

Oracle® Essbase

Essbase Stack Deployment on Oracle Cloud Infrastructure



F17139-34
December 2025



Oracle Essbase Essbase Stack Deployment on Oracle Cloud Infrastructure,

F17139-34

Copyright © 2019, 2025, 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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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, MySQL, and NetSuite 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

1 Get Started with Oracle Essbase Setup and Administration

About Oracle Essbase	1
About Components and Terminology	1
Differences Between Essbase Deployment Options	4
Differences Between Essbase 11g and Oracle Essbase 21c	5
Administrator Access Requirements	11
Typical Workflow for Administrators	12

2 Set Up Oracle Essbase

Before You Begin with Oracle Essbase	1
Changes to Cloud Identity for Essbase	5
Create Dynamic Groups	6
Set Up Policies	6
Set Up Cloud Identity Access for Essbase	9
Plan a Federated Partition Environment	10
Set Relational Database Connectivity	10
Create a Confidential Identity Application	11
Create a Vault and Secrets, and Encrypt Values	15
Supported Compute Shapes	17
Deploy Essbase	18
Create Stack	18
Upgrade Stack	24
Before Upgrade of Stack	26
Collect Source Instance Metadata	28
After Upgrade of Stack	29
Complete Post-Deployment Tasks	29
Secure Your Network	30
Modify the Confidential Identity Application	30
Set Up SSL Certificates	31
Test Connectivity to Essbase	33
Link Essbase Instance to Autonomous Database Data Studio	35
Complete System and Security Hardening and Cleanup Tasks	36
Encryption at Rest for Essbase Applications	37

Set Up Multiple Virtual Hosts	38
Troubleshoot Deployment Errors	40

3 Manage the Oracle Essbase Stack on Oracle Cloud Infrastructure

Start, Stop, and Destroy an Essbase Stack Instance	1
Prepare to Work with an Essbase Stack Instance	1
Start the Essbase Stack	2
Stop the Essbase Stack	4
Destroy the Essbase Stack	4
Access Oracle Essbase Using SSH	6
Use Commands to Start, Stop, and View Status of Processes	9
Restart the Essbase Compute Instance	10
Resize Block Storage Volumes	10
Patch and Roll Back	12
Patch and Roll Back - for Version 19.3.0.2.3 and Earlier	12
Patch and Roll Back - for Version 19.3.0.3.4 and Later	13
Monitor and Diagnose Essbase Operations	15
Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service	15
Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service	16
Collect Diagnostic Information on the Essbase Node	16
Access the WebLogic Console	17
Reset or Update Admin Password	19
Fix Expired RCU Passwords that Use Autonomous Database as Repository	23
Manage Essbase Temporary Files	25

4 Migrate Essbase Applications

About Migration Tools and Use Cases	1
Migrate From Essbase 11g On-Premise	2
Prepare to Migrate from Essbase 11g On-Premise	2
Migrated Essbase 11g Artifacts	7
Convert Non-Unicode Aggregate Storage Application to Unicode Mode	9
Migrate Essbase 11g Users and Groups	14
Export 11g Users and Groups to Essbase 21c Configured with IAM or IDCS	18
Migrate an Essbase 11g On-Premises Application	22
11g LCM Export Utility Options	24
Migrate Essbase to Latest Cloud Version	25
Prepare to Migrate Essbase Applications and Users	26
Migrated 21c Artifacts	27
Migrate Applications Using Command Line Interface	29
Migrate Applications Using Migration Utility	29

Migrate from FCCS or PBCS	32
Post-Migration and Advanced Topics	33
Selective Export of Artifacts	33
Selective and Ordered Import of Artifacts	34
Test the Migrated Essbase Instance	36
Post Upgrade Tasks for CLI	36
Upgrade Aggregate Storage Outline Version	36
Export Essbase 11g On-Premise Cubes	40
Download the 11g Excel Export Utility	40
Review Member Names Before you Import an Application Workbook Created by the 11g Excel Export Utility	41

5 Manage Users and Roles

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

6 Back Up and Restore Essbase

About Backup and Restore	1
Back Up and Restore Applications	1
Back Up Cube Files Using LCM	3
Restore Cube Files Using LCM	3
Back Up and Restore an Essbase Instance	3
Back Up and Restore: Version 21.1 and Prior	4
Oracle-Scripted Backup and Restore	5
Non-Oracle-Scripted Backup and Restore	10
Back Up and Restore: Version 21.2 and Later	15
Install and Configure Oracle Instant Client and Tools	22
Required Policies	23

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 <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

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](#)
- [Differences Between Essbase Deployment Options](#)
- [Differences Between Essbase 11g and Oracle Essbase 21c](#)
- [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 operation within your business or government 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

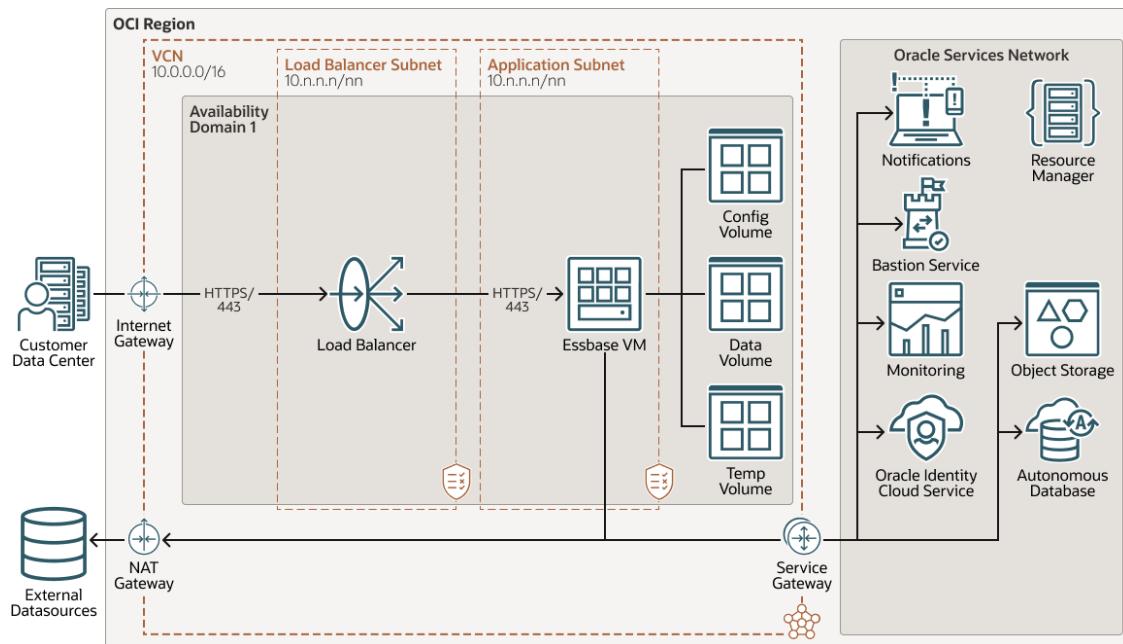
Learn about the Oracle Cloud Infrastructure components and terminology related to your setup and configuration of Oracle Essbase.

Essbase Topology within OCI

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. Sample CIDRs are used in the diagram for illustrative purposes.

Note

As an additional resource, see an Interactive Architecture Diagram (IAD) version of the diagram below, at [Oracle Essbase 21c Technical Architecture](#). The diagram contains clickable links to display additional, related documentation content.



Essbase Topology Components

- **Virtual Cloud Network and Subnet Components in OCI Region**
 - **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. See [Networking Components Overview](#)
 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 Subnet:** 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.
- **Gateways in OCI Region:** These are among the optional virtual routers you can add to your VCN.
 - **NAT Gateway:** NAT Gateway along with associated subnets and partitions, provides cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections. 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.
 - **Service Gateway:** This gateway provides a path for private network traffic between your VCN and supported services in the Oracle Services Network
 - **Internet Gateway:** This gateway provides direct internet access

- **Oracle Services Network Components in OCI Region**
 - **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.
 - **Bastion Service:** (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](#).
- **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.
- **Resource Manager:** Oracle Cloud Infrastructure Resource Manager provisions resources on Oracle Cloud Infrastructure for Essbase setup and configuration. See [Overview of Resource Manager](#) in Oracle Cloud Infrastructure documentation.

Additional Terminology Used with Essbase Deployment

- **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 Cloud Infrastructure Resource Manager to provision resources on Oracle Cloud Infrastructure, and then configures Essbase. See [Overview of Marketplace](#) in Oracle Cloud Infrastructure documentation.
- **Stack:** A stack is a collection of related cloud resources provisioned by Oracle Cloud Infrastructure Resource Manager. The stack includes an Oracle Autonomous Database 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:
 - * Oracle Autonomous Database using either the Autonomous Transaction Processing or Data Warehouse workload types. 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.
- **Vault**:

 **Note**

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 a Vault and Secrets, and Encrypt Values](#).

- **Node Manager**: This Java utility 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**: This 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**: This includes host business applications, application components, Web services, and their associated resources.

Differences Between Essbase Deployment Options

Review this topic to learn the differences between Essbase 21c deployment options.

For an Essbase 21c independent deployment, you install and configure Essbase using installation and configuration tools available on Oracle Software Delivery Cloud.

If you select to use Essbase 21c deployment on Oracle Cloud Infrastructure, you do not need to run the installation and configuration tools. The deployment process sets up Essbase on your Oracle Cloud Infrastructure (OCI) tenancy. Access the deployment stack listings from Oracle Cloud Marketplace.

Feature or Component	Independent Deployment	Stack Deployment on OCI
Integration with EPM System Foundation Services	Yes	No
Built-in integration with Identity Cloud Service	No	Yes
Support for failover configuration	Yes	No
Essbase Administration Services (EAS Lite)	Yes	No
Support for federated partitions to Autonomous AI Lakehouse	No	Yes
Support on Windows	Yes	No
Support for centralized Smart View URL for multiple Essbase instances	Yes	No
Support for Smart View for Office (Mac and Browser).	No	Yes
Encrypted applications	No	Yes
Support for Ask Essbase, a documentation-based chat interface.	No	Yes

Differences Between Essbase 11g and Oracle Essbase 21c

To understand the latest Oracle Essbase platform, review these differences in features and functionality from Essbase 11g On-Premise to Oracle Essbase 21c.

For differences in Essbase deployment types, see [Differences Between Essbase Deployment Options](#).

Architecture/Core Engine

Feature or Component	Essbase 11g	Essbase 21c	Notes
Application Architecture	EPM System, Foundation Services	Essbase runs on a middle-tier WebLogic platform, either in Fusion Middleware or Oracle Cloud Infrastructure.	In independent deployments, Fusion Middleware enables fast performance, optimized memory usage, high concurrency, and failover. In deployments on Oracle Cloud Infrastructure, you do not have to run Essbase install/configuration tools. OCI enables fast performance, optimized memory usage, and flexible deployment.
Relational Database for Essbase metadata	No	Yes	Repository Creation Utility (RCU) schemas hold information about Essbase platform artifacts and components. This metadata is stored in a supported relational database (RDBMS) of your choosing. Note that your Essbase applications and cubes aren't stored in these schemas. For independent deployments, applications are in your selected <Application Directory> location on the server where you install Essbase. For deployments on Oracle Cloud Infrastructure, applications are in the data block volume on your Essbase compute instance.
Query Engine Options (aggregate storage, block storage, hybrid mode)	Block storage default	Hybrid mode default	Hybrid mode is the default for block storage cubes, providing robust dependency analysis, fast aggregation, ability to process more calculations, and several tuning options. See Adopt Hybrid Mode for Fast Analytic Processing .

Feature or Component	Essbase 11g	Essbase 21c	Notes
Failover	Yes	Yes (independent deployment on Linux only)	Failover is integrated with WebLogic to support a central request leasing system that determines which node is active and which nodes are waiting on standby.
Unicode	Yes	Yes	Essbase 21c uses UTF-8 encoding. Convert all native-encoded Essbase 11g On-Premise applications to Unicode before running the LCM export operation.

Interfaces/Tools

Feature or Component	Essbase 11g	Essbase 21c	Notes
Administrative Essbase Web Interface	No	Yes	The Essbase web interface enables you to manage applications, users/groups, and Essbase artifacts. It includes a rich outline editor, scripting editors, a data analysis interface where you can save grid layouts, and a load rules editor with built-in data previews. A centralized Jobs interface lets you initiate requests, and monitor active and recent requests. Cube Designer and Smart View, as well as utilities for migration, automation, and administration, are available to download from the Console.
Cube Designer, Application Workbooks	No	Yes	The Cube Designer extension for Microsoft Excel is a client interface for designing and building Essbase cubes from application workbooks. This interface offers a flexible and portable cube design and administration system. Structured workbooks simplify everyday cube design, optimization, and portability. Cube Designer infers patterns found in unstructured workbooks, to help you shape raw data into hierarchically organized cubes. Cube Designer and application workbooks also offer the benefit of offline cube development and the ease of iterative builds. See Work with Cubes in Cube Designer.
Migration, Backup/Restore	Yes	Yes	Essbase makes it convenient to migrate applications across Essbase releases and host servers, using a choice of utilities via the Console, depending on the migration path; see About Migration Tools and Uses Cases. You must also maintain regular backups of the Essbase RCU schemas stored in a relational database, as well as user roles. Consult the back up and restore instructions for your deployment type.
Catalog	No	Yes	The Catalog is a central place to store files and artifacts associated with Essbase applications and users. It includes user and shared directories, and an instructive Gallery of sample cubes.
Gallery	No	Yes	Included in the Catalog is a Gallery of cube templates, in the form of Excel application workbooks. Import these workbooks to build a diverse variety of sample cubes. The samples are instructional for learning about different use cases for Essbase applications and features, as well as learning how to build and design cubes from structured and unstructured workbooks.
Smart View	Yes	Yes	See Working with Oracle Smart View for Office.
Provider Services	Yes	Yes	Provider Services is built in to Essbase 21c. Provider Services requests time out by default after 10 minutes. Configuration options are available in the Console. You must update client service URLs to the new format, as described in migration instructions.

Feature or Component	Essbase 11g	Essbase 21c	Notes
Essbase Administration Services	Yes	Yes (Independent deployment only)	EAS Lite is an option (with independent deployment) for continued management of your applications, in case your company is not ready to adopt the web interface. The features and functionality of EAS are limited to what was available in Release 11g and do not encompass the modern platform features.
Essbase Studio	Yes	No	Use connections and Datasources to connect to your outside sources of data, use the improved Rules editor to preview and shape imports to your cube, and use Drill Through reports to access data you won't import.

Security/Authentication

Feature or Component	Essbase 11g	Essbase 21c	Notes
Communication/Network Security	Security file	TLS, load balancer	All data is encrypted in transit layer using Transport Layer Security (TLS). You can optionally implement a load balancer. On independent deployments, this could be Oracle HTTP Server or another load balancer you choose. On Oracle Cloud Infrastructure deployments, you can select to configure a load balancer through OCI.
Authentication Method	EPM Shared Services or native	Oracle Identity Cloud Service (OCI deployment only) EPM Shared Services (independent deployment) WebLogic Embedded LDAP	On independent deployments, for user and group authentication, choose WebLogic Embedded LDAP, or if you already use EPM Shared Services, you can continue to use it with Essbase 21c (independent deployment). If you select WebLogic, it is strongly recommended to federate users to an external authentication provider, such as Microsoft Active Directory, which is suitable for production environments. On OCI deployments, (including when using MSAD) integrate with Oracle Identity Cloud Service (IDCS).
User/Group Roles and Application Permissions	Essbase roles, Shared Services roles, and application roles	Simplified: 3 roles, 4 permissions	Essbase user roles are User, Power User, and Service Administrator. Application permissions, granted separately, are Application Manager, Database Manager, Database Update, and Database Access. All roles from EPM Shared Services can be mapped to the new roles and permissions, or, you can continue to use EPM Shared Services (independent deployment only).
Filters	Yes	Yes	Filters help you implement fine-tuned, cell-level access controls to your cubes. Using dynamic filters with built in functions/variables, you can make filters extensible and adaptable to a changing user base and real-time source data. You can use LoginAs to test the filters in the Essbase web interface. In Essbase 11g On-Premise, only one filter can be granted per user per cube. In Essbase 21c, new filter assignments are combined with existing filter assignments.
Security file	Yes	No	There is no need for essbase.sec in Essbase 21c.

Data Connectivity/Interoperability

Feature or Component	Essbase 11g	Essbase 21c	Notes
Connections and Datasources	No	Yes	Essbase administration tasks often require connectivity to remote source data or hosts. With reusable connections and Datasources, you no longer have to code the connection details into artifacts like rule files or filters, or enter them each time you perform other connection-dependent tasks.
Network Connectivity	Yes	Yes	Essbase APIs use TLS/SSL for secure connectivity both internally between components and externally with other applications. You can connect from any software using Essbase Runtime Client (RTC) over secure HTTP without needing to open additional TCP/IP ports to enable client connectivity.
Partitions	Yes	Yes	Linked partitions aren't supported in Essbase 21c. Use @XREF and @XWRITE to analyze data across cubes. See Link Cubes Using Partitions and XREF/XWRITE . The administrator must set up user security for the source and target cubes. See Security for Partitioned Databases . For OCI deployments, federated partitions are available with Autonomous AI Lakehouse, enabling transparent integration of relational data and the Essbase cube.
Drill Through	Yes	Yes	When you need more data than what is in the cube, use Drill Through reports to access external data sources. Performance is improved for drill through connections to Oracle Database. The flexibility of Drill Through report design is improved, allowing diversified selection of multiple cells or ranges of cells. Selections can be recursive, non-recursive, level 0, contiguous, or non-contiguous.

Calculation/Data Flow

Feature or Component	Essbase 11g	Essbase 21c	Notes
Calculation Scripts (block storage and hybrid mode)	Yes	Yes	Full library of calculation functions and commands suits most analytical applications. Build your own custom defined functions and macros. Calculation tracing helps analyze and debug calc script performance and member formula processing. Tuple-based calculation helps optimize and refine calculation scope, limiting it to focus on the active Smart View grid. Hybrid calculation can be used, and offers tuning options.
Aggregate Storage Calculation	Yes	Yes	MDX Insert helps you perform custom calculations and allocations. You can automate the creation and maintenance of default aggregate views.
Hybrid query processor enabled by default for block storage cubes	No	Yes	Essbase 21c processes dynamic dependencies in queries using hybrid mode. In hybrid mode, Essbase evaluates formula dependencies prior to resolving queries, ensuring minimal processing time and accurate results. Hybrid mode is not the default mode for calculation scripts, but you can enable it using the HYBRIDBSOINCALCSCRIPT configuration setting. In Essbase 11g, the default query engine is not hybrid mode, but you can enable it using ASODYNAMICAGGINBSO.

Note

Do not use two-pass calculation with hybrid mode cubes. Only use solve order.

Feature or Component	Essbase 11g	Essbase 21c	Notes
Parallel Calculation	Yes	Yes	See Using Parallel Calculation.
Data Load/Dimension Build Rules	Yes	Yes	Load rules editor has built-in data previews. You can import from the Catalog or from outside sources. Rule file columns can employ functions like Sum, Min, Max, Count, and Avg, to help you shape your import. Performance is improved for SQL-based loading. Batch Outline Editing is available from Java or REST API. Command-line interface (CLI) supports streaming data load from a variety of sources. Aggregate storage data load optimizations include buffer, merge, and cache tuning options. You can migrate rule files from Essbase Studio and edit them in the Essbase web interface.
Custom-Defined Calculation Functions and Macros	Yes	Yes	Build your own custom-defined functions and macros in Java.
MDX	Yes	Yes	In addition to MDX's well-known utility as a multidimensional query language, you can use its Insert and Export directives to shape, copy, move, and update any custom slice of multidimensional data. Sub Select lets you filter the volume of queried data.
Two pass calc	Yes	Not recommended	Do not use two-pass calculation with hybrid mode cubes. Only use solve order.
Scenario Management	No	Yes	Scenario management offers the ability to build private work areas or "sandboxes" in which users can model different assumptions within the data to see the effect on aggregated results, without affecting the cube.

Development, Automation, and Audit

Feature or Component	Essbase 11g	Essbase 21c	Notes
Automation and Developer Tools	Yes	Yes	REST API helps you automate management of Essbase resources over secured HTTP. Java API, Command Line Interface (CLI), MaxL administrative language, and Report Writer are also available.
Accelerated Development and Audit Capabilities	No	Yes	Calculation tracing lets you monitor and debug calculation scripts. Query tracing can be used to monitor and debug query performance. Audit trail enables you to track changes made to data. Solve order can be adjusted while you're working in Smart View.
Shadow Applications	No	Yes	To perform cube modifications and restructures with limited down time, you can create a shadow application that is a copy of the primary application. The primary application continues to serve read-only operations, such as queries, while you perform modifications on the shadow application. You can make the shadow application visible or hidden. Available in REST API – see the Create Shadow Application endpoint.

Configuration

Feature or Component	Essbase 11g	Essbase 21c	Notes
Configuration Tool	No	Yes (independent deployment only)	For independent deployments, you can configure your Essbase environment any time using the integrated configuration tool, either at the end of your Essbase installation, or by launching the Configuration Tool later, after installation. Silent mode is available if you're configuring frequently (for example, in a development environment).

Feature or Component	Essbase 11g	Essbase 21c	Notes
Configuration Settings	Yes	Yes	Most configuration parameters you need for application tuning should be set per application, using the Essbase web interface. Server-level configuration is handled by the configuration tool (for independent deployments), but you can also change server configuration defaults using <code>essbase.cfg</code> , if needed. Provider Services configuration options for network-related preferences are available in the Console.

Other Notable and Legacy Features

Feature or Component	Essbase 11g	Essbase 21c	Notes
Expanded Analysis Methods	No	Yes	You can perform ad hoc data queries/grid analyses on cube data from the administrative Essbase web interface, as a built-in alternative to connecting via Smart View. You can save your grid layouts, run report scripts, and run and save named MDX queries. See Analyze Data in the Web Interface .
Member ID	No	Yes	A unique Member ID is automatically assigned to every member in the cube outline.
Logging	Yes	Yes	Application logs are in Oracle Diagnostic Logging (ODL) format. You can download them from the Essbase web interface. You can use Performance Analyzer to analyze Essbase logs to generate usage and performance statistics.
Implied Sharing	True by default	You choose	In Essbase 11g On-Premise, the implied share setting could be changed for an application using the <code>IMPLIED_SHARE</code> setting. When you migrate an Essbase 11g On-Premise application to Essbase 21c, your <code>IMPLIED_SHARE</code> setting is preserved. For each application, you set implied share behavior only once, at application creation time (it's not meant to be changed). See <code>IMPLIED_SHARE_ON_CREATE</code> configuration property, which is <code>FALSE</code> by default (unless you migrate an application that has <code>IMPLIED_SHARE</code> set to <code>TRUE</code>).
Report Scripts	Yes	Yes	See Report on Data .
Currency Conversion	Yes	Yes	See Designing and Building Currency Conversion Applications .
ESSCMD	Yes	Yes	ESSCMD is still supported, but is not updated with the latest features. Try CLI and MaxL .
Varying Attributes	Yes	Only in EAS Lite (for independent deployments)	-
Hybrid Analysis	Yes	No	Not applicable in Essbase 21c. Use connections and Datasources to connect to your outside sources of data, use the improved Rules editor to preview and shape imports to your cube, and use Drill Through reports to access data you won't import.
Incremental Restructuring	Yes	No	Restructuring is not as time-consuming in Essbase 21c. If needed, you can work in a shadow application to limit down-time.
Dynamic Calc and Store	Yes	No	Dynamic Calc and Store members are treated as Dynamic Calc.
Direct I/O, cache memory locking	Yes	No	Not applicable

Feature or Component	Essbase 11g	Essbase 21c	Notes
Disk volumes	Yes	Yes	Set up disk volumes using standard operating system mounts, instead of DISKVOLUMES configuration property.
Repair Invalid Block Headers	Yes	No	Not applicable
Data Compression Types: Zlib, None	Yes	No	Not applicable
Isolation modes	Yes	No	Essbase 21c manages block storage data transactions in uncommitted mode. Essbase releases a block after it is updated, and commits blocks when the transaction is completed.

