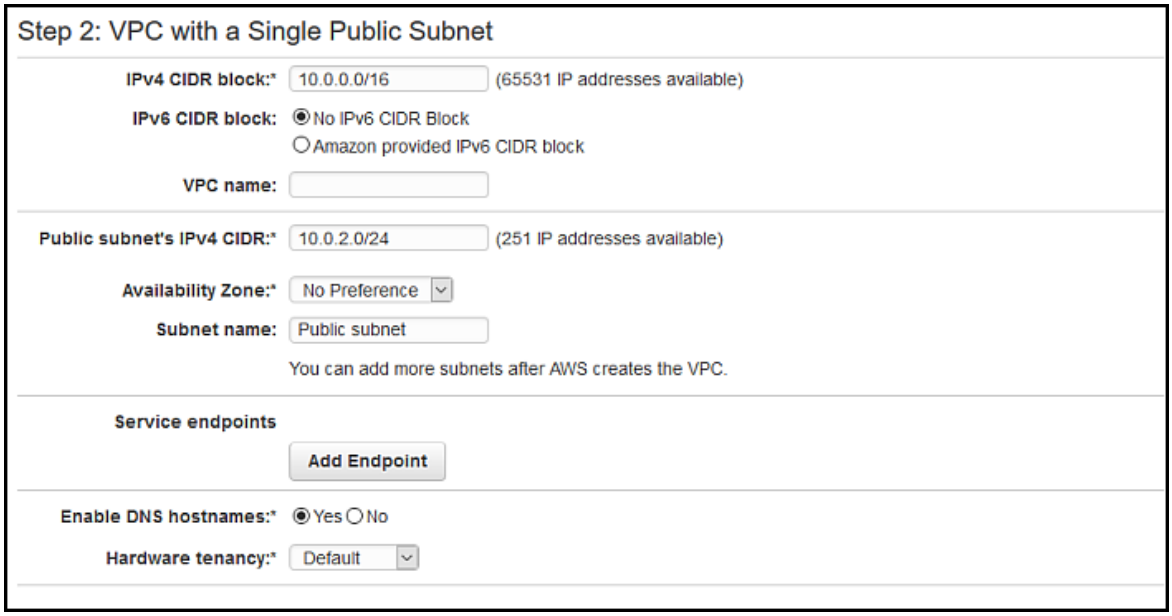

For more information about creating VPCs, see https://aws.amazon.com/documentation/vpc/ .

4.2 Naming Volumes for Future Reference

SAP ASE subscription versions running on AWS include five volumes.

The five volumes are:

Purpose	Size	System Mount Point
Operating system disk	20 GB	/
SAP ASE installation	10GB	`/opt/sap`
SAP ASE master devices	10 GB	`/ase/data/system_dbs`
User database data	User-determined	`/ase/data/userdata_dbs`
User database log	User-determined	`/ase/data/userlog_dbs`

The operating system, SAP ASE installation, and master device volumes are all deleted when you terminate the AWS EC2 instance or delete the stack. However, the user database data and log volumes are retained with your account under AWS volumes when you terminate the AWS EC2 instance or delete the stack.

Because AWS retains these volumes, provide them with a descriptive name for future reference (for example, using the size and creation date) so that users can identify them after the AWS instance or stack are gone.

Generally, naming your volumes involves this process:

1. While you are building the instance, indicate the size of user data and log volumes on the *Specify Details* page of the *Create Stack* wizard. For example, the following depicts 200 GB data and 40 GB log volumes for the `testdb` user database:

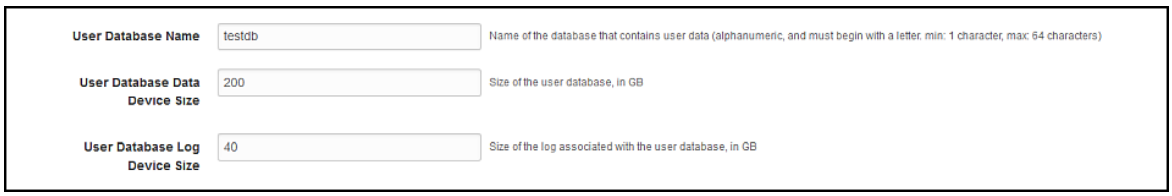

2. Once the instance in created, provide your volumes with descriptive names.
   1. Select *EC2 Dashboard* on AWS console.
   2. Select *Volumes* from the *EC2 Dashboard*.
   3. Select the pencil icon from the line provided for your volume and enter a descriptive name in the field:

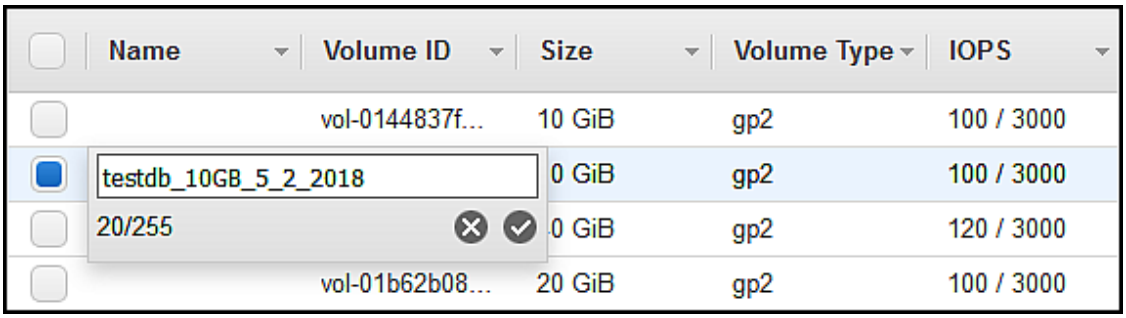

4.3 Configuring SAP ASE for Automatic Backups

SAP ASE on AWS allows you to configure the server for automatic backups.

Auto Backup is configured when you first install SAP ASE on AWS. Once enabled, SAP ASE performs a backup of a user database once a day, and performs a backup of transaction logs once every hour.

