



PUBLIC

Document Version: 1.0 – 2026-03-17

SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Migration Guide

Content

1	SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Migration Guide.	4
2	Prerequisites.	7
3	Discover.	8
3.1	Retrieving Server Information.	8
3.2	Retrieving Database Information.	12
4	Pre-Migration.	15
4.1	Determining Resource Size.	16
4.2	Checking Database Consistency.	18
4.3	Backing Up the Database and Configuration Files.	19
4.4	Estimating Migration Duration.	19
4.5	Checking for Incompatible Features.	20
	Unsupported Features.	21
4.6	Checking on Data Cache Binding Information.	23
4.7	Managing Login Accounts and Roles.	25
	Logins with System-Reserved Login Names or Server User IDs.	26
	Obtaining Login Accounts.	47
	Obtaining User-Defined Roles.	51
4.8	Encrypting Your Database.	55
4.9	Dumping the Database and Transactions.	57
5	Migration.	60
5.1	Provisioning an SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Instance.	60
5.2	Restoring Non-Default Configurations.	61
5.3	Creating Data Caches.	61
5.4	Migrating the Database Encryption Key.	63
	Migrating the DEK with transfer encryption key.	63
5.5	Creating a Database.	64
5.6	Synchronizing Login Accounts and Roles.	65
	Creating Login Accounts.	66
	Creating User-Defined Roles.	67
5.7	Uploading Dump Files	69
5.8	Loading the Database and Transactions.	70
5.9	Bringing the Database Online.	71
5.10	Restoring Cache Bindings.	72
6	Post-Migration.	73

7	Troubleshooting.....	76
----------	-----------------------------	-----------

1 SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Migration Guide

To provide the combination of on-premises stability and cloud-based innovation, SAP has developed a migration solution that allows you to migrate your existing databases to the cloud. This guide explains how to migrate the databases from on-premises SAP ASE to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

You need to make a migration plan before starting the migration. You can break down the migration plan into four phases:

1. Discover
2. Pre-Migration
3. Migration
4. Post-Migration

The following are typical activities that happen in each migration phase.

Migration Phase	Activities
Discover	• Check whether the minimal migration requirements are met. • If the server is not running on Linux, migrate it to Linux first. • If the SAP ASE version is earlier than 16.0 SP04, upgrade it to 16.0 SP04 or higher. • Retrieve server and database information Retrieve information about your on-premises server and databases, including total memory, master database size, device names and sizes, page size, character set, sort order, and other relevant configuration details.

Migration Phase	Activities
Pre-Migration	1. Determine resource size and choose proper specification for the ASE instance in the cloud. It is recommended to choose ASE instance specification in the cloud equivalent to the on-premises server. You can choose a more powerful configuration for better performance. 2. Check database consistency with the dbcc command. Checking Database Consistency [page 18] 3. Back up the configuration files. 4. Estimate migration duration with the sp_migration_estimate stored procedure. 5. Check whether the enabled features in the on-premises SAP ASE server and database are available in the cloud edition. If there are any features that are unavailable in the cloud edition, disable the features in the on-premises SAP ASE server and database before taking the database dump. Checking for Incompatible Features [page 20] 6. Check information on data cache bindings. 7. Obtain login accounts and roles. If there is any conflict between the user-defined logins and system reserved logins of the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud, resolve the conflict. Logins with System-Reserved Login Names or Server User IDs [page 26] 8. Check whether the on-premises database is encrypted. If the database is not encrypted, encrypt the database. 9. Dump the database and transactions.

Migration Phase	Activities
Migration	1. Provision an SAP ASE instance in the cloud with the chosen specification. Make sure the page size, character set, and sort order of the ASE instance are the same as that of the on-premises server. 2. If required, restore non-default configurations by using sp_configure. 3. Create the same data caches as the on-premises database. 4. Migrate the database encryption key (DEK) and create a database: 1. Upload the DEK and assign it to the SAP ASE instance using the SAP ASE Cloud Portal. Migrating the DEK with transfer encryption key [page 63] 2. Create the database encrypted with the migrated DEK in the cloud. 3. Configure the database options in the cloud by using the sp_dboption procedure. You should perform this step on all hosts in a HADR cluster. 5. If required, create the same logins and roles in the SAP ASE instance as in the on-premises server. 6. Upload dump files of the database to the instance specific backup bucket in IBM Cloud Object Store. 7. Restore the database. 8. Bring the database online. 9. Grant privileges to user-defined roles. 10. Bind the data caches and restart the SAP ASE instance through a support request or self-service because the bindings and cache configurations need to be reflected in the standby and DR nodes.
Post-Migration	Perform post-migration activities: 1. Validate the following: • non-default database and server configuration • database options • data cache bindings • login accounts and roles 2. Run a consistency check with the dbcc command. 3. Run a functional test 4. Tune and validate the performance of the ASE instance.

2 Prerequisites

Migration of SAP ASE databases with versions lower than 16.0 SP04 is not supported. You should migrate your SAP ASE to at least version 16.0 SP04 before migrating your databases to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

Supported Operating Systems

You can only load a database dump file from Linux x86-64 to IBM Cloud.

This is because the cloud edition supports only the Linux platform, and database encryption is mandatory for the SAP ASE service. However, dumping and loading encrypted databases across platforms is not supported.

If your SAP ASE server is running on platforms other than Linux x86-64, migrate it to a Linux x86-64 on-premises server before migrating to IBM Cloud. Note that the source database should not be encrypted. See [Dumping and Loading Databases Across Platforms](#) in the *SAP Adaptive Server Enterprise System Administration Guide: Volume 2* for more information.

Supported SAP ASE Versions

If your SAP ASE is version 16.0 SP04 PL04 or lower, upgrade to SP04 PL05 or later to install new stored procedures available in the `installmaster` script (`sp_serverinfo_report`, `sp_dbinfo_report`, `sp_featurecompat`, and `sp_migration_estimate`) into the `sybsystemprocs` database. These procedures are used in migration tasks.

3 Discover

Before starting your migration, you should retrieve and review information about your on-premises SAP ASE server and databases to be migrated. This information can help you to prepare a proper migration plan.

[Retrieving Server Information \[page 8\]](#)

Collecting the on-premises SAP ASE server information can help you choose compute, memory, and storage for your instance, estimate end-to-end migration time, and create a migration schedule.

[Retrieving Database Information \[page 12\]](#)

Collecting the on-premises SAP ASE database information can help you choose the database size in the cloud and find out existing features, logins, and roles in the current database.

3.1 Retrieving Server Information

Collecting the on-premises SAP ASE server information can help you choose compute, memory, and storage for your instance, estimate end-to-end migration time, and create a migration schedule.

Use the `sp_serverinfo_report` to obtain your server level information, including:

- sizing-based resource configurations
- set of attributes at server installation
- server level non-default configurations
After provisioning the SAP ASE instance on IBM Cloud, you need to check these non-default configurations one by one and restore them if needed.
- list of databases
- list of login accounts
- list of roles
- security-related checks
- SAP Adaptive Server Enterprise, cloud edition by IBM Cloud compatibility validation
- SAP ASE server date and time validation.

The following is an example of what a server information report looks like:

```
sp_serverinfo_report
go
=====
Section 1: Server level configurations.
=====
[Basic configurations]
Server Version: Adaptive Server Enterprise/16.0 SP04 PL05/EBF 30841 SMP/Auto-82
                  8808-1/x86_64/SLES 12.5/ase160sp04pl05x/0/64-bit/OPT/Wed Nov 29
                  14:11:00 2023
Page Size: 16k
Default character set id: 190 (utf8)
Default sortorder id: 25 (binary)
Lock scheme: datarows
Kernel mode: threaded
```

```

[Processors]
Max online engines: 4
Currently online engines: 4
[Physical Memory]
Max memory: 1048576 (Unit: 2k)
Statement cache size: 0 (Unit: 2k)
Procedure cache size: 15000 (Unit: 2k)
User log cache size: 65536 (Unit: bytes)
Session tempdb log cache size: 16384 (Unit: bytes)
[User Defined Caches]
User defined caches lists as below.
  name                      size
  -----
  default data cache        8.00 MB
[Storage Disk]
User created devices lists as below.
  vdevno  dbid  segmap  device_name          size
  -----
    1    31514      7 sybspdb0ccus0_19299  0.215 GB
    2         2      7 tempdb0             0.145 GB
    3    31513      7 sybsystemdb0         0.293 GB
    4         4      3 ldev1              0.047 GB
    5         4      4 ldev2              0.098 GB
  device_name          size
  -----
  sybsystemdb0         0.293 GB
  tempdb0              0.145 GB
  sybspdb0ccus0_19299  0.215 GB
  ldev2                0.098 GB
  ldev1                0.047 GB
[User databases]
There is (are) 1 user database(s).
  name          dbid      suid      encrkeyid      status5
  -----
  tdb1          4          6 NULL          32
[Login accounts]
There are 16 login accounts.
  suid      name
  -----
    -9 testlogin1
     1 sa
     2 probe
     3 user0
     4 user1
     5 user2
     6 user3
     7 user4
     8 user5
     9 user6
    10 user7
    11 user8
    12 user9
    13 user10
    14 user11
    15 saptu
[Roles]
There are 22 roles.
  srid      name
  -----
     0 sa_role
     1 sso_role
     2 oper_role
     3 sybase_ts_role
     4 navigator_role
     5 replication_role
     7 dtm_tm_role
     8 ha_role
    10 mon_role

```

```

11 js_admin_role
12 messaging_role
13 js_client_role
14 js_user_role
15 webservices_role
16 keycustodian_role
17 hadr_admin_role_gp
18 replication_maint_role_gp
19 saptu_role
20 sapsupport_role
21 cust_sa_role
22 cust_sso_role
32 sa_serverprivs_role

```

Section 2: Non-default configurations.

There are 27 non-default configurations list as below.

ParamName	CfgValue	RunValue
cpu grace time	1500	1500
default character set id	190	190
default sortorder id	25	25
enable concurrent dump tran	1	1
enable console logging	1	1
enable functionality group	1	1
enable inline default sharing	1	1
enable large pool for load	1	1
enable mem scale	1	1
enable pci	1	1
enable permissive unicode	1	1
enable select into in tran	1	1
enable semantic partitioning	1	1
enable utility lvl 0 scan wait	1	1
exception on rounding overflow	1	1
expand numeric truncated scale	1	1
lock scheme	datarows	datarows
max memory	1048576	1048576
max online engines	4	4
max pci slots	4	4
number of open databases	21	21
procedure cache size	15000	15000
quoted identifier enhancements	1	1
restricted parameter markers	1	1
select for update	1	1
streamlined dynamic SQL	1	1
suppress js max task message	1	1

Section 3: Security related check.

3.1 Check if the reserved login accounts and suids are existed

Error: There is (are) 2 login account(s) with reserved suids

-4/-5/-6/-7/-8/-9/-10/-11/-12/-13 or names

saptu/sapsupport/aseadmin/sappr/sapping/sapexporter/sapaudlog/sapamcadmin/sapdradmin/sapdrmaint on ASE Cloudservice.

suid name

-9 testlogin1

15 saptu

3.2 Check if the reserved roles are existed

3.3 Check granular permission

enable granular permissions: 0

3.4 Check if enabled to store a key externally (HSM)

external keystore: NULL

```

Section 4: ASE Cloudservice compatibility validation.
=====
Check the deprecated features of ASE service.
==== Following is the list of features that are not supported for ====
==== SAP Adaptive Server Enterprise with cloudservice version.      ====
Server level features :
- Simplified native access plan
- SAP ASE as a companion server in HA subsystem
- SAP ASE 'process' kernel mode
- Sending audit records to syslog
- disk mirroring
- XP server
Usage:
sp_featurecompat 'info' : Info of all the features deprecated.
sp_featurecompat 'list' : Info of all the features deprecated for use on the
server.
sp_featurecompat 'info', [feature] : Disable instructions for related feature.
sp_featurecompat ['help'] : Usage information for stored procedure.
Following features are in use :
Msg 17454, Level 16, State 1:
Procedure 'sp_featurecompat', Line 174:
Error : SAP Adaptive Server Enterprise is currently configured to use the
'XP_server' feature. This feature is deprecated/restricted on ASE Cloudservice.
=====
Section 5: ASE server date and time validation.
=====
The current system date and time: Nov 30 2023  6:33AM.
Date and time in Universal Coordinated Time: Nov 30 2023 11:33AM.
Warning: ASE inherits timezone from hosting OS, users are expected to convert
the data which is timezone relevant if the timezone of target ASE is different
from the one of source ASE.
Collecting Adaptive server information completed with 2 error(s).
(return status = 0)

```

Provisioning Related Information

Make a note of the following information you have retrieved by using the `sp_serverinfo_report` procedure, which are highlighted in bold in the previous output.

- max online engines
- max memory
- page size
- character set
- sort order

You need to provide this information when you provision the SAP ASE instance. You can keep the same configurations for the instance or scale up compute capacity, memory, or storage for better performance.

Related Information

[Prerequisites \[page 7\]](#)

3.2 Retrieving Database Information

Collecting the on-premises SAP ASE database information can help you choose the database size in the cloud and find out existing features, logins, and roles in the current database.

Run the `sp_dbinfo_report` stored procedure to get database level information for the on-premises SAP ASE. The database information includes:

- properties specified by `create/alter database`
- properties setup by `sp_dboption`
- properties stored in the `sysattributes` catalog
- device fragments information
- cache and binding information
- encryption status and DEK information
- user-related information
- information on user-defined roles
- cross-database reference information
- SAP Adaptive Server Enterprise, cloud edition by IBM Cloud service compatibility validation
- validation of date and time related columns.