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.
IAM or IDCS system administrator	The OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) identity domain administrator. This administrator is responsible for setting up Essbase cloud identity access, including creating and modifying a confidential identity application.
IAM or IDCS Essbase admin user	An OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) 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	<p>Log into Oracle Cloud Infrastructure and select Essbase in Oracle Cloud Marketplace.</p> <p>Enter the required metadata, and select the options that you prefer in Oracle Resource Manager setup wizard.</p>	Set up Oracle Essbase
Perform post-deployment 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 Essbase Applications
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 and restore	Perform 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 OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS).

Note

For important changes about identity workflows, refer to [Changes to Cloud Identity for Essbase](#).

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 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

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

Prerequisite Metadata	Links to overviews and tasks / examples	Record values needed for deployment
IAM or IDCS system administrator user ID (or create it later with REST API for IAM or IDCS)	Administrator Access Requirements	IAM or IDCS system admin user ID
IAM or IDCS Essbase administrator user name , as defined during stack creation	Administrator Access Requirements	Initial Essbase administrator name
Sufficient quota in target region and target availability domain	Regions and Availability Domains	Region Availability domain
Vault secret and OCIDs for secrets created to encrypt your Essbase administrator password, Database system administrator password, and IAM or IDCS application client secret, using Oracle Cloud Infrastructure Vault provisioning encryption	Create a Vault and Secrets, and Encrypt Values , Also see, Overview of Vault in Oracle Cloud Infrastructure documentation.	Vault encryption key OCID
[Optional] Load balancer to specify shape and subnets	Overview of Load Balancing	Load balancer shape and subnets
Compute shape for Essbase compute instance	Supported Compute Shapes	Compute shape
[Optional] Use existing network setup option - if used, you must configure network infrastructure with security lists, or ask Oracle Support to create an instance	Creating a Virtual Cloud Network	Existing virtual cloud network Name
Essbase administrator user name , as defined during stack creation.	Administrator Access Requirements . Also see Create Stack step on entering Essbase System Admin User Name.	Also known as Essbase 911 user name (administrator who manages the WebLogic server on which Essbase runs)
Essbase Node Public IP	-	If creating a private Essbase subnet, record the Essbase node private IP and the VCN on which it is deployed.
Essbase URL for Essbase web interface and confidential application use	-	Essbase URL
[Optional] Enable Monitoring service	Monitor Operations and Resources Using Oracle Cloud Infrastructure Monitoring Service	-
[Optional] Enable and set up Notifications service	Get Event Notifications Using Oracle Cloud Infrastructure Notifications Service	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
1. Select one of these database options: <ul style="list-style-type: none"> • New Oracle Autonomous Database • Existing Oracle Autonomous Database Serverless or Oracle Autonomous Database on Dedicated Exadata Infrastructure that you deployed from Oracle Cloud Infrastructure Console • Existing Database System that you deployed from Oracle Cloud Infrastructure 	Creating an Autonomous Database Creating a Database System How to Use Network Components Set up Rules for Database Connectivity Set Relational Database Connectivity	<ul style="list-style-type: none"> • For Oracle-created database: admin name and password • For existing Autonomous database: admin name and password, and compartment • For existing Database System: admin name and password, database name and home, and compartment
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: <ul style="list-style-type: none"> • Essbase administrator password • Database system administrator password • IAM or IDCS application client secret 	Create a Vault and Secrets, and Encrypt Values	- Vault encryption key OCID. The OCIDs need to be noted for Essbase administrator password, Database system administrator password, and IAM or IDCS 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 & Security>Groups>Administrator>Create User), close browser, change password using email you receive, and log in as new admin.	Administrator Access Requirements	-
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 & Security>Compartments>Create Compartment) where you want to deploy Essbase.	Choosing a Compartment	Compartment ID (OCID) and name

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

Prerequisite Tasks	Links to overviews and tasks / examples	Record values needed for deployment
6. Create dynamic group - In OCI console, create a dynamic group to allow 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 Common Policies How Policies Work.	Policy statements (enter in text worklist doc for entry convenience)
8. Set up Essbase access - In IAM or IDCS, set up Essbase access.	Set Up Cloud Identity Access for Essbase	-
9. Set up confidential application to register Essbase - In IAM or IDCS, for each compartment in which you plan to deploy Essbase, create and activate a confidential application.	Create a Confidential Identity Application	Confidential application name IAM or IDCS Instance GUID IAM or IDCS Application Client ID IAM or IDCS Application Client Secret
11. Complete post-deployment tasks	Complete Post-Deployment Tasks	-

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 Oracle Cloud Infrastructure 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:
 IAM or 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:
 IAM or IDCS instance GUID:

IAM or IDCS application client ID:
IAM or IDCS application client secret:
IAM or 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):

Changes to Cloud Identity for Essbase

When you deploy Essbase on Oracle Cloud Infrastructure using the Essbase Marketplace listing, cloud identity may be managed by OCI Identity and Access Management (IAM) or by Oracle Identity Cloud Service (IDCS), depending on your tenancy.

If you are not sure which identity management is used on your tenancy, refer to Documentation to Use for Cloud Identity.

If you already use Essbase on OCI, and your tenancy has been updated to use IAM identity domains, you don't need to create a new confidential application. Your identity application is migrated as an integrated application within an identity domain named OracleIdentityCloudService.

If you are starting a new cloud deployment on an OCI tenancy with IAM, Oracle recommends creating a unique identity domain for Essbase instead of using the default identity domain. Create the dynamic groups, the initial Essbase administrator user, and the confidential identity application all within this identity domain.

Note

If you originally deployed Essbase using IDCS before the tenancy migrated to an IAM domain, you may need to update the policies. This may apply if you are performing a new deployment and your group or dynamic group is in an identity domain other than Default.

For licensing information, refer to [IAM Identity Domain Object Limits](#).

To configure clients to be able to access the signing certificate for the identity domain, go to the **Settings** for the identity domain. Under **Domain settings - Locale**, click **Edit domain settings**. Then, under **Access signing certificate**, toggle on **Configure client access**, and click **Save changes**.

Policies are global to the tenancy (not domain specific), but the syntax for specifying group names or dynamic group names in policy statements may need to be updated if your group or dynamic group is in an identity domain other than Default. Review the Oracle Cloud Infrastructure documentation on policy syntax: Create an IAM Policy in an Identity Domain.

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).

The following steps are for OCI tenancies that use Oracle Identity Cloud Service (IDCS) for identity management.

If you are starting a new cloud deployment on an OCI tenancy with IAM instead of IDCS, Oracle recommends creating a unique identity domain for Essbase instead of using the default identity domain. Create the dynamic group within this domain. In this dynamic group, set up the same rules as described below for IDCS.

If you are not sure which identity management is used on your tenancy, refer to Documentation to Use for Cloud Identity.

1. On the Oracle Cloud Infrastructure console, navigate to **Identity & Security**, and click **Domains** under **Identity**.
2. Select the identity domain where you want to create the dynamic group, and navigate to **Dynamic groups**.
3. Click **Create dynamic group**.
4. Provide a **Name** and **Description** (optional) for the dynamic group, and record the name for future use.
5. In the **Matching rules** section, click **Rule builder** to define the resources for the dynamic group.
6. From the **Include instances that match** dropdown menu:
 - Select **Any of the following** to allow the dynamic group to have more relaxed/broad access. A compute instance that matches any of the statements you add to the rule will be included in the group's resources.
 - Select **All of the following** to limit the dynamic group's resources more strictly. Only those compute instances that match *all* of the statements in the rule will be included.
7. From the **Match instances with** dropdown menu, select **Compartment OCID**. For the corresponding **Value** field, paste the OCID of the compartment in which you're creating the Essbase stack. With this option, any compute instances in the named compartment will work using this dynamic group. The other option, **Instance OCID**, specifies matching just one compute instance ID.
8. Click **Add rule**.
9. You can enter tags (optional) to organize and track resources in your tenancy.
10. Click **Create** to complete the dynamic group creation.

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 Vault

To create policies:

1. On the Oracle Cloud Infrastructure console, navigate to **Identity & Security**, click **Policies** under **Zero Trust Packet Routing**, select the compartment created for Essbase, and then click **Create policy**.
2. Provide a **Name** and **Description** for the policy.
3. Click **Add policy statements** to add an "Allow" policy statement for each instance in the compartment. Copy the statements from your text worklist file and specify the group_name in each statement.
4. Choose a policy builder type **Policy template builder**, **Simple policy builder**, or **Manual policy builder** based on how you want to build your policies. Configure the policy details as required, and then click **Add** to save the policy.
5. When done, click **Create policy** to complete the process.

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

 ⓘ Note

If a group or dynamic group is in an identity domain other than Default, you must qualify references to it in policies using syntax that identifies the identity domain. For example,

```
Allow group domain-name/group_name to verb resource-type in
compartment compartment-name
```

```
Allow dynamic-group domain-name/group_name to verb resource-type in
compartment compartment-name
```

Review the Oracle Cloud Infrastructure documentation on policy syntax: [Create an IAM Policy in an Identity Domain](#).

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 domain-name/group_name to use ons-topic in compartment
dev where request.permission='ONS_TOPIC_PUBLISH'
```

- Oracle Cloud Infrastructure Monitoring integration:

```
Allow dynamic-group domain-name/group_name to use metrics in compartment
dev where target.metrics.namespace='oracle_essbase'
```

- Encryption at rest for Essbase applications:

