# Oracle® Essbase

# Essbase Stack Deployment on Oracle Cloud Infrastructure





Oracle Essbase Essbase Stack Deployment on Oracle Cloud Infrastructure,

F17139-19

Copyright © 2019, 2022, Oracle and/or its affiliates.

Primary Author: Essbase Information Development Team

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

# Contents

Get Started with Oracle Essbase Setup and Administr			
About Oracle Essbase	1-1		
About Components and Terminology	1-1 1-4		
Administrator Access Requirements			
Typical Workflow for Administrators	1-5		
Set Up Oracle Essbase			
Before You Begin with Oracle Essbase	2-1		
Create Dynamic Groups	2-5		
Set Up Policies	2-6		
Set Up Essbase Access in Identity Cloud Service	2-8		
Set Relational Database Connectivity	2-9		
Create a Confidential Identity Cloud Service Application	2-9		
Supported Compute Shapes	2-10		
Create Vault, Secrets, and Encrypt Values	2-11		
Deploy Essbase	2-13		
Complete Post-Deployment Tasks	2-17		
Modify the Confidential Identity Cloud Service Application	2-17		
Set Up the SSL Certificate	2-18		
Secure Your Network	2-20		
Test Connectivity to Essbase	2-20		
Complete System Hardening and Cleanup Tasks	2-22		
Troubleshoot Deployment Errors	2-23		
Manage the Oracle Essbase Stack on Oracle Cloud Ir	nfrastructure		
Start, Stop, and Delete an Essbase Stack Instance	3-1		
Prepare to Work with an Essbase Stack Instance	3-1		
Start the Stack	3-2		
Stop the Stack	3-2		
Delete the Stack	3-3		
Use Commands to Start, Stop, and View Status of Processes	3-4		



Restart the Essbase Compute Instance	3-5
Access Oracle Essbase Using SSH	3-6
Resize Block Storage Volumes	3-9
Patch and Roll Back	3-10
Patch and Roll Back - for Version 19.3.0.2.3 and Earlier	3-10
Patch and Roll Back - For Version 19.3.0.3.4 and Later	3-11
Back Up and Restore Essbase	3-13
About Backup and Restore	3-13
Back Up and Restore Individual Applications	3-13
Back Up and Restore an Essbase Instance	3-14
Install and Configure Oracle Instant Client and Tools	3-23
Monitor and Diagnose Essbase Operations	3-23
Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service	3-24
Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service	3-24
Collect Diagnostic Information on the Essbase Node	3-25
Access the WebLogic Console	3-25
Reset or Update Admin Password	3-27
	<del></del>
Selective and Ordered Import of Artifacts	
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users	4-2
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications	4-2 4-3
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts	4-2 4-3 4-6
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options	4-1 4-2 4-3 4-6 4-8 4-9
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups	4-2 4-3 4-6 4-8 4-9
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility	4-2 4-3 4-6 4-8 4-9 4-11
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes	4-2 4-3 4-6 4-8 4-9 4-11 4-12
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes Download the Cube Export Utility	4-2 4-3 4-6 4-8 4-9 4-11 4-12
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility  Export Essbase 11g On-Premises Cubes	4-2 4-3 4-6 4-8 4-9 4-11
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes Download the Cube Export Utility Review Member Names Before you Import an Application Workbook Created by	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility  Export Essbase 11g On-Premises Cubes  Download the Cube Export Utility  Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility  Export Essbase 11g On-Premises Cubes  Download the Cube Export Utility  Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility  Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode  Upgrade Aggregate Storage Outline Version	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13 4-16
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility  Export Essbase 11g On-Premises Cubes  Download the Cube Export Utility  Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility  Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode  Upgrade Aggregate Storage Outline Version	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13 4-13 4-16 4-18
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes Download the Cube Export Utility Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode Upgrade Aggregate Storage Outline Version  Migrate Essbase Cloud Service Applications and Users	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13
Selective and Ordered Import of Artifacts  Migrate Essbase 11g On-Premise Applications and Users  Prepare to Migrate from Essbase 11g On-Premises Applications  Migrated Essbase 11g On-Premises Artifacts  LCM Utility Export Options  Migrate Essbase 11g On-Premises Users and Groups  Migrate an Essbase 11g On-Premises Application Using LCM Utility  Export Essbase 11g On-Premises Cubes  Download the Cube Export Utility  Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility  Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode  Upgrade Aggregate Storage Outline Version  Migrate Essbase Cloud Service Applications and Users  Prepare to Migrate Cloud Service Applications and Users	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13 4-16 4-18 4-18
Selective and Ordered Import of Artifacts Migrate Essbase 11g On-Premise Applications and Users Prepare to Migrate from Essbase 11g On-Premises Applications Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes Download the Cube Export Utility Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode Upgrade Aggregate Storage Outline Version Migrate Essbase Cloud Service Applications and Users Prepare to Migrate Cloud Service Applications and Users Migrated Cloud Service Artifacts	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13 4-16 4-18 4-19
Migrated Essbase 11g On-Premises Artifacts LCM Utility Export Options Migrate Essbase 11g On-Premises Users and Groups Migrate an Essbase 11g On-Premises Application Using LCM Utility Export Essbase 11g On-Premises Cubes Download the Cube Export Utility Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode Upgrade Aggregate Storage Outline Version Migrate Essbase Cloud Service Applications and Users Prepare to Migrate Cloud Service Applications and Users Migrated Cloud Service Artifacts Migrate Cloud Service Applications Using CLI Tool	4-2 4-3 4-6 4-8 4-9 4-11 4-12 4-13 4-13 4-16 4-18 4-18 4-19 4-21



# 5 Manage Users and Roles

About Users and Roles	5-1
User Roles and Application Permissions	5-1
Provision Application Permissions	5-2



# Accessibility and Support

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

#### **Access to Oracle Support**

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <a href="http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info">http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs</a> if you are hearing impaired.



1

# Get Started with Oracle Essbase Setup and Administration

Let's explore Oracle Essbase and what you need to know to get started with administration.

#### Topics:

- About Oracle Essbase
- About Components and Terminology
- Administrator Access Requirements
- Typical Workflow for Administrators

### **About Oracle Essbase**

Oracle Essbase is a business solution that uses a proven, flexible, best-in-class architecture for analysis, reporting, and collaboration. Essbase delivers instant value and greater productivity for your business users, analysts, modelers, and decision-makers, across all lines of business within your organization.

When you deploy Essbase on Oracle Cloud Infrastructure, you complete some initial steps for setup and configuration. For general information about getting started with Essbase and its features, see the Getting Started With Oracle Essbase documentation.

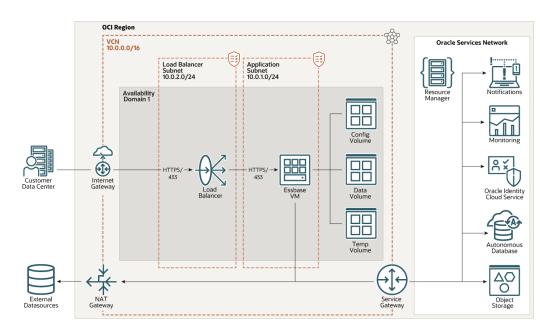
# **About Components and Terminology**

Learn about the Oracle Cloud Infrastructure components and terminology related to your setup and configuration of Oracle Essbase.

#### **Essbase Topology**

This diagram displays an example of a default, full topology of Essbase created using an Oracle Cloud Infrastructure via Marketplace deployment with Oracle Identity Cloud Service integration.





#### **Essbase Components and Terminology**

Virtual Cloud Network and Subnets: Essbase scripts assign compute instances
and load balancers to specific subnets in a virtual cloud network (VCN). A VCN in
Oracle Cloud Infrastructure covers a single, contiguous Classless Inter-Domain
Routing (CIDR) block of your choice. A subnet is a subdivision of a VCN that
consists of a contiguous range of IP addresses that don't overlap with other
subnets in the VCN. A VCN includes one or more subnets, route tables, security
lists, gateways, and Dynamic Host Configuration Protocol (DHCP) options.

Essbase scripts can automatically create a VCN and subnets for the new stack deployment, or you can create your own. By default, subnets are public. Any compute instances assigned to a private subnet can't be directly accessed from outside of Oracle Cloud.

- Load Balancer: Load balancer is an optional component that provides an extra
  layer of security, allowing the Essbase compute node to be isolated on a private
  subnet. Load balancer is recommended for supporting SSL, and provides an
  easier interface to manage the outbound SSL certificate and host name settings.
  Load balancer routes all requests from clients to a single Essbase instance. See
  Overview of Load Balancing in Oracle Cloud Infrastructure documentation.
- NAT Gateways, Subnets, and Partitions: If you set up a NAT gateway, when
  using public and private subnets, the NAT gateway needs to be added to ingress
  rules in load balancer security rules for partitions to work.
- Bastion: (optional) An OCI Bastion service instance is needed when an Essbase node is created on a private network, without a public IP. Previously, in 19.3.0.4.5, a Bastion host and compute node was created in the Essbase stack. Now, OCI Bastion service is employed. You may access a stack with private IP by making use of OCI Bastion service. Additionally, you must enable Oracle Cloud Agent (OCA) Bastion plugin on the compute node that you want to access. In order to do that, open the Essbase compute instance in OCI console, go to Oracle Cloud Agent tab and enable the Bastion toggle switch. You then need to create the Bastion provided under Identity & Security. For more information on OCA plugin, see Manage Plugins with Oracle Cloud Agent.



Bastion provides administrative access to a domain on a private subnet. Oracle recommends Bastion as a way to control external access (for example, SSH) to VCN hosts. Usually, a Bastion in a VCN public subnet controls access to VCN private subnet hosts. You can put the load balancer and Essbase on existing private subnets.

See Simplify Secure Access with OCI Bastion Service.

- Marketplace: Essbase is deployed in Oracle Cloud Marketplace, an online store
  available in the Oracle Cloud Infrastructure console. When you select Essbase from
  Marketplace, it prompts you for some basic information, directs you to Oracle Resource
  Manager to provision resources on Oracle Cloud Infrastructure, and then configures
  Essbase. See Overview of Marketplace in Oracle Cloud Infrastructure documentation.
- Resource Manager: Oracle Resource Manager provisions resources on Oracle Cloud Infrastructure for Essbase setup and configuration. See Overview of Resource Manager in Oracle Cloud Infrastructure documentation.
- Stack: A stack is a collection of related cloud resources provisioned by Resource Manager. The stack includes an Autonomous Transaction Processing instance, a compute instance, block storage volumes, object storage bucket, load balancer, and additional network components. It can include but isn't limited to the following Oracle Cloud Infrastructure components:
  - Compute instance, running the administration server and the managed server. The
    compute shape is the resources allocated to a compute instance. See Compute
    Shapes and Supported Compute Shapes.
  - Virtual cloud network (VCN), described above, which you can provide, or specify in Resource Manager to provision one for you.
  - Load balancer, described above.
  - Bastion, described above. You can use the Bastion service to gain administrative access to a domain on a private subnet.
  - Database for Essbase metadata. You have the following options for deploying the Essbase stack:
    - \* Autonomous Database using the Autonomous Transaction Processing workload. See Overview of Autonomous Database in Oracle Cloud Infrastructure documentation.
    - \* Oracle Cloud Infrastructure Database System. You must deploy Oracle Cloud Infrastructure before starting the Essbase listing. You can deploy a Virtual Machine database system. See Overview for Database System documentation.
- Notifications Service: Oracle Cloud Infrastructure Notifications service broadcasts
  messages to distributed components for applications hosted on Oracle Cloud
  Infrastructure and externally. See Notifications Overview in Oracle Cloud Infrastructure
  documentation.
- **Monitoring Service**: Oracle Cloud Infrastructure Monitoring service enables you to actively and passively monitor your cloud resources using the Metrics and Alarms features. See Monitoring Overview in Oracle Cloud Infrastructure documentation.
- Oracle Identity Cloud Service: Oracle Identity Cloud Service provides identity
  management, single sign-on (SSO), and identity governance for applications on premise,
  in the cloud, or on mobile devices. Employees and business partners can access
  applications at any time, from anywhere, and on any device in a secure manner. See
  About Oracle Identity Cloud Service in Administering Oracle Identity Cloud Service
  documentation.



#### Vault:



Prior to 19.3.0.3.4, this was referred to as Key Management, and metadata names were listed as KMS.

Oracle Cloud Infrastructure Vault enables you to manage sensitive information when creating a server domain. A vault is a container for encryption keys or secrets. Previously, in 19.3.0.4.5, you may have encrypted required passwords for a new domain using a key, and then Resource Manager used the same key to decrypt the passwords when creating the domain. We use secrets, created with the Vault UI. In your vault, you enter the password, and the latest version of it is stored as part of your key, in the vault. You refer to it using the OCID of the secret. See Overview of Vault in Oracle Cloud Infrastructure documentation, and see Create Vault, Secrets, and Encrypt Values.

- Node Manager: Java utility that runs as a separate process from Oracle WebLogic Server and allows you to perform common operations for a Managed Server, regardless of its location with respect to its Administration Server.
- Administration Server: Operates as the central control entity for the configuration
  of the entire domain. It maintains the domain's configuration documents and
  distributes changes in the configuration documents to Managed Servers. The
  Administration Server serves as a central location from which to monitor all
  resources in a domain. Each domain must have one server instance that acts as
  the Administration Server.
- Managed Server: Host business applications, application components, Web services, and their associated resources.

# Administrator Access Requirements

Learn about different kinds of administrative access requirements for setting up Oracle Essbase on Oracle Cloud Infrastructure.

Administrator Type	Description
Oracle Cloud Infrastructure administrator	The administrator who is subscribed to the Oracle Cloud Infrastructure tenancy. This administrator sets up policies for access to resources in the compartment, and also designates the other administrators (listed below) during stack deployment.
IDCS system administrator	The Oracle Identity Cloud Service identity domain administrator. This administrator is responsible for setting up Essbase access in Identity Cloud Service, including creating and modifying a confidential IDCS application.



Administrator Type	Description
IDCS Essbase admin user	An Identity Cloud Service user selected to act as the initial Service Administrator for Essbase. During provisioning and deployment of the stack, the Oracle Cloud Infrastructure administrator provides the ID of this user. This user can log in to Essbase and provision other users.
Essbase admin user	This user is specified by the Oracle Cloud Infrastructure administrator during provisioning and deployment of the stack. This is a native Essbase administrator used during deployment. This user ID provides an alternate way to log in to Essbase for administration of the Essbase environment/deployment.

# Typical Workflow for Administrators

Use this workflow as a high-level guide to administrator tasks for Oracle Essbase.

Task	Description	More Information
Address prerequisites prior to deployment	Know and perform the prerequisites, values, and tasks needed before deployment, including access, security, and resources	Before You Begin with Oracle Essbase
Deploy and configure Essbase and its required stack	Log into Oracle Cloud Infrastructure and select Essbase in Oracle Cloud Marketplace.	Set up Oracle Essbase
	Enter the required metadata, and select the options that you prefer in Oracle Resource Manager setup wizard.	
Perform post-deploment tasks	Complete the required post-deployment tasks, including security, access, and resource cleanup	Complete Post-Deployment Tasks
Migrate existing Essbase deployments, data and content	Move data from Essbase 11g On- Premise or cloud services to Essbase.	Migrate Applications and Users
Set up users	Set up roles for your users and assign them appropriate privileges.	Manage Users and Roles
Monitor performance and collect diagnostics	Monitor the operation of Essbase.	Collect Diagnostic Information on the Essbase Node
Patch an instance	Apply a patch or roll back a patch.	Patch and Roll Back
Back up an instance	Perform regular backups to protect content and allow you to restore it.	Back Up and Restore Essbase



# Set Up Oracle Essbase

Let's explore the process to deploy and configure Oracle Essbase.

#### **Topics:**

- Before You Begin with Oracle Essbase
- Deploy Essbase
- Complete Post-Deployment Tasks

# Before You Begin with Oracle Essbase

Before you begin to set up Oracle Essbase deployment, here are pre-requisite lists of metadata you must gather and tasks that you must complete.

The quick-start process, which deploys Essbase on Oracle Cloud Infrastructure using Marketplace, uses default settings on Marketplace. The process assumes, and at times provides, less prohibitive access to infrastructure components. You're recommended to use the information here, as well as the default settings, only as a guide, and to determine the appropriate security and access requirements for your organization. You can also use Oracle Cloud Infrastructure documentation as a reference.

When Essbase is deployed on Oracle Cloud Infrastructure, you receive access to the required services based on your defined role and policies, including Oracle Cloud Infrastructure Compute Service and Oracle Identity Cloud Service.

A text worklist is provided at the end of this page, which you can copy to a text file and use for storing names, IDs, and other values needed during setup.

Table 2-1 Pre-deployment metadata

Prerequisite Metadata	Links to overviews and tasks / examples	Record values needed for Verify completion deployment	
Account and environment for Oracle Cloud Infrastructure (OCI)	-	Account	
Command Line Interface (CLI) tool installed from OCI	CLI Quickstart	-	
User name for OCI administrator	Administrator Access Requirements	Admin user name and password	
Identity Cloud Service (IDCS) system administrator user ID (or create it later with REST API for IDCS)	Administrator Access Requirements	IDCS system admin user ID	



Table 2-1 (Cont.) Pre-deployment metadata

Links to overviews and tasks / examples	Record values needed for deployment	Verify completion
Administrator Access Requirements	Initial Essbase administrator name	
Regions and Availability Domains	Region Availability domain	
Create Vault, Secrets, and Encrypt Values, Also see, Overview of Vault in Oracle Cloud Infrastructure documentation.	Vault encryption key OCID	
Overview of Load Balancing	Load balancer shape and subnets	
Supported Compute Shapes	Compute shape	
Creating a Virtual Cloud Network	Existing virtual cloud network Name	
Administrator Access Requirements	Also known as Essbase 911 user name (administrator who manages the WebLogic server on which Essbase runs)	
-	If creating a private Essbase subnet, record the Essbase node private IP and the VCN on which it is deployed.	
-	Essbase URL	
Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service		
Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service	Notifications topic OCID	
	tasks / examples  Administrator Access Requirements  Regions and Availability Domains  Create Vault, Secrets, and Encrypt Values, Also see, Overview of Vault in Oracle Cloud Infrastructure documentation.  Overview of Load Balancing  Supported Compute Shapes  Creating a Virtual Cloud Network  Administrator Access Requirements  -  Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service  Get Event Notifications Using Oracle Cloud Infrastructure Notifications	Administrator Access Requirements  Regions and Availability Domains  Region Availability domain  Create Vault, Secrets, and Encrypt Values, Also see, Overview of Vault in Oracle Cloud Infrastructure documentation.  Creating a Virtual Cloud Network  Administrator Access Requirements  Also known as Essbase 911 user name (administrator who manages the Webl_ogic server on which Essbase node private IP and the VCN on which it is deployed.  Essbase URL  Monitor Operations and Resources Using Oracle Cloud Infrastructure Notifications Usal Mail Initial Essbase administrator name Availability domain  Vault encryption key OCID  Availability domain  Vault encryption key OCID  Notifications topic OCID  Notifications topic OCID