> **i Note**
>
> This version of SAP ASE on AWS supports Auto Backup for a single user database which was created during the SAP ASE initialization on AWS.

Configuring Auto Backup creates a file named `autobackup.sh` in the `/ase/config` directory of the AWS instance. Once SAP ASE is initialized, it executes the first `dump database` within a day, and executes the first `dump transaction` within an hour. However, since `dump transaction` requires an initial `dump database` command to run first, all `dump transaction` commands will fail until the first `dump database` completes.

Generally, it takes about one day before the Auto Backup feature performs all of its backups smoothly.

Each successful `dump database` or `dump transaction` operation generates a dump file in the `/ase/data/autobackup_dbs` directory on the EC2 instance. The file name format is:

```
<database_name>.<database_type>.<year_month_day_minute_second>.dmp
```

Where `<database_type>` is one of:

- `dmpdb` – for database dumps
- `dmptx` – for transaction dumps

For example, these are dumps of the `pubs2` database and the transaction log in the `/ase/data/autobackup_dbs` directory:

```
ls /ase/data/autobackup_dbs
pubs2.dmpdb.20180727080004.dmp
pubs2.dmptx.20180730031501.dmp
```

After the dump file is generated, it is immediately copied to the archive location on S3 storage (specified when you configure SAP ASE for Auto Backup). Once this dump file is copied to the S3 location, it is deleted from `/ase/data/autobackup_dbs`.

SAP ASE does not delete the dump file if it is not successfully copied to the archive location. The next execution of `dump database` or `dump transaction` command attempts to generate another dump file in `/ase/data/autobackup_dbs`; if there is enough space, this next dump file is successfully created, and is subsequently transferred to the archive location and deleted from the local directory. If the local directory does not have sufficient space for the new dump file due to the file left by the previous dump, the next dump command fails until space is made available.

By default, the dump history feature is enabled when you configure SAP ASE on AWS. After each database dump, SAP ASE generates a dump history file named `dumphist` in the `/ase/config/ASE-16_0/` directory. After each dump, the `autobackup.sh` script saves a copy of this file to the S3 archive folder, and rename it using this convention:

```
dumphist.<YYYYmmddHHMMSS>
```

For example:

```
dumphist.20180729141503
```

4.3.1 Considerations When Using Auto Backup

There are a number of items to consider when using Auto Backup.

Auto Backup Unavailable Due to Database Initialization

SAP ASE on AWS creates user databases with `async_init` enabled, so the database is immediately available when it is created or altered, not when the database initialization is complete. The initialization occurs transparently to the database user because the initializer works ahead of any commands that require space in the database. You will see this message in the `dumpdb.log` if you perform a dump of the database before the database initialization is complete:

```
Detect DATABASE is with async_init status, Abort dump this time and retry at
next cron job
```

Because Auto Backup schedules `dump database` and `dump transaction` commands with a cron job, in rare circumstances these commands may overlap: `dump transaction` is about to start but `dump database` is performing a dump operation. SAP ASE issues this message to `dumptran.log` in this situation:

```
Detect DUMP DATABASE cmd while DUMP TRAN, Abort dump tran once and retry at next
cron job
```

Granular Permissions

Granular permissions are enabled by default if Auto Backup is choosen on SAP ASE on AWS. Do not disable granular permissions. Auto Backup relies on granular permissions for `dump database` and `dump transaction` permissions. If you disable granular permissions, `dump database` and `dump transaction` fail due to a permission error.

See:

- Initialize Databases Asynchronously
- Granular Permissions

4.4 Security Group

AWS requires your instance to have a security group, which acts as a virtual firewall and regulates traffic to your instance. The SAP ASE cloud formation template (CFT) creates a security group when you start SAP ASE. You

can have more than one security group associated with each instance, and you can modify existing security groups later.

SAP ASE uses these ports for IP addresses, which you provide for the `SSHLocation` parameter:

- 22 – for ssh
- 5000 – for SAP ASE
- 4283 – for SAP ASE cockpit

See here for more information about creating security groups.

See "Amazon EC2 Security Groups for Linux Instances" in the Amazon EC2 *User Guide for Linux Instances* at https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html for more information about creating security groups.

4.5 Launching an SAP ASE Instance

Use the one-click launch to provision your SAP ASE on AWS.

Prerequisites

Sign up for AWS and create the IAM user, key pair, virtual private cloud, and the Elastic IP address before launching SAP ASE on AWS. See Preparing for the Installation [page 11].

> **i Note**
>
> When you create an AWS account, you are automatically signed up for all AWS services. However, you are charged only for the services you use. Before you create launch SAP ASE on AWS, make sure you read the *SAP ASE What's New Guide*.

Procedure

1. Go to https://aws.amazon.com/marketplace .
2. If you have not done so yet, sign in to AWS.
3. Enter **SAP ASE** in the search bar or browse by category from the *Databases* icon:

4. Select the link for the 16.0 SP04 PL01 or later premium version.
5. Click *Continue to Subscribe*.
6. Click *Continue to Configuration*.
7. Specify the *CloudFormation* fulfillment option, your *Region*, and click *Continue to Launch*.
8. Select the *Launch CloudFormation* action and select *Launch*.
9. Click *Next* to deploy the SAP ASE template.
10. Enter values for:

    SAP ASE public cloud edition version 16.0 SP04 PL01

    - *Stack name*
    - *Instance Type* – size of the EC2 instance on which SAP ASE runs. SAP ASE supports these sizes:
        - m5.2xlarge
        - m5.4xlarge
        - m5.12xlarge
        - x1.16xlarge
    - *Instance Name* – name of the instance you are creating.
    - *KeyName* – name of the key pair you are using to encrypt and decrypt login information.
    - *SSH Location* – IP address range for using SSH to connect with this instance. Either provide a whitelist of allowed IP addresses. Although it is possible to use the value 0.0.0.0/0 to allow all users, you should avoid doing this.
    - *vpc* – name of the existing VPC used to create the security group.
    - *subnet* – name of the subnet from which you are launching this instance. This subnet must belong to the VPC you are using. See the VPCs and their associated subnets by doing the following:
        1. From the *AWS Services* page, select *Networking and Content Delivery* > *VPC*.
        2. Select *Subnets* from the *VPC Dashboard*.
    - *SAP ASE Server Name* – name of the SAP ASE you are launching.
    - *Master Password* – password the system administrator (sa) uses to log in to SAP ASE and the SAP ASE Cockpit.
    - *Language* – language the server users. Default is `us_english`.
    - *User Database Name* – name of the user database into which you will migrate your on-premises data.
    - *User Database Data Device Size* – size of the user database for your migrated data.
    - *User Database Log Device Size* – size of the log device for the user database.
    - *Auto Backup feature* – enables or disables the Auto Backup feature. The default value is *enable*. If you select *disable*, you need not enter values for the *Auto Backup user password*, *Volume size for dump files*, and *S3 bucket for dump files* parameters.
    - *Auto Backup user password* – password used to backup databases. Minimum of 8 characters, maximum of 41.
    - *Volume size for dump files* – size, in gigabytes, of the volume containing the Auto Backup dump files. The default is 245 gigabytes. This volume must be larger than the size of the user database.
    - *S3 bucket for dump files* – path to the S3 bucket that contains the Auto Backup dump files. This must be a valid and existing S3 bucket path using this format:

      ```
      s3://abcde/fghij
      ```

      The default value, `s3://example.path`, is just an example; you must enter a valid path. The S3 bucket path cannot end with a `/` (that is, a forward slash). The installation program does not check the

existence of this bucket. See https://docs.aws.amazon.com/AmazonS3/latest/dev/BucketRestrictions.html for s3 bucket naming conventions.

- *Memory Scale* – offers increased throughput and reduced latency for OLTP workloads and reporting queries. See MemScale Option [page 9] for more information. Select values for:
  - *Latch-Free Indexes* – enable or disable. Allows you to create indexes that do not require latches during traversals or scans. .
  - *Simplified Native Access Plan (SNAP)* – enable or disable. Compiles lava execution plans into native code.
  - *Transactional Memory* – enable or disable. Reduces spinlock contention within the lock manager and the buffer manager.
  - *IMRS – DRC/MVCC/HCB* – enable or disable. Configures databases to store transactionally active data, providing a storage repository for hot data and extends the table's standard, page-based storage. Selecting this feature:
    - Enables the DRC and MVCC features at the database level.
    - Sets the `enable HCB index` configuration parameter to 1. However, to use HCB indexes, you must create an hash-cached BTree (HCB) index and set the value for `HCB index memory pool` once your instance is available.

    You cannot enable or disable the datarow caching (DRC), multiversion concurrency control (MVCC), or HCB features individually.

    > **i Note**
    >
    > You cannot configure both IMRS - DRC/MVCC/HCB and IMRS - On-Disk MVCC on the same server.

  - *IMRS - On-Disk MVCC* – enable or disable. Provides snapshot isolation in the IMRS. Enabled at the database level.
  - *IMDB Database Size* – size, in GB, of the IMDB database. The default value of 0 means the IMDB database is not created, and the *IMDB database name* parameter (below) is ignored. If the size is greater than 30% of the total system memory, the database size is kept at 30% of total system memory.
  - *IMDB Database Name* – name of the IMDB database.
- *Existing Elastic IP* – an existing Elastic IP address (Optional). If you do not specify one, AWS selects an available IP address for you.

11. Click *Next*.
12. (Optional) *Tags* – enter a tag for your server, which helps you find your server in the list: – name of the EC2 stack you are creating.
    a. Enter the phrase "Name" under the *Key* heading.
    b. Enter the server name under the *Value* heading.
13. (Optional) *Permissions* – name of the EC2 stack you are – enter the IAM role used to create the stack.
14. (Optional) *Rollback Triggers*:
    a. Enter the number of minutes the CloudFormation monitors roll back triggers.
    b. Enter the *Type* and *Amazon Resource Name* (ARN) for the trigger.
15. (Optional) Expand *Advanced*:
    - *Notification options* – determines if AWS notifies you about these events:
      - *No notification* – notifications are not sent.

    - *New Amazon SNS topic* – email and topic heading to which you want Amazon Simple Notification Service (SNS) messages sent.
    - *Existing Amazon SNS topic* – an existing SNS topic.
    - *Existing ARN* – an existing Amazon Resource Name (ARN).
  - *Termination protection* – enable or disable protection from the stack being deleted.
  - *Timeout* – specify whether to use a timeout. If you enable this option, enter the amount of time, in minutes, before AWS times out if the instance cannot be created.
  - *Rollback on failure* – enable or disable an instance from rolling back after a failure.
  - *Stack policy* – enter or upload a policy to enforce.
16. Select *Next*.
17. Review the server configuration. Select *Previous* to make any changes if necessary, then select *Create* to create the server.

4.6 Assigning an Elastic IP Address

Elastic IP addresses are static IP addresses that can be assigned to your instance. You are automatically assigned an Elastic IP address when you build your server.

Procedure

1. Log in to the *AWS Console*.
2. From the EC2 *Dashboard*, select *Elastic IPs* (under the *Network and Security* heading).
3. Right-click on the Elastic IP address then select *Associate address*. If no Elastic IP address is available (that is, not associated with an instance), allocate a new address:
   1. Select *Allocate new address*.
   2. Click *Allocate* to add this address to the list of Elastic IP addresses.
   3. Click *Close*.

5 Administering an SAP ASE Instance on AWS

SAP ASE on AWS allows you to perform common activities like taking regular backups, adding disk space, and changing configurations.

5.1 Configuring SAP ASE on AWS for HADR

Configuring SAP ASE on AWS for HADR requires that you restart your instances.

Procedure

1. Launch two SAP ASE instances on AWS. These instances must have:
   - Different host names for each SAP ASE server (for example, aseh1 and aseh2)
   - The same name for their cluster ID database (for example, `db1`).

     > **i** Note
     >
     > You must include the cluster ID database when you set up the SAP ASE instance. You cannot add this database later. The cluster ID database is the only database you can include during the set up for this SAP ASE subscription, and the cluster ID database name must be 3 characters in length and must start with an alphabet character.

