



Essbase on OCI: Backup and Restore

This paper covers disaster recovery and version upgrade planning for Essbase instances deployed on OCI .

June 2020 | Version 1.2
Copyright © 2020, Oracle and/or its affiliates
Confidential - Public

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

Disclaimer	1
Purpose Statement	3
Backup and restore terminology and Use cases	3
Defining Restore Use Cases	3
Understanding Essbase Backups	4
Use case: Disaster Recovery	5
Default Essbase Deployment: Recover Compute	5
Non-Default Essbase Deployment: Recover Compute	10
Essbase Backups: Step-by-step	14
Backup Entire ATP and Block Volumes: Step-by-Step	14
Backing up Essbase using Oracle-Provided Scripts	14
Backing up Essbase using DB client and OCI console	14
Backup ATP Schema and Block Volumes: Step-by-Step	16
Before You Begin: Install and Configure Oracle Instant Client and Tools	16
Backing up Essbase schema using Data Pump and OCI console	18
Various other notes	20
Managing Volume Groups	20
Resolving Network Components	20

PURPOSE STATEMENT

This document provides recommendations for the life cycle management of Essbase deployments on Oracle Cloud Infrastructure (OCI). We'll review use cases and steps for full Essbase instance backup in support of disaster recovery and version upgrades. Essbase deployments require regular backup and healthy redundancy planning. For deployments on OCI, block volumes and relational database schemas must be thoughtfully managed.

Application-level backups within an Essbase service also have a very important place in the overall Essbase life cycle management, but they are not the focus of this paper. Here, we are focused on recovering/upgrading the entire Essbase instance (all applications) as of some point-in-time.

Task-level steps described in this whitepaper are as of Essbase version 19.3.0.2.3. Should a step not work exactly as described, please consult the Essbase documentation.

BACKUP AND RESTORE TERMINOLOGY AND USE CASES

When you use the Oracle Marketplace to deploy Essbase on OCI, you don't *install* it in the traditional sense, as you would to an on-premise server machine. Instead, you deploy Essbase as part of an image on a cloud compute node (or a 'compute' for short). This compute, together with the other necessary OCI components (VCN, relational database, block and object storage, etc.), is called an Essbase stack.

Certain components from every stack you create contain information that makes your Essbase deployment unique to your end user community. You will need to back up these unique stack components at appropriate intervals to meet your recovery objectives. In the event of an OCI compute or Availability Domain (AD) failure, you can recover your Essbase instance by building a new stack and restoring into it your Essbase backup. Similarly, you can perform necessary Essbase version updates using backups.

Restore, in this paper, will be a generic term used to describe the action of deploying a new Essbase stack and attaching to it the appropriate block volume and relational database backups.

If your restored Essbase stack should be the latest release, you will use Oracle Marketplace to create it.

If you wish to recover the same version of Essbase you ran prior to restoring, then you can take an older deployment script from [GITHUB](#).

Important: Think of the pre-restore Essbase stack as a **source** (of block volumes, block volume backups, RDB schemas, RDB backups) and the post-restore Essbase stack as a **target**. **Your restored target instance should reflect the source instance as of some point in time.**

Defining Restore Use Cases

The restore requirements for your Essbase deployment will define the required backup tasks. There are two common restore scenarios:

Disaster Recovery: should you experience a hardware failure with the compute node in which Essbase is running, you can easily recover your instance in any AD in the same region.

Version Upgrade: between major versions, Oracle will require out-of-place Essbase binary/configuration uptake, esp. when Fusion Middleware changes or Essbase schema changes are being introduced.

Although not every restore use cases is driven by disaster recovery, restoring Essbase always requires the deployment of a new *target* stack [for the sake of recovering or updating the compute image]. When defining the *target* stack, you can re-use some *source* stack components, along with your *source* backups.

Although not highlighted in this paper, service level agreements (SLAs) with the Essbase user community are an important consideration for restore planning. What is the acceptable downtime window during any recovery or update? There are things you can do to manage the overall time it takes to restore an instance, such as cloning block volumes and building target stacks in advance so that the *target* OCI environment is in passive standby.

Note: Detailed steps for the **version upgrade** restore use case will be published with the first major release after 19.3.x.x. The remainder of this paper will focus on **disaster recovery**.

Understanding Essbase Backups

Suitable and reliable backups are the backbone of your restore use case. Any backup of Essbase on OCI will depend on some details of your Essbase stack. A complete backup must protect all information that makes your Essbase deployment unique. These items are:

- Relational database schemas for every Essbase stack, which store some application, user and configuration information.
 - i) A single database schema for Essbase, called `<instance prefix>_Essbase`
 - ii) Eight database schemas for WebLogic, with the same `<instance prefix>_<schemaname>`
- Essbase application and database information stored on a block volume mounted as `/u01/data`.
- WebLogic domain and configuration information stored on a block volume mounted as `/u01/config`. [Essbase is a managed service within a WebLogic domain.]

Your backup strategy must ensure that the information above is captured at appropriate intervals to align with your restore use case. Also, the timing of relational database backups and Essbase block volume backups must be synchronized to ensure that the *target* instance will have the required consistency for Essbase to run properly.

Consider this... Although it is possible to restore *source* block volumes directly to a *target* stack, you may not know the status of any write operations that were in process in the case of compute disaster recovery. Stopping Essbase before you take a backup allows you to know the exact state of your Essbase data when you use it to restore an Essbase instance.

Block Volume Backup Planning for Your Restore Use Case

When you restore Essbase in a *target* stack, you will always restore and attach a *source* data block volume backup (and in some cases, a config block volume backup). To consistently restore a *target* Essbase instance, backups of your *source* Essbase data block volume

- should be from the same point-in-time as your *source* ATP database backup.
- should be taken at a time when the Essbase services were stopped.

Optionally, you can work directly with the *source* block volumes if you are sure about the Essbase data they contain.

Note: Block volumes are restored from backup into an AD of your choosing. Clones can be introduced into your backup processes to minimize Essbase downtime. Block volumes (or clones) cannot be moved between ADs. Clones can be used to quickly replicate a *source* block volume, then can be backed up and likewise restored into an AD of your choosing. When clones are being used, Essbase must only be stopped during the cloning process.

Database Backup Planning For Your Restore Use Case

When you deploy an Essbase stack, the default behavior is to deploy a new ATP database as part of the stack. If you deploy a second Essbase stack in the same region, you can optionally choose to deploy without a second ATP database by selecting the “use existing database” option. This results in license cost savings, as a single ATP database is scalable enough to support schemas for multiple Essbase instances. Schemas associated with each instance of Essbase have a common and unique prefix. The number of Essbase instances deployed in your ATP database, combined with your recovery use case, defines whether you can back up the entire database or if you need to perform schema-level backups.

Relational Database Backups for Default Essbase Deployments

Default Essbase deployments consist of a single Essbase stack with RCU schemas deployed into a single relational database. For Essbase default deployments, the relational database can be backed up without exporting schemas for individual Essbase instances. This is possible because restoring the database only impacts one Essbase instance.

Relational Database Backups for Non-Default Essbase Deployments

Non-default Essbase deployments consist of multiple Essbase stacks, each with unique RCU schemas, deployed into a single relational database. Essbase non-default deployments require you to refine your relational database backup strategy. Restoration of a relational database backup is “all or nothing”, meaning that the relational database schemas for all Essbase instances would be reset to the same timestamp should a restore action be required on the database.

Best practice for Essbase backups is to synchronize the relational database backup with the block volume backups. But it is not a best practice to require all Essbase instances to be on the same backup and restore schedule. For non-default deployments, isolation of the data and metadata for a single Essbase instance is accomplished by using database schema backups.

