

# Oracle® Communications Billing and Revenue Management

## Cloud Native Deployment Guide



Release 15.0

F86189-07

April 2026

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2019, 2026, Oracle and/or its affiliates.

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

## Preface

---

Audience	i
Documentation Accessibility	i
Diversity and Inclusion	i

## Part I Overview of BRM Cloud Native

---

### 1 Overview of the BRM Cloud Native Deployment

---

About the BRM Cloud Native Deployment	1
BRM Cloud Native Deployment Architecture	1
Images and Containers	3
Images and Containers with Non-WebLogic Server Pattern	3
Images and Containers for Applications Using WebLogic	3

### 2 About the BRM Cloud Native Deployment Packages

---

Overview of the BRM Cloud Native Deployment Package	1
About BRM Pods	1
About Client Pods and Images	3
About BRM PVCs and Pods	5

## Part II Getting Started with BRM Cloud Native Deployment

---

### 3 About Configuring and Deploying Your BRM Cloud Native Environment

---

About Configuring and Deploying BRM Cloud Native	1
High-Level Installation Tasks	2

### 4 Setting Up Prerequisite Software

---

BRM Cloud Native Prerequisite Tasks	1
Software Compatibility	2

Creating a Kubernetes Cluster	2
Installing Podman	3
Installing Helm	4
Creating and Configuring Your BRM Database	4
Installing an External Provisioner	5
Installing WebLogic Kubernetes Operator	6
Installing an Ingress Controller	7
Setting Up ECE Cloud Native Ingress and Egress Flows	7

## 5 Preparing Your BRM Cloud Native Environment

---

Tasks for Preparing Your BRM Cloud Native Environment	1
Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files	1
Pulling BRM Images from the Oracle Container Registry	3
Downloading BRM Images from Oracle Software Delivery Website	6
Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center	8

## Part III Configuring and Deploying BRM Cloud Native

---

### 6 Deploying the BRM Database Schema

---

Deploying BRM with a New Database Schema	1
Deploying BRM with an Existing Schema	7

### 7 Configuring BRM Server, PDC, and PCC Services

---

About Configuring BRM Cloud Native Services	1
Creating Secrets for Docker Registry Authorization	1
Configuring Global Values	2
Specifying the BRM Services to Deploy	3
Configuring the BRM Server	5
Configuring BRM for a Multischema Database	12
Configuring Pricing Design Center	15
Adding PDC Keys for oc-cn-helm-chart	15
Adding PDC Keys for oc-cn-op-job-helm-chart	22
Setting Up SSO for PDC Cloud Native	29
Configuring Pipeline Configuration Center	30
Adding Pipeline Configuration Center Keys for oc-cn-op-job-helm-chart	30
Adding Pipeline Configuration Center Keys for oc-cn-helm-chart	33
About PCC Volume Mounts	35
Creating a WebLogic Domain and Installing the PCC Application	35

Setting Up SSO for PCC	36
Setting Up Local Users and Groups for PCC	37
Starting and Stopping WebLogic Servers	38
Configuring SSL in PCC	38

## 8 Configuring REST Services

---

Configuring BRM REST Services Manager	1
Generating an SSL Certificate for BRM REST Services Manager	1
Configuring the SDK (Optional)	2
Configuring the Oracle Unified Directory HTTPS Port	2
Connecting to a Separate BRM Cluster	2
Adding BRM REST Services Manager Keys	3
Sample override-values for IDCS Security Type	5
Sample override-values.yaml for OAM Security Type	6
Sample BRM RSM override-values for Separate BRM Cluster	6
Configuring PDC REST Services Manager	7
Adding PDC REST Services Manager Keys	8
Configuring OAuth Authentication in PDC REST Services Manager	9
Configuring Requests to the Enterprise Product Catalog	11
Enabling TLS in PDC REST Services Manager	12
Enabling T3S in PDC REST Services Manager	13
Configuring Mapping of TMF620 priceType to BRM Events	14

## 9 Configuring the Billing Care, Billing Care REST API, and Business Operations Center Services

---

About Configuring Business Operations Center, Billing Care, and Billing Care REST API	1
Configuring Business Operations Center	2
Adding Business Operations Center Keys for oc-cn-op-job-helm-chart	3
Adding Business Operations Center Keys for oc-cn-helm-chart	6
Updating Infranet.properties for Business Operations Center	7
Adding Custom Configuration to Deployment Workflow for Business Operations Center	8
About Business Operations Center Volume Mounts	8
Creating a WebLogic Domain and Installing the Business Operations Center Application	8
Setting Up SSO for Business Operations Center	9
Setting Up Local Users and Groups for Business Operations Center	10
Starting and Stopping WebLogic Servers	11
Adding Authorization for a Custom Job Resource	12
Configuring Billing Care	12
Adding Billing Care Keys for oc-cn-op-job-helm-chart	13
Adding Billing Care Keys for oc-cn-helm-chart	15

Updating Infranet.properties for Billing Care	17
Adding Custom Configuration to Deployment Workflow for Billing Care	18
About Billing Care Volume Mounts	18
Creating a WebLogic Domain and Installing the Billing Care Application	19
Setting Up SSO for Billing Care	19
Setting Up Local Users and Groups for Billing Care	20
Starting and Stopping WebLogic Servers	21
Configuring the Billing Care REST API	22
Adding Billing Care REST API Keys for oc-cn-op-job-helm-chart	22
Adding Billing Care REST API Keys for oc-cn-helm-chart	24
Updating Infranet Properties for the Billing Care REST API	26
Adding Custom Configuration to Deployment Workflow for Billing Care REST API	27
About Billing Care REST API Volume Mounts	27
Creating a WebLogic Domain and Installing the Billing Care REST API	28
Setting Up Local Users and Groups for Billing Care REST API	28
Starting and Stopping WebLogic Servers	29

## 10 Configuring ECE Services

---

Adding Elastic Charging Engine Keys	1
Enabling SSL in Elastic Charging Engine	7
Connecting ECE Cloud Native to an SSL-Enabled Database	7
About Elastic Charging Engine Volume Mounts	10
Loading Custom Diameter AVP	10
Generating CDRs for Unrated Events	10
Scaling the cdrgateway and cdrFormatter Pods	14
Configuring ECE to Support Prepaid Usage Overage	15
Recording Failed ECE Usage Requests	15
Loading BRM Configuration XML Files	16
Setting Up Notification Handling in ECE	16
Creating an Apache Kafka Notification Topic	16
Creating an Oracle WebLogic Notification Queue	17
Configuring ECE for a Multischema BRM Environment	19

## 11 Deploying BRM Cloud Native Services

---

Deploying BRM Cloud Native Services	1
-------------------------------------	---

## 12 Deploying into Oracle Cloud Infrastructure

---

Deploying into Oracle Cloud Infrastructure	1
--	---

## 13 Uninstalling Your BRM Cloud Native Deployment

---

Uninstalling Your BRM Cloud Native Deployment	1
Uninstalling Selected BRM Cloud Native Services	1

## Part IV Customizing BRM Cloud Native

---

### 14 Customizing BRM Cloud Native Services

---

Customizing BRM Server	1
Customizing Billing Care	4
Customizing ECE	5

### 15 Building Your Own Images

---

Building BRM Server Images	1
Building Your BRM Server Base Image	2
Building Images of BRM Server Components	3
Building Web Services Manager Images	4
Building and Deploying Web Services Manager for Apache Tomcat Image	4
Building and Deploying Web Services Manager for WebLogic Server Image	6
Containerization of Email Data Manager	9
Containerization of Roaming Pipeline	10
Building and Deploying Vertex Manager	11
Deploying with Vertex Communications Tax Q Series	11
Deploying with Vertex Sales Tax Q Series	13
Building BRM REST Services Manager Images	14
Building PDC REST Services Manager Images	15
Building PDC Images	16
Building Pipeline Configuration Center Images	17
Pulling the Fusion Middleware Infrastructure Image	18
Building the Pipeline Configuration Center Image	18
Building Billing Care Images	18
Building the Billing Care Image	18
Building the Billing Care REST API Image	19
Building Business Operations Center Images	19

## Part V Upgrading BRM Cloud Native

---

## 16 Upgrading Your BRM Cloud Native Environment

---

Tasks for the BRM Cloud Native Upgrade	1
Upgrading Your Database Schema	2
Upgrading Your BRM Cloud Native Services	3
Upgrading Your ECE Cloud Native Services	5
Upgrading ECE Cloud Native to the Latest Interim Patch	6
Upgrading Your PDC Cloud Native Services	7
Upgrading BRM REST Services Manager	10
Upgrading Your Business Operations Center Cloud Native Services	11
Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.7.0 or Earlier to 15.0.x.0.0	11
Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.8.0 to 15.0.x.0.0	13
Upgrading Your PCC Cloud Native Services	14
Upgrading Your PCC Cloud Native Services from 12.0.0.x.0 or Earlier to 15.0.0.0.0	14
Upgrading Your PCC Cloud Native Services from 15.0.0.0.0 to 15.0.x.0.0	15
Upgrading Your Billing Care and Billing Care REST API Cloud Native Services	17
Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.7.0 or Earlier to 15.0.x.0.0	17
Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.8.0 to 15.0.x.0.0	19

## 17 Performing Zero-Downtime Upgrades

---

Performing a Zero-Downtime Upgrade of BRM	1
Performing a Zero Downtime Upgrade of PDC	4

## 18 Performing Zero-Downtime Upgrades of Disaster Recovery Cloud Native Systems

---

About the Zero-Downtime Upgrade of an Active-Active Disaster Recovery System	1
Tasks for Upgrading a BRM Cloud Native Active-Active System	1
Switching Off Site 2	2
Uninstalling BRM and ECE from Site 2	3
Upgrading Your BRM Database Schema in Site 2	4
Installing BRM 15.0.x Cloud Native on Site 2	5
Dropping the ECE Persistence Database Schema from Site 2	6
Installing ECE 15.0.x Cloud Native on Site 2	7
Failing Over Site 1 to Site 2	7
Uninstalling BRM and ECE from Site 1	9
Installing BRM Cloud Native on Site 1	9
Dropping the ECE Persistence Database Schema from Site 1	11

Installing ECE 15.0.x Cloud Native on Site 1	12
Federating ECE Cache Data Between Site 1 and Site 2	12

## 19 Migrating from On-Premise BRM to BRM Cloud Native

---

Migrating to BRM Cloud Native	1
Migrating from PDC On Premises to PDC Cloud Native	1

## Part VI Troubleshooting BRM Cloud Native Deployments

---

### 20 Troubleshooting Your BRM Cloud Native Deployment

---

Problems with the Helm Installation	1
Helm Installation Fails with Time-Out Error	2
BRM Cloud Native Deployment Out of Memory Errors	3
PDC Messages Stuck in Rating Engine Queues	3
PDC Interceptor Pod is Started But Went to Error State	3
eceTopology.conf Errors While Restarting Pods	4
SyncPDC Pod Fails to Start and Enters CrashLoopBackOff State	5

# Preface

This guide describes how to install and administer Oracle Communications Billing and Revenue Management (BRM) Cloud Native Deployment Option.

This guide has been updated to include changes and new feature content added for release 15.0.1.

## Audience

This document is intended for DevOps administrators and those involved in installing and maintaining an Oracle Communications Billing and Revenue Management (BRM) Cloud Native Deployment.