   - The same size for their cluster ID database.
     The instances should use the same VPC. However, if you use different VPCs, their CIDR values cannot overlap, and you will need to set peer connections between the VPCs after configuring the SAP ASE subscriptions.
2. Set the inbound rules for the security groups for each instance. The TCP port numbers must be the same as those that you will use in the HADR resource file used with `setuphadr` in a subsequent step. The default port numbers are:
   - RMA port number range – RMA requires a range of ports for communication. Typically, you should reserve a range that is 10 places prior to the port number which you reserve for RMA. For example, 6990 – 7001
   - Replication Server port number – 5005
   - Backup Server – 5001

   The resource file lists the port numbers as:

```
primary.backup_server_port=5001
primary.rma_tds_port=7001
primary.rma_rmi_port=7000
```

```
primary.srs_port=5005
```

3. Use a standard connection tool (for example, PuTTY) to connect to the EC2 instance running SAP ASE.
4. As root, configure each instance. Install the `libgcc.i686` and `glibc.i686` packages. Issue:

```
yum install libgcc.i686
yum install glibc.i686
```

> **i Note**
>
> If `yum` reports a version mismatch error similar to this:
>
> ```
> --> Finished Dependency Resolution
> Error: Multilib version problems found.
> ```
>
> Perform the upgrade first then re-run the `yum` install:
>
> ```
> yum upgrade glibc
> yum install glibc
> yum upgrade libgcc
> yum install libgcc
> ```

5. Connect to the SAP ASE instance and switch to the sybase user:

```
sudo su - sybase
```

6. If auto backup is enabled, you must disable it before configuring the system for HADR.
   a. Disable any auto backups scheduled by commenting them out of existing `cron` jobs. For example, this comments out and disables two auto backup jobs:

```
crontab -e
## dump database at 03:30 every day
#30 03 * * * /ase/config/autobackup.sh database >> /ase/config/dumpdb.log
2>&1
## dump transaction on 15th minute every hour
#15 * * * * /ase/config/autobackup.sh transaction >> /ase/config/
dumptran.log 2>&1
```

7. Create the `/ase/data/dm` directory on both hosts:

```
cd /ase/data
mkdir dm
```

8. Create links for the log and runserver files (`setuphadr` uses the settings from `<$SYBASE>` to determine the location of the runserver file to restart SAP ASE):

```
ln -s /ase/config/ASE-16_0/install/* /opt/sap/ASE-16_0/install
```

9. Edit `/opt/sap/DM/RMA-16_0/instances/AgentContainer/config/bootstrap.prop` to add the line below:

```
rsge.bootstrap.tds.port.number=<rma_port>
```

   For example:

```
rsge.bootstrap.tds.port.number=7001
```

10. Source the `SYBASE.sh` (BASH shell) or `SYBASE.csh` (C shell) file to set your environment variables. For example:

```
source /opt/sap/SYBASE.sh
```

11. Create a resource file named `/ase/data/primary.rs` on the primary instance, copying in this template:

```
##<primary.rs>
cluster_id=db1   ---> Cluster ID database name for the primary SAP ASE. This
is the database you specified at setup
setup_site=primary
primary.site_role=primary
companion.site_role=companion
is_secondary_site_setup=false
synchronization_mode=sync
ase_sa_user=sa
ase_sa_password=Syb12345   ---> sa user password, defined during SAP ASE set
up
hadr_maintenance_user=DR_maint
hadr_maintenance_password=Sybase123
rma_admin_user=DR_admin
rma_admin_password=Sybase123
participating_database_1=db1  ---> Cluster ID database name for the primary
SAP ASE set up. The same database as specified by cluster_id=
materialize_participating_database_1=true
primary.ase_host_name=aseh1   ---> Host name for the primary SAP ASE
primary.site_name=aseh1       ---> Primary HADR site name, often set as the
current host name
primary.ase_release_directory=/opt/sap
primary.ase_server_name=ase1  --> Primary SAP ASE server name
primary.ase_server_port=5000
primary.backup_server_name=ase1_BS -->Primary SAP ASE backup server name
primary.backup_server_port=5001
primary.backup_server_dump_directory=/ase/data/dm  --> Location in which to
save primary server's interim dump files
primary.rma_tds_port=7001
primary.rma_rmi_port=7000
primary.srs_port=5005
primary.device_buffer_dir=/ase/data/dm  --> Location for the primary
Replication Server device buffer
primary.device_buffer_size=1000
primary.simple_persistent_queue_dir=/ase/data/dm  --> Location in which to
save the primary Replication Server queue
primary.simple_persistent_queue_size=1000
companion.ase_host_name=aseh2  ---> Host name for the companion server
companion.site_name=aseh2      ---> HADR site name for the companion server
companion.ase_release_directory=/opt/sap
companion.ase_server_name=ase2  --> SAP ASE server name for the companion
server
companion.ase_server_port=5000
companion.backup_server_name=ase2_BS  --> SAP ASE backup server name for the
companion server
companion.backup_server_port=5001
companion.backup_server_dump_directory=/ase/data/dm  --> Location in which to
save companion server's interim dump files
companion.rma_tds_port=7001
companion.rma_rmi_port=7000
companion.srs_port=5005
companion.device_buffer_dir=/ase/data/dm  --> Location for companion
Replication Server device buffer
companion.device_buffer_size=1000
companion.simple_persistent_queue_dir=/ase/data/dm  --> Location in which to
save the companion Replication Server queue
companion.simple_persistent_queue_size=1000
```

12. Edit the resource file for your site, replacing the sample values and removing the descriptive text (indicated with "-->").
13. Source the `SYBASE.sh` (BASH shell) or `SYBASE.csh` (C shell) file to set your environment variables. For example:

```
source /opt/sap/SYBASE.sh
```

14. Use the `/opt/sap/ASE-16_0/bin/setuphadr` utility to run the resource file. For example:

```
/opt/sap/ASE-16_0/bin/setuphadr primary.rs
```