Table 2-2 Pre-deployment tasks to be completed

Prerequisite Tasks	Links to overviews and tasks / examples	Record values needed for deployment	Verify completion
Select one of these database options:     Oracle Autonomous Database deployed by Oracle     Existing Oracle Autonomous Database (shared (ADB-S) or dedicated (ADB-D)) that you deployed from Oracle Cloud Infrastructure Console     Existing Database System that you deployed from Oracle Cloud Infrastructure	<ul> <li>Creating an Autonomous Database</li> <li>Creating a Database System         <ul> <li>How to Use Network Components</li> <li>Set Up Rules for Database Connectivity</li> <li>Set Relational Database Connectivity</li> </ul> </li> </ul>	<ul> <li>For Oracle-created database: admin name and password</li> <li>For existing Autonomous database: admin name and password, and compartment</li> <li>For existing Database System: admin name and password, database name and home, and compartment</li> </ul>	
2. As shown in the table above, if you haven't already done so, use Oracle Cloud Infrastructure Vault provisioning secrets and OCIDs for secrets created for:  Essbase administrator password  Database system administrator password  Identity Cloud Service application client secret	Create Vault, Secrets, and Encrypt Values	Vault encryption key OCID. The OCIDs need to be noted for Essbase administrator password, Database system administrator password, and Identity Cloud Service application client secret	
3. Log in to OCI tenancy as Administrator - In OCI console, log in to your tenancy as the admin subscribed to that tenancy. [Optional]: Create OCI admin user - Add user to admin group (Identity>Groups>Admin istrator>Create User), close browser, change password using email you receive, and log in as new admin.	Administrator Access Requirements	-	



Table 2-2 (Cont.) Pre-deployment tasks to be completed

Prerequisite Tasks	Links to overviews and tasks / examples	Record values needed for deployment	Verify completion
4. Create an SSH key pair - In Oracle Cloud Infrastructure (OCI) console, create SSH public key and corresponding private key to access Essbase compute instances.	Creating a Key Pair	SSH public key Path to private key	
5. Create compartment(s) - In OCI console, choose or create a compartment (Identity>Compartments > Create Compartment) where you want to deploy Essbase.	Choosing a Compartment	Compartment ID (OCID) and name	
6. Create dynamic group - In OCI console, create a dynamic group to allow resources to allow OCI resources to be created and networked together dynamically without explicit approvals. You can associate groups with policies.	Create Dynamic Groups	Dynamic group name	
7. Set up policies - In OCI console, set up policies to enable you to manage or create resources in OCI.	Set Up Policies See also Common Policies and How Policies Work.	Policy statements (enter in text worklist doc for entry convenience)	
8. Set up Essbase access - In Oracle Identity Cloud Service (IDCS), set up Essbase access.	Set Up Essbase Access in Identity Cloud Service	-	
9. Set up confidential application to register Essbase - In Oracle Identity Cloud Service, for each compartment in which you plan to deploy Essbase, create a confidential Oracle Identity Cloud Service application, and activate the confidential application.		Confidential application name IDCS Instance GUID IDCS Application Client ID IDCS Application Client Secret	

### **Storing Recorded Metadata for Deployment**



Copy and paste the following to a text file, for your convenience, and enter the relevant values, to be used during deployment. If you fail to record any needed deployment metadata, an Oracle Cloud Infrastructure administrator can collect them from the Variables page or Application Information page of the Oracle Resource Manager. Make sure to protect and dispose of the metadata text file appropriately.

```
Region:
Target availability domain:
Oracle Cloud Infrastructure administrator user:
OCI administrator group name:
IDCS system administrator:
DB system admin password:
Existing virtual cloud network name (optional):
Compute node shape:
Essbase node OCID (OCI ID of the compute node):
Load balancer shape and subnets (optional):
SSH-2 RSA key pair, stored locally:
Path to private key:
Compartment ID:
Compartment name:
Dynamic group name:
Policy statements:
Confidential application details
Name:
 IDCS instance GUID (IDCS host):
 IDCS application client ID:
 IDCS application client secret:
IDCS Essbase admin user and password (Initial Essbase admin):
Essbase admin user (also known as WebLogic user--defined during stack
creation):
SSH access details:
  Essbase node private IP (when creating private subnet):
  Essbase node public IP (when not creating private subnet):
  Essbase URL (for web interface):
  Essbase IP (for confidential application):
Notification topic OCID (optional):
```

# Create Dynamic Groups

You create dynamic groups of Oracle Cloud Infrastructure compute instances, and associate them with policies. You must provide a unique, unchangeable name for the dynamic group. Oracle assigns a unique Oracle Cloud ID (OCID).

- On the Oracle Cloud Infrastructure console, navigate to the left icon under the Governance and Administration section, click Identity, and then click Dynamic Groups.
- 2. Click Create Dynamic Group.
- 3. Enter a name for the dynamic group. Record the name for future use.
- 4. Enter a description (optional).
- 5. In the Matching Rules section, click Launch Rule Builder to define the rule.
- In Add instances that match the following rules. Rules to consider for match, select the option Any of the following rules, allows the dynamic group broad access if you

- have multiple rules. The other option, **All of the following rules**, enables the dynamic group to service specific compartment/instance combinations when multiple rules are specified.
- 7. In the Create Matching Rule dialog, select a resource. In ATTRIBUTE field, select the option Match Instances in Compartment ID. For the VALUE field, paste the compartment ID you noted for the compartment in which you're creating the Essbase stack. With this option, any instances in the compartment will work using this dynanic group. The other option, Match instances with ID, specifies matching just one instance ID.
- 8. Click Add Rule.
- 9. You can enter tags (optional) to organize and track resources in your tenancy.
- 10. Click Create Dynamic Group.

For more information on dynamic groups, see Managing Dynamic Groups.

### Set Up Policies

A policy is a document that specifies who can access which Oracle Cloud Infrastructure resources that your company has, and how.

Before deploying the Essbase stack on a compartment in Oracle Cloud Infrastructure, the tenant administrator must set up policies to access or create the following resources in the selected compartment:

- Marketplace applications
- Resource Manager stacks and jobs
- Compute instances, networks, and load balancers
- Database for storing Essbase metadata
- Managing and using virtual keys for Oracle Cloud Infrastructure Vault

#### To create policies:

- On the Oracle Cloud Infrastructure console, navigate to the left icon under the Governance and Administration section, click Identity, select Policies, select the root compartment, and then click Create Policy.
- 2. Provide a name and description for the policy.
- 3. Add a policy statement (Allow) for each instance in the compartment. Copy them from your text worklist file. Specify the group\_name in the policy statement.
- 4. When done, click Create.

Create a policy each, for both groups and dynamic groups, as necessary.

For Bastion policy information, see Bastion Policies.

Set up policies that are appropriate for your organization's security setup. The following is an example of a policy template, with each row being a policy statement.

```
Allow group group_name to manage orm-stacks in compartment compartment_name
Allow group group_name to manage orm-jobs in compartment compartment_name
Allow group group name to manage virtual-network-family in compartment
```



```
compartment name
Allow group group name to manage instances in compartment compartment name
Allow group group name to manage volume-family in compartment
compartment name
Allow group group name to manage load-balancers in compartment
compartment name
Allow group group name to manage buckets in compartment compartment name
Allow group group name to manage objects in compartment compartment name
Allow group group name to manage autonomous-database-family in compartment
compartment name
Allow group group name to use instance-family in compartment compartment name
Allow group group name to manage autonomous-backups in compartment
compartment name
Allow group group name to manage buckets in compartment compartment name
Allow group group name to manage vaults in compartment compartment name
Allow group group name to manage keys in compartment compartment name
Allow group group name to manage secret-family in compartment
compartment name
Allow group group name to manage app-catalog-listing in compartment
compartment name
```

Some policies may be optional, depending on expected use. For example, if you're not using a load balancer, you don't need a policy that allows management of load balancers.

To allow instances within the compartment to invoke functionality without requiring further authentication, you must have group policies for the instances in the compartment. To do this, create a dynamic group, and set the policies for that dynamic group, such as shown in the following example:

```
Allow dynamic-group group name to use autonomous-database in compartment
compartment name
Allow dynamic-group group name to use secret-family in compartment
compartment name
Allow dynamic-group group name to use keys in compartment compartment name
Allow dynamic-group group name to read buckets in compartment
compartment name
Allow dynamic-group group name to manage objects in compartment
compartment name
Allow dynamic-group group name to inspect volume-groups in compartment
compartment name
Allow dynamic-group group name to manage volumes in compartment
compartment name
Allow dynamic-group group name to manage volume-group-backups in compartment
compartment name
Allow dynamic-group group name to manage volume-backups in compartment
compartment name
Allow dynamic-group group name to manage autonomous-backups in compartment
compartment name
Allow dynamic-group group name to manage database-family in compartment
compartment name
```

The following policies are optional, but necessary for the following integrations:

Oracle Notification Service integration:

allow dynamic-group group\_name to use ons-topic in compartment dev where request.permission='ONS TOPIC PUBLISH'

Oracle Cloud Infrastructure Monitoring integration:

allow dynamic-group group\_name to use metrics in compartment dev where target.metrics.namespace='oracle essbase'

# Set Up Essbase Access in Identity Cloud Service

To set up security and access, you integrate Essbase with Oracle Identity Cloud Service. You provision Essbase users using Essbase roles, rather than Oracle Identity Cloud Service roles.

To prepare security access for Essbase, you must log in to Oracle Identity Cloud Service as the identity domain administrator and complete a few tasks.

Before you can provision users and groups in Essbase, you need, during creation of the Essbase stack, to provide the name of a user in Oracle Identity Cloud Service who will be the initial Service Administrator for Essbase.

This Service Administrator can then log in to the Essbase web interface to provision other users.

You also need to provide access to the signing certificate.

Complete the following tasks in Identity Cloud Service before deploying the Essbase stack.

- Log in to Identity Cloud Service as the identity domain administrator. To get to the Identity Cloud Service console from Oracle Cloud Infrastructure, click Identity, then Federation, and click on the URL link next to Oracle Identity Cloud Service Console.
- 2. In the Identity Cloud Service console, expand the navigation drawer icon, click **Settings**, and then click **Default Settings**.
- **3.** Turn on the switch under **Access Signing Certificate** to enable clients to access the tenant signing certificate without logging in to Identity Cloud Service.
- 4. Scrill up and click **Save** to store your changes.
- 5. If not already created, create a user in Identity Cloud Service who will be the initial Essbase Service Administrator.

#### **About Single Sign-On (SSO)**

If you use single sign-on (SSO) with Identity Cloud Service, your Essbase login screen routes to Identity Cloud Service.

If you use SSO that is external to Identity Cloud Service, you configure Identity Cloud Service to point to the external security provider. The Essbase login screen routes to Identity Cloud Service, which routes to the external login screen. After logging in, you're directed back to the Essbase web interface.



### Set Relational Database Connectivity

Before you can configure Essbase, you need network connectivity to a relational database where the Essbase and Fusion Middleware RCU schemas reside.

Oracle recommends deploying a distinct pluggable database (PDB) for Essbase. You can read about Oracle's multitenant architecture here: Introduction to Multitenant Architecture.

- No other applications should have access to the Essbase repository schemas generated by the Repository Creation Utility (RCU).
- No one else other than the designated administrator should have permission to access the schemas or their tables.
- No one else should have the credentials to assign or change roles to access the PDB.
- Every change to the PDB should be logged.

# Create a Confidential Identity Cloud Service Application

Before deploying the Essbase stack, create a confidential application in Oracle Identity Cloud Service and register Essbase with it.

- Open the Oracle Identity Cloud Service Console. From Oracle Cloud Infrastructure, select Identity, Federation, Identity Provider Details. In the Identity Provider Information tab, click the Oracle Identity Cloud Service Console link.
- 2. In the Identity Cloud Service console, expand the Navigation Drawer icon, and then select **Applications**.
- 3. Select +Add.
- 4. Select **Confidential Application**, as described in Add a Confidential Application.
- 5. In the **App Details** step, enter a name only, and then select **Next**. Tip: you may use the same name as the compartment, as you need one confidential application per compartment. Record the name for your information.
- **6.** In the **Client** step, select the option **Configure this application as a client now**.
- 7. In the Authorization section,
  - Select the following allowed grant types: Client Credentials and Authorization Code.
  - If you don't plan to provision a load balancer, select Allow non-HTTPS URLs.
    - a. For the Essbase Redirect URL, enter a temporary/mock redirection URL (it ends with \_uri):

```
http://temp/essbase/redirect uri
```

b. For the Essbase Post Logout Redirect URL, enter a temporary/mock URL:

```
http://temp/essbase/jet/logout.html
```

 Otherwise, if you're provisioning a load balancer, enter the following: above URL, but using https:, as shown.



a. For the **Essbase Redirect URL**, enter a temporary/mock redirection URL:

https://temp/essbase/redirect uri

b. For the Essbase Post Logout Redirect URL, enter a temporary/mock URL:

https://temp/essbase/jet/logout.html

- 8. Under Token Issuance Policy, in the section Grant the client access to Identity Cloud Service Admin APIs, click Add, find and select the Identity Domain Administrator role, and select Add.
- Scroll to the top of the page and click Next until you reach the Authorization section.
- 10. Click Finish.
- 11. From the Application Added popup window, record the following Identity Cloud Service details: IDCS Application Client ID and IDCS Application Client Secret. Record these values to use during your Essbase deployment.
- 12. Record the IDCS Instance GUID from the following location: in the Identity Cloud Service Console, select your ID icon in the top right corner (the icon contains your initials), select **About**, and record the **IDCS Instance GUID** value. If you don't have access, ask your administrator to provide it. Example: idcs-123456789a123b123c12345678d123e1. Alternatively, the **IDCS Instance GUID** is at the front of the IDCS url in the browser take the host portion of the url.
- 13. Select **Activate** in the title bar, next to your application's name.

Oracle Cloud Services accounts provides Oracle Identity Foundation, which enables basic identity services functionality. This includes user management, group management, basic reporting, and authentication for Oracle applications. See: Oracle PaaS and IaaS Universal Credits Service Descriptions. For information regarding features available in various Oracle Identity Cloud Service versions, see: About Oracle Identity Cloud Service Pricing Models.

# Supported Compute Shapes

Oracle Essbase offers compute sizes (OCPUs) to suit different scenarios. The larger the compute size, the greater the processing power. If you're not sure which size to use, contact your sales team to discuss sizing guidelines.

Essbase can be Oracle Compute Unit (OCPU)-intensive depending on your application. The minimum number of OCPUs recommended for production deployments is 4 OCPUs. To help you decide which compute size best suits your deployment, consider how many active users you expect to perform concurrent activities such as:

- Users running queries in hybrid mode
- Users running calculations in block storage mode
- Users running reports or queries in aggregate storage cubes

You can configure storage size when you deploy Essbase. Determine the storage size needed, or consult with your sales team to determine that your storage needs are met, based on the number of applications that you plan to deploy.



Essbase currently supports the following shapes:

- VM.Standard.2.n
- VM.Standard.E2.n
- VM.Standard.E3.Flex
- VM.Standard.E4.Flex
- BM.Standard.2.52
- BM.Standard.E2.64
- BM.Standard.E3.128
- BM.Standard.E4.n

For a description of the difference between VM and BM shapes, and a discussion on how to decide which to use, see <a href="https://cloud.oracle.com/compute/fag">https://cloud.oracle.com/compute/fag</a>.

### Create Vault, Secrets, and Encrypt Values

Oracle Cloud Infrastructure Vault enables you to manage sensitive information when creating a server domain. A vault is a container for encryption keys and secrets.

In 21c and 19.3.0.5.6, you use secrets, created with the Vault UI. In your vault, you enter the password, and the latest version of it is stored as part of your key, in the Vault. You refer to it using the OCID of the secret. See Overview of Vault in Oracle Cloud Infrastructure documentation,



If you're using an existing vault and have an encryption key already created, you can skip **Create a New Vault** and **Create a New Encryption Key** sections and move to the **Create a New Secret** section. Otherwise, you must create a vault and key first.

When you use Vault to encrypt credentials during provisioning, you need to create a secret. Passwords chosen for Essbase administrator and Database must meet the Resource Manager password requirements.

Secrets need to be added for the following fields:

- Essbase Administrator Password
- IDCS application client secret
- Database system administrator password

#### Create a New Vault:



These steps explain how to create a virtual Vault, which is a lower-cost option than a private vault. New entities are needed only if they have not already been created.



- 1. Sign in to the Oracle Cloud Infrastructure Console.
- 2. In the navigation menu, select **Security**, and click **Vault**.
- 3. Select your **Compartment**, if not already selected.
- 4. Click Create Vault.
- 5. For Name, enter OracleEssbaseVault.
- 6. For the lower-cost option, leave unchecked the option to make it a private vault.
- 7. Click Create.

#### Note:

The Vault Crypto Endpoint value can be retrieved for any future use, by clicking at any time on the newly created vault, as listed on the Vaults page.

#### Create a New Encryption Key

Go to the Vaults page, and create a new encryption key as follows.

- 1. Select the newly created vault, for example, OracleEssbaseVault from the previous section.
- 2. Select Master Encryption Keys in the left panel.
- **3.** Select **Create Key** and provide a name for the key, such as OracleEssbaseEncryptionKey.
- 4. Click **Create Key**. This key is used during secret creation.

#### **Create a New Secret**

Go to the Vaults page. For each password, create a secret as follows.

- 1. Click Secrets.
- 2. Click Create Secrets.
- 3. Enter the Name for the secret and a relevant Description.
- 4. Select the new encryption key (created in the previous section), or an existing one. For example, OracleEssbaseEncryptionKey.
- **5.** Enter the password text in **Secret Contents**.
- 6. Click Create Secret.
- For each created secret, click the related password and copy the OCID value for it, for later use in configuration.

#### To encrypt your Oracle Essbase Administrator password (for older versions):

1. Convert the administrator password that you want to use for the Essbase domain to a base64 encoding.

For example, from a Linux terminal, use this command:

```
echo -n 'OracleEssbase Password' | base64
```



- 2. Run the encrypt oci command using Oracle Cloud Infrastructure command line interface. Provide the following parameters:
  - Vault Encryption Key OCID
  - Vault Crypto Endpoint
  - base64-encoded password

```
oci kms crypto encrypt --key-id Key_OCID --endpoint Cryptographic_Endpoint_URL --plaintext Base64_OracleEssbase_Password
```

From the output, copy the encrypted password value for use in the deploy process, as shown here:

```
"ciphertext": "Encrypted Password"
```

You also use vault encryption to encrypt your Database Password and your Client Secret.

# **Deploy Essbase**

Deploy Oracle Essbase from Oracle Cloud Marketplace.

As the Oracle Cloud Infrastructure administrator, you use Oracle Cloud Infrastructure to set up Essbase. Oracle Cloud Marketplace uses Oracle Resource Manager to provision the network, compute instances, Autonomous Transaction Processing database for storing Essbase metadata, and Load Balancer.

During this process, you'll need to provide other administrator user IDs. Review Administrator Access Requirements to understand what these administrator accounts can do.

- Read prerequisites and requirements that you need to know or do before deployment.
   See Before You Begin with Oracle Essbase.
- Sign into Oracle Cloud Infrastructure console as the Oracle Cloud Infrastructure administrator.
- 3. From the navigation menu, select **Marketplace**.
- 4. On Oracle Marketplace page,
  - a. In the title bar, select or accept the region from which to run the deployment.
  - b. In the Category dropdown menu, select **Database Management**.
  - c. Under All Applications, select Oracle Essbase.
  - d. Select the stack version, or accept the default.
  - From the dropdown menu, select the target Compartment that you created for Essbase, in which to create the stack instance.
  - f. Select the check box to indicate you accept the Oracle Standard Terms and Restrictions.
  - g. Click Launch Stack.
- 5. In **Stack Information**, on the Create Stack page.
  - a. For My Configuration, the terraform configuration source files to be uploaded, select Zip File (instead of the default Folder option). If necessary, drop or browse to the stack zip file. The stack name is displayed.



- **b.** Enter the stack description and other stack information, as necessary.
- c. Click Next.
- 6. In **General Settings**, on the Configure Variables page, you configure variables for the infrastructure resources that the stack creates. [Optional] Enter **Stack Display Name** value to identify your stack deployment for all generated resources, for example <code>essbase\_<userid></code>. Provide a meaningful stack display name. This name is used as a dimension for filtering Essbase metrics that correspond to components in this stack. If not entered, the display name is generated. The target compartment you previously selected is shown.