## Documentation Accessibility

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 customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

## Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

# Part I

## Overview of BRM Cloud Native

This part provides an overview of the Oracle Communications Billing and Revenue Management (BRM) cloud native deployment. It contains the following chapters:

- [Overview of the BRM Cloud Native Deployment](#)
- [About the BRM Cloud Native Deployment Packages](#)

# 1

## Overview of the BRM Cloud Native Deployment

Learn about configuring Oracle Communications Billing and Revenue Management (BRM) to run as a cloud native application in a containerized and orchestrated deployment architecture.

Topics in this document:

- [About the BRM Cloud Native Deployment](#)
- [BRM Cloud Native Deployment Architecture](#)

### About the BRM Cloud Native Deployment

Oracle Communications Billing and Revenue Management (BRM), along with the following BRM applications, are available in a cloud native deployment option, supporting a Kubernetes-orchestrated containerized multi-service architecture to facilitate continuous integration, continuous delivery, and DevOps practices. This allows you to harness the benefits of the cloud with BRM's services.

- Oracle Communications Pricing Design Center (PDC)
- Oracle Communications Elastic Charging Engine (ECE)
- Oracle Communications Pipeline Configuration Center (PCC)
- Oracle Communications Billing Care
- Oracle Communications Business Operations Center

#### Note

You can also deploy Oracle Communications Offline Mediation Controller on a cloud native environment. See "About the Offline Mediation Controller Cloud Native Deployment" in *Offline Mediation Controller Cloud Native Installation and Administration Guide* for more information.

You can set up your own BRM cloud native environment or build your own images of BRM and its applications. You use the cloud native deployment package to automate the deployment of BRM products and speed up the process to get services up and running, with product deployments preconfigured to communicate with each other through Helm charts.

### BRM Cloud Native Deployment Architecture

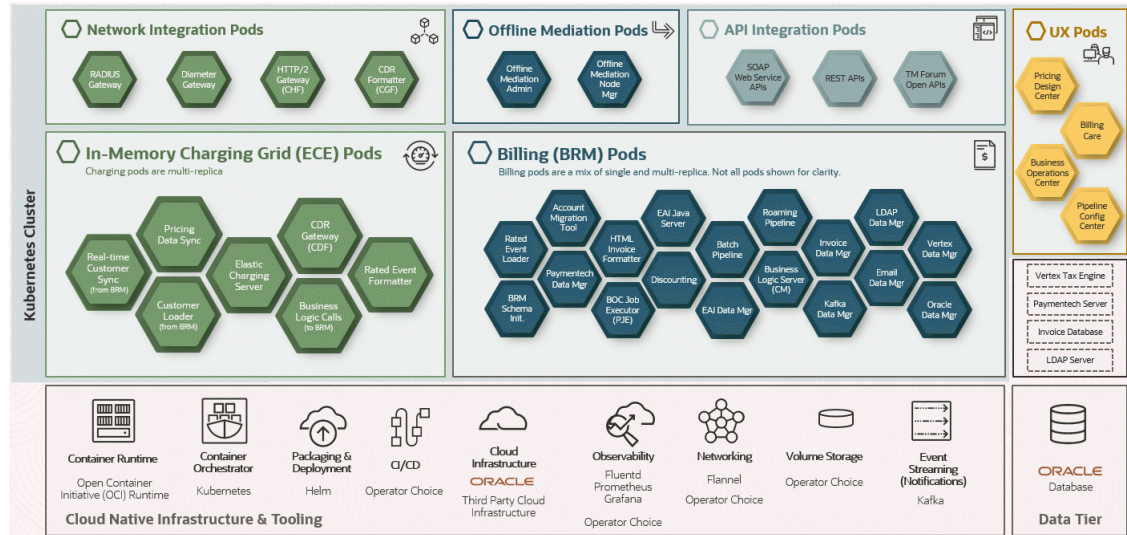
In the BRM cloud native architecture, each BRM service runs as a container and deploys as a Kubernetes pod, which is the fundamental building block of Kubernetes. Many core BRM services can be deployed and managed as multiple replicas within a Kubernetes replica set.

[Figure 1-1](#) shows the pods and other components in a typical BRM cloud native deployment.

**Note**

Not all pods are shown for clarity. Pod names are descriptive and may differ from actual names in some cases.

**Figure 1-1 BRM Cloud Native Deployment Architecture**



In this figure:

- Pricing Design Center, Billing Care, Business Operations Center, and Pipeline Configuration Center are client applications. They connect to the CM, which represents the business logic layer of BRM, by using the Portal Communications Protocol (PCP).
- The CM communicates with other pods, which represent the data management layer of BRM, by using the PCP protocol.
- All PCP protocol communication is encrypted using TLS.
- The data managers (DMs) interact with other downstream products that run the business logic. The downstream products can be containers or an on-premise system.
- ECE rates events and applies charges.
- Rating files for the batch pipeline are fed in through a Kubernetes PersistentVolumeClaim (PVC). The batch pipeline output is also available in a PVC for consumption by the Rated Event (RE) Loader pod.

**Note**

BRM services are not exposed outside of the cluster to improve security. Only PDC, PCC, Web Services Manager, Billing Care, and Business Operations Center are exposed.

## Images and Containers

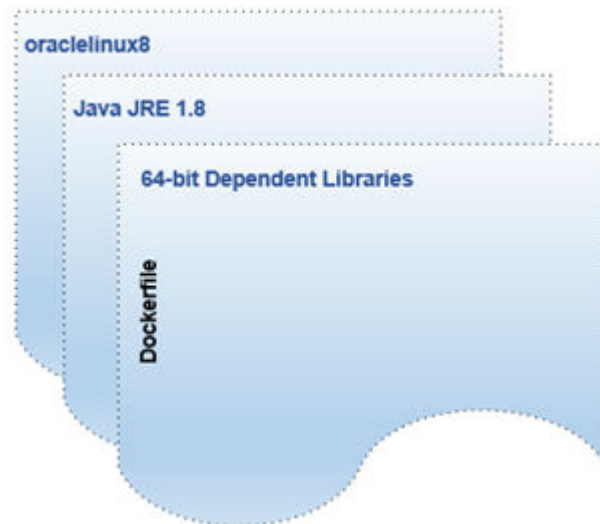
BRM has a multi-service architecture, with each service provided as an image for deploying as a run-time container in a Kubernetes cluster on cloud infrastructure. An image consists of read-only layers, each representing a Dockerfile instruction. The layers are stacked, and each is a delta of the changes from the previous layer. BRM cloud native deployment images are built by stacking multiple layers, extending an operating system image with a dependent library image, and then with an image packaging the application.

### Images and Containers with Non-WebLogic Server Pattern

BRM cloud native images that do not use WebLogic Server, such as the BRM base image, use the layering pattern shown in [Figure 1-2](#).

If you want to build your own BRM images, you must layer the images in this pattern.

**Figure 1-2 Base Image Layering with Non-WebLogic Server Pattern**



### Images and Containers for Applications Using WebLogic

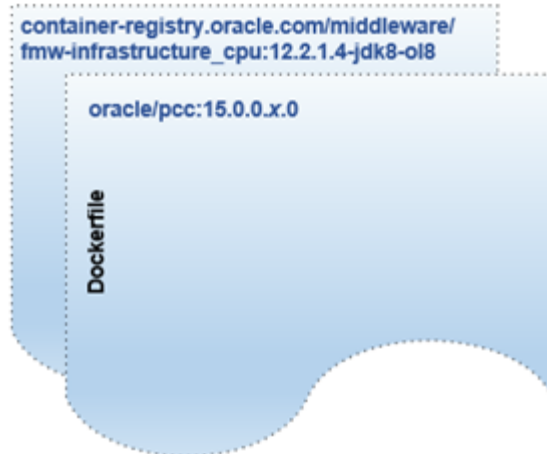
Some applications that use WebLogic have images based on the WebLogic image, and some applications use an external WebLogic image.

The following images are based on the WebLogic image:

- Web Services Manager image
- Pipeline Configuration Center (PCC)

[Figure 1-3](#) shows the layering pattern for a PCC application image, but a similar image stack also applies to the Web Services Manager image.

If you want to build your own PCC or Web Services Manager images, you must layer the images in this pattern.


**Figure 1-3 Images Based on a WebLogic Image**

In this figure:

- Fusion Middleware Infrastructure 12.2.1.x is the base image. This image is available from the Oracle Container Registry (<https://container-registry.oracle.com>).  
The Fusion Middleware Infrastructure image is based on Oracle Linux and Oracle JDK 8 (Server JRE). It's regularly patched with critical security fixes until the release date.
- The PCC application image extends the Fusion Middleware Infrastructure image, which provides WebLogic Server and JRF for OPSS for authorized access to the application.

[Figure 1-4](#) shows the layering pattern for a Billing Care image, but a similar image stack also applies to the Billing Care REST API and Business Operations Center images.

If you want to build your own Billing Care, Billing Care REST API, or Business Operations Center images, you must layer the images in this pattern.

**Figure 1-4 Images Using an External WebLogic Image**


In this figure:

- Oracle Linux 8 is the base image. This image is available from the Oracle Container Registry (<https://container-registry.oracle.com>). The Oracle Linux image is regularly patched with critical security fixes until the release date.
- The Billing Care image extends the Oracle Linux image. It references an external image (not part of the stack) for WebLogic functions, like WebLogic Server and JRF for OPSS for authorized access to the application.

[Figure 1-5](#) shows the layering pattern for a Billing Care image, but a similar image stack also applies to the Billing Care REST API and Business Operations Center images.

If you want to build your own PDC images, you must layer the images in this pattern.

**Figure 1-5 PDC Images Using an External WebLogic Image**



In this figure:

- The Server JRE image is the base image. This image is available from the Oracle Container Registry (<https://container-registry.oracle.com>).
- The PDC image extends the Server JRE image. It references an external image (not part of the stack) for WebLogic functions, like WebLogic Server and JRF for OPSS for authorized access to the application.

# 2

## About the BRM Cloud Native Deployment Packages

Learn about the Helm charts and images in the Oracle Communications Billing and Revenue Management (BRM) cloud native deployment package that help you deploy and manage pods of BRM product services in Kubernetes.

Topics in this document:

- [Overview of the BRM Cloud Native Deployment Package](#)
- [About BRM Pods](#)
- [About Client Pods and Images](#)
- [About BRM PVCs and Pods](#)

### Overview of the BRM Cloud Native Deployment Package

The BRM cloud native deployment package includes the following:

- Ready-to-use images and Helm charts to help you orchestrate containers in Kubernetes.
- Sample Dockerfiles and scripts that you can use as a reference for building your own images.

You can use the images and Helm charts to help you deploy and manage pods of BRM product services in Kubernetes. Communication between pods of services of BRM products is preconfigured in the Helm charts.

### About BRM Pods

[Table 2-1](#) lists the pods for BRM whose containers are created and services are exposed through them.