    See the log in `/opt/sap/ASE-16_0/init/logs/setuphadr<date>.<00n>` for configuration and progress details.
15. Connect to the SAP ASE companion instance and switch to the sybase user:

```
sudo su - sybase
```

16. Create a resource file named `/ase/data/companion.rs` on the companion instance, copying in this template:

```
##<companion.rs>
cluster_id=db1   ---> Cluster ID database name for the primary SAP ASE. This
is the database you specified at setup
setup_site=companion
primary.site_role=primary
companion.site_role=companion
is_secondary_site_setup=true
synchronization_mode=sync
ase_sa_user=sa
ase_sa_password=Syb12345   ---> sa user password, defined during SAP ASE set
up
hadr_maintenance_user=DR_maint
hadr_maintenance_password=Sybase123
rma_admin_user=DR_admin
rma_admin_password=Sybase123
participating_database_1=db1  ---> Cluster ID database name for the primary
SAP ASE set up. The same database as specified by cluster_id=
materialize_participating_database_1=true
primary.ase_host_name=aseh1   ---> Host name for the primary SAP ASE
primary.site_name=aseh1       ---> Primary HADR site name, often set as the
current host name
primary.ase_release_directory=/opt/sap
primary.ase_server_name=ase1  --> Primary SAP ASE server name
primary.ase_server_port=5000
primary.backup_server_name=ase1_BS --> Primary SAP ASE backup server name
primary.backup_server_port=5001
primary.backup_server_dump_directory=/ase/data/dm  --> Location in which to
save primary server's interim dump files
primary.rma_tds_port=7001
primary.rma_rmi_port=7000
primary.srs_port=5005
primary.device_buffer_dir=/ase/data/dm  --> Location for the primary
Replication Server device buffer
primary.device_buffer_size=1000
primary.simple_persistent_queue_dir=/ase/data/dm  --> Location in which to
save the primary Replication Server queue
primary.simple_persistent_queue_size=1000
companion.ase_host_name=aseh2  ---> Host name for the companion server
companion.site_name=aseh2      ---> HADR site name for the companion server
companion.ase_release_directory=/opt/sap
companion.ase_server_name=ase2  --> SAP ASE server name for the companion
server
```

```
companion.ase_server_port=5000
companion.backup_server_name=ase2_BS  --> SAP ASE backup server name for the
companion server
companion.backup_server_port=5001
companion.backup_server_dump_directory=/ase/data/dm  --> Location in which to
save companion server's interim dump files
companion.rma_tds_port=7001
companion.rma_rmi_port=7000
companion.srs_port=5005
companion.device_buffer_dir=/ase/data/dm  --> Location for companion
Replication Server device buffer
companion.device_buffer_size=1000
companion.simple_persistent_queue_dir=/ase/data/dm  --> Location in which to
save the companion Replication Server queue
companion.simple_persistent_queue_size=1000
```

17. Edit the resource file for your site, replacing the sample values.
18. Source the `SYBASE.sh` file to configure your environment variables:

    ```
    source /opt/sap/SYBASE.sh
    ```

19. Use the `/opt/sap/ASE-16_0/bin/setuphadr` utility to run the resource file. For example:

    ```
    /opt/sap/ASE-16_0/bin/setuphadr companion.rs
    ```

    See the log in `/opt/sap/ASE-16_0/init/logs/setuphadr<date>.<00n>` for configuration and progress details. See Installation for information about solving installation issues.
20. Connect to each instance to check the status of the HADR system, issuing this command:

    ```
    select hadr_mode(),hadr_state()
    ```

Next Steps

See the HADR Users Guide for more information about managing the SAP ASE HADR system.

5.2 Backing Up and Restoring Your Data

A typical backup operation for SAP ASE includes dumping the database and transaction log.

This example dumps a database named `my_database`:

```
dump database my_database to '/local_dump_dir/my_database.dump'
```

This example dumps the `my_database` transaction log:

```
dump transaction my_database to '/local_dump_dir/my_database.$timestamp.dumptran'
```

> **Note**
>
> These examples require that the file system under directory `/local_dumpdir` have enough space to accommodate the database dump. If necessary, create an EBS volume larger than the required database dump size and mount it to `/local_dumpdir`.

See Backing Up and Restoring User Databases.

For security reasons, do not use the sa account to perform backup operations within SAP ASE. Instead, create a standalone account to perform the backup operations. For example use the following to create an account named `backupuser` and grant the `oper_role` to that user:

```
create login backupuser with password $password
go
grant role 'oper_role' to 'backupuser'
go
```

Using the AWS S3 repository to archive backup dump files can consume large amounts of space. For example, this command copies the `myfile.txt` file to an S3 bucket:

```
aws s3 cp myfile.txt s3://mybucket/myfile.txt
```

Before you can access an S3 bucket from your EC2 instance, you may need to configure `awscli` by specifying your AWS security credentials and default region with the `aws` configure parameter. See the *AWS CLI User Guide* at http://docs.aws.amazon.com/cli/latest/userguide/cli-chap-getting-started.html .

You can also use the SAP ASE cockpit to back up and restore SAP ASE databases. See Backup and Restore.

5.3 Administering Auto Backup

SAP ASE on AWS supports automatic backup for a single database if it is created during the SAP ASE initialization.

The database and transaction log configured for dumps are specified in the `/ase/config/autobackup_cronjob` file.

You can issue the `autobackup.sh` utility (located in `/ase/config`) from the command line. You cannot indicate the individual database or transaction log for which you are creating a dump. The `autobackup.sh` makes dumps for the database you specify when you configure the server for Auto Backup.

However, you can run the `autobackup.sh` utility manually after you configure SAP ASE. The syntax is:

```
autobackup.sh dmptype(database|transaction)
```

Where `dmptype` specifies one of:

- `database` – you are performing a dump of a user database. `autobackup.sh` allows you to perform backups only for the user database created when SAP ASE launched, and only for the database you specified when you configured SAP ASE for Auto Backup. The size of the user database determines the size of the backup created by `autobackup.sh`.