#### 7. In Essbase Instance:

- **a.** Select an **Essbase Availability Domain** in which to create the Essbase compute instance.
- **b.** Select the **Essbase Instance Shape** for the Essbase compute instance.
  - If **VM.Standard.En.Flex** compute shape is selected (new for 19.3.0.5.6), additional entry fields are displayed:
  - **Essbase Instance OCPUs** enter number of OCPUs to be used for the Essbase compute instance: values of 1- 64 are allowed; but four or more are recommended for production workloads);
- c. Enter the **Data Volume Size** or accept the default. The minimum value is 256GB.
- **d.** Paste the value of the **SSH Public Key** that you created, to access the Essbase compute instance.
- e. In the Essbase System Admin User Name field, enter an Essbase administrator user name you can optionally use the Identity Cloud Service user name. It provides an additional way to log in to Essbase, and is also the administrator used to Access the WebLogic Console on which Essbase runs. If you don't enter an Identity Cloud Service user in this field, you must provide one in the IDCS Essbase Admin User field later in the stack definition, in the Security Configuration section. If you enter an Identity Cloud Service user in this field, the Identity Cloud Service System Administrator User ID is optional in the Security Configuration section.
- f. In the Essbase System Admin User Password field, enter an OCID for the Vault secret that contains the password for the Essbase system admin user. See Create Vault, Secrets, and Encrypt Values. (For 19c through 19.3.0.4.5, this field contained the encrypted admin password you created using the CLI tool).
- g. [Optional] Enter Essbase Instance Timezone.

#### 8. In Monitoring Configuration

- [Optional] Enter **Notification Topic OCID**, to which messages are published. For information on how to enable notifications, see Notifications Overview.
- [Optional] Select Enable Monitoring to support publishing of metrics to the Monitoring Service.

#### 9. In Identity Configuration:

a. For Identity Provider, select IDCS. To set up security and access, you integrate Essbase with Identity Cloud Service as part of the stack deployment. The Embedded LDAP option isn't recommended or supported for production workloads.



- b. Enter the IDCS Instance GUID, IDCS Application Client ID, and IDCS Application Client Secret values, which you recorded as pre-deployment requirements, after you created a confidential Identity Cloud Service Application.
- c. Enter IDCS Essbase Admin User value. This can't be the same user ID as the Essbase administrator. Additionally, this ID must already exist in the Identity Cloud Service tenancy. If you don't provide this user ID during stack creation, or if it's mapping to the initial Essbase administrator doesn't done correctly, you can later use the Identity Cloud Service REST API to create this user and link it to Essbase. See REST API for Oracle Identity Cloud Service.
- **10.** In **Database Configuration**, select from the following options and then perform the configuration tasks:

#### **Database options and considerations:**

- If you didn't already, review recommendations and rules regarding database connectivity. See Set Relational Database Connectivity.
- If you plan to use the Oracle Autonomous Database deployed automatically by the stack, you must provide the OCID for the Vault secret that contains the DB password. (For 19c through 19.3.0.4.5, this field contained the encrypted DB password you created using the CLI tool). Select the database license or accept the default.
- If you plan to use an existing Oracle Autonomous database, and select Use Existing Database, provide the OCID for the Vault secret that contains the DB password. This is the same password that you have provided during the DB service creation. Specify the compartment in which Autonomous Transaction Processing was created. (For 19c through 19.3.0.4.5, this field contained the encrypted DB password you created using the CLI tool).
- If you plan to use Federated Partitions to Autonomous Data Warehouse, you must provide an Autonomous Data Warehouse on Shared Infrastructure instance to host the Essbase RCU schema and fact table sources for the Federated Partition. Select Use Existing Database option to deploy to your instance of Autonomous Data Warehouse on Shared Infrastructure.
- To use an existing Oracle Cloud Infrastructure Database System for the internal Essbase repository, select the option **Database System** for Database Type, and specify the compartment and database details. The database must be accessible to the created compute node. If the database has a private IP, use the existing network option where the network is set up, to allow for traffic between the subnet that hosts the compute node and the subnet that hosts the database. See Bare Metal and Virtual Machine Database Systems.

#### 11. In Network Configuration:

- a. If you chose Use Existing Network, select the name of the existing virtual cloud network. You can still create a new instance of the Autonomous Transaction Processing database.
- b. If you want to create a new virtual cloud network, enter a **Virtual Network CIDR** value to assign the VCN. See Overview of Networking.
- **c.** Select the target network compartment, virtual network, and application subnet.
- **d.** If you want to create a private Essbase subnet, enter an **Application Network CIDR** to assign to the subnet for the target Essbase compute node.
- e. Select a subnet strategy: use an existing public subnet or select **Create a Private Essbase Subnet** for the Essbase node.



f. [Optional] Select Public Essbase Node Visibility to enable a public IP address for the Essbase instance. If selected, the subnet provided must allow for public IP address.

#### 12. In Load Balancer Configuration:

- **a.** Select **Provision Load Balancer** to provision it in Oracle Cloud Infrastructure with a demo certificate. This is not recommended for production workloads.
- b. Select Public Load Balancer Visibility to enable a public IP address for the Load Balancer, and to add an extra layer of security. Select a load balancer shape.

#### 13. In Bastion Configuration:

#### For 19c through 19.3.0.4.5 if Public Essbase Node Visibility is not set:

- a. Select **Provision Bastion** to enable the creation of a bastion.
- Select a Bastion Availability Domain to provide the target availability domain of the bastion.
- c. Select a Bastion Instance Shape. You must have the capacity of the target shape in the given availability domain for the bastion compute instance to be created successfully. Your bastion shape value doesn't need to match the compute node shape.

#### For 19.3.0.5.6:

When you deploy a stack with private IP, a Bastion is used to access it and you're required to enable Oracle Cloud Agent (OCA) Bastion plugin on the compute node. In order to do that, open compute instance in OCI, go to Oracle Cloud Agent tab, and enable the Bastion toggle switch. For more information on OCA plugin, see Manage Plugins with Oracle Cloud Agent. Bastion creation and configuration doesn't need to be done during deployment. It can be done later, when access is needed. See Access Oracle Essbase Using SSH.

- 14. On the Review page, review the information that you provided, and click. Then run the apply job, by clicking Apply. The Job Information tab in Oracle Resource Manager shows the status until the job finishes and the stack is created. This can be modified, as job status only shows us the status of OCI resources created and allotted. To check stack configuration, use Monitoring by providing notification of OCID or SSH into the image.
- 15. Check for any log errors. If necessary, see Troubleshoot Deployment Errors.
- **16.** On the Review page, the value for **essbase\_url** is used in the browser to access Essbase. The **essbase node public ip** is for accessing SSH.
- 17. After you complete deployment, then complete the post-deployment tasks, including: modifying your created Identity Cloud Service application, testing connectivity to Essbase, and the other listed tasks.

You can modify the created resources and configure variables later. Logs are created that can be forwarded to Oracle Support, if necessary for troubleshooting. After deployment, you're ready to assign users to roles and permissions in the Essbase web interface. You can also perform additional network and security configuration.

#### **Reviewing or Collecting Output After Deployment**

If you didn't keep a record of all of the deployment output, an Oracle Cloud Infrastructure administrator can collect them from the Variables page or Application Information of the Oracle Resource Manager, as well as in the client configuration details of the Identity Cloud Service confidential application.



#### Viewing the Deployment

Log into Oracle Cloud Infrastructure console, go to Resource Manager for your compartment, and view the details for the Essbase stack you created. From there, if you click on the apply job, you can see the deployment log and output details. If you selected to use a load balancer, its public IP is in the essbase\_url. For 19c through 19.3.0.4.5, if you have deployed a bastion host, the outputs include a <code>bastion\_host\_public\_ip</code> and there isn't an <code>essbase\_node\_public\_ip</code>.

#### Viewing the Variables

In addition to using the log to find and record deployment details, you can also view most of them in the Variables page or the Application Information page of the Resource Manager. If you selected to use a load balancer, **create\_load\_balancer** is true.

#### Viewing the Confidential Application Configuration

To locate the client secret, which is masked in Resource Manager, an Identity Cloud Service Administrator can go to the Identity Cloud Service Console, select the confidential application, and view its configuration.

# Complete Post-Deployment Tasks

After you deploy Oracle Essbase on Oracle Cloud Infrastructure using Marketplace, complete the following tasks.

- Modify the Confidential Identity Cloud Service Application
- Set Up the SSL Certificate
- Secure Your Network
- Test Connectivity to Essbase
- Complete System Hardening and Cleanup Tasks
- Troubleshoot Deployment Errors

For information on setting up a trusted certificate authority, see Managing SSL Certificates.

### Modify the Confidential Identity Cloud Service Application

After deploying the Essbase stack from Oracle Cloud Marketplace, update your confidential Identity Cloud Service application with the correct Essbase URLs.

- Log in to Identity Cloud Service as the identity domain administrator. To get to the Identity Cloud Service console from Oracle Cloud Infrastructure, click Identity, then Federation, and click on the URL link next to Oracle Identity Cloud Service Console.
- 2. In the Identity Cloud Service console, expand the Navigation Drawer icon, and then click **Applications**.
- 3. Locate and select your confidential application.
- 4. Select Configuration and expand Client Configuration.
- 5. Update the Essbase Redirect URL to reflect the actual Essbase URL.

https://192.0.2.1/essbase/redirect\_uri



If you deployed a load balancer, include the port number:

```
https://192.0.2.1:443/essbase/redirect uri
```

Note that if you deployed a load balancer, the IP in the Essbase URL will be for the load balancer.

6. Update the Essbase Post Logout Redirect URL to reflect the essbase\_url. For example:

```
https://192.0.2.1/essbase/jet/logout.html
```

7. Scroll up and save the updated confidential application.

# Set Up the SSL Certificate

After you deploy the Essbase stack, Oracle highly recommends that you update the SSL certificate, using the Oracle Cloud Infrastructure console or APIs, to one that has been signed with a trusted certificate authority.

For information on setting up a trusted certificate authority, see Managing SSL Certificates.

If you select to provision the Oracle Cloud Infrastructure Load Balancer during the Essbase stack provisioning process, the Load Balancer is configured with a demo certificate you can use for SSL access. The demo certificate is self-signed.

When you use a self-signed certificate, including the provided demo certificate, you must perform additional configuration to enable the use of partitions, as well as Essbase C- and Java-based clients. MaxL is a C-based client. You also need to ignore hostname verification on the WebLogic part of the Essbase stack. **Caution**: use of self-signed certificates should be only temporary, until you can obtain a trusted CA certificate.

#### **Steps for Using Partitions with Self-Signed Certificates**

When you use a self-signed certificate, you must perform additional configuration and also disable peer certificate verification, to enable the use of partitions.

- Access the Essbase node using SSH, as described in Access Oracle Essbase Using SSH.
- 2. Change to oracle user.

```
sudo su - oracle
```

3. Open essbase.cfg for editing.

```
vi /u01/config/domains/essbase_domain/config/fmwconfig/essconfig/
essbase/essbase.cfg
```

4. Add the following variable to the bottom of the file.

```
env: API DISABLE PEER VERIFICATION 1
```



#### Steps for Using MaxL with Self-Signed Certificates

- When you use a self-signed certificate, you must perform configurations to enable the use of MaxL.
  - Either use MaxL client, following the instructions in Manage Essbase Using the MaxL Client.
  - OR use MaxL on the server, using the startMAXLsh file at the following path on the server:

```
/ \verb"u01/config/domains/essbase_domain/esstools/bin"
```

- 2. In order to use self-signed certificate, peer verification should be disabled, by setting the environment variable API\_DISABLE\_PEER\_VERIFICATION=1.
  - In Linux, edit the MaxL startup script (startMAXL.sh) and add the following line:

```
export API_DISABLE_PEER_VERIFICATION=1
```

• In Windows, edit start maxl script (startMAXL.bat) and add the following line:

```
set API DISABLE PEER VERIFICATION=1
```

#### Steps for Using Java-based Clients with Self-Signed Certificates

When you use a self-signed certificate and a Java client, you must configure your Java client.

- 1. From an external host:
  - a. When Load Balancer was configured:
     Download the certificate provided with the Oracle Cloud Infrastructure Load Balancer.

```
echo -n | openssl s_client -connect < LOAD BALANCER IP>:443 | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/lbr.cert
```

b. When Load Balancer wasn't configured: Download the certificate as follows.

```
echo /p' > /tmp/lbr.cert
```

- 2. Import the certificate to the Java keystore. For example, if you're working from the Essbase node, and assuming you downloaded the certificate to /tmp/lbr.cert on the Essbase server.
  - a. Log in as user **opc**. Access the Essbase node using SSH.
  - b. Run commands to add lbr.cert to the keystore. For example (your path details may differ):

```
sudo /usr/java/default/bin/keytool -import -alias mysert - file /tmp/
lbr.cert -keystore /usr/java/default/jre/lib/security/cacerts -
storepass new2mepass
```

```
Trust this certificate? [no]: yes Certificate was added to keystore
```



- 3. Restart the Java process, if the Java client is WebLogic.
- 4. Stop and restart the Essbase stack instance.
- 5. Set up WebLogic to ignore hostname verification, as described in the next section.

#### Steps for Configuring WebLogic for Use with Self-Signed Certificates

If you decide to use a self-signed certificate, you must set up the WebLogic component of the Essbase stack to ignore hostname verifications.

- 1. Access the Essbase node using SSH.
- 2. Change to oracle user.

```
sudo su - oracle
```

3. Open the setDomainEnv.sh file for editing.

```
vi /u01/config/domains/essbase domain/bin/setDomainEnv.sh
```

4. Add the following line to the JAVA OPTIONS="\${JAVA OPTIONS}" string:

```
-Dweblogic.security.SSL.ignoreHostnameVerification=true
```

When you're finished, it should look like this:

```
JAVA_OPTIONS="${JAVA_OPTIONS} -
Dweblogic.security.SSL.ignoreHostnameVerification=true"
```

- Save the file.
- 6. Stop and restart the Essbase stack instance.

### Secure Your Network

After you deploy the Essbase stack on Oracle Cloud Infrastructure, take steps to secure your network.

See Ways to Secure Your Network.

# Test Connectivity to Essbase

After you deploy Oracle Essbase on Oracle Cloud Infrastructure, test your connectivity to Essbase by logging in to the Essbase web interface, Cube Designer, the Essbase command-line tool (CLI), and MaxL.

#### Log In to the Essbase web interface

- Find the value of essbase\_url of the Essbase instance, as described in Prepare to Work with an Essbase Stack Instance.
- 2. Start the Essbase stack, as described in Start the Stack.
- 3. In a browser, enter the value of essbase\_url of the Essbase instance.
- 4. Log in as the initial Essbase service administrator, using the IDCS System Administrator User ID that you provided during stack deployment.



When you first log in, a pop-up message requests application permissions. Click **Allow**. This message should not appear again.

#### Log in to Cube Designer

- Find the value of essbase\_url of the Essbase instance, as described in Prepare to Work with an Essbase Stack Instance.
- 2. Start the Essbase stack, as described in Start the Stack.
- 3. Follow the steps to set up Cube Designer, as described in Set up Cube Designer.
- 4. Log in as the initial Essbase service administrator, using the IDCS System Administrator User ID that you provided during stack deployment.

#### Log in to Essbase CLI

- 1. Find the public IP address of the Essbase instance, as described in Prepare to Work with an Essbase Stack Instance. For example, https://l92.0.2.1/essbase.
- 2. Start the Essbase stack, as described in Start the Stack.
- Follow the steps to set up CLI, as described in Download and Use the Command-Line Interface.
- 4. If you are connecting over VPN and you're using a load balancer, set a proxy. When proxies are needed, you must set them for each shell session. Example:

```
set HTTPS_PROXY=www-proxy-example.com:80
```

**5.** Change directories to the location where you downloaded the client. Example:

```
cd ../../temp/client/cli
```

 Run the CLI shell, esscs.bat or esscs.sh, and log in as the initial Essbase service administrator, which is the IDCS System Administrator User ID that you provided during stack deployment.

In the following example, as the password is not entered, the administrator will be prompted to provide it next. The URL is the **essbase\_url** from the job outputs resulting from the stack deployment.

```
esscs login -u admin1@example.com -url https://192.0.2.1/essbase
```

Any Identity Cloud Service users provisioned to work with Essbase can log in to CLI. The actions they can perform in CLI are determined by their roles and permissions. For details, see Understand Your Access Permissions in Essbase.

#### Log in to MaxL Client

- 1. Find the public IP address of the Essbase instance, as described in Prepare to Work with an Essbase Stack Instance. For example, https://l92.0.2.1/essbase.
- 2. Start the Essbase stack, as described in Start the Stack.
- Follow the steps to set up MaxL, as described in Manage Essbase Using the MaxL Client.



4. Change directories to the location where you downloaded the client. Example:

```
cd ../../temp/client/maxl
```

- 5. Run the startMAXL batch or shell script. A command prompt opens, and the MaxL Client starts up.
- 6. Log in by providing your credentials and the Essbase URL in the MaxL **login** statement.

In the following example, the user logging in, admin1@example.com, is the initial Essbase service administrator, which is the IDCS System Administrator User ID that you provided during Essbase stack deployment. As the password is not entered in this example, the administrator will be prompted to provide it next. The URL is the essbase\_url from the job outputs resulting from the Essbase deployment.

```
login admin1@example.com on "https://192.0.2.1/essbase";
```

Any Identity Cloud Service user provisioned to work with Essbase can log in to MaxL, as long as they are provisioned as a power user or administrator.

# Complete System Hardening and Cleanup Tasks

After you deploy Oracle Essbase on Oracle Cloud Infrastructure using Marketplace, complete the following tasks for security.

#### **Change Database Administrator Password**

If you created an Autonomous Transaction Processing database or Oracle Cloud Infrastructure database during creation of the Essbase stack, Oracle highly recommends that you update the database administrator password, using the Oracle Cloud Infrastructure console. This password isn't used during the normal run time of the Essbase stack, but it may need to be provided for maintenance tasks. Afterward, delete the stack without running the destroy action.

#### **Secure Your Network**

See Ways to Secure Your Network.

#### Remove Vault

For 19c versions up to 19.3.0.4.5 only, delete any vault or key (formerly prefixed by KMS) that you created during Essbase stack provisioning, in Encrypt Values Using Vault.

#### **Delete Stack**

(Optional) Once the Essbase stack is created, you can delete the stack without running the destroy action. This helps to declutter the Oracle Resource Manager interface.

See Delete the Stack.



# **Troubleshoot Deployment Errors**

When you deploy the Oracle Essbase stack, you may encounter some of the following errors in the log displayed in the Oracle Resource Manager console.

Error Code	Message / More Information
ESSPROV-00000	Unknown error occurred.
ESSPROV-00001	Essbase System Admin username should be alphanumeric and length should be between 5 and 128 characters.
ESSPROV-00002	Essbase System Admin password should start with a letter and length should be between 8 and 30 characters, and should contain at least one number, and optionally, any number of the special characters (\$ # _). For example, Ach1z0#d.
ESSPROV-00003	Database Admin password should start with a letter and length should be between 12 and 30 characters, and should contain at least one number, and at least one of the special characters (\$ # _). For example, BEstr0ng_#12.
	See Complete System Hardening and Cleanup Tasks.
ESSPROV-00005	Missing value for IDCS Application Client ID.
	See Create a Confidential Identity Cloud Service Application.
ESSPROV-00006	Missing value for IDCS Application Client Secret.
	See Create a Confidential Identity Cloud Service Application.
ESSPROV-00007	Missing value for IDCS Instance GUID.
	See Create a Confidential Identity Cloud Service Application.
ESSPROV-10001	Permission denied accessing the target autonomous database.
	There may be a mismatch in the target dynamic policies for the compute instance. See Set Up Policies.
ESSPROV-10002	Permission denied downloading the wallet for the target autonomous database.
	There may be a mismatch in the target dynamic policies for the compute instance. See Set Up Policies.
ESSPROV-10003	Permission denied decrypting the target encrypted value with the given encryption key or secret
	There may be a mismatch in the target dynamic policies for the compute instance. See Set Up Policies.
ESSPROV-10010	Unable to validate the given IDCS Client ID and/or IDCS Client Secret.
ESSPROV-10011	Unable to connect to the IDCS endpoint. Validate that the IDCS Tenant GUID is valid.