```
Allow dynamic-group domain-name/group_name to use vaults in compartment compartment_name  
Allow dynamic-group domain-name/group_name to use keys in compartment compartment_name
```

Set Up Cloud Identity Access for Essbase

To set up cloud identity access, you integrate Essbase with OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS). You provision Essbase users using Essbase roles, rather than IAM or IDCS roles.

To prepare security access for Essbase, you must log in to IAM or IDCS 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 IAM or IDCS 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 IAM or IDCS before deploying the Essbase stack.

- [IAM](#)
- [IDCS](#)

IAM

1. Log in to the Oracle Cloud Infrastructure Console. Select **Identity & Security**.
2. Click **Domains**.
3. Select the identity domain that exists for the Essbase stack to use.
4. Within the identity domain, go to **Users management**, and click **Create** to add a user who will act as the initial Essbase Service Administrator, if one does not already exist.
5. To configure clients to be able to access the signing certificate for the identity domain, go to the **Settings** for the identity domain. Under **Domain settings - Locale**, click **Edit domain settings**. Then, under **Access signing certificate**, toggle on **Configure client access**, and click **Save changes**.
6. Create a confidential identity application, as described in [Create a Confidential Identity Application](#).

IDCS

1. 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. Scroll 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.
6. Create a confidential identity application, as described in [Create a Confidential Identity Application](#).

About Single Sign-On (SSO)

If you use single sign-on (SSO) with IAM or IDCS, your Essbase login screen routes to IAM or IDCS.

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

Plan a Federated Partition Environment

To combine the benefits of Essbase and Autonomous AI Lakehouse, consider implementing a federated cube. Here is how to plan the environment and your deployment strategy.

1. Read the federated partition prerequisites before you use the Marketplace listing to deploy Essbase as a stack on your Oracle Cloud Infrastructure (OCI) tenancy: [Prerequisites for Federated Cubes](#)
2. As is always advised when deploying Essbase on OCI via Marketplace, review the deployment instructions in [Deploy Essbase](#). But before you deploy or upgrade the Essbase stack, refer also to these deployment instructions specifically for the federated partition environment: [Deploy Essbase from Marketplace for Federated Cubes](#).
3. Size your Autonomous AI Lakehouse to be able to handle being used as a repository for both the RCU schema as well as your user schema where the fact table resides. You shouldn't use any of the RCU schemas as the fact table. Refer also to [Provision Autonomous AI Lakehouse for Federated Cubes](#).
4. After you are finished creating the federated partition, one or more individuals must configure DBMS_CLOUD credentials to allow connectivity from Essbase to Autonomous AI Lakehouse. Refer to [Set up Credentials for Federated Partitions](#) for more information.

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.

Essbase integration with Autonomous AI Lakehouse via federated partitions is enabled out-of-the-box when you use the Marketplace listing to deploy Essbase as a stack on your Oracle Cloud Infrastructure (OCI) tenancy.

While working with federated partitions, you should size your Autonomous AI Lakehouse to be able to handle both the RCU schema as well as your user schema where the fact table resides. You shouldn't use any of the RCU schema for the fact table. Refer also to Provision Autonomous Data Warehouse for Federated Partitions.

Create a Confidential Identity Application

Before deploying the Essbase stack, create a confidential application in OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS), and register Essbase with it.

If you are not sure which identity management is used on your tenancy (IAM or IDCS), refer to Documentation to Use for Cloud Identity.

- [IAM](#)
- [IDCS](#)

IAM

1. Open the Oracle Cloud Infrastructure Console. Select **Identity & Security**.
2. Click **Domains** and select your respective domain.

 **Note**

Ensure that an identity domain exists for the Essbase stack to use.

If you already deploy Essbase on OCI, and your tenancy has been updated to use IAM identity domains, you don't need to create a new confidential application. Your identity application is migrated into an IAM identity domain.

If you are starting a new cloud deployment on an OCI tenancy with IAM, Oracle recommends creating a unique identity domain for Essbase, instead of using the default identity domain. Then create the [dynamic groups](#) and the confidential identity application within the identity domain.

For licensing information, refer to [IAM Identity Domain Object Limits](#).

3. Assuming you have an identity domain for Essbase, click **Integrated applications** within that domain.

① Note

If you do not already have a confidential application for Essbase, click **Add application**, select **Confidential Application**, and click **Launch workflow**.

Enter a name for the application, and click **Submit**.

4. Select the **Confidential Application** you launched to configure the **Client configuration**.
5. Click **OAuth configuration**, then click **Edit OAuth configuration** to open the setup window for **Resource server configurations** and **Client configurations**.
6. In the **Client configuration** tile, select **Configure this application as a client now**.
7. After completing the following required fields in the **Authorization** section, proceed to the **Token issuance policy** section.
 - Select the following allowed grant types: **Client Credentials** and **Authorization Code**.
 - If you don't plan to provision a load balancer, select the **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`
8. Under **Token issuance policy**, add the following application roles: **Identity Domain Administrator**, **Security Administrator**, and **Cloud Gate**.
 - a. Click the **Add app roles** toggle switch to enable application roles.
 - b. Click **Add app roles** and select **Identity Domain Administrator**, **Security Administrator**, and **Cloud Gate**, then click **Add** to assign the selected roles.

Token issuance policy

Authorized resources

Choose which authorized resources a client application can access. Select 'All' to access those resources where an explicit association between the client and the resource exists.

All

Confidential

Add resources



Add resources if you want your application to access the APIs of other applications.

Add app roles



Add the application roles to assign to this application. For example, add the Identity Domain Administrator role if the Identity domain administrator will be available to the application.

App roles

<input type="checkbox"/> App roles
<input type="checkbox"/> Identity Domain Administrator
<input type="checkbox"/> Security Administrator

9. Click **Submit** to complete the **Confidential Application** configuration. Then, go to the **Integrated Applications** page, select **Activate** in the title bar next to your application's name and click **Activate application**.
10. From the **OAuth configuration**, record the following details, which are under **General Information**: **Client ID** and **Client secret**.

① Note

Click **show or copy** under **Show secret** to view or copy the secret.

Record these values to use during your Essbase stack deployment. Store the Client secret in the Vault (see [Create a Vault and Secrets, and Encrypt Values](#)).

11. Record the IDCS Instance GUID from the following location: in the OCI Console, navigate to **Identity & Security**, and click **Domains**. Click the domain name that contains your confidential application. In the **Details** tab, copy the **Domain URL** value. From this value, extract only the host name; for example, `idcs-123456789a123b123c12345678d123e1`. If you don't have access to this value, ask your identity domain administrator to provide it. Save this value and enter it when prompted for the **IDCS Instance GUID** during your Essbase stack deployment.
12. To configure clients to access the signing certificate for the identity domain, go to the **Settings** for the identity domain you created for Essbase. Under **Domain settings** -

Locale, click **Edit domain settings**. Then, under **Access signing certificate**, toggle on **Configure client access**, and click **Save changes**.

IDCS

1. 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 Identity Cloud Service console, in the **Applications and Services** tile, click **+Add**.
3. Select **Confidential Application**.
4. In the **App Details** step, enter a name only, and click **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.
5. In the **Client** step, select the option **Configure this application as a client now**.
6. 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`
 - If you're provisioning a load balancer, use **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`
7. Under **Token Issuance Policy**, in the section **Grant the client access to Identity Cloud Service Admin APIs**, click **Add**.
8. In the **Add App Role** dialog, find and add the following roles: **Identity Domain Administrator**, **Cloud Gate**, and **Security Administrator**.
9. 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 details: IDCS Application Client ID and IDCS Application Client Secret. Record these values to use during your Essbase deployment. Store the Client Secret in the Vault (see [Create a Vault and Secrets, and Encrypt Values](#)).
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.

Create a Vault and 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 Essbase 21c and 19.3.0.5.6, you use secrets, created with the Vault in the **Identity & Security** area of Oracle Cloud Infrastructure Console. In your Vault, you enter the password, and the latest version of it is stored as part of your key, in the Vault.

Refer to [Key Management](#) in Oracle Cloud Infrastructure documentation.

Note

If you're using an existing **Vault** and already have an encryption key, you can skip the **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 encryption key first.

The Vault and key should be in same region as the Essbase stack, if you want to use them for Essbase application encryption.

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
- IAM or IDCS application client secret
- Database system administrator password

Create a New Vault:

Note

These steps explain how to create the lower-cost Vault option. Create new entities only if they do not already exist.

1. Sign in to the Oracle Cloud Infrastructure Console.
2. In the navigation menu, select **Identity & Security**, and under the **Key Management & Secret Management** section, click **Vault**.
3. Under **Applied filters**, select your compartment for Essbase, if not already selected.
4. Click **Create Vault**.
5. Under **Create in Compartment**, ensure your Essbase compartment is selected.

6. For **Name**, enter a name, such as OracleEssbaseVault.
7. For the lower-cost option, leave the virtual private vault option unchecked.
8. Click **Create Vault**.

 **Note**

You can retrieve the Vault **Cryptographic Endpoint** value at any time by selecting the newly created vault on the **Vaults** page.

Create a New Encryption Key

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

1. Click the name of the newly created Vault; for example, OracleEssbaseVault from the previous section.
2. Under **Master Encryption Keys**, click **Create Key**.
3. Provide a name for the key, such as OracleEssbaseEncryptionKey.
4. Click **Create Key**. You will use this key during **Secrets** creation.

Create a New Secret

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

1. Click **Secrets**.
2. Click **Create Secret**.
3. Enter a unique and easily identifiable **Name** for the secret – for example, dbadminpassword – and a relevant **Description**.
4. Select the new **Encryption Key** from the dropdown menu (created in the previous section), or an existing one. For example, OracleEssbaseEncryptionKey.
5. Select either **Automatic secret generation** or **Manual secret generation** to enable the **Secret Contents** field. Then, enter the password in the field.
6. Click **Create Secret**.
7. For each created secret, take note of its name, vault, OCID, and other details, so you can easily identify it later during configuration.

 **Note**

Please ensure that each secret has a unique name, even if they are stored in different vaults.

To encrypt your Oracle Essbase Administrator password (for 19C):

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 `oci kms crypto encrypt` command using Oracle Cloud Infrastructure command line interface. Provide the following parameters:

- Vault Encryption Key OCID
- Vault Crypto Endpoint
- base64-encoded password

Example:

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

3. 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.

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.Optimized3.Flex
- VM.Standard3.Flex
- VM.Standard.E3.Flex
- VM.Standard.E4.Flex
- VM.Standard.E5.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 <https://cloud.oracle.com/compute/faq>.

Deploy Essbase

Deploy Oracle Essbase from Oracle Cloud Marketplace.

We can deploy an Essbase stack in the following ways:

- [Create Stack](#)
- [Upgrade Stack](#)

Create Stack

Create the Oracle Essbase stack, 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 AI 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.

1. Read prerequisites and requirements that you need to know or do before deployment. See [Before You Begin with Oracle Essbase](#).
2. Sign into Oracle Cloud Infrastructure console as the Oracle Cloud Infrastructure administrator.
3. From the navigation menu, select **Marketplace**.
4. On the Marketplace page,
 - a. In the **Search** field, type Essbase and search.
 - b. Select Oracle Essbase.
 - c. From the dropdown menu, select the stack **Version**, or accept the default.
 - d. From the dropdown menu, select the target **Compartment** you created for Essbase to create the stack instance.
 - e. Select the check box to accept the **Oracle standard Terms and Restrictions**.
 - f. Click **Launch Stack**.

Note

This applies to users deploying through **Resource Manager** under **Developer Services**.

For **My Configuration**, to upload the terraform configuration source files, select **Zip File** under Stack configuration (instead of the default Folder option). If necessary, drop or browse to the stack zip file. The stack name is displayed.

5. In the **Stack Information** section on the Create Stack page.
 - a. Enter the stack description and other stack information, as necessary.
 - b. Click **Next**.

① Note

Do not select the **Use Custom Terraform Providers** check box during the Essbase stack deployment process.

6. On the **Configure variables** page under **General Settings**, you configure variables for the infrastructure resources that the stack creates.
 - [Optional] Enter **Unique Resource Identifier** to identify your resources generated during stack deployment. For example, `essbase_<userid>`. Use a meaningful **Unique Resource Identifier**, as it will be used as a dimension for filter Essbase metrics related to the components in this stack. If you don't enter a value, the system generates one automatically.
 - The **Target Compartment** you created for Essbase, where the stack instance will be deployed, is displayed here.
 - If you want to create a new node (VM), based on an existing deployment using a new image, select **Upgrade Stack**. This selection updates the stack configuration page to display only the fields required for the image update. If selected, skip the following step.
7. In **Essbase Instance** - when **Upgrade Stack** was NOT selected in **General Settings**, the following options are displayed and available for your selection:
 - a. Select an **Essbase Availability Domain** in which you want to create the Essbase compute instance.
 - b. Select the **Essbase Instance Shape** for the Essbase compute instance.
If you select new Intel flex shapes **VM.Standard.E(n).Flex**, **VM.Optimized3.Flex**, or **VM.Standard3.Flex**, additional entry field, **Essbase Instance OCPUs**, is displayed. Enter the **Essbase Instance OCPUs** or accept the default. The minimum value is 4 OCPUs.
 - c. Enter the **Data Volume Size** or accept the default. The minimum value is 256GB. You can also select [Optional] **Config Volume Size**, [Optional] **Temp Volume Size**; minimum value is 64 GB.
 - d. Paste the value of the **SSH Public Key** that you want to use, 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/IAM Essbase Admin User** field later in the stack definition, in the [Identity Configuration](#) section. If you have already used **Identity Cloud Service** user name in **Essbase System Admin User Name** field, you must provide a different **Identity Cloud Service** user name in the **IDCS/IAM Essbase Admin User** field, under the [Identity Configuration](#) section.
 - f. [Optional] Enter **Essbase Instance Timezone**.
 - g. [Optional, but required for Federated Partitions].
You can change the **Catalog Storage Type** from the default (**Local Filesystem**) to **Object Storage Bucket**. However, once you change to Object Storage Bucket, you cannot change back to Local Filesystem after deployment. Default Local Filesystem path is `/u01/data/essbase/catalog`, for `/users`, `/shared`, and `/system` directory, which has datasource, wallet, and connection details.

If you choose Object Storage Bucket as catalog storage type, Essbase catalog storage is integrated with Oracle cloud storage and a new Object Storage bucket is created. This bucket can be used for storing objects, which include Essbase artifact files (such as data or rules) to be used in Essbase jobs (such as data load and dimension build). Essbase buckets are visible on **Resource Manager > Stacks > Stack Details>Application Information** page under **Storage Details**.

Under the Storage Details header, Catalog Bucket Name is found. Essbase artifacts integrate with Object Storage Bucket or external object stores, for various import/export use cases. Note that Application folder is on disk, so related artifacts (data, rules, jobs) are stored on `/u01/data/essbase/app`. Currently, only `/users` and `/shared` folders are integrated with Object Storage Bucket, so only those folder files are stored on Object Storage Bucket inside the newly created buckets, if Object Storage Bucket is opted as catalog storage during stack creation.

 **Note**

Do not create or delete folders in user directories using Bucket object or by using ssh. Perform all file catalog operations using Essbase web interface or REST API.

- h. [Optional] Starting from release 21.5.x and onward, you can deploy the Smart View Server to configure Oracle Smart View for Office (Mac and Browser). To do this, select the checkbox **Deploy Smart View Server**. This action deploys Oracle Smart View for Office (Mac and Browser) for Essbase and starts the Smart View server along with the Essbase server. See [Deploying and Administering Oracle Smart View for Office \(Mac and Browser\)](#) and [Working with Oracle Smart View for Office \(Mac and Browser\)](#)

 **Note**

Smart View for Office (Mac and Browser) can only be deployed during a fresh Essbase deployment. If you are upgrading from a previous release, the option to deploy Smart View for Office (Mac and Browser) is not available.

8. In Essbase Instance - if **Upgrade Stack** was selected above in General Settings, then follow the process in [Upgrade Stack](#). When done, skip to the [Review Page Step](#) below.
9. 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](#).
10. 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/IAM Instance GUID** and **IDCS/IAM Application Client ID** values, which you recorded as pre-deployment requirements (see [Table 2-2](#)), after you created a confidential Identity Cloud Service Application.
 - c. Enter **IDCS/IAM Essbase Admin User** value. 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 wasn'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](#).

d. The default for **IDCS/IAM Endpoint URL** is `identity.oraclecloud.com`. Change this only if your **IDCS/IAM Endpoint URL** is different.

11. In Secret Selection:

- If you want to select secrets stored in different compartments, select the checkbox **Show Advanced Secret Selection Options**. On selecting the checkbox, three additional fields **Compartment for Essbase System Admin Password**, **Compartment for Database Admin Password**, and **Compartment for IDCS/IAM Application Client Secret** will be available. If the checkbox is not selected, you will be able to choose secrets only from the same compartment as that of the deployment.
 - a. In the **Compartment for Essbase System Admin Password** field, choose the compartment where you have stored the secret for Essbase System Admin Password.
 - b. In the **Compartment for Database Admin Password** field, choose the compartment where you have stored the secret for Database Admin Password.
 - c. In the **Compartment for IDCS/IAM Application Client Secret** field, choose the compartment where the secret for IDCS/IAM Application Client is stored. This field is available only if you have selected **IDCS** as the **Identity Provider** under the **Identity Configuration** section.
- In the **Essbase System Admin Password** field, choose the secret that contains the password for the Essbase System Admin. See [Create a Vault and Secrets, and Encrypt Values](#).
- In the **Database Admin Password** field, choose the secret that contains the Database Admin password.
- In the **IDCS/IAM Application Client Secret** field, choose the secret that contains the password for IDCS/IAM Application Client. This field is available only if the **Identity Provider** is chosen as **IDCS** under the **Identity Configuration** section.

12. 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.

13. In Database Configuration, after you have reviewed database connectivity recommendations and rules (in [Set Relational Database Connectivity](#)), select from the following options.

Database options and considerations:

- Starting with 21.4, you have the option of creating a secure database. You can have restricted access through the VCN you configure or through a private endpoint. These

options are only available when you create a new database on a new network (when you did not select **Use Existing Network** or **Use existing database**).

- a. Select **Show Advanced Database Options**.
- b. [Optional] If you want to restrict database access to a private endpoint in an OCI VCN, select **Private endpoint access only for database**.
- c. Accept the default new database type **Autonomous Transaction Processing** (ATP) or select **Autonomous Data Warehouse** (ADW). The database will have a private endpoint, a private IP, and the database only communicates through the created OCI VPN.

- If you want to configure Autonomous Database with Essbase Launch URL, select **Add Essbase URL to Autonomous Database**. The link to your Essbase instance will appear in Data Studio suite of tools on the Autonomous Database instance, after deployment.
- If you want hardened network and security rules, see the section: Harden Network and Autonomous Database Security Rules in [Complete System and Security Hardening and Cleanup Tasks](#).
- If you plan to use the Oracle Autonomous Database deployed automatically by the stack, select the database license or accept the default.
- If you plan to use an existing Oracle Autonomous database, and select **Use Existing Database**, specify the compartment in which Autonomous Transaction Processing was created.
- If you plan to use Federated Partitions, you must provide an Autonomous AI Lakehouse Serverless instance to host the Essbase RCU schemas as well as the schema that holds the fact table. Select **Use Existing Database** option to deploy to your instance of Autonomous AI Lakehouse Serverless.
- 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.

14. 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.

15. In Bastion Configuration:

Under **Bastion Configuration (Option available only until 19.3.0.4.5)**, if **Public Essbase Node Visibility** is not set:

- a. Select **Provision Bastion** to enable the creation of a bastion.
- b. 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 and later:

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](#).

16. Click **Next**.
17. **Review Page Step** - on the **Review** page, you review the information that you provided. If you want to immediately provision the resources defined in the configuration, run the apply job on the new stack, by selecting **Run Apply**. Click **Create** (or Upgrade) to create the stack. 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.
18. Check for any log errors. If necessary, see [Troubleshoot Deployment Errors](#).
19. If **Upgrade Stack** was selected, and the job is completed, add the private IP of the Target node to the Source node load balancer (if it is there). For detailed steps for this, see the load balancer details in [Upgrade Stack](#) and [After Upgrade of Stack](#). Otherwise, you can access Essbase using the URL:
`https://essbase_node_public_ip/essbase`
20. 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 `bastion_host_public_ip` and there isn't an `essbase_node_public_ip`.
- **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.

Upgrade Stack

Use this option to upgrade an existing stack to the listing version. Supported stack features, including Essbase VM, are upgraded to the listing version.

Note

Starting with version 21.8.0.0.1, you cannot upgrade stacks from Essbase 19c using the Marketplace listing. To migrate from Essbase 19c to Essbase 21c, refer to [Migrate Essbase to Latest Cloud Version](#).

Note

The Upgrade Stack option is supported only for Essbase on OCI deployments where Autonomous AI Database is provisioned as the repository database during deployment.

Note

Please also refer to these related topics before proceeding with the Upgrade Stack process: [Before Upgrade of Stack](#), and [After Upgrade of Stack](#).

Upgrade Stack option creates an Oracle Essbase Target instance based on the existing stack configuration of the Source instance on Oracle Cloud Infrastructure (OCI) stack deployment.

Supported Versions for Upgrade

Supported Versions	Required Tasks for Source Instance Before Upgrade
21.x	Collect extended-metadata.

Target Instance Creation

As part of the Target instance creation, the following occurs.

- Source instance is shut down
- Clones of the Source instance's block volumes are created
- Target instance is created using the latest published image of Essbase on OCI:
 - Using the same VCN as subnet of the Source instance
 - Using the block volume clones of the Target instance
 - Using the same database as the Source instance

After the initialization of the Target instance, Essbase is started on the Target instance. Essbase applications and data from the Source instance are then available for use on the Target instance. You can use the same ssh key, passwords and credentials used for logging in to the Source instance to log in and use the Target instance.

Supported Features

Note

Features that are specific to the new upgrade version, and require additional configuration of the stack, are not available on Target instances created using Upgrade Stack. For example, Catalog Object Storage is not supported in 21.4, 21.5 and Smart View for Office 365 is not supported in 21.5.

Upgrade Stack Process

1. After you selected the Upgrade Stack option for deployment from Marketplace, the appropriate version of the source instance and compartment, and entered the name of your choice for the stack, perform the following steps.
 2. Make sure that you reviewed [Before Upgrade of Stack](#) before proceeding.
 - a. Choose **Source Instance Essbase Stack Version** from the drop-down list. Use "**Essbase 21c - 21.2 or above**", if source instance is 21.2 or higher. Choose "**Essbase 21c - 21.1**", if source instance is 21.1.
 - b. Choose **Source Instance** from the drop down where the previous version was installed.
 - c. Copy the JSON output that you collected earlier (in [Upgrade Stack](#)) from the OCI command ran on the Source instance to be upgraded. Enter it in the **Source Instance Metadata** input field.
 - d. [Optional] Select **Specify the Private IP** for the target to specify a private IP for the new node, and then enter Target IP Value. Enter the available IP address from the same subnet as the source.
 - e. Choose **Target Instance Shape** from the drop down list. Choose "**Upgrade to recommended Flex Shape**", if you want to upgrade it to the recommended shape for this release. The recommended shape for 21.6 is VM.Standard.E5.Flex. Choose "**Keep the same shape as Source Instance**"(default) if you want to retain the same shape as the Source Instance's shape.
 - f. If you want to select secrets stored in different compartments, select the checkbox **Show Advanced Secret Selection Options**. On selecting the checkbox, three additional fields **Compartment for Essbase System Admin Password**, **Compartment for Database Admin Password**, and **Compartment for IDCS/IAM Application Client Secret** will be available. If the checkbox is not selected, you will be able to choose secrets only from the same compartment as deployment.
 - g. In the **Compartment for Essbase System Admin Password** field, choose the compartment where you have stored the secret for Essbase System Admin Password.
 - h. In the **Compartment for Database Admin Password** field, choose the compartment where you have stored the secret for Database Admin Password.
 - i. In the **Compartment for IDCS/IAM Application Client Secret** field, choose the compartment where the secret for IDCS/IAM Application Client is stored.
 - j. Click **Next**.

Note

After you create the stack with the new image, the Source instance will go to a stopped state. The Source instance node should not be started while the new node is running.

Note

The new node will use the same database, clones of the disk volumes, SSH key, shape, and so on, as used by the old (source) node.

3. Continue with the **Review Page Step**, towards the end of the [Create Stack](#) topic, and then proceed until the end of that topic.

Before Upgrade of Stack

Here are some tasks and considerations before you perform an upgrade.

- Stop the Essbase server using stop.sh on the Source instance. This is recommended before you perform Upgrade Stack.
- Record Load Balancer details, if your Source instance used Load Balancer. These details are used after Upgrade Stack, to add the new target node private IP to the load balancer backend set.
- Record provisioned LDAP user/groups, if your Source instance included any, and make them available after the upgrade.
- Name the Source instance components according to original deployment naming conventions, if they changed. Check the latest Release Notes for Known Issues that apply to your Source instance.
- Add dynamic group policies:

```
Allow dynamic-group <dyn group> to manage instance-family in compartment
<name>
Allow dynamic-group <dyn group> to manage virtual-network-family in
compartment <name>
```

- Check whether the steps mentioned in [Source Instance Restored From Backup](#) apply to your upgrade - if so, they must be .

Source Instance Restored From Backup

If your source instance is restored from backup or uses a disaster recovery instance, do the following.

1. Stop Essbase Source instance.

```
/u01/config/domains/essbase_domain/esstools/bin/stop.sh
```

2. Unmount data volume.

```
sudo umount /u01/data with opc user
```

3. Disconnect data volume using iSCSI commands.

4. Detach using OCI console:
 - a. Temporarily disable /etc/fstab config and data block volume entries.
 - i. ssh to target compute (as opc user).
 - ii. Run the following.

```
sudo vi /etc/fstab
```
 - iii. Insert # in front of the /u01/config and /u01/data entries.
 - iv. Save the file.
 - b. Detach data and config block volumes from the target Essbase compute. Note the iSCSI caution and be sure to unmount and disconnect each volume before detaching using the OCI console.
 - i. To unmount:
 - i. ssh to target the compute (as opc user).
 - ii. Run the following.

```
lsblk
```
 - iii. Run the following.

```
sudo umount /u01/data
```
 - iv. Repeat these steps to unmount the data block volume.
 - ii. To disconnect iSCSI:
 - i. In the OCI console, select the target compute.
 - ii. Select resources > attached block volumes.
 - iii. From the Actions menu  for the config block volume, select **iSCSI Commands & Information**.
 - iv. Copy iSCSI commands for disconnect.
 - v. ssh to the target compute (as opc user).
 - vi. Paste the disconnect commands you copied and press enter.
 - vii. Repeat these steps to disconnect the data volume.
 - iii. To detach:
 - i. In the OCI console, select the target compute.
 - ii. Select resources > attached block volumes.
 - iii. From the Actions menu  for the config block volume, select **Detach**.
 - iv. Repeat these steps to detach the data block volume.
5. Rename (clone) the data volume as tag-data-volume on OCI console page.
6. Reattach the data volume using the following OCI CLI command (as opc user).

- for 21.1

```
oci compute volume-attachment attach-iscsi-volume --instance-  
id $instanceid  
--volume-id $datavolumeid --display-name data-volume-attachment --auth  
instance_principal
```

- for 21.2 and higher

```
oci compute volume-attachment attach-iscsi-volume --instance-  
id $instanceid  
--volume-id $datavolumeid --display-name data-volume --auth  
instance_principal
```

7. Copy the ICSI commands from OCI console page, enter them in the session, and connect using iSCSI.
8. Update /etc/fstab with data volume UUID.

```
sudo vi /etc/fstab
```

9. Mount data volume.

```
sudo systemctl daemon-reload
```

```
sudo mount -a
```

10. Start server.

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

11. Proceed to collect metadata, as shown in [Collect Source Instance Metadata](#).

Collect Source Instance Metadata

Here are steps to collect metadata for a source instance.

1. Access the source instance (to be upgraded) using SSH. Change to oracle user and stop the Essbase service using the commands:

```
sudo su oracle  
/u01/config/domains/essbase_domain/esstools/bin/stop.sh
```

2. After stop.sh is completed run the following command:

```
oci compute instance get --instance-id $(oci-metadata -j | jq -r  
.instance.id') --auth instance_principal | jq '.data."extended-metadata"'
```

Note

You can also run the above command as opc user.

3. The output from the above is in JSON format. Copy the entire output contents (including the braces {}), to be used later in Source Instance Metadata field during new upgrade stack creation.

After Upgrade of Stack

Here are some tasks and considerations after you perform an upgrade.

- If any errors occur on the initial startup of the Target instance, try restarting the Essbase server processes using stop.sh and start.sh, or rebooting the compute instance.

ⓘ Note

While the Target instance is running, do not start the Source instance. Simultaneous operation of both the Target instance and the Source instance may lead to synchronization issues.

- If you had LDAP users/groups provisioned in the Source instance, you must manually back them up and restore them on the Target instance
- If you had Load Balancer in the Source instance, you must manually update the Load Balancer with the new private IP of the Target instance.
 1. Go to **Networking > Load Balancers**, and select the load balancer you want to edit.
 2. (In the left-side menu), under **Resources**,
 - a. Select **Backend Sets**, and select the set **essbase**.
 - b. Select **Backends**, and select **Add backends**.
 3. Select the **Target instance**, set port to 443, and click **Add**.
 4. Remove the private IP of the Source instance, from the backend list.
- After the upgrade is completed, you can ssh to the new instance with the same private key.
- Go to the confidential application – client configuration (of the previous stack) and update it with the new IP address, if necessary.

If you previously required: dynamic-group policy for “manage instance-family”, it can be removed.

Complete Post-Deployment Tasks

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

- [Secure Your Network](#)
- [Modify the Confidential Identity Application](#)
- [Set Up Multiple Virtual Hosts](#)
- [Set Up SSL Certificates](#)
- [Test Connectivity to Essbase](#)
- [Complete System and Security Hardening and Cleanup Tasks](#)
- [Encryption at Rest for Essbase Applications](#)
- [Troubleshoot Deployment Errors](#)

For information on setting up a trusted certificate authority, see [Managing SSL Certificates](#).

Secure Your Network

After you deploy the Essbase stack on Oracle Cloud Infrastructure, take steps to secure your network. Also read and follow the other options in this chapter to secure your communication and network.

For TLS Everywhere secure communication topology diagram and components, see [About Securing Your Communication and Network](#).

Also see [Ways to Secure Your Network](#)

Modify the Confidential Identity Application

After deploying the Essbase stack from Oracle Cloud Marketplace, update your confidential application in OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) with the correct Essbase URLs.

If you need to support multiple access points to the same Essbase instance integrated with IAM or IDCS, follow the instructions in [Set Up Multiple Virtual Hosts](#) to register multiple Redirect URIs for Essbase, instead of the following steps.

- [IAM](#)
- [IDCS](#)

IAM

1. Log in to the Oracle Cloud Infrastructure Console. Select **Identity & Security**.
2. Click **Domains**, and click the name of the identity domain that is reserved for the Essbase stack.
3. Click **Integrated applications** within that domain.
4. Locate and select your confidential application.
5. In the **Client Configuration**, update the Essbase **Redirect URL** to reflect the actual Essbase URL.

`https://192.0.2.1/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`

If you deployed a load balancer, include port 443 in the post logout redirect:

`https://192.0.2.1:443/essbase/jet/logout.html`

7. Save the updated confidential application.

IDCS

1. 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`

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`

If you deployed a load balancer, include port 443 in the post logout redirect:

`https://192.0.2.1:443/essbase/jet/logout.html`

7. Scroll up and save the updated confidential application.

Set Up SSL Certificates

After you deploy the Essbase stack, Oracle highly recommends that you update the SSL certificates, 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 want use MaxL with self-signed certificates, see [Manage Essbase Using the MaxL Client](#).

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. 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.

1. 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.

Example location for a stack deployment on OCI:

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

The example location is likely different for an independent deployment:

```
/scratch/user/oracle_home/user_projects/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 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, download the certificate provided with Essbase.

```
echo -n | openssl s_client -connect <ESSBASE-ENDPOINT>:443 | sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/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 -noprompt -import -trustcacerts -alias mycert -file /tmp/lbr.cert -keystore /usr/java/default/jre/lib/security/cacerts -storepass
```

Note that the `-storepass` value is the default, or existing, cacerts keystore password.

3. Restart the Java process, if the Java client is WebLogic.
4. Stop and restart the Essbase 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"
```

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

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.

Note

For Essbase login URLs, see Essbase, REST, and Smart View Client URLs.

Log In to the Essbase web interface

1. 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 Essbase 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 IAM or 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

1. 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 Essbase 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 IAM or 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://192.0.2.1/essbase`.
2. Start the Essbase stack, as described in [Start the Essbase Stack](#).
3. 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
```

6. Run the CLI shell, `esscs.bat` or `esscs.sh`, and log in as the initial Essbase service administrator, which is the IAM or 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://192.0.2.1/essbase`.
2. Start the Essbase stack, as described in [Start the Essbase Stack](#).
3. 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 IAM or 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.

Link Essbase Instance to Autonomous Database Data Studio

The registration script configures Autonomous Database with Essbase Launch URL. The link to your Essbase instance will appear in Data Studio suite of tools on the Autonomous Database instance.

This feature is available if you have created a new deployment of Essbase 21.7.0.0.1 or above from Oracle Cloud Marketplace with Autonomous Data Warehouse Serverless as the repository database (RCU). Only ONE Essbase instance link can be added per Autonomous Database.

The option is available during deployment. For more information, see [Create Stack](#).

Alternatively, you can add the link by following the steps mentioned below.

About Data Studio

The Data Studio Tools allow you to load data from cloud and other diverse sources, analyzes it and gain insights from it. You can share the result of the analysis with other users. It is a one-stop application of your analytics tool from multiple data sources. This tool makes sure that there is seamless transition between different applications. The multiple ways of navigation do not impact the progress of your work.

The Data Studio comprises of the Data Load, the Data Analysis, the Data Insights, Catalog, Data Marketplace and the Data Share tool.

For more information, see [The Data Studio Overview Page](#).

Run Registration Script

To link Essbase to Autonomous Database in case the option was not selected during deployment, you need to run the Essbase URL registration script.

1. Log in to the Essbase compute instance as user **opc**.
2. Switch to user **oracle**:

```
su oracle
```

3. Navigate to `/u01/vmtools/config/adb_url`:

```
cd /u01/vmtools/config/adb_url
```

4. Run `configure-essbase-url.sh`, using **vault** option, to register the current instance's URL on its repository database.

Example:

```
./configure-essbase-url.sh --vault
```

5. The link to your Essbase instance should now appear in Data Studio suite of tools.
6. If needed, remove the URL link by running the script with **drop** option.

Example:

```
./configure-essbase-url.sh --vault --drop
```

Complete System and Security 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 AI Database or other Oracle Cloud Infrastructure-supported 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 Database Access

If you want secure database access, through the VCN you configure or through a private endpoint, using a new database and VCN, see the options in the Configure Database section in [Deploy Essbase](#).