You should run `sp_dbinfo_report` for all databases you need to migrate. For example, to check database information on a database called `tdb1`, you can run:

```
sp_dbinfo_report tdb1
go
=====
Section 1: Database properties.
=====
name      db_size      owner      dbid      created
-----
tdb1      148.0 MB user3      4 Nov 30, 2023
(1 row affected)
database_property      status/value
-----
abort tran on log full enabled
durability             full
inrowlen              NULL
lobcomp_lvl           0
=====
Section 2: The related sysattributes data.
=====
There is no sysattributes data.
=====
Section 3: The device fragments information.
=====
device_fragments size      usage      created      free_kbytes
-----
ldev1            48.0 MB data only Nov 30 2023  5:53AM      35200
ldev2            100.0 MB log only  Nov 30 2023  5:53AM      not applicable
-----
-----
-----
-----
-----
log only free kbytes = 101984
Summary of device fragments:
```

```

Total size of all the data only devices:      0.047 GB.
Total size of all the log only devices:      0.098 GB.
Total size of all the mixed devices:         0.000 GB.
=====
Section 4: Cache and binding information.
=====
There is no data cache belongs to database tdb1.
=====
Section 5: Encryption status and DEK information.
=====
Database encryption status: 0.
Error: Database tdb1 is not encrypted.
=====
Section 6: Logins related information.
=====
Error: There is (are) 2 login account(s) with reserved suids
-4/-5/-6/-7/-8/-9/-10/-11/-12/-13 or names
saptu/sapsupport/aseadmin/sappr/sapping/sapexporter/sapaudlog/sapamcadadmin/
sapdradmin/sapdrmaint on ASE Cloudservice.
suid  uid  gid  name
-----
    -9   14    0 testlogin1
    15   15    0 saptu
Total number of login accounts which are explicitly added by sp_adduser in
database tdb1: 13.
suid  name      fullname      procid      lpid
-----
    3 user0      NULL          NULL        NULL
    4 user1      NULL          NULL        NULL
    5 user2      NULL          NULL        NULL
    6 user3      NULL          NULL        NULL
    7 user4      NULL          NULL        NULL
    8 user5      NULL          NULL        NULL
    9 user6      NULL          NULL        NULL
   -9 testlogin1 NULL          NULL        NULL
   10 user7      NULL          NULL        NULL
   11 user8      NULL          NULL        NULL
   12 user9      NULL          NULL        NULL
   13 user10     NULL          NULL        NULL
   15 saptu      NULL          NULL        NULL
=====
Section 7: Roles related information.
=====
Total number of roles which are granted database-wide privileges in
database tdb1: 0.
=====
Section 8: Cross-database reference information.
=====
There is no cross-database referential integrity constraint.
=====
Section 9: Date and time related columns validation.
=====
The current system date and time: Nov 30 2023  6:35AM.
Date and time in Universal Coordinated Time: Nov 30 2023 11:35AM.
Total number of columns which are date and time related in database tdb1: 0.
Collecting database information completed with 2 error(s).
(return status = 0)

```

Provisioning Related Information

Make a note of the following device sizes, which are highlighted in the previous output:

- total size of all the data-only devices

- total size of all the log-only devices
- total size of all the mixed devices

For the mixed device, you can use the `sp_spaceused` procedure to assess the log size and the data size used in the device. You should sum up the data size of all databases to determine the required data size for your ASE instance and sum up the log size of all databases to determine the required log size. You will later choose the storage size of your ASE instance based on this information.

To retrieve `create database` command information, you can specify the `listonly = create_sql` option when loading the database.

4 Pre-Migration

This section describes tasks you should perform before migrating your data to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

You should have the following files ready by the end of pre-migration phase:

- login creation script
- role creation script
- cache creation script
- command to restore the database level cache binding
- DEK creation script
- output of the `load database . . . listonly=create_sql` command
- output of the `load database . . . listonly=load_sql` command
- dump files

1. [Determining Resource Size \[page 16\]](#)

Choose a proper resource size for your SAP ASE instance in the cloud based on the on-premises server and database information.

2. [Checking Database Consistency \[page 18\]](#)

It is recommended to run DBCC commands to check the consistency of database before taking the database dump.

3. [Backing Up the Database and Configuration Files \[page 19\]](#)

Some optional operations are recommended before making any configuration changes to the on-premises database. You can back up the database and configuration file to avoid any potential risk.

4. [Estimating Migration Duration \[page 19\]](#)

You can use the `sp_migration_estimate` stored procedure to estimate the end-to-end migration duration, which can help you determine the maintenance time of the database.

5. [Checking for Incompatible Features \[page 20\]](#)

You should disable all deprecated and restricted features that are being used by the database before migration. Use the `sp_featurecompat` procedure to find these features and instructions about how to disable them.

6. [Checking on Data Cache Binding Information \[page 23\]](#)

Binding information for databases is stored in `master . . sysattributes`. Binding information for objects is stored in the `sysattributes` table in the database itself. Loading databases will not restore the data cache binding information for databases.

7. [Managing Login Accounts and Roles \[page 25\]](#)

You need to ensure the server user IDs for all users and user-defined roles are synchronized between the on-premises SAP ASE database and the database that is to be migrated SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

8. [Encrypting Your Database \[page 55\]](#)

The dump-and-load migration method requires the source database (on-premises) and the target database (in the cloud) to have the same database encryption key (DEK). Since databases created

in SAP Adaptive Server Enterprise, cloud edition by IBM Cloud are encrypted by default, you must encrypt your on-premises database before dumping for migration.

9. [Dumping the Database and Transactions \[page 57\]](#)

Use the `dump database` command to make a copy of the entire database, which is to be loaded in the cloud, and then use the `dump transaction` command to make a copy of the transactions since the database dump.

4.1 Determining Resource Size

Choose a proper resource size for your SAP ASE instance in the cloud based on the on-premises server and database information.

SAP ASE cloud edition provides the following T-shirt sizes based on different configurations of compute, memory, and storage:

T-Shirt Size	Compute in GB	Memory in GB	Storage in GB
XS (available only in non-HA service)	2	8	256
S	4	32	512
M	8	64	1,024
L	16	128	1,536
XL	32	256	2,048
XXL	64	512	3,072
XXXL	128	1,024	5,120

You can add additional storage on top of your chosen size. But you cannot increase CPU or memory other than switching to the nearest T-shirt size.

Compute and Memory Type

You should choose a compute and memory type that is at least equal to your current configuration. For example, if your server information report looks like the following:

```
[Processors]
Max online engines: 6
Currently online engines: 1
[Physical Memory]
Max memory: 47 (Unit: G)
```

You can:

- **Compute**

Choose a compute no less than `Max online engines` whose value is 6 in this example.

- **Memory Type**

The max memory is 47 GB. Since the ASE service allocates 90% of memory as the max memory, you'll need a memory of 53 GB ($47/0.9$) at least for this example. Therefore, you should choose the memory type whose memory size is no less than 53 GB.

In this case, you can choose size M for your SAP ASE instance.

Storage

In ASE, a device is used as fragments. Device fragments have the following usages: data only, log only, and data and log. For the data and log mixed device, you can use the `sp_spaceused` procedure to assess the log size and the data size used in the mixed device. You should sum up the data size of all databases to determine the required data size for your ASE instance and sum up the log size of all databases to determine the required log size.

When provisioning, you can choose the total storage of the ASE instance and use the recommended data size and log size. If the recommended size does not work for you, you can also choose a data size and a log size separately. Either way, the total storage, data storage, or log storage should be at least equal to that of your on-premises SAP ASE server. For example, if your on-premises server has 500 GB data and 60 GB log, you can choose a 512 GB data size and a 64 GB log size for your ASE instance in the cloud.

You can use `sp_dbinfo_report <database name>` to get the device fragments information of your databases to be migrated. For example, you have two databases `tdb1` and `tdb2` on your server which you're planning to migrate to the cloud.

`sp_dbinfo_report tdb1` gets the following output in the report:

```
=====
Section 3: The device fragments information.
=====
Total size of all the data only devices:      20.0 GB
Total size of all the log only devices:       10.0 GB
Total size of all the mixed devices:          0.0 GB
```

`sp_dbinfo_report tdb2` gets the following output in the report:

```
=====
Section 3: The device fragments information.
=====
Total size of all the data only devices:      30.0 GB
Total size of all the log only devices:       15.0 GB
Total size of all the mixed devices:          0.0 GB
```

Now you can have the following sum-up:

```
Total size of all the data only devices:      50.0 GB
Total size of all the log only devices:       25.0 GB
Total size of all the mixed devices:          0.0 GB
```

So in this case, you can choose either of the following:

- a total storage that is no less than 75 GB
- a data size that is no less than 50 GB and a log size that is no less than 25 GB

Parent topic: [Pre-Migration \[page 15\]](#)

Next: [Checking Database Consistency \[page 18\]](#)

4.2 Checking Database Consistency

It is recommended to run DBCC commands to check the consistency of database before taking the database dump.

- Run `dbcc checkstorage (<dbname>)` to check the following items:

- Allocation of text valued columns
- Page allocation and consistency
- OAM page entries
- An OAM page for each partition exists
- Pointer consistency
- Text-valued columns and text-column chains

See *Configuring SAP ASE to Use dbcc checkstorage and checkverify* in the *SAP Adaptive Server Enterprise System Administration Guide: Volume 2* for more information.

- If you have installed the `dbccdb` database, run `sp_dbcc_summaryreport "<dbname>"` to generate a summary report on the specified database.
- Run `dbcc checkdb (<dbname>)` to check:
 - Index and data pages are linked correctly
 - Indexes are sorted properly
 - Pointers are consistent
Data rows on each page have entries in the row-offset table; these entries match the locations for the data rows on the page
 - Data rows on each page have entries in the row-offset table in the page that match their respective locations on the page
 - Partition statistics for partitioned tables are correct

Fix the reported problems before starting the database dump. The output can also be used to compare and verify on-cloud database.

See *dbcc checkstorage* and *dbcc checkdb* in the *SAP Adaptive Server Enterprise System Administration Guide: Volume 2* for details.

Parent topic: [Pre-Migration \[page 15\]](#)

Previous: [Determining Resource Size \[page 16\]](#)

Next: [Backing Up the Database and Configuration Files \[page 19\]](#)

4.3 Backing Up the Database and Configuration Files

Some optional operations are recommended before making any configuration changes to the on-premises database. You can back up the database and configuration file to avoid any potential risk.

Backing Up the Database

Use the `dump database` command to make a backup copy of the entire database.

```
dump database <db_name> to <file_name>
```

See `dump database` in the *SAP Adaptive Server Enterprise Reference Manual: Commands* for details on parameters you can use in this command.

Backing Up the Configuration File

Use `sp_configure` to generate a new copy of all the values in `sysconfigures`.

```
sp_configure "configuration file", 0, "restore", "<file_name>"
```

Parent topic: [Pre-Migration \[page 15\]](#)

Previous: [Checking Database Consistency \[page 18\]](#)

Next: [Estimating Migration Duration \[page 19\]](#)

4.4 Estimating Migration Duration

You can use the `sp_migration_estimate` stored procedure to estimate the end-to-end migration duration, which can help you determine the maintenance time of the database.

The syntax is as follows:

```
sp_migration_estimate <db_name>
```

The output of `sp_migration_estimate` looks like the following:

```
sp_migration_estimate testdb
go
          ASE Migration Estimation
=====
On-prem duration:
```

```

Encrypt database:          1.00 hour
Dump database:            0.10 hour
Upload database dump:      4.00 hour
-----
On-prem total:            5.10 hour
On-cloud duration:
Provision ASE:             0.50 hour
Download database dump:    0.50 hour
Load database:            1.00 hour
-----
On-cloud total:           2.00 hour
Total duration:           7.10 hour

```

The output consists of two parts, the on-premises duration part and the cloud duration part. All the above estimation of time consumption for the on-premises part is based on the following hardware environment. You can adjust the estimation according to your actual hardware environment.

- Upload speed: 60 Mbit/s
- CPU: Intel(R) Xeon(R) CPU E7- 4870 @ 2.40GHz
- Disk: HDD

Parent topic: [Pre-Migration \[page 15\]](#)

Previous: [Backing Up the Database and Configuration Files \[page 19\]](#)

Next task: [Checking for Incompatible Features \[page 20\]](#)

4.5 Checking for Incompatible Features

You should disable all deprecated and restricted features that are being used by the database before migration. Use the `sp_featurecompat` procedure to find these features and instructions about how to disable them.

Context

Run `sp_featurecompat` without parameters to learn how to use this procedure:

```

sp_featurecompat
go
Usage:
sp_featurecompat 'info' : Info of all the features deprecated.
sp_featurecompat 'list' : Info of all the features deprecated for use on the
server.
sp_featurecompat 'info', [feature] : Disable instructions for related feature.
sp_featurecompat ['help'] : Usage information for stored procedure.
(return status = 0)

```

For an overview of deprecated and restricted features, run `sp_featurecompat 'info'` to display all features that are not supported in the cloud edition. For more information, see [Unsupported Features \[page 21\]](#).