Restore Use Case	Essbase Deployment Type	RDB Backup Type	Block Volume(s) to Recover
Disaster Recovery	Default	Database	Data and Config
Disaster Recovery	Non-Default	Schema	Data

Next, we will look at disaster recovery and the backup steps required to support it. Some assumptions have been made to limit the size and scope of this paper. All examples in this paper assume the following:

- ATP is the relational database into which Essbase schemas are deployed
- Oracle Identity Cloud Service (IDCS) is the security provider for the Essbase deployment
- The ESSBASE SYSTEM ADMIN USER NAME (stored in WebLogic – only non-IDCS user in the system) is the same between the *source* and the *target* Essbase stack
- The *source* instance backup has [at the time of restore] at least one valid IDCS user with an Essbase system administrator role.

USE CASE: DISASTER RECOVERY

If the compute node on which your Essbase services are deployed fails, you can recover it in any AD in the same region. When you build the *target* stack and recover the compute, you can re-use many parts of the *source* stack by using “use existing database” and “use existing network” options. If you have a load balancer in the *source* deployment, you can change its backend networking and use it with the *target* stack. You can also attach the block volumes (or their backup/clone) that were previously attached to the *source* compute.

When restoring your Essbase environment, do not be confused by the term “use existing database”. Although you are free to use a pre-existing relational database, every new Essbase stack will deploy a unique set of Essbase schemas. Deploying into an existing database will not “hook” your *target* Essbase services to your *source* Essbase schemas – rather, it will create new *target* schemas. Mapping *source* schema data and metadata into the new *target* database is accomplished when you restore your *source* database backup into the *target*.

Remember that a failed compute node will cause failure of your Essbase services. The backups referenced as part of this backup and restore discussion must be taken at regular intervals in advance of any compute failure.

Default Essbase Deployment: Recover Compute

Default deployments can be restored from backup using the OCI console.

Backups should be taken prior to compute failure following one of the methods here: [Backup Entire ATP and Block Volumes: Step-by-Step](#). It doesn’t matter whether you choose to backup using Oracle-provided scripts or using a database client and the OCI console.

1) Deploy a *target* Essbase stack using Oracle Marketplace

- Use the *source* IDCS confidential application
- Use the *source* ATP database and password
- Use the *source* VCN and application subnet
- If your *source* stack has a load balancer, do not deploy a *target* Load Balancer– you can change the Backend Set after deploying the *target* stack
- If your *source* stack has a bastion host, deploy a bastion host with the *target* stack [*source* bastion host can be deleted after successful recovery]
- Use the same Essbase system admin user name and password in the *target* as you used in the *source*
- Use the same IDCS Essbase admin user in the *target* stack as you used in the *source* stack. If this is not possible, make sure the *source* Essbase instance has at least one valid IDCS user with the Essbase system administrator role. After you restore, you must login to the *target* instance as a valid IDCS user who had Essbase system administrator role on the *source* instance.

The Essbase services in the *target* stack will be started automatically by the Oracle Resource Manager (ORM) apply job.

Note: After deploying this stack, two sets of Essbase schemas will be present in your ATP (*source* and *target*). Remember that your ATP backup, taken prior to this stack deployment contains only one set of Essbase schemas (*source*).

2) Stop the *target* Essbase services. [Do not stop the Essbase compute in OCI.]

Hint: It goes without saying that the Essbase services on the *source* compute are stopped, because this use case involves *source* compute hardware failure. If you are simulating these steps, make sure you also stop the *source* compute's Essbase services.

- ssh to the *target* Essbase compute as opc user
- execute the following command: `sudo systemctl stop essbase.service`
<https://www.oracle.com/pls/topic/lookup?ctx=en/database/other-databases/essbase/19.3&id=GUID-E442392E-04F3-4E0D-AA98-868B57EFB907>

3) Restore the *source* ATP from Backup.

Be sure to select a *source* ATP backup that was taken during a time when your *source* Essbase services were stopped; also, be sure that you have *source* Essbase 'data' and 'config' block volumes from the same time.

- Using the OCI console main menu, choose [Autonomous Transaction Processing](#)
- Select the *source* ATP (note: *source* and *target* schemas are in the *source* ATP)
- Under [Resources](#), select [Backups](#)
- Find the timestamp of the Backup you want to restore, and select [Restore](#) using the menu on the right side of the page

Important: Before restoring your ATP for the first time, use the following link and **carefully review steps 4 and 5**. After you complete a restoration of your ATP, you may render other backups invalid, depending on their timestamp.

<https://www.oracle.com/pls/topic/lookup?ctx=en/cloud/paas/atp-cloud&id=GUID-78C28C41-AA87-4AD7-BEB6-693235C73F3C>

- After the restore finishes, audit your data using a database client like SQL Developer and then be sure to Stop and Start the *source* ATP, which will be in a read only state after a successful restore. You can look at the `ESSBASE_APPLICATION` table within the `<targetprefix>_ESSBASE` schema to verify the restored applications.

4) Detach 'data' and 'config' block volumes from the *target* Essbase compute.

- As opc user, ssh to your *target* compute and unmount the `/u01/data` volume
`sudo umount /u01/data`
- Using the OCI console main menu, choose [Compute->Instances](#)
- Select the *target* Compute in which your Essbase instance is deployed
- In [Compute Details->Resources](#), select [Attached Block Volumes](#)
- Using the actions menu to the right of the 'data' volume, select [Detach](#)
- ssh to the *target* Compute as opc user and issue the detach ISCSI commands as indicated.
- Select [Continue Detachment](#) in OCI after issuing ISCSI detach commands.
- Repeat for the *target* 'config' block volume

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/detachingavolume.htm>

5) Create 'data' and 'config' block volumes from *source* block volume backup.

Block volume backups are stored in object storage; they are not able to be attached to compute instances without first creating a new volume in block storage. Clones, on the other hand, are standalone block volumes and can be directly attached to a compute instance. Skip this step if you cloned your block volumes. You may have backed up the block volume group; if so, you can create the volume group.

- Using the OCI console main menu, choose [Block Storage->Block Volume Backups](#)
- Find the desired 'data' volume backup; using the actions menu on the right, select [Create Block Volume](#)
- Provide a name for the Block Volume and select the correct AD and compartment of your *target* Essbase compute.
- Select [Create Block Volume](#)
- Repeat for the *source* 'config' block volume backup

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/creatingavolume.htm>

6) Attach block volumes created in step 5) to the target Compute

- Using the OCI console main menu, choose [Compute->Instances](#)
- Select the *target* Compute (the instance you are recovering)
- Under [Resources](#), select [Attached Block Volumes](#); there should be no attached volumes
- Select [Attach Block Volume](#)
- Use the default selections for iSCSI, READ/WRITE. Select the correct compartment and then select the *source* 'data' block volume you just created from backup and select [Attach](#).
- Repeat for the *source* 'config' block volume you just created from backup

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/attachingavolume.htm>

7) Connect to newly attached target block volumes and mount them

- After the restored 'data' and 'config' volumes are attached to the *target* compute, use the console menu for each attached volume to retrieve the iSCSI attach commands.
- As *opc* user, connect to the *target* Compute and execute iSCSI attach commands.
- As *opc* user, [lsblk](#) to show attached volumes and their disk labels.

```
[opc@essbase26jzoa4p-1 ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdb   8:16  0  1T  0 disk
├─sdb1 8:17  0 1024G 0 part
sdc   8:32  0  512G 0 disk
├─sdc1 8:33  0  512G 0 part
sda   8:0    0  46.6G 0 disk
├─sda2 8:2    0   8G  0 part [SWAP]
├─sda3 8:3    0  38.4G 0 part /
└─sda1 8:1    0  200M 0 part /boot/efi
```

- As *opc* user, mount the newly attached volumes as `/u01/data` and `/u01/config`

```
sudo systemctl daemon-reload
```

```
sudo mount /dev/sdc1 /u01/config
```

```
sudo mount /dev/sdb1 /u01/data
```

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/connectingtoavolume.htm>

8) As *opc* user, update the *target* compute `/etc/fstab` file with the UUID of the newly attached 'data' and 'config' volumes.