Harden Network and Autonomous Database Security Rules

When you deploy a secured Autonomous AI Database, additional database details are displayed in the database information in OCI Console. Under the Autonomous database information tab for the database, under the Network heading, **Subnet: application** indicates that the Autonomous AI Database is allocated to an application. The private endpoint configuration displays the Private Endpoint IP and URL and the Network Security Groups link. The Security Groups link opens the collection of Security Rules for ingress and egress traffic, which you can view and customize. Here you can add and specify CIDR values for your deployment for specific IP addresses, to further harden security. Click the menu link in the right column for a security rule, and then click Edit. When you associate an Autonomous AI Database to a network security group, it is considered as secure, as defined by the security rules.

Secure Your Network

See [Ways to Secure Your Network](#).

Plan for RCU Password Expiry

Essbase RCU passwords expire in one year unless you set them to never expire. See [Fix Expired RCU Passwords that Use Autonomous Database as Repository](#) for more information.

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.

(Optional) Delete Stack

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.

Encryption at Rest for Essbase Applications

To prevent unauthorized access to Essbase applications, you can encrypt them at rest using the Oracle Cloud Infrastructure (OCI) Vault.

Encryption at rest prevents non-authorized users from accessing Essbase application data or metadata by copying files. Only authorized and logged in Essbase users can read or use the application.

By default, the security feature is enabled in Essbase Marketplace instances. However, applications are not encrypted by default.

Encryption Terms

The following terminology is helpful for understanding application encryption:

OCI Vault - a service that securely stores and manages master encryption keys and secrets. As a prerequisite to Essbase stack deployment on OCI, an administrator must create the Vault and keys. You can find the Vault in the OCI Console, in the Identity & Security section.

Data Encryption Key (DEK) - A data encryption key is a piece of information that securely encrypts or decrypts data.

Master Encryption Key (MEK) - A master encryption key encrypts the DEK. As a prerequisite to Essbase stack deployment on OCI, an administrator must create the MEK in the Vault, to be used for storing required secrets.

Refer to [Create a Vault and Secrets, and Encrypt Values](#).

Encryption Precautions for the Essbase Administrator or Application Manager

Before encrypting or decrypting any Essbase application, back it up. A pre-encrypted backup is essential, because if the encryption request fails to complete, the application may become corrupted.

Refer to: [Back Up and Restore Applications](#)

If the master encryption key is deleted from the Vault, there is no way to recover the encrypted application. Therefore, it is important to fully review and understand key management practices in OCI before making the decision to encrypt Essbase applications.

Refer to: [Backing Up and Restoring Vaults and Keys in Oracle Cloud Infrastructure](#) documentation.

Policies Required for Encryption of Applications

Review [Set Up Policies](#) for the minimum required Oracle Cloud Infrastructure policies that must be implemented to encrypt Essbase applications.

Encryption Tasks for the Essbase Administrator or Application Manager

Using the REST API, you can encrypt applications on Essbase instances deployed to OCI using the Marketplace listing.

You must be an Essbase system administrator or application manager to use the REST API endpoints.

- Use Get Encryption Info to list the types of encryption that are supported by Essbase and available for the application.
- Use Encrypt Application to encrypt an application.
- Use Decrypt Application to decrypt an application.

Migration of Encrypted Applications

Encrypted applications can be migrated using Lifecycle Management (LCM) export/import, as well as Migration Utility. If you migrate encrypted applications to a different Essbase instance, the target instance must have access to the MEK and its corresponding Vault.

To migrate applications using LCM export and import,

1. Download and set up the Command-Line Interface (CLI). Refer to [Download and Use the Command-Line Interface](#).
2. Issue the `lcmexport` command to back up the application from your source Essbase instance to an LCM zip file, providing a password to protect the encrypted application. The password must be between 6-15 characters, and should not contain any of the following special characters: `?=.,*!@#&()[]{}:;`'~/~$^+<>-`

 **Caution**

If this password is forgotten, there is no way to retrieve it, and the application cannot be imported.

Example:

```
esscs lcmexport -a Sample -z Sample_lcmexport.zip -v -restEncryPassword  
enCrYpa55123%
```

3. Issue the `lcmimport` command to restore the application from the LCM zip file to your target Essbase instance, providing the password you selected on `lcmexport` to protect the encrypted application.

Example:

```
esscs lcmimport -z Sample_lcmexport.zip -o -ta Sample2 -restEncryPassword  
enCrYpa55123%
```

If you need to migrate users and groups as well as applications, use a different utility. Refer in that case to [Migrate Applications Using Migration Utility](#).

Set Up Multiple Virtual Hosts

If you need to support multiple access points to the same Essbase instance integrated with Oracle Identity Cloud Service, follow these instructions to register multiple Redirect URLs for Essbase in your confidential IDCS application.

You must register multiple Redirect URLs with Identity Cloud Service. However, you use IDCS REST API for this task, instead of the Identity Cloud Service Console.

Caution: After you perform the following steps, do not edit the confidential application using the IDCS console. This would overwrite the application and cause the multiple Redirect URLs to be lost.

1. Use the IDCS REST API to generate an OAuth token using clientid and client secret. The following example uses the `jq` command line tool to process the cURL command, creating an access token as an environment variable:

```
export ACCESS_TOKEN=`curl -u "<clientid>:<clientsecret>" -H 'Content-Type=UTF-8' --request POST https://<idcsguid>.identity.oraclecloud.com/oauth2/v1/token -d 'grant_type=client_credentials&scope=urn:opc:idm:_myscopes_' | jq -raw-output '.access_token'
```

2. Get the APP ID for the application that you want to modify. The following example queries the list of applications, adding a filter for the target application name (assuming the filter returns exactly one application):

```
export APP_ID=`curl -X GET -H "Authorization: Bearer $ACCESS_TOKEN" -H "Content-Type:application/scim+json" https://<idcsguid>.identity.oraclecloud.com/admin/v1/Apps?filter=displayName+co+"Essbase+21c" | jq -raw-output '.Resources[0].id'
```

3. Create a `patch.json` file defining the replacements for `redirectUris`, `postLogoutRedirectUris`, and `logoutUris`. Note that the port numbers are added to the URIs.

```
"schemas": [
  [
    "urn:ietf:params:scim:api:messages:2.0:PatchOp"
  ],
  "Operations": [
    {
      "op": "replace",
      "path": "redirectUris",
      "value": [
        [
          "https://192.0.2.1:443/essbase/redirect_uri" |
          "https://192.0.2.1:443/essbase/redirect_uri",
          [
            "https://192.0.2.18:443/essbase/redirect_uri" |
            "https://192.0.2.18/essbase/redirect_uri"
          ]
        ]
      ],
      "op": "replace",
      "path": "postLogoutRedirectUris",
      "value": [
        [
          "https://192.0.2.1:443/essbase/jet/logout.html" |
          "https://192.0.2.1:443/essbase/jet/logout.html",
          [
            "https://192.0.2.18:443/essbase/jet/logout.html" |
            "https://192.0.2.18/essbase/jet/logout.html"
          ]
        ]
      ],
      "op": "replace",
      "path": "logoutUri",
      "value": [
        [
          "https://192.0.2.1:443/essbase/jet/logout.html" |
          "https://192.0.2.1:443/essbase/jet/logout.html",
          [
            "https://192.0.2.18:443/essbase/jet/logout.html" |
            "https://192.0.2.18/essbase/jet/logout.html"
          ]
        ]
      ]
    }
  ]
]
```

```

        "value": "/essbase/jet/logout.html"      }      ]
    }

```

- Issue the patch command with patch.json.

```

curl -X PATCH -H "Authorization: Bearer $ACCESS_TOKEN" -H "Content-Type:application/scim+json" [https://<idcsguid>.identity.oraclecloud.com|https://%3Cidcsguid%3E.identity.oraclecloud.com/]//admin/v1/Apps/$APP_ID --data @patch.json

```

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 and Security Hardening and Cleanup Tasks .
ESSPROV-00005	Missing value for IAM or IDCS Application Client ID. See Create a Confidential Identity Application .
ESSPROV-00006	Missing value for IAM or IDCS Application Client Secret. See Create a Confidential Identity Application .
ESSPROV-00007	Missing value for IAM or IDCS Instance GUID. See Create a Confidential Identity 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 .

Error Code	Message / More Information
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 IAM or IDCS Client ID and/or Client Secret.
ESSPROV-10011	Unable to connect to the IAM or IDCS endpoint. Validate that the 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 Destroy an Essbase Stack Instance](#)
- [Access Oracle Essbase Using SSH](#)
- [Use Commands to Start, Stop, and View Status of Processes](#)
- [Restart the Essbase Compute Instance](#)
- [Resize Block Storage Volumes](#)
- [Monitor and Diagnose Essbase Operations](#)
- [Access the WebLogic Console](#)
- [Reset or Update Admin Password](#)
- [Fix Expired RCU Passwords that Use Autonomous Database as Repository](#)
- [Manage Essbase Temporary Files](#)

Start, Stop, and Destroy an Essbase Stack Instance

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

Topics:

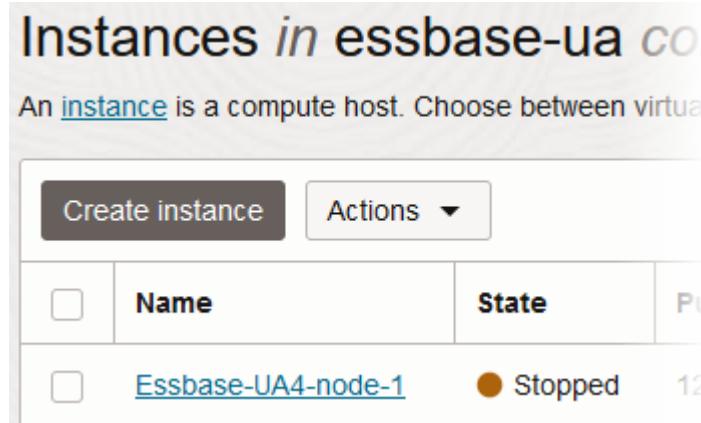
- [Prepare to Work with an Essbase Stack Instance](#)
- [Start the Essbase Stack](#)
- [Stop the Essbase Stack](#)
- [Destroy the Essbase Stack](#)

Prepare to Work with an Essbase Stack Instance

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

Navigate to the Essbase compute instance in the Oracle Cloud Infrastructure Console:

1. Log in to the Oracle Cloud Infrastructure Console at [cloud.oracle.com](#).
2. Open the navigation menu and click **Compute, Instances**.
3. Choose your compartment.
4. Click the name of the Essbase instance you want to manage.



	Name	State	Public IP
<input type="checkbox"/>	Essbase-UA4-node-1	● Stopped	128.111.111.111

5. Under **Primary VNIC**, find the **Public IP Address**.

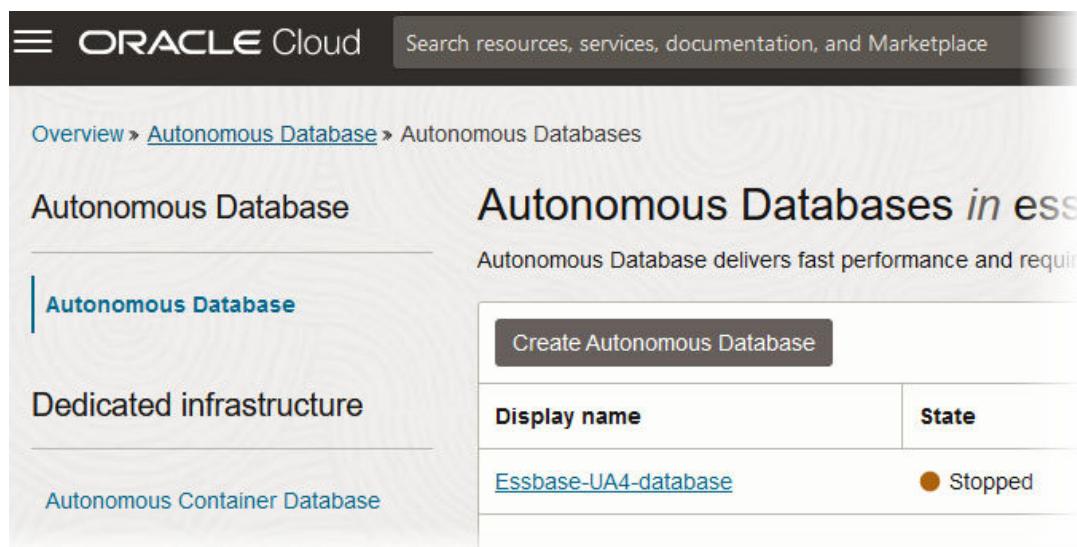
Start the Essbase Stack

The Essbase stack on Oracle Cloud Infrastructure includes the Autonomous AI Database instance, the compute instance, two block storage volumes, object storage bucket, load balancer, and additional network components.

To start the Essbase stack on OCI, you must start the Autonomous AI Database instance first, and then the compute instance. Starting the compute instance starts Essbase.

Start the Autonomous AI Database instance:

1. Log in to the Oracle Cloud Infrastructure Console at cloud.oracle.com.
2. Open the navigation menu and click **Oracle Database**.
3. Under **Autonomous Database**, select **Autonomous Data Warehouse** or **Autonomous Transaction Processing**, depending on which workload type is deployed with Essbase.
4. Choose your compartment.
5. In the list of databases, click the display name of the database you wish to administer.



Display name	State
Essbase-UA4-database	● Stopped

6. Click Start.

Essbase-UA4-database		● Primary
Actions		
Database actions	Database connection	Performance hub
Manage resource allocation	More actions	
		Start
		Restart
		Restore
		Create clone

Start the Essbase compute instance:

1. Confirm in the Oracle Cloud Infrastructure Console that the Autonomous AI Database is available.



2. Navigate to the compute instance on the Oracle Cloud Infrastructure Console.
3. Click **Start**.

Instances in essbase-ua

An [instance](#) is a compute host. Choose between virtual and physical instances.

Create instance		Actions
<input checked="" type="checkbox"/>	Name	Reboot
<input checked="" type="checkbox"/>	Essbase-UA4	Start
1 selected		Stop

The Essbase compute instance takes a few minutes to start.

Stop the Essbase Stack

The Essbase stack on Oracle Cloud Infrastructure includes the Autonomous AI Database instance, the compute instance, two block storage volumes, object storage bucket, load balancer, and additional network components.

To stop the Essbase stack on OCI, you must stop the compute instance first, and then optionally stop Autonomous AI Database. Stopping the compute instance stops Essbase.

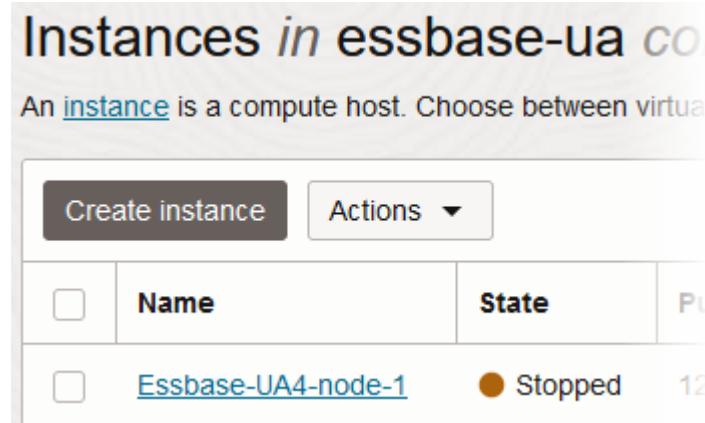
Stop the compute instance:

1. Navigate to the compute instance on the Oracle Cloud Infrastructure Console.
2. Select the instance name, and from the **Actions** menu, click **Stop**.

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

Optionally, stop the Autonomous AI Database instance, if it is not being used:

1. Confirm in the Oracle Cloud Infrastructure Console that the Essbase compute instance is already stopped.



	Name	State	Pl
<input type="checkbox"/>	Essbase-UA4-node-1	Stopped	12

2. Open the navigation menu and click **Oracle Database**.
3. Under **Autonomous Database**, select **Autonomous Data Warehouse** or **Autonomous Transaction Processing**, depending on which workload type is deployed with Essbase.
4. Select your compartment.
5. In the list of databases, click the display name of the database you want to administer.
6. From the **More actions** menu, click **Stop**.

Billing for storage continues when Autonomous AI Database is stopped. Only the compute instance and the Autonomous AI Database instance are stopped when you stop the stack.

Destroy the Essbase Stack

Destroying the Essbase stack terminates all resources that were created when the stack was created, including the Essbase compute node, the Autonomous AI Database instance, and network components.

Note

The Destroy action is permanent and cannot be reversed.

Before you destroy 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. 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.
 - a. ssh into the compute instance, and run the following utility to drop the schemas.

```
/u01/vmtools/sysman/drop-schema.sh
```

- b. In the object storage bucket for the database backups, manually delete all items.

2. Log in to the OCI Console and click **Developer Services**.
3. Under **Resource Manager**, click **Stacks**.
4. Click on your stack.
5. Click **Destroy**.
6. When prompted for confirmation, click **Destroy**.
7. 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 destroy the stack. For example, terminate the networking load balancer (if used), terminate the virtual cloud network (VCN), and terminate the Autonomous AI Database.

If an email address is associated with your user profile, you receive an email notification.

Proceed with destroying the stack regardless of whether the job passes or fails.

Note

There are cases in which the Destroy job may fail:

- if private endpoint is chosen for the database; re-run Destroy job
- if catalog storage is done in object storage bucket: delete the objects in the bucket and then re-run Destroy job
- if backup operation was performed; delete the backup from the bucket and re-run Destroy job

Refer to destroy job logs for more information.

8. After the destroy job is complete, go back to the stack details, and under **More actions**, click **Delete Stack**.
9. When prompted for confirmation, click **Delete**.

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 `essbase_node_public_ip` address, you can access it directly with `ssh` 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

1. In OCI Console > Identity & Security > Bastion.
2. 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.
5. Select the **Target Subnet** on which Essbase node is deployed.
6. Enter **CIDR** – IP or range of IPs from which access will be allowed.
7. Click **Create Bastion**. Bastion is created.

Create a Session

1. In OCI Console > Identity & Security > Bastion > click relevant created Bastion name link.
2. 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.
6. 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 <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.

```
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](#)

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:

```
$ ./status.sh
.....
Name          Type           Machine          Status
----          ----           -----          -----
essbase_server1    Server        myhost.example.com    RUNNING
AdminServer      Server        myhost          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 essbase_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 essbase_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**.

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.

```
$ sudo xfs_growfs -d /u01/data
meta-data=/dev/sdc          isize=256  agcount=4, agsize=16777216
blksize=4096
=                         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
naming =version 2          bsize=4096  ascii-ci=0  ftype=1
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.

Note

If you created Connections or Datasources after applying the patch, and then roll back the patch, you must re-create the Connections or Datasources.

- [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.

1. 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. If you do not have access to patches in My Oracle Support, please open a service request with 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/patch 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/patch 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.
 - If the Essbase instance is running, first stop Essbase system service, as user **opc**, as described above.
 - Run the rollback command:

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

For example: \$ORACLE_HOME/OPatch/patch rollback -id 30081463

- 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 stack deployment on Oracle Cloud Infrastructure.

You cannot roll back to a version that is older than the one used to configure the Essbase instance.

Note

If you created connections or Datasources after applying the patch, and then roll back the patch, you must re-create them. This step is required for drill through and other connection-dependent features to continue working as before.

Note

If Essbase is configured to use Object Storage buckets or Smart View for Office 365, rolling back to any Essbase version prior to 21.4 is not supported.

Note

If you are updating a stack image, see the stack deployment instructions, under the General Settings page steps, in [Deploy Essbase](#).

1. 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
```

6. 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/patch 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/patch 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.

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](#).

 **Note**

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
```

where the following options are:

The syntax is as follows:

```
./collect-diagnostics.sh [--filename] [--vault | -v] [--timeout | -T]
```

Example:

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

The script has the following parameters and options:

--filename sets the target zip file name to generate.

--vault or -v sets the script to take the required credentials (admin password) stored in the vault, accessed using the OCID, instead of prompting the user for the password.

--timeout or -T sets the timeout duration (in seconds) to collect diagnostic information. (default = 600)

If option --vault or -v is not used, you are prompted to enter the database admin (clear text) password. The protected password is not displayed as you enter it.

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:

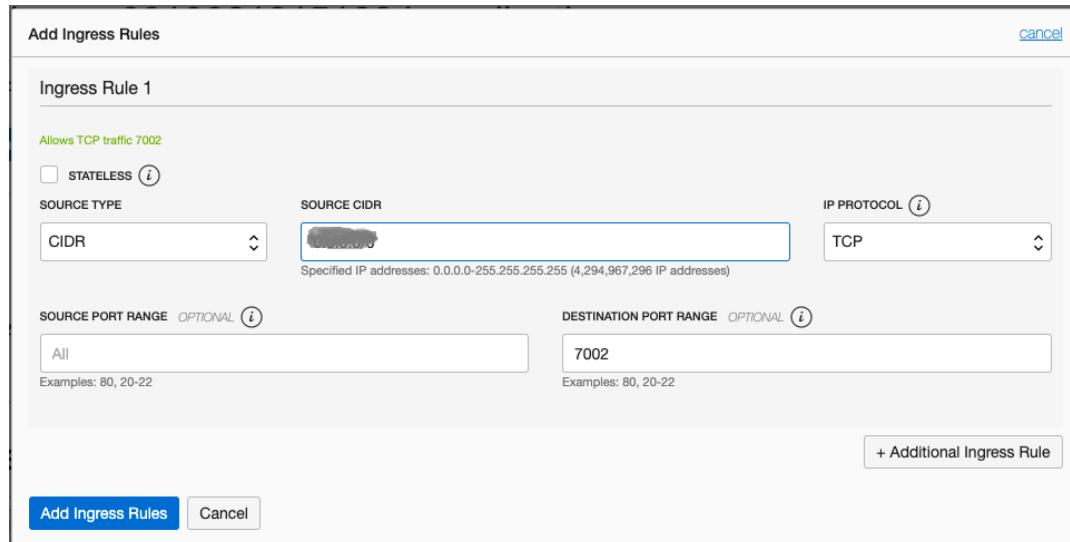
- Essbase instances are configured by default with no access to the administrative T3 WebLogic port. Oracle highly recommends that all access to the T3 port remain disabled, and that you secure it immediately. If necessary for business purposes, access to the T3 port should only be allowed from a certain 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 <prefix>-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.



The screenshot shows the 'Add Ingress Rules' dialog box. The 'Ingress Rule 1' section is displayed. It includes fields for 'Source CIDR' (set to 0.0.0.0-255.255.255.255), 'IP Protocol' (TCP), 'Source Port Range' (All), and 'Destination Port Range' (7002). There is also a note indicating 'Specified IP addresses: 0.0.0-255.255.255.255 (4,294,967,296 IP addresses)'. At the bottom, there are 'Add Ingress Rules' and 'Cancel' buttons, along with a '+ Additional Ingress Rule' link.

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

1. Follow the steps to access the WebLogic console on OCI - see [Access WebLogic Console on OCI](#).

① 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:\xxxxxx\privateKey.pem
  ProxyCommand ssh targetbastion -W %h:%p

Host targetbastion
  HostName 140.xx.xx.xx
  User opc
  StrictHostKeyChecking=no
  IdentityFile C:\xxxxxx\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:
 - i. 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.oraclecloud.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.oraclecloud.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:

```
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

1. 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:

```
export DOMAIN_HOME=/u01/config/domains/essbase_domain
```

4. 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 -----


[oracle@essxx-1 essbase_domain]$ mv $DOMAIN_HOME/servers/AdminServer/
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

Fix Expired RCU Passwords that Use Autonomous Database as Repository

Use these steps to update RCU schema passwords for Essbase Marketplace Listing Version 19.3.0.3.4 and above on OCI Marketplace deployed with Autonomous AI Lakehouse.