Procedure

1. Run `sp_featurecompat 'list'` to find out deprecated features that are in use at the server level. For example:

```
1> sp_featurecompat 'list'
2> go
Following features are in use :
Msg 17454, Level 16, State 1:
Procedure 'sp_featurecompat', Line 139:
Error : SAP Adaptive Server Enterprise is currently configured to use the
'disable disk mirroring' feature. This feature is deprecated/restricted on SAP
Adaptive Server Enterprise, cloud edition by IBM Cloud.
Msg 17454, Level 16, State 1:
Procedure 'sp_featurecompat', Line 156:
Error : SAP Adaptive Server Enterprise is currently configured to use the
'audit
trail type' feature. This feature is deprecated/restricted on SAP
Adaptive Server Enterprise, cloud edition by IBM Cloud.
(return status = 1)
```

2. Run `sp_featurecompat 'info'` to learn how to disable a feature:

```
sp_featurecompat 'info', '<deprecated_feature_name>'
```

For example, if you want to know how to disable the `audit trail type` feature, run:

```
sp_featurecompat 'info', 'audit trail type'
go
Disable the feature as follows:
sp_configure 'audit trail type', 0, 'table'
(return status = 0)
```

Follow all instructions given in the output to disable the feature.

3. Run `sp_featurecompat 'list'` to see if all deprecated and restricted features are disabled. If not, repeat the previous steps to disable them.

Task overview: [Pre-Migration \[page 15\]](#)

Previous: [Estimating Migration Duration \[page 19\]](#)

Next task: [Checking on Data Cache Binding Information \[page 23\]](#)

4.5.1 Unsupported Features

This topic lists features that you can use in the on-premises version of SAP ASE but are not available in SAP Adaptive Server Enterprise, cloud edition by IBM Cloud. These features are in addition to those already documented as unavailable in SAP Note [2922454 - SAP Adaptive Server Enterprise \(SAP ASE\) on Cloud Platforms](#).

Feature	Unsupported Components/Functions/Utilities
SAP ASE High Availability (HA) – External Vendor	• <code>installhasvss</code> • <code>sybha</code>
Process mode kernel	<code>sp_configure "kernel mode", process</code>
XP Server	• XP Server utility • Extended stored procedures • Start XP Server • <code>xp_cmdshell</code> • <code>xp_logevent</code> • <code>xp_enumgroups</code> • <code>create procedure</code> with the <code>as external name</code> parameter • <code>create xpserver</code> • <code>drop xpserver</code> • <code>set xpserver port</code> • <code>show xpserver</code>
Disk Mirroring	Disk mirror commands: • <code>disk mirror</code> • <code>disk unmirror</code> • <code>disk remirror</code> • <code>waitfor mirrorexit</code>
Simplified Native Access Plan (SNAP)	<code>sp_configure 'simplified native access plan'</code>
Audit Records	Sending audit records to syslog using <code>sp_configure 'audit trail type'</code> .

4.6 Checking on Data Cache Binding Information

Binding information for databases is stored in `master..sysattributes`. Binding information for objects is stored in the `sysattributes` table in the database itself. Loading databases will not restore the data cache binding information for databases.

Context

When you bring a database online by using `online database` after a load, SAP ASE verifies all cache bindings for the database and database objects. If a cache does not exist, SAP ASE writes a warning to the error log, and the binding in `sysattributes` is marked as invalid. Here is an example of the error log message:

```
Cache binding for database '5', object '208003772', index '3' is being marked
invalid in Sysattributes.
```

Invalid cache bindings are not deleted. If you create a cache of the same name and restart SAP ASE, the binding is marked as valid and the cache is used. If you do not create a cache with the same name, you can bind the object to another cache; otherwise, the object will automatically use the default cache.

If you want to retain cache configuration and cache binding settings in the cloud, you need to create the needed caches before making the database online and bind the cache to the database. The detailed steps are as follows.

Procedure

1. Get the list of data caches from the database information report:

The following is an example of the cache and binding information from the report:

```
=====
=
Section 4: Cache and binding information.
=====
=
Cache Name      Config Size    Run Size      Overhead      Cache Type
-----
testdb_cache    10.00 MB      10.00 MB      0.30 MB      Mixed
testidx_cache   30.00 MB      30.00 MB      0.69 MB      Mixed
testtab_cache   20.00 MB      20.00 MB      0.49 MB      Mixed
Memory Available For      Memory Configured
Named Caches              To Named Caches
-----
60.00 MB                  60.00 MB
-----
----- Cache Binding Information -----
Cache Name      Entity Name    Type          Index
Name            Status
-----
-----
testidx_cache   testdb.dbo.rrtab  index
idx1            V
```

testtab_cache table	testdb.dbo.rrtab V
testdb_cache database	testdb V

This report can be used to double check after the database is migrated to the cloud. Because cache-binding information of objects is stored in the database itself, the information is retained after the load.

2. Use the `ddlgen` utility to generate all the cache and pool configurations in the server, as shown in the following examples:

```
ddlgen -U<login> -P<password> -S<server>:<port> -TC -N%
```

or,

```
ddlgen -U<login> -P<password> -S<server>:<port> -TC -N<cache name>
```

In the generated configurations, save the information on all required caches, for example, `testidx_cache`, `testtab_cache`, and `testdb_cache`. The cache information looks like the following example:

```
-----
-- DDL for Cache 'testdb_cache'
-----
print '<<<< CREATING Cache - "testdb_cache" >>>>'
```

```
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name = 'testdb_cache'
AND parent = 19)
    exec sp_cacheconfig 'testdb_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testdb_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testdb_cache' , '10240K' , mixed
go
-----
-- DDL for Cache 'testidx_cache'
-----
print '<<<< CREATING Cache - "testidx_cache" >>>>'
```

```
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name =
'testidx_cache' AND parent = 19)
    exec sp_cacheconfig 'testidx_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testidx_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testidx_cache' , '30720K' , mixed
go
exec sp_poolconfig 'testidx_cache' , '10240K' , '32K'
go
-----
-- DDL for Cache 'testtab_cache'
-----
print '<<<< CREATING Cache - "testtab_cache" >>>>'
```

```
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name =
'testtab_cache' AND parent = 19)
```

```
exec sp_cacheconfig 'testtab_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testtab_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testtab_cache', '20480K', mixed
go
```

You can use it later in the cloud.

- Generate the commands to restore the cache binding for the database.

```
select "sp_bindcache "+ char_value + ", " + db_name(object) from
sysattributes where class = 3 and object_type = "D"
go
```


sp_bindcache testdb_cache, testdb
(1 row affected)

```

-----
-----
-----
-----
-----
-----
-----
sp_bindcache testdb_cache, testdb
(1 row affected)

```

Task overview: [Pre-Migration \[page 15\]](#)

Previous task: [Checking for Incompatible Features \[page 20\]](#)

Next: [Managing Login Accounts and Roles](#) [page 25]

4.7 Managing Login Accounts and Roles

You need to ensure the server user IDs for all users and user-defined roles are synchronized between the on-premises SAP ASE database and the database that is to be migrated SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

Parent topic: [Pre-Migration](#) [page 15]

Previous task: [Checking on Data Cache Binding Information \[page 23\]](#)

Next task: [Encrypting Your Database](#) [page 55]

4.7.1 Logins with System-Reserved Login Names or Server User IDs

SAP Adaptive Server Enterprise, cloud edition by IBM Cloud introduces several new login accounts with reserved server user IDs (suid) by default.

In rare cases, these login names or suids may be used in your on-premises server. You need to check whether these login names or suids are used in the source server or not. If yes, drop these login accounts or suids in the on-premises side before migrating databases.

Login Account	SUID	Description
saptu	-4	Internal technical user that performs all operations of provisioning
sapsupport	-5	SAP human operators
aseadmin	-6	Customer database administrator
sappr	-7	Used to rotate passwords of SAP users.
sapping	-8	Used by the ping agent
sapexporter	-9	Used for telemetry
sapaudlog	-10	Used for audit log service
sapamcadmin	-11	AMC admin for IBM human user
sapdradmin	-12	DR admin log of HADR in IBM setup
sapdrmaint	-13	DR maintenance login of HADR in IBM setup

4.7.1.1 Conflicts Detection

There are two kinds of conflicts between the user-defined logins and SAP reserved logins of the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud service.

You can use the `sp_serverinfo_report` and `sp_dbinfo_report` stored procedures to check these conflicts. If no conflicts are detected, you can skip the following detection steps.

Conflicts of Login Names

If there is any login by the name `saptu`, `sapsupport`, `aseadmin`, `sappr`, `sapping`, `sapexporter`, `sapaudlog`, `sapamcadmin`, `sapdradmin`, or `sapdrmaint`, create a new login with a common name to replace this old login, and then drop the old one.

You can use the following SQL commands to check if there's any SAP reserved login name existing in the on-premises server.

```
use master
```

```

go
select * from syslogins
where name in ('saptu', 'sapsupport', 'aseadmin', 'sappr', 'sapping',
'sapexporter', 'sapaudlog', 'sapamcadmin', 'sapdradmin', 'sapdrmaint')
go
use testdb
go
select * from sysusers
where sysusers.suid =
    (select master..syslogins.suid
     from master..syslogins
     where master..syslogins.name
         in ('saptu', 'sapsupport', 'aseadmin', 'sappr', 'sapping',
'sapexporter', 'sapaudlog', 'sapamcadmin', 'sapdradmin', 'sapdrmaint'))
go

```

If there is any row returned, it indicates that there's at least one conflicted login name. Follow instructions in [Resolving Login Name Conflicts \[page 27\]](#) to resolve the conflict.

Conflicts of Login suids

If there is any login with an suid of -4, -5, -6, -7, -8, -9, -10, -11, -12, or -13 update the suid for this login. You can run the following SQL commands to check if there's any SAP reserved login suid existing in the on-premises server.

```

use master
go
select * from syslogins where suid >= (-13) and suid <= (-4)
go
use testdb
go
select * from sysusers where suid >= (-13) and suid <= (-4)
go

```

If there is any row returned, it indicates that there's at least one conflicted login suid. Follow instructions in [Resolving Login suids Conflicts \[page 37\]](#) to resolve the conflict.

4.7.1.2 Resolving Login Name Conflicts

If there is a conflict of login names, create a new login with a normal name to replace the old login, and then drop the old one.

Context

This topic shows you how to resolve login name conflicts, assuming there's a login name `sapsupport` in the on-premises server. It owns a table called `t1` and a stored procedure called `p1` in the `master` database. It also owns a table called `t2` and a stored procedure called `p2` in the `testdb` database. The `sapsupport` login has the `mon_role` role, and has an alias user `sapsupport_alias` both in the `master` and `testdb` databases.

1. Obtain the following information on the `sap support` login.
 1. Get the information on logins `sap support` and `sap support_alias` in the `master` and `testdb` databases.

28 PUBLIC Pre-Migration

```

-----
7
sapsupport_alias 6
sapsupport
(1 row affected)

```

```

1> use testdb
2> go
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
suid      name
altsuid   altname
-----
7
sapsupport_alias 6
sapsupport
(1 row affected)

```

2. Check the roles that are assigned to the sapsupport login:

```

1> use master
2> go
1> sp_displaylogin sapsupport
2> go
Suid: 6
Loginame: sapsupport
Fullname:
Default Database: master
Default Language:
Auto Login Script:
Configured Authorization:
        mon_role (default ON)
Locked: NO
Date of Last Password Change: Aug 13 2020  2:09AM
Password expiration interval: 0
Password expired: NO
Minimum password length: 6
Maximum failed logins: 0
Current failed login attempts:
Authenticate with: AUTH_DEFAULT
Login Password Encryption: SHA-256
Last login date:
Exempt inactive lock: 0
(return status = 0)

```

3. Run sp_helprotect to display the permissions of sapsupport in the master and testdb databases:

```

1> use master
2> go
1> sp_helprotect sapsupport
2> go
grantor      grantee
type         action
object
column                                predicate      grantable
-----
-----
-----
-----

```

```

    dbo          mon_role
Grant          Select
dbo.monBucketPool
All                                NULL                FALSE
.....
    dbo          public
Grant          Select
dbo.sysxtypes
All                                NULL                FALSE
    dbo          sapsupport      Grant          Create
Procedure
All                                NULL
FALSE
    dbo          sapsupport      Grant          Create
Table
All                                NULL
FALSE
(0 rows affected)
(return status = 0)

```

```

1> use testdb
2> go
1> sp_helprotect sapsupport
2> go
    grantor      grantee      type
action          object
column          predicate      grantable
-----
    dbo          public      Grant
Select          dbo.sysalternates
All                                NULL                FALSE
.....
    dbo          public      Grant
Select          dbo.sysxtypes
All                                NULL                FALSE
    dbo          sapsupport      Grant
Create Procedure
All                                NULL                FALSE
    dbo          sapsupport      Grant
Create Table
All                                NULL                FALSE
(0 rows affected)
(return status = 0)

```

- Find all the tables and procedures created by sapsupport in the master and testdb databases:

```

1> use master
2> go
1> sp_helpuser sapsupport, display_object
2> go
    Object_name      Object_type      Create_date
-----
    p1              stored procedure      Aug 13 2020
2:17AM
    t1              user table          Aug 13 2020
2:17AM
(1 row affected)
User_type
-----
(return status = 0)

```

```

1> use testdb
1> go

```

```

1> sp_helpuser sapsupport, display_object
2> go
Object_name          Object_type          Create_date
-----
p2                   stored procedure     Aug 13 2020
2:18AM
t2                   user table           Aug 13 2020
2:18AM
(1 row affected)
User_type
-----
(return status = 0)

```

2. Create a new login called sapsupport_new to replace sapsupport:

```

1> use master
2> go
1> create login sapsupport_new with password abcd123
2> go
1> select suid, name from syslogins
2> go
suid      name
-----
1 sa
2 probe
4 ASELOGIN03
3 ASELOGIN01
5 ASELOGIN02
6 sapsupport
7 sapsupport_alias
8 sapsupport_new
(8 rows affected)