- `transaction` – you are performing a dump of the transaction logs.

You can load databases directly from the dump files located in `/ase/data/autobackup_dbs`, or you can copy the archived dump files from the `s3` location to a directory on your EC2 instance (for example, `/ase/data/autobackup_dbs`), and use these to restore your database to a previous version. For example, this series restores the `pubs2` database from database and transaction log dumps in `/ase/data/autobackup_dbs` directory:

```
LOAD DATABASE pubs2 FROM
'/ase/data/autobackup_dbs/pubs2.dmpdb.20180729153007.dmp'
go
LOAD TRAN pubs2 FROM
'/ase/data/autobackup_dbs/pubs2.dmptx.20180729161501.dmp'
go
LOAD TRAN pubs2 FROM
'/ase/data/autobackup_dbs/pubs2.dmptx.20180729171501.dmp'
go
LOAD TRAN pubs2 FROM
'/ase/data/autobackup_dbs/pubs2.dmptx.20180729181501.dmp'
go
LOAD TRAN pubs2 FROM
'/ase/data/autobackup_dbs/pubs2.dmptx.20180729191501.dmp'
go
```

Optionally, use the `until_time` parameter to load database dumps that were performed up to a specific time. For example, this loads the database dumps for `pubs2` performed up to 12:00 on July 7, 2018:

```
load database pubs2 with listonly=load_sql ,until_time='2018-07-29 12:00:00'
```

Use the `sp_dump_history list` system procedure to display the history of dumps:

```
sp_dump_history list
go
DUMP
----
 Dump_Type Dbid Database_name Stripes Dump_instant
File                                                    Server_name
Compression_lvl Password Status  Label Dump_date
 --------- ---- ------------- ------- --------------------------
------------------------------------------------------------ -----------
--------------- -------- ------- ----- --------------------------
 DATABASE     1 master              1 Jul 27 2018  6:03:16:913AM /ase/data/
autobackup_dbs/master.dmpdb.20180727060315.dmp *                          0
no       Success *     Jul 27 2018  6:03:17:450AM
 DATABASE     1 master              1 Jul 27 2018  7:23:59:933AM /ase/data/
autobackup_dbs/master.dmpdb.20180727072358.dmp *                          0
no       Success *     Jul 27 2018  7:24:00:323AM
 DATABASE     1 master              1 Jul 27 2018  7:54:29:910AM /ase/data/
autobackup_dbs/master.dmpdb.20180727075428.dmp *                          0
no       Success *     Jul 27 2018  7:54:30:550AM
```

Issue `crontab -l` at the command line prompt as user sybase to see the currently executing backup schedule. For example:

```
crontab -l
## dump database at 03:30 every day
30 03 * * * /ase/config/autobackup.sh database >> /ase/config/dumpdb.log 2>&1
## dump transaction on 15th minute every hour
15 * * * * /ase/config/autobackup.sh transaction >> /ase/config/dumptran.log 2>&1
```

Issue `crontab -e` to edit the `autobackup_cronjob` script to change the backup schedule.

To stop the Auto Backup feature, comment out each line in the `autobackup_cronjob` script by adding a pound sign (`#`) at the beginning of each line. Execute this command after commenting out each line to make the changes take effect:

```
crontab autobackup_cronjob
```

5.4 Connecting with the SAP ASE Cockpit

You can monitor SAP ASE on AWS with the cockpit.

To connect with the SAP ASE cockpit, use the following URL:

```
https://[<instance_EIP | public_IP>]:4283/cockpit
```

Replace `<instance_EIP | public_IP>` with the actual IP address. For example, if your Elastic IP address is 35.167.127.98, the cockpit URL is:

```
https://35.167.127.98:4283/cockpit
```

Use these values to log in to the SAP ASE cockpit:

- *System* – server name. This value is autofilled.
- *User name* – user name "sa" created during installation.
- *Password* – password you specified during installation.

The SAP ASE Cockpit documentation is available at SAP Adaptive Server Enterprise Cockpit.

Error When Displaying Cockpit Error Log

Occasionally, the SAP ASE cockpit issues this error message when you attempt to display the error log:

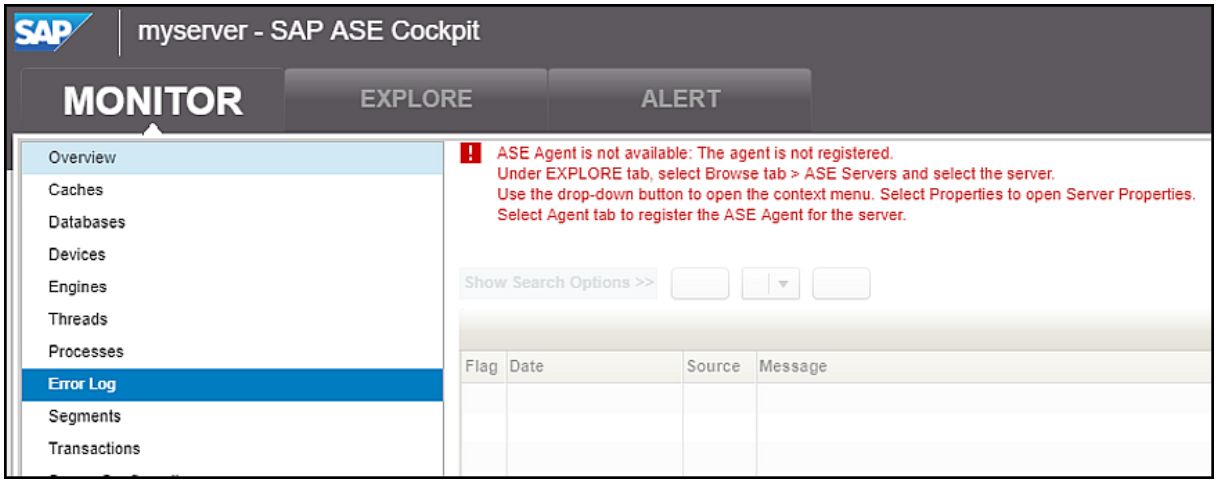

This error is caused by a resource limitation on the smaller-sized AWS EC2 systems (for example, t2.medium).