Essbase RCU passwords expire in one year unless you set them to never expire.

Steps to Update the RCU Passwords

1. SSH to the Essbase compute instance as opc user.
2. Run sudo su - oracle.
3. Shut down Essbase:

```
/u01/config/domains/essbase_domain/esstools/bin/stop.sh
```

4. From the directory /u01/config/domains/essbase_domain/bin/>,

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

make a copy of setStartupEnv.sh.

```
cp setStartupEnv.sh setStartupEnv_ORG.sh
```

5. Run the following command:

```
/u01/vmtools/sysman/rotate-schema-credentials.sh
```

This will prompt for the database admin password.

 **Note**

Make sure that the value for -DODBC_URL remains unchanged (in SERVER_SYSTEM_PROPERTIES setting).

 **Note**

If the following error is encountered when running rotate-schema-credentials.sh script, install cx_oracle and database_utils module to Essbase VM.

```
./rotate-schema-credentials.sh
```

Traceback (most recent call last):

File "/u01/vmtools/sysman/rotate-schema-credentials.py", line 13, in

<module>

import database_utils

File "/u01/vmtools/scripts/database_utils.py", line 16, in <module>

import cx_Oracle

ModuleNotFoundError: No module named 'cx_Oracle'

Usually this package is available as part of python install available on all OCI pods. For some reason, this package is missing. If this error is encountered, the package can be installed as follows:

```
$ pip install cx_oracle
```

If it already installed, the following message is returned:

Defaulting to user installation because normal site-packages is not writeable

Requirement already satisfied: cx_oracle in /usr/local/lib64/python3.6/site-packages (8.2.1)

6. Start Essbase using:

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

To enforce that user password never expires in the future

1. SSH to the Autonomous AI Lakehouse database instance as opc user.
2. sudo su - oracle

Run the following in SQL Developer / SQLP**To get the list of users**

```
Select * from DBA_USERS where username like 'SCHEMA_PREFIX%'
```

 ⓘ Note

The schema prefix can be obtained from the Terraform job log. Look for the rcu_schema_prefix in the Outputs section.

Create a profile with an unlimited life time for password

```
CREATE PROFILE <><>Profile_Name>>
LIMIT FAILED_LOGIN_ATTEMPTS 5
PASSWORD_LIFE_TIME unlimited
PASSWORD_REUSE_TIME 60
PASSWORD_REUSE_MAX 5
PASSWORD_VERIFY_FUNCTION verify_function;
```

Assign the users to profile with an unlimited life time for password

```
alter user <><>SCHEMA_PREFIX_OPSS>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_STB>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_WLS>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_WLS_RUNTIME>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_MDS>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_IAU_APPEND>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_IAU_VIEWER>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_IAU>> profile <><>Profile_Name>>;
alter user <><>SCHEMA_PREFIX_ESSBASE>> profile <><>Profile_Name>>;
```

A restart of Essbase services is required.

Manage Essbase Temporary Files

Essbase temporary files are stored in /u01/tmp and /tmp.

Temporary files of Essbase can be deleted if they are causing a build-up of used space and there is a space shortage.

Files in /tmp can be deleted according to user discretion.

Files in /u01/tmp can be deleted or cleared as required by space constraints.

Migrate Essbase Applications

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

Topics:

- [About Migration Tools and Use Cases](#)
- [Migrate From Essbase 11g On-Premise](#)
- [Migrate Essbase to Latest Cloud Version](#)
- [Post-Migration and Advanced Topics](#)

About Migration Tools and Use Cases

Here is an explanation of various tools and uses cases for migrating Essbase applications to Essbase 21c.

Migration Tools

The following migration tools are available on the Essbase web interface Console.

Note

In relation to migration, references to 19c in this document applies specifically to migration from Essbase 19c on OCI Marketplace deployments. Otherwise, all migration content in this document applies to migration from Essbase 21c or higher independent deployments.

- [11g Excel Export Utility](#): Exports Essbase 11g On-Premise applications to application workbooks. You can use the application workbooks to re-create the applications on the current Essbase version.
- [11g LCM Export Utility](#): Exports artifacts from Essbase 11g On-Premise as a .zip file, which you can import in to Essbase 12c or higher. This Life Cycle Management (LCM) utility can also be used to export from, and import to, Essbase 11g On-Premise. This utility packages into a zip everything you need to support migration to the current version. Download `EssbaseLCMUtility.zip`, and see the enclosed README for usage details. With the 11g LCM Export Utility, you can create applications by exporting applications and cubes. You then import them using the Essbase Command Line Interface.
- **Command Line Interface (CLI)**: A scripting interface that uses REST APIs to perform most common Essbase administrative actions. CLI includes an LCM import command you can use for migrating 11g LCM Export Utility .zip files exported from Essbase 11g On-Premise. The LCM export and LCM import commands also facilitate migrating applications from instances on versions 12c or higher. Use cases include:
 - Migration of applications from different instances of the same version (e.g. Dev > Test > Prod)

- Upgrade utility for migration to higher versions of Essbase

Both LCM commands are based on the REST API. You can access these commands in Essbase and run them from the following:

- Jobs in the Essbase web interface: **Export LCM** and **Import LCM** jobs. These are used for 21c to 21c migration.
- Command Line Interface (CLI): using lcmexport and lcmimport commands. lcmimport can be used to import data zip file originating from Essbase 11g, or from 19c or higher. lcmexport can be used for Essbase 21c to 21c migration. Using its login command, authenticate Essbase.
- REST API: Execute Job operation (using jobtypes lcmExport and lcmImport)

 **Note**

LCM export command can be used to export an application to a zip file and then LCM import is used to import it to the target version of Essbase. LCM Import command can alternatively be used to migrate an application originating from Essbase 11g On-Premise (exported from 11g using the 11g LCM Export Utility to an export zip file and then imported from the zip file it to the target Essbase command-line).

- **Migration Utility:** Manage migration of an entire Essbase instance, for Essbase 12c or higher. In addition to migrating application artifacts, this utility also helps you migrate user role assignments and users/groups from supported identity providers. Download `migrationTools.zip`, and see the enclosed README for usage details. The tool supports migration from Essbase 19c or Essbase 21c. Use the Migration Utility between major versions and between un-patchable releases to ensure database schema integrity after migration. Not supported for use with EPM Shared Services security mode.

Migrate From Essbase 11g On-Premise

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 11g LCM Export Utility, to export an on-premises cube to a zip file, and then use the Command Line Interface (CLI) to import the zip file that imports Essbase 11g On-Premise applications, folders, and elements.

Topics:

- [Prepare to Migrate from Essbase 11g On-Premise](#)
- [Migrate Essbase 11g Users and Groups](#)

Prepare to Migrate from Essbase 11g On-Premise

If you have an existing Essbase 11g On-Premise application and cube to migrate to Essbase 21c, review the following considerations and prerequisites.

Differences Between Essbase 11g On-Premise and Essbase 21c

- Free-form data exports and imports for cubes with typed measures behaves differently in Essbase 21c. For the latest information, see Loading, Clearing, and Exporting Text and Date Measures.
- Before you migrate, review the [Differences Between Essbase 11g and Oracle Essbase 21c](#).

Task Flow for Migrating

Note

- If you used EPM Shared Services in Essbase 11g On-Premise to configure an external security provider, then step 1 below isn't required. Based on the security mode chosen during configuration, you may need to configure your target EPM instance or WebLogic to use the same external security provider as used in Essbase 11g On-Premise.
- If you're using Oracle Identity Cloud Service, configure it to use the same external security provider as used in Essbase 11g On-Premise.

1. Migrate users and groups from source Essbase 11g On-Premise EPM Shared Services to OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) (for OCI deployment) or EPM Shared Services / WebLogic LDAP (for Independent deployments). See [Migrate Essbase 11g Users and Groups](#).
2. If you're exporting non-Unicode Essbase 11g On-Premise applications, you must convert the applications to Unicode.
 - a. Use Alter System on the server, and then on a backup copy of the Essbase application, prior to running 11g LCM Export Utility, to enable Essbase itself to support the Unicode application.
 - b. For non-Unicode, **block** storage applications, export the application using – converttoutf8 option in the export command. See [11g LCM Export Utility Options](#).
For non-Unicode, **aggregate** storage applications, follow the manual Unicode conversion instructions in [Convert Non-Unicode Aggregate Storage Application to Unicode Mode](#).
3. Migrate 11g applications:
 - a. Export Essbase applications using the 11g LCM Export Utility, downloaded from target Essbase 21c instance, and running the utility on the computer where Essbase 11g On-Premise is installed.

Note

To use the 11g LCM Export Utility, Java Development Kit (JDK) 8 or higher must be installed, and the following variables must be set: JAVA_HOME environment variable, and EPM_ORACLE_HOME and EPM_ORACLE_INSTANCE variables in the shell terminal.

- b. Import Essbase applications using the Essbase Command Line Interface (CLI) Utility, downloaded from the Essbase 21c instance. Run the utility for each exported zip file.

Supported Essbase Versions and Paths

The following releases have been tested for migration: 11.1.2.3.0nn, 11.1.2.4.0nn, 12.2.1, and later.

Note

Migration from 11.1.2.3 is supported, however, not every application can be migrated from 11.1.2.3 to 21c. We recommend that you upgrade from 11.1.2.3.0.n.n to 11.1.2.4.0.n.n before attempting the migration to 21c.

The following migration paths are not supported:

- Essbase 19c or 21c on OCI (Marketplace deployment) to Essbase 21c on Windows (independent deployment)
- Essbase 21c on Linux (independent deployment) to Essbase 21c on Windows (independent deployment)

Migrated 11g Artifacts

Review the 11g artifacts that are supported for migration. See [Migrated 21c Artifacts](#).

Unsupported Application and Database Settings

The following application- or database-level settings aren't supported in migration: disk volumes.

CONFIGURATION NOTES

• 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. Note: Hybrid mode is not the default if using calculation scripts - if your calculation scripts contain many sparse dependencies, consider enabling hybrid mode for calculation scripts as well.

• Implied Sharing

If you use the IMPLIED_SHARE configuration setting in your Essbase 11g On-Premise application, your implied sharing setting is migrated, for minimal disruption. For more details about implied sharing defaults in Essbase 21c, see the IMPLIED_SHARE_ON_CREATE configuration topic.

• Warning Regarding the UPPERCASECONNECTION esssql.cfg Setting

If your environment has an esssql.cfg file containing the no longer supported UPPERCASECONNECTION setting, you may get a warning like the following while performing data load operations:

```
WARNING - 1021037 - SQL Config file syntax error [UpperCaseConnection],  
ignored.
```

The way to fix this is to manually remove the UPPERCASECONNECTION setting from esssql.cfg, which is located in

```
<DOMAIN_HOME>/config/fmwconfig/essconfig/essbase/esssql.cfg
```

and then restart the Essbase servers.

- **Text and Date Measures**

Starting with Essbase 21.3 and higher, the configuration ALLOWOUTOFRANGELOAD is deprecated and the behavior will be the same as "ALLOWOUTOFRANGELOAD TRUE" in previous versions. Out of range and missing values of typed measures will be loaded to, and exported from, cubes that have textual measures. In the following data export file example, the <OutOfRange_Name> is "Invalid" and "NoColor" for Missing values.

```
"Color1" "Color2"  
"Shoes" "Massachusetts" "Q1" "#Txt:NoColor" "#Txt:Yellow"  
"Q2" "#Txt:Green"  
"Connecticut" "Q1" "#Txt:Green" "#Txt:Invalid"
```

For information on loading out of range values, refer to Loading, Clearing, and Exporting Text and Date Measures.

- **Configuration Settings**

Some default configuration values are different than they were in Essbase 11g On-Premise. Check the Configuration Reference.

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.

To modify the default values for application-level configuration settings, you use the Essbase web interface, as described in Set Application-Level Configuration Properties.

Oracle recommends managing most configurations at the application level. When you migrate applications, your application-level configuration is preserved during the LCM export and import processes. Some configurations, however, are only applicable to the Essbase Server. Most of these server configurations you specify while configuring Essbase during deployment, but you can also change server configuration defaults using essbase.cfg, if needed.

GENERAL NOTES

- **Imported SQL Timestamp Data Types**

SQL timestamp data types that formerly displayed in ODBC format [yyyy-mm-dd hh:mm:ss] now display in a different format [dd-MON-yy hh.mm.ss.mmm a], in the rules editor of the Essbase web interface and everywhere else that timestamp data types are imported. To load timestamp in the format of your choosing during data loads or dimension builds, you can convert timestamps to your chosen string format by using a SQL conversion function in the query section of the load rule. The following SQL query example uses the format function: `SELECT introdate, format(introdate, 'yyyy-MM-dd hh:mm:ss') FROM tbc.dbo.product`

- **Upgrading EPM Applications - Calculations and Filters**

After upgrading applications to Essbase 21c, you can't provision members to calculations or filters from EPM Shared Services console. You must use the Essbase web interface to assign members. Refer to: Assigning Filters and Access to Calculations.

- **Partitions**

When you perform the LCM 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.

After you roll back an OPatch, you may need to recreate transparent and replicated partitions, and re-validate the partitions.

- **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

- **Location Aliases**

LCM doesn't support Location alias credentials migration. After you migrate your applications from Essbase 11g On-Premise, you must replace your location aliases. You can use the following automated method, or a manual method using MaxL.

Automated method to replace location aliases

1. Unzip the LCM exported zip.
2. Go to {ApplicationName}\Databases\{dbName}\Location Aliases.
3. Open the file under this directory. This is an XML format file, where userName and password field are empty. You can provide the credentials.
4. Zip the directory again.
5. Import the application using the zip directory.

Sample xml file

```
<?xml version="1.0" encoding="UTF-8"?>
<java version="$VERSION$" class="java.beans.XMLDecoder">
<object class="oracle.essbase.lcm.essbase.EssbaseLocationAlias">
<void property="aliasAppName">
<string>
{appName}
</string>
</void>
<void property="aliasCubeName">
<string>
{dbName}
</string>
</void>
<void property="aliasHostName">
<string>localhost</string>
</void>
<void property="password">
<string>password</string>
</void>
<void property="userName">
<string>lauser</string>
</void>
</object>
</java>
```

Manual method using MaxL to replace location aliases

An alternative option manual method uses MaxL. After you import source applications by performing the CLI LCMImport job, re-create location aliases using Create Location Alias.

- **Custom-Defined Functions and Macros**

FOR INDEPENDENT DEPLOYMENTS - If you have custom .jar file that you use for custom-defined calculation functions and macros, these are not migrated by the 11g LCM Export Utility. You must move them manually. To do this,

1. Note the location of your <Essbase Path> and <Application directory> (ARBORPATH) in Essbase 21c. Refer to Environment Locations in the Essbase Platform if needed.
2. Migrate global (system-level) functions by copying your .jar files from <Essbase Path>/java/udf on your source instance to <Essbase Path>/java/udf on your target Essbase instance.
3. Migrate local (application-level) functions by copying your .jar files from the application directory on your source instance to the application directory on your target Essbase instance. In other words, copy the .jar files from <ARBORPATH>/app/<app_name> on your 11g server to <Application directory>/app/<app_name> on your 21c server.
4. On your target Essbase instance, add JVMMODULELOCATION to essbase.cfg, providing as an argument the path to the JVM library on your system.

- **Client Service URLs**

FOR INDEPENDENT DEPLOYMENTS - In Essbase 11g, Provider Services is the middle-tier data-source provider to Oracle Essbase for Java API, Smart View, and XML for Analysis (XMLA) clients.

Provider Services functionality is integrated with WebLogic. Update the client URLs to the current format.

Clients	Former URL for Connecting Provider Services to the Specified Client	New URL in Essbase 21c
Java API	http:// server_name:port/aps/JAPI	http://server_name:port/ essbase/japi
Smart View	http:// server_name:port/aps/ SmartView	http://server_name:port/ essbase/smartview
XML for Analysis (XMLA)	http:// server_name:port/aps/XMLA	http://server_name:port/ essbase/xmla

- **Outline Solve Order and Enabled Typed Measures**

FOR OCI MARKETPLACE DEPLOYMENTS - After you migrate an Essbase application from Essbase 11g On Premises server to Essbase deployed on OCI via Marketplace, you must enable typed measures in the outline before you can change the outline solve order.

In Essbase deployed on OCI, the **Typed Measures Enabled** outline property is set to FALSE by default. In order to change the solve order, the Typed Measures Enabled property first needs to be changed to TRUE. To change this property, see About Typed Measures in *Database Administrator's Guide for Oracle Essbase*.

Migrated Essbase 11g Artifacts

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

Note

After upgrading EPM applications to Essbase 21c, you can't provision members to calculations or filters from EPM Shared Services console. You must use the Essbase web interface to assign members. See topics: [Assigning Filters and Access to Calculations](#).

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. If users and groups are already migrated before importing the application, then User/Group - Calc associations would also be imported.
Custom-defined functions and custom-defined macros	no	Any jar files used by Custom-defined functions must be manually migrated.
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. If users and groups are already migrated before importing the application, then User/Group - filter associations would also be imported.
Linked Reporting Objects (LROs)	no	-
Location aliases	no	Location aliases are migrated with the cube. LCM doesn't support location alias credentials migration. After you migrate your applications from Essbase 11g On-Premise, you must replace your location aliases. See Location Aliases section in Prepare to Migrate from Essbase 11g On-Premise .
Log files	no	-

Artifact	Supported for migration	Exceptions/Comments
MDX Reports	no	MDX scripts, triggers, and macros are not supported for migration from 11g Essbase, and therefore are NOT migrated, and not upgraded to Essbase 21c in EPM.
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	no	LCM export operations and Migration Utility export do not support migration of scenarios for applications.
Substitution variables	yes	Application- and cube-level substitution variables are migrated. Server-level substitution variables are migrated if you use the optional <code>-include-server-level</code> option.
Users	no	-
User roles	-	User roles are migrated if you use the <code>-exportepmroles</code> option.

Convert Non-Unicode Aggregate Storage Application to Unicode Mode

Before exporting an aggregate storage application from Essbase 11g On-Premise in preparation for migration to Essbase 21c, convert it to unicode mode.

Note

Other references for ESSCMDQ workaround information are: [How to Use ESSCMDQ to Compact Outlines \(Doc ID 1534496.1\)](#) and [Essbase 21c Compress Dimension Option Not Visible in JET UI \(Doc ID 2853804.1\)](#).

Client-side configuration - To set up ESSCMD/ESSCMDQ on a Windows Client

1. From the Essbase 21c web interface, download the Essbase Client for Windows and the MaxL Client.
2. Extract the MaxL Client to a `MaxLClient` directory.

3. Extract the ESSCMD client to an EssbaseClient directory.
4. Copy the `startMAXL.bat` script from the `MaxLClient` directory to the `EssbaseClient` directory. Rename the script to `startEsscmd.bat`.
5. Edit the `startEsscmd.bat` file, and add a new setting to call `esscmd`, instead of `esscmd.sh`:

```
"%ESSBASEPATH%\bin\esscmd" %*
```

6. Save the file. Run it as administrator and test logging into Essbase 21c Marketplace or Independent deployment.

Login syntax:

```
login https://<servername_or_IP>/essbase/agent <user> <password>
```

7. After you confirm that it works, download ESSCMDQ for Windows from <https://www.oracle.com/middleware/technologies/esscmdq-sampleapps-downloads.html>.
8. Extract `ESSCMDQ.exe` to the `EssbaseClient/bin` directory.
9. Copy the `startEsscmd.bat` script to `startEsscmdQ.bat`.
10. Edit the `startEsscmdQ.bat` file to call `esscmdq` instead of `esscmd`:

```
"%ESSBASEPATH%\bin\esscmdq" %*
```

11. Save the file. Run it as administrator and test logging into Essbase 21c Marketplace or Independent deployment.

Login syntax:

```
login https://<servername_or_IP>/essbase/agent <user> <password>
```

Note

- Do not use an IDCS/MSAD userid to connect. Use a 'native' user to login.
- If using a proxy server, you may need to add the following settings to the `startEsscmd/q.bat` scripts:
 - `set HTTP_PROXY=<proxyserver>:<port>`
 - `set HTTPS_PROXY=<proxyserver>:<port>`

Server-side configuration - To set up directly on the Essbase server

1. Convert the copied aggregate storage 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 aggregate storage application to Unicode mode using MaxL Shell:

1. Log in to the source Essbase 11g On-Premise 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 \(Aggregate Storage\)](#).

 **Note**

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

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

1. Download ESSCMDQ
 - a. 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,
For Linux

`$ESSBASEPATH/bin`

For example:

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

For Windows

`%ESSBASEPATH%\bin`

For example:

```
.\Middleware\EPMSystem11R1\products\Essbase\EssbaseServer\bin\ESSCMDQ.exe
```

To know the values of environment variables in your source EPM 11g installation, check the environment file.

For Linux:

```
./Middleware/user_projects/<epm_instance>/EssbaseServer/essbaseserver1/bin/setEssbaseEnv.sh
```

For Windows:

```
./Middleware/user_projects/<epm_instance>/EssbaseServer/essbaseserver1/bin/setEssbaseEnv.bat
```

By default, `<epm_instance>` would be `epmsystem1`.