- `sudo blkid` to show the UUID for the newly attached 'data' and 'config' block volumes

```
[opc@essbase26jzoa4p-1 ~]$ sudo blkid
/dev/sda3: UUID="53bb89bf-50e3-4358-9f5e-e9c5ac8d535c" TYPE="xfs" PARTUUID="5a281070-157b-4c4a-ad99-8f2c3c773211"
/dev/sdb1: UUID="abdeaa9e-abc6-4ece-b1c6-362d7ff49037" TYPE="xfs" PARTUUID="9aa6827a-b81d-432f-9894-a1f46174324f"
/dev/sdc1: UUID="d84f0bc2-0392-44fe-8cc9-109b5d07f08f" TYPE="xfs" PARTUUID="f5cd0e89-7777-42f8-a35c-90bdc0d65f83"
/dev/sda1: SEC_TYPE="msdos" UUID="F831-6C7E" TYPE="vfat" PARTLABEL="EFI System Partition" PARTUUID="ce519aa3-9581-493e-97c2-f5bc164b5d3f"
```

- `sudo vi /etc/fstab` to update the 'data' and 'config' lines and replace the UUID, if changed.

```
# /etc/fstab
# Created by anaconda on Thu Feb 20 00:14:59 2020
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=53bb89bf-50e3-4358-9f5e-e9c5ac8d535c / xfs defaults,_netdev,netdev
UUID=F831-6C7E /boot/efi vfat defaults,uid=0,gid=0,umask=0077,shortname=windows-1.1
UUID=02789b15-4058-485e-86a6-38522a5c00c6 swap swap defaults,_netdev,x-initrd
#####
## ORACLE CLOUD INFRASTRUCTURE CUSTOMERS
##
## If you are adding an iSCSI remote block volume to this file you MUST
## include the '_netdev' mount option or your instance will become
## unavailable after the next reboot.
## SCSI device names are not stable across reboots; please use the device UUID instead of /dev/p
##
## Example:
## UUID="94c5aade-8bb1-4d55-ad0c-388bb8aa716a" /data1 xfs defaults,noatime,_netdev
##
## More information:
## https://uocs.us-phoenix-1.oraclecloud.com/Content/Block/Tasks/connectingtoavolume.htm
UUID=abdeaa9e-abc6-4ece-b1c6-362d7ff49037 /u01/data xfs defaults,_netdev,nofail 0 2
UUID=d84f0bc2-0392-44fe-8cc9-109b5d07f08f /u01/config xfs defaults,_netdev,nofail 0 2
#
-- INSERT --
```


9) As **opc** user, update the **target** compute **/etc/hosts** file to map the **source** domain information to the **target**

- Determine the domain name of your **source** compute
- Using the OCI console main menu, choose [Resource Manager->Stacks](#)
- Select your **source** stack to edit its details
- Under [Resources->Jobs](#), select the successful apply job
- Under [Job Details->Resources](#), scroll to the bottom of the log to review the job outputs

```

Outputs:

bastion_host_id = ocid1.instance.oc1.phx.anyhq1js63ksoticaydc7j5haxry1q5efencwqqui24cpbq3odiusr6cadgq
bastion_host_public_ip = 129.146.203.34
essbase_external_url = https://129.146.88.13/essbase
essbase_node_domain_name = essbaseuckqf41e-1.app.essuckqf41e.oraclelevcn.com
essbase_node_id = ocid1.instance.oc1.phx.anyhq1js63ksoticejnzspqh3njo7kcifmkckmrvckc6hk4a2mrrg4vuukfq
essbase_node_private_ip = 10.0.1.2
essbase_node_public_ip =
essbase_post_logout_redirect_url = https://129.146.88.13:443/essbase/jet/logout.html
essbase_redirect_url = https://129.146.88.13:443/essbase/redirect_uri
essbase_url = https://129.146.88.13/essbase
idcs_client_id = a9c014cb3c8b493d9a6588391dfc4a81
idcs_client_tenant = idcs-66c51de862f441718c91b68832884ea8
rcu_schema_prefix = ESSUCKQF41E
    
```

- Record the [essbase_node_domain_name](#)
- Edit the **/etc/hosts** file on the **target** compute and add the **source** domain information

`sudo vi /etc/hosts`

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.1.4 essbase26jzoo4p-1.app.essuckqf41e.oraclelevcn.com essbase26jzoo4p-1.essbaseuckqf41e-1.app.essuckqf41e.oraclelevcn.com essbaseuckqf41e-1
    
```

10) As **opc** user, update the **target** compute **/etc/oci-hostname.conf** file

- Set the **PRESERVE_HOSTINFO** setting to 3

`sudo vi /etc/oci-hostname.conf`

```

# This configuration file controls the hostname persistence behavior for Oracle Linux
# compute instance on Oracle Cloud Infrastructure (formerly Baremetal Cloud Services)
# Set PRESERVE_HOSTINFO to one of the following values
# 0 -- default behavior to update hostname, /etc/hosts and /etc/resolv.conf to
#     reflect the hostname set during instance creation from the metadata service
# 1 -- preserve user configured hostname across reboots; update /etc/hosts and
#     /etc/resolv.conf from the metadata service
# 2 -- preserve user configured hostname across instance reboots; no custom
#     changes to /etc/hosts and /etc/resolv.conf from the metadata service,
#     but dhclient will still overwrite /etc/resolv.conf
# 3 -- preserve hostname and /etc/hosts entries across instance reboots;
#     update /etc/resolv.conf from instance metadata service
PRESERVE_HOSTINFO=3
    
```

11) Update **source** load balancer networking for use with the **target** compute [skip if no **source** load balancer]

- Using the OCI console main menu, choose [Networking->Load Balancers](#)
- Select the **source** load balancer
- Select [Resources-> Backend Sets](#)
- Select the Backend Set **essbase**
- Select [Resources-> Backends](#)
- Select [Add Backends](#) and select the **target** compute.
- Remove the Backend corresponding to the **source** compute

Backends							
IP Address	Port	Weight	Drain	Offline	Backup	Health	
<input type="checkbox"/>	10.0.1.4	443	1	False	False	False	OK

- There is no need to update the IDCS confidential application URLs, as the same load balancer IP is now routing to the **target** Essbase instance.

12) If your source stack did not have a load balancer, update your IDCS Confidential Application login/logout URLs with the target IP address

- Login to IDCS and edit the source confidential application to ensure that the target IP address is substituted for the source IP address that was previously used.

13) Start the Essbase services.

- ssh to the target Essbase compute as opc user
- execute the following command: `sudo systemctl start essbase.service`

14) Resolve startup issues, if any.

- If it takes longer than you expect for Essbase to start, open another bash window to connect and check the status. If you see that the services didn't find the JDK, then link the source JDK being used by the mounted config volume to the JDK on the target system.

```
[opc@essbase26jzoa4p-1 ~]$ sudo systemctl start essbase.service
```

```
[opc@essbase26jzoa4p-1 ~]$ sudo systemctl status essbase.service [in 2nd bash window]
```

```
● essbase.service - Oracle Essbase service
   Loaded: loaded (/etc/systemd/system/essbase.service; enabled; vendor preset: disabled)
   Active: activating (start) since Fri 2020-03-06 22:48:24 GMT; 23s ago
 Main PID: 4782 (code=exited, status=0/SUCCESS);   : 16336 (start.sh)
 CGroup: /system.slice/essbase.service
 └─16336 /bin/sh /u01/config/domains/essbase_domain/esstools/bin/start.sh
 └─16337 /bin/bash /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_essbase.sh
 └─16347 /bin/bash /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_essbase.sh
 └─16349 /bin/sh /u01/oracle/oracle_common/common/bin/wlst.sh /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_servers.py
 essbasemdbvmvrl-1 /u01/config/domains/essbase_domain
 └─16352 /bin/sh /u01/oracle/oracle_common/common/bin/fmwconfig_common.sh wlst_internal.sh
 /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_servers.py essbasemdbvmvrl-1 /u01/config/domains/essbase_domain
 └─16353 /bin/sh /u01/oracle/oracle_common/common/bin/wlst_internal.sh /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_servers.py
 essbasemdbvmvrl-1 /u01/config/domains/essbase_domain
 └─16409 /usr/java/jdk1.8.0_241/bin/java -DORACLE_HOME=/u01/oracle/oracle_common -Xms32m -Xmx1024m WebLogic.WLST
 /u01/oracle/essbase/modules/oracle.essbase.sysman/scripts/start_servers.py essbasemdbvmvrl-1 /u01/config/doma...