**Table 2-1 BRM Pods**

Pod Name	Replica Type (1)	Container Port	Container Port Name	Service Type
cm	Multi-replica	11960 (cm) 11961 (perflib, metrics) 11932 (eai-java-server, metrics)	cm-pcp-port cm-perflib-port eai-prom-port	ClusterIP
dm-kafka	Multi-replica	12010 12012 (metrics)	dm-pcp-port dm-prom-port	ClusterIP
dm-oracle	Multi-replica	12950 12951 (perflib, metrics)	dm-pcp-port dm-perflib-port	ClusterIP
brm-sdk	Single-replica	N/A	N/A	ClusterIP

Table 2-1 (Cont.) BRM Pods

Pod Name	Replica Type (1)	Container Port	Container Port Name	Service Type
init-db	Single-replica	N/A	N/A	ClusterIP
dm-vertex	Single-replica	31247	dm-vertex-port	ClusterIP
dm-eai	Multi-replica	11970	dm-pcp-port	ClusterIP
dm-invoice	Multi-replica	27777	dm-pcp-port	ClusterIP
dm-ldap	Multi-replica	12850	dm-pcp-port	ClusterIP
dm-prov-telco	Multi-replica	20315	dm-pcp-port	ClusterIP
ecs	Multi-replica	9999	N/A	NodePort
configloader	Single-replica	9999	N/A	NodePort
pricingupdate r	Single-replica	9999	N/A	NodePort
brmgateway	Multi- instance, 1 instance per schema	9999	N/A	NodePort
customerupdat er	Multi- instance, 1 instance per schema	9999	N/A	NodePort
diametergate way	Multi- instance, each instance can have 1 replica	9999	N/A	NodePort
emgateway	Multi-replica	9999	N/A	NodePort
httpgateway	Multi-replica	9999	N/A	NodePort
radiusgatewa y	Multi-replica	9999	N/A	NodePort
ratedeventfor matter	Multi- instance, 1 instance required for each role on each schema (2)	9999	N/A	NodePort
realtime-pipe	Multi-replica	24000	rtp	ClusterIP
batch- wireless-pipe	Single-replica	24001	batchpipe	ClusterIP
roampipe	Single-replica	24002	roampipe	ClusterIP
rel-daemon	Multi-replica	N/A	N/A	ClusterIP
rem	Multi- instance, 1 instance per schema	8080	rem-metrics	ClusterIP
batch- controller	Multi-replica	N/A	N/A	ClusterIP

Table 2-1 (Cont.) BRM Pods

Pod Name	Replica Type (1)	Container Port	Container Port Name	Service Type
formatter	Multi-replica	22272	formatter-port	ClusterIP
fusa-simulator	Single-replica	9780 (answer_s, online simulator) 8780 (answer_b, online simulator)	answer-s-port answer-b-port	ClusterIP
dm-fusa	Multi-replica	15772	dm-fusa-port	ClusterIP
dm-email	Multi-replica	17777	cm-pcp-port	ClusterIP
pje	Multi-replica	31960	pje-pcp-port	ClusterIP
amt	Single-replica	N/A	N/A	ClusterIP
brm_apps_jobs	N/A	N/A	N/A	N/A
config_job	N/A	N/A	N/A	ClusterIP

**Note:**

- The BRM cloud native pods support the following replica types:
  - Multi-replica:** You can scale these pods at the Kubernetes level to the number of replicas needed.
  - Single-replica:** You cannot scale these pods.
  - Multi-instance:** You can scale these pods but not at the Kubernetes level. Instead, a new instance needs to be created and scaling is usually limited.
- The `ratedeventformatter` pod requires one instance for each role on each schema. For example, suppose you have two schemas on two sites. In that case, you create primary and secondary instances for each schema in site 1 and primary and secondary instances for each schema in site 2, for a total of 8 instances.

## About Client Pods and Images

[Table 2-2](#) lists the pods and images for PDC, PDC REST Services Manager, Pipeline Configuration Center, Billing Care, Business Operations Center, and BRM REST Services Manager.

**Note**

For the list of pods and images for Offline Mediation Controller, see "About Offline Mediation Controller Pods and Images" in *Offline Mediation Controller Cloud Native Installation and Administration Guide*.

Table 2-2 Client Pods and Images

Pod	Replica Type <sup>(1)</sup>	Image	Container Port	Service Type	Access URL
pdcc	Single-replica	<b>pdcc:15.0.x.0.0</b>	8001 (HTTP) <sup>(2)</sup>	NodePort	<i>host:port/pdcc</i>
pdcrsm	Multi-replica	<b>pdcrsm:15.0.x.0.0</b>	31000 8080	Container Port, ClusterIP	<i>host:port/</i> <b>productCatalogManagement</b>
pcc	Single-replica	<b>pcc:15.0.x.0.0</b>	7012 (HTTPS)	NodePort	<i>host:port/pcc</i>
billingcare	Multi-replica <sup>(3)</sup>	<b>billingcare:15.0.x.0.0</b>	7011 (admin-server) 8001 (managed-serverN and cluster-1) 8080 (monitoring-exporter, if monitoring is enabled)	ClusterIP	<i>host:port/bc</i>
bcws	Multi-replica <sup>(3)</sup>	<b>bcws:15.0.x.0.0</b>	7011 (admin-server) 8001 (managed-serverN and cluster-1) 8080 (monitoring-exporter, if monitoring is enabled)	ClusterIP	<i>host:port/bcws</i>
boc	multi-replica	<b>boc:15.0.x.0.0</b>	7011 (admin-server) 8001 (managed-serverN and cluster-1) 8080 (monitoring-exporter, if monitoring is enabled)	ClusterIP	<i>host:port/</i> <b>opsdashboard</b>
brm-rest-services-manager	Multi-replica	<b>brm-rest-services-manager:15.0.x.0.0</b>	30000/30001	NodePort	<i>host:port/brm</i>
wsm-wls	Multi-replica	<b>brm_wsm_wls:15.0.x.0.0</b>	7001 (admin-server) 8001 (managed-serverN and cluster-1) 8080 (monitoring-exporter, metrics)	default default metrics	<i>host:port/</i> <b>BrmWebServices</b>
N/A	N/A	<b>webhook:15.0.x.0.0</b>	N/A	N/A	N/A

**Note:**

- The BRM cloud native pods support the following replica types:
  - Multi-replica:** You can scale these pods at the Kubernetes level to the number of replicas needed.
  - Single-replica:** You cannot scale these pods.
  - Multi-instance:** You can scale these pods but not at the Kubernetes level. Instead, a new instance needs to be created and scaling is usually limited.
- If the PDC user sets the **t3ChannelPort** and **t3sChannelPort** keys in the **values.yaml** file, the HTTP, HTTPS, t3Channel, and t3sChannel ports will be NodePort.
- The Billing Care and Billing Care REST API pods support multi-replica managed services with scaling done through WebLogic Kubernetes Operator.

## About BRM PVCs and Pods

[Table 2-3](#) lists the PVCs and pods in a BRM cloud native deployment.

**Table 2-3 List of PVCs in BRM Server**

PVC Name	Pods
bcws-domain-domain-pvc	bcws-domain-deployer bcws-domain-admin-server bcws-domain-managed-server <i>N</i>
bcws-domain-batch-payment-pvc	bcws-domain-deployer bcws-domain-admin-server bcws-domain-managed-server <i>N</i>
billingcare-domain-domain-pvc	billingcare-domain-deployer billingcare-domain-admin-server billingcare-domain-managed-server <i>N</i>
billingcare-domain-batch-payment-pvc	billingcare-domain-deployer billingcare-domain-admin-server billingcare-domain-managed-server <i>N</i>
boc-domain-domain-pvc	boc-domain-deployer boc-domain-admin-server boc-domain-managed-server <i>N</i>
common-semaphore	batch-wireless-pipe realtime-pipe roampipe
ctqdir	dm-vertex
custom-job-file	brm-apps-job cm
data	batch-wireless-pipe roampipe
fusa-temp	dm-fusa
oms-rel-archive	rel-daemon
oms-rel-input	rel-daemon
oms-rel-reject	rel-daemon

**Table 2-3 (Cont.) List of PVCs in BRM Server**

PVC Name	Pods
oms-uel-archive	batch-controller
oms-uel-input	batch-controller
oms-uel-reject	batch-controller
outputcdr	batch-wireless-pipe rel-daemon
outputreject	batch-wireless-pipe rel-daemon
pipelinelog	batch-wireless-pipe roampipe
pdcc-app-pvc	PDC pod (PDC Application Container)
pdcc-brm-pvc	PDC pod (PDC BRM Integration Pack)
roamoutputcdr	roampipe rel-daemon
roamoutputreject	roampipe rel-daemon
service-order	dm-prov-telco brm-apps-job
virtual-time	All pods

[Table 2-4](#) lists the services associated with ECE.

**Table 2-4 ECE Services**

Service Name	Service Type	Port	Description
ece-brmgateway	NodePort	External port	BRM Gateway service
ece-dgw	NodePort	External port	Diameter Gateway service
ece-emg	NodePort	External port	EM Gateway service
ece-http	NodePort	External port	HTTP Gateway service
ece-rgw	NodePort	External port	RADIUS Gateway service

# Part II

## Getting Started with BRM Cloud Native Deployment

This part provides information about getting started with your Oracle Communications Billing and Revenue Management (BRM) cloud native deployment, including installing the prerequisite software and downloading the deployment package. It contains the following chapters:

- [About Configuring and Deploying Your BRM Cloud Native Environment](#)
- [Setting Up Prerequisite Software](#)
- [Preparing Your BRM Cloud Native Environment](#)

# 3

## About Configuring and Deploying Your BRM Cloud Native Environment

Learn about the high-level steps for configuring and deploying your Oracle Communications Billing and Revenue Management (BRM) cloud native environment.

Topics in this document:

- [About Configuring and Deploying BRM Cloud Native](#)
- [High-Level Installation Tasks](#)

### About Configuring and Deploying BRM Cloud Native

You install the BRM cloud native deployment package by configuring and deploying its Helm charts. The Helm charts include YAML template descriptors for all Kubernetes resources and a **values.yaml** file that provides default configuration values for each chart.

Installing a Helm chart generates valid Kubernetes manifest files by replacing default values from the **values.yaml** file with custom values from your **override-values.yaml** file, and creates Kubernetes resources. Helm calls this a new release. You use the release name to track and maintain this installation.

**Note**

This documentation uses the **override-values.yaml** file name for ease of use, but you can name the file whatever you want.

The BRM cloud native deployment package includes the Helm charts in [Table 3-1](#).