```

3. Create a new user called sapsupport_new for sapsupport_new.
 - In the master database, create the user, and then transfer the ownership of objects owned by user sapsupport to sapsupport_new. You can run sp_helpuser to verify the ownership.

```

1> use master
2> go
1> sp_adduser sapsupport_new
2> go
New user added.
(return status = 0)
1> alter all sapsupport.* modify owner sapsupport_new
2> go
1> sp_helpuser sapsupport, display_object
2> go
Object_name          Object_type          Create_date
-----
(1 row affected)
User_type
-----
(return status = 0)
1> sp_helpuser sapsupport_new, display_object
2> go
Object_name          Object_type          Create_date
-----
p1                   stored procedure     Aug 13 2020
2:17AM
t1                   user table           Aug 13 2020
2:17AM
(1 row affected)
User_type

```

```
-----
(return status = 0)
```

- Repeat the previous step in the testdb database.

```
1> use testdb
2> go
1> sp_adduser sapsupport_new
2> go
New user added.
(return status = 0)
1> alter all sapsupport.* modify owner sapsupport_new
2> go
1> sp_helpuser sapsupport, display_object
2> go
Object_name          Object_type          Create_date
-----
(1 row affected)
User_type
-----
(return status = 0)
1> sp_helpuser sapsupport_new, display_object
2> go
Object_name          Object_type          Create_date
-----
p2
2:18AM               stored procedure     Aug 13 2020
t2
2:18AM               user table           Aug 13 2020
(1 row affected)
User_type
-----
(return status = 0)
```

4. Obtain the privileges and aliases of sapsupport in the testdb database.

```
1> use testdb
2> go
1> sp_helprotect sapsupport
2> go
grantor      grantee      type
action      column      predicate      object      grantable
-----
dbo          public      Grant
Select      All          NULL          dbo.sysalternates FALSE
.....
dbo          public      Grant
Select      All          NULL          dbo.sysxtypes FALSE
dbo          sapsupport  Grant
Create Procedure All          NULL          FALSE
dbo          sapsupport  Grant
Create Table All          NULL          FALSE
(0 rows affected)
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
slog2
3> where salt.suid = slog1.suid
4> and salt.altsuid = slog2.suid
```

```

5> go
    suid      name
altsuid      altname
-----
7
sapsupport_alias
sapsupport
(1 row affected)
6

```

- Drop the sapsupport user from the testdb database:

```

1> use testdb
2> go
1> sp_dropuser sapsupport
2> go
The dependent aliases were also dropped.
User has been dropped from current database.
(return status = 0)
1> sp_helprotect sapsupport
2> go
Msg 17674, Level 16, State 1:
Procedure 'sp_helprotect', Line 321:
No such object or user exists in the database.
(return status = 1)

```

- Obtain the privileges and aliases of sapsupport in the master database:

```

1> use master
2> go
1> sp_helprotect sapsupport
2> go
grantor      grantee
type         action
object
column      predicate      grantable
-----
-----
dbo          mon_role
Grant       Select
dbo.monBucketPool
All         NULL          FALSE
.....
dbo          public
Grant       Select
dbo.sysxtypes
All         NULL          FALSE
dbo          sapsupport
Procedure   Grant         Create
All         NULL          FALSE
dbo          sapsupport
Table       Grant         Create
All         NULL          FALSE
(0 rows affected)
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
slog2
3> where salt.suid = slog1.suid
4>       and salt.altsuid = slog2.suid
5> go
    suid      name
altsuid      altname
-----
-----

```

```

7
sapsupport_alias
sapsupport
(1 row affected)

```

6

7. Drop the sapsupport user from the master database:

```

1> use master
2> go
1> sp_dropuser sapsupport
2> go
The dependent aliases were also dropped.
User has been dropped from current database.
(return status = 0)
1> sp_helprotect sapsupport
2> go
Msg 17674, Level 16, State 1:
Procedure 'sp_helprotect', Line 321:
No such object or user exists in the database.
(return status = 1)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
slog2
3> where salt.suid = slog1.suid
4> and salt.altsuid = slog2.suid
5> go
suid      name
altsuid   altname
-----
(0 rows affected)

```

8. Check the roles that are assigned to the sapsupport login:

```

1> sp_displaylogin sapsupport
2> go
Suid: 6
Loginname: sapsupport
Fullname:
Default Database: master
Default Language:
Auto Login Script:
Configured Authorization:
mon_role (default ON)
Locked: NO
Date of Last Password Change: Aug 13 2020 2:09AM
Password expiration interval: 0
Password expired: NO
Minimum password length: 6
Maximum failed logins: 0
Current failed login attempts:
Authenticate with: AUTH_DEFAULT
Login Password Encryption: SHA-256
Last login date: Aug 13 2020 2:17AM
Exempt inactive lock: 0
(return status = 0)

```

9. Grant roles and privileges and add an alias to the sapsupport_new login based on the information obtained in previous steps.
 1. Grant privileges to sapsupport_new in the master and testdb databases:

```

1> use master
2> go
1> grant create table to sapsupport_new
2> go

```

```

1> grant create procedure to sapsupport_new
2> go
1> sp_helprotect sapsupport_new
2> go

```

grantor	grantee	action	object	column	predicate	grantable
dbo	public	Grant	Execute	dbo.sp_autoformat	NULL	FALSE
dbo	public	Grant	Select	dbo.sysxtypes	NULL	FALSE
dbo	sapsupport_new	Grant	Create	Procedure	NULL	FALSE
dbo	sapsupport_new	Grant	Create	Table	NULL	FALSE

```

(0 rows affected)
(return status = 0)

```

```

1> use testdb
2> go
1> grant create table to sapsupport_new
2> go
1> grant create procedure to sapsupport_new
2> go
1> sp_helprotect sapsupport_new
2> go

```

grantor	grantee	action	object	column	predicate	grantable
dbo	public	Grant	dbo.sysalternates	SELECT	NULL	FALSE
dbo	public	Grant	dbo.sysxtypes	SELECT	NULL	FALSE
dbo	sapsupport_new	Grant	Create Procedure	ALL	NULL	FALSE
dbo	sapsupport_new	Grant	Create Table	ALL	NULL	FALSE

```

(0 rows affected)
(return status = 0)

```

2. Add an alias to sapsupport_new in the master and testdb databases:

```

1> use master
2> go
1> sp_addalias sapsupport_alias, sapsupport_new

```

```

2> go
Alias user added.
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
    suid      name
altsuid      altname
-----
              7
sapsupport_alias                                8
sapsupport_new
(1 row affected)

```

```

1> use testdb
2> go
1> sp_addalias sapsupport_alias, sapsupport_new
2> go
Alias user added.
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
    suid      name
altsuid      altname
-----
              7
sapsupport_alias                                8
sapsupport_new
(1 row affected)

```

3. Grant roles to sapsupport_new in the master database according to the information you obtained in [Step 8 \[page 34\]](#):

```

1> use master
2> go
1> grant role mon_role to sapsupport_new
2> go
1> sp_displaylogin sapsupport_new
2> go
Suid: 8
Loginame: sapsupport_new
Fullname:
Default Database: master
Default Language:
Auto Login Script:
Configured Authorization:
        mon_role (default ON)
Locked: NO
Date of Last Password Change: Aug 13 2020  6:25PM
Password expiration interval: 0
Password expired: NO
Minimum password length: 6
Maximum failed logins: 0
Current failed login attempts:
Authenticate with: AUTH_DEFAULT
Login Password Encryption: SHA-256
Last login date:
Exempt inactive lock: 0
(return status = 0)

```

10. Modify objects, such as stored procedures, triggers, or views, which refer the old `sapsupport` login.

For example, if there is a SQL statement as below in any stored procedure,

```
select * from 'sapsupport.t1'
```

You need to modify this stored procedure as below:

```
select * from 'sapsupport_new.t1'
```

11. Once the `sapsupport_new` login is equivalent to the old `sapsupport` login, drop the old login.

```
1> use master
2> go
1> drop login sapsupport
2> go
```

4.7.1.3 Resolving Login suids Conflicts

If there is a conflict of login suids, update the `suid` for this login.

Context

This topic shows you how to resolve the login suids conflicts, assuming there's a login called `joeuser` with a `suid` of `-5` in the on-premises server. It owns a table called `t1` and a stored procedure called `p1` in the `master` database. It also owns a table called `t2` and a stored procedure called `p2` in the `testdb` database. And the user `joeuser` has an alias user `joeuser_alias` both in the `master` and `testdb` databases.

Procedure

1. Obtain the following information regarding the login `joeuser`.
 - Get the information of the `joeuser` and `joeuser_alias` logins in the `master` and `testdb` databases.

```
1> use master
2> go
1> select * from syslogins where suid >= (-13) and suid <= (-4)
2> go
suid          status accdate          totcpu
totio         spacelimit  timelimit  resultlimit
dbname
name
password

language
pwwdate          audflags
fullname
```

```

srvname      lastlogindate          logincount
procid       crdate               locksuid    lockreason
lockdate     lpid                 crsuid
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
0         -5           0              Aug 12 2020   6:27PM             0
0 ``master``
joeuser
0xc0072fe340f1a33a1953722f6ff9334e7f8d8294f03bc292cfc0b0f09d9c67e4433ca3241
0b52a149dd7

NULL
Aug 12 2020 6:27PM            0
NULL
NULL
NULL                     NULL                Aug 12 2020
6:27PM                   NULL                  NULL
NULL                      1
(1 row affected)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
4> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
5> and salt.altsuid >= (-13) and salt.altsuid <= (-4)
6> go
suid        name
altsuid      altname
-----
-----
-----
10
joeuser_alias                               -5
joeuser
(1 row affected)
```

- Run `sp_helprotect` to print the permissions of the user `joeuser` in the `master` and `testdb` databases.

```

1> use master
2> go
1> sp_helprotect joeuser
2> go
  grantor      grantee      type      action
object
column                predicate      grantable
-----
dbo
Procedure          joeuser      Grant      Create
All
FALSE
dbo
Table              joeuser      Grant      Create
All
FALSE
.....
dbo
dbo.sysxtypes      public      Grant      Select
All
(0 rows affected)
(return status = 0)

```

```

1> use testdb
2> go
1> sp_helprotect joeuser
2> go
  grantor      grantee      type      action
object
predicate      grantable      column
-----
dbo
Create Procedure  joeuser      Grant
All
NULL
FALSE
dbo
Create Table     joeuser      Grant
All
NULL
FALSE
.....
dbo
dbo.sysxtypes    public      Grant      Select
NULL
FALSE
All
(0 rows affected)
(return status = 0)

```

- Find all tables and procedures created by `joeuser` in the `master` and `testdb` databases.

```

1> use master
2> go
1> sp_helpuser joeuser, display_object
2> go
  Object_name      Object_type      Create_date
-----
p1
7:22PM            stored procedure  Aug 12 2020
t1
7:21PM            user table        Aug 12 2020
(1 row affected)
User_type

```

```
-----
(return status = 0)
```

```
1> use testdb
2> go
1> sp_helpuser joeuser, display_object
2> go
Object_name          Object_type          Create_date
-----
p2                    stored procedure     Aug 12 2020
7:29PM
t2                    user table           Aug 12 2020
7:25PM
(1 row affected)
User_type
-----
(return status = 0)
```

2. Create a temporary login joeuser_tmp for joeuser with suid -5:

```
1> use master
2> go
1> create login joeuser_tmp with password abcd123
2> go
```

3. Create a user joeuser_tmp for the joeuser_tmp login in the master and testdb databases.

- In the master database, create a user joeuser_tmp, and then transfer the ownership of objects owned by joeuser to joeuser_tmp. Use sp_helpuser to verify that the ownership is changed.

```
1> use master
2> go
1> sp_adduser joeuser_tmp
2> go
New user added.
(return status = 0)
1> alter all joeuser.* modify owner joeuser_tmp
2> go
1> sp_helpuser joeuser, display_object
2> go
Object_name          Object_type          Create_date
-----
(1 row affected)
User_type
-----
(return status = 0)
1> sp_helpuser joeuser_tmp, display_object
2> go
Object_name          Object_type          Create_date
-----
p1                    stored procedure     Aug 12 2020
7:22PM
t1                    user table           Aug 12 2020
7:21PM
(1 row affected)
User_type
-----
(return status = 0)
```

- Repeat the previous step in the testdb database:

```
1> use testdb
2> go
1> sp_adduser joeuser_tmp
```

```

2> go
New user added.
(return status = 0)
1> alter all joeuser.* modify owner joeuser_tmp
2> go
1> sp_helpuser joeuser, display_object
2> go
Object_name          Object_type          Create_date
-----
(1 row affected)
User_type
-----
(return status = 0)
1> sp_helpuser joeuser_tmp, display_object
2> go
Object_name          Object_type          Create_date
-----
p2
7:29PM              stored procedure      Aug 12 2020
t2
7:25PM              user table             Aug 12 2020
(1 row affected)
User_type
-----
(return status = 0)

```

4. Obtain the information of privileges and aliases in the `testdb` database.

Note that all the information of privileges for user `joeuser` in the `testdb` database needs to be synchronized.

```

1> use testdb
2> go
1> sp_helprotect joeuser
2> go
grantor      grantee      type      action
object
predicate      grantable
-----
dbo          joeuser      Grant
Create Procedure
All          NULL          Grant      FALSE
dbo          joeuser      Grant
Create Table
All          NULL          Grant      FALSE
.....
dbo          public      Grant      Select
dbo.sysxtypes
NULL          FALSE
(0 rows affected)
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
slog2
3> where salt.suid = slog1.suid
4>      and salt.altsuid = slog2.suid
5>      and salt.altsuid = -5
6> go
suid      name
altsuid    altname
-----
-----