- c. Make a copy of the existing script

For Linux

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

as

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

For Windows

```
.\Middleware\user_projects\epmsystem1\EssbaseServer\essbaseserver1\bin\startEsscmd.bat
```

as

```
.\Middleware\user_projects\epmsystem1\EssbaseServer\essbaseserver1\bin\startEsscmdQ.bat
```

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:

For Linux

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

For Windows

```
set ESSCMDQ_UTF8MODE=1  
set 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.
4. Copy the application folder from ARBORPATH/app directory to client directory. For ASOBck app, for example:

For Linux

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

For Windows

```
%ARBORPATH%\app\ASOBck as %ARBORPATH%\client\ASOBck
```

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

① Note

ESSCMDQ is interactive, so parameters for each command can be found in interactive mode. To see what a parameter means, enter the command, such as OpenOtl, and then press Enter to see the menu explaining the parameter. Or enter the ESSCMDQ command and press Enter, without any parameters, and the parameter menu is displayed.

Linux example

```
./Middleware/user_projects/epmsystem1/EssbaseServer/essbaseserver1/bin/  
startEsscmdQ.sh  
openotlex2 1 1 appName dbName outlineName Y Y Locale N 0  
  
writeotlex 0 1 1 appName dbName outlineName 2
```

Windows example

```
.\Middleware\user_projects\epmsystem1\EssbaseServer\essbaseserver1\bin\star  
tEsscmdQ.bat  
openotlex2 1 1 appName dbName outlineName Y Y Locale N 0  
  
writeotlex 0 1 1 appName dbName outlineName 2
```

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

For example (Linux syntax)

```
mkdir $ARBORPATH/client  
cp -r $ARBORPATH/app/ASOBck $ARBORPATH/client  
  
./Middleware/user_projects/epmsystem1/EssbaseServer/essbaseserver1/bin/  
startEsscmdQ.sh  
openotlex2 1 1 ASOBck Basic Basic Y Y "Japanese_Japan.MS932@Binary" N 0  
writeotlex 0 1 1 ASOBck Basic Basic 2  
exit
```

6. 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 (Linux syntax):

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

```
#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:

For Linux

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

For Windows

```
.\Middleware\user_projects\epmsystem1\EssbaseServer\essbaseserver1\bin\star  
tEsscmdQ.bat
```

and restructure each cube.

```
#Please replace hostname, username, password, appname and cubename with  
appropriate values  
login 'hostname' 'username' 'password'  
select appname cubename  
openotl 2 1 appname cubename outlinename y y 0  
writeotl 0 2 1 appname cubename outlinename  
restructotl 1  
closeotl 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
```

Migrate Essbase 11g Users and Groups

The task flow for migrating Essbase users and groups from Essbase 11g On-Premise to Essbase 21c varies depending on your identity provider and details about your applications.

Prerequisites and Considerations

- If, for your Essbase 11g On-Premise instance, you stored users and groups natively in EPM Shared Services, you need to export those users and groups, to import them into your security provider for Essbase 21c. Reformatting of the exported user and group files may be required if you're migrating from Shared Services into WebLogic Embedded LDAP.

Shared Services security is recommended only for Essbase customers who also use EPM applications and have user overlap between EPM applications and “stand-alone” Essbase applications. Essbase customers who don’t use any EPM applications are recommended to migrate to Essbase using the default WebLogic security, and not Shared Services. WebLogic security can be federated with many external authentication identity providers. See [WebLogic Authentication](#).

- 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.
- If you’re using native (default) identity providers, migrate users and groups from Essbase 11g On-Premise by exporting them to a CSV file and importing them after you install and configure Essbase 21c. If you’re using federated/external providers, integrate these with Essbase 21c.
- When you import user names, the following special characters are not allowed in the name.

; , = + * ? [] |< > \ " ' / [Space] [Tab]

The name length is limited to 50 characters.

Scenarios for Migrating users and groups

Scenario 1 - Exporting users/groups from Essbase 11g and importing them into Essbase 21c, which is configured in Shared Services mode

Exporting users/groups

- If the native Shared Services directory is used in the source EPM instance, then export users and groups using Shared Services Console. See [Migrating Native Directory \(Security\)](#). You should only select users and groups, and don’t migrate roles, while exporting from EPM Shared Services. User roles are migrated by the Essbase 11g LCM Export Utility.
- If a source Shared Services instance is configured to use an external security provider, then no explicit user/group export is required.

Importing users/groups

- If native Shared Services directory is used in the target EPM instance, then import users and groups using Shared Services Console.
- If the source EPM instance was configured to use an external security provider (including when you use MSAD or other LDAP-based user directories), then configure the target Shared Services instance with the same provider details. See [Configuring OID, Active Directory, and Other LDAP-based User Directories](#).

Options for the Shared Services Administrator

If you want to import the Shared Services administrator user from an Essbase 11g On-Premise instance to an Essbase 21c instance configured to use EPM Shared Services for authentication, the following considerations can help you avoid pitfalls.

Before any user migration steps, ensure you have a dedicated EPM Foundation-only instance of Shared Services that you have configured with Essbase 21c. This Shared Services instance already has a Shared Services administrator.

⚠ Caution

If the Shared Services administrator in the source 11g EPM instance isn't the same administrator that you configured in the target EPM Foundation-only instance, and you include this source administrator in the CSV file when you import users to your EPM Foundation-only Shared Services configured with Essbase 21c, your target Shared Services administrator will be overwritten by the 11g Shared Services administrator.

To handle this, select an option:

- **Only Use 11g Administrator**

Allow the source Shared Services administrator to overwrite the target administrator.

1. Export all users from Essbase 11g On-Premise to a CSV file, including the Shared Services administrator.
2. Import all users to the target, empty EPM Foundation-only instance.
3. Log in on the target instance as the 11g Shared Services administrator, and assign roles and permissions to users. Your Shared Services administrator user in the target instance retains the same user ID, but has the password you specified during configuration on the target instance.

- **Only Use Target Administrator**

Remove the 11g Shared Services administrator from the export file, so that the target administrator will not be overwritten.

1. Export all users from Essbase 11g On-Premise to a CSV file, including the Shared Services administrator.
2. From the export CSV file, remove the row containing the administrator.
3. Import the rest of the users to the target, empty EPM Foundation-only instance.
4. Log in as the target Shared Services administrator, and assign roles and permissions to users.

- **Keep Both Administrators**

Take steps to migrate the 11g Shared Services administrator without affecting the target administrator.

1. Export all users from Essbase 11g On-Premise to a CSV file, including the Shared Services administrator.
2. Edit the CSV file to remove the **internal_id** value associated with the source Shared Services administrator. This removes the Shared Services and Essbase administrator role, but keeps the user ID and password intact.
3. Import the users to the target EPM Foundation-only instance. The Shared Services administrator's user ID is migrated, but no longer has administrator role.
4. Log in as the target Shared Services administrator and grant whichever role you want to give to the 11g Shared Services administrator user ID you just migrated.

Scenario 2 - Exporting users/groups from Essbase 11g and Importing them into Essbase 21c, which is configured in WebLogic security mode

Exporting users/groups

- If native Shared Services directory is used in the source EPM instance, then export users and groups using Shared Services Console. See Migrating Native Directory (Security).

- If source Shared Services instance is configured to use an external security provider, then no explicit user/group export is required.

Importing users/groups

- If native Shared Services directory was used in the EPM 11g instance, then you may need to manually convert the file exported from Shared Services to a format that WebLogic security mode can understand.
 1. Open the users/groups zip files exported by Shared Services, and extract the files "resource\Native Directory\Users.csv" and "resource\Native Directory\Groups.csv" as target 21c files.
 2. Manually assign groups as follows: user group associations should be extracted from the source Essbase 11g CSV file, added to the target Essbase 21c CSV file, and then later imported into the Essbase 21c interface.
 3. Manually re-order the columns in these target 21c CSV files to a format that contains the user ID, first and last names (optional), email address (optional), password (optional) and role type (User, Power User, or Service Administrator).
 4. Please specify the role type field in these target CSV files as "User".
 5. Import the modified target CSV files using the Essbase 21c interface, logged in as a Service Administrator. Go to the Applications home page > Security > Import. Browse to the .csv files, and click **Import**.
- If source EPM instance was configured to use an external security provider, then please configure WebLogic with the same security provider details. See [Configuring Authentication Providers](#).
- Your existing Essbase 11g On-Premise instances use Shared Services security, with users and groups stored natively in Shared Services or with users and groups stored in an external identity provider.

During configuration of Essbase 21c, you chose a security mode: either embedded WebLogic or Shared Services. Regardless of the selected security mode, if your Essbase users and groups exist in an external identity provider, you should integrate Essbase 21c with that provider. See [WebLogic Authentication](#) and [EPM Shared Services Authentication](#), and their subtopics on external identity providers.

ⓘ Note

Reformatting of the exported user and group files may be required if you're migrating from Shared Services into WebLogic Embedded LDAP.

ⓘ Note

Shared Services security is recommended only for Essbase customers who also use EPM applications and have user overlap between EPM applications and "stand-alone" Essbase applications. Essbase customers who don't use any EPM applications are recommended to migrate to Essbase using the default WebLogic security, and not Shared Services. WebLogic security can be federated with many external authentication identity providers. See [WebLogic Authentication](#).

User Roles for Access

Assignment of user roles behavior differs from Essbase 11g On-Premise if you choose Essbase to run in WebLogic security mode. 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, you must 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.

The following Essbase security artifacts are migrated using the 11g LCM Export Utility: Essbase server-level roles, application-level roles, filter associations, and calc associations. If you choose to migrate to an Essbase instance that uses WebLogic security, LCM handles provisioning users and groups with the corresponding new roles. Note that this mapping isn't applicable if your target Essbase instance is configured to Shared Services security and the same 11g roles would remain in Essbase.

Table 4-1 Default role mapping

Source 11g EPM Shared Services Roles	Target 21c WebLogic Embedded LDAP 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
Filter	Database Access	Application
Read	Database Access	Application
Server Access	User	Server
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.

Required access for tasks:

- 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 11g LCM Export Utility and their corresponding operations: Service Administrator role for all applications; Create or Delete Application roles for only those applications created by the user.
- For importing: A user with at least Power User role (in WebLogic security mode) or Create or Delete application roles (in EPM security mode) can create applications (during import) and manage applications can create applications (during import) and manage applications.

Scenario 3 - Exporting users/groups from Essbase 11g On-Premise and Importing them into Essbase 21c, which is configured with IAM or IDCS as Identity Provider

See [Export 11g Users and Groups to Essbase 21c Configured with IAM or IDCS](#).

Export 11g Users and Groups to Essbase 21c Configured with IAM or IDCS

If you're using native default EPM 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 OCI Identity and Access

Management (IAM) or Oracle Identity Cloud Service (IDCS) to use the same external security provider that Shared Services used.

Notes

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.

Required User Roles for Access

Note that the following Essbase security artifacts are migrated using the 11g Export 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 4-2 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
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.

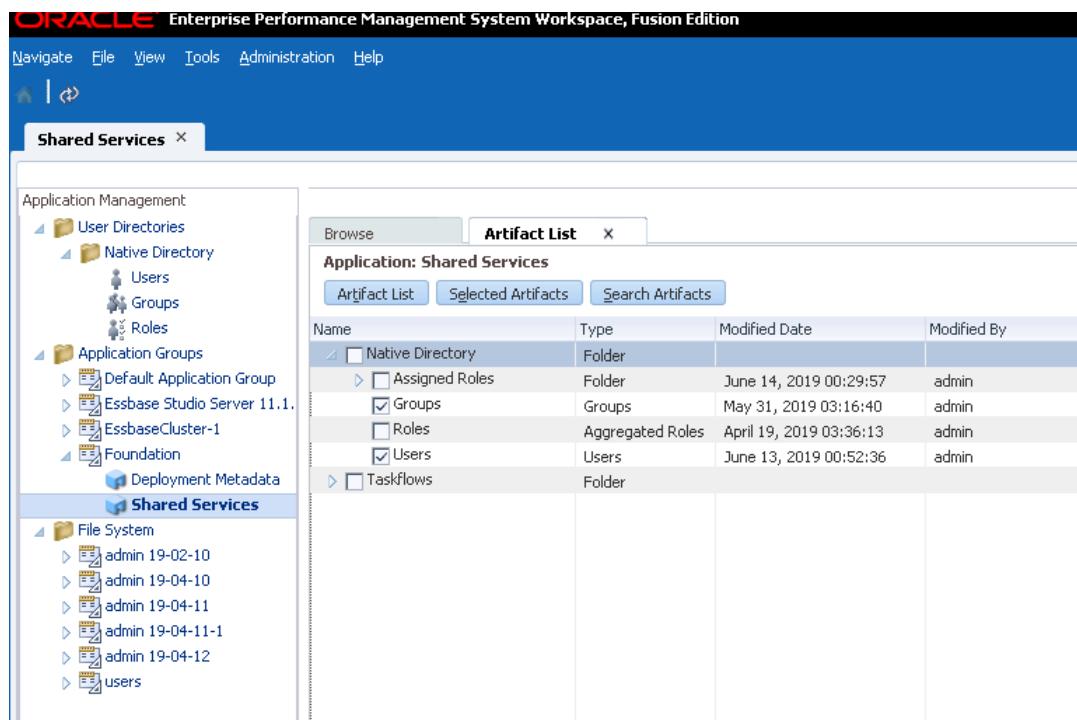
The following access is required:

- 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 11g Export Utility and their corresponding operations: Service Administrator for all applications; Power User for all applications created by the Power User.

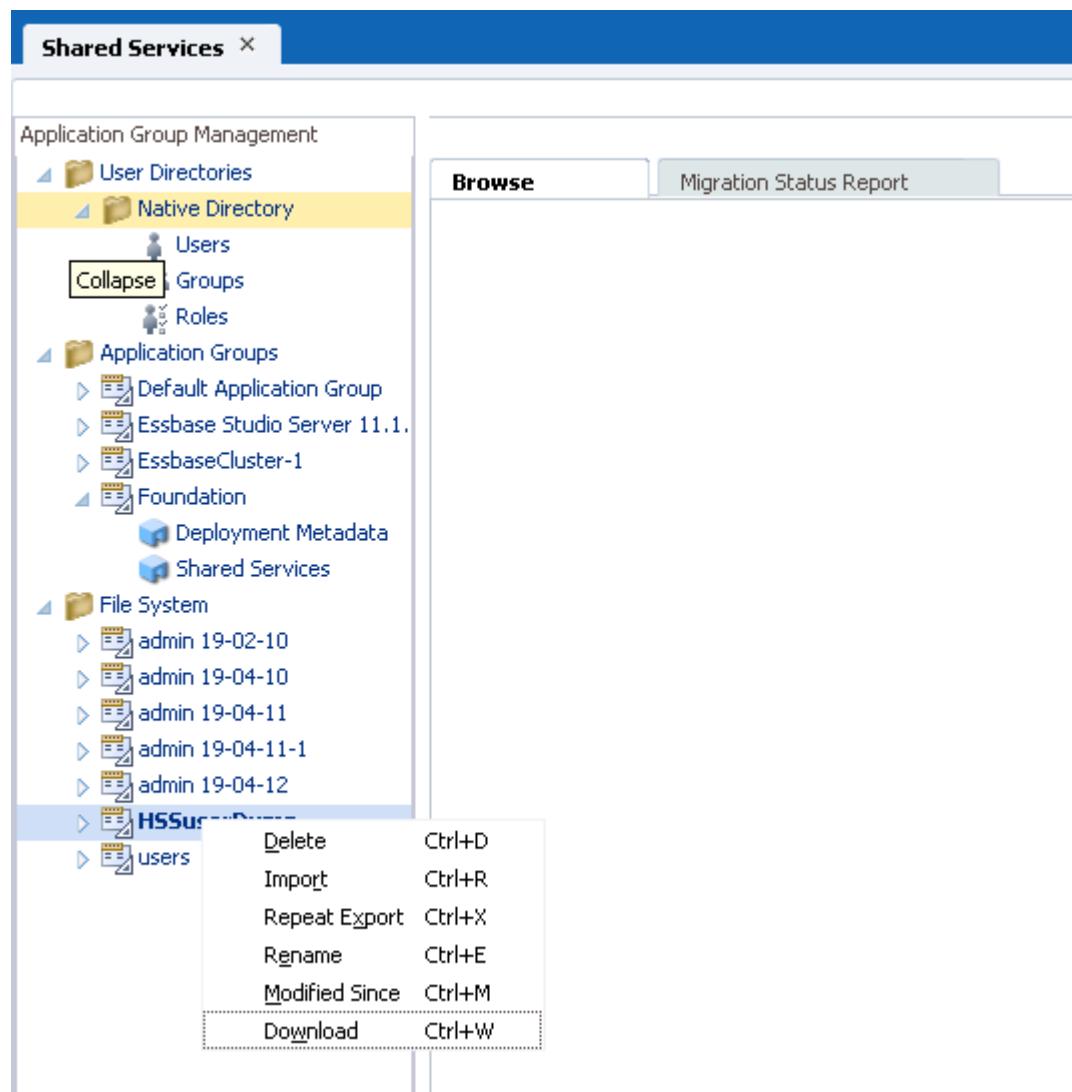
- For importing: A user with at least Power User role can create applications (during import) and manage applications.

- Launch the Shared Services Console. Navigate to **Application Groups, Foundation**, and then **Shared Services**.

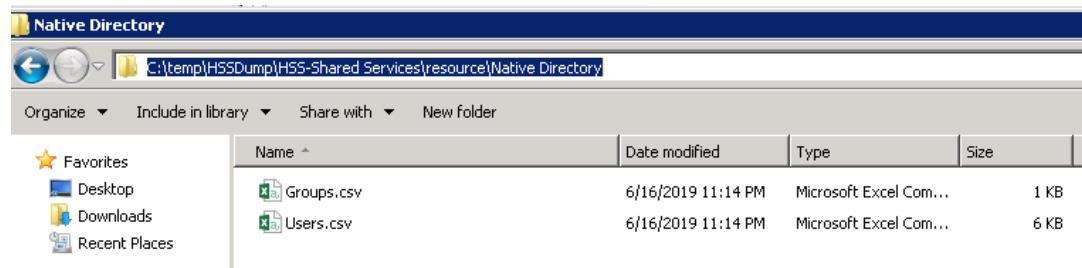


Name	Type	Modified Date	Modified By
Native Directory	Folder		
Assigned Roles	Folder	June 14, 2019 00:29:57	admin
Groups	Groups	May 31, 2019 03:16:40	admin
Roles	Aggregated Roles	April 19, 2019 03:36:13	admin
Users	Users	June 13, 2019 00:52:36	admin
Taskflows	Folder		

2. Select **Groups** and **Users** check boxes, and click **Export**.
3. 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 IAM or IDCS format. You must manually reorganize or map the files to IAM or IDCS format, as shown below, using a CSV or text editor.

Exported users and groups	Shared Services format	IAM or IDCS format
Users.csv	id,provider,login_name,first_name,last_name,description,email,internal_id,password,active	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,Password,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,Department,Manager Name
Groups.csv	id,provider,name,description,internal_id (Following data for each group:) #group_children id,group_id,group_provider,user_id,user_provider	Display Name,Description,User Members

For columns that are empty, enter placeholder text.

7. Import the users and groups to IAM or IDCS 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

You can export applications from versions 11.1.2.3.0nn, 11.1.2.4.0nn, 11.1.2.4.5nn, or 12.2.1, using the 11g LCM Export Utility, and then import them to the target version using the CLI Utility. Note that you can also run LCMIImport from the Essbase Jobs tab to import applications.

This is the workflow to migrate from 11g:

1. **Download the 11g LCM Export Utility:** In the Essbase web interface, click **Console**, expand **Command Line Tools**, and download the 11g LCM Export 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 export parameters:** If `-exportepmroles` option is enabled, you must set the following parameters before you run the LCM export.

- For Linux:

```
export EPM_ORACLE_HOME=/scratch/Oracle/Middleware/EPMSystem11R1
export EPM_ORACLE_INSTANCE=/scratch/Oracle/Middleware/user_projects/
epmsystem1
```

- For Windows:

```
set EPM_ORACLE_HOME=C:\Oracle\Middleware\EPMSystem11R1
set EPM_ORACLE_INSTANCE=C:\Oracle\Middleware\user_projects\epmsystem1
```

3. **Set up the 11g LCM Export Utility:** Before running the utility, you must set and export the environment variables EPM_ORACLE_HOME and EPM_ORACLE_INSTANCE in the shell terminal. These variables must be the same as those used in the source EPM 11g environment. For details, see [About Middleware Home, EPM Oracle Home, and EPM Oracle Instance](#). Also, 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. Also see [11g LCM Export Utility Options](#).
4. **Check the `/tmp` directory:** You may need to change the location of the `tmp` directory. If it is full, the 11g LCM Export Utility may fail.
5. **Run the export:** When you export non-Unicode block storage applications, use `--converttoutf8` option in the LCM export command. When you export non-Unicode aggregate storage applications, manually convert them using the steps in [Convert Non-Unicode Aggregate Storage Application to Unicode Mode..](#)

LCM export syntax: At the command prompt, enter the following command syntax to export one or more applications to a .zip file:

```
EssbaseLCM.bat| .sh export -server essbasehost:port -user username -  
password password -application  
    appname -zipfile zipfilename [-nodata] [-include-server-level] [-  
converttoutf8] [-forceutf8]  
    [-generateartifactlist] [-exportepmroles] [-allApp] [-exportdata]  
    [-cube] [-filetype] [-partitions] [-filters]
```

Specify these options:

- exportepmroles : (optional) Exports Enterprise Performance Management (EPM) roles
- include-server-level : (optional) Includes server-level artifacts, such as server-level substitution variables or server-level roles
- generateartifactlist : (optional) Generates 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 exports additional artifacts: user and group server-level roles, application-level roles, calculations, and filter associations.