**Table 3-1 BRM Cloud Native Helm Charts**

Chart Name	Description	When to Use
<b>oc-cn-init-db-helm-chart</b>	<p>This chart initializes and upgrades the database schema for the BRM server.</p> <ul style="list-style-type: none"><li>• In initialize mode, it:<ul style="list-style-type: none"><li>– Creates tables, views, procedures, indexes, and other database objects needed by BRM Server</li><li>– Loads seed data</li></ul></li><li>• In upgrade mode, it modifies the existing database schema to match the current release's data model.</li></ul>	<p>Use this chart in initialize mode when preparing a new BRM setup and have an empty database schema.</p> <p>Use this chart in upgrade mode when upgrading your schema to the latest release.</p>

Table 3-1 (Cont.) BRM Cloud Native Helm Charts

Chart Name	Description	When to Use
<b>oc-cn-op-job-helm-chart</b>	<p>This chart does the following:</p> <ul style="list-style-type: none"> <li>Creates WebLogic Server domains for PDC, Billing Care, the Billing Care REST API, and Business Operations Center.</li> <li>Installs PDC, Billing Care, the Billing Care REST API, and Business Operations Center in their respective domains.</li> <li>Populates persistent volumes with domain and application files for sharing between WebLogic Server runtimes.</li> <li>(Release 15.0.1 or later) Creates the following PDC groups: <ul style="list-style-type: none"> <li><b>PricingDesignAdmin:</b> This group's users have administrative privileges on PDC. They can perform operations on all PDC UI screens, pricing components, and setup components.</li> <li><b>PricingAnalyst:</b> This group's users have administrative privileges for pricing components and view-only privileges for setup components.</li> <li><b>PricingReviewer:</b> This group's users have view-only privileges for all pricing and setup components.</li> </ul> </li> </ul>	<p>If you want to use Billing Care, Business Operations Center, Pricing Design Center, or the Billing Care REST API, install this chart before you install <b>oc-cn-helm-chart</b>.</p>
<b>oc-cn-helm-chart</b>	<p>This chart does the following:</p> <ul style="list-style-type: none"> <li>Deploys BRM server, PDC, and PCC.</li> <li>Starts the WebLogic servers for Billing Care, the Billing Care REST API, and Business Operations Center.</li> <li>Exposes web clients as services outside of the cluster.</li> <li>Shares persistent volumes between its services through persistent volume claims.</li> </ul>	<p>Install this chart to use the services of the BRM server, PDC, PCC, Billing Care, Business Operations Center, or the Billing Care REST API.</p>
<b>oc-cn-ece-helm-chart</b>	<p>This chart does the following:</p> <ul style="list-style-type: none"> <li>Deploys ECE and its services.</li> <li>Sets up the connection with the BRM server and PDC.</li> <li>Configures sharing of persistent volumes with the BRM server.</li> </ul>	<p>Install this chart to use ECE as your convergent charging solution.</p>

## High-Level Installation Tasks

You install BRM cloud native on your system by performing these high-level tasks:

1. Install all prerequisite software for your BRM cloud native environment.  
See "[Setting Up Prerequisite Software](#)".
2. Prepare your deployment environment by downloading the BRM cloud native deployment package, extracting the Helm charts, and loading the BRM component images.  
See "[Preparing Your BRM Cloud Native Environment](#)".
3. Configure and deploy the BRM database schema in your cloud native environment.  
See "[Deploying the BRM Database Schema](#)".
4. Configure the BRM cloud native services that you want to include in your system, including:
  - BRM server, PDC, and PCC services. See "[Configuring BRM Server, PDC, and PCC Services](#)".
  - BRM and PDC REST services. See "[Configuring REST Services](#)".
  - BRM client services such as Billing Care and Business Operations Center. See "[Configuring the Billing Care, Billing Care REST API, and Business Operations Center Services](#)".
  - ECE services. See "[Configuring ECE Services](#)".
5. Deploy the BRM cloud native services in your cloud native environment.  
See "[Deploying BRM Cloud Native Services](#)".

# 4

## Setting Up Prerequisite Software

Learn about prerequisite tasks to perform before installing the Oracle Communications Billing and Revenue Management (BRM) cloud native deployment package, such as installing Podman and Helm.

Topics in this document:

- [BRM Cloud Native Prerequisite Tasks](#)
- [Software Compatibility](#)
- [Creating a Kubernetes Cluster](#)
- [Installing Podman](#)
- [Installing Helm](#)
- [Creating and Configuring Your BRM Database](#)
- [Installing an External Provisioner](#)
- [Installing WebLogic Kubernetes Operator](#)
- [Installing an Ingress Controller](#)
- [Setting Up ECE Cloud Native Ingress and Egress Flows](#)

### **Caution**

Oracle does not provide support for any prerequisite third-party software installation or configuration. The customer must handle any installation or configuration issues related to non-Oracle prerequisite software.

## BRM Cloud Native Prerequisite Tasks

As part of preparing your environment for BRM cloud native, you choose, install, and set up various components and services in ways that are best suited for your cloud native environment. The following shows the high-level prerequisite tasks for BRM cloud native:

1. Ensure you have downloaded the latest supported software that is compatible with BRM cloud native.
2. Create a Kubernetes cluster.
3. Install Podman and a container runtime supported by Kubernetes.
4. Install Helm.
5. Create and configure a BRM database.
6. Install and configure an external provisioner.
7. If you plan to deploy Pricing Design Center (PDC), Billing Care, the Billing Care REST API, Web Services Manager, or Business Operations Center:

- Install and configure WebLogic Kubernetes Operator.
  - Install an ingress controller.
8. If you plan to deploy Elastic Charging Engine (ECE), install and set up an ingress controller and an egress controller.
  9. If you plan to deploy the Billing Care REST API or the BRM REST Services Manager API, install Oracle Access Management. See the "Install Oracle Access Management 12c" tutorial for installation instructions.
  10. If you plan to integrate your BRM cloud native deployment with a Kafka Server, install the Apache Kafka software. See "[Apache Kafka Quickstart](#)" on the Apache Kafka website for installation instructions.
  11. If you plan to integrate your BRM cloud native deployment with Oracle Analytics Publisher, install Oracle Analytics Publisher. See "[Installing the Oracle Analytics Server Software](#)" in *Oracle Analytics Installing and Configuring Oracle Analytics Server* for installation instructions.

#### Note

The Oracle Analytics Publisher software was previously named Oracle Business Intelligence (BI) Publisher.

Prepare your environment with these technologies installed, configured, and tuned for performance, networking, security, and high availability. Make sure backup nodes are available in case of system failure in any of the cluster's active nodes.

The following sections provide more information about the required components and services, the options you can choose from, and how you must set them up for your BRM cloud native environment.

## Software Compatibility

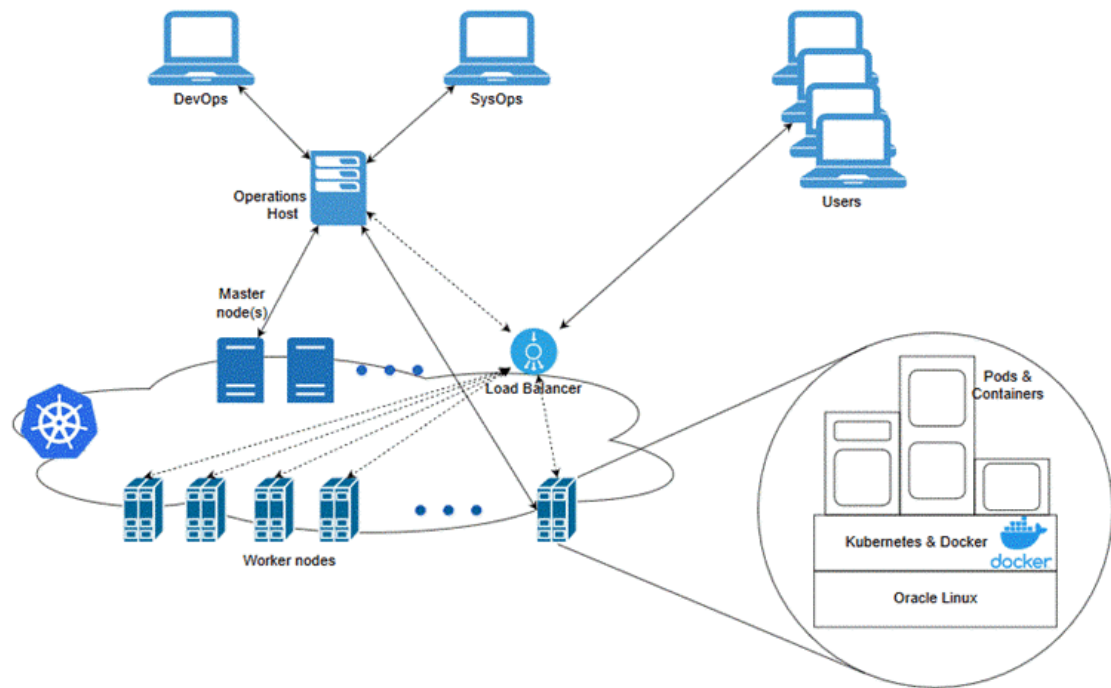
To run, manage, and monitor your BRM cloud native deployment, ensure you use the latest versions of all compatible software. See "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix*.

## Creating a Kubernetes Cluster

Kubernetes is an open-source system for automating the deployment, scaling, and management of containerized applications. It groups containers into logical units for easy management and discovery. When you deploy Kubernetes, you get a physical cluster with machines called nodes. A reliable cluster must have multiple worker nodes spread over separate physical infrastructure, and a very reliable cluster must have multiple primary nodes spread over separate physical infrastructure.

[Figure 4-1](#) illustrates the Kubernetes cluster and the components that it interacts with.

Figure 4-1 Overview of the Kubernetes Cluster



Set up a Kubernetes cluster for your BRM cloud native deployment, securing access to the cluster and its objects with the help of service accounts and proper authentication and authorization modules. Also, set up the following in your cluster:

- **Volumes:** Volumes are directories that are accessible to the containers in a pod and provide a way to share data. The BRM cloud native deployment package uses persistent volumes for sharing data in and out of containers, but does not enforce any particular type. You can choose from the volume type options available in Kubernetes.
- **A networking model:** Kubernetes assumes that pods can communicate with other pods, regardless of which host they land on. Every pod gets its own IP address, so you do not need to explicitly create a link between pods or map container ports to host ports. Several implementations are available that meet the fundamental requirements of Kubernetes' networking model. Choose the networking model depending on the cluster requirement.

For more information about Kubernetes, see "[Kubernetes Concepts](#)" in the Kubernetes documentation.

## Installing Podman

You use the Podman platform to containerize BRM products. Install Podman if you want to do one of these:

- Use the prebuilt images provided with the BRM cloud native deployment package.
- Build your own BRM images by writing your own Dockerfiles using the sample Dockerfiles from the BRM cloud native deployment package.

You can use Podman or any container runtime that supports the Open Container Initiative if it supports the Kubernetes version specified in "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix*.

## Installing Helm

Helm is a package manager that helps you install and maintain software on a Kubernetes system. In Helm, a package is called a *chart*, which consists of YAML files and templates rendered into Kubernetes manifest files. The BRM cloud native deployment package includes Helm charts that help create Kubernetes objects, such as ConfigMaps, Secrets, controller sets, and pods, with a single command.

The following shows sample steps for installing and validating Helm:

1. Download the Helm software from <https://github.com/helm/helm/releases>.

For the list of supported Helm versions, see "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix*.

2. Extract the Helm files from the archive:

```
tar -zxvf helm-version-linux-amd64.tar.gz
```

where *version* is the Helm version number.

3. Find the helm binary in the unpacked directory and move it to your desired directory. For example:

```
mv linux-amd64/helm /usr/local/bin/helm
```

4. Check the version of Helm:

```
helm version
```

Helm leverages **kubeconfig** for users running the **helm** command to access the Kubernetes cluster. By default, this is **\$HOME/.kube/config**. Helm inherits the permissions set up for this access into the cluster. If role-based access control (RBAC) is configured, you must grant Helm users sufficient cluster permissions.

For more information about installing Helm, see "[Installing Helm](#)" in the Helm documentation.

## Creating and Configuring Your BRM Database

You must install an Oracle database accessible through the Kubernetes network so BRM cloud native pods can perform database operations. The Oracle database you use can be:

- On-premises, which can be either physical or VM
- Cloud-based, such as Bare Metal, VM, or DBaaS on Oracle Cloud Infrastructure

You can use an existing BRM database or create a new one. See *BRM Compatibility Matrix* for the latest supported database versions.