```

```

10
joeuser_alias
joeuser
(1 row affected)
-5

```

- Drop user joeuser from the testdb database:

```

1> use testdb
2> go
1> sp_dropuser joeuser
2> go
The dependent aliases were also dropped.
User has been dropped from current database.
(return status = 0)

```

- Obtain the privileges and aliases of joeuser in the master database:

```

1> use master
2> go
1> sp_helprotect joeuser
2> go

```

grantor	grantee	type	action	object	column	predicate	grantable
dbo	joeuser	Grant	Create	Procedure			
dbo	joeuser	Grant	Create	Table			FALSE
dbo	public	Grant	Select	dbo.sysxtypes			FALSE

```

.....
(0 rows affected)
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
slog2
3> where salt.suid = slog1.suid
4> and salt.altsuid = slog2.suid
5> and salt.altsuid = -5
6> go

```

suid	name	altsuid	altname
10	joeuser_alias	-5	

```

joeuser
(1 row affected)

```

- Drop user joeuser from the master database:

```

1> use master
2> go
1> sp_dropuser joeuser
2> go
The dependent aliases were also dropped.
User has been dropped from current database.
(return status = 0)

```

8. Obtain the information about roles assigned to the login joeuser:

```
1> sp_displaylogin joeuser
2> go
Suid: -5
Loginame: joeuser
Fullname:
Default Database: master
Default Language:
Auto Login Script:
Configured Authorization:
        mon_role (default ON)
Locked: NO
Date of Last Password Change: Aug 12 2020  6:27PM
Password expiration interval: 0
Password expired: NO
Minimum password length: 6
Maximum failed logins: 0
Current failed login attempts:
Authenticate with: AUTH_DEFAULT
Login Password Encryption: SHA-256
Last login date: Aug 12 2020  7:18PM
Exempt inactive lock: 0
(return status = 0)
```

9. Drop the old login joeuser:

```
1> use master
2> go
1> drop login joeuser
2> go
```

10. Re-create the login joeuser again. Now it should have a new suid other than the reserved -5.

```
1> use master
2> go
1> create login joeuser with password abcd123
2> go
1> select suid, name from syslogins
2> go
suid      name
-----
1 sa
2 probe
3 ASELOGIN01
4 ASELOGIN02
12 joeuser
10 joeuser_alias
9 ASELOGIN04
11 joeuser_tmp
(8 rows affected)
```

11. Add the re-created login to databases in which the original login was present.

Re-create the user joeuser in the master and testdb databases.

```
1> use master
2> go
1> sp_adduser joeuser
2> go
New user added.
(return status = 0)
1> sp_addalias joeuser_alias, joeuser
2> go
Alias user added.
(return status = 0)
```

```

1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
   slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
      suid          name
altsuid          altname
-----
          10
joeuser_alias                                12
joeuser
(1 row affected)

```

```

1> use testdb
2> go
1> sp_adduser joeuser
2> go
New user added.
(return status = 0)
1> sp_addalias joeuser_alias, joeuser
2> go
Alias user added.
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1, master.dbo.syslogins
   slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
      suid          name
altsuid          altname
-----
          10
joeuser_alias                                12
joeuser
(1 row affected)

```

12. Transfer the ownership of objects owned by joeuser_tmp to joeuser in the master and testdb databases.

```

1> use master
2> go
1> alter all joeuser_tmp.* modify owner joeuser
2> go
1> use testdb
2> go
1> alter all joeuser_tmp.* modify owner joeuser
2> go

```

13. Grant roles and privileges, and add aliases to joeuser based on the obtained information in previous steps.

1. Grant privileges to joeuser:

```

1> use master
2> go
1> grant create table to joeuser
2> go
1> grant create procedure to joeuser
2> go
1> sp_helprotect joeuser
2> go
      grantor          grantee          type          action
object
column                predicate          grantable

```

```

-----
-----
-----
dbo      joeuser      Grant      Create
Procedure
All
FALSE
dbo      joeuser      Grant      Create
Table
All
FALSE
.....
dbo      public      Grant      Select
dbo.sysxtypes
All
NULL
FALSE
(0 rows affected)
(return status = 0)

```

```

1> use testdb
2> go
1> grant create table to joeuser
2> go
1> grant create procedure to joeuser
2> go
1> sp_helprotect joeuser
2> go
grantor      grantee      type      action
object      predicate      grantable      column
-----
-----
dbo      joeuser      Grant
Create Procedure
All
NULL
FALSE
dbo      joeuser      Grant
Create Table
All
NULL
FALSE
.....
dbo      public      Grant      Select
dbo.sysxtypes
All
FALSE
(0 rows affected)
(return status = 0)

```

2. Add an alias to joeuser:

```

1> use master
2> go
1> sp_addalias joeuser_alias, joeuser
2> go
Alias user added.
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
suid      name
altsuid      altname
-----
-----
10
joeuser_alias
joeuser
(1 row affected)

```

```

1> use testdb
2> go
1> sp_addalias joeuser_alias, joeuser
2> go
Alias user added.
(return status = 0)
1> select salt.suid, slog1.name, salt.altsuid, slog2.name as altname
2> from sysalternates salt, master.dbo.syslogins slog1,
master.dbo.syslogins slog2
3> where salt.suid = slog1.suid and salt.altsuid = slog2.suid
4> go
      suid      name
altsuid      altname
-----
              10
joeuser_alias                                12
joeuser
(1 row affected)

```

3. Grant roles to joeuser:

```

1> use master
2> go
1> grant role mon_role to joeuser
2> go
1> sp_displaylogin joeuser
2> go
Suid: 12
Loginname: joeuser
Fullname:
Default Database: master
Default Language:
Auto Login Script:
Configured Authorization:
      mon_role (default ON)
Locked: NO
Date of Last Password Change: Aug 12 2020 11:20PM
Password expiration interval: 0
Password expired: NO
Minimum password length: 6
Maximum failed logins: 0
Current failed login attempts:
Authenticate with: AUTH_DEFAULT
Login Password Encryption: SHA-256
Last login date:
Exempt inactive lock: 0
(return status = 0)

```

14. Drop the temporary user from each user database and master database:

Drop the temporary user joeuser_tmp from the testdb database:

```

1> use testdb
2> go
1> sp_dropuser joeuser_tmp
2> go
User has been dropped from current database.
(return status = 0)

```

Drop the temporary user joeuser_tmp from the master database.

```

1> use master
2> go
1> sp_dropuser joeuser_tmp
2> go

```

```
User has been dropped from current database.  
(return status = 0)
```

15. Drop the temporary login joeuser_tmp from the master database.

```
1> use master  
2> go  
1> drop login joeuser_tmp  
2> go
```

4.7.2 Obtaining Login Accounts

Get the list of server user IDs (suid) for the on-premises database so that you can recreate them after the database is migrated to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

Prerequisites

Make sure all the reserved login accounts and server user IDs are not used in the on-premises side. See [Logins with System-Reserved Login Names or Server User IDs \[page 26\]](#) for details.

Context

For security reasons, login accounts are created in the master database and users are created in the user databases. The syslogins system table contains one row for each valid login account and a relevant suid. The sysusers system table contains one row for each allowed user in the database, and each user maps a suid which is copied from syslogins.

Dumping and loading databases does not change the value of suid in the sysusers system table. After database loading is complete, you need to make sure the login account in the loaded user database matches with the same login account in the on-premises database. If the login account does not exist in the cloud side, you need to create this login account.

Procedure

1. Get the list of the login accounts and users that are used in the database to be migrated.

For example, login accounts ASELOGIN01 and ASELOGIN02 can be found in the on-premises side:

```
1> use master  
2> go  
1> select * from syslogins where suid >= 3  
2> go  
suid      status accdate      totcpu      totio  
spacelimit timelimit resultlimit
```


ASELOGIN02	5	public	ASELOGIN02
dbo	1	public	sa
usedb_user	3	public	NULL

(1 row affected)
(return status = 0)

You can also get the information by running the `sp_dbinfo_report` command.

```
=====
=
Section 6: Logins related information.
=====
=
There is no reserved login account with reserved suids
-4/-5/-6/-7/-8/-9/-10/-11/-12/-13 or names
saptu/sapsupport/aseadmin/sappr/sapping/sapexporter/sapaudlog/sapamcadmin/
sapdradmin/sapdrmaint.
Total number of login accounts which are explicitly added by sp_adduser in
database tdb1: 2
suid  name      fullname  procid  lpid
-----
3 ASELOGIN01 NULL      NULL    NULL
4 ASELOGIN02 NULL      NULL    NULL
```

2. Run the `ddlgen` utility to generate DDLs of logins, which you will use to create logins in the cloud.

`ddlgen` is located in the `$SDK/tools/bin` directory. Use `ddlgen`:

- with the `-X OEP` option to get the DDLs of logins with encrypted password and suid options
- with `-D` option only to generate the DDLs of logins which are explicitly add by `sp_adduser` in the specified database.

Note

You need to set the environment variables `SYBROOT` and `SAP_JRE8` before running `ddlgen`. For example:

```
setenv SYBROOT $SDK_PATH
setenv SAP_JRE8 $SYBASE/shared/SAPJRE-8_1_062_64BIT
```

The syntax is as follows:

```
ddlgen -U<user_name> -P<password> -S<server_name> -I<interfaces_file_name>
-TL -D <database_name> -N % -X OEP
```

The output is as follows:

```
-----
-- DDL for Login 'ASELOGIN01'
-----
print '<<<< CREATING Login - "ASELOGIN01" >>>> '
go
IF EXISTS (SELECT 1 FROM master.dbo.syslogins WHERE name = 'ASELOGIN01')
    drop login 'ASELOGIN01'
go
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Login ASELOGIN01'
    SELECT syb_quit()
END
go
-- The DDL is generated with encrypted password - '0xxxxxxx'
-- and suid when this login was created.
```

```

create login ASELOGIN01 with encrypted password
0xc007b6d4abb6d76913fdcf4369b7cd24b83d11f3fdefed3f66e310ca5130f845b6d93aafa52e
c73a8364
    suid 3
    default database 'master'
    authenticate with 'ANY'
    exempt inactive lock false
go
exec sp_locklogin 'ASELOGIN01', 'unlock'
go
-----
-- DDL for Login 'ASELOGIN02'
-----
print '<<<< CREATING Login - "ASELOGIN02" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.syslogins WHERE name = 'ASELOGIN02')
    drop login 'ASELOGIN02'
go
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Login ASELOGIN02'
    SELECT syb_quit()
END
go
-- The DDL is generated with encrypted password - '0xxxxxxx'
-- and suid when this login was created.
create login ASELOGIN02 with encrypted password
0xc0073aafa52ec73a83640dedd2979b9d09cb9e56bf783a02b547c46bc12793bd36ca0085d9a8
278e2ba9
    suid 4
    default database 'master'
    authenticate with 'ANY'
    exempt inactive lock false
go
exec sp_locklogin 'ASELOGIN02', 'unlock'
go
-----
-- Dependent DDL for Object(s)
-----
sp_passwordpolicy 'set', 'rsa keysizes', '2048'
go

```

⚠ Caution

The generated login creation commands contain sensitive information. From the security perspective, make sure you store these commands in a safe place. The ASE service can protect sensitive information by locking logins or roles with a configured number of failed attempts, so a brute force method cannot be used to crack the password. But if these passwords are accessed from outside of the ASE service, the ASE service does not have control on them and a brute force method can be used to crack the passwords.

4.7.3 Obtaining User-Defined Roles

Obtain user-defined roles from the on-premises server before migration to recreate them after the database is migrated to the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud service.

Context

For security reasons user-defined roles are created in the master database and are used to grant permissions to objects in the user databases. The `sysrvroles` system table in the master database contains one row for each valid role and each role has a unique `srid` (server role ID). The `sysroles` system table contains one row for each granted role in the database and each role has a `lrid` (logical role ID) which maps to a role name that is copied from `master..sysrvroles`.

Dumping and loading databases does not change the value of `lrid` in the `sysroles` system table of the database and the role name in the `sysusers` system table. After migration you just need to create the same user-defined roles in the cloud service.