# Manage the Oracle Essbase Stack on Oracle Cloud Infrastructure

Let's explore the processes to manage the Oracle Essbase stack in Oracle Cloud Infrastructure.

#### **Topics:**

- Start, Stop, and Delete an Essbase Stack Instance
- Use Commands to Start, Stop, and View Status of Processes
- Restart the Essbase Compute Instance
- Access Oracle Essbase Using SSH
- Resize Block Storage Volumes
- · Patch and Roll Back
- · Back Up and Restore Essbase
- Monitor and Diagnose Essbase Operations
- Access the WebLogic Console
- · Reset or Update Admin Password

# Start, Stop, and Delete an Essbase Stack Instance

You stop, start, and delete the Essbase stack instance in the Oracle Cloud Infrastructure console.

#### Topics:

- Prepare to Work with an Essbase Stack Instance
- Start the Stack
- · Stop the Stack
- Delete the Stack

# Prepare to Work with an Essbase Stack Instance

Before you can work with the Essbase stack instance, navigate to the instance and find the public IP address.

Navigate to the Essbase Instance in the Oracle Cloud Infrastructure console:

- Log in to the Oracle Cloud Infrastructure control panel on cloud.oracle.com.
- 2. Open the navigation menu . Under Core Infrastructure, click Compute and then click Instances.
- Choose your compartment.



- 4. Select the Essbase instance you want to manage.
- **5.** Navigate to the compute instance in the Oracle Cloud Infrastructure console.
- Under Primary VNIC Information, find the Public IP Address.

#### Start the Stack

The stack includes the Autonomous Transaction Processing instance, the compute instance, two block storage volumes, object storage bucket, load balancer, and additional network components.

When starting the stack, you must start the Autonomous Transaction Processing instance first, and then the compute instance. Starting the compute instance starts Essbase.

Start the Autonomous Transaction Processing instance:

- 1. Log in to the Oracle Cloud Infrastructure control panel, cloud.oracle.com.
- Open the navigation menu. Under Database, select Autonomous Transaction Processing.
- 3. Choose your compartment.
- In the list of autonomous databases, click on the display name of the database you wish to administer.
- 5. Click Start.

Start the Essbase compute:

- Confirm in the Oracle Cloud Infrastructure interface that the autonomous database is available.
- 2. Navigate to the compute instance on the Oracle Cloud Infrastructure console. See Prepare to Work with an Essbase Stack Instance.
- 3. Click **Start**. This takes a few minutes.

# Stop the Stack

The stack includes the Autonomous Transaction Processing instance, the compute instance, two block storage volumes, object storage bucket, load balancer, and additional network components.

When stopping the stack, you must stop the compute instance first, and then the Autonomous Transaction Processing instance. Stopping the compute instance stops Essbase.

Stop the compute instance:

- 1. Navigate to the compute instance in Oracle Cloud Infrastructure console. See Prepare to Work with an Essbase Stack Instance.
- 2. Click Stop.

This shuts down the instance. After the instance is stopped, the **Start** button is enabled.

When you stop your compute instance, billing stops for CPU usage.



Stop the Autonomous Transaction Processing instance:

- Confirm in the Oracle Cloud Infrastructure interface that the compute instance is stopped.
- Open the navigation menu. Under Database, select Autonomous Transaction Processing.
- Choose your compartment.
- 4. In the list of autonomous databases, click on the display name of the database you wish to administer.
- Click Stop.
   Billing for storage continues when Autonomous Transaction Processing is stopped.

Only the compute instance and the Autonomous Transaction Processing instance are stopped when you stop the stack.

### Delete the Stack

Deleting the stack terminates all resources that were created when the stack was created, including the Essbase compute node, the Autonomous Transaction Processing instance, and network components.

Before you delete the stack, ensure that the bucket is empty. For instructions, see "Using the Console" in Managing Objects. The bucket name is backup *databasename*.

When using a shared database, delete schemas created by Essbase prior to destroying the stack, or your database may contain orphaned database content. To delete the schemas, ssh to the compute instance and run /u01/vmtools/sysman/drop-schema.sh. You're prompted to enter your database administrator password.

- 1. a. If using a shared, existing database, it's recommended to delete schemas created by Essbase, prior to destroying the stack, or your database may contain orphaned database content. To do this, ssh into the compute instance and run the following: /u01/vmtools/sysman/drop-schema.sh and use the drop-schema.sh tool to drop the schemas from the existing database.
  - **b.** In the object storage bucket for the database backups, manually delete all items.
- 2. Click the navigation menu , select **Resource Manager**, and then click **Stacks**.
- 3. Click on your stack.
- 4. Click Terraform Actions, and then select Destroy.
- 5. When prompted for confirmation, click **Destroy**.
- 6. Periodically monitor the progress of the Destroy job until it is finished. If the destroy job fails, manually terminate the resources that run as part of the stack on Oracle Cloud Infrastructure before you delete the stack. For example, terminate the networking load balancer (if used), terminate the virtual cloud network (VCN), and terminate the Autonomous Transaction Processing database.
  - If an email address is associated with your user profile, you receive an email notification.
  - Proceed with deleting the stack regardless of whether the job passes or fails.
- 7. Click Delete Stack.
- 8. When prompted for confirmation, click **Delete**.



# Use Commands to Start, Stop, and View Status of Processes

You can control deployed Essbase server processes using script commands to start, stop, and view component status. Use these scripts to manage Essbase servers without stopping the compute instance.

- You use the start.sh or stop.sh scripts to start or stop the WebLogic Administration Server and the Managed Server components. Use these scripts to restart these servers if they crash.
- You use the systemctl <start|stop|status> essbase\_domain-nodemanager service
  to start or stop the Node Manager service or to check its status. It's important that
  the Node Manager service is already running when you are using the start.sh/
  stop.sh/status.sh scripts.

#### **Control the AdminServer and Managed Servers**

The opc user has rights to use sudo to execute the commands shown in these examples. The oracle user does not, but it does have access to the paths of the Essbase cube objects and artifacts. Therefore, Oracle recommends using the opc user, with sudo permissions of oracle, to run the scripts that manage the WebLogic AdminServer and any Managed Server components, as shown in these examples.

**1.** As opc user, assume the privileges of the oracle user. Example:

```
sudo su - oracle
```

2. Navigate to the location of the scripts in the domain tools directory, <Domain Root>/<Domain Name>/esstools/bin/. Example:

```
cd /u01/config/domains/essbase_domain/esstools/bin
```

3. To check status of servers, run the status script.

```
./status.sh
```

In this example, AdminServer and Essbase managed server are running:

If the status script fails with an error, for example, that the connection to Node Manager was refused, check the log to locate the cause of the error.



4. To stop AdminServer and all managed servers in the Essbase platform, run the stop script without any arguments:

```
./stop.sh
```

To stop a specific server, use the -i option and provide a server name:

```
../stop.sh -i ess server1
```

5. To start AdminServer and all managed servers in the Essbase platform, run the start script without any arguments:

```
./start.sh
```

To start a specific server, use the -i option and provide a server name:

```
./start.sh -i ess_server1
```

#### **Control the Node Manager Process**

To start or stop the Node Manager or view its status,

To start the Node Manager, as opc user:

```
sudo systemctl start essbase_domain-nodemanager
```

Note that this command will reconnect the Node Manager to existing administration processes. This command doesn't stop the AdminServer or the managed server(s).

To stop the Node Manager:

```
sudo systemctl stop essbase domain-nodemanager
```

Note that this command doesn't stop the AdminServer or the managed server(s).

To restart the Node Manager:

```
sudo systemctl restart essbase domain-nodemanager
```

Note that this command doesn't restart the AdminServer or the managed server(s).

• To show status of the Node Manager:

```
sudo systemctl status essbase domain-nodemanager
```

# Restart the Essbase Compute Instance

Restart the Essbase compute instance in the Oracle Cloud Infrastructure console.

- 1. Navigate to the compute instance in the Oracle Cloud Infrastructure console.
- 2. Click Reboot.



# Access Oracle Essbase Using SSH

Use Secure Shell (SSH) client software to connect to the Essbase instance deployed on Oracle Cloud Infrastructure to perform administrative tasks.

To access the private compute node, it depends on your environment. If your network configuration uses FastConnection or VPN with IPSec, you must provide the network setup that allows you to SSH to the private compute node.

You can log in securely to your Essbase instance from a remote host by using a secure shell (SSH) connection.

Before creating the Essbase instance on Oracle Cloud Infrastructure, generate at least one SSH key pair, and ensure that the private key is available on each host that you'll use to access Essbase instances.

To connect to Essbase using SSH, you use the SSH private key, which is part of the key pair created as a prerequisite to deploying the stack. After you've created the stack, you have a few different IP addresses, depending on the network topology.

You can use any SSL utility, for example openSSL, to generate SSH keys and to log in to your Essbase instance.

#### **Public IP**

If the Essbase compute node has an <code>essbase\_node\_public\_ip</code> address, you can access it directly with <code>ssh</code> of the Essbase public IP using the following:

```
$ ssh -i <path_to_private_key> opc@<essbase_public_ip>
```

#### **Private IP**

If you have deployed stack in a private subnet, the Essbase compute node will have a private IP only. In this case, you must create Bastion provided under Identity & Security. Use this Bastion to access the Essbase compute node. SSH commands can be copied from the Bastion session when you create the session.

(For 19c through 19.3.0.4.5) If you deployed a private subnet, the Essbase compute node has a private IP, but no public IP. In this case, you must use the public IP of the bastion host to access the Essbase compute node by proxy. Examples are provided in this topic. SSH syntax may vary, depending on your SSH client and operating system.

#### **Create a Bastion**

- In OCI Console > Identity & Security > Bastion.
- On Bastion page, click Create Bastion.
- 3. Enter Bastion name.
- **4.** In **Configure Networking**, select the **Target Virtual Cloud Network** (VCN), the VCN in the relevant compartment the VCN on which Essbase node is deployed.
- Select the Target Subnet on which Essbase node is deployed.
- Enter CIDR IP or range of IPs from which access will be allowed.
- Click Create Bastion. Bastion is created.

#### **Create a Session**



- 1. In OCI Console > Identity & Security > Bastion > click relevant created Bastion name link.
- Click Create Session.
- 3. On the Create Session page, select session type value.
- 4. Enter username.
- 5. Enter computer instance to which you want to connect.
- Specify the SSH Key.
- 7. In **Advanced options**, you can specify the IP to which you want to connect.
- 8. Click Create Session. The session is created.
- 9. Open the compute instance and enable Bastion Service toggle for the instance. Note by default, a Bastion session is available for three hours, and can be defined otherwise.

#### **SSH Command to Access Essbase Node**

- 1. Go to OCI Console > Identity & Security > Bastion.
- 2. Select the relevant Bastion.
- 3. Under Resources, click **Sessions** to view the Sessions page.
- 4. On Sessions page, open drop-down menu to the right of the relevant created session, and click Copy SSH Command. You will use this command to ssh and log into the Essbase node.
- 5. Replace ssh\_key paths in the copied command. Edit the SSH command as follows: Replace <pri>privateKey> with the path to the private key (from the SSH key pair used to create the session).

#### Bastion Host SSH Tips (For 19c through 19.3.0.4.5)

Essbase doesn't have a public IP address when you deploy a private subnet using a bastion host. Use these guidelines to help you configure your system for SSH access to the Essbase compute node on Oracle Cloud Infrastructure. These examples utilize a bash shell. Bash commands you enter are in bold.

1. Change to the hidden directory, .ssh, usually located in your user directory (check the documentation for your specific SSH client).

```
cd ~/.ssh
```

2. Modify (or create) the config file in the .ssh directory. Though the following example invokes the UNIX vi editor, you can use any text editor.

```
vi config
```

3. In the config file, enter HostName and IdentityFile details for the bastion host. For HostName, provide the IP address of the bastion host, and for IdentityFile, provide the location of the private key that matches the public key you provided to Resource Manager during the Essbase deployment. Format:

```
Host bastion
   HostName <bastion_host_public_ip>
   IdentityFile <path to private key>
```



4. Add to the config file an additional host entry for the private Essbase subnet. For HostName, provide the essbase\_node\_private\_ip, and for IdentityFile, provide the location of the private key that matches the public key you provided to Resource Manager during the Essbase deployment. For ProxyCommand, set up SSH access for the opc user to access the bastion host by proxy. Example:

```
Host essbase
  HostName <essbase_node_private_ip>
  IdentityFile <path_to_private_key>
  ProxyCommand ssh opc@bastion -W %h:%p
```

Here is an example of a completed config file.

```
Host bastion
   HostName 192.0.2.111
   IdentityFile C:/temp/ids/my_key

Host essbase
   HostName 10.0.1.2
   IdentityFile C:/temp/ids/my_key
   ProxyCommand ssh opc@bastion -W %h:%p
```

- 5. Save the config file and exit the editor.
- **6.** In your command window, log in over SSH to Essbase, by proxy of the bastion host, as the **opc** user.

```
ssh opc@essbase
```

7. Switch to user **oracle** to explore the Essbase compute and complete any administrative tasks.

```
sudo su oracle
```

8. Change to the home directory of the Essbase compute node on Oracle Cloud Infrastructure.

cd /

9. View the directories.

```
{\tt ls} bin boot dev etc home lib lib64 media mnt opt proc root run sbin srv tmp u01 usr var
```

10. Explore the Essbase directories. The applications are in the app directory, and the file catalog is in the catalog directory.

```
cd /u01/data/essbase
ls
app catalog hybrid jagentId.id
```

#### **Related topics**



#### Connecting to an Instance

Simplify Secure Access with OCI Bastion Service

# Resize Block Storage Volumes

Resize block volumes in the compute instance, in the Oracle Cloud Infrastructure console.

- 1. Navigate to the block volume in the Oracle Cloud Infrastructure console. See Prepare to Work with an Essbase Stack Instance.
- 2. Click Edit Block Size Or Performance, Wait until the available state.
- 3. Enter the new block size in GB and click **Save Changes**.
- 4. Wait until it's in an available state.
- 5. SSH to the compute instance. See Access Oracle Essbase Using SSH.
- 6. Run sudo lsblk as user opc to determine the volume files system that should be increased.

```
$sudo lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

sdd 8:48 0 64G 0 disk /u01/tmp

sdb 8:16 0 64G 0 disk /u01/config

sdc 8:32 0 256G 0 disk /u01/data

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
```

- 7. Run sudo dd iflag=direct if=/dev/sdc of=/dev/null count=1
- 8. Run echo "1" | sudo tee /sys/class/block/sdc/device/rescan
- 9. After running, check again using sudo lsblk that the partition size is updated.

```
$ sudo lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

sdd 8:48 0 64G 0 disk /u01/tmp

sdb 8:16 0 64G 0 disk /u01/config

sdc 8:32 0 378G 0 disk /u01/data

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
```

**10.** If the target volume is a partition, run sudo growpart, defining the physical disk name and partition number.

```
sudo growpart /dev/sdc 1
```

- 11. After running, check again using lsblk that the partition size is updated.
- **12.** Run sudo xfs growfs to update the size of the target file system for the volume.



```
blks
                               sectsz=4096 attr=2,
projid32bit=1
                              crc=0
                                           finobt=0 spinodes=0
rmapbt=0
                              reflink=0
data
                              bsize=4096 blocks=67108864,
imaxpct=25
                              sunit=0
                                         swidth=0 blks
                              bsize=4096 ascii-ci=0 ftype=1
naming =version 2
log
        =internal
                              bsize=4096 blocks=32768, version=2
                              sectsz=4096 sunit=1 blks, lazy-
count=1
realtime =none
                              extsz=4096 blocks=0, rtextents=0
data blocks changed from 67108864 to 99090432
```

13. Run df -h to validate that the target file system has been resized.

```
$ df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sdc 378G 33M 378G 1% /u01/data
```

#### **Related Topics**

- Oracle blog: Size Up Your Volumes While They're Online!
- Editing Volume Size and Performance, in Oracle Cloud Infrastructure Documentation

## Patch and Roll Back

You can patch and roll back an Oracle Essbase cloud instance on Oracle Cloud Infrastructure.

- Patch and Roll Back For Version 19.3.0.3.4 and Later
- Patch and Roll Back for Version 19.3.0.2.3 and Earlier

## Patch and Roll Back - for Version 19.3.0.2.3 and Earlier

You can patch and roll back an Oracle Essbase cloud instance on Oracle Cloud Infrastructure.

For version 19.3.0.2.3 and earlier.

- Access the Essbase node using SSH, as described Access Oracle Essbase Using SSH.
- 2. Stop Essbase instance, as user opc user.

```
sudo systemctl stop essbase.service
```

3. Set environment variables, as user **oracle**.

```
export ORACLE_HOME=/u01/oracle
export OPATCH NO FUSER=true
```



- 4. Download the OPATCH patch file from My Oracle Support.
- 5. Copy the downloaded patch zip file to cloud instance, using a tool such as WinSCP, SCP, Filezilla, and others.

If you use WinSCP, for example, on the WinSCP Login page, under Session, enter the Host name, Port number, User name opc, Private key file, and click Login. The file is copied.

- 6. Ensure that user **oracle** has **read** access to the OPATCH file.
- 7. Apply OPATCH, as user **oracle**. You can apply OPATCH directly while in zip format. Provide the absolute path of the zip file.

```
$ORACLE HOME/OPatch/opatch apply /tmp/Marketplace/<OPATCH_ID>.zip
```

Note: If you want, you can optionally unzip the zip file and apply the patch. Provide the absolute path of the unzipped folder.

```
unzip ./<OPATCH_ID>.zip
$ORACLE HOME/OPatch/opatch apply /tmp/Marketplace/<OPATCH ID>/
```

8. Start the Essbase cloud instance, as user **opc**.

```
sudo systemctl start essbase.service
```

Verify the log to see that all binaries patched successfully. Log in to Essbase instance and verify the build number in **admin** → **About**.

- 9. [OPTIONAL] If you need to roll back the applied patch, do the following steps.
  - **a.** If the Essbase instance is running, first stop Essbase system service, as user **opc**, as described above.
  - b. Run the rollback command:

```
$ORACLE_HOME/OPatch/opatch rollback -id <OPATCH_ID>
```

For example: \$ORACLE\_HOME/OPatch/opatch rollback -id 30081463

c. Restart the Essbase instance after rollback. If you're unable to login in, clear the Essbase related browser cache or open a new browser.

### Patch and Roll Back - For Version 19.3.0.3.4 and Later

You can patch and roll back an Oracle Essbase cloud instance on Oracle Cloud Infrastructure.



If you are patching from 19.3.0.4.5 to 19.3.0.5.6, continue to follow the backup and restore steps in Back Up and Restore: Version 19.3.0.4.5 and Prior.

Access the Essbase node using SSH. See Access Oracle Essbase Using SSH.



2. As opc user, sudo to oracle user:

```
sudo su oracle
```

3. Then, as **oracle** user, access scripts in esstools, with this path:

```
cd /u01/config/domains/essbase domain/esstools/bin
```

4. Stop Essbase instance, as **oracle** user, using the script:

```
./stop.sh
```

5. Set environment variables, as **oracle** user.

```
export ORACLE_HOME=/u01/oracle
export OPATCH NO FUSER=true
```

- Download the OPATCH patch file from My Oracle Support. If you do not have access to patches in My Oracle Support, please open a service request with Oracle Support.
- 7. Copy the downloaded patch zip file to cloud instance, using a tool such as WinSCP, SCP, Filezilla, or others.