To create and configure a new BRM database:

1. When you install and create your database, pay particular attention to the following requirements:
  - Install **Oracle Enterprise Edition**
  - Install the following Oracle components: **Oracle XML DB**, **Oracle XML Developer's Kit (XDK)**, and **Oracle JServer**
  - To partition the tables in your BRM database, install the **Oracle Partitioning** component
  - Set the **Character Set** to **AL32UTF8**

- Set the **National Character Set** to **UTF8**
- 2. Set your `LD_LIBRARY_PATH` environment variable to **\$ORACLE\_HOME/lib**.
- 3. (Optional) Set up TLS authentication in the BRM database. See "[Configuring Transport Layer Security Authentication](#)" in *Oracle Database Security Guide*. Also, ensure that you:
  - Create a TLS certificate or obtain one from a certificate provider
  - Install the certificate in the Oracle Database Server
- 4. Configure your database manually or let the BRM installer configure the database for you.

### Note

For production systems, your database administrator must create the database manually.

You can do one of the following:

- **Use the BRM installer to configure a demonstration database for you**  
The BRM installer can automatically configure your database for demonstration or development systems. The BRM installer configures your database by:
  - Creating the following tablespaces: **pin00** (for data), **pinx00** (for indexes), and **PINTEMP** (for a temporary tablespace)
  - Creating a BRM user named **pin**
  - Granting connection privileges to the **pin** user
- **Configure a demonstration database manually**  
You can configure your database manually so it contains additional or larger tablespaces. For more information, see "Configuring Your Database Manually for Demonstration Systems" in *BRM Installation Guide*.
- **Configure a production database manually**  
For production systems, your database administrator must create the tablespaces for the BRM data and indexes. For information on estimating your database size and creating tablespaces, see "Planning Your Database Configuration" in *BRM Installation Guide*. To set the default storage model values for your cloud native installation, set the values in the **configmap\_partition\_cfg.yaml** file instead of the **pin\_tables.values** file.
- 5. Grant the BRM schema user select permission on the `V$SESSION` database table. To do so, connect to the Oracle database with SQL\*Plus as the **system** user and then enter this command:

```
SQL> GRANT SELECT ON TABLE V$SESSION TO brmSchemaUser;
```

The installers for PDC, Billing Care, and all other products automatically create the tablespaces and users that are required for those products.

## Installing an External Provisioner

An external provisioner creates shared, persistent storage for the containers in your BRM cloud native environment. It stores:

- Input data, such as pricing XML files

- Output data, such as archive files and reject files from Rated Event Loader and Universal Event Loader
- Data that needs to be shared between containers, such as `pin_virtual_time`

Install and set up an external provisioner with ReadWriteMany access in your system that provisions volumes dynamically.

## Installing WebLogic Kubernetes Operator

Oracle WebLogic Kubernetes Operator helps you to deploy and manage WebLogic domains in your Kubernetes environment. It consists of several parts:

- The operator runtime
- The model for a Kubernetes customer resource definition (CRD)
- A Helm chart for installing the operator

In the BRM cloud native environment, you use WebLogic Kubernetes Operator to maintain the domains and services for Billing Care, the Billing Care REST API, Web Services Manager, PDC, and Business Operations Center.

The following shows sample steps for installing WebLogic Kubernetes Operator on your BRM cloud native environment:

1. Add the Helm repository for WebLogic Kubernetes Operator:

```
helm repo add weblogic-operator https://oracle.github.io/weblogic-kubernetes-operator/charts
```

2. Create a new namespace for WebLogic Kubernetes Operator. For example, this `kubectl` command creates the namespace `operator`:

```
kubectl create namespace operator
```

3. Install WebLogic Kubernetes Operator:

```
helm install weblogic-operator weblogic-operator/weblogic-operator --namespace operator --version version
```

where `version` is the version of WebLogic Kubernetes Operator, such as **2.5.0** or **3.0.0**. See "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix* for a list of supported versions.

If the installation is successful, you will see something similar to this:

```
NAME: weblogic-operator
LAST DEPLOYED: Tue Oct 6 08:29:03 2020
NAMESPACE: weblogic-operator
STATUS: deployed
REVISION: 1
TEST SUITE: None
```

4. Check the pod:

```
kubectl get pods
```

You should see something similar to this:

NAME	READY	STATUS	RESTARTS	AGE
weblogic-operator-849cc6bdd8-vkx7n	1/1	Running	0	57s

For more information about WebLogic Kubernetes Operator, see "[Introduction](#)" in the WebLogic Kubernetes Operator documentation.

## Installing an Ingress Controller

Using an ingress controller exposes BRM services outside the Kubernetes cluster and allows clients to communicate with BRM.

The ingress controller monitors the ingress objects and acts on the configuration embedded in these objects to expose BRM HTTP and T3 services to the external network. Adding an external load balancer provides highly reliable single-point access to the services exposed by the Kubernetes cluster. In this case, the ingress controller exposes the services on behalf of the BRM cloud native instance. Using a load balancer removes the need to expose Kubernetes node IPs to the larger user base, insulates users from changes (in terms of nodes appearing or being decommissioned) to the Kubernetes cluster, and enforces access policies.

If you are using Billing Care, the Billing Care REST API, or Business Operations Center, you must add a load balancer to your BRM cloud native system that has:

- Path-based routing for the WebLogic Cluster service.
- Sticky sessions enabled. That is, if the load balancer redirects a client's login request to Managed Server 1, all subsequent requests from that client are redirected to Managed Server 1.
- TLS enabled between the client and the load balancer to secure communications outside of the Kubernetes cluster.

Business Operations Center and Billing Care use HTTP and rely on the load balancer to terminate HTTPS.

See "[Ingress](#)" in the WebLogic Kubernetes Operator documentation for more information about setting up an ingress controller and sample load balancers.

## Setting Up ECE Cloud Native Ingress and Egress Flows

Ingress and egress controllers expose ECE services outside the Kubernetes cluster, allowing external networks to communicate with ECE. For example, an ingress controller can route requests from Diameter and 5G HTTP clients to the httpgateway and diametergateway pods for processing. Likewise, an egress controller can send CDR records from the cdrformatter pod to the ECE database.

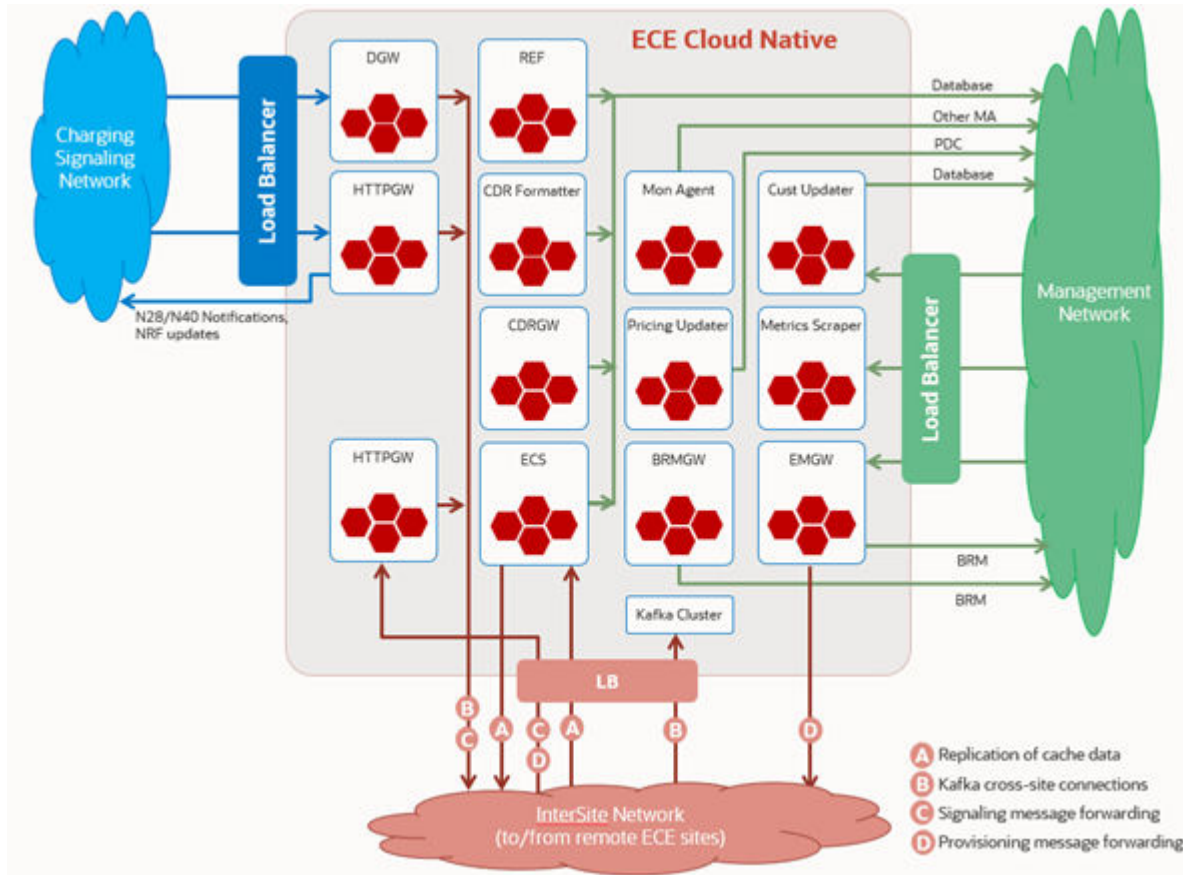
You can expose external network IPs for ingress traffic from external clients using the following:

- A load balancer exposing the IPs on the external network. The load balancer sends ingress traffic to Kubernetes services or node ports.
- Kubernetes service IPs or worker node IPs residing on the external network.

You can route egress traffic through an external network IP that is hosted by a worker node interface.

[Figure 4-2](#) shows an ECE cloud native deployment with sample ingress and egress flows.

Figure 4-2 ECE Cloud Native Ingress and Egress Flows



In this figure:

- The ingress flows traverse a load balancer, but you can use an alternate ingress flow to meet your business requirements.
- The egress flows are depicted generically. Network source addressing and routing may vary based on your business requirements.
- ECE cloud native uses logical external networks. The number and content of these networks may vary depending on your business requirements.

[Table 4-1](#) describes the egress flow from each ECE pod to an endpoint.

Table 4-1 ECE Cloud Native Egress Flows

ECE Pod	Egress Endpoints
brmgateway	ECE database BRM database <b>Note:</b> The BRM database is accessed during installation.
cdrformatter	ECE database <b>Note:</b> The ECE database is used for CDR management.
cdrgateway	ECE database <b>Note:</b> The ECE database is used for CDR management.

Table 4-1 (Cont.) ECE Cloud Native Egress Flows

ECE Pod	Egress Endpoints
configloader	ECE database
customerupdater	ECE database BRM database
diametergateway	Remote HTTP Gateway (for active-active deployments only) Remote Kafka server (for active-active deployments only) <b>Note:</b> This pod does not initiate Diameter connections to Diameter signaling clients.
ece-customerloader-job	ECE database BRM database
ece-persistence-job	ECE database
ece-persistence-upgrade-job	ECE database
ecs	ECE database BRM database Remote ECE Coherence Federation (for active-active and active-standby deployments only) <b>Note:</b> The BRM database is accessed during customer loading.
emgateway	BRM database Remote HTTP Gateway (for active-active deployments only) <b>Note:</b> This pod forwards requests to the remote HTTP Gateway. This is optional for active-active deployments.
httpgateway	Charging signaling clients Remote HTTP Gateway (for active-active deployments only) Remote Kafka server (for active-active deployments only) <b>Note:</b> This pod sends HTTP/2 requests to 5G clients. <b>Note:</b> The httpgateway pod's egress to a remote ECE HTTP Gateway is needed only if it is processing 5G charging traffic.
monitoringagent	Monitoring Agents Remote Monitoring Agent (for active-active and active-standby deployments only)
pricingupdater	Pricing Design Center
ratedeventformatter	ECE database BRM database (for Rated Event Manager plug-in only) <b>Note:</b> Direct access to the BRM database occurs only when the Rated Event Manager plug-in is configured to write rated events directly to the BRM database.