Procedure

1. In the on-premises SAP ASE, review the list of user-defined roles that are used in the database to be migrated.

In the following example, the roles `usrdef_role01` and `usrdef_role02` exist in the on-premises side and are used to grant permissions in the `tdb2` database.

```
1> use master
2> go
1> select * from sysrvroles where srid > 32
2> go
srid      name
password

locksuid  lockreason  pwdate      status logincount
lockdate
-----
33 usrdef_role01
0xc00709d58947c8eadb6f214400acb5fd5cf875e325f656bf57b27f9ff38d66f53ad69f240573ec03feff
0          NULL      NULL      Jul 21 2020  6:26PM      0
34 usrdef_role02
0xc0079f240573ec03feffa260149927fbcccc017776793f3c83864bf8a5e2efbeec46283cad25241db7b9
0          NULL      NULL      Jul 21 2020  6:26PM      0
```

```

(2 rows affected)
1> use tdb2
2> go
1> select su.suid, su.uid, su.gid, su.name, ssr.srid
2> from sysusers su, sysroles sr, master.dbo.syssrvroles ssr
3> where su.uid = sr.lrid
4>       and sr.id = ssr.srid
5>       and ssr.srid > 32
6> go

```

suid	uid	gid	name	srid
-2	16401	16401	usrdef_role01	33
-2	16402	16402	usrdef_role02	34

```

(2 rows affected)
1> sp_helprotect usrdef_role01
2> go

```

grantor object	grantee column	predicate	type grantable	action
dbo	usrdef_role01		Grant	Delete
dbo.tb1	All	NULL	FALSE	
dbo	usrdef_role01		Grant	Insert
dbo.tb1	All	NULL	FALSE	
dbo	usrdef_role01		Grant	Update
dbo.tb1	All	NULL	FALSE	

```

(1 row affected)
(return status = 0)
1> sp_helprotect usrdef_role02
2> go

```

grantor object	grantee column	predicate	type grantable	action
dbo	usrdef_role02		Grant	Delete
dbo.tb1	All	NULL	FALSE	
dbo	usrdef_role02		Grant	Insert
dbo.tb1	All	NULL	FALSE	
dbo	usrdef_role02		Grant	Update
dbo.tb1	All	NULL	FALSE	

```

(1 row affected)
(return status = 0)

```

You can also get the information by running the `sp_dbinfo_report tdb1` command:

```

=====
=
Section 7: User-defined roles related information.
=====
=
Total number of roles which are granted database-wide privileges in
database tdb1: 2

```

suid	uid	gid	name	srid
-2	16401	16401	usrdef_role01	33
-2	16402	16402	usrdef_role02	34

2. Run the `ddlgen` utility to generate DDLs of roles, which you will use to create roles in the cloud.

`ddlgen` is located in the `$SDK/tools/bin` directory. Use `ddlgen`:

- with the `-X OEP` option to get the DDLs of roles with encrypted password options

- with the -D option to get the roles which are granted database-wide privileges in the specified database

Note

You need to set the environment variables SYBROOT and SAP_JRE8 before running ddlgen. For example:

```
setenv SYBROOT $SDK_PATH
setenv SAP_JRE8 $SYBASE/shared/SAPJRE-8_1_062_64BIT
```

The syntax is as follows:

```
ddlgen -U<user_name> -P<password> -S<server_name> -I<interfaces_file_name>
-TRO -D <database_name> -N % -X OEP
```

The output is similar to this:

```
-----
-- DDL for Role 'usrdef_role01'
-----
print '<<<< CREATING Role - "usrdef_role01" >>>> '
go
use master
go
IF EXISTS (SELECT 1 FROM master.dbo.sysrvroles where name = 'usrdef_role01')
    DROP ROLE usrdef_role01 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role01'
    SELECT syb_quit()
END
go
create role usrdef_role01 with encrypted password
0xc00775bc19efc136539ef52ac99d9460c4381a4d58b7e3f0ed9712f7599080b9d3cd86de0d53
05bc857e
go
use tdb1
go
Grant Insert on dbo.tb1 to usrdef_role01 Granted by dbo
go
Grant Delete on dbo.tb1 to usrdef_role01 Granted by dbo
go
Grant Update on dbo.tb1 to usrdef_role01 Granted by dbo
go
-----
-- DDL for Role 'usrdef_role02'
-----
print '<<<< CREATING Role - "usrdef_role02" >>>> '
go
IF EXISTS (SELECT 1 FROM master.dbo.sysrvroles where name = 'usrdef_role02')
    DROP ROLE usrdef_role02 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role02'
    SELECT syb_quit()
END
go
create role usrdef_role02 with encrypted password
0xc00786de0d5305bc857e8c2ffe8ec7752d2bb242d1c62b1aa2fa9603870bc7341a5aca23045b
f9f73555
go
use tdb1
go
```

```
Grant Insert on dbo.tb1 to usrdef_role02 Granted by dbo
go
Grant Delete on dbo.tb1 to usrdef_role02 Granted by dbo
go
Grant Update on dbo.tb1 to usrdef_role02 Granted by dbo
go
```

⚠ Caution

The generated role creation commands contain sensitive information. From the security perspective, make sure you store these commands in a safe place. The ASE service can protect sensitive information by locking logins or roles with a configured number of failed attempts, so a brute force method cannot be used to crack the password. But if these passwords are accessed from outside of the ASE service, the ASE service does not have control on them and a brute force method can be used to crack the passwords.

3. Create a role creation script according to the previous output. You can use this script to create roles in the cloud.

Role creation script:

```
-----
-- DDL for Role 'usrdef_role01'
-----
print '<<<< CREATING Role - "usrdef_role01" >>>>'
go
use master
go
IF EXISTS (SELECT 1 FROM master.dbo.syssrvroles where name = 'usrdef_role01')
    DROP ROLE usrdef_role01 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role01'
    SELECT syb_quit()
END
go
create role usrdef_role01 with encrypted password
0xc00775bc19efc136539ef52ac99d9460c4381a4d58b7e3f0ed9712f7599080b9d3cd86de0d53
05bc857e
go
-----
-- DDL for Role 'usrdef_role02'
-----
print '<<<< CREATING Role - "usrdef_role02" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.syssrvroles where name = 'usrdef_role02')
    DROP ROLE usrdef_role02 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role02'
    SELECT syb_quit()
END
go
create role usrdef_role02 with encrypted password
0xc00786de0d5305bc857e8c2ffe8ec7752d2bb242d1c62b1aa2fa9603870bc7341a5aca23045b
f9f73555
go
```

4.8 Encrypting Your Database

The dump-and-load migration method requires the source database (on-premises) and the target database (in the cloud) to have the same database encryption key (DEK). Since databases created in SAP Adaptive Server Enterprise, cloud edition by IBM Cloud are encrypted by default, you must encrypt your on-premises database before dumping for migration.

Prerequisites

- You have the `keycustodian_role` or the `sso_role` to create the master KEY and DEK.
- You have the `manage database encryption key` privilege if you are using granular permissions.
- You are the database owner or a user with the `sa` privilege to run `alter database ... encrypt`.
- You have a valid SAP ASE encryption feature license (ASE_ENCRYPTION). If not currently in use, the database encryption feature can be enabled as a temporary 90-day trial. Otherwise, please contact your sales representative to arrange a trial.

Context

See the following topics in *Database Encryption* for SAP ASE for more information about encryption.

- Encryption Status and Progress
- Performance Considerations
- Suspend the Encryption Process
- Resume the Encryption Process

Procedure

1. Configure the SAP ASE server for database encryption.
 1. Install the SAP ASE encryption feature license ASE_ENCRYPTION. See the *Installation Guide* for SAP ASE.
 2. Set the `enable encrypted columns` configuration parameter to true:

```
sp_configure 'enable encrypted columns', 1
```

2. Create the key-encryption key (KEK) in the master database.
 1. Create the encryption key:

```
create encryption key master with passwd "<your_password>"
```

2. Configure the password:

```
set encryption passwd "<your_password>" for key master
```

3. Create the DEK. To create a single DEK, run:

```
create encryption key <key_name> for database encryption
```

The DEK can be protected by a master key or by dual control.

4. Use the `alter database` command to encrypt your database. For example, this encrypts the `pubs2` database with a DEK named `important_key`:

```
alter database pubs2 encrypt with important_key
```

Note that database encryption requires at least two worker processes. You can use `sp_configure "number of worker processes"` to set the number of worker processes.

Encrypting an existing database with `alter database` requires that the server reads existing data and encrypts the contents, which can take a considerable amount of time (depending on database size).

However, the encryption process occurs in the background. You can run `sp_helpdb` to check the encryption process. The progress is displayed as a percentage while the database encryption is in progress. The database is shown as encrypted once the encryption is completed.

5. Set the KEK passwords in memory before you can bring the encrypted databases online. You can set the password in the following ways:

- Manually

1. Run the following command to set the password for the KEK:

```
set encryption passwd '<the_password>' for key master
```

2. Bring the encrypted databases online with the `online database` command.

```
online database <database_name>
```

- Without supervision – Use an unattended startup to automatically set the KEK password if you are unavailable to set the passwords:

1. Enable unattended startup mode:

```
sp_configure "automatic master key access",1
```

Enabling this configuration parameter stores the KEK encryption keys in the KEK startup file. If this file does not exist, a file is created. However, the encryption key information is not written to the file until `automatic_startup` key copies of the master or dual master keys are created. Use `sp_encryption` to override the name and location of this file.

2. Make copies of the KEK (and dual master key, if used) using `alter encryption key`. The encrypted databases come online unattended because SAP ASE loads the KEKs from the KEK startup file:

```
alter encryption key master with passwd "your_password" add encryption  
for automatic_startup
```

Task overview: [Pre-Migration \[page 15\]](#)

Previous: [Managing Login Accounts and Roles \[page 25\]](#)

Next task: [Dumping the Database and Transactions \[page 57\]](#)

4.9 Dumping the Database and Transactions

Use the `dump database` command to make a copy of the entire database, which is to be loaded in the cloud, and then use the `dump transaction` command to make a copy of the transactions since the database dump.

Context

You can put the database into a read-only state before dumping the database to ensure data consistency. This read-only state means that no changes can be made to your on-premises database until you restore the database in the cloud. The duration of down time depends on the amount of data to be backed up and stored, and the speed of your internet connections.

You can also dump the entire database without putting the database into a read-only state and use `dump transactions` and `load transactions` afterward to synch the data. This method can limit your database downtime to a minimum. This guide will show you how to use this method to dump and load your database and transactions.

Procedure

1. Perform a dump of the database:

```
dump database <db_name> to <file_name>
```

See *dump database* in *Reference Manual: Commands for SAP ASE* for more information about the `dump database` command.

For migration purposes, use `stripes` parameter to enhance the performance. Striping allows you to use multiple dump devices for a single dump or load command. Dumps are made concurrently on all devices, reducing the time required to dump an individual database. You can use the stored procedure `sp_config_dump` to set:

- `@stripe_dir` – a file system directory in which files are archived during the dump operation. Archived files are typically named using this format: `database_name.dump_type.date-timestamp.stripeID`. `@stripe_dir` defaults to the directory where the backup server is started.
- `@num_stripes` – the number of stripe devices used during the dump operation. The default is 1.

The following example creates a new dump configuration called `dmp_cfg2`, which specifies that the dump operation creates 5 stripes in the `/opt/bin/Sybase/dumps` stripe directory:

```
sp_config_dump 'dmp_cfg2',  
              @stripe_dir='/opt/bin/Sybase/dumps', @num_stripes='5',
```

You can use the same stored procedure to change the existing dump configuration. The following command changes the stripe number of the existing dump configuration `dmp_cfg2`:

```
sp_config_dump 'dmp_cfg2',@num_stripes = 6
```

The following command can be used to dump a database with a specified dump configuration:

```
dump database pubs2 using config = 'dmp_cfg2'
```

See `sp_config_dump` in the *Reference Manual: Procedures* for SAP ASE for details about `sp_config_dump`. Alternatively, you can specify the stripe device name in the `dump database` command. For example:

```
dump database mydb to "mydb.dump.stripe"  
stripe on "mydb.dump.stripe1"  
stripe on "mydb.dump.stripe2"
```

The previous command creates three dump devices. So, the backup server dumps approximately one-third of the database to each device.

Note

Compression parameter of the `dump database` command is not recommended for migration because compression does not work on encrypted databases. Database encryption is mandatory in the cloud edition. Compression cannot reduce the size of the dump file in this case.

2. Run the `dump transaction` command to synch up transactions that happened after the database dump.

For example:

```
DUMP TRAN testdb TO 'testdb.tran'  
go
```

3. Run the `load database` command and provide the path of the dump database file to retrieve all database options:

```
load database <database_name> from '<dump_database_file_path>' with  
listonly=create_sql
```

The following example shows the fragment of the returned result that lists all the options created in the on-premises database:

```
master..sp_dboption <database_name>, 'select into/bulkcopy/pllsort', true  
go  
master..sp_dboption <database_name>, 'allow nulls by default', true  
go
```

Note

Make sure the `sp_configure enable_dump_history` parameter is turned on to use the `listonly = create_sql` and `listonly = load_sql` options.

4. Use the `load database` command to regenerate the load command sequence required to restore the database using the latest available dumps:

```
load database <src_dbname> with listonly=load_sql
```

The following example shows the fragment of the returned result that lists all the load commands sequence required to restore the `testdb` database:

```
LOAD DATABASE testdb FROM  
'testdb.dump.stripe'  
STRIPE ON 'testdb.dump.stripe1'  
STRIPE ON 'testdb.dump.stripe2'
```

```
go
LOAD TRAN testdb FROM 'testdb.tran'
go
```

Task overview: [Pre-Migration \[page 15\]](#)

Previous task: [Encrypting Your Database \[page 55\]](#)

5 Migration

This section describes the steps you need to take when migrating the database from on-premises SAP ASE to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

1. [Provisioning an SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Instance \[page 60\]](#)
Provision an SAP ASE instance on IBM Cloud.
2. [Restoring Non-Default Configurations \[page 61\]](#)
Restore the non-default configurations collected by running the `sp_serverinfo_report` stored procedure.
3. [Creating Data Caches \[page 61\]](#)
You should create the same data caches in the cloud as in the on-premises database before loading the database.
4. [Migrating the Database Encryption Key \[page 63\]](#)
You can use the `transfer encryption key` command to transfer the database encryption key (DEK) from your on-premises database to the SAP ASE instance in the IBM Cloud.
5. [Creating a Database \[page 64\]](#)
Create a database in the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud before you upload the dump file.
6. [Synchronizing Login Accounts and Roles \[page 65\]](#)
Create the same login accounts and roles in the ASE instance as in the on-premises server.
7. [Uploading Dump Files \[page 69\]](#)
After you have generated the dump files for your database, upload the dump files to the instance specific backup bucket in IBM Cloud Object Store (COS).
8. [Loading the Database and Transactions \[page 70\]](#)
Use SAP ASE Cloud Portal to restore the database dump files to the cloud.
9. [Bringing the Database Online \[page 71\]](#)
You should bring the database online after loading the database and transaction dumps successfully.
10. [Restoring Cache Bindings \[page 72\]](#)
You need to restore cache bindings after the database is online.