If you use WinSCP, for example, on the WinSCP Login page, under Session, enter the Host name, Port number, User name opc, Private key file, and click Login. The file is copied.

- 8. Ensure that **oracle** user has **read** access to the OPATCH file.
- **9.** Apply OPATCH, as **oracle** user. You can apply OPATCH directly, while in zip format. Provide the absolute path of the zip file.

```
$ORACLE HOME/OPatch/opatch apply /tmp/Marketplace/<OPATCH ID>.zip
```

Note: You can optionally unzip the zip file and apply the patch. Provide the absolute path of the unzipped folder.

```
unzip ./<OPATCH_ID>.zip
$ORACLE_HOME/OPatch/opatch apply /tmp/Marketplace/<OPATCH_ID>/
```

10. Start the Essbase cloud instance, as oracle user, and access scripts in esstools with this path: cd /u01/config/domains/essbase\_domain/esstools/bin. Start Essbase instance, as oracle user. using the script: ./start.sh.

```
cd /u01/config/domains/essbase_domain/esstools/bin
    ./start.sh
```

Verify the log to see that all binaries patched successfully. Log in to Essbase instance and verify the build number in **admin** → **About**.

- 11. [OPTIONAL] If you need to roll back the applied patch, do the following steps.
  - **a.** If the Essbase instance is running, first stop Essbase system service, as **oracle** user, as described above.



#### b. Run the rollback command:

\$ORACLE HOME/OPatch/opatch rollback -id <OPATCH ID>

For example: \$ORACLE HOME/OPatch/opatch rollback -id 30081463

c. Start the Essbase instance after rollback. If you're unable to login in, clear the Essbase related browser cache or open a new browser.

## Back Up and Restore Essbase

Periodically, you may need to restore your instance from a backup. The topics below explain how to do this for stack deployments on Oracle Cloud Infrastructure where you have a single Essbase instance and have deployed using Autonomous Transaction Processing Database.

## About Backup and Restore

Essbase backup and restore planning is required at both the application and instance level to have full flexibility to manage the life cycle of your Essbase instances, and also to provide disaster recovery.

Essbase allows both Recovery Point Objective (RPO) and Recovery Time Objective (RTO) to be customized for your user base, supporting both instance-level and application-level backups. For example, if your RPO for most applications is twenty-four hours, then you must do an instance-level backup once a day. What if you have one application that needs to be backed up more often, say once every four hours? You can back up that application every four hours using LCM export. If you need to restore the instance, you can restore from the instance-level backup and then update that one application using the latest LCM export.

Backups of individual applications protect you from application failures or application artifact corruption, and can easily be migrated between servers. When you restore a single application, there is no disruption to user activity with other applications in your instance. Essbase application backups are taken using L:CM export and import commands.

Essbase instance backups protect you against unplanned hardware or Essbase agent failures, which affect all applications on the instance. All user activity on the instance is affected for the duration of the recovery process, and all applications are restored to the point in time of the instance backup. Instance backups are also helpful when you are retiring older hardware.

## Back Up and Restore Individual Applications

Routine backup of individual Essbase applications ensures that you can recover from any sort of disruption to a single application without affecting the other applications running on the same Essbase instance.

LCM export and import operations allow you to move applications on and off your Essbase instance or between Essbase instances. Export and import operations can be run sequentially in instance migration use cases, exporting from the source instance and importing into the target instance. However, to protect from unexpected application-level failures a routine backup cadence is required. The frequency of your application backups should correspond to the acceptable loss (recovery-point objective or RPO) specified by the application's users. At the time of any application-specific failure, the latest LCM export file can be used to recover all or part of the application.



By default, LCM export exports your entire application and its cubes without an itemized inventory of artifacts. When executing an LCM export, consider generating an artifact list. Only if you have included this artifact list on export will you have the option to selectively import specific components of your application and its cubes. To ensure backup consistency, make sure that the application is stopped before taking an LCM export.

#### Note:

For example, if using the CLI LcmExport command, you can use the optional -generateartifactlist parameter.

To back up applications, do an LCM export operation. To restore them, do an LCM import operation. You have various options, detailed below:

- Back up your 19c applications using one of the methods listed below to initiate the LCM export operation. To do an LCM export, you need at least user role with Application Manager permission, or, you must be the power user who created the application.
  - Command-line interface (CLI): LcmExport command
  - Essbase web interface Export LCM job
  - REST API: Execute Job operation (using jobType Icmexport)
- If you will use the CLI to perform the LCM export, first download it to your compute, from the Console in the Essbase web interface, and set it up. See: Access Tools and Tasks from the Console and Download and Use the Command-Line Interface.
- 3. Restore applications from backup using one of the methods listed below to initiate the LCM Import operation.
  - Command-line interface (CLI): LcmImport
  - Essbase web interface Import LCM job
  - REST API: Execute Job operation (using jobType lcmimport)

#### Note:

If your application is on 11g, you must first migrate to the current release using the <code>EssbaseLCMUtility.zip</code> before you can back up using LCM export and LCM import commands. Download the utility from the Console in the Essbase web interface, and see the enclosed README for usage instructions.

## Back Up and Restore an Essbase Instance

Use Essbase instance backups to restore all applications on your instance to a common point in time. Instance backups are primarily for disaster recovery, but are also appropriate when you want to migrate or restore all applications at once.

Backup and restore should be performed on the same version of Essbase.



Certain components from every Essbase stack you create contain information that makes your Essbase deployment unique. You will need to back up these unique stack components at appropriate intervals to meet your recovery objectives.

In the event of an Oracle Cloud Infrastructure compute or availability domain failure, you can recover your Essbase instance by building a new stack and restoring into it your Essbase backup. The newly deployed stack should be the same version of Essbase as the instance that failed. You may need to use GitHub to restore to the same version if you have not migrated versions in a timely manner.

For Essbase on Oracle Cloud Infrastructure, restoring from disaster necessitates deploying a new Essbase stack and attaching to it the appropriate block volume and relational database schema backups. Think of the pre-restore Essbase stack as a source (of block volumes, block volume backups, relational database schemas, relational database 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.

Backups of Essbase on Oracle Cloud Infrastructure depend on some details of your Essbase stack. A complete backup must protect all information that makes your Essbase deployment unique. Items you may be instructed to back up include:

- Relational database schemas for every Essbase stack, which store some application, user and configuration information.
  - A single database schema for Essbase, called <instance prefix> Essbase.
  - 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/confg. (Essbase is a managed service within a WebLogic domain.)

Make sure that your backup strategy captures information at appropriate intervals to align with your RTO/RPO use case. As Essbase is not an autonomous database, Essbase services are stopped during backup.

Back Up and Restore: Version 19.3.0.4.5 and Prior

If you are using version 19.3.0.4.5 or an earlier release, follow these steps to back up and restore an Essbase instance.

#### Create a Manual Backup

- SSH to the Essbase host. See Access Oracle Essbase Using SSH.
- 2. Become user oracle.

```
sudo su - oracle
```

3. Run this command:

```
crontab -e
```

4. Add cron entry. For example, to run every day at 00:01, as the host timezone of the computer instance, enter:

```
1 0 * * * /u01/vmtools/backup.sh > /dev/null
```

Note that the backup script should be run in background mode or cronjob. In interactive mode, the script may time out and the instance may not start up.

5. Save the file.



#### **Restore from Backup**

Restore Essbase from backup in the Oracle Cloud Infrastructure console.

- 1. Choose the date and time you want to restore.
- 2. Check that the block storage backup and Autonomous Transaction Processing backup exist with the same creation time stamp.
- 3. Navigate to the block storage backups.
  - a. Open the navigation menu.
  - Under Core Infrastructure, go to Block Storage and click on Block Volume Backups.
- Create two new block storage volumes (Data and Config) from the selected backups.
- **5.** Add newly created block volumes to the instance's **Volume Group**.
- 6. Stop the service.
- 7. Restore Autonomous Transaction Processing.
- 8. Navigate to the compute on the Oracle Cloud Infrastructure console.
- 9. On the Resources tab, click Attached Block Volumes.
- 10. For each disk:
  - a. Click on the menu \*\*\* next to the disk.
  - b. Click Detach.
  - c. Click Continue Detachment and Okav.
  - **d.** Wait for the old disk to be detached.
- **11.** After both disks are detached, attach the newly created volumes to the instance, keeping the same device path.
- 12. Wait for the disks to be attached.
- 13. Wait for Autonomous Transaction Processing to be restored.
- **14.** After Autonomous Transaction Processing is restored, and the status is AVAILABLE NEED ATTENTION, stop and then restart it.
- 15. Restart the instance.

#### What's Not Recovered in Disaster Recovery or Restore to New Instance

An Essbase connection to an Essbase instance created in a source stack doesn't work after disaster recovery or backup/restore is made to a new Essbase Marketplace instance.

For example, let's say an Essbase connection is created to an Essbase instance (Stack A) and a block backup is run. Later, a disaster occurs and Stack A no longer works. We create another Essbase instance (Stack B) and restore to it the blocks backed up from Stack A. After the restore/recovery to the new instance, the connection to stack A no longer works.

#### Set Up a Bucket to Store Manual Backups

This is a one-time action.



You are not required to set up a bucket to store manual backups if your instance is configured using Autonomous Transaction Processing - Dedicated (ATP-D).

- If you have not already done so, generate an auth token for the Oracle Cloud Infrastructure Object Storage user to access the bucket you create in the next step. See create an auth token to learn how to do this. You will need this auth token for the database credential you create in step 5.
- SSH to the Essbase host. See Access Oracle Essbase Using SSH.
- 3. Become user oracle.

```
sudo su - oracle
```

4. Change directory to /u01/vmtools.

```
cd /u01/vmtools
```

- 5. Run script ./configure-backup-storage.sh and input the required fields:
  - Database Admin Password
  - Oracle Cloud Infrastructure Admin User Name
  - Oracle Cloud Infrastructure Admin Token this is the token you created in step 1.

#### Back Up and Restore: Version 19.3.0.5.6

Beginning with version 19.3.0.5.6, Oracle-scripted backups are for a single Essbase instance. This means that you can use Oracle scripts to perform backups even if you deployed multiple Essbase instances using a single Autonomous Database.

The Oracle-provided scripts support only Autonomous Database deployments (Autonomous Transaction Processing or Autonomous Data Warehouse).

We've made some assumptions to limit the size and scope of this chapter. All examples assume that:

- Autonomous Transaction Processing 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; it is the only non-IDCS user in the system) is the same between the source and the target Essbase stacks.
- At the time of restore, the source instance backup has at least one valid Oracle Identity Cloud Service user with an Essbase system administrator role.

#### Oracle-Scripted Backup of an Essbase Instance

The Oracle-provided scripts perform the following tasks:

- Configure your database to work with your object storage.
- Stop the Essbase services.
- Back up your database.
- Back up your Essbase data block volume.
- Start the Essbase services.

Before beginning a backup, you may want to gracefully bring users off the system because the script stops your Essbase services. See Alter Application (especially enable/disable) and Alter System (especially logoff/ kill). If you use disable commands, you must reverse them with enable commands after completing your backup.



To initiate a backup of an Essbase instance deployed using an Autonomous Database, schedule the backup for a convenient time when users are not in the system and follow these steps:



If using a private IP, provision an Oracle Cloud Infrastructure Bastion Service instance and use it as a proxy.

- Ensure that the required policies to manage backups are in place. See Set Up Policies.
- 2. ssh to your Essbase compute (as opc user).
- 3. sudo su oracle
- 4. cd /u01/vmtools/backup/
- **5.** Run the script to configure your database to work with object storage. This is a one-time action per Essbase instance.

```
./configure-backup-storage.sh
```

After you launch the script, you are prompted for three inputs:

a. Enter database admin password:

Type the clear text password. Because the password is protected information, you will not see the text as you type at the command prompt.

b. Enter OCI Username:

To find the Oracle Cloud Infrastructure username:

- i. In Oracle Cloud Infrastructure, go to the upper right hand corner and click the profile menu icon  $\Omega$ .
- ii. Click on the user.
- iii. Copy the profile name at the top of the page.

  Copy the full username including oracleidentitycloudservice/ (not just the e-mail address).
- iv. Go back to the command prompt, paste and press Enter.
- c. Enter OCI Token:

Enter the auth token. Because the auth token is protected information, you will not see the text you type or paste at the command prompt.

If you have not already done so, generate an auth token for the Oracle Cloud Infrastructure Object Storage user.

6. Run the backup script.

You are prompted for the database admin password. Type the clear text password. Because the password is protected information, you will not see the text as you type at the command prompt.

This backs up the data block volume.

./run-backup.sh



#### Note:

The Essbase services are stopped by the backup.sh script and restarted after it finishes.

7. Capture object storage native URI.

At the end of the backup script, there will be two "moving export file" entries. You can copy and save the first one, which contains your object storage native URI. The object storage native URI is required when you issue the Data Pump Import command during restore. (You must modify the moving export file entry to replace the .dmp file name with essbase%u.dmp.)

#### Restore an Essbase Instance

If you haven't done so already, install and configure Oracle Instant Client and tools. You will need to run Data Pump and SQL\*Plus.

When recovering from disaster, you'll need to deploy a new target Essbase instance before you can restore a non-Oracle-scripted backup. The new target instance should be the same version that was deployed on your failed compute. After your new target instance is deployed, you can recover from backup using the target.

If you have not experienced a disaster, but want to migrate or roll back your source instance (same host recovery), use the restore steps below with these exceptions:

- Skip steps 1 through 6.
- Do not include the REMAP\_SCHEMA=<sourceEBprefix>\_ESSBASE:<targetEBprefix>\_ESSBASE parameter in step 8b.
- Because you are recovering on the same host from which you took your backup, the target instance referenced in the steps below will be your source instance.

The following steps will allow you to restore a single Essbase instance without impacting other Essbase instances that may be deployed in the same Autonomous Database.



If using a private IP, provision an Oracle Cloud Infrastructure Bastion Service instance and use it as a proxy.

- 1. Deploy a target Essbase stack using Oracle Marketplace.
  - Use the source Oracle Identity Cloud Service confidential application.
  - Use the source Autonomous Database and password. Optionally, create a new Autonomous Database. This is necessary if you are restoring in a new region.
  - Use the source virtual cloud network and application subnet. Optionally, use a new network. This is necessary if you are restoring in a new region.
  - 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. Optionally, create a new load balancer. This is necessary if you are restoring in a new region.



- If your source stack has a bastion, deploy a bastion with the target stack (the source bastion 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 Oracle Identity Cloud Service 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 Oracle Identity Cloud Service user with the Essbase system administrator role. After you restore, you must login to the target instance as a valid Oracle Identity Cloud Service user who had Essbase system administrator role on the source instance.
- 2. Manage the target instance login URL:
  - a. If you have a source load balancer, manage its 'essbase' backend set for use with the target compute. After allowing the load balancer time to refresh its connection to the target compute, login to the Essbase web interface to make sure your target instance is deployed correctly before proceeding to restore from backup.
    - i. Remove the source compute backend.
    - ii. Add the target compute backend. Use port 443.

There is no need to update the Oracle Identity Cloud Service confidential application URLs, as the same load balancer IP is now routing to the target Essbase instance.

- b. If your source stack did not have a load balancer, update your Oracle Identity Cloud Service Confidential Application Redirect and Post Logout Redirect URLs with the target IP address.
- 3. ssh to your target Essbase compute (as opc user).
- 4. sudo su oracle
- 5. cd /u01/vmtools/backup/
- 6. Run the script to configure your database to work with object storage. This is a one-time action per Essbase instance. Execute the configure-backup script for the target.

```
./configure-backup-storage.sh
```

After you launch the script, you are prompted for three inputs:

- a. Enter database admin password:
  - Type the clear text password. Because the password is protected information, you will not see the text as you type at the command prompt.
- b. Enter OCI Username:

To find the Oracle Cloud Infrastructure username:

- i. In Oracle Cloud Infrastructure, go to the upper right hand corner and click the profile menu icon ①.
- ii. Click on the user.
- iii. Copy the profile name at the top of the page. Copy the full username including oracleidentitycloudservice/ (not just the e-mail address).



- iv. Go back to the command prompt, paste and press Enter.
- c. Enter OCI Token:
  - Enter the auth token. Because the auth token is protected information, you will not see the text you type or paste at the command prompt.
- **d.** Record the directory name and the credential name, which will be printed on the screen after the script has finished executing.
- 7. Stop the target Essbase services (as oracle user). Do not stop the node manager or the Essbase compute in Oracle Cloud Infrastructure
- 8. Restore the target database schema from source schema backup that was taken by the backup.sh script.
  - When restoring the Autonomous Database for your target stack, you will import your source schema backup into the target Essbase schema using the REMAP\_SCHEMA option.
  - **a.** Make sure your Instant Client is configured to point to the Autonomous Database containing your target Essbase schemas.
  - b. Using the Oracle Instant Client, issue the following Data Pump import command

```
impdp admin@<database name>_high directory=<directory name>
credential=<credential name> dumpfile=<object storage native URI>
REMAP_SCHEMA=<source essbase prefix>_ESSBASE:<target essbase
prefix> ESSBASE table exists action=replace
```

#### Notes:

- The <database name> is the name of the database to which you connect to do the import.
- The <directory name> is the directory name you recorded after running the configure-backup-storage.sh Script.
- The <credential name> is the credential name you recorded after running the configure-backup-storage.sh script.
- The <object storage native uri> is the same one you captured during backup:
  - The namespace (/n/) is your tenancy's namespace.
  - The bucket (/b/) is the source instance object storage bucket.
  - The object (/o/) is the <backup folder name>/<dump file name> of the source instance object storage. The.dmp file name used in the impdp statement should be essbase%u.dmp. The %u is a wildcard that will pick up multiple.dmp files, in case more than one is created.
- 9. 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.
- **10.** Temporarily disable the /etc/fstab data block volume entry.
  - a. ssh to target compute (as opc user).
  - b. sudo vi /etc/fstab
  - c. Insert a # in front of the /u01/data entry.



- d. Save the file.
- 11. Detach data block volume from the target Essbase compute. Note the iSCSI caution and be sure to unmount and disconnect the volume before detaching using the Oracle Cloud Infrastructure console.
  - a. To unmount:
    - i. ssh to target compute (as opc user).
    - ii. lsblk
    - iii. sudo umount /u01/data
  - b. To disconnect iSCSI:
    - i. In the Oracle Cloud Infrastructure console, select the target compute.
    - ii. Select resources > attached block volumes.
    - iii. From the Actions menu i for the data block volume select iSCSI Commands & Information.
  - Copy iSCSI commands for disconnect.
  - d. ssh to the target compute (as opc user).
  - e. Paste the disconnect command you copied and press enter.
  - f. To detach:
    - i. In the Oracle Cloud Infrastructure console, select the target compute.
    - ii. Select resources > attached block volumes.
    - iii. From the Actions menu for the data block volume select **Detach**.
- **12.** Restore data block volume from source data block volume backup. Be sure to select the same Availability Domain as your target compute instance.
- **13.** Attach the data block volume created in the previous step to the target compute. The volume Attachment Type should be iSCSI. You don't need to select the device path.
- **14.** Use the iSCSI commands listed in the OCI console to connect your newly attached target block volumes.
- 15. Update /etc/fstab /u01/data entry and mount the new data block volume.
  - a. ssh to the target compute (as opc user).
  - b. lsblk to identify the name of the newly attached config and data volumes.
  - c. sudo blkid and record the UUID of the newly attached data volume.
  - d. sudo vi /etc/fstab
  - e. Uncomment the data block volume entries.
  - f. Replace the UUID in the existing data volume entry with the UUID of the newly attached data volume. Be sure not to change the mount point /u01/data. See Traditional fstab Options.
  - g. After saving the /etc/fstab file, issue the following commands:
    - i. sudo systemctl daemon-reload
    - ii. sudo mount -a



- h. lsblk to verify the mount points.
- 16. Start the Essbase services (as oracle user).