# 5

## Preparing Your BRM Cloud Native Environment

Learn how to prepare your system for the Oracle Communications Billing and Revenue Management (BRM) cloud native deployment by downloading the BRM cloud native Helm charts and BRM images.

Topics in this document:

- [Tasks for Preparing Your BRM Cloud Native Environment](#)
- [Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)
- [Pulling BRM Images from the Oracle Container Registry](#)
- [Downloading BRM Images from Oracle Software Delivery Website](#)
- [Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center](#)

### Tasks for Preparing Your BRM Cloud Native Environment

Prepare your system for the BRM cloud native deployment by performing the following high-level tasks:

1. Downloading the Helm charts for the BRM cloud native deployment. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Downloading the BRM cloud native images in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. If you plan to deploy Pricing Design Center (PDC), Billing Care, the Billing Care REST API, or Business Operations Center, downloading the Oracle WebLogic cloud native image. See "[Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center](#)".

### Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files

To download the BRM cloud native Helm charts and Docker files:

1. Go to <https://edelivery.oracle.com>.
2. Sign in to the Oracle Software Delivery website using an Oracle account.
3. Search for and select **Oracle Communications Billing and Revenue Management Cloud Native Deployment Option 15.0.x.0.0**, and then click **Continue**.
4. Select the following and then click **Continue**:

- Oracle Communications Billing and Revenue Management Cloud Native Deployment Option 15.0.x.0.0-CN
  - Oracle Communications Elastic Charging Engine Cloud Native Deployment Option 15.0.x.0.0-CN
  - Oracle Communications Pricing Design Center Cloud Native Deployment Option 15.0.x.0.0-CN
5. Accept the Oracle standard terms and restrictions and then click **Continue**.
  6. Select the following packages and then click **Download**:
    - Oracle Communications Cloud Native Helm Chart 15.0.x.0.0
    - Oracle Communications Elastic Charging Engine Cloud Native Deployment Option Helm Chart 15.0.x.0.0
    - Oracle Communications Cloud Native Operator Job Helm Chart 15.0.x.0.0
    - Oracle Communications Cloud Native Database Initializer Helm Chart 15.0.x.0.0
    - Oracle Communications Cloud Native Docker Build Files 15.0.x.0.0
    - Oracle Communications Elastic Charging Engine Cloud Native Docker Files 15.0.x.0.0
    - Oracle Communications Cloud Native Pricing Design Center 15.0.x.0.0

Each package is downloaded as a separate Zip file.

7. Extract the following Helm chart and Docker archive files from each Zip file:
  - BRM Helm Chart: **oc-cn-helm-chart-15.0.x.0.0.tgz**
  - ECE Helm Chart: **oc-cn-ece-helm-chart-15.0.x.0.0.tgz**
  - Operator Job Helm Chart: **oc-cn-op-job-helm-chart-15.0.x.0.0.tgz**
  - Database Initializer Helm Chart: **oc-cn-init-db-helm-chart-15.0.x.0.0.tgz**
  - BRM Dockerfiles: **oc-cn-docker-files-15.0.x.0.0.tgz**
  - ECE Dockerfiles: **oc-cn-ece-docker-files-15.0.x.0.0.tgz**

8. Extract the Helm charts and Dockerfiles from the archive files by running these commands:

```
tar xvzf oc-cn-helm-chart-15.0.x.0.0.tgz
```

```
tar xvzf oc-cn-ece-helm-chart-15.0.x.0.0.tgz
```

```
tar xvzf oc-cn-op-job-helm-chart-15.0.x.0.0.tgz
```

```
tar xvzf oc-cn-init-db-helm-chart-15.0.x.0.0.tgz
```

```
tar xvzf oc-cn-docker-files-15.0.x.0.0.tgz
```

```
tar xvzf oc-cn-ece-docker-files-15.0.x.0.0.tgz
```

[Table 5-1](#) lists the files and directories extracted from the archive files.

Table 5-1 Extracted Files

Archive File	Extracted Directories
<b>oc-cn-helm-chart-15.0.x.0.0.tgz</b>	<b>oc-cn-helm-chart</b> directory: Contains the BRM Helm chart files. <b>sample_configurations</b> directory: This directory contains the default configuration XML files, such as <b>bus_params_AR.xml</b> and <b>pin_config_export_gl.xml</b> .
<b>oc-cn-ece-helm-chart-15.0.x.0.0.tgz</b>	<b>oc-cn-ece-helm-chart</b> directory: Contains the ECE Helm chart files.
<b>oc-cn-op-job-helm-chart-15.0.x.0.0.tgz</b>	<b>oc-cn-op-job-helm-chart</b> directory: Contains the WebLogic Operator Job Helm chart files.
<b>oc-cn-init-db-helm-chart-15.0.x.0.0.tgz</b>	<b>oc-cn-init-db-helm-chart</b> directory: Contains the Database Initializer Helm chart files.
<b>oc-cn-docker-files-15.0.x.0.0.tgz</b>	<b>oc-cn-docker-files</b> directory: Contains the Dockerfiles for BRM, PDC, PDC REST Services Manager, Pipeline Configuration Center, Business Operations Center, and Billing Care.
<b>oc-cn-ece-docker-files-15.0.x.0.0.tgz</b>	<b>docker_files</b> directory: Contains the Dockerfiles for ECE.

## Pulling BRM Images from the Oracle Container Registry

To pull BRM cloud native images, such as the Connection Manager (CM) image and the Data Manager (DM) image, from the Oracle Container Registry, do the following:

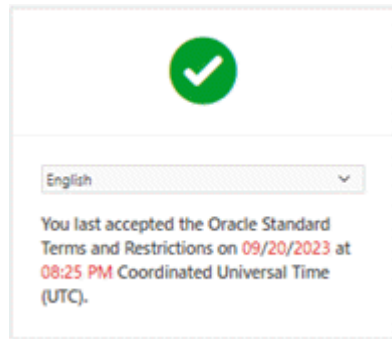
1. In a web browser, go to <https://container-registry.oracle.com>.
2. Sign in to the Oracle Container Registry using an Oracle account.

### Note

To pull images for licensed software on the Oracle Container Registry, you must have an Oracle account. You can create an Oracle account at <https://profile.oracle.com/myprofile/account/create-account.jsp>.

3. Select the **Oracle Communications Cloud Scale Monetization** container.  
The Oracle Communications Cloud Scale Monetization page appears.
4. Select one of the repository names from [Table 5-2](#).  
The repository page appears.
5. Accept the Oracle terms and restrictions by:
  - a. (For non-CPU images) Selecting your desired language.
  - b. Clicking **Continue**.
  - c. Scrolling to the bottom of the terms and restrictions pages and clicking **Accept**.

If successful, you will see something similar to this:



- On your host system, log in to the Oracle Container Registry using the Podman command-line interface (CLI):

```
podman login container-registry.oracle.com
```

- When prompted for a user name and password, enter your Oracle credentials.

- Pull the BRM cloud native image from the registry:

```
podman pull container-registry.oracle.com/communications_monetization/imageName:tag
```

where:

- imageName* is the name of a software image listed in [Table 5-2](#).
- tag* is the tag name for the BRM cloud native image, such as **15.0.x.0.0**.

For example, to pull the CM cloud native image from the registry:

```
podman pull container-registry.oracle.com/communications_monetization/cm:15.0.x.0.0
```

The image is pulled from the Oracle Container Registry and stored locally, where it is ready to be used to deploy containers.

- Confirm the images have been pulled from the Oracle Container Registry:

```
podman images
```

If successful, you will see something similar to this:

```

REPOSITORY                                     TAG
IMAGE ID          CREATED
container-registry.oracle.com/communications_monetization/cm      15.0.x.0.0
133dd3580b87      2 seconds ago
container-registry.oracle.com/communications_monetization/dm_kafka 15.0.x.0.0
136dd3593b47      3 seconds ago

```

- Log out of the registry to prevent unauthorized access and to remove any record of sign-in credentials that Podman might store for future operations:

```
podman logout container-registry.oracle.com
```

[Table 5-2](#) lists the image names for the BRM cloud native components.

**Table 5-2 BRM Cloud Native Images**

Component Name	Image Name
Batch Controller	<b>batch_controller</b>
Batch Pipeline	<b>batch_pipeline</b>

**Table 5-2 (Cont.) BRM Cloud Native Images**

<b>Component Name</b>	<b>Image Name</b>
Billing Care	<b>billingcare</b>
Billing Care REST API	<b>bcws</b>
BRM Applications	<b>brm_apps</b>
BRM REST Services Manager	<b>brm-rest-services-manager</b>
BRM SDK	<b>brm_sdk</b>
Business Operations Center	<b>boc</b>
Connection Manager	<b>cm</b>
Database Initializer	<b>init_db</b>
Database Upgrade	<b>upgrade</b>
Elastic Charging Engine	<b>oc-cn-ece</b>
Email Data Manager	<b>dm_email</b>
Enterprise Application Integration Data Manager	<b>dm_eai</b>
Enterprise Application Integration Java Server	<b>eai_js</b>
Invoice Data Manager	<b>dm_invoice</b>
Invoice Formatter	<b>formatter</b>
Kafka Data Manager	<b>dm_kafka</b>
LDAP Data Manager	<b>dm_ldap</b>
Oracle Data Manager	<b>dm_oracle</b>
Paymentech Data Manager	<b>dm_fusa</b>
Paymentech Simulator	<b>answer</b>
Performance Libraries	<b>perflib</b>
Pipeline Configuration Center	<b>pcc</b>
Pricing Design Center	<b>pdc</b>
Pricing Design Center REST Services Manager	<b>pdcrsm</b>
Provisioning Data Manager	<b>dm_prov_telco</b>
Rated Event Loader	<b>rel_daemon</b>
Rated Event Manager	<b>rem</b>
Real-Time Pipeline	<b>realtimepipe</b>
Roaming Manager	<b>roam_pipeline</b>
Vertex Data Manager	<b>dm_vertex</b>
Webhook	<b>webhook</b>
Web Services Manager	<b>brm_wsm_wls</b> <b>brm_wsm_wl_init</b>

# Downloading BRM Images from Oracle Software Delivery Website

To download BRM cloud native images, such as the Billing Care image, from the Oracle Software Delivery website:

1. Go to <https://edelivery.oracle.com>.
2. Sign in to the Oracle Software Delivery website using an Oracle account.
3. Search for and select **Oracle Communications Billing and Revenue Management Cloud Native Deployment Option 15.0.x.0.0** and then click **Continue**.
4. Select the following and then click **Continue**:
  - Oracle Communications Billing and Revenue Management Cloud Native Deployment Option 15.0.x.0.0-CN
  - Oracle Communications Pricing Design Center Cloud Native Deployment Option 15.0.x.0.0-CN
  - Oracle Communications Elastic Charging Engine Cloud Native Deployment Option 15.0.x.0.0-CN
5. Accept the Oracle standard terms and restrictions and then click **Continue**.
6. Select the packages listed in [Table 5-3](#) and then click **Download**.  
Each package is downloaded to a separate Zip file.
7. Extract the package files listed in [Table 5-3](#) from each Zip file.