To resolve the error, preform the following once server is running and all dataserver background processes are complete:

1. Use PuTTY to connect to your EC2 instance.
2. Issue this command as the user `sybase`:

   ```
   java TestRegisterASEAP <password> <IP_address> <server_name>
   ```

   Where:
   - `<password>` – is the password for the system administrator.
   - `<IP_address>` – is the internal IP address of the instance (use ipconfig to view this address).
   - `<server_name>` – is the name of the server, as it appears in the interfaces file.

5.5 Connecting to the EC2 Instance

Connect to the EC2 instance to perform monitoring and administrative tasks.

Use a standard connection tool (for example, PuTTY) to connect from your on-premises machine to the EC2 instance running SAP ASE.

Provide the following information:

- User – ec2-user
- Connection type – SSH
- Host name – the Elastic IP address or public IP address
- Port number – 22
- Key name – path to the private key that you created from the AWS generated public key. See http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html and http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html for more information.

5.6 Connecting Clients and Developing Applications for AWS

SAP ASE for AWS includes the Windows OCS SDK for creating clients, and configuring them to connect.

Downloading the Client SDK from SAP

If you are already an SAP customer, you can download the Open Client SDK.

Prerequisites

Download the SDK from https://support.sap.com/en/my-support/software-downloads.html.

Procedure

1. Go to https://support.sap.com.
2. Log in by selecting the person icon at the top of the page.
3. Select *Download Software*, then select the *Support Packages and Patches* tab.
4. Search for, or navigate to *SAP OPEN SERVER*.
5. Select the appropriate version and platform, then click *Continue*.
6. Select your platform from the drop-down list.
7. Select the software to download.
8. Click *Save*.
9. Browse to the location to which you are downloading the file.
10. Extract the compressed file.
11. Run the `setup.exe` installer (located in `extracted_files\opernserver_<version>\<ebf>_<version>`).
12. Follow the directions on the wizard to install the client.

Downloading the Client SDK from AWS

SAP ASE for AWS includes the Open Client SDK, which you can download to your on-premises server.

Procedure

1. Copy the appropriate version of the SDK to your S3 location from the following locations on the EC2 instance:
   - Linux – /ase/software/SDK/linux/SDKASE160003P_2-21012007.TGZ
   - Windows – /ase/software/SDK/windows/SDKASE160003P_2-21012005.ZIP
2. Use a utility like `ftp` or `WINscp` to download it to your on-premises server.
3. To install the Open Client, see SDK for SAP Adaptive Server Enterprise.

Connecting Clients to SAP ASE on AWS

Connect clients to SAP ASE on AWS the same way that you connect clients to your on-premises servers.

Context

These steps describe how to connect `isql`; however, the same basic steps apply to any client you want to connect.

Procedure

1. Change to the `$SYBASE_OCS/bin` directory.
2. Issue the following command, specifying the Elastic IP address (or host name) and port number:

   ```
   isql -U<login_name> -P<password> -S<IP_address>:<port_number>
   ```

   For example:

   ```
   isql -Usa -Psybase123 -S10.189.27.155:5000
   ```

   You can also use Interactive SQL, a GUI-based `isql` utility, to connect to SAP ASE running on AWS.

   If you require an encrypted connection, use the `isql -X` parameter.

   See the Utility Guide for more information about Interactive SQL and `isql -X`.

5.7 Monitoring SAP ASE on AWS

There are a number of ways you can monitor SAP ASE on AWS.

- Configure AWS CloudWatch to monitor and generate alerts for SAP ASE events. See https://aws.amazon.com/cloudwatch/ .
- Use the `sp_sysmon` system procedure to provide in-depth information about monitoring SAP ASE. See Performance and Tuning Series: Monitoring with sp_sysmon.
- Use the SAP ASE cockpit to manage and monitor SAP ASE. See SAP Adaptive Server Enterprise Cockpit.

5.8 Performance and Tuning

SAP ASE for AWS is configured for OLTP performance using the TPCC database. However, you can configure the server for the needs of your site.

Use the `sp_sysmon` system procedure to view the current performance characteristics of SAP ASE. See Performance and Tuning Series: Monitoring with sp_sysmon.

See the Performance and Tuning Series for general performance information and recommendations.

5.9 Lifecycle Management

For products within the Amazon Machine Image (AMI), which contains the information for launching an instance, the definition for the end of mainstream maintenance (EoMM) is the same as what is published for each version of the product in the Product Availability Matrix (see **https://apps.support.sap.com/sap/support/pam** .

SAP plans to release a new AMI following each release of an SAP ASE SP and PL, unless otherwise specified. The AMI release for each SP or PL release will be made available around the same time as the corresponding SP or PL releases are made available for on-premise versions of SAP ASE Enterprise Edition.

SAP intends to have two categories of AMIs:

- Innovation – an SP release for which subsequent PL releases may contain enhancements.
- Safe harbor – an SP release for which subsequent PL releases contain only bug fixes, but no enhancements.

For either category, when SAP makes a new AMI available with a new PL, the previous AMI with an older PL becomes unavailable for new users to launch. These AMIs, with the older PL, become archived. However, users who launched instances on these AMIs can continue to use them.

For more information about the product life cycle, see https://i7p.wdf.sap.corp/sap(bD1lbiZjPTAwMQ==)/bc/bsp/sno/ui_entry/entry.htm?param=69765F6D6F64653D3030312669765F7361706E6F7465735F6E756D6265723D3235333133323626 .

5.10 SAP ASE on AWS Internals

The SAP ASE subscription on AWS is designed to separate the SAP ASE software from the data, allowing for modular operation and easy upgrades.

Storage Structure for SAP ASE Subscription on AWS

The structure of SAP ASE running on AWS:

Name	Size, in Gigabytes	Location	Purpose
Host root volume	20	Mounted on `/`	Includes the operating system, users, and SAP ASE deployment structures.
SAP ASE software	10	Mounted on `/opt/sap`	Contains the SAP ASE release.
SAP ASE system databases	10	Mounted on `/ase/data/system_dbs`	Includes the master device.
User database data	User-specified	Mounted on `/ase/data/userdata_dbs`	User database data. Size is specified during deployment.
User database log	User-specified	Mounted on `/ase/data/userlog_dbs`	User database log. Size is specified during deployment.