#### Note:

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

## Install and Configure Oracle Instant Client and Tools

To back up the schema and block volumes you will need to use Oracle Database Instant Client and Oracle Data Pump to export the schemas related to a specific Essbase instance.



You cannot connect to the server hosting your database in the cloud. Oracle Instant Client is a convenient tool for establishing local conections to cloud database instances.

- Download an Oracle Instant Client version compatible with your Autonomous Transaction Processing database version. See Oracle Instant Client Downloads.
  - Choose a basic package.
  - Be sure to download and install the corresponding Visual Studio redistributable.
  - Choose 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.
- The installation instructions for Oracle Instant Client are at the bottom of the platform installation download page.
- 3. Connect your Instant Client to your Autonomous Transaction Processing database.

# Monitor and Diagnose Essbase Operations

You can monitor, set notifications, and collect diagnostic information on Essbase operations.

#### Topics:

- Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service
- Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service
- Collect Diagnostic Information on the Essbase Node



## Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service

Oracle Cloud Infrastructure Monitoring service enables you to actively and passively monitor your cloud resources using the Metrics and Alarms features.

You use Monitoring service, which collects metrics for Essbase processes and volumes deployed in the stack, to create alarms and triggers related to your CPU, memory, and storage utilization. See Monitoring Overview in Oracle Cloud Infrastructure documentation. You must set the relevant additional policies.

When Monitoring is enabled, the compute instance starts a background process to collect metrics and publish them to the Oracle Cloud Infrastructure Monitoring service using the "oracle\_essbase" namespace.

Monitoring queries can be customized and run in the Metrics Explorer, which can be accessed from the Oracle Cloud Infrastructure console, in Monitoring > Metrics Explorer. For information on viewing default metrics and building queries, see Viewing Default Metric Charts and Building Metric Queries in the Oracle Cloud Infrastructure documentation.

For metrics dimensions, various dimensions are provided that can be used to filter the metrics that will be displayed. There are some noteworthy dimensions, including:

 stackDisplayName – corresponds to the name of the stack that was provided on the stack definition page.

The following metrics are also provided.

- Volume collected for each block volume attached to the compute instance
  - VolumeTotalSize total size in bytes for a given volume
  - VolumeUsedSize number of bytes used for a given volume
  - VolumeFreeSize number of bytes free for a given volume
  - VolumeUsedPercent percentage of the volume used
- Process collected for all Essbase processes running on the compute instance
  - CpuUtilization CPU utilization for a process
  - MemoryUtilization memory utilization for a process

# Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service

Use notifications to get notified when event rules are triggered, or alarms are breached, or to directly publish a message. This feature is optional.

The use of Oracle Cloud Infrastructure Notifications service is optional. This service can be used by subscribers to be notified of life cycle events. Notifications service broadcasts messages to distributed components through a publish-subscribe pattern, delivering secure, highly reliable, low latency, and durable messages, for applications hosted on Oracle Cloud Infrastructure and externally. See Notifications Overview in Oracle Cloud Infrastructure documentation. For information on managing and creating alarms, see Building Metrics and Managing Alarms.





If you don't use Monitoring service, the alternative method to monitor is to ssh into the machine and monitor the essbase-init-log file.

If notifications are enabled, messages are published on the given topic for the following events:

- Compute instance configuration started
- Compute instance configuration completed or failed
- Backup started
- · Backup completed or failed

## Collect Diagnostic Information on the Essbase Node

You can SSH to the Essbase node on Oracle Cloud Infrastructure to collect diagnostic information for troubleshooting purposes.

To get diagnostics,

1. Connect to the Essbase node using SSH.

See Access Oracle Essbase Using SSH.

2. Change to user oracle.

```
sudo su - oracle
```

3. Change directory to /u01/vmtools/diagnostics

```
cd /u01/vmtools/diagnostics
```

4. Run the diagnostics collection script, providing as the argument a path and a filename without any extension. The script can be run without a password.

```
./collect-diagnostics.sh /tmp/diagnostics
```

The diagnostics are collected in a compressed file (for example, diagnostics.zip).

# Access the WebLogic Console

As an Essbase administrator, when you are managing your Oracle Essbase stack on Oracle Cloud Infrastructure, you may need access to the WebLogic console to perform some administrative tasks.

The Essbase stack on Oracle Cloud Infrastructure runs from a managed WebLogic server. When you start or stop the Essbase stack, it starts and stops the WebLogic server as well as the Essbase applications.

#### Caution:

Essbase instances are configured by default with no access to the administrative T3
 WebLogic port. Oracle highly recommends that all access to the T3 port remain disabled,



and that you secure it immediately. If necessary for business purposes, access to the T3 port should only be allowed from a certain fixed set of IPs, using SecIPList or a restricted classless inter-domain routing block (CIDR); for example, xx.xx.0.0/16.

- ALL ports should be closed to the public internet. There should be only two network options:
  - 1. VPN You open ports in your private network.
  - 2. Public access ssh (port 22) through bastion (service instance), 443 through Load Balancer, host ports should not be open to the public internet.

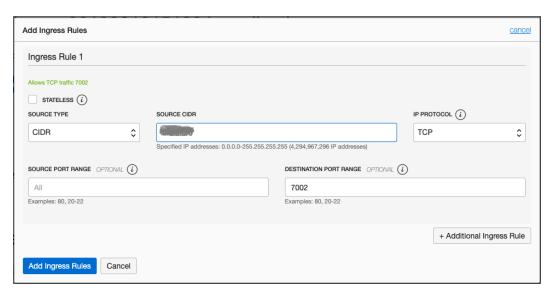
The WebLogic AdminServer runs on port 7002. To access it,

**1.** Expose the port on the target compute node. To do this, SSH into the target machine as the opc user and run the following commands:

```
[opc@essbase-1 ~]$ sudo firewall-cmd --add-port=7002/tcp --
zone=public
# To make this survive restarts of the firewall service
[opc@essbase-1 ~]$ sudo firewall-cmd --add-port=7002/tcp --
zone=public --permanent
[opc@essbase-1 ~]$ sudo systemctl restart firewalld
[opc@essbase-1 ~]$ sudo firewall-cmd --list-all
public
  target: default
  icmp-block-inversion: no
  interfaces:
  sources:
  services: sshdhcpv6-client http https
 ports: 7002/tcp
 protocols:
 masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

2. Enable the security list for the subnet to allow limited access to the port from a source network. The quick start creates a virtual cloud network (VCN) named cprefix>-vcn, and a security list named <prefix>-app-security-list. Add an ingress rule as shown below. Enter an IP address in the Source CIDR field that only your admin users are allowed to access from their client laptop or desktop.





3. Log in to the target machine using the Console URL https://<Essbase IP address>:7002/console, with the same Essbase administrator login that you used during configuration of the stack.

# Reset or Update Admin Password

You can use these steps to reset or change your Admin password.

#### Steps to Update the Admin Password on OCI node

 Follow the steps to access the WebLogic console on OCI - see Access WebLogic Console on OCI.



#### Note:

If the Admin user doesn't want to enable permanent access to WebLogic port, they can create a temporary tunnel as mentioned below, to access the WebLogic console:

a. • For a **stack with public IP**, open command shell and run the following command to open tunnel:

```
ssh -L 7002:private IP:7002 target
```

where target is alias for ssh opc@target.

#### For example:

```
ssh -L 7002:10.xx.xx.xx:7002 target
```

#### The .ssh/config file looks like this:

```
Host target

Hostname private_ip

User opc

StrictHostKeyChecking=no
IdentityFile private_key_for_stack

ProxyCommand ssh targetbastion -W %h:%p

Host targetbastion
HostName public_ip
User opc
StrictHostKeyChecking=no
IdentityFile private_key_for_stack
```

#### For example:

```
Host target
    Hostname 10.xx.xx.xx
User opc
    StrictHostKeyChecking=no
    IdentityFile C:\xxxxx\privateKey.pem
    ProxyCommand ssh targetbastion -W %h:%p
Host targetbastion
    HostName 140.xx.xx.xx
User opc
    StrictHostKeyChecking=no
    IdentityFile C:\xxxxx\privateKey.pem
```

- (For 19c through 19.3.0.4.5) If stack has only a private IP, then
  in Host targetbastion, set bastion host with a public IP.
  Otherwise, use stack public IP.
- (For 19.3.0.5.6) for a stack with a private IP only, create
  Bastion and create a session for the node, as described in
  Access Oracle Essbase Using SSH. Then copy the ssh
  command and do the following:

- Replace <privateKey> -> path to the private key (from the SSH key pair used to create the session).
- ii. Remove -p 22 from command wherever it occurs.
- iii. Add -L 7002:private ip stack:7002 after ProxyCommand.

#### Sample ssh command:

```
ssh -i <privateKey> -o ProxyCommand="ssh -i <privateKey> -
W %h:%p -p 22
ocid1.bastionsession.oc1.**.***@host.bastion.***.oci.oracle
cloud.com" -p 22 opc@private ip stack
```

#### New edited ssh command:

```
ssh -i <privateKey> -o ProxyCommand="ssh -i <privateKey> -
W %h:%p
ocid1.bastionsession.oc1.**.***@host.bastion.***.oci.oracle
cloud.com" -L 7002:private_ip_stack:7002
opc@private_ip_stack
```

#### For example:

```
ssh -i C:\***\privateKey.pem -o ProxyCommand="ssh -i
C:\***\privateKey_org.pem -W %h:%p
ocid1.bastionsession.oc1.phx.***@host.bastion.us-
phoenix-1.oci.oraclecloud.com" -L 7002:10.100.100.100:7002
opc@10.100.100.100
```

**b.** You can then log in to WebLogic console (with open tunnel):

```
https://localhost:7002/console
```

- 2. Log in to WebLogic Server Admin Console.
- 3. On the home page, under **Domain Structure**, select **Security Realms**.
- 4. Under Summary of Security Realms > Realms, select myrealm.
- 5. Under Settings for myrealm, select Users and Groups tab.
- 6. Under **Users**, select or click on your WebLogic admin username, for example, WebLogic or admin (for Marketplace).
- 7. Select Passwords tab.
- 8. Enter the new password twice, and click **Save**.
- 9. Go to essbase\_domain > environment > servers.
- **10.** In the **Control** tab, select both servers, and then perform shutdown. You can do a force shutdown if no tasks are being processed.
- 11. Follow the steps to ssh using the private key. See Access Oracle Essbase Using SSH.



12. After ssh, change user to oracle:

```
sudo su oracle
```

13. Run the following command:

```
\verb|sh|/u01/config/domains/essbase_domain/esstools/bin/start.sh|
```

**14.** In the prompt for WebLogic user and password, enter the admin's user name and changed password. It is registered in the appropriate boot.properties file.

#### Steps to Reset the Admin Password for Instance on OCI

- Follow the steps to ssh using the private key. See Access Oracle Essbase Using SSH.
- 2. After ssh, change user to oracle:

```
sudo su oracle
```

3. Set domain home variable:

```
\verb|export DOMAIN_HOME| = | \verb|u01/config/domains/essbase_domain| \\
```

Then switch to domain directory:

```
cd $DOMAIN HOME
```

5. Stop the servers using:

```
$DOMAIN HOME/esstools/bin/stop.sh
```

6. Move old AdminServer data to different location:

```
mv $DOMAIN_HOME/servers/AdminServer/data $DOMAIN_HOME/servers/
AdminServer/data old
```

- 7. Set the environment variables:
  - . \$DOMAIN\_HOME/bin/setDomainEnv.sh
- 8. Reset the password using the following command. Remember to substitute the appropriate username and password.

```
cd $DOMAIN_HOME/security
java weblogic.security.utils.AdminAccount <adminuser>
<newpassword> .
```

**9.** Start Essbase services. You'll be prompted for the admin/password on startup and your changed password will be registered in the appropriate boot.properties file.

```
$DOMAIN HOME/esstools/bin/start.sh
```



#### For example:

```
opc@testhost> ssh -i <privatekey> -o ProxyCommand="ssh -i <privatekey> -
W %h:%p -p 22 ocid1.bastionsession.XXXXX" -p 22 opc@10.XX.XX.XX
Last login: Fri Feb 11 07:38:13 2022 from 10.XX.XX.XX
Welcome to Oracle Essbase on OCI 19.3.0.5.6-SNAPSHOT
Running Oracle Essbase 19.3.0.5.6 (Build 042)
Effective kernel version is 5.4.17-2136.302.7.2.2.el7uek.x86 64
[opc@essxx-1 ~]$ sudo su oracle
[oracle@essxx-1 opc]$ export DOMAIN HOME=/u01/config/domains/
essbase domain
[oracle@essxx-1 opc]$ cd $DOMAIN HOME
[oracle@essxx-1 essbase domain] $ $DOMAIN HOME/esstools/bin/stop.sh
Stopping domain; Using domainHome: /u01/config/domains/essbase domain ...
----stop script output ----
Stopping all managed servers and system components ...
Stopping ess server1 (Original State: RUNNING) ...
Stopped ess server1
----stop script output ----
data $DOMAIN HOME/servers/AdminServer/data old
[oracle@essxx-1 essbase domain]$ . $DOMAIN HOME/bin/setDomainEnv.sh
[oracle@essxx-1 essbase domain] $ cd $DOMAIN HOME/security
[oracle@essxx-1 security]$ java weblogic.security.utils.AdminAccount
admin pwdxxx.
[oracle@essxx-1 security] $ $DOMAIN HOME/esstools/bin/start.sh
--- start output---
Requesting credentials ...
Enter Weblogic login details at prompt
Weblogic Username: admin
Weblogic Password:
--- start output---
NodeManager (essxx-1:9556): RUNNING
Name
               Type
                              Machine
                                                       Status
ess server1
               Server
                              essxx-1.app.essxx.oraclevcn.com RUNNING
AdminServer
               Server
                              essxx-1
                                              RUNNING
```



4

# Migrate Applications and Users

If you have existing applications from an Essbase 11g On-Premise installation, or a cloud service instance, you can migrate them.

#### Topics:

- Selective and Ordered Import of Artifacts
- Migrate Essbase 11g On-Premise Applications and Users
- Migrate Essbase Cloud Service Applications and Users
- Test the Migrated Essbase Instance

# Selective and Ordered Import of Artifacts

You can control import of Essbase artifacts using a selection list text file, for on-premise migrations (using the standalone LCM utility) and for cloud service migrations (using the CLI tool).

A selection list text file contains a list of all artifacts in the exported zip that are grouped by section. You can generate the file during export using lomexport command. At the end of the file is an IMPORT section that contains the list of artifact entries to be imported.

You can edit the file and delete, or comment, the rows of artifacts that you want to skip in the import, using lcmimport command. You provide the text file as an argument in lcmimport operation. You can also control the order of import.

#### Sample selection list text file

#### How to use this feature

- When we use LCM utility for export, you can specify the optional argument generateartifactlist to generate a text file containing a list of exported artifacts.
- To skip a complete category of files, such as .rul files, comment the corresponding
   IMPORT section at the end of the text file.
- To skip specific files, delete or comment those entries in the text file.

 To control the import order, rearrange the entries under any specific category into the order that you prefer them to be imported. Files are then imported in the order listed under that category. During import, specify this file using

-al,-artifactlist

- Note that the lcmimport command has an -overwrite option.
  - If -overwrite is true, the import operation recreates the entire application. It
    only imports the artifacts or files that are listed in the text file.
  - If -overwrite is false, the import operation imports just the artifacts or files that aren't commented in the text file. It doesn't impact other artifacts already present in the target application.

#### Sample use cases

#### Import only the data from exported zip

You have an exported zip of Sample app and want to just import the data from Sample/Basic.

- In the text file generated during lcmexport, comment all the import entries,
   except import @Databases/Basic.
- Also comment /Sample/Databases/Basic/Basic outline under @Databases/Basic, just to import data alone.
- Note that -overwrite option is not valid for this use case ("data only" import).
   The reason is that during import, LCM will drop the entire application and import it as blank. Then, only data is attempted to be imported, without the outline, therefore making the application invalid.

#### Import outline only

You want to update the Sample.Basic cube with just the outline from the exported zip.

- In the IMPORT section at the end of the text file, comment all entries except "import @Databases/Basic".
- Also comment "/Sample/Databases/Basic/Data" under "@Databases/Basic", just to import the outline.

#### Import single cube for an application with multiple cubes

Sample application has three cubes named Basic, Basic1, Basic2, and you want to just import Basic.

- In the IMPORT section at the end of the text file, comment all entries except "Basic" cube (import @Databases/Basic, import @Databases/Basic/Xml\_files, etc.).
- Without the -overwrite option, it imports or overrides only the Basic cube, whereas other cubes (Basic1, Basic2) in that application, remain as they are without any impact.
- With the -overwrite option, it drops and recreates the application, with just the Basic cube.

# Migrate Essbase 11g On-Premise Applications and Users

You can migrate Essbase 11g On-Premise applications, cubes (databases), and users.



Moving all elements to the same data center, particularly for large volumes of data, removes uncertainty about added network latency. Files and databases are local to Essbase and are accessed as efficiently as if they were based on-premises. You can use the Lifecycle Management (LCM) utility, to export an on-premises cube to a zip file, and then use the Command Line Tool (CLI) to import the zip file that imports Essbase 11g On-Premise applications, folders, and elements.

#### Topics:

- Prepare to Migrate from Essbase 11g On-Premises Applications
- Migrated Essbase 11g On-Premises Artifacts
- · LCM Utility Export Options
- Migrate Essbase 11g On-Premises Users and Groups
- Migrate an Essbase 11g On-Premises Application Using LCM Utility
- Export Essbase 11g On-Premises Cubes
- Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode
- Upgrade Aggregate Storage Outline Version

## Prepare to Migrate from Essbase 11g On-Premises Applications

If you have an existing Essbase 11g On-Premise application and cube to migrate, review the following considerations and prerequisites.

#### **Task flow for Migrating**

Note that if Shared Services in Essbase 11g On-Premise was configured to use an external security provider, then steps 1 and 2 below aren't required. Configure Oracle Identity Cloud Service to use the same external security provider as used in Essbase 11g On-Premise.

- 1. Export users and groups from Essbase 11g On-Premise Shared Services.
- 2. Import users and groups to Oracle Identity Cloud Service.
- 3. Export Essbase applications using the Life Cycle Management (LCM) standalone utility, by running the utility on the computer where Essbase 11g On-Premise is installed.
- 4. Import Essbase applications using the Essbase command-line tool (CLI).

#### **Required User Roles for Access**

Note that the following Essbase security artifacts are migrated using the LCM utility: Essbase server-level roles, application-level roles, filter associations, and calc associations. LCM handles provisioning users and groups with the corresponding new roles.

Table 4-1 Default role mapping

Source EPM System Security Mode Roles	Target WebLogic Security Roles	Level
Administrator	Service Administrator	Server
Application Manager	Application Manager	Application
Calc	Database Update	Application
Create/Delete application	Power User	Server
Database Manager	Database Manager	Application



Source EPM System Security Mode Roles	Target WebLogic Security Roles	Level
Filter	Database Access	Application
Read	Database Access	Application
Server Access	User	Server
Start/Stop Application	Database Access	Application
Write	Database Update	Application

Note that Filter role in Essbase 11g On-Premise doesn't allow Read access, but allows access to members restricted by the filter. Now, there's no Filter role, and the lowest role access is Database Access, which allows Read access to all members. To restrict access to selective members, use a group filter that restricts global access.

To export and import, the following access is required:

- For exporting: A user with at least Application Manager role, for the application created, can export applications, folders, and artifacts.
  - In addition, the following roles can use the LCM utility and their corresponding operations: Service Administrator for all applications; Power User for all applications created by the Power User.
- For importing: A user with at least Power User role can create applications (during import) and manage applications.

#### Lifecycle Management Utility

With the Lifecycle Management (LCM) utility, you can create applications by exporting on-premises applications and cubes. You then import them using the Essbase command-line tool (CLI).

To use the LCM utility, you must have installed Java Development Kit 8 or higher, and have set the JAVA\_HOME environment variable. You must also set EPM\_ORACLE\_HOME and EPM\_ORACLE\_INSTANCE variables in the shell terminal.

#### **Convert Applications to Unicode**

You must convert all Essbase 11g On-Premise applications to Unicode as follows:

- Use Alter System on the server, and then on a backup copy of the Essbase application, prior to running LCM utility export, to enable Essbase itself to support the Unicode application.
- For non-Unicode, Block Storage Cube applications, export the application using "converttoutf8" option in the export command. See LCM Utility Export Options.
  For non-Unicode, Aggregate Storage Cube applications, follow the manual
  Unicode conversion instructions in Upgrade Aggregate Storage Cube Outline
  Version.

#### **Supported Essbase Versions**

The following releases have been tested for migration: 11.1.2.3.0nn, 11.1.2.4.0nn, 12.2.1, and later. Older versions can be used - work with Oracle Support if you need any assistance.



#### **Hybrid Mode**

The default calculation and query processor is hybrid mode. Hybrid mode enables block storage cubes to have dynamic, upper-level sparse members, and fully dynamic query and calculation. You can query data immediately after updating it, without running batch calculations. In hybrid mode, there is no impact to your cubes if you choose not to apply Dynamic Calc to upper-level sparse members.

#### **Implied Sharing**

Implied sharing isn't applicable. All stored intersections have data, regardless of their child count.

#### **Configuration Settings**

Configuration is now done at the application level, and default configuration values are different.

- IGNORECONSTANTS setting is now TRUE by default. Calculations in hybrid mode don't assign constants.
- INDEXCACHESIZE and DATACACHESIZE settings now control cache sizes for all Essbase cubes (except for aggregate storage cubes). Formerly, these settings only affected newly created or migrated cubes.
   You can't change the cache sizes using MaxL. You can only change the cache sizes using these configuration settings.
- GRIDSUPPRESSINVALID is now TRUE by default. Invalid intersections aren't displayed in Smart View grids.
- QRYGOVEXECTIME is now set by default to 300 seconds, meaning that queries time out if not completed in that timeframe.

In addition to the above noted configuration changes, you can modify the default values for the application-level configuration settings.

Oracle recommends managing all configuration settings at the application level. Application-level configuration is preserved during the LCM utility export and import processes.

#### **Application Files and Artifacts**

It is recommended to convert all application-level files and artifacts, such as calculation scripts, rule files, and text files, to database-level files and artifacts before you export them from Essbase 11g On-Premise instances, and before you migrate them. Artifacts are supported at the database level.

You can import rule files and execute them.

If you encounter file upload size restrictions between external clients and Essbase, then you may need to split large files into smaller files and then concatenate them together after uploading them, using an SSH connection to the server. This option is only available for Essbase users in a customer-managed environment.

#### **Outlines**

Outlines are encrypted on the deployment servers. If you need to export and import outlines between deployment servers, the LCM command-line utility and application workbooks are the supported methods.

#### **Users and Groups**



If you want filters and calculation assignments of existing users to be migrated, ensure that Essbase has the same set of users and groups already available.

Assignment of user roles behavior differs from Essbase 11g On-Premise. Database Access is now the lowest role, and has, by default, read access to data values in all cells. To restrict access to data values in Essbase, you must now create a NONE filter and assign it to users and groups. This was not a requirement in Essbase 11g On-Premise, where Filter was the lowest role, and has, by default, no access to data values in all cells.

#### **Unsupported Application and Database Settings**

The following application- and database-level settings aren't applicable in Essbase instances:

- Enable/disable Commands (enabled by default)
- Enable/disable Connects (enabled by default)
- Enable/disable Updates (enabled by default)
- Data and index cache size controls (defaults are fixed, but can be changed per application using INDEXCACHESIZE and DATACACHESIZE configuration settings)
- Minimum permission levels (create security filters prior to LCM export instead)
- Set lock timeout
- Currency conversion
- Disk volumes

#### **Partitions**

When you perform the LCM utility import operation, import the source applications before the target applications. If you don't import source applications prior to target applications, then the partition definition won't work, and you must re-create the partition definition after importing source applications.

#### **Size Requirements**

Ensure that pre-existing applications you plan to migrate will fit at the resource level you procure. Estimate sizing requirements, and procure the most relevant combination of CPU, memory, and storage.

#### **Application Creation Options Other than LCM**

In addition to using LCM to migrate exported applications, you can also create applications in the following ways:

- Import using Excel application workbooks
- In Smart View, use the Cube Designer extension
- MaxL create application statement

## Migrated Essbase 11q On-Premises Artifacts

The following table describes which global, application-level, and cube-level Essbase artifacts you can migrate from Essbase 11g On-Premise, using the Lifecycle Management (LCM) utility. A .zip file, created by the LCM utility, contains the artifacts of the exported application.



Artifact	Supported for migration	Exceptions/Comments
Application and cube metadata	yes	Application metadata includes application type and settings. Cube metadata includes cube properties and settings.
Calculation scripts	yes	Application- and cube-level calculations are migrated. To see the calculation scripts, you must move application-level scripts to the cube level using the file catalog.
Data	yes	To be migrated, data must be in the cube directory in the file catalog.
Disk volumes	no	Disk volume definitions are not applicable.
Drill through definitions	yes	-
Excel workbooks and files	yes	-
Filters	yes	Cube-level filters and user- created filters are migrated.
Linked Reporting Objects (LROs)	no	-
Location aliases	yes	Location aliases are migrated with the cube.
Log files	no	-
Outlines and formulas	yes	-
Partitions	yes	Replicated and transparent partitions are migrated.
		Only partition definitions from the target cube are exported to the file system.
		When migrating the partitioned cubes, you must import the source cube before the target cube; otherwise, partition definitions may not be restored.
Report scripts	yes	Report scripts are migrated at both application and cube levels.
Rule files, text files, .csv files	yes	Application- and cube-level files are migrated.
Scenarios	NA	Scenarios do not apply for Essbase 11g On-Premise applications.
Substitution variables	yes	Application- and cube-level substitution variables are migrated. Server-level substitution variables are migrated if you use the optional – include-server-level option.
Users	no	-
User roles	-	User roles are migrated if you use the -exportepmroles option.



# LCM Utility Export Options

You have the following options to use LCM utility to export from Essbase 11g On-Premise.

- You can specify -converttoutf8 option during export to automatically convert the Essbase 11g On-Premise application to Unicode, before exporting it to a .zip file. Note that this will convert the source application to Unicode; it is recommended to run a backup before specifying this option.
- You can specify the arguments (-server, -user, -password, -application, -zipfile) in any order.
- To prompt for a password, do not include the -password password option.
- To skip the export of cube data during export, specify -nodata, which is an optional argument). By default, all cube data is exported.

Command	Description
-server <essbasehost:port></essbasehost:port>	Server host name and port number
-user <username></username>	Server user name
-password <password></password>	Server password
-application <appname></appname>	Application name
-zipfile zipfilename	Export zip file
-nodata	(Optional) Skip data export
-overwrite	(Optional) Overwrite file in local directory with exported file
-converttoutf8	(Optional) Convert application to Unicode. Prompts you to type Y to confirm
-forceutf8	(Optional) Same as -converttoutf8, but without any prompt. Can be used in automation scripts
-exportepmroles	(Optional) Exports Essbase roles from Enterprise Performance Management (EPM) source
-generateartifactlist	(Optional) Generate file with list of artifacts from zip export

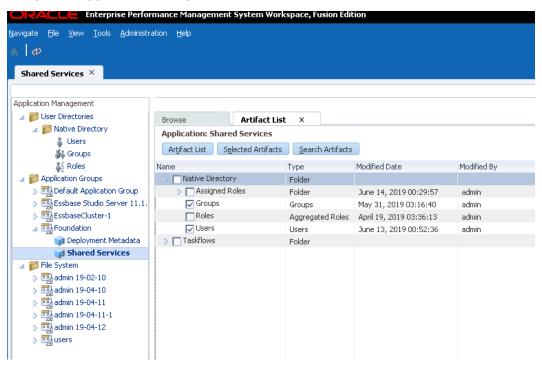


(Optional) Include server-level artifacts, such as server-level substitution variables and server-level roles
u

### Migrate Essbase 11g On-Premises Users and Groups

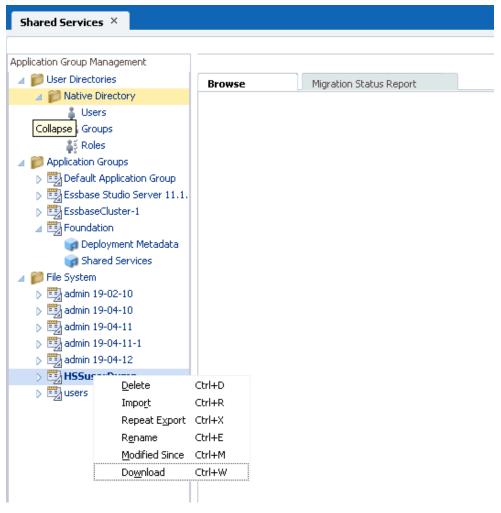
If you're using native default Shared Services security, these steps are required. If Shared Services uses an external security provider, or you're using federated setup in Shared Services, the following steps are optional. You should configure Oracle Identity Cloud Service to use the same external security provider that Shared Services used.

 Launch the Enterprise Performance Management (EPM) Shared Services user interface. Navigate to Application Groups, Foundation, and then Shared Services.

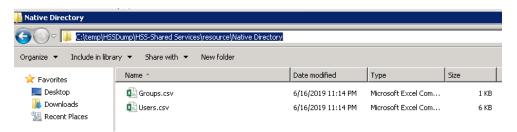


- 2. Select **Groups** and **Users** check boxes, and click **Export**.
- On the Export to File System dialog box, enter a name for the target File Systems Folder and click Export.
- 4. After the export completes, the exported Shared Services-based security content appears under the target File System folder. Right-click on the exported Shared Services content, and click Download to download the files locally.





5. Expand the downloaded zip file. For each folder, you may see files like those shown below.



6. The format of the generated Users.csv and Groups.csv files are not compatible with Identity Cloud Service format. You must manually reorganize or map the files to Identity Cloud Service format, as shown below, using a CSV or text editor.



Exported users and groups	Shared Services format	Identity Cloud Service format
Users.csv	id,provider,login_name,first_na me,last_name,description,email ,internal_id,password,activ	User ID,Last Name,First Name,Middle Name,Honorific Prefix,Honorific Suffix,Display Name,Nick Name,Profile URL,Title,User Type,Locale,Preferred Language,TimeZone,Active,Pa ssword,Work Email,Home Email,Primary Email Type,Work Phone,Mobile No,Work Street Address,Work City,Work State,Work Postal Code,Work Country,Federated,Employee Number,Cost Center,Organization,Division,D epartment,Manager Name
Groups.csv	id,provider,name,description,int ernal_id	Display Name,Description,User Members
	(Following data for each group:)	
	#group_children	
	id,group_id,group_provider,user _id,user_provider	

For columns that are empty, enter dummy text.

7. Import the users and groups to Identity Cloud Service according to the instructions in Import a Batch of Users into a Cloud Account with Identity Cloud Service.

### Migrate an Essbase 11g On-Premises Application Using LCM Utility

Use the Life Cycle Management (LCM) standalone utility to migrate Essbase 11g On-Premise applications from versions 11.1.2.3.0nn, 11.1.2.4.0nn, 11.1.2.4.5nn, 12.2.1, and later.

The workflow to migrate using the LCM utility is as follows:

- 1. **Download the LCM utility**: In the Essbase web interface, click **Console**, expand **Command Line Tools**, and download the Life Cycle Management utility (EssbaseLCMUtility.zip). The downloaded utility must be copied to and run on the same machine as the Essbase 11g On-Premise or 12.2.1 installation, for Enterprise Performance Management (EPM) roles to be exported.
- 2. Set up the LCM utility: In the uncompressed downloaded file, run EssbaseLCM.bat (Windows) or EssbaseLCM.sh (Linux), based on the platform on which you want to run the utility.
- 3. Run the export: To export each Essbase 11g On-Premise or 12.2.1 application and its artifacts to the specified .zip file, use the LCM utility with the parameters below, and run the export command.

At the LCM utility command prompt, enter the following command syntax to export the application to a .zip file:

export -server <hostname> <port> -user username -password password application appname -zipFile zipfilename [-nodata] [-include-serverlevel] -exportepmroles [-generateartifactlist]



#### Specify these options:

-exportepmroles : (optional) Export Enterprise Performance Management (EPM) roles

-include-server-level: (optional) If you want to include server-level artifacts, such as server-level substitution variables or server-level roles

-generateartifactlist: (optional) Generate artifact list

#### For example:

```
EssbaseLCM.sh export -server localhost:1423 -user admin -password password -application Sample -zipfile Sample.zip -include-server-level -exportepmroles -generateartifactlist
```

This export additional artifacts: user and group server-level roles, application-level roles, calculations, and filter associations.

**4. Run the import**: To import the application, use the Essbase command-line tool (CLI) to upload the .zip file to target application.

The syntax for the CLI lcmimport command is as follows:

```
lcmImport [-verbose] -zipfilename filename [-overwrite] [-
targetappName targetApplicationName] [-artifactList]
```

When partitions exist in the source between a source application or database, and a target application or database, only partitions from the target are exported to the file system. When partitions exist between cubes being migrated, you must import the data source before the data target. Otherwise, partition definitions may not be restored.

Roles are set only if the users are available in Oracle Identity Cloud Service. You can override default role mapping by changing the mapping in CSSMappings.xml provided with LCM utility.

- 5. Upgrade Aggregate Storage Cube outline version: After import of Aggregate Storage Cube applications, the outline must be converted to Unicode using ESSCMDQ. See Upgrade Aggregate Storage Cube Outline Version.
- **6. Validate**: Log in to the Essbase web interface to see the application and cube on the Applications page.

### Export Essbase 11g On-Premises Cubes

If you have applications and cubes that were created in a supported on-premises instance of Essbase, then you can use the cube export utility, which is a command-line tool, to export the metadata and data of a cube into an application workbook. Then you can import the application workbook to create a cube in the cloud service.

Using the cube export utility, you can export applications and cubes created in Essbase on-premises instances: 11.1.2.4.0nn, 11.1.2.4.5nn, 12.2.1, and later. You can't export cubes on these releases to application workbooks.

#### See:

Download the Cube Export Utility



 Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility

### Download the Cube Export Utility

The cube export utility is supported on Windows and UNIX/Linux.

To download the cube export utility from Essbase:

- 1. In the Essbase web interface, on the Applications page, click **Console**.
- 2. On the Console page, click **Download**

next to Export Utility.

3. Save the cube export utility, which is named <code>dbxtool.zip</code>, to a local drive.

# Review Member Names Before you Import an Application Workbook Created by the Cube Export Utility

When importing an application workbook that was created using the cube export utility, you should carefully review member names in the application workbook. Member names are exported to the application workbook as is. If a member name ends with a backslash (for example, mbrname\) or mbr\name\), then the member name is exported to the application workbook as is (mbrname\) or mbr\name\). During the import process, however, the trailing backslash is interpreted as an escape character and the member is rejected (not added to the cube outline).

When the import process is completed, a dialog box provides status details, such as whether a dimension build was successful or if errors were encountered.

For each dimension in which one or more member names are rejected, an error file is created. The error file is named <code>err\_DimName.txt</code> or <code>err\_Dim\_DimName.txt</code>. For example, if the Year dimension has any rejected member names, the error file name is <code>err\_Year.txt</code> or <code>err\_Dim\_Year.txt</code>.

In the dimension error file, each rejected member name is listed, as shown:

```
\Record #98 - Error in association transaction [RB6300] to [Curr_EUR] (3362) "OTHER", "RB6300", "N", "", "", "Ballsport L", "", "", "Curr EUR"
```

The rejected member record text files are available on the Files page. Review the text files and correct the issues in the application workbook.

# Convert Non-Unicode Aggregate Storage Cube Application to Unicode Mode

These are the steps to perform before exporting an Aggregate Storage Cube Essbase 11g application.

#### Workflow

 Convert the copied ASO application to Unicode mode using MaxL shell, as described below.



2. Change the ESSLANG value within the source outline from native encoding to UTF-8, as described below.

#### Convert the copied ASO application to Unicode mode using MaxL shell:

- Log in to the source Essbase 11g instance using MaxL shell.
- **2.** Execute MaxL statement alter application <copied\_app> set type unicode mode to convert the application to Unicode.

#### For example:

MaxL> alter application SampleBck set type unicode mode;

For more details on the MaxL, see Alter Application.

# Change the ESSLANG value within the source outline from native encoding to UTF-8.



All of the following operations must be performed on the copied application and not the source application.

- 1. Download ESSCMDQ,
  - Download platform-specific "11.1.2.4.010+" version of ESSCMDQ from Download ESSCMDQ to your source EPM 11g instance.
  - b. Unzip the files directly to the same directory where ESSCMD is present in the installation, \$ESSBASEPATH/bin.

#### For example:

./Middleware/EPMSystem11R1/products/Essbase/EssbaseServer/bin/ESSCMDQ

# To know the values of environment variables \$ESSBASEPATH and \$ARBORPATH in your source EPM 11g installation, check the file

./Middleware/user\_projects/<epm instance>/EssbaseServer/essbaseserver1/bin/setEssbaseEnv.sh

By default, instance> would be epmsystem1.

c. Make a copy of the existing script

./Middleware/user\_projects/epmsystem1/EssbaseServer/
essbaseserver1/bin/startEsscmd.sh

#### as

./Middleware/user\_projects/epmsystem1/EssbaseServer/essbaseserver1/bin/startEsscmdQ.sh.



Within this newly created script, change the call from ESSCMD to ESSCMDQ.

d. Just before this last line (just before the call to ESSCMDQ), add the following lines:

```
export ESSCMDQ_UTF8MODE=1
export ESSLANG=.UTF-8@Binary
```

- 2. Make sure that you have stopped the copied application before converting the outline.
- 3. Now create a "client" folder under \$ARBORPATH.
- Copy the application folder from \$ARBORPATH/app directory to client directory. For ASOBck app,

```
$ARBORPATH/app/ASOBck as $ARBORPATH/client/ASOBck
```

**5.** Execute the following commands in ESSCMDQ after launching:

```
./Middleware/user_projects/epmsystem1/EssbaseServer/essbaseserver1/bin/startEsscmdQ.sh
openotlex2 1 1 appName dbName outlineName Y Y Locale N 0
```

#### and then:

```
writeotlex 0 1 1 appName dbName outlineName 2
```

Note that Locale should be the native ESSLANG value used in the source Essbase 11g environment.

#### For example:

```
#Create client directory and copy the application folder
mkdir $ARBORPATH/client
cp -r $ARBORPATH/app/ASOBck $ARBORPATH/client

#Launch ESSCMDQ
./Middleware/user_projects/epmsystem1/EssbaseServer/essbaseserver1/bin/
startEsscmdQ.sh

#Execute the below within ESSCMDQ
#Note - Specify the native ESSLANG value used in the Essbase 11g
environment in place of "Japanese_Japan.MS932@Binary" below.
openotlex2 1 1 ASOBck Basic Basic Y Y "Japanese_Japan.MS932@Binary" N 0
writeotlex 0 1 1 ASOBck Basic Basic 2

#Exit ESSCMDQ
exit
```

Make sure that no errors are displayed while executing the above commands. Then, for each cube, copy just the outline file from the client directory back to the application directory.

#### For example:

#Now copy back the converted outline only for each cube. For ASOBck app cp \$ARBORPATH/client/ASOBck/Basic/Basic.otl \$ARBORPATH/app/ASOBck/Basic/

```
Basic.otl
```

#Note: The artifact files (.txt or .csc), which were created in native locale, may need to be converted to UTF-8 manually using third party tools which help in converting text encoding.

#### 7. Launch ESSCMDQ again using:

```
./Middleware/user_projects/epmsystem1/EssbaseServer/
essbaseserver1/bin/startEsscmdQ.sh
```

#### and restructure each cube.