**Table 5-3 BRM Cloud Native Packages and Package Files**

BRM Package Name	Package File Name
Oracle Communications Cloud Native Batch Controller	<b>oc-cn-brm-batch-controller-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Batch Pipeline	<b>oc-cn-brm-batch-pipeline-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Billing Care	<b>oc-cn-billingcare-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Billing Care REST API	<b>oc-cn-bcws-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native BRM Applications	<b>oc-cn-brm-apps-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native BRM REST Services Manager	<b>oc-cn-brm-rest-services-manager-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native BRM SDK	<b>oc-cn-brm-sdk-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Business Operations Center	<b>oc-cn-boc-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Connection Manager	<b>oc-cn-brm-cm-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Database Upgrade	<b>oc-cn-brm-upgrade-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Elastic Charging Engine Cloud Native Deployment Option	<b>oc-cn-ece-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Email Data Manager	<b>oc-cn-brm-dm-email-15.0.x.0.0.tar</b>

**Table 5-3 (Cont.) BRM Cloud Native Packages and Package Files**

BRM Package Name	Package File Name
Oracle Communications Cloud Native Enterprise Application Integration Data Manager	<b>oc-cn-brm-dm-eai-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Enterprise Application Integration Java Server	<b>oc-cn-brm-eai-js-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Fusa Data Manager	<b>oc-cn-brm-dm-fusa-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Fusa Simulator	<b>oc-cn-brm-fusa-simulator-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Invoice Data Manager	<b>oc-cn-brm-dm-invoice-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Invoice Formatter	<b>oc-cn-brm-invoice-formatter-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Kafka Data Manager	<b>oc-cn-brm-dm-kafka-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native LDAP Data Manager	<b>oc-cn-brm-dm-ldap-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Oracle Database Manager	<b>oc-cn-brm-dm-oracle-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Performance Profiling Toolkit	<b>oc-cn-brm-perflib-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Pipeline Configuration Center	<b>oc-cn-pcc-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Pricing Design Center	<b>oc-cn-pdc-15.0.x.0.0.tar</b> <b>oc-cn-pdc-rsm-15.0.x.0.0.tar</b> <b>oc-cn-pdc-rsm-jars-15.0.x.0.0.tar.gz</b>
Oracle Communications Cloud Native Provisioning Data Manager	<b>oc-cn-brm-dm-prov-telco-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Rated Event Loader	<b>oc-cn-brm-rel-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Rated Event Manager	<b>oc-cn-brm-rem-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Real-Time Pipeline	<b>oc-cn-brm-realtime-pipeline-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Roaming Manager	<b>oc-cn-brm-roam-pipeline-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Vertex Data Manager	<b>oc-cn-brm-dm-vertex-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Webhook	<b>oc-cn-brm-webhook-15.0.x.0.0.tar</b>
Oracle Communications Cloud Native Web Services Manager	<b>oc-cn-brm-wsm-wls-15.0.x.0.0.tar</b> <b>oc-cn-brm-wsm-wl-init-15.0.x.0.0.tar</b>

8. Load each package file as an image into the Podman system using the following command:

```
podman load --input fileName
```

where *fileName* is the package file name listed in [Table 5-3](#).

For example, to load the Kafka DM image in the Podman system, enter this command:

```
podman load --input oc-cn-brm-dm-kafka-15.0.x.0.0.tar
```

If you use an internal registry to access images from different Kubernetes nodes, push the images from your local system to the registry server. For example, if the registry is identified by `RepoHost:RepoPort`, you'd push the Kafka DM image to the registry using the Podman CLI like this:

1. Tag the Kafka DM image with the registry server:

```
podman tag dm-kafka:15.0.x.0.0 RepoHost:RepoPort/dm-kafka:15.0.x.0.0
```

2. Push the Kafka DM image to the registry server:

```
podman push RepoHost:RepoPort/dm-kafka:15.0.x.0.0
```

## Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center

If you use PDC, Billing Care, Billing Care REST API, or Business Operations Center, pull the Oracle WebLogic image from the Oracle Container Registry into your private repository.

To load the Oracle WebLogic image into your private repository:

1. In a web browser, go to <https://container-registry.oracle.com>.
2. Sign in to the Oracle Container Registry using an Oracle account.

### Note

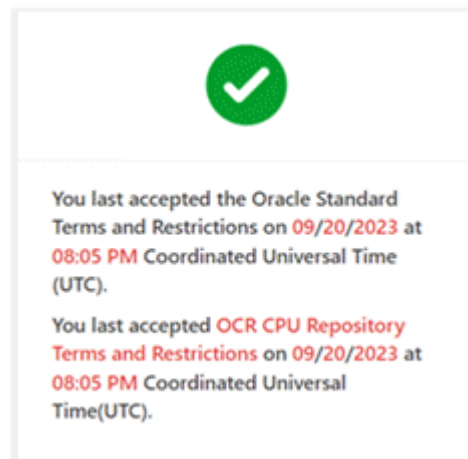
To pull images for licensed software on the Oracle Container Registry, you must have an Oracle account. You can create an Oracle account at <https://profile.oracle.com/myprofile/account/create-account.jspx>.

3. Click the **Middleware** container, and then click the **fmw-infrastructure\_cpu** CPU repository.

The repository page appears.

4. Accept the Oracle terms and restrictions by clicking **Continue** and, on the next page, clicking **Accept**.

If successful, you will see something similar to this:



5. On the host system, log in to the Oracle Container Registry using the Podman CLI:

```
podman login container-registry.oracle.com
```

6. When prompted for a user name and password, enter your Oracle credentials.
7. Pull the WebLogic image into your local system using the following command:

```
podman pull container-registry.oracle.com/middleware/fmw-infrastructure_cpu:12.2.1.4-jdk8-ol8
```

8. Tag the image with the registry server using the following command, where *RepoHost* and *RepoPort* are the host and port of your private repository:

```
podman tag container-registry.oracle.com/middleware/fmw-infrastructure_cpu:12.2.1.4-jdk8-ol8 RepoHost:RepoPort/middleware/fmw-infrastructure_cpu:12.2.1.x-jdk8-ol8
```

9. Push the image to the registry server using the following command, where *RepoHost* and *RepoPort* are the host and port of your private repository:

```
podman push RepoHost:RepoPort/middleware/fmw-infrastructure_cpu:12.2.1.4-jdk8-ol8
```

# Part III

## Configuring and Deploying BRM Cloud Native

This part provides information about configuring and deploying Oracle Communications Billing and Revenue Management (BRM) in your cloud native environment. It contains the following chapters:

- [Deploying the BRM Database Schema](#)
- [Configuring BRM Server, PDC, and PCC Services](#)
- [Configuring REST Services](#)
- [Configuring the Billing Care, Billing Care REST API, and Business Operations Center Services](#)
- [Configuring ECE Services](#)
- [Deploying BRM Cloud Native Services](#)
- [Deploying into Oracle Cloud Infrastructure](#)
- [Uninstalling Your BRM Cloud Native Deployment](#)

# 6

## Deploying the BRM Database Schema

Learn how to deploy a new or existing database schema in the Oracle Communications Billing and Revenue Management (BRM) cloud native environment.

Topics in this document:

- [Deploying BRM with a New Database Schema](#)
- [Deploying BRM with an Existing Schema](#)

### Deploying BRM with a New Database Schema

To deploy BRM with a new BRM and pipeline database schema:

1. If you have not already done so, create a BRM database and schema users for the BRM Server and pipeline. See "[Creating and Configuring Your BRM Database](#)".
2. Create a new Kubernetes namespace for **oc-cn-init-db-helm-chart**:  

```
kubectl create namespace InitDbNameSpace
```
3. Create an **override-values.yaml** file for **oc-cn-init-db-helm-chart**, and add keys from [Table 6-1](#).

#### Note

This documentation uses the **override-values.yaml** file name for ease of use, but you can name the file whatever you want.

4. If you want to deploy the BRM cloud native schema into a multischema database, do the following in your **override-values.yaml** file:
  - a. Set the **ocbrm.db.skipPrimary** key to **false**.
  - b. For each secondary schema in your system, add an **ocbrm.db.multiSchemas.secondary $N$**  block, where  $N$  is **1** for the first secondary schema, **2** for the next secondary schema, and so on.
  - c. In each **ocbrm.db.multiSchemas.secondary $N$**  block, set the following keys:

#### Note

If the **host**, **port**, and **service** keys are not defined, the secondary schema uses the same host name, service, and port number as the primary schema.

- **deploy**: Set this to **true** to deploy this secondary schema.
- **host**: Set this to the host name of the secondary schema. This key is optional.
- **port**: Set this to the port number for the secondary schema. This key is optional.

- **service:** Set this to the service name for the secondary schema. This key is optional.
  - **schemauser:** Set this to the schema user name.
  - **schemapass:** Set this to the schema password.
  - **schematablespace:** Set this to the name of the schema tablespace, such as pin01.
  - **indextablespace:** Set this to the name of the index tablespace, such as pinx01.
- d. Enable account migration between your schemas by setting the **ocbrm.isAmt** key to **true**.

See "Understanding Account Migration" in *BRM Moving Accounts Between Database Schemas* for more information.

This shows example entries for a BRM database with one primary schema, two secondary schemas, and account migration enabled:

```
ocbrm:
  isAmt: true
ocbrm:
  db:
    host: hostname
    port: 12345
    service: serviceName
    schemauser: pin01
    schemapass: password
    schematablespace: pin01
    indextablespace: pinx01
    skipPrimary: false
    multiSchemas:
      secondary1:
        deploy: true
        schemauser: pin02
        schemapass: password
        schematablespace: pin02
        indextablespace: pinx02
      secondary2:
        deploy: true
        schemauser: pin03
        schemapass: password
        schematablespace: pin03
        indextablespace: pinx03
```

5. To commit seed data to the BRM database, modify the `init-db` container (**configmap\_create\_obj\_2.yaml**) to add the corresponding `PCM_OP_CREATE_OBJ` input list. For example:

```
<PCM_OP $PIN_OPNAME=$PIN_CONF_INIT_OPNAME; $PIN_OPFLAGS=0>
0 PIN_FLD_POID          POID [0] $DB_NUMBER /config/recharge_card_type 0 0
0 PIN_FLD_NAME          STR [0] "-"
0 PIN_FLD_PROGRAM_NAME STR [0] "load_pin_recharge_card_type"
0 PIN_FLD_HOSTNAME      STR [0] "-"
0 PIN_FLD_VERSION       STR [0] "1"
0 PIN_FLD_ACCOUNT_OBJ   POID [0] 0.0.0.1 /account 1 0
</PCM_OP>
```

6. Initialize the BRM database schema by running this command from the **helmcharts** directory:

```
helm install InitDbReleaseName oc-cn-init-db-helm-chart --namespace InitDbNameSpace
--values OverrideValuesFile
```

where:

- *InitDbReleaseName* is the release name for **oc-cn-init-db-helm-chart** and is used to track this installation instance.
- *InitDbNameSpace* is the namespace for **oc-cn-init-db-helm-chart**.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-init-db-helm-chart**.

The init-db Helm chart creates an init-db pod and job for each schema.

7. After the init-db Helm chart deploys successfully, delete the Helm chart from your namespace by running this command from the **helmcharts** directory:

```
helm delete InitDbReleaseName -n InitDbNameSpace
```

[Table 6-1](#) lists the keys that directly impact the BRM database schema and pipeline database schema. Add these keys to your **override-values.yaml** file with the same path hierarchy.