5.1 Provisioning an SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Instance

Provision an SAP ASE instance on IBM Cloud.

To ensure the migration is successful:

- The page size, character set, and sort order in your created instance should be synchronized with the source SAP ASE server. Get the information in [Retrieving Server Information \[page 8\]](#).
- The resource size should be determined appropriately. See [Determining Resource Size \[page 16\]](#).

Parent topic: [Migration \[page 60\]](#)

Next: [Restoring Non-Default Configurations \[page 61\]](#)

5.2 Restoring Non-Default Configurations

Restore the non-default configurations collected by running the `sp_serverinfo_report` stored procedure.

For example, if you need to enable the compression feature, you can run the `sp_configure 'enable compression', 1` command.

Parent topic: [Migration \[page 60\]](#)

Previous: [Provisioning an SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Instance \[page 60\]](#)

Next task: [Creating Data Caches \[page 61\]](#)

5.3 Creating Data Caches

You should create the same data caches in the cloud as in the on-premises database before loading the database.

Procedure

1. Run the `sp_helpcache` command to check whether the data caches bound by the database and its tables and indexes exist in the cloud. For example:

```
sp_helpcache
go
Cache Name          Config Size  Run Size  Overhead  Cache Type
-----
default data cache 7168.00 Mb  7168.00 Mb  144.58 Mb Default
Memory Available For      Memory Configured
Named Caches              To Named Caches
-----
7168.00 Mb              7168.00 Mb
-----
Cache Binding Information: -----
Cache Name      Entity Name      Type      Index Name      Status
-----
(return status = 0)
```

Note that there is only a default data cache in the cloud.

2. Compare the cache information from the cloud service with that of the on-premises server, and create the same data caches in the cloud with the script you obtained in the pre-migration phase.

For example, we notice that there are three data caches used by the `testdb` database and its tables and indexes in the on-premises server before migration. Now create the three data caches with the script generated by the `DDLgen` utility previously as follows:

```
-----
-- DDL for Cache 'testdb_cache'
-----
print '<<<< CREATING Cache - "testdb_cache" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name = 'testdb_cache'
AND parent = 19)
    exec sp_cacheconfig 'testdb_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testdb_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testdb_cache' , '10240K' , mixed
go
-----
-- DDL for Cache 'testidx_cache'
-----
print '<<<< CREATING Cache - "testidx_cache" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name =
'testidx_cache' AND parent = 19)
    exec sp_cacheconfig 'testidx_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testidx_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testidx_cache' , '30720K' , mixed
go
exec sp_poolconfig 'testidx_cache' , '10240K' , '32K'
go
-----
-- DDL for Cache 'testtab_cache'
-----
print '<<<< CREATING Cache - "testtab_cache" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.sysconfigures WHERE name =
'testtab_cache' AND parent = 19)
    exec sp_cacheconfig 'testtab_cache', '0'
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Cache testtab_cache'
    SELECT syb_quit()
END
go
exec sp_cacheconfig 'testtab_cache' , '20480K' , mixed
go
-----
```

3. Use `sp_helpcache` to verify whether all the needed caches are in the cloud.

For example, the three data caches created in the previous step looks like this:

```
sp_helpcache
go
Cache Name          Config Size  Run Size    Overhead    Cache Type
-----
```

default data cache	7168.00 Mb	7168.00 Mb	144.58 Mb	Default
testdb_cache	10.00 Mb	10.00 Mb	0.30 Mb	Mixed
testidx_cache	30.00 Mb	30.00 Mb	0.69 Mb	Mixed
testtab_cache	20.00 Mb	20.00 Mb	0.49 Mb	Mixed
Memory Available For	Memory Configured			
Named Caches	To Named Caches			
-----		-----		
7228.00 Mb		7228.00 Mb		
----- Cache Binding Information: -----				
Cache Name	Entity Name	Type	Index Name	Status
-----	-----	----	-----	-----
(return status = 0)				

Task overview: [Migration \[page 60\]](#)

Previous: [Restoring Non-Default Configurations \[page 61\]](#)

Next: [Migrating the Database Encryption Key \[page 63\]](#)

5.4 Migrating the Database Encryption Key

You can use the `transfer encryption key` command to transfer the database encryption key (DEK) from your on-premises database to the SAP ASE instance in the IBM Cloud.

DEK is used to encrypt the database in SAP ASE and it is encrypted with the master key (KEK) of the master database which in turn is encrypted with either a password or an HSM key. The dump of an encrypted database remains encrypted. In order to load the database to the cloud, the database in the cloud should be encrypted with the same DEK as that of the database dump. Therefore, migrate the DEK from the on-premises server to the cloud and create the target encrypted database with that DEK before loading the dump.

Parent topic: [Migration \[page 60\]](#)

Previous task: [Creating Data Caches \[page 61\]](#)

Next task: [Creating a Database \[page 64\]](#)

5.4.1 Migrating the DEK with transfer encryption key

Use the SAP ASE Cloud Portal to migrate the DEK from on-premises to the cloud.

Context

The recommended version of SAP ASE is 16.0 SP04 PL06 HF1 or above.

Procedure

1. Export the DEK encrypted with your user password from the on-premises SAP ASE by using the `transfer encryption key` command. For example:

```
transfer encryption key dek for database encryption with password 'mysecret9'
go
0x0100011000001000001080000000000000000000000001001bc572fba81246b2934a7541e675cf06
7e54396202e90604c540184f3f8cc595e17d73ccf4299fb7c6e635752f31002d1f76900c13175c
e5
```

2. Log in to the SAP ASE Cloud Portal.
3. On the Database Encryption Keys page, click [Upload a DEK](#).
4. In the Upload a DEK window, enter a name for the DEK, for example, dek1.
5. Enter the same password you used during the export in step 1 and paste the encryption key generated from the same step into the relevant field.
6. Click [Confirm](#) to upload the key.
7. In the Customer Encryption Keys list, locate the encryption key you just uploaded and click the [Transfer](#) button at the right-hand side of the list to install the encryption key on your SAP ASE instance.
8. In the Install DEK window, select the name of the SAP ASE instance to install the DEK and click [Confirm](#).
9. If your SAP ASE instance is part of an Always On/HADR cluster, repeat steps seven and eight for each SAP ASE in the cluster.

5.5 Creating a Database

Create a database in the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud before you upload the dump file.

Context

To successfully migrate a database from an on-premises SAP ASE to an SAP ASE in IBM Cloud, ensure that the SAP ASE database you created in IBM Cloud:

- has a device size equal to or larger than the on-premises database
- contains the same database options as the on-premises database

Procedure

1. Check the device sizes of the on-premises database.

Check the total sizes of the data-only devices, the log-only devices, and the mixed devices in the on-premises database by referring to the **device fragments** section in the retrieved database information. See [Retrieving Database Information \[page 12\]](#) for the details on how to retrieve database information.

2. In AMC, on the [Devices Overview](#) screen, click the [Create](#) button to create a data device and a log device for your database. The size of the data device and the log device should not be smaller than the size of the data device and the log device (or plus the size of the mixed device, if any) of the on-premises database. You can also create a mixed device if necessary.
3. On the [Database Overview](#) screen, click the [Create](#) button to create a database in IBM Cloud. When you create the database, ensure that you select the database encryption key that you migrated from your on-premises environment in [Migrating the Database Encryption Key \[page 63\]](#).
4. Set all database options to match the options set in the on-premises database. In this example, set the `select into/bulkcopy/pllsort` and `allow nulls by default` options that you retrieved with the `load database...with listonly=create_sql` command in the pre-migration phase.

```
sp_dboption <database_name>, 'select into/bulkcopy/pllsort', true
go
master..sp_dboption <database_name>, 'allow nulls by default', true
go
```

Task overview: [Migration \[page 60\]](#)

Previous: [Migrating the Database Encryption Key \[page 63\]](#)

Next: [Synchronizing Login Accounts and Roles \[page 65\]](#)

5.6 Synchronizing Login Accounts and Roles

Create the same login accounts and roles in the ASE instance as in the on-premises server.

Refer to the login accounts and roles you obtained in the pre-migration phase to recreate them in the cloud.

Parent topic: [Migration \[page 60\]](#)

Previous task: [Creating a Database \[page 64\]](#)

Next task: [Uploading Dump Files \[page 69\]](#)

5.6.1 Creating Login Accounts

Create the same login accounts in the cloud service as in the on-premises SAP ASE.

Context

This topic assumes that there's no other login account in the cloud except the system reserved login accounts before migration.

Procedure

Create the same login accounts in the cloud service.

For example, in the cloud side, create two logins ASELOGIN01 with `suid = 3` and ASELOGIN02 with `suid = 4`.

```
1> use master
2> go
1> create login ASELOGIN01 with password sybase123 suid 3
2> go
1> create login ASELOGIN02 with password sybase123 suid 4
2> go
```

If you have generated the DDLs of logins by running the `ddlgen` utility, you can also run the DDLs script to create logins directly.

After creating login ASELOGIN01 and ASELOGIN02, the output of `syslogins` and `sysusers` looks like this:

```
1> select * from syslogins
2> go
suid      status  acctdate      totcpu      totio
spacelimit  timelimit  resultlimit  dbname      name
password
language  pwdate      audflags  fullname  srvname  logincount  procid
lastlogindate  crdate      locksuid  lockreason  lockdate
lpid  crsuid
-----
-----
-----
-----
-----
-----
1          0          Jun 22 2020  1:58AM          438
243955      0          0          0 master      sa
0xc007e18ecc53461276067497e855fa4a3cf6408bfed92646053a10ef1e2eb4adb136c310b4b30b9
a3b21      NULL Jun 22 2020  1:58AM          0      NULL      NULL      NULL      NULL
Jun 28 2020  7:37PM Jun 22 2020  1:58AM          NULL      NULL
NULL      -1      NULL
2          0          Jun 22 2020  2:06AM          0
0          0          0          0 sybsystemdb  probe
0xc00778c29b6365c4bde856944873429c5c2ac0f9e9c3174bacda2c66ab71a91995c6399600cae7b
f2228      NULL Jun 22 2020  2:06AM          0      NULL      NULL      0
NULL      NULL      NULL Jun 22 2020  2:06AM          NULL      NULL      NULL Jun 22 2020
2:06AM      -1      NULL
```

```

0          3      0          Jun 28 2020 11:09PM          0
0          0          0          0 master ASELOGIN01
0xc007bd81493fb0d34ff6e161d09f296221274758355992acf84d9dce3bdcf3da5e3327732c062a8
bbe42      NULL Jun 28 2020 11:09PM          0      NULL      NULL      NULL
NULL
NULL NULL          1          NULL Jun 28 2020 11:09PM          NULL      NULL
NULL NULL          4      0          Jun 28 2020 11:09PM          0
0          0          0          0 master ASELOGIN02
0xc00727732c062a8bbe42a3851ee3dc159fc0dbdd0ef14e4fbaf987ae2ff10de9064193a569ac434
3b914      NULL Jun 28 2020 11:09PM          0      NULL      NULL      NULL
NULL
NULL NULL          1          NULL Jun 28 2020 11:09PM          NULL      NULL
NULL NULL          1
(4 rows affected)

```

5.6.2 Creating User-Defined Roles

The topic shows you how to create the same roles in the ASE service as the on-premises server.

Context

If the required user-defined roles do not exist in the cloud service, you will see error messages like the following ones when loading the database. It is recommended to create the roles before loading the database:

User defined role 'usrdef_role01' used in database 'tdb2' does not exist in this server. Use CREATE ROLE to create this role after the database online.
 User defined role 'usrdef_role02' used in database 'tdb2' does not exist in this server. Use CREATE ROLE to create this role after the database online.

In the cloud service, the srid of each newly created user-defined role doesn't need to match the srid in the on-premises side. You only need to create the user-defined roles with the same names.

The roles existing in the on-premises server before migration are usrdef_role01 and usrdef_role02 according to the example in the [Obtaining User-Defined Roles \[page 51\]](#) topic.

Procedure

1. Check whether the needed user-defined roles—usrdef_role01 and usrdef_role02—exist in the cloud side.

```

1> use master
2> go
1> select * from sysrvroles where srid > 32
2> go
srid      name
password

          pwdate      status logincount
locksuid  lockreason lockdate
-----
-----
-----

```

```

-----
-----
-----
33 usrdef_role03
0xc007c4be5cff6c9d8f03a0ccb3b3b47537f319c0d06f3d4627b432822c5ebb13f1de75bc19ef
c136539e

0          NULL          NULL          Jul 21 2020   6:35PM          0
(1 rows affected)

```

The output shows that these two roles do not exist in the cloud service.