Mar 06 22:48:35 essbasemdbvmvrl-1 startup.sh[16336]: <Mar 6, 2020 10:48:35 PM GMT> <Info> <Security> <BEA-090905> <Disabling the CryptoJ JCE Provider self-integrity check for better startup performance. To enable thi...fication=true.>
Mar 06 22:48:35 essbasemdbvmvrl-1 startup.sh[16336]: <Mar 6, 2020 10:48:35 PM GMT> <Info> <Security> <BEA-090906> <Changing the default Random Number Generator in RSA CryptoJ from ECDRBG128 to HMACDRBG. To disable th...faultPRNG=true.>
Mar 06 22:48:35 essbasemdbvmvrl-1 startup.sh[16336]: <Mar 6, 2020 10:48:35 PM GMT> <Info> <Security> <BEA-090909> <Using the configured custom SSL Hostname Verifier implementation: WebLogic.security.util.s.SSLWLSHostn...stnameVerifier.>
Mar 06 22:48:35 essbasemdbvmvrl-1 startup.sh[16336]: Node manager not running. Starting it...
Mar 06 22:48:36 essbasemdbvmvrl-1 startup.sh[16336]: NMPProcess: NODEMGR_HOME is already set to /u01/config/domains/essbase_domain/nodemanager
Mar 06 22:48:36 essbasemdbvmvrl-1 startup.sh[16336]: NMPProcess: The JDK wasn't found in directory /usr/java/jdk1.8.0_221.
Mar 06 22:48:36 essbasemdbvmvrl-1 startup.sh[16336]: NMPProcess: Please edit the startNodeManager.sh script so that the JAVA_HOME
Mar 06 22:48:36 essbasemdbvmvrl-1 startup.sh[16336]: NMPProcess: variable points to the location of your JDK.
```

- Link the source JDK to the JDK the target is expecting

```
[opc@essbase26jzoa4p-1 ~]$ sudo ln -s /usr/java/jdk1.8.0_241 /usr/java/jdk1.8.0_221
```

```
[opc@essbase26jzoa4p-1 ~]$ sudo systemctl start essbase.service
```

```
[opc@essbase26jzoa4p-1 ~]$ sudo systemctl status essbase.service
```

```
● essbase.service - Oracle Essbase service
   Loaded: loaded (/etc/systemd/system/essbase.service; enabled; vendor preset: disabled)
   Active: active (running) since Fri 2020-03-06 23:05:07 GMT; 22s ago
 Process: 17749 ExecStart=/u01/vmtools/startup.sh (code=exited, status=0/SUCCESS)
 Main PID: 17914 (java)
 CGroup: /system.slice/essbase.service
 └─17863 /bin/sh /u01/config/domains/essbase_domain/bin/startNodeManager.sh
 └─17864 /bin/sh /u01/oracle/wlserver/server/bin/startNodeManager.sh
 └─17914 /usr/java/jdk1.8.0_221/bin/java -server -Xms32m -Xmx200m -Djdk.tls.ephemeralDHKeySize=2048 -
 Dcoherence.home=/u01/oracle/wlserver/./coherence -Dbea.home=/u01/oracle/wlserver/.. -DNodeManagerHome=/u01/config/domain...
 └─18014 /bin/sh /u01/config/domains/essbase_domain/bin/startWebLogic.sh
 └─18064 /usr/java/jdk1.8.0_221/bin/java -server -Xms256m -Xmx1024m -cp /u01/oracle/wlserver/server/lib/WebLogic-launcher.jar -
 Dlaunch.use.env.classpath=true -DWebLogic.Name=AdminServer -Djava.security.policy=/u01/oracle/w...
 └─18253 /bin/sh /u01/config/domains/essbase_domain/bin/startWebLogic.sh
 └─18303 /usr/java/jdk1.8.0_221/bin/java -server -Xms512m -Xmx12059m -cp /u01/oracle/wlserver/server/lib/WebLogic-launcher.jar -
 Dlaunch.use.env.classpath=true -DWebLogic.Name=ess_server1 -Djava.security.policy=/u01/oracle/...

Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: Started ess_server1
Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: Finished starting servers
Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: Status of Domain: /u01/config/domains/essbase_domain
Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: NodeManager (essbasemdbvmvrl-1:9556): RUNNING
Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: Name      Type      Machine      Status
Mar 06 23:05:06 essbasemdbvmvrl-1 startup.sh[17749]: ----      ----      -
```

```
Mar 06 23:05:06 essbasemdbvmvml-1 startup.sh[17749]: ess_server1 Server essbasemdbvmvml-1.app.essmdbvmvml.oraclevcn.com RUNNING
Mar 06 23:05:06 essbasemdbvmvml-1 startup.sh[17749]: AdminServer Server essbasemdbvmvml-1.app.essmdbvmvml.oraclevcn.com RUNNING
Mar 06 23:05:07 essbasemdbvmvml-1 systemd[1]: essbase.service: Supervising process 17914 which is not our child. We'll most likely not notice when
it exits.
Mar 06 23:05:07 essbasemdbvmvml-1 systemd[1]: Started Oracle Essbase service.
```

- 15) After successfully recovering into the *target* Essbase stack, you can delete the failed *source* compute node and do further cleanup to un-needed block volumes and backups.

Non-Default Essbase Deployment: Recover Compute

Remember that non-default deployments have schemas for multiple Essbase instances in a single ATP. Still, if a compute failure takes place for one of the Essbase instances, you will deploy the new *target* Essbase schemas in the same ATP. You will then use relational database schema backups to import the *source* Essbase schema into the *target*. *Source* continues to refer to the schema belonging to the failed Essbase compute node, and *target* refers to a newly deployed stack built with “use existing database” and “use existing VCN” options.

Before beginning this recovery, backups should be taken following this method: [Backup ATP Schema and Block Volumes: Step-by-Step](#). To perform schema backups using data pump, you will need to first [install and configure the Oracle Instant Client](#).

1) Deploy a *target* Essbase stack using Oracle Marketplace

- Use *source* IDCS confidential application
- Use *source* ATP database and password
- Use *source* VCN and application subnet
- If your *source* stack has a load balancer, do not deploy a *target* Load Balancer– you can change the Backend Set after deploying the *target* stack
- If your *source* stack has a bastion host, deploy a bastion host with the *target* stack [*source* bastion host can be deleted after successful recovery]
- Use the same Essbase system admin user name and password in the *target* as you used in the *source*
- Use the same IDCS Essbase admin user in the *target* stack as you used in the *source* stack. If this is not possible, make sure the *source* Essbase instance has at least one valid IDCS user with the Essbase system administrator role. After you restore, you must login to the *target* instance as a valid IDCS user who had Essbase system administrator role on the *source* instance.

The Essbase services in the *target* stack will be started automatically by the ORM apply job.

Note: After deploying this stack, you will have one additional set of Essbase schemas in your ATP, because the failed instance has both *source* and *target* schemas. After completing the recovery, you can drop the schemas from the failed *source* instance.

2) Stop the *target* Essbase services. [Do not stop the Essbase compute in OCI.]

Hint: It goes without saying that the Essbase services on the *source* compute are stopped, because this use case involves *source* compute hardware failure. If you are simulating these steps, make sure you also stop the *source* compute’s Essbase services.

- ssh to the *target* Essbase compute as opc user
- execute the following command: `sudo systemctl stop essbase.service`
<https://www.oracle.com/pls/topic/lookup?ctx=en/database/other-databases/essbase/19.3&id=GUID-E442392E-04F3-4E0D-AA98-868B57EFB907>

3) Restore the *target* ATP Schema from *source* Schema Backup.

When restoring the ATP for your *target* stack, you will import your *source* schema backup into the target Essbase schema using the REMAP_SCHEMA option.

Be sure to select a *source* schema backup that was taken during a time when your *source* Essbase services were stopped; also, be sure that you have *source* Essbase ‘data’ block volume from the same time. In the case of compute recovery for non-default deployments, we will not replace the *target* ‘config’ block volume.

Note: the ATP will already be configured for use with object storage

- Make sure your instant client is configured to point to the ATP containing your *source* and *target* Essbase schemas
- Using the Oracle instant client, issue the following data pump import command