```
#Please replace hostname, username, password, appname and cubename
with appropriate values
login 'hostname' 'username' 'password'
select appname cubename
openot1 2 1 appname cubename outlinename y y 0
writeot1 0 2 1 appname cubename outlinename
restructot1 1
closeot1 0
unlockobj 1 appname cubename outlinename
logout
exit
```

#### For example:

```
login localhost:1423 user password
select ASOBck Basic
openotl 2 1 ASOBck Basic Basic y y 0
writeotl 0 2 1 ASOBck Basic Basic
restructotl 1
closeotl 0
unlockobj 1 ASOBck Basic Basic
logout
exit
```

### **Upgrade Aggregate Storage Outline Version**

These are the steps to perform to upgrade an aggregate storage outline to Essbase 21c.

- 1. Note that these steps can be performed only after importing the outline.
  - Download platform- specific Essbase 21c ESSCMDQ from Download ESSCMDQ to your target Essbase system.
  - **b.** Unzip the files directly to the same directory where ESSCMD is present in your installation.
  - c. Make a copy of the existing script:

```
./Oracle/domains/esscs/esstools/bin/startESSCMD.sh
```



#### as

./Oracle/domains/esscs/esstools/bin/startESSCMDQ.sh

#### Within this newly created script, change the call from:

<Essbase\_Product\_Home>/products/Essbase/EssbaseServer/bin/
startESSCMD.sh

#### to

<Essbase\_Product\_Home>/products/Essbase/EssbaseServer/bin/
startESSCMDQ.sh

#### d. Make a copy of the script:

<Essbase\_Product\_Home>/products/Essbase/EssbaseServer/bin/
startESSCMD.sh

#### as:

<Essbase\_Product\_Home>/products/Essbase/EssbaseServer/bin/
startESSCMDQ.sh

#### e. Edit

<Essbase\_Product\_Home>/products/Essbase/EssbaseServer/bin/
startESSCMDQ.sh

#### and change the last line from

<EssbaseBasePath>/bin/ESSCMD

#### to:

<EssbaseBasePath>/bin/ESSCMDQ

#### 2. Just before this last line, add the following lines:

```
export ESSCMDQ_UTF8MODE=1
export ESSLANG=.UTF-8@Binary
```

#### 3. After launching this script:

<Essbase Product Home>/products/Essbase/EssbaseServer/bin/startESSCMDQ.sh

#### execute the following commands in ESSCMDQ:

login hostname username password;



```
select "appname" "cubename";

openotl 2 1 "appname" "cubename" "outlinename" y y 0;
setopgversion 0 "111241";
getopgversion 0;

writeotl 0 2 1 "appname" "cubename" "outlinename";
restructotl 1;
closeotl 0;
unlockobj 1 "appname" "cubename" "outlinename";
logout;
exit;
```

# Migrate Essbase Cloud Service Applications and Users

You can migrate applications and cubes (databases) from cloud service instances. Learn how to prepare for migration, and review some use cases for migrating.

You can use the Essbase command-line tool (CLI) to migrate your source application and artifacts from cloud deployments and releases. This is used to migrate applications one-at-a-time.

You can use the Migration Utility tool to migrate multiple applications, artifacts, users, and groups at one time, from Oracle Analytics Cloud – Essbase.

- Prepare to Migrate Cloud Service Applications and Users
- Migrated Cloud Service Artifacts
- Migrate Cloud Service Applications Using CLI Tool
- Migrate Cloud Service Applications Using Migration Utility
- Migrate from FCCS or PBCS

### Prepare to Migrate Cloud Service Applications and Users

Here are some considerations and requirements when migrating cloud service applications from Oracle Analytics Cloud - Essbase.

You can use the Essbase command-line tool (CLI) to migrate your source application and artifacts across Oracle Analytics Cloud - Essbase deployments and releases. This is used to migrate applications one-at-a-time.

You can use Migration Utility tool to migrate multiple applications, artifacts, and users at one time, across Essbase cloud services.

- If you're migrating across Essbase cloud deployments and releases, from v17.3.3 (or earlier), use the scripts for migrating to Essbase. See Scripts for Administration Tasks. This also applies to export and import of provisioned application roles and scripts.
- Restoring an application or database from a prior backup, after the application or database was re-created using LCM import, isn't supported.
- Global variables, email configuration settings, and file scanner settings must be set on the target instance before using any of the migration tools.



- Oracle Identity Cloud Service roles aren't supported in Essbase.
- After migration from Oracle Analytics Cloud Essbase, Identity Cloud Service provisioning doesn't continue to be honored. However, Essbase server/application level role assignments are migrated.
- All Identity Cloud Service users and groups listed in the Security provisioning page in the Essbase web interface can be provisioned with one of the Essbase roles: User, Power User, and Service Administrator.
- Migration Utility can migrate users and groups from embedded LDAP (or from Identity Cloud Service) to Identity Cloud Service in addition to all Essbase applications.
- If you're migrating users and groups from an LDAP source to an Essbase instance, Identity Cloud Service doesn't support nested groups. Therefore, group associations to other parent groups, from an LDAP source instance, aren't migrated to Identity Cloud Service targets, when using Migration Utility.
- Any users or groups that exist with the same name in the target environment as in the source environment, aren't updated in the target.
- To run Essbase CLI or Migration Utility, use the Identity Cloud Service user that you
  provisioned to be the initial Essbase Service Administrator during the Essbase
  deployment and setup.
- When you run Migration utility for SSL connection, include the host (-Dhttps.proxyHost) and port (-Dhttps.proxyPort) proxy settings in the command line.

The required user roles are as follows:

- For exporting: Application manager for the application created. In addition, the following roles can use LCM utility and CLI tool: Service Administrator for all applications; Power User for all applications created by the Power User.
- For importing: Power User or Service Administrator, for creating new applications during import. If you use the Power User role, then the target applications are owned by the Power User used in the migration.

## Migrated Cloud Service Artifacts

The following table describes which global, application-level, and cube-level Essbase artifacts you can migrate between cloud service instances.

Artifact	Supported For Oracle Analytics Cloud migration	Exceptions/Comments
Application and cube metadata	yes	Application metadata includes application type and settings. Cube metadata includes cube properties and settings.
Application-level configuration files	yes	If these files exist, they're migrated.
Calculation scripts	yes	Application- and cube-level calculations are migrated.



Artifact	Supported For Oracle Analytics Cloud migration	Exceptions/Comments
Catalog server	no	Files listed under Files in the web interface under Applications/ <appname> are migrated. Other files stored under Shared/Users folders aren't migrated. You can manually download them in the web interface and restore them.</appname>
Connections and Datasources	yes	Using Migration Utility, system- and application-level connections and Datasources are migrated. Using CLI tool, connections and Datasources created at the application level are migrated. With both tools, you must include
		the following argument in lcmexport operations: - include-server-level (or its abbreviation -isl).
Data	yes	To be migrated, data must be in the cube directory on the cloud instance.
Disk volumes	NA	Disk volume definitions aren't applicable to Essbase cloud instances.
Drill through definitions	yes	Drill through definitions are migrated.
Excel workbooks and files	yes	Excel workbooks and files are migrated.
Filters	yes	Cube-level and user-created filters are migrated.
Global variables	yes	-
Layouts	yes	Cube-level layouts are migrated.
Linked Reporting Objects (LROs)	yes	LROs are included for backward compatibility with migrated on-premises applications.
Location aliases	yes	Location aliases are migrated with the cube.
Log files	no	Log files aren't migrated.
Named queries	yes	Cube-level named queries are migrated.
Outlines and formulas	yes	Formulas containing @XREF aren't migrated.
Partitions	yes	Replicated and transparent partitions are migrated.
		Only partition definitions from the target cube are exported to the file system.



Artifact	Supported For Oracle Analytics Cloud migration	Exceptions/Comments
Report scripts	yes	Application- and cube-level report scripts are migrated. The scripts are included for backward compatibility with migrated onpremises applications.
Rule files, text files, .csv files	yes	Application-and cube-level files are migrated.
Scenarios	yes	If a cube is scenario-enabled and has a Sandbox dimension, the related scenarios are migrated.
Substitution variables	yes	Application- and cube-level substitution variables are migrated. If you have global (server)-level substitution variables, you must convert them to application-level variables prior to migration, or recreate them in the Console after migration.
Users and groups	-	Users and groups are migrated using Migration Utility; they aren't migrated when using CLI tool.
User roles	yes	User roles can be migrated only from one Essbase cloud instance to another.
Wallet files	yes	Wallet files are migrated for the specified application.

# Migrate Cloud Service Applications Using CLI Tool

You can use Command-Line Interface (CLI) tool to migrate a source application and artifacts across Essbase cloud deployments and releases. The tool is used to migrate applications one-at-a-time.

The standard migration workflow using the Essbase Command Line Interface tool command-line tool (CLI) is as follows:

- 1. Download the tool and use the lcmexport commands to export individual applications one-by-one from source to a zip file.
- 2. Use the lcmimport command to import each individual application from a zip file to Oracle Essbase.

Use the following workflow when the source deployment is from Oracle Analytics Cloud (either Oracle Cloud Infrastructure or Oracle Cloud Infrastructure Classic) using Oracle Identity Cloud Service Federated Security.

- 1. Create a new Identity Cloud Service application. See Create a Confidential Identity Cloud Service Application.
- 2. Configure your Identity Cloud Service instance to point to the same external security provider used in Oracle Analytics Cloud.



- Download the CLI tool, and using the lcmexport command, export individual
  applications one-by-one from source to a zip file. Specify the -include-serverlevel option to include server-level roles.
- 4. Manually re-create, in the new Identity Cloud Service instance, any non-federated users (Identity Cloud Service local users) that you had in the source instance. This is necessary if you want to use tools like CLI, MaxL, or REST API.
- 5. Using the CLI lcmimport command, import each individual application from a zip file to Oracle Essbase.

When partitions exist in the source between a source application or database, and a target application or database, only partitions from the target are exported to the file system. When partitions exist between cubes being migrated, you must import the data source before the data target. Otherwise, partition definitions may not be restored.

# Migrate Cloud Service Applications Using Migration Utility

You can use Migration Utility to migrate source applications and elements from Oracle Analytics Cloud - Essbase deployments and Essbase OCI releases. The utility migrates multiple applications at one time. It also migrates artifacts, rules, users and groups.

As an Essbase Service Administrator user, you can use Migration Utility to migrate an entire instance (all applications, users and groups, and other artifacts) from one cloud instance to another in a single process. Note that the Essbase Command-line tool (CLI) commands, lcmimport and lcmexport, require you to migrate applications one-at-a-time, and do not migrate users and groups.

**Before** using the Migration Utility, the following is a prerequisite: Set Up the SSL Certificate.

Here are some use cases for migrating with Migration Utility.

- WebLogic LDAP users can migrate the users from LDAP in the source to Identity Cloud Service in the target.
- Use this utility for basic deployments that aren't customized. If your deployment includes customizations, such as custom Single Sign On solutions, use CLI tool instead of Migration Utility.
- When the source deployment is from
  - Oracle Analytics Cloud on Oracle Cloud Infrastructure, using Identity Cloud Service Native Security, or
  - Oracle Analytics Cloud on Oracle Cloud Infrastructure Classic, using Identity Cloud Service Native Security, or
  - Oracle Analytics Cloud on Oracle Cloud Infrastructure Classic, using Embedded LDAP
  - Essbase Marketplace on Oracle Cloud Infrastructure, using Identity Cloud Service Security, or
  - Essbase Marketplace on Oracle Cloud Infrastructure, using Embedded LDAP

then before using the Migration utility steps below to export, first create a new Identity Cloud Service application as outlined in Create a Confidential Identity Cloud Service Application. Also, before using the Migration Utility steps below to



import, change the host and Identity Cloud Service details in import.properties to point to the target Essbase instance.

#### To migrate cloud service applications and users using Migration Utility

- Before you use the utility, if you haven't already, patch your source Essbase instance to the latest version.
- If it isn't already installed, download and install Java SE Development Kit (JDK) 8 from Oracle Technology Network.
- 3. Set the JAVA\_HOME environment variable name on your system to point to the JDK installation folder. If the installation path contains any spaces, enclose the path in the variable value, within quotation marks, for example, "C:\Program Files\Java\jdk1.8.0 171".
- 4. Sign in to the target Essbase, and navigate to the Console tab.
- **5.** In the Console, go to Desktop Tools, and expand Command Line Tools.
- 6. Click **Download** next to **Migration Utility**.
- 7. Download the Migration Utility zip file to a local drive. For best results, choose a path that has no spaces, for example, C:\Oracle.
- 8. Extract the zip file, and see the extracted files (properties, jar, and readme) in the migrationTools folder.
- Before you run the import or export commands provided with Migration Utility, you must edit the properties files.
  - a. Edit the properties strings in the export.properties file:
    - userName Essbase administrator user name.
    - password Essbase administrator password.
    - host Essbase host or IP address.
    - port Essbase port. Enter the value of 80 for LDAP source. Otherwise, accept the default value of 443 (SSL/TLS) for Identity Cloud Service source.
    - proxy proxy server URL.
    - skiprole import of roles from Oracle Analytics Cloud Essbase is skipped. Value must be left empty or true.
  - b. Edit the properties strings in import.properties file:
    - userName Essbase administrator user name.
    - password Essbase administrator password.
    - host Essbase host or IP address.
    - port Essbase HTTP listening port. Default value 443 (SSL/TLS).
    - userPassword Initial password assigned for all new users.
    - proxy proxy server URL (optional)
  - c. Edit Identity Cloud Service information in import.properties. Obtain these values from the Service Console for Oracle Identity Cloud Service.
    - idcsHost Identity Cloud Service host
    - idcsTenant Identity Cloud Service tenant



- clientId Client identifier for OAuth authorization
- clientSecret Client secret for OAuth authorization
- appld Application identifier
- **10.** To run Migration Utility, use the following java command to export all applications, users, and groups from the Essbase source instance catalog to a tar file.

```
java -jar -Dhttps.proxyHost=proxy-url> -Dhttps.proxyPort=<nn>
    migrationTools.jar export export.properties <new tar file>
```

#### for example:

**11.** After you export from the source instance, use the following java command to import the data tar file to the target instance.

12. After you run the import, the data is stored in the Essbase catalog of the target instance. If any exported applications already exist on the target, they aren't overwritten.

### Migrate from FCCS or PBCS

You can migrate applications and databases from Oracle Financial Consolidation and Close Cloud Service or from Oracle Planning and Budgeting Cloud Service.

If applications imported from Oracle Financial Consolidation and Close Cloud Service to Essbase fail to load due to disk volume errors, see Knowledge base Doc ID 2707616.1, at My Oracle Support.

- Export from Oracle Planning and Budgeting Cloud Service or Oracle Financial Consolidation and Close Cloud Service using either the product interface, or the EPM Automate Utility command line tool using exportsnapshot. See EPM Automate Utility Commands in Working with EPM Automate for Oracle Enterprise Performance Management Cloud.
- 2. Run the Essbase command-line tool (CLI) to import the Essbase application and cubes from the exported .zip file using CLI command lcmimport.

# Test the Migrated Essbase Instance

After migrating your instance to Oracle Cloud Infrastructure, test thoroughly to ensure it's production-ready.

#### **Essbase post-migration tasks:**

 If you have any artifacts in LCM that are not supported for migration, they can be manually migrated.



- Test that the migrated data loads and dimension builds work as expected.
- Run a Smart View report to check connectivity and data.
- After the results look OK, scan the application logs for errors, warnings, and suspicious messages.



5

# Manage Users and Roles

Essbase integrates with security layers managed by Oracle to create a highly secure environment for the cloud. Service Administrators can assign appropriate user roles and application permissions in Essbase.

#### **Topics:**

- · About Users and Roles
- User Roles and Application Permissions
- Provision Application Permissions

### **About Users and Roles**

The following are common use cases for assigning access to users:

- Users can view and access cubes (databases) for which they were assigned access to the related applications.
- Power Users can create enterprise-level cubes and grant other users access to applications for which they have an Application Manager role.
- Service Administrators can assign users at all levels and manage all aspects of the applications, cubes, and users.
- Service Administrators can assign a Database Update role for users who need to update data in a cube.

To provide access to Essbase users, the following steps are required:

- Assign Essbase user role
- Assign Essbase application-level permissions

Access to Essbase is restricted by security, and managed by Oracle Identity Cloud Service. You create users and user groups in the Oracle Identity Cloud Service administration interface.

Oracle Identity Cloud Service doesn't support creating nested groups (assigning a group to a parent group).

# **User Roles and Application Permissions**

Users can work with applications and cubes according to their assigned roles and permissions. Roles and permissions help you manage the business activities users are permitted to perform within an Essbase instance, and the application data that they can access.

User roles are incremental; access granted to lower-level roles is inherited by higher-level roles. For example, Service Administrators, in addition to the access that only they have, inherit the access granted to Power User and User roles. You assign user roles in the Security page (available only to Service Administrators).

Table 5-1 User Roles

User Role	Description
Service Administrator	Full access to administer users, applications, and cubes.
Power User	Ability to create and delete applications and cubes that were created by this user. Ability to be granted access to, and to perform, some administrative tasks in applications and cubes created by others and provisioned to this user.
User	Ability to access any provisioned application, or a cube that has a minimum access permission. This user role has no access to administrative tasks in applications or cubes.

Users can access most Essbase features and functionality only after being assigned an application permission in addition to their user role. Application permissions determine more than simply which users and groups can see an application or cube. They also determine whether the user can view data, update data, or manage the cube or application.

Application permissions can be assigned to users and groups using the Permissions tab within the application inspector (available to Service Administrators, application managers, and some power users).

**Table 5-2 Application Permissions** 

Application Permission	Description
Application Manager	Ability to create, delete, and modify cubes and application settings within the assigned application; assign users to an application; create and delete scenarios, and give permission to run calculation scripts.
Database Manager	Ability to manage cubes, cube elements, locks, and sessions within the assigned application; create and delete scenarios, execute calculation scripts, and assign permissions to run calculation scripts.
Database Update	Ability to read and update data values based on assigned scope. Ability to create and delete scenarios. The permission to execute calculation scripts necessitates write access; however, filters may be assigned with None or Read permission to block access to certain cells.
Database Access	Ability to access scenarios, read data values in all cells, and access specific data and metadata, unless further overridden by filters. Can update values in specific cells, if granted write access to those cells through filters.

# **Provision Application Permissions**

If you're a Service Administrator or Power User, you can provision application access permissions, which are incremental. Upper-level permissions include the privileges of lower-level permissions.

Users can have a unique permission for each application or cube. The permissions, from least privileged to highest, are:



- Database Access
- Database Update
- Database Manager
- Application Manager
- 1. In the Essbase web interface, on the Applications page, select an application row, and then in the **Actions** menu, select **Inspect**.
- 2. On the **Permission** tab, use the + to open a menu for selecting users or groups to provision for access to the application.
- 3. Use the radio buttons to select the appropriate role(s) for the relevant users and groups.
- 4. Click Close.