2. Create the user-defined roles for the migrated database in the cloud service if these roles do not exist.

In this example, create roles `usrdef_role01` and `usrdef_role02` as follows:

```

1> use master
2> go
1> create role usrdef_role01 with passwd '<password>'
2> go
1> create role usrdef_role02 with passwd '<password>'
2> go

```

If you have generated the DDLs of roles by running the `ddlgen` utility, you can also run the DDLs script to create roles directly. For example, run the following script you created in the pre-migration phase:

```

-----
-- DDL for Role 'usrdef_role01'
-----
print '<<<< CREATING Role - "usrdef_role01" >>>>'
go
use master
go
IF EXISTS (SELECT 1 FROM master.dbo.sysrvroles where name = 'usrdef_role01')
    DROP ROLE usrdef_role01 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role01'
    SELECT syb_quit()
END
go
create role usrdef_role01 with encrypted password
0xc00775bc19efc136539ef52ac99d9460c4381a4d58b7e3f0ed9712f7599080b9d3cd86de0d53
05bc857e
go
-----
-- DDL for Role 'usrdef_role02'
-----
print '<<<< CREATING Role - "usrdef_role02" >>>>'
go
IF EXISTS (SELECT 1 FROM master.dbo.sysrvroles where name = 'usrdef_role02')
    DROP ROLE usrdef_role02 with override
IF (@@error != 0)
BEGIN
    PRINT 'Error dropping Role usrdef_role02'
    SELECT syb_quit()
END
go
create role usrdef_role02 with encrypted password
0xc00786de0d5305bc857e8c2ffe8ec7752d2bb242d1c62b1aa2fa9603870bc7341a5aca23045b
f9f73555
go

```

After creating the roles `usrdef_role01` and `usrdef_role02`, the output of `sysssrvroles` looks like this:

```
1> select * from sysssrvroles where srid> 32
2> go
srid          name
password

locksuid      lockreason    pwdate      status logincount
lockdate
-----
33 usrdef_role03
0xc007c4be5cff6c9d8f03a0ccb3b3b47537f319c0d06f3d4627b432822c5ebb13f1de75bc19ef
c136539e
0          NULL          NULL          Jul 21 2020  6:35PM          0
34 usrdef_role01
0xc00775bc19efc136539ef52ac99d9460c4381a4d58b7e3f0ed9712f7599080b9d3cd86de0d53
05bc857e
0          NULL          NULL          Jul 21 2020  6:39PM          0
35 usrdef_role02
0xc00786de0d5305bc857e8c2ffe8ec7752d2bb242d1c62b1aa2fa9603870bc7341a5aca23045b
f9f73555
0          NULL          NULL          Jul 21 2020  6:39PM          0
(3 rows affected)
```

5.7 Uploading Dump Files

After you have generated the dump files for your database, upload the dump files to the instance specific backup bucket in IBM Cloud Object Store (COS).

Prerequisites

- Install Python 3.
- Install IBM Aspera Transfer SDK.
- Obtain access to the SAP ASE common services IBM Schematics Workspace.
- Obtain an IBM Service ID API Key with the following access policies:
 - Schematics: Service Access: Manager
 - Secrets Manager: Service Access: SecretsReader
- Generate and download the `upload_to_cos.py` script from the [About Uploading Backup Items](#) page in SAP ASE Cloud Portal.

Procedure

1. Open the downloaded IBM Aspera Transfer SDK folder in your IDE.
2. Move the `upload_to_cos.py` script to the `/connectors/python` folder in the IBM Aspera SDK.
3. Move the dump files that you want to upload to COS into the `/connectors/python` directory in the IBM Aspera SDK.
4. Navigate to the python directory in the SDK and then run the following command to install all dependencies:

```
pip3 install -r requirements.txt
```

5. Update the following variable on line 94 of the `upload_to_cos.py` script:
FILE_NAME: Specify the name of the dump file (including extension) to upload to COS. The file must exist in the `/connectors/python` directory.
6. In the root directory of the IBM Aspera SDK folder, start the Aspera Transfer SDK daemon:

```
./bin/asperatransferd
```

7. In a separate terminal window, after the daemon is running, navigate to the python directory and run the following command:

```
python3 upload_to_cos.py
```

Results

Your dump files are successfully uploaded to COS.

Task overview: [Migration \[page 60\]](#)

Previous: [Synchronizing Login Accounts and Roles \[page 65\]](#)

Next task: [Loading the Database and Transactions \[page 70\]](#)

5.8 Loading the Database and Transactions

Use SAP ASE Cloud Portal to restore the database dump files to the cloud.

Prerequisites

- You have uploaded the dump files for the database to the instance specific backup bucket in IBM Cloud Object Store. See more details in [Uploading Dump Files \[page 69\]](#).

- You have created the SAP Adaptive Server Enterprise, cloud edition by IBM Cloud instance. See more details in [Provisioning an SAP Adaptive Server Enterprise, Cloud Edition by IBM Cloud Instance \[page 60\]](#).
- You have migrated the database encryption key (DEK) to the SAP ASE instance on IBM Cloud. See more details in [Migrating the DEK with transfer encryption key \[page 63\]](#).
- You have created the migrated database on IBM Cloud. You have also encrypted the database with the migrated DEK and restored the database options. See more details in [Creating a Database \[page 64\]](#).
- You have restored database option for the migrated database.
- You have the `load any database` privilege in the SAP ASE instance.

Procedure

1. Log in to the SAP ASE Cloud Portal.
2. Click the Backups tab of the SAP ASE instance and locate the dump file you uploaded.
3. Click the [Restore](#) button at the right-hand side of the row to load the database and transactions.
4. In the [Restore <dump_file>.dmp](#) window, enter the database name to be restored and select the restore DEK you migrated to the cloud.
5. Click [Confirm](#) to initiate restore.
6. You can check the restore status on the Database Restores tab.

Task overview: [Migration \[page 60\]](#)

Previous task: [Uploading Dump Files \[page 69\]](#)

Next task: [Bringing the Database Online \[page 71\]](#)

5.9 Bringing the Database Online

You should bring the database online after loading the database and transaction dumps successfully.

Procedure

Use the `online database` command, for example:

```
online database testdb
```

If you see the following message:

Database 'testdb' is now online.

The migrated database is ready for use.

Task overview: [Migration \[page 60\]](#)

Previous task: [Loading the Database and Transactions \[page 70\]](#)

Next: [Restoring Cache Bindings \[page 72\]](#)

5.10 Restoring Cache Bindings

You need to restore cache bindings after the database is online.

1. Run the following script to bind data caches to the migrated database.

```
1> sp_bindcache testdb_cache, testdb
2> go
(return status = 0)
1> sp_helpcache
2> go
```

Cache Name	Config Size	Run Size	Overhead	Cache Type
default data cache	7168.00 Mb	7168.00 Mb	144.58 Mb	Default
testdb_cache	10.00 Mb	10.00 Mb	0.30 Mb	Mixed
testidx_cache	30.00 Mb	30.00 Mb	0.69 Mb	Mixed
testtab_cache	20.00 Mb	20.00 Mb	0.49 Mb	Mixed

Memory Available For Named Caches Memory Configured To Named Caches

7228.00 Mb 7228.00 Mb

----- Cache Binding Information: -----

Cache Name	Entity Name	Type	Index Name	
testdb_cache	testdb	database		V
testidx_cache	testdb.dbo.rrtab	index	idx1	V
testtab_cache	testdb.dbo.rrtab	table		V

```
(return status = 0)
```

2. Check the binding information in the database information report.
You can find cache-related information of the target database in the database information report and compare it with the information you obtained before the migration to double-check whether all cache configuration and cache binding settings are restored in the cloud.

Parent topic: [Migration \[page 60\]](#)

Previous task: [Bringing the Database Online \[page 71\]](#)

6 Post-Migration

This section describes the actions you need to take after migrating the database.

It is recommended to perform some testing or tuning before switchover.

Validation

- Server and Database Configurations
You can run `sp_serverinfo_report` and `sp_dbinfo_report` to generate a server information report and a database information report. Compare the two reports with the reports against your on-premises system to see whether the following configurations are restored in the cloud:

- Non-default configurations
- Database options
- Database cache bindings
- Login accounts and roles
 - Check the user information of the target database in the cloud service.
For example, the user information of database `tdb1` is as follows:

```
1> use tdb1
2> go
1> select * from sysusers where suid > 0
2> go
suid      uid      gid      name      environ user_status
-----
          1          1      0 dbo      NULL          0
          3          4      0 ASELOGIN01 NULL          0
          4          5      0 ASELOGIN02 NULL          0
(3 rows affected)
1> sp_helpuser
2> go
Users_name      ID_in_db      Group_name      Login_name
-----
ASELOGIN01          4      public      ASELOGIN01
ASELOGIN02          5      public      ASELOGIN02
dbo              1      public      sa
usedb_user        3      public      NULL
(1 row affected)
(return status = 0)
```

- Check whether all the user-defined roles exist in the cloud after migration.
For example, check the roles `usrdef_role01` and `usrdef_role02` in the database `tdb2`.

```
1> use tdb2
2> go
1> select * from sysroles where id > 32
2> go
id      lrid      type      status
-----
      34      16401      0          0
      35      16402      0          0
```

```

(2 rows affected)
1> select su.suid, su.uid, su.gid, su.name, ssr.srid
2> from sysusers su, sysroles sr, master.dbo.syssrvroles ssr
3> where su.uid = sr.lrid
4>       and sr.id = ssr.srid
5>       and ssr.srid > 32
6> go
  suid      uid      gid
name
-----
usrdef_role01  -2      16401      16401
usrdef_role02  -2      16402      16402
(2 rows affected)
1> sp_helprotect usrdef_role01
2> go
 grantor      grantee      type      action
object      column      predicate      grantable
-----
dbo          usrdef_role01      Grant      Delete
dbo.tb1      All          NULL          FALSE
dbo          usrdef_role01      Grant      Insert
dbo.tb1      All          NULL          FALSE
dbo          usrdef_role01      Grant      Update
dbo.tb1      All          NULL          FALSE
(1 row affected)
(return status = 0)
1> sp_helprotect usrdef_role02
2> go
 grantor      grantee      type      action
object      column      predicate      grantable
-----
dbo          usrdef_role02      Grant      Delete
dbo.tb1      All          NULL          FALSE
dbo          usrdef_role02      Grant      Insert
dbo.tb1      All          NULL          FALSE
dbo          usrdef_role02      Grant      Update
dbo.tb1      All          NULL          FALSE
(1 row affected)
(return status = 0)

```

Note

Since the `sysusers.uid` of each database matches the `master..syssrvroles.name` but not `master..syssrvroles.srid`, therefore, you do not need to update `master..syssrvroles.srid` for newly created user-defined roles in the cloud.

- Consistency Check
You can use `dbcc` commands to check whether a database is in consistent state, for example, `dbcc checkdb` or `dbcc checkalloc`.
If you have run these commands in the on-premises server before the migration and have the reports archived, you can compare the reports to validate whether the data is migrated as expected.

Functional Test

You can set up a test environment to perform some functional tests to make sure that the system is working as expected.

Performance Tuning

Since the ASE service is not exactly as the on-premises system, you need to verify the performance of their whole system before bringing the system online. If there is any gap, tune the performance in the cloud. See the *SAP Adaptive Server Enterprise Performance and Tuning Series* guides for details.

Note

The operations described in this topic are beyond the service offered by SAP for database migration. We only recommend conducting the operations yourself for a better performance in your migrated database.

7 Troubleshooting

This topic focuses on troubleshooting topics for migration from on-premises SAP ASE to SAP Adaptive Server Enterprise, cloud edition by IBM Cloud.

SAP ASE `dump` and `load` commands are used in the database migration. If errors occur when loading the database, see *SAP Adaptive Server Enterprise Troubleshooting Guide* for troubleshooting errors or warnings.

To migrate a database from on-premises SAP ASE to the cloud edition successfully, you need to make sure the following criteria are met when dumping the database:

- The version of the on-premises SAP ASE is equal to or higher than 16.0 SP04.
- The charset and the sort order of the on-premises SAP ASE are the same as in the cloud edition.
- The page size of the on-premises SAP ASE is the same as that of the cloud edition.
- The DEK of the on-premises SAP ASE is the same as that of the cloud edition.
- The database size in the cloud edition is not smaller than that in the on-premises SAP ASE.
- The feature set of the on-premises SAP ASE and the migrated database(s) are compatible with the ASE service.

You are also recommended to make sure vCPU and memory in the cloud edition are equivalent or better than that in the on-premises SAP ASE. Otherwise there might be a performance degradation or resource-related errors reported.

To check the specification of vCPU and memory in the cloud instance and compare to those to the on-premises instance, use `sp_serverinfo_report` and `sp_dbinfo_report <database name>` to generate reports in both the on-premises and the cloud sides.

This topic helps you to troubleshoot issues in both the on-premises side and the cloud side.

On-Premises Troubleshooting

In the on-premises SAP ASE, the following two stored procedures need to be executed to collect the server-level and database-level information:

- `sp_serverinfo_report`
- `sp_dbinfo_report`

Check if there are any warnings or error messages raised by these stored procedures.

If there are no errors at the server level, you get the following message:

`Collecting ASE server information completed with no error.`

If warnings or error messages are raised, refer to [Retrieving Server Information \[page 8\]](#) for the steps to resolve the problem.

If there are no errors at the database level, you get the following message:

`Collecting database information completed with no error.`

If warnings or error messages are raised, refer to [Retrieving Database Information \[page 12\]](#) for the steps to resolve the problem.

You also need to run the stored procedure `sp_featurecompat` to check all the deprecated features and disable them all. See [Checking for Incompatible Features \[page 20\]](#) for details.

Cloud Troubleshooting

In the cloud edition, warnings or error messages may be raised when executing the `load` command. The following message is for migration:

Error 9616

Database '%.*s' is at an older version '%d' than ASE server version '%d'

Severity: 16.

Explanation: Migration from on-premises to cloud only supports migrating databases equal to or higher than SAP ASE 16.0 SP04. If the database dump file to be loaded comes from a lower version, you see the following message.

Action: Upgrade the source database before migration. For how to upgrade the databases, refer to *SAP Adaptive Server Enterprise Installation and Upgrade Guide*.

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 an 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.

© 2026 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.