```
impdp admin/E55xxxxxx#@<database name>_low directory=data_pump_dir credential=<yourcredname>
dumpfile=https://objectstorage.us-phoenix-
1.oraclecloud.com/n/idpbzw1f44my/b/backup_<target_ATP_database
name>/o/essbaseschema_ESSHE4U6C90.dmp
REMAP_SCHEMA=<sourceEBprefix>_ESSBASE:<targetEBprefix>_ESSBASE parallel=16 partition_options=merge
table_exists_action=replace transform=segment_attributes:n transform=dwcs_cvt_iots:y
transform=constraint_use_default_index:y
exclude=index,cluster,indextype,materialized_view,materialized_view_log,materialized_zonemap,db_link
```

```
kzubetz@kzubetz-lap1 /usr/bin
$ impdp admin/E @ESSOT8QU27W_low directory=data_pump_dir credential=DEF_CRED_kz dumpfile=https://objectstorage.us
-phoenix-1.oraclecloud.com/n/idpbzw1f44my/b/backup_essot8qu27w/o/essbaseschema_ESSHE4U6C90.dmp REMAP_SCHEMA=ESSHE4U6C90_ESS
BASE:ESSPT915WMM_ESSBASE parallel=16 partition_options=merge table_exists_action=replace transform=segment_attributes:n tran
sform=dwcs_cvt_iots:y transform=constraint_use_default_index:y exclude=index,cluster,indextype,materialized_view,materializ
ed_view_log,materialized_zonemap,db_link


Import: Release 18.0.0.0.0 - Production on Tue Apr 28 15:13:26 2020
Version 18.5.0.0.0

Copyright (c) 1982, 2018, Oracle and/or its affiliates. All rights reserved.

Connected to: Oracle Database 18c Enterprise Edition Release 18.0.0.0.0 - Production
Master table "ADMIN"."SYS_IMPORT_FULL_01" successfully loaded/unloaded
Starting "ADMIN"."SYS_IMPORT_FULL_01": admin/*****@ESSOT8QU27W_low directory=data_pump_dir credential=DEF_CRED_kz dumpfil
e=https://objectstorage.us-phoenix-1.oraclecloud.com/n/idpbzw1f44my/b/backup_essot8qu27w/o/essbaseschema_ESSHE4U6C90.dmp REMA
P_SCHEMA=ESSHE4U6C90_ESSBASE:ESSPT915WMM_ESSBASE parallel=16 partition_options=merge table_exists_action=replace transform=se
gment_attributes:n transform=dwcs_cvt_iots:y transform=constraint_use_default_index:y exclude=index,cluster,indextype,materia
lized_view,materialized_view_log,materialized_zonemap,db_link
Processing object type SCHEMA_EXPORT/USER
ORA-31684: Object type USER:"ESSPT915WMM_ESSBASE" already exists

Processing object type SCHEMA_EXPORT/SYSTEM_GRANT
Processing object type SCHEMA_EXPORT/ROLE_GRANT
```

- After the schema import finishes, audit your data using a database client like SQL Developer. You can look at the **ESSBASE_APPLICATION** table within the **<targetprefix>_ESSBASE** schema and see that the *target* schema (which was empty prior to schema import) has the *source* applications.



4) Detach 'data' block volume from the target Essbase compute.

- As *opc* user, connect to your *target* compute and unmount the **/u01/data** volume
`sudo umount /u01/data`
- Using the OCI console main menu, choose **Compute->Instances**
- Select the *target* Compute in which your Essbase instance is deployed
- In **Compute Details->Resources**, select **Attached Block Volumes**
- Using the actions menu to the right of the 'data' volume, select **Detach**
- ssh to the Compute as *opc* user and issue the detach ISCSI commands as indicated.
- Select **Continue Detachment** in OCI after issuing ISCSI detach commands.

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/detachingavolume.htm>

5) Create 'data' block volume from source block volume backup.

Block volume backups are stored in object storage; they are not able to be attached to compute instances without first creating a new volume in block storage. Clones, on the other hand, are standalone block volumes and can be directly attached to a compute instance. Skip this step if you cloned your block volumes. You may have backed up the block volume group; if so, you can create the volume group.

- Using the OCI console main menu, choose [Block Storage->Block Volume Backups](#)
- Find the desired 'data' volume backup; using the actions menu on the right, select [Create Block Volume](#)
- Provide a name for the Block Volume and select the correct AD and compartment of your *target* Essbase compute.
- Select Create Block Volume

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/creatingavolume.htm>

6) Attach 'data' block volume created in step 5) to the target Compute

- Using the OCI console main menu, choose [Compute->Instances](#)
- Select the *target* Compute (the instance you are recovering)
- Under [Resources](#), select [Attached Block Volumes](#); there should be no attached 'data' volume
- Select [Attach Block Volume](#)
- Use the default selections for ISCSI, READ/WRITE. Select the correct compartment and then select the *source* 'data' block volume you just created from backup and select [Attach](#).

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/attachingavolume.htm>

7) Connect to newly attached target block volume and mount it

- After the restored 'data' volume is attached to the *target* compute, use the actions menu for each attached volume to retrieve the ISCSI attach commands.
- As opc user, connect to the *target* Compute and execute ISCSI attach commands.
- As opc user, `lsblk` to show attached volumes and their disk labels.

```
[opc@essbase26jzoa4p-1 ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdb   8:16  0   1T  0 disk
└─sdb1 8:17  0 1024G 0 part
sdc   8:32  0   512G 0 disk
└─sdc1 8:33  0   512G 0 part
sda   8:0    0  46.6G 0 disk
├─sda2 8:2    0    8G  0 part [SWAP]
├─sda3 8:3    0  38.4G 0 part /
└─sda1 8:1    0   200M 0 part /boot/efi
```

- As opc user, mount the newly attached volume as `/u01/data`
`sudo systemctl daemon-reload`
`sudo mount /dev/sdb1 /u01/data`

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Tasks/connectingtoavolume.htm>

8) As opc user, update the target compute `/etc/fstab` file with the UUID of the newly attached 'data' volume.

- `sudo blkid` to show the UUID for the newly attached 'data' block volume

```
[opc@essbase26jzoa4p-1 ~]$ sudo blkid
/dev/sda3: UUID="53bb89bf-50e3-4358-9f5e-e9c5ac8d535c" TYPE="xfs" PARTUUID="5a281070-157b-4c4a-ad99-8f2c3c773211"
/dev/sdb1: UUID="abdeaa9e-abc6-4ece-b1c6-362d7ff49037" TYPE="xfs" PARTUUID="9aa6827a-b81d-432f-9894-a1f46174324f"
/dev/sdc1: UUID="d84f0bc2-0392-44fe-8cc9-109b5d07f08f" TYPE="xfs" PARTUUID="f5cd0e89-7777-42f8-a35c-90bd0d65f83"
/dev/sda1: SEC_TYPE="msdos" UUID="F831-6C/E" TYPE="vfat" PARTLABEL="EFI System Partition" PARTUUID="ce519aa3-9581-493e-97c2-f5bc164b5d3f"
```

- `sudo vi /etc/fstab` to update the 'data' line and replace the UUID, if changed.