SAP ASE Operation Files

The SAP ASE software is located under `/opt/sap`, and the SAP ASE installation and operation files are located under `/ase/config`:

File name and path	Function
`/ase/config/interfaces`	Connection and port information for the SAP ASE, Backup Server, and XP Server. **Note** `/ase/config/interfaces` is symbolically linked to `/opt/sap/interfaces`.

File name and path	Function
/ase/config/ASE-16_0/<server_name>.cfg	SAP ASE internal configuration. **Note** <server_name> is user-specified during deployment.
/ase/config/ASE-16_0/<server_name>.krg	SAP ASE shared memory information (internal to SAP ASE).
/ase/config/ASE-16_0/install/RUN_<server_name>	SAP ASE runserver startup script.
/ase/config/ASE-16_0/install/<server_name>.log	SAP ASE error log file.
/ase/config/ASE-16_0/install/RUN_<server_name>_BS	Backup Server runserver startup script.
/ase/config/ASE-16_0/install/<server_name>_BS.log	Backup Server error log file.

How SAP ASE is Deployed on AWS

At deployment, the master and user-specified databases are created by running the following script as the `sybase` user:

```
/home/sybase/aws_srvbuild
```

The default configuration for a newly created server is as follows:

- Page size – 16 K
- Default character set – utf8
- Default sort order – binary

Starting and Stopping Services on AWS

When an AWS SAP ASE instance is created, the `systemd` unit file for managing the ASE service file is deployed by running:

```
systemctl enable /etc/systemd/system/ase.service
```

The ASE service is used to start or stop SAP ASE when you start or stop an instance. By default, the ASE service is enabled so that SAP ASE automatically starts when the AWS instance starts. The `asesvcs` script performs the starting and stopping. For example:

```
/ase/config/ASE-16_0/install/asesvcs start
```

This command subsequently calls the following commands to start SAP ASE and Backup Server:

```
/opt/sap/ASE-16_0/bin/startserver -f /ase/config/ASE-16_0/install/RUN_myserver
/opt/sap/ASE-16_0/bin/startserver -f /ase/config/ASE-16_0/install/RUN_myserver_BS
```

For information about any issues that occur during startup, see `/var/log/messages`.

5.10.1 Stopping the Host Instance

You may occasionally need to shut down the instance during times of low database demand. Use the AWS console to do this.

Procedure

1. It is important that you perform a polite shutdown of SAP ASE before shutting down the instance.
   a. Use `isql` to log in to an SAP ASE account with system administrator privileges:

      ```
      isql -Usa -P<password> -S<server_name>
      ```

   b. Shut down the server:

      ```
      shutdown
      ```

      The default for `shutdown` uses the `with wait` option, which allows SAP ASE to finish executing SQL statements or procedures, perform a `checkpoint` in each database, disable new logins, and perform other shutdown tasks.

      Issuing the `shutdown` command prints a message similar to:

      ```
      Server SHUTDOWN by request.The SQL Server is terminating this process.
      CT-LIBRARY error:
      ```

      This is normal behavior. If the message indicates that SAP ASE is waiting for processes to complete, and you must stop SAP ASE immediately, you can use `shutdown with nowait` which neither waits for currently executing statements to finish, nor performs checkpoints in every database.

      > **i Note**
      >
      > Use `shutdown with nowait` only when necessary.

2. Select the EC2 console.
3. Select *EC2 Dashboard* > *Instances*, then select your instance.

4. Select *Actions* > *Instance State*, then select *Stop*.

5.10.2 Restarting a Stopped Host Instance

Use the AWS console to restart an instance you have stopped.

Procedure

1. Select the EC2 console.
2. Select *EC2 Dashboard* > *Instances*, then select your instance.
3. Select *Actions* > *Instance State*, then select *Stop*.

   The instance automatically starts SAP ASE and Backup Server with its services.

5.10.3 Stopping SAP ASE, Backup Server, and XP Server

You can stop and start servers running on the instance from the command line.

Procedure

1. Log into SAP ASE with `isql`:

   ```
   $SYBASE/$SYBASE_OCS/bin/isql -Usa -P<password> -S<server_name>
   ```

2. Issue `sp_helpserver` to view the running servers.
3. Use the `shutdown` command to shut down each server, except SAP ASE. This example shuts down the Backup Server:

   ```
   shutdown SYB_BACKUP
   ```

4. Issue the `shutdown` command to shut down SAP ASE.

Next Steps

Use the runserver start-up scripts located in `/ase/config/ASE-16_0/install` to restart a stopped SAP ASE and Backup Server.

> **i Note**
>
> XP Server starts automatically when you first use it, and does not require an explicit start-up script.

See the SAP ASE Cockpit documentation for information about starting and stopping servers from the cockpit: SAP Adaptive Server Enterprise Cockpit.

5.10.4 Enabling XP Server on AWS

Perform configuration steps to enable XP Server for SAP ASE on AWS.

Procedure

Validate the XP Server addition. Connect to SAP ASE with `isql` and issue this command:

```
xp_cmdshell "ls"
```

This command may take longer to run because XP Server starts only when it is first called.

> **i Note**
>
> By default, the context for `xp_cmdshell` is set to 1, which means commands are executed on behalf of the user who is logged into SAP ASE (in this case, sa), and you may see error messages similar to the following when you issue `xp_cmdshell`:
>
> ```
> User access denied. Failed to change the user context.
> ```
>
> Because the sa user does not exist in the host, you can change the context to 0 so the `xp_cmdshell` commands are run under the XP Server account, which is the `sybase` user. To change the context to 0, enter the following command:
>
> ```
> sp_configure 'xp_cmdshell', 0
> ```
>
> For more information about the XP Server context see XP Server.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap

THE BEST RUN SAP