6. **Run the import:** To import one or more applications, use the Essbase Command Line Interface Utility (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][-include-server-level][-artifactlist artifactList][-  
restEncryPassword]
```

See: [LcmImport: Restore Cube Files](#).

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 11g LCM Export Utility.

7. Federated partitions are not migrated, so when moving your application and cube to another server or version, you need to delete the federated partition and recreate it in the new environment. See [Federated Cube Maintenance and Troubleshooting](#).
8. **Upgrade aggregate storage outline version:** After import of aggregate storage applications, the outline must be upgraded using ESSCMDQ. See [Upgrade Aggregate Storage Outline Version](#).
9. **Validate:** Log in to the Essbase web interface to see the application and cube on the Applications page.

11g LCM Export Utility Options

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

Syntax

```
EssbaseLCM.bat|.sh export -server essbasehost:port -user username -password  
password -application  
    appname -zipfile zipfilename [-nodata] [-include-server-level] [-  
converttoutf8] [-forceutf8]  
    [-generateartifactlist] [-exportepmroles] [-allApp] [-exportdata]  
    [-cube] [-filetype] [-partitions] [-filters]
```

Notes

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

Command Options and Descriptions

Option	Description
<code>-server <essbasehost:port></code>	Server host name and port number.
<code>-user <username></code>	Server user name.
<code>-password <password></code>	Server password. Skip if you want to be prompted for the password.
<code>-application <appname></code>	Name of application to back up.
<code>-zipfile zipfilename</code>	Optional. Name of compressed file to hold backup files.
<code>-nodata</code>	Optional. Do not include data in the backup.
<code>-overwrite</code>	Optional. Overwrite existing backup file.

Option	Description
-converttouft8	(Optional) Convert block storage application to Unicode. Prompts you to type Y to confirm.
-forceuft8	(Optional) Same as -converttouft8, but without any prompt. Can be used in automation scripts.
-exportepmroles	(Optional) Exports Essbase roles from Enterprise Performance Management (EPM) source.
-generateartiflist	Optional. Generate a text file containing a complete list of the exported artifacts. You can use this text file to manage the import of artifacts. For example, you can rearrange the order of artifacts in the list to control the order in which they are imported. You can skip importing some artifacts by removing or commenting out items in the list.
-include-server-level	(Optional) Include server-level artifacts, such as server-level substitution variables and server-level roles. Requires system administrator role.
-allApp	Optional (and case-sensitive). If used instead of -application, exports all applications to a single zip file. Lcmimport can accept single-application zip files or multiple-application zip files.
-exportdata	Optional. Only export data.
-cube	Optional. Export a single cube. This option can be specified along with the options to export only: data, files of certain types, partitions, or filters.
-filetype	Optional. Only export files of the specified type. Supported file types include OTL (outline), TXT (text), RUL (rule), CSC (calc script), DTR (drill through report definition), and Excel (only .xls files are exported. No .xlsx files are exported).
-partitions	Optional. Only export partition definitions. Lifecycle Management (LCM) import operations (and Migration Utility import) are not supported for migration of federated partitions. Federated partitions must be recreated manually on the target.
-filters	Optional. Only export security filters.

Migrate Essbase to Latest Cloud Version

You can migrate applications from Essbase 19c to Essbase 21c. In addition to upgrading versions, you may also need to migrate applications/cubes from one instance to another. Learn how to prepare for migration, and review some use cases for migrating.

You can use the Essbase Command Line Interface (CLI) to migrate your source application and artifacts.

You can use the Migration Utility to migrate multiple applications, artifacts, users, and groups at one time, from Essbase 19c and higher.

- [Prepare to Migrate Essbase Applications and Users](#)
- [Migrated 21c Artifacts](#)
- [Migrate Applications Using Command Line Interface](#)

- [Migrate Applications Using Migration Utility](#)
- [Migrate from FCCS or PBCS](#)

Prepare to Migrate Essbase Applications and Users

Here are some considerations and requirements when migrating applications to Essbase 21c, whether from independent Essbase deployments, or Essbase on OCI via Marketplace deployments.

Considerations and requirements

- Lifecycle Management (LCM) import and Migration Utility import do not support migration of federated partitions. Federated partitions must be recreated manually on the target, after migration from 21c to 21c.
- You can use the Essbase Command Line Interface (CLI) to migrate your source application and artifacts across deployments and releases.
- In Independent deployments, when you have a large number of applications and using LCM is not feasible to export your applications, see: [Migrate Multiple Essbase Instances to a Single Shared Services Instance](#).
- Export and import using Migration Utility does not support migration of applications, if the Essbase instance is configured in EPM Shared Services mode; you must use CLI with LCM export and import commands.
- 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.
- 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 CLI or Migration Utility, use the OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) 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.
- Solve order applies to dynamic members in the outline and to dynamic calculation execution of dimensions and members. Adjust the solve order of dimensions and members to indicate their calculation priority. Solve order is recommended instead of using Two Pass calculation, because it's more flexible. You can set solve order for dimensions or members, or you can use the default Essbase solve order. The minimum solve order you can set is 0, and the maximum is 127. A higher solve order means the member is calculated later; for example, a member with a solve order of 1 is solved before a member with a solve order of 2. See [Solve Order in Hybrid Mode](#).

- Free-form data exports and imports for cubes with typed measures behaves differently in 21c. For the latest information, see Loading, Clearing, and Exporting Text and Date Measures.

Required user roles

- For exporting: Application Manager for the application created. In addition, the following roles can use LCM commands and CLI: 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 21c Artifacts

The following table describes which global, application-level, and cube-level Essbase artifacts you can migrate between cloud service instances. These artifacts migrate as described from Essbase 21c on either OCI/Marketplace or independent deployments, and from Essbase 19c on OCI/Marketplace.

Artifact	Exceptions/Comments
Application and cube metadata	Application metadata includes application type and settings. Cube metadata includes cube properties and settings.
Application-level configuration files	If these files exist, they're migrated.
Calculation scripts	Application- and cube-level calculations are migrated.
Catalog server	Files in the application directory are migrated. Files stored under <code>shared</code> and <code>users</code> folders aren't migrated. You can manually download them from the web interface and restore them.
Connections and Datasources	Using Migration Utility, system- and application-level connections and Datasources are migrated. Using CLI Utility, connections and Datasources created at the application level are migrated. With both tools, you must include the following argument in <code>lcmexport</code> operations: <code>-include-server-level</code> (or its abbreviation <code>-isl</code>).
Data	To be migrated, data must be in the cube directory on the cloud instance.
Disk volumes	Disk volume definitions aren't applicable to Essbase cloud instances.
Drill through definitions	Drill through definitions are migrated. Do not rename drill through report definitions (.dtr files). Drill through report definitions that are renamed may not be editable and may not work as expected.
Excel workbooks and files	Excel workbooks and files are migrated.
Filters	Cube-level and user-created filters are migrated.
Global variables	-
Layouts	Cube-level layouts are migrated.
Linked Reporting Objects (LROs)	LROs are included for backward compatibility with migrated Essbase 11g On-Premise applications.

Artifact	Exceptions/Comments
Location aliases	Location aliases are migrated with the cube. LCM doesn't support Location alias credentials migration. After you migrate your applications from Essbase 11g On-Premise, you must replace your location aliases. See Location Aliases section in Prepare to Migrate from Essbase 11g On-Premise .
Log files	Log files aren't migrated.
MDX reports	Cube-level named queries are migrated. See Analyze Data with MDX Reports.
Outlines and formulas	Formulas containing @XREF aren't migrated.

 **Note**

To migrate an existing cube outline to Essbase 21c, use a migration tool. See About Migration Tools and Use Cases. Do not copy the .otl file to migrate an outline to 21c.

Partitions	Replicated and transparent partitions are migrated. Only partition definitions from the target cube are exported to the file system. Lifecycle Management (LCM) import and Migration Utility import do not support migration of federated partitions. Federated partitions must be recreated manually on the target, after migration from 21c to 21c.
Report scripts	Application- and cube-level report scripts are migrated. The scripts are included for backward compatibility with migrated Essbase 11g On-Premise applications.
Rule files, text files, .csv files	Application-and cube-level files are migrated.
Scenarios	If a cube is scenario-enabled and has a Sandbox dimension, the related scenarios are migrated.
Substitution variables	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	User roles can be migrated only from one Essbase cloud instance to another.
Wallet files	Wallet files are migrated for the specified application.

Migrate Applications Using Command Line Interface

You can use Command-Line Interface (CLI) to migrate source applications and artifacts across Essbase cloud deployments and releases.

The standard migration workflow using the Command Line Interface (CLI) is as follows:

1. Download the tool and use the `lcmexport` commands to export single or multiple applications from source to a zip file.
2. Use the `lcmimport` command to import single or multiple applications 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 Applications Using Migration Utility

As an Essbase Service Administrator, you can use the Migration Utility (`migrationTools.jar`) to migrate an entire Essbase instance (all applications, users and groups, and other artifacts) from one cloud instance to another in a single process.

You can use Migration Utility to migrate from source applications and elements from Essbase deployments on OCI via Marketplace. The utility migrates multiple applications at one time. It also migrates artifacts, rules, users and groups.

Note that the Essbase Command-line Interface (CLI) commands, `lcmimport` and `lcmexport` do not migrate users and groups.

Before using the Migration Utility, complete the following prerequisite: [Set Up the SSL Certificate](#).

Use Cases for Migration Utility

Here are some use cases for migrating with the Migration Utility.

- Migrate Essbase users from WebLogic LDAP in the source to OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS) in the target.
- Use Migration Utility for basic deployments that aren't customized. If your deployment includes customizations, such as custom single sign-on solutions, use CLI instead of Migration Utility.
- When the source Essbase deployment is from
 - Essbase Marketplace on Oracle Cloud Infrastructure, using OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS), or
 - Essbase Marketplace on Oracle Cloud Infrastructure, using Embedded LDAP

then before using the Migration utility steps below to export, first create a new confidential application as outlined in [Create a Confidential Identity Application](#). Also, before using the Migration Utility to import, change the host and IAM or IDCS details in `import.properties` to point to the target Essbase instance.

Steps to Migrate Cloud Applications and Users using Migration Utility

1. Before you use the Migration Utility, if you haven't already, patch your source Essbase instance to the latest version.
2. 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 web interface, and navigate to the Console.
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.
9. Before you run the Migration Utility export, edit the `export.properties` file, located in the same directory as the utility. Provide the details of the Essbase source instance that you are migrating from.

Export Property	Description
<code>userName</code>	Required. Essbase administrator user name.
<code>password</code>	Required. Essbase administrator password.
<code>host</code>	Required. IP address of the load balancer, if a load balancer was configured during Essbase deployment. Otherwise, the host name or IP address that corresponds with the Essbase public IP address.
<code>port</code>	Optional. Source Essbase instance's HTTP listening port. If source instance is configured in secure (TLS) mode and <code>isTLS=true</code> port 443 is assumed. If source instance is configured with LDAP identity provider, set port as 80 .
<code>isTLS</code>	Optional: true or false . If true, source instance is configured in secure (TLS) mode using the specified port value.
<code>proxy</code>	Optional. HTTP proxy host address, if the connection to the source Essbase instance requires a proxy. Proxy host and proxy port are separated by a colon. If the proxy port is omitted, its value is assumed to be 80. Example with port omitted: <code>www-proxy-example.com</code> . Example with port included: <code>www-proxy-example.com:80</code>
<code>skipdata</code>	Optional. Default is false . Set to true if only schemas of the Essbase application should be exported.
<code>skipUser</code>	Optional: true or false . Set to true only if there is no need to export users and/or groups. Default value is false.
<code>idcsHost</code>	Required only when migrating users and/or groups from an instance deployed on OCI via Marketplace. The source's identity domain host in IAM or IDCS. From the Identity Domain URL (example: <code>https://idcs- <string_of_numbers_and_letters>.identity.oraclecloud.com</code>), extract only the domain part after the string. Example: <code>identity.oraclecloud.com</code>

Export Property	Description
idcsTenant	Required only for migrating users and/or groups from an instance deployed on OCI via Marketplace. The source host's identity domain tenancy identifier in IAM or IDCS. From the Identity Domain URL (example: <a href="https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com">https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com</string_of_numbers_and_letters>), extract only idcs- <string_of_numbers_and_letters>.</string_of_numbers_and_letters>
clientId	Required only for migrating users and/or groups from an instance deployed on OCI via Marketplace. The source host's client identifier for OAuth authorization, found in the integrated confidential identity application.
clientSecret	Required only for migrating users and/or groups from an instance deployed on OCI via Marketplace. The source host's client secret for OAuth authorization, found in the integrated confidential identity application.
appId	Required only for migrating users and/or groups from an instance deployed on OCI via Marketplace. Identifier of the integrated confidential identity application on the OCI identity domain (IAM or IDCS). This identity application should be configured with application roles grantable to users and groups. This property enables granting roles to migrated users and groups.

10. Edit the `import.properties` file, located in the same directory as the utility. Provide the details of the Essbase target instance that you are migrating to.

Import Property	Description
userName	Required. Target instance's Essbase administrator user name.
password	Required. Target instance's Essbase administrator password.
host	Required. Target instance's IP address of the load balancer, if a load balancer was configured during Essbase deployment. Otherwise, the host name or IP address that corresponds with the Essbase public IP address.
port	Optional. Target Essbase instance's HTTP listening port. If target instance is configured in secure (TLS) mode and <code>isTLS=true</code> port 443 is assumed. If source instance is configured with LDAP identity provider, set port as 80 .
isTLS	Optional: true or false . If true, target instance is configured in secure (TLS) mode using the specified port value.
userPassword	Optional. Initial password to be assigned to any new users the utility migrates to the integrated identity application in IAM or IDCS. To change the passwords after migration, use the OCI Console.
proxy	Optional. HTTP proxy host address, if the connection to the target Essbase instance requires a proxy. Proxy host and proxy port are separated by a colon. If the proxy port is omitted, its value is assumed to be 80. Example with port omitted: <code>www-proxy-example.com</code> . Example with port included: <code>www-proxy-example.com:80</code>
skipUser	Optional: true or false . Set to true only if there is no need to import users and/or groups. Default value is false .
idcsHost	Required only for migrating users and/or groups to an instance deployed on OCI via Marketplace. Target identity domain host. From the Identity Domain URL (example: <a href="https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com">https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com</string_of_numbers_and_letters>), extract only the domain part after the string. Example: <code>identity.oraclecloud.com</code>
idcsTenant	Required only for migrating users and/or groups to an instance deployed on OCI via Marketplace. The target host's identity domain tenancy identifier. From the Identity Domain URL (example: <a href="https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com">https://idcs-<string_of_numbers_and_letters>.identity.oraclecloud.com</string_of_numbers_and_letters>), extract only idcs- <string_of_numbers_and_letters>.</string_of_numbers_and_letters>

Import Property	Description
clientId	Required only for migrating users and/or groups to an instance deployed on OCI via Marketplace. The target host's client identifier for OAuth authorization, found in the integrated confidential identity application.
clientSecret	Required only for migrating users and/or groups to an instance deployed on OCI via Marketplace. The target host's client secret for OAuth authorization, found in the integrated confidential identity application.
appId	Required only for migrating users and/or groups to an instance deployed on OCI via Marketplace. Identifier of the integrated confidential identity application on the OCI identity domain (IAM or IDCS). This identity application should be configured with application roles grantable to users and groups. This property enables granting roles to migrated users and groups.
overwrite	Optional. Set to true if applications to be imported already exist on the target, and you want to overwrite them. Default value is false .
certificates	Optional. Full path to all trust certificate files for all TLS/SSL servers. For Windows, escape backslashes with another backslash. Example: C:\\migrationTools\\\\all_certs.pem

11. Run Migration Utility in Java with the export command to export all applications, users, and groups from the Essbase source instance catalog to a compressed file.

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

Example:

```
java -jar -Dhttps.proxyHost=www-proxy-abcdef.example.com -
Dhttps.proxyPort=80 migrationTools.jar export export.properties
export.tar.gz
```

12. Run Migration Utility in Java with the import command to import the compressed file to the target instance.

```
java -jar -Dhttps.proxyHost=<proxy-url> -Dhttps.proxyPort=<nn>
migrationTools.jar import import.properties <existing_tar_file>
```

13. After you run the import, the data is migrated to the Essbase catalog of the target instance.

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](#).

1. 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 Interface (CLI) Utility to import the Essbase application and cubes from the exported .zip file using CLI command lcmimport.

Post-Migration and Advanced Topics

After migrating or upgrading to Essbase 21c, test that the applications are working as expected. You may need to take additional steps to import connections and artifacts. Aggregate storage outlines may require post-migration upgrade steps.

- [Selective and Ordered Import of Artifacts](#)
- [Selective Export of Artifacts](#)
- [Test the Migrated Essbase Instance](#)
- [Post Upgrade Tasks for CLI](#)
- [Upgrade Aggregate Storage Outline Version](#)
- [Export Essbase 11g On-Premise Cubes](#)

Selective Export of Artifacts

You can control export by selectively exporting individual Essbase artifacts for cloud service migrations, using the CLI utility.

You can export individual artifacts, for example, exporting only rules, or outlines, or data, and so on. You can also use this process for periodic backup of individual artifacts.

Selective Migration, Backup, and Export

You can selectively export the following:

- Cube - Export just a single cube from an application.
- Essbase Files - Export particular files, such as Outline, Rule files, Drill Through reports, and others.
- Data - Export only data. This can be useful for periodic backup of data.
- Partitions - Export only partitions.
- Filters - Export only security filters of a cube.

Supported Essbase File Types

The following Essbase file types are currently supported for selective export.

- OTL
- TXT
- RUL
- CSC
- DTR
- EXCEL

CLI Utility Command Options

The following command options have been added or changed to support selective export.

-c or --cube <cube name>

Exports a single cube. This option can be used by itself or combined with one of the following options.

-d or --exportdata	Exports data only.
-ft or --filetype <filetype>	Export files of the specified type. File type keywords, such as OTL or RUL, are case-sensitive and must be entered in upper case.
-ep or --exportpartitions	Export partitions only.
-ef or --exportfilters	Export security filters only.

Examples

- Export data only, from a single cube, for example, Demo.Basic:

```
./esscs.sh lcmexport -a Demo -c Basic -z data.zip -d
```

- Export a single cube only:

```
./esscs.sh lcmexport -a Demo -z data.zip -c basic
```

- Export only text files, from all cubes under Demo:

```
./esscs.sh lcmexport -a Demo -z data.zip -ft txt
```

- Export only partitions, from all cubes under Demo:

```
./esscs.sh lcmexport -a Demo -z data.zip -ep
```

- Export only the data from all of the applications:

```
./esscs.sh lcmexport -aa -z all_data.zip -d
```

- Export only outlines, from all cubes in an instance:

```
./esscs.sh lcmexport -aa -z cubes.zip -ft OTL
```

- Export all partitions of an instance:

```
./esscs.sh lcmexport -aa -z all_partitions.zip -ep
```

- Export all security filters of a single cube:

```
./esscs.sh lcmexport -a Sample -z all_partitions.zip -c Basic -ef
```

Selective and Ordered Import of Artifacts

You can control import of Essbase artifacts using a selection list text file, for 11g migrations (using the 11g LCM Export Utility) and for cloud service migrations (using the CLI utility).

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 lcmexport 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

```
@Provisions
/Sample/Provisions/CalcAssociation.csv

@Databases/Basic/Calc_scripts
/Sample/Databases/Basic/Calc scripts/Default Calc
/Sample/Databases/Basic/Calc scripts/CalcAll.csc

# -----IMPORT-----
import @Provisions
import @Databases/Basic/Calc_scripts
# -----IMPORT-----
```

How to use this feature

- When we use 11g LCM Export Utility, 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.

Test the Migrated Essbase Instance

After migrating your Essbase instance, 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 data loads and dimension builds work as expected in migrated applications.
- 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.

Post Upgrade Tasks for CLI

After migrating or upgrading to Essbase 21c (Release 21.4 or higher), users of CLI must take the following actions, if applicable to their usage.

The Service Administrator needs to recreate any saved local connections that were created using the `createlocalconnection` command.

All affected CLI users must reset any stored passwords they created using the `setpassword` command.

Upgrade Aggregate Storage Outline Version

These are the steps, for Linux and for Windows, to upgrade an aggregate storage outline to Essbase 21c.

Note

Other references for ESSCMDQ workaround information are: [How to Use ESSCMDQ to Compact Outlines \(Doc ID 1534496.1\)](#) and [Essbase 21c Compress Dimension Option Not Visible in JET UI \(Doc ID 2853804.1\)](#).

Server-based upgrade steps - for Linux

1. Note that these steps can be performed only after importing the outline.
 - a. 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 ESSCMDQ:

```
./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
```

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

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

g. Create cube directory under the database directory.

```
mkdir $ARBORPATH/client/{appname}/{cubename}/cube
```

2. After launching this script:

```
<Essbase_Product_Home>/products/Essbase/EssbaseServer/bin/startESSCMDQ.sh
```

execute the following commands in ESSCMDQ:

```
#Login to Essbase 21c instance
login hostname username password;
#Download outline to client directory location.
qgetObject3 2 1 "appname" "cubename" "outlinename" $ARBORPATH/client/
{appname}/{cubename}/cube/{outlinename}.otl;
#Specify aggregate storage appname and cubename below
select "appname" "cubename";

#Update outline version
openotl 1 1 "appname" "cubename" "outlinename" "y" "y" 0;
setopgversion 0 "111241";
writeotl 0 "appname" "cubename" "outlinename"
lockobj 1 "appname" "cubename" "outlinename";
qPutObject3 2 1 "appname" "cubename" "outlinename" $ARBORPATH/client/
{appname}/{cubename}/cube/{outlinename}.otl N;
restructotl 1;
closeotl 0;
unlockobj 1 "appname" "cubename" "outlinename";
```

Server-based upgrade steps - for Windows

1. Note that these steps can be performed only after importing the outline.
 - a. 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.bat
```

as ESSCMDQ:

```
.\Oracle\domains\esscs\esstools\bin\startESSCMDQ.bat
```

Within this newly created script, change the call from:

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMD.bat
```

to

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMDQ.bat
```

d. Make a copy of the script:

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMD.bat
```

as:

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMDQ.bat
```

e. Edit

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMDQ.bat
```

and change the last line from

```
"%ESSBASEPATH%\bin\ESSCMD.exe" %*
```

to:

```
"%ESSBASEPATH%\bin\ESSCMDQ.exe" %*
```

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

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

g. Create cube directory under the database directory.

```
mkdir %ARBORPATH%\client\{appname}\{cubename}\cube
```

2. After launching this script:

```
<Essbase_Product_Home>\products\Essbase\EssbaseServer\bin\startESSCMDQ.bat
```

execute the following commands in ESSCMDQ:

```
#Login to Essbase 21c instance  
login hostname username password;  
#Download outline to client directory location.q  
qgetobject3 2 1 "appname" "cubename" "outlinename" %ARBORPATH%\client\
```

```
{appname}/{cubename}/cube/{outlinename}.otl;
#Specify aggregate storage appname and cubename below
select "appname" "cubename";

#Update outline version
openotl 1 1 "appname" "cubename" "outlinename" "y" "y" 0;
setopgversion 0 "111241";
writeotl 0 "appname" "cubename" "outlinename"
lockobj 1 "appname" "cubename" "outlinename";
qPutObject3 2 1 "appname" "cubename" "outlinename" $ARBORPATH/client/
{appname}/{cubename}/cube/{outlinename}.otl N;
restructotl 1;
closeotl 0;
unlockobj 1 "appname" "cubename" "outlinename";
```

Note

1. If you get the following error while running writeotl, it can be ignored.

```
Writeotl:
Error: Invalid location => 0
sts: 0
```

2. While executing “qgetobject3” command, you must provide the full path, which can be downloaded from server (Linux syntax).