```

# /etc/fstab
# Created by anaconda on Thu Feb 20 00:14:59 2020
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=53bb89bf-50e3-4358-9f5e-e9c5ac8d535c / xfs defaults,_netdev,_net
UUID=F831-6C7E /boot/efi vfat defaults,uid=0,gid=0,umask=0077,shortname=
UUID=02789b15-4058-485e-86a6-38522a5c00c6 swap swap defaults,_netdev,x-initrd
#####
## ORACLE CLOUD INFRASTRUCTURE CUSTOMERS
##
## If you are adding an iSCSI remote block volume to this file you MUST
## include the '_netdev' mount option or your instance will become
## unavailable after the next reboot.
## SCSI device names are not stable across reboots; please use the device UUID instead of /dev/p
##
## Example:
## UUID="94c5aade-8bb1-4d55-ad0c-388bb8aa716a" /data1 xfs defaults,noatime,_netdev
##
## More information:
## https://uocs.us-phoenix-1.oraclecloud.com/Content/Block/Tasks/connectingtoavolume.htm
UUID=abdeaa9e-abc6-4ece-b1c6-362d7ff49037 /u01/data xfs defaults,_netdev,nofail 0 2
UUID=d84f0c2-0392-44fe-8cc9-109b5d07f08f /u01/config xfs defaults,_netdev,nofail 0 2
-- INSERT --

```

9) Update *source* load balancer networking for use with the *target* compute [skip if no *source* load balancer]

- Using the OCI console main menu, choose [Networking->Load Balancers](#)
- Select the *source* load balancer
- Select [Resources-> Backend Sets](#)
- Select the Backend Set *essbase*
- Select [Resources-> Backends](#)
- Select [Add Backends](#) and select the *target* compute.
- Remove the Backend corresponding to the *source* compute

Backends							
Add Backends Search...							
<input type="checkbox"/>	IP Address	Port	Weight	Drain	Offline	Backup	Health
<input type="checkbox"/>	10.0.1.4	443	1	False	False	False	OK

- There is no need to update the IDCS confidential application URLs, as the same load balancer IP is now routing to the *target* Essbase instance.

10) If your *source* stack did not have a load balancer, update your IDCS Confidential Application login/logout URLs with the *target* IP address

- Login to IDCS and edit the *source* confidential application to ensure that the *target* IP address is substituted for the *source* IP address that was previously used.

11) Start the Essbase services.

- ssh to the *target* Essbase compute as *opc* user
- execute the following command: `sudo systemctl start essbase.service`

12) After successfully recovering *source* information into the *target* Essbase stack, you can delete the failed *source* compute node and do further cleanup to un-needed block volumes and backups. At this time, you can also drop the failed *source* schemas from your ATP.

ESSBASE BACKUPS: STEP-BY-STEP

Essbase backups support your recovery needs. Depending on your recover use case, you will follow one of two backup methods. They differ in the treatment of the relational database backup. Sometimes you will restore the entire relational database, and at other times, you will restore only selected schemas from the relational database. The restore use cases above indicate which backup method to follow.

Backup Entire ATP and Block Volumes: Step-by-Step

If your default deployment uses Autonomous Transaction Processing (ATP) as the relational database, then you can use Oracle-provided scripts to take synchronized backup of your ATP and Essbase block volumes. Alternatively, you can use a database client and the OCI console to perform the same steps. If you deployed using DBaaS, then you must take manual backups, but the concepts for when to use the full database vs schema backup apply to the various 'restore' use cases.

Backing up Essbase using Oracle-Provided Scripts

Backup scripts are provided for default Marketplace deployments that use ATP as the relational database. The provided scripts perform the following tasks:

- 1) **Configure your ATP to work with your object storage**
- 2) **Stop the Essbase services**
- 3) **Back up your ATP**
- 4) **Back up your Essbase 'data' and 'config' block volumes**
- 5) **Start the Essbase services**

To initiate a backup of an instance created using the default Marketplace image deployment, schedule the backup for a convenient time when your users are not in the system and follow these steps:

- Using documentation, ensure that the required policies to manage backups are in place
<https://www.oracle.com/pls/topic/lookup?ctx=en/database/other-databases/essbase/19.3&id=GUID-64D56F3C-ED1A-4CF4-8CF7-AACB669EBB87>
- ssh to your Essbase compute as opc user (If you deployed a Bastion host, you will use it as a proxy)
- cd /u01/vmtools
- ./configure-backup-storage.sh
- ./backup.sh

Note: Essbase services will automatically be stopped and started by the backup.sh script.

Backing up Essbase using DB client and OCI console

Alternatively, you can use a database client and the OCI console to perform backups of Essbase default Marketplace deployments. Remember that an Essbase backup consists of block volume backups and ATP backup.

1) **Configure your ATP instance to work with your object storage**

Use a DB client application to configure your ATP instance for use with your OCI account and a default object storage bucket. SQL Developer is a nice choice for DB client, because it allows Cloud Wallet connections and can connect to multiple ATP instances at the same time, if needed.

- Create a connection to your ATP instance using SQL Developer [consider your proxy needs if connected via a corporate network – not described here]:

Database Type Oracle
Authentication type Default
Username admin
Password [database password you provided during Essbase deployment]
Connection Type Cloud Wallet
Configuration File [Browse to location of downloaded ATP wallet .zip file]
Service [select the _low service for your ATP wallet]

<https://www.oracle.com/pls/topic/lookup?ctx=en/cloud/paas/atp-cloud&id=GUID-9132B1DF-7847-45F9-B63B-F91B2B8C88F7>

- Create an auth token using the OCI Console.

Identity->Users->User Details->Auth Tokens->"Generate Token"

Copy and record the token value in a secure location. You will not be able to retrieve it again.

<https://docs.cloud.oracle.com/en-us/iaas/Content/Registry/Tasks/registrygettingauthtoken.htm>

- Using SQL Developer, connect to your ATP instance and execute the following SQL:
 - i) Create cloud credential to identify your OCI user

```
BEGIN
DBMS_CLOUD.CREATE_CREDENTIAL(
  credential_name => '<name of you choosing>',
  username => '<your OCI username>',
  password => '<OCI auth token value>'
);
END;
```
 - ii) Associate the cloud credential with your ATP instance

```
ALTER DATABASE PROPERTY SET DEFAULT_CREDENTIAL = 'ADMIN.<credential_name from above>;'
```
 - iii) Set the default OCI Object Storage Namespace for your ATP database

```
ALTER DATABASE PROPERTY SET default_bucket='https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/idpbzwxxy';
```

Note: us-phoenix-1.oraclecloud.com should be replaced with your OCI region

Note: idpbzwxxy should be replaced with the **Namespace** of your object storage.

<https://docs.cloud.oracle.com/en-us/iaas/Content/Object/Tasks/understandingnamespaces.htm>

<https://www.oracle.com/pls/topic/lookup?ctx=en/cloud/paas/atp-cloud&id=GUID-D95E5D6A-C470-4A68-9545-CC99D937E7D1>

2) Stop the Essbase services. [Do not stop the Essbase compute in OCI.]

- ssh to the Essbase compute as opc user
- execute the following command: `sudo systemctl stop essbase.service`

<https://www.oracle.com/pls/topic/lookup?ctx=en/database/other-databases/essbase/19.3&id=GUID-E442392E-04F3-4E0D-AA98-868B57EFB907>

3) Back up your ATP

- In the OCI console, navigate to the details of your ATP instance.
- Select Backups from the lower left, then [Create Manual Backup](#).
- Provide a Name for the backup and click [Create](#).

<https://www.oracle.com/pls/topic/lookup?ctx=en/cloud/paas/atp-cloud/atpug&id=GUID-D95E5D6A-C470-4A68-9545-CC99D937E7D1>

Your backup is automatically stored in the default object storage space for your ATP instance.

4) Backup your Essbase 'data' and 'config' block volumes

- In the OCI console, navigate to [Block Storage->Block Volume Groups](#)
- Find the volume group that corresponds to your Essbase compute (use of optional RESOURCE DISPLAY NAME PREFIX when creating your stack allows you to identify all stack components using a common prefix)
- Using the actions menu to the right of the created date, select [Create Volume Group Backup](#)
- Provide a Name for the backup group and click [Create](#).
- Alternatively, you can create a clone of the volume group. Read about backups and clones here:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Concepts/blockvolumebackups.htm>

5) Start the Essbase services.

- ssh to the Essbase compute as opc user
- execute the following command: `sudo systemctl start essbase.service`

<https://www.oracle.com/pls/topic/lookup?ctx=en/database/other-databases/essbase/19.3&id=GUID-E442392E-04F3-4E0D-AA98-868B57EFB907>

6) Note the timestamp of your ATP and Block Volume backups. Because your Essbase services are stopped, these backups can be used to consistently restore Essbase if the need arises.

Backup ATP Schema and Block Volumes: Step-by-Step

Before You Begin: Install and Configure Oracle Instant Client and Tools

Backing up non-default Marketplace deployments requires use of the Oracle Database Instant Client and Oracle Data Pump to export the schemas related to a specific Essbase instance instead of taking a complete ATP backup. The Oracle Database Instant Client is recommended for this purpose.

1) Download and install an Instant Client version compatible with your ATP version. [Oracle Instant Client Downloads](#)

- Essbase 19c deployments recently began using ATP version 19c. [The examples in this paper were created with ATP version 18c.]
- Be sure to download and install the corresponding Visual Studio Redistributable
- You will also need the Tools Package which includes Data Pump
- Optionally, you can download SQL*Plus Package, but SQL developer will work as well.
- See the installation instructions on the platform install download page for the installation steps required after you download Oracle Instant Client and the Tools Package. Below is an example:

Instant Client Installation for Microsoft Windows 64-bit

See the [Instant Client Home Page](#) for more information about Instant Client packages. Client-server version interoperability is detailed in [Doc ID 207303.1](#). For example, Oracle Call Interface 19, 18 and 12.2 can connect to Oracle Database 11.2 or later. Some tools may have other restrictions.

1. Download the appropriate Instant Client packages for your platform. All installations require the Basic or Basic Light package.
2. Unzip the packages into a single directory such as `C:\oracle\instantclient_18_5`
3. Add this directory to the `PATH` environment variable. If you have multiple versions of Oracle libraries installed, make sure the new directory occurs first in the path. Restart any terminal windows or otherwise make sure the new `PATH` is used by your applications.
4. Download and install the correct Visual Studio Redistributable from Microsoft. Instant Client 19 requires the [Visual Studio 2017 redistributable](#). Instant Client 18 and 12.2 require the [Visual Studio 2013 redistributable](#). Instant Client 12.1 requires the [Visual Studio 2010 redistributable](#).
5. If you intend to co-locate optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora`, `ldap.ora`, or `oraaccess.xml` with Instant Client, then create a subdirectory such as `C:\oracle\instantclient_19_3\network\admin`. This is the default Oracle client configuration directory for applications linked with this Instant Client.

Alternatively, Oracle client configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.
6. Start your application.

ODBC users should follow the [ODBC Installation Instructions](#).


2) Configure your Instant Client for use with multiple ATP instances:

- Download the DB Wallet from each ATP instance you plan to connect to.
<https://www.oracle.com/pls/topic/lookup?ctx=en/cloud/paas/atp-cloud/atpug&id=GUID-B06202D2-0597-41AA-9481-3B174F75D4B1>
- Create a network folder under your instant client folder; for example, C:\temp\instantclient_18_5\network
- Unzip the wallet contents for each ATP instance in a folder under your network folder; for example:

System (C:) > temp > instantclient_18_5 > network > embWL_admin

Name	Date modified	Type	Size
ewallet.sso	4/3/2020 6:44 PM	SSO File	7 KB
ewallet.p12	4/3/2020 6:44 PM	Personal Information...	7 KB
keystore.jks	4/3/2020 6:44 PM	JKS File	4 KB
ojdbc.properties	4/3/2020 6:44 PM	PROPERTIES File	1 KB
readme.md	4/3/2020 6:44 PM	MD File	1 KB
sqlnet.ora	4/3/2020 7:05 PM	ORA File	1 KB
tnsnames.ora	4/3/2020 6:44 PM	ORA File	2 KB
truststore.jks	4/3/2020 6:44 PM	JKS File	4 KB

- Edit the sqlnet.ora file for each unzipped wallet folder and update the WALLET_LOCATION:



- Set a system environment variable appropriate for your instant client install location
`TNS_PATH= C:\temp\instantclient_18_5\network`
- When you are ready to connect to an ATP instance using the Instant Client, SET/EXPORT a TNS_ADMIN variable for your shell (examples in this document use bash on windows); then connect.
`$ export TNS_ADMIN=$TNS_PATH/<unzipped wallet folder name>`
`$ sqlplus admin@<"ATP wallet TNS connection"_low>`

```
/usr/bin
kzubetz@kzubetz-lap1 /usr/bin
$ export TNS_ADMIN=$TNS_PATH/embWL_admin/
kzubetz@kzubetz-lap1 /usr/bin
$ sqlplus admin@ess8frzlire_low

SQL*Plus: Release 18.0.0.0.0 - Production on Fri Apr 3 20:16:40 2020
Version 18.5.0.0.0

Copyright (c) 1982, 2018, Oracle. All rights reserved.

Enter password:
Last Successful login time: Fri Apr 03 2020 20:13:58 -07:00

Connected to:
Oracle Database 18c Enterprise Edition Release 18.0.0.0.0 - Production
Version 18.4.0.0.0

SP2-0310: unable to open file "LOGIN.SQL"
SQL>
```

Backing up Essbase schema using Data Pump and OCI console

1) Configure your ATP instance to work with your object storage

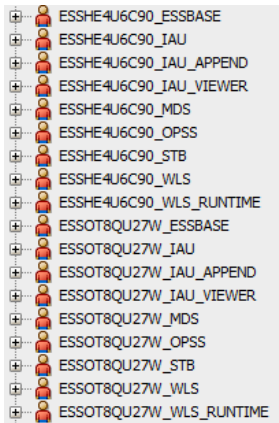
Same as [Backing up Essbase using DB client and OCI console](#), step 1). Or, optionally, you can use SQL*Plus and Oracle Instant Client.

2) Stop the Essbase services. [Do not stop the Essbase compute in OCI.]

Same as [Backing up Essbase using DB client and OCI console](#), step 2).

3) Backup the Essbase schema from your ATP instance using Data Pump and move the .dmp file to OCI Object Storage

Each Essbase instance has 9 associated schemas in your ATP. All 9 schemas have a common `rcu_schema_prefix`, which is reported in the outputs of the Oracle Resource Manager (ORM) apply job. When using IDCS for security, you only need to back up the `<prefix>_ESSBASE` schema that corresponds to the Essbase instance you want to back up. Remember that your ATP instance may have Essbase schemas from multiple instances.



- Make sure your `TNS_ADMIN` variable points to the wallet location of the ATP instance you plan to export from. Use Oracle Instant Client to issue the following Data Pump command:

```
$ export TNS_ADMIN=TNS_PATH/mywallet_admin
```

```
$ expdp admin/E55base2020###@ESSOT8QU27W_low directory=data_pump_dir  
schemas=ESSHE4U6C90_ESSBASE logfile=ESSHE4U6C90.out dumpfile=ESSHE4U6C90.dmp
```

Note that the ATP prefix in the `connection information` may not be the same as the `schema prefix` for the Essbase schema you are backing up.