```
qgetobject3 2 1 asos Sample Sample /scratch/12cesscs/data/client/
asos/Sample.otl
```

Export Essbase 11g On-Premise Cubes

If you have applications and cubes that were created in a supported on-premises instance of Essbase, then you can use the 11g Excel 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 11g Excel 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 11g Excel Export Utility](#)
- [Review Member Names Before you Import an Application Workbook Created by the 11g Excel Export Utility](#)

Download the 11g Excel Export Utility

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

To download the 11g Excel 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 **11g Excel Export Utility**.
3. Save the 11g Excel Export Utility, which is named `dbxtool.zip`, to a local drive.

Review Member Names Before you Import an Application Workbook Created by the 11g Excel Export Utility

When importing an application workbook that was created using the 11g Excel 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 `err_DimName.txt` or `err_Dim_DimName.txt`. For example, if the Year dimension has any rejected member names, the error file name is `err_Year.txt` or `err_Dim_Year.txt`.

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.

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

To provide Essbase access to users, you must first assign an Essbase role to the users from your identity provider, and then assign them application-level permissions.

Access to Essbase is restricted by security, and user accounts are managed in the identity domain, using OCI Identity and Access Management (IAM) or Oracle Identity Cloud Service (IDCS).

The following are common use cases for assigning Essbase 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.

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	Has same permission as User Role, with added ability to create applications and cubes. Has Application Manager permission for the applications and cubes this user created, as well as the ability to delete them. Any additional permission must be granted, the same as for User Role.
User Role	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 unless that permission is granted at the application level.

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, load, update, and clear data values based on assigned scope. Ability to create and delete scenarios. Ability to run provisioned calculation scripts.
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

To provision application roles using the Essbase web interface,

1. On the Home page, select an application, and then click **Permissions**.
2. On the **Permission** page, click **Add** 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**.

Back Up and Restore Essbase

Let's explore the process for backing up and restoring Essbase, both individual applications and Essbase instances, on Oracle Cloud Infrastructure.

Topics

- [About Backup and Restore](#)
- [Back Up and Restore Applications](#)
- [Back Up and Restore an Essbase Instance](#)

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 LCM 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.

You can use the backup and restore process in this chapter for your disaster recovery strategy and solution. See also [Essbase on OCI Backup and Restore](#), a reference paper on disaster recovery and version upgrade planning, particularly for Essbase instances deployed on OCI.

Back Up and Restore Applications

Routine backup of Essbase applications ensures that you can recover from any sort of disruption to applications 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.

 Note

Use LCM export option `--application` to export a single application to a zip file. Use option (case-sensitive) `--allApp` (or `-aa`), instead of `--application`, to export all applications to a single zip file. The `--allApp` option does not export users and groups.

By default, LCM export exports your 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:

1. 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 lcmexport`)
2. 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 `EssbaseLCMUtility.zip` 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.

ⓘ Note

Location aliases are migrated with the cube. LCM doesn't support Location alias credentials migration. After you migrate your applications from 11g you must replace your location aliases. See Location Aliases section in [Prepare to Migrate from Essbase 11g On-Premise](#).

Back Up Cube Files Using LCM

Use one of these tools to initiate an LCM Export operation, which backs up application and cube artifacts to a Lifecycle Management (LCM) .zip file:

 ⓘ Note

This topic applies if you are backing an application from a release after 11g. If your application is on 11g, you must first [migrate](#) that application using the 11g Export utility, before you can back it up. For backups, you use one of the tools listed in this topic to initiate the LCM Export operation. To migrate from 11g, you must use the 11g Export utility.

Requires at least user role with Application Manager permission, or, you must be the power user who created the application.

- Command-line interface (CLI): `LcmExport`
- Essbase web interface Export LCM job
- REST API: Execute Job operation (using jobType `Icmexport`)

Restore Cube Files Using LCM

Use one of these tools to initiate an LCM Import operation, which restores application and cube artifacts from a Lifecycle Management (LCM) .zip file:

- Command-line interface (CLI): `LcmImport`
- Essbase web interface Import LCM job
- REST API: Execute Job operation (using jobType `Icmimport`)

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/config. (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 21.1 and Prior

There are two backup/restore use cases for Essbase version 21.1 and older. Your choice of schema database and the number of Essbase instances you have deployed in it defines how you will need to take backups.

Note

The steps in Oracle-Scripted and Non-Oracle Scripted Backup processes, mentioned in this documentation are only applicable for Autonomous AI Transaction Processing database.

[Oracle-Scripted Backup and Restore](#):

When you deploy an Essbase stack, the default behavior is to deploy a new Autonomous AI Transaction Processing database as part of the stack. If you have a single Essbase stack with schemas deployed into a single Autonomous AI Transaction Processing database, you can take your backups using the scripts provided with Essbase 19c or 21.1.

[Non-Oracle-Scripted Backup and Restore](#):

If you did not use an Autonomous AI Transaction Processing database or you have deployed multiple Essbase stacks into a single relational database, you cannot use the backup scripts that come with Essbase 19c or 21.1. Instead, you will back up the data block volume and Essbase database schema as instructed in this chapter.

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

- Autonomous AI Transaction Processing is the relational database into which Essbase schemas are deployed.
- OCI Identity and Access Management (IAM) or 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-IAM or 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 IAM or IDCS user with an Essbase system administrator role.

Oracle-Scripted Backup and Restore

When a single Essbase instance is deployed to an Autonomous AI Transaction Processing database, you can use scripts provided by Oracle to back up the entire database and the Essbase block volumes.

- [Oracle-Scripted Backup of an Essbase Instance](#)
- [Restore an Essbase Instance from an Oracle-Scripted Backup](#)

Oracle-Scripted Backup of an Essbase Instance

These 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 and config block volumes.
- 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 a single instance deployed using a single Autonomous AI Transaction Processing database, schedule the backup for a convenient time when users are not in the system and follow these steps:

Note

If Essbase node has no Public IP, use the Bastion as a proxy.

1. Ensure that the required [policies](#) to manage backups are in place.
2. Using the Oracle Cloud Infrastructure console main menu, navigate to your source compute and record the Internal FQDN and Hostname from the compute's Instance Details page. Should your source compute be unreachable at the time you restore, you will need this information.
3. ssh to your Essbase compute (as opc user).

① Note

If a Bastion is used, you may need to extend the default session length of 3 hours to meet the backup time needed for your quantity of data, or run the backup in background.

4. Run `sudo su - oracle`.
5. Run `cd /u01/vmtools`.
6. Run the script to configure your database to work with object storage. This is a one-time action.

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

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

- 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 .
- ii. Click on the user.
- iii. Copy the profile name at the top of the page.
- 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.

7. Run the backup script.

This backs up the data and config block volumes.

```
./backup.sh
```

① Note

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

Restore an Essbase Instance from an Oracle-Scripted Backup

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 your backup was taken using the scripts provided by Oracle, you can easily recover from disaster. Simply use the following steps to restore your entire database and your block volumes to a new target instance.

Note

If Essbase node has no Public IP, use the Bastion as a proxy.

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 the following exceptions:

- Skip steps 1, 2, and 11.
- 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.

1. Deploy a target Essbase stack using Oracle Marketplace.
 - Use the source Oracle Identity Cloud Service confidential application.
 - Use the source Autonomous AI Transaction Processing database and password.
 - Use the source virtual cloud network and application subnet.
 - If your source stack has a load balancer, do not deploy a target load balancer. You can change the backend set after deploying the target stack.
 - 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.
 - b. If your source stack did not have a load balancer, update your Oracle Identity Cloud Service Confidential Application with the new Redirect and Post Logout Redirect URLs with the target IP address.
Login to Oracle Identity Cloud Service and edit the source confidential application to ensure that the target IP address is substituted for the source IP address that was previously used.
3. [Stop](#) the target Essbase services first (as oracle user) and then stop the WebLogic services and the Node Manager (as opc user). Do not stop the Essbase compute in Oracle Cloud Infrastructure.

① Note

If you are simulating recovery, make sure you also stop the source compute's Essbase services.

4. **Restore** the source Autonomous AI Transaction Processing database from backup. Be sure to select a source Autonomous AI Transaction Processing backup that was taken during a time when your source Essbase services were stopped; also, be sure that you have a source Essbase block volume backup from the same time.
After the restore finishes, audit your data using a database client like SQL Developer. For example, you can look at the ESSBASE_APPLICATION table within the <targetprefix>_ESSBASE schema to verify the restored applications.
5. Temporarily disable /etc/fstab config and data block volume entries.
 - a. ssh to target compute (as opc user).
 - b. sudo vi /etc/fstab
 - c. Insert a # in front of the /u01/config and /u01/data entries.
 - d. Save the file.
6. Detach data and config block volumes from the target Essbase compute. Note the iSCSI caution and be sure to unmount and disconnect each volume before detaching using the OCI console.
 - a. To unmount:
 - i. ssh to target the compute (as opc user).
 - ii. lsblk
 - iii. sudo umount /u01/config
 - iv. Repeat these steps to unmount the data block volume.
 - b. To disconnect iSCSI:
 - i. In the OCI console, select the target compute.
 - ii. Select **resources > attached block volumes**.
 - iii. From the Actions menu ⋮ for the config block volume, select **iSCSI Commands & Information**.
 - iv. Copy iSCSI commands for disconnect.
 - v. ssh to the target compute (as opc user).
 - vi. Paste the disconnect commands you copied and press enter.
 - vii. Repeat these steps to disconnect the data volume.
 - c. To detach:
 - i. In the OCI console, select the target compute.
 - ii. Select **resources > attached block volumes**.
 - iii. From the Actions menu ⋮ for the config block volume select **Detach**.
 - iv. Repeat these steps to detach the data block volume.

7. Restore data and config block volumes from source block volume backup. Be sure to select the same Availability Domain as your target compute.
8. Attach block volumes created in the previous step to the target compute. The volume Attachment Type should be iSCSI.
9. Use the iSCSI commands listed in the OCI console to connect your newly attached target block volumes.
10. Update `/etc/fstab` entries and mount the new block volumes.
 - 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 two newly attached volumes.
 - d. `sudo vi /etc/fstab`
 - e. Uncomment the config and data block volume entries.
 - f. Replace the UUID in the existing config and data volume entries with the UUID of the newly attached data and config volumes. Be sure not to change the mount points `/u01/config` and `/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.
11. Manage hosts:
 - a. As `opc` user, update the target compute `/etc/hosts` file to map the source domain information to the target.
Determine the fully qualified domain name of your source compute:
 - i. Using the OCI console main menu, navigate to your source compute and record the Internal FQDN and Hostname from the compute's Instance Details page. If the source compute is not accessible, retrieve the record you saved during backup. See step 2 in [Oracle-Scripted Backup of an Essbase Instance](#).
 - ii. Edit the `/etc/hosts` file on the target compute and append the source domain information to the entry containing the target compute.
12. `sudo vi /etc/hosts`

The entry will have the format:

`<Target Compute IP> <target compute FQDN> <target compute hostname>`
`<source compute FQDN> <source compute hostname>`

 - b. As `opc` user, update the target compute `/etc/oci-hostname.conf` file. Set the `PRESERVE_HOSTINFO` setting to 3.
`sudo vi /etc/oci-hostname.conf`
12. Start the node manager (as `opc` user) first, and then start the Essbase services (as `oracle` user).

(i) Note

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

Non-Oracle-Scripted Backup and Restore

If you did not use Autonomous AI Transaction Processing for your RCU schemas or have multiple Essbase instances deployed with a single Autonomous AI Transaction Processing database, you cannot use the Oracle-provided backup scripts. You will need to individually back up the <Essbase prefix>_Essbase schema and data block volume for each instance.

- [Non-Oracle-Scripted Backup of an Essbase Instance](#)
- [Restore an Essbase Instance from a Non-Oracle-Scripted Backup](#)

Non-Oracle-Scripted Backup of 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.

Back up the Essbase Schema Using Data Pump and Oracle Cloud Infrastructure Console:

1. Configure your Autonomous AI Transaction Processing instance to work with your object storage.
Use a database client application to configure your Autonomous AI Transaction Processing instance for use with your Oracle Cloud Infrastructure account and a default object storage bucket. SQL Developer is a nice choice for a database client, because it allows Cloud Wallet connections and can connect to multiple Autonomous AI Transaction Processing instances at the same time, if needed.
 - a. Create a connection to your Autonomous AI Transaction Processing instance using SQL Developer (consider your proxy needs if connected via a corporate network).
 - b. Create an Auth Token using the Oracle Cloud Infrastructure Console.
Copy and record the token value in a secure location. You will not be able to retrieve it again.

(i) Note

Previously, after this step, backups were manually configured for Autonomous Database from the OCI console. This manual operation is no longer in use.

2. [Stop](#) the Essbase services. You do not need to stop the node manager. (Do not stop the Essbase compute in Oracle Cloud Infrastructure.)
3. Back up the Essbase schema from your Autonomous AI Transaction Processing instance using Data Pump.
Each Essbase instance has nine associated schemas in your Autonomous AI Transaction Processing database. All nine schemas have a common `rcu_schema_prefix`, which is reported in the outputs of the Oracle Resource Manager (ORM) apply job. When using Oracle Identity Cloud Service for security, you only need to back up the <prefix>_ESSBASE schema that corresponds to the Essbase instance you want to back up. Remember that your Autonomous AI Transaction Processing instance may have Essbase schemas from multiple instances.

- a. Make sure your TNS_ADMIN variable points to the wallet location of the Autonomous AI Transaction Processing instance from which you are exporting. Use Oracle Instant Client to issue the following Data Pump command:

```
$ expdp admin_<database name>_low directory=data_pump_dir  
schemas=<source schema prefix>_ESSBASE logfile=<logfile name>.out  
dumpfile=<dump file name>.dmp
```

 **Note**

The Autonomous AI Transaction Processing prefix in the connection information may not be the same as the schema prefix for the Essbase schema that you are backing up if you have deployed multiple Essbase instances to the same Autonomous AI Transaction Processing database.

- b. Schema backups are physically stored on disk in the Data Pump directory of your Autonomous AI Transaction Processing database.
4. Use the PUT OBJECT procedure to move the .dmp file to Oracle Cloud Infrastructure object storage.
 - a. Move the .dmp file into an Object Storage bucket of your choosing (use SQL Developer, unless you installed the Instant Client SQL*Plus package). The /o/xxxxxxxxxx.dmp portion of the Put Object uri indicates the name you want to assign for the .dmp file in your Object Storage. The file_name must match the .dmp filename you assigned when you created the export on disk using data pump.

 **Note**

The object storage bucket name is case sensitive.

```
BEGIN DBMS_CLOUD.PUT_OBJECT(credential_name => <your credential name>,  
object_uri => <your object uri>, directory_name => <your data pump  
directory name>, file_name => <your data pump export file name>); END;/
```

- b. Refresh your Object Storage Bucket to see the .dmp file.
5. Back up the Essbase data volume. Make sure you back up the block volume attached to the Essbase instance for which you just exported the Essbase schema. See Overview of Block Volume Backups and Backing Up a Volume. For step 2 under Using the Console in Backing Up a Volume, instead of selecting the block volume, use the  menu and select **Create Manual Backup**.

 **Note**

Block Volume Backups are snapshots, so you can restart your Essbase services as soon as the Autonomous AI Transaction Processing schema backup is complete and the block volume backup has been initiated. You don't need to wait for the block volume backup to complete before restarting your Essbase services.

6. Take note of the timestamp of your Autonomous AI Transaction Processing .dmp file in object storage and block volume backups. Because your Essbase services were stopped, these backups can be used to consistently restore Essbase if the need arises.
7. [Start](#) the Essbase services.

Restore an Essbase Instance from a Non-Oracle-Scripted Backup

If you have experienced a disaster, you'll need to deploy a new target Essbase instance before you can restore a non-Oracle-scripted backup. After deploying a new target instance, you can restore the Essbase schema and data block volume from backup. The new target instance should be the same version that was deployed on your failed compute. If you haven't upgraded recently, you may need to retrieve an image from [GitHub](#). After your new target instance is deployed, you can recover from backup using the target.

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

- Skip steps 1 and 2.
- Do not include the `REMAP_SCHEMA=<sourceEBprefix>_ESSBASE:<targetEBprefix>_ESSBASE` parameter in step 4b.
- 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 AI Transaction Processing database. These steps are also appropriate if you did not use Autonomous AI Transaction Processing for your database schemas.

Note

If your Essbase node has no Public IP, use the Bastion 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 AI Transaction Processing database and password. Optionally, create a new Autonomous AI Transaction Processing 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.
 - 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.
 - b. If your source stack did not have a load balancer, update your Oracle Identity Cloud Service Confidential Application with the new Redirect and Post Logout Redirect URLs with the target IP address.
Login to Oracle Identity Cloud Service and edit the source confidential application to ensure that the target IP address is substituted for the source IP address that was previously used.
3. [Stop](#) the target Essbase services (as oracle user). Do not stop the node manager or the Essbase compute in Oracle Cloud Infrastructure.

 **Note**

We assume that the Essbase services on the source compute are stopped, because this use case involves source compute server or hardware failure. If you are simulating these steps, make sure you also stop the source compute's Essbase services.

4. Restore the target database schema from source schema backup.
When restoring the Autonomous AI Transaction Processing database for your target stack, you will import your source schema backup into the target Essbase schema using the REMAP_SCHEMA option.
Be sure to select a source schema backup that was taken during a time when your source Essbase services were stopped. Also, be sure that you have the source Essbase data block volume from the same time.

 **Note**

If you used the same Autonomous AI Transaction Processing database for your target instance as you had used for your source instance, then it is already configured for use with object storage. If you created a new Autonomous AI Transaction Processing database, then you will need to configure it for use with object storage. See step 1. in the **Non-Oracle-Scripted Backup of an Essbase Instance** section above.

- a. Make sure your Instant Client is configured to point to the Autonomous AI Transaction Processing 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=data_pump_dir  
credential=<yourcredname>dumpfile=<your object storage native URI>
```

```
REMAP_SCHEMA=<sourceEBprefix>_ESSBASE:<targetEBprefix>_ESSBASE
table_exists_action=replace
```

 **Note**

The object (/o/) in your Object Storage Native URI will be the .dmp file name of your source backup in your object storage.

5. 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.
6. 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.
7. 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 OCI 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 OCI console, select the target compute.
 - ii. Select **resources > attached block volumes**.
 - iii. From the Actions menu  for the data block volume select **iSCSI Commands & Information**.
 - c. 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 OCI console, select the target compute.
 - ii. Select **resources > attached block volumes**.
 - iii. From the Actions menu  for the data block volume select **Detach**.
8. Restore the source data block volume from backup. Be sure to select the same Availability Domain as your target compute.
9. Attach the block volume created in the previous step to the target compute. The volume Attachment Type should be iSCSI.
10. Use the iSCSI commands listed in the OCI console to connect your newly attached target data block volume.

11. 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 data volume.
 - c. `sudo blkid` and record the UUID of the newly attached data volume.
 - d. `sudo vi /etc/fstab`
 - e. Uncomment the data block volume entry.
 - 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.
12. [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 block volumes and backups.

Back Up and Restore: Version 21.2 and Later

In releases 21.2 and later, 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 AI Database.

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

- [Oracle-Scripted Backup of an Essbase Instance](#)
- [Restore an Essbase Instance Using an Oracle-Scripted Backup](#)

 **Note**

If you didn't deploy with Autonomous AI Database, you cannot use the Oracle scripts to back up your Essbase instance. Instead, follow the same instructions provided for prior releases in [Non-Oracle-Scripted Backup and Restore](#).

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

- Autonomous AI Transaction Processing is the relational database into which Essbase schemas are deployed.
- OCI Identity and Access Management (IAM) or 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-IAM or 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 IAM or IDCS 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, gradually disconnect active user sessions, using MaxL alter application disable commands and/or connects (to prevent any new user activity), followed by alter system logout session and/or kill request (if you need to terminate any active sessions that don't need to complete). If you use disable commands in the application, remember to re-enable commands for that application after performing the backup.

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

 **Note**

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

1. Ensure that the required [policies](#) to manage backups are in place.
2. ssh to your Essbase compute (as **opc** user).
3. Switch to the **oracle** user.

```
sudo su oracle
```

4. Change directories to `/u01/vmtools/backup/`.

```
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:

ⓘ Note

As the Oracle Cloud Infrastructure user must be able to configure Autonomous AI Database to work with object storage, the OCI user requires these policies:

```
Allow group group_name to use buckets in compartment compartment_name
Allow group group_name to manage objects in compartment compartment_name
```

Refer also to [Set Up Policies](#).

To find the OCI username:

- i. In Oracle Cloud Infrastructure, go to the upper right hand corner and click the profile menu icon .
- 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.

6. Run the backup script, `run-backup.sh`. This script also backs up the data block volume.

⚠ Caution

This backup script does not back up data and files in object storage if the stack catalog was configured with object storage in 21.4 or higher. You must take a backup of the data and files.

You are prompted for the database admin password if the vault option is not provided. Enter the clear text password. Because the password is protected information, you will not see the text as you enter it at the command prompt.

The syntax is as follows:

```
./run-backup.sh [--vault | -v] [--cpu n | -C n] [--parallel n]
```

Example:

```
./run-backup.sh --vault
```

The script has the following options:

`--vault` or `-v` sets the script to take the required credentials (admin password) stored in the vault, accessed using the OCID, instead of prompting you for the password.

--parallel sets the degree of parallelizing for the backup. The default is to use all available parallel processes. This parameter is supported only for stacks created in Essbase 21.4+.

--cpu or -c also sets the degree of parallelizing for the backup; this parameter is shown here only for backward compatibility.

If option --vault or -v is not used, you are prompted to enter the database admin (clear text) password. The protected password is not displayed as you enter it.

 **Note**

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

7. Capture the 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 Using an Oracle-Scripted Backup

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 AI Database.

 **Note**

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 AI Database and password. Optionally, create a new Autonomous AI 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 IAM or IDCS Essbase admin user in the target stack as you used in the source stack. If this is not possible, make sure the source Essbase instance has at least one valid IAM or IDCS user with the Essbase system administrator role. After you restore, you must log in to the target instance as a valid IAM or IDCS user who had Essbase system administrator role on the source instance.

2. Manage the target instance log in 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, log in 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.
 - b. If your source stack did not have a load balancer, update your IAM or IDCS Confidential Application Redirect and Post Logout Redirect URLs with the target IP address.
3. ssh to your target Essbase compute (as opc user).
4. Switch to the **oracle** user.

```
sudo su oracle
```

5. Change directories to /u01/vmtools/backup/.

```
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-storage script for the target.

```
./configure-backup-storage.sh [--vault | -V]
```

where `--vault` or `-V` option will pull the required credentials from the vault instead of prompting for the password. By specifying this, the script will automatically take the password / OCID secret stored in the vault. The vault is identified to the instance from the metadata.

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 .
- 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 AI 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 AI 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
exclude=TABLE: "IN( 'ESSBASE_SCHEMA_VERSION' )"
```

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 such as 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 the 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  for the data block volume select **iSCSI Commands & Information**.
 - c. 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 volume using the following OCI CLI command (as **opc** user):

```
oci compute volume-attachment attach-iscsi-volume --instance-id $instanceid --volume-id $datavolumeid --display-name data-volume --auth instance_principal
```

where `instanceid` is the OCID of the compute instance and `datavolumeid` is the OCID of the data volume. The `--display-name` parameter must be provided with this value.

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).

① Note

After successfully recovering into the target Essbase stack, you can delete the failed source compute node and do further cleanup to unneeded 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.

① Note

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

1. 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.
2. 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.

Required Policies

Before you undertake any back up and restoration activity, the policies outlined in this topic must be in place.

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

See [Set Up Policies](#) and [How Policies Work](#).

21.1 Required Policies

```
allow dynamic-group group_name to use autonomous-database 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 manage autonomous-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 volume-group-backups in compartment  
compartment_name
```

```
allow dynamic-group group_name to manage volumes in compartment  
compartment_name
```

```
allow dynamic-group group_name to inspect volume-groups in compartment  
compartment_name
```

```
allow dynamic-group group_name to manage buckets in compartment  
compartment_name
```

Required Policies for 21.2 and Higher

```
allow dynamic-group group_name to use autonomous-database 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 manage volume-backups in compartment  
compartment_name
```

```
allow dynamic-group group_name to use volumes in compartment compartment_name
```