```
kzubetz@kzubetz-1apl /usr/bin  
$ export TNS_ADMIN=${TNS_PATH}/kzbaseATP_admin/  
  
kzubetz@kzubetz-1apl /usr/bin  
$ expdp admin/E55base2020###@ESSOT8QU27W_low directory=data_pump_dir schemas=ESSHE4U6C90_ESSBASE logfile=ESSHE4U6C90.out  
dumpfile=ESSHE4U6C90.dmp  
  
Export: Release 18.0.0.0.0 - Production on Tue Apr 28 14:36:45 2020  
Version 18.5.0.0.0  
  
Copyright (c) 1982, 2018, Oracle and/or its affiliates. All rights reserved.  
  
Connected to: Oracle Database 18c Enterprise Edition Release 18.0.0.0.0 - Production  
Starting "ADMIN"."SYS_EXPORT_SCHEMA_01": admin/*****@ESSOT8QU27W_low directory=data_pump_dir schemas=ESSHE4U6C90_ESSBASE logfile=ESSHE4U6C90.out dumpfile=ESSHE4U6C90.dmp  
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA  
Processing object type SCHEMA_EXPORT/TABLE/INDEX/STATISTICS/INDEX_STATISTICS  
Processing object type SCHEMA_EXPORT/TABLE/STATISTICS/TABLE_STATISTICS  
Processing object type SCHEMA_EXPORT/STATISTICS/MARKER
```

- Schema backups are physically stored on disk in the Data Pump Directory of your ATP.

```
Dump file set for ADMIN.SYS_EXPORT_SCHEMA_01 is:  
/u03/dbfs/A44C828407B0B83E0536410000A532D/data/dpdump/ESSHE4U6C90.dmp  
Job "ADMIN"."SYS_EXPORT_SCHEMA_01" successfully completed at Tue Apr 28 21:39:08 2020 elapsed 0 00:02:14
```

- Move the .dmp file into an Object Storage bucket of your choosing (use SQL Developer, unless you installed the Instant Client SQL*Plus package). The /o/xxxxxxxx.dmp portion of the PUT_OBJECT uri indicates the name you want to assign for the .dmp file in your Object Storage. The file_name must match the .dmp filename you assigned when you created the export on disk using data pump.

```
BEGIN
  DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_kz',
    object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/idpbzw1f44my/b/backup_essot8qu27w/o/essbaseschema_ESSHE4U6C90.dmp',
    directory_name => 'DATA_PUMP_DIR',
    file_name => 'ESSHE4U6C90.dmp');
END;
/
```

```
ksubeta@ksubeta-lap1 /user/bin
$ sqlplus admin@ESSOT8QU27W_low_

SQL*Plus: Release 18.0.0.0.0 - Production on Tue Apr 28 14:48:09 2020
Version 18.5.0.0.0

Copyright (c) 1982, 2018, Oracle. All rights reserved.

Enter password:
Last Successful login time: Tue Apr 28 2020 14:47:28 -07:00

Connected to:
Oracle Database 18c Enterprise Edition Release 18.0.0.0.0 - Production
Version 18.4.0.0.0

SP2-0310: unable to open file "LOGIN.SQL"
SQL> BEGIN
  2  DBMS_CLOUD.PUT_OBJECT(credential_name => 'DEF_CRED_kz',
  3  object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/idpbzw1f44my/b/backup_essot8qu27w/o/essbaseschema_ESSHE4U6C90.dmp',
  4  directory_name => 'DATA_PUMP_DIR',
  5  file_name => 'ESSHE4U6C90.dmp');
  6 END;
  7 /

PL/SQL procedure successfully completed.
```

- Refresh your Object Storage Bucket to see the dmp file.

Object Storage » Bucket Details

backup_essot8qu27w

Edit Visibility Move Resource Re-encrypt Add Tags Delete

Bucket Information Tags

Visibility: Private Encrypt
 Namespace: idpbzw1f44my Created:
 Storage Tier: Standard Compart
 Approximate Count: 9 objects ⓘ Approxim
 ETag: d956106d-92c3-4dcf-a04d-5d3c110f8235 Emit Obj
 OCID: ...3tteerka [Show](#) [Copy](#) Object Ve

Resources

Objects

Upload Objects Restore Delete

<input type="checkbox"/>	Name	Size	L
<input type="checkbox"/>	essbaseschema_ESSHE4U6C90.dmp	3.11 MiB	T

4) Backup the Essbase 'data' and 'config' block volumes

Same as [Backing up Essbase using DB client and OCI console](#) step 4). [Make sure you back up the block volumes attached to the Essbase instance for which you just exported RCU schemas.]

5) Start the Essbase services.

Same as [Backing up Essbase using DB client and OCI console](#) step 5).

6) Note the timestamp of your ATP .dmp file in object storage and block volume backups. Because your Essbase services were stopped, these backups can be used to consistently restore Essbase if the need arises.

VARIOUS OTHER NOTES

Managing Volume Groups

After detaching and attaching volume groups, you may want to perform some maintenance to ensure that the 'config' and 'data' volumes running in your target Essbase instance are so-called "matched pairs".

You can use the OCI console to manage [Block Volume Groups](#).

- Select the [Block Volume Group](#) for your target Essbase instance.
- Remove any Block Volumes in the group that are no longer attached to your compute.
- Add Block Volumes that are not in the group but are attached to your compute.

Resolving Network Components

To minimize disruption to your end user community, you should consider the use of some resolving network component to maintain a consistent URL endpoint for your users. You should work with your networking team to assign an appropriate corporate domain endpoint for your Essbase instance(s).

On a Windows box, this can be simulated as follows:

- **As an administrator**, edit the `Windows\System32\drivers\etc\hosts` file, adding your target IP address and some domain name (`essbase.com`)

```
1 # Copyright (c) 1993-2009 Microsoft Corp.
2 #
3 # This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
4 #
5 # This file contains the mappings of IP addresses to host names. Each
6 # entry should be kept on an individual line. The IP address should
7 # be placed in the first column followed by the corresponding host name.
8 # The IP address and the host name should be separated by at least one
9 # space.
10 #
11 # Additionally, comments (such as these) may be inserted on individual
12 # lines or following the machine name denoted by a '#' symbol.
13 #
14 # For example:
15 #
16 #       102.54.94.97       rhino.acme.com           # source server
17 #       38.25.63.10      x.acme.com             # x client host
18
19 # localhost name resolution is handled within DNS itself.
20 #   127.0.0.1       localhost
21 #   ::1            localhost
22 150.136.30.175     essbase.com
```

- Edit your *target* IDCS application and replace the IP address with the domain you used above.

Client Configuration

Register Client No Client

Allowed Grant Types Resource Owner Client Credentials

Allow non-HTTPS URLs

* Redirect URL


Logout URL

Post Logout Redirect URL

* Client Type Trusted Confidential Public

Certificate

- Connect to Essbase using the modified `essbase_external_url` from your apply job output.



If you are using Essbase version 19.3.0.0.1, there are additional steps to those directly above.

1) Connect to your compute as opc user

2) Edit `auth_openidc_essbase.conf`

```
sudo vi /etc/httpd/conf.d/auth_openidc_essbase.conf
```

search for your load balancer IP and replace it with your domain name [example: essbase.com]

3) Add your domain name to `/etc/hosts`

```
sudo vi /etc/hosts
```

add a line with load balancer IP and your domain name

4) Update `/etc/oci-hostname.conf`

```
sudo vi /etc/oci-hostname.conf
```

change PRESERVE_HOSTINFO setting to 3 (to retain changes to `/etc/hosts`)

5) Restart httpd:

```
sudo systemctl restart httpd
```

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2020, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Essbase on OCI: Backup and Restore
September, 2020
Author: Karen Zubetz