### Caution

Some keys hold sensitive data. They must be handled carefully with controlled access to the file containing its values. Encode all of these values in Base64. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 6-1 oc-cn-init-db-helm-chart Keys**

Key	Purpose
<b>imageRepository</b>	The registry server where you have pushed images. Typically, in the format "RepoHost:RepoPort". The value is added as a prefix to all image names when you install or upgrade Helm charts. This key is empty by default.
<b>imagePullSecrets</b>	The name of the Secret containing credentials for accessing images from your private image server. This is added to each pod to give it permission to pull the image from your private registry server. See " <a href="#">Creating Secrets for Docker Registry Authorization</a> " for more information.
<b>uniPass</b>	Use this key to apply a uniform password to all BRM cloud native services, including: <ul style="list-style-type: none"> <li>• Database Schemas</li> <li>• BRM Root Login</li> <li>• Oracle Wallets</li> <li>• WebLogic User</li> </ul> To override this password for a specific service, specify a different password in the service's key. <b>Note:</b> Use this key for test or demonstration systems only.
<b>db.sslMode</b>	The type of SSL connection required for connecting to the database: <b>TWO_WAY</b> , <b>ONE_WAY</b> , or <b>NO</b> . The default is <b>ONE_WAY</b> .
<b>ocbrm.imagePullPolicy</b>	Specify when to pull images: <ul style="list-style-type: none"> <li>• <b>IfNotPresent:</b> Pulls an image only if one is not present locally. This is the default.</li> <li>• <b>Always:</b> Always pulls an image.</li> </ul>

Table 6-1 (Cont.) oc-cn-init-db-helm-chart Keys

Key	Purpose
<b>ocbrm.isAmt</b>	Specify whether account migration is enabled in your BRM multischema database: <ul style="list-style-type: none"> <li><b>True:</b> Account migration is enabled.</li> <li><b>False:</b> Account migration is disabled. This is the default.</li> </ul>
<b>ocbrm.isIPV6Enabled</b>	Specify whether IPV6 is enabled in your Kubernetes environment: <ul style="list-style-type: none"> <li><b>True:</b> IPV6 is enabled.</li> <li><b>False:</b> IPV6 is disabled. This is the default.</li> </ul>
<b>ocbrm.ece_deployed</b>	Specify whether to deploy ECE in your BRM cloud native environment: <ul style="list-style-type: none"> <li><b>True:</b> ECE configurations will be loaded into your environment. This is the default.</li> <li><b>False:</b> ECE configurations will not be loaded into your environment.</li> </ul>
<b>ocbrm.pdc_deployed</b>	Specify whether to deploy PDC in your BRM cloud native environment: <ul style="list-style-type: none"> <li><b>True:</b> PDC configurations will be loaded into your environment.</li> <li><b>False:</b> PDC configurations will not be loaded into your environment. This is the default.</li> </ul>
<b>ocbrm.existing_rootkey_wallet</b>	Specify whether you are using an existing BRM database or an existing root key wallet: <ul style="list-style-type: none"> <li><b>True:</b> Uses your existing BRM database or root key wallet.</li> <li><b>False:</b> Performs a fresh database initialization. This is the default. When set to <b>false</b>, the <code>init-db</code> Helm chart runs <b>lds-config-job</b> to load default strings into the BRM database during deployment. Manually delete <b>lds-config-job</b> after BRM deploys successfully.</li> </ul>
<b>ocbrm.is_upgrade</b>	Specify whether to upgrade the existing database schema: <ul style="list-style-type: none"> <li><b>True:</b> Upgrades your existing database schema.</li> <li><b>False:</b> Deploys a new BRM database schema. This is the default.</li> </ul>
<b>ocbrm.isSSEnabled</b>	Specify whether SSL is enabled in the BRM database. The default is <b>true</b> .
<b>ocbrm.cmSSLTermination</b>	Specify whether to disable SSL between the CM and DM/EM. <ul style="list-style-type: none"> <li><b>true:</b> The CM is the SSL endpoint. Set it to <b>true</b> when using a custom TLS certificate for the CM. When set to <b>true</b>, also set the <b>isSSEnabled</b> key to <b>true</b>.</li> <li><b>false:</b> SSL is enabled across all components, from client to CM to DMs and EMs. This is the default.</li> </ul>
<b>ocbrm.customSSLWallet</b>	Whether to use a custom TLS certificate for the CM. <ul style="list-style-type: none"> <li><b>true:</b> A custom TLS certificate is used for the CM. When set to <b>true</b>, also set the <b>cmSSLTermination</b> key to <b>true</b> and move the Oracle wallet (<code>brm_custom_wallet</code>) containing the custom TLS certificate to the top level of the Helm chart.</li> <li><b>false:</b> The default TLS certificate is used for the CM. This is the default.</li> </ul>
<b>ocbrm.root_key_rotate</b>	Whether to rotate the BRM root key. The default is <b>false</b> . For more information, see "Rotating the BRM Root Key" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>ocbrm.brm_root_pass</b>	The BRM root password. The value can be per user. You must generate a Secret. Add this key to your <b>override-values.yaml</b> file with the same path hierarchy.
<b>ocbrm.rotate_brm_role_passwords</b>	Whether to rotate the BRM role passwords. The default is <b>false</b> .

Table 6-1 (Cont.) oc-cn-init-db-helm-chart Keys

Key	Purpose
<b>ocbrm.brm_role_pass.*</b>	<p>The initial passwords for each BRM role. The roles grant users the permission to access different BRM components, such as Customer Center or Pipeline Configuration Center (PCC).</p> <p><b>Note:</b> You must set all of these passwords when the <a href="#">unipass</a> key is not set.</p> <p>You set passwords for the following roles:</p> <ul style="list-style-type: none"> <li>• <b>acct_recv.0.0.0.1:</b> Accounts receivable</li> <li>• <b>bc_client.0.0.0.1:</b> Billing Care</li> <li>• <b>bill_inv_pymt_sub.0.0.0.1:</b> Invoice payments</li> <li>• <b>billing.0.0.0.1:</b> Billing</li> <li>• <b>boc_client.0.0.0.1:</b> Business Operations Center</li> <li>• <b>collections.0.0.0.1:</b> Collections</li> <li>• <b>crypt_utils.0.0.0.1:</b> Encryption utilities</li> <li>• <b>cust_center.0.0.0.1:</b> Customer Center</li> <li>• <b>cust_mgnt.0.0.0.1:</b> Customer management</li> <li>• <b>invoicing.0.0.0.1:</b> Invoicing</li> <li>• <b>java_client.0.0.0.1:</b> Java clients</li> <li>• <b>load_utils.0.0.0.1:</b> Load utilities</li> <li>• <b>payments.0.0.0.1:</b> Payments</li> <li>• <b>pcc_client.0.0.0.1:</b> PCC</li> <li>• <b>rerating.0.0.0.1:</b> Rerating</li> <li>• <b>rsm.0.0.0.1:</b> BRM REST Services Manager</li> <li>• <b>super_user.0.0.0.1:</b> Super User</li> <li>• <b>ui_client.0.0.0.1:</b> All thick clients</li> <li>• <b>ece.0.0.0.1:</b> ECE</li> </ul> <p><b>Note:</b> After you deploy or upgrade the database schema, you should rotate the role passwords regularly to improve security. To do so, see "Rotating BRM Role Passwords" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>ocbrm.wallet.*</b>	<p>Specify the passwords for these wallets:</p> <ul style="list-style-type: none"> <li>• <b>client:</b> The password for the client wallet.</li> <li>• <b>server:</b> The password for the server wallet.</li> <li>• <b>root:</b> The password for the root wallet.</li> </ul> <p>You must generate Secrets for these keys.</p>

Table 6-1 (Cont.) oc-cn-init-db-helm-chart Keys

Key	Purpose
ocbrm.db.*	<p>The details for connecting to a shared database. The keys in this section take precedence over other database connection keys.</p> <p>Add these keys only if your system uses a shared database:</p> <ul style="list-style-type: none"> <li>• <b>host</b>: The host name of the machine on which the BRM database is configured.</li> <li>• <b>port</b>: The port on which the BRM database is configured.</li> <li>• <b>service</b>: The BRM database service name.</li> <li>• <b>sslMode</b>: The type of SSL connection required for connecting to the database: <b>TWO_WAY</b>, <b>ONE_WAY</b>, or <b>NO</b>. The default is <b>ONE_WAY</b>.</li> <li>• <b>walletPassword</b>: The password for accessing the database certificates from the TrustStore and KeyStore.</li> <li>• <b>walletType</b>: The type of file specified as the TrustStore or KeyStore for database SSL connections: <b>SSO</b> or <b>pkcs12</b>.</li> <li>• <b>enable_partition</b>: Whether partitioning is enabled in your BRM database: <b>Yes</b> or <b>No</b>. The default is <b>Yes</b>.</li> <li>• <b>storage_model</b>: The size of the BRM database tablespaces: <b>Test</b> (less than 700 MB), <b>Small</b> (less than 1.5 GB), <b>Medium</b> (less than 30 GB), or <b>Large</b> (greater than 30 GB).</li> <li>• <b>schemauser</b>: The user name for the primary BRM schema. The user name should be pre-created with all of the required grants.</li> <li>• <b>schemapass</b>: The password for the database schema.</li> <li>• <b>schematablespace</b>: The name of the schema tablespace.</li> <li>• <b>indextablespace</b>: The name of the index tablespace.</li> <li>• <b>nls_lang</b>: The language, territory, and character set. Set this to <b>American_America.characterset</b>, where <i>characterset</i> is either <b>UTF8</b> or <b>AL32UTF8</b>.</li> </ul> <p><b>Note:</b> You must use <b>American_America</b> as the language and territory, regardless of your locale.</p> <ul style="list-style-type: none"> <li>• <b>pipelineschemauser</b>: The BRM pipeline schema user name, which should be pre-created with all of the required grants.</li> <li>• <b>pipelineschemapass</b>: The BRM pipeline schema password. You must generate a Secret. Add this key to your <b>override-values.yaml</b> file with the same path hierarchy.</li> <li>• <b>pipelineschematablespace</b>: The name of the tablespace for the BRM pipeline schema. This field is case-sensitive.</li> <li>• <b>pipelineindextablespace</b>: The name of the index tablespace for the BRM pipeline schema. This field is case-sensitive.</li> <li>• <b>skipPrimary</b>: Whether to skip the deployment of the primary schema: <b>False</b> (deploy the primary schema) or <b>True</b> (do not deploy the primary schema).</li> </ul> <p>Set <b>skipPrimary</b> to <b>true</b> only if you are adding a schema to an existing BRM multischema system. See "Adding Schemas to a Multischema System" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>

Table 6-1 (Cont.) oc-cn-init-db-helm-chart Keys

Key	Purpose
<b>ocbrm.db.multiSchemas.secondaryN</b>	<p>The details for connecting to your secondary schemas, where <i>N</i> is <b>1</b> for the first secondary schema, <b>2</b> for the next secondary schema, and so on.</p> <p>Add this block only if your BRM database contains multiple schemas. This section will be commented out by default:</p> <pre>multiSchemas:   secondary1:     deploy: true     host: localhost     port: 1521     service: pindb     schemauser:     schemapass:     schematablespace:     indextablespace:</pre>

## Deploying BRM with an Existing Schema

To deploy BRM with an existing schema and with default Oracle images:

1. Copy the root-key wallet files from the **\$BRM\_WALLET/client** directory to the **oc-cn-helm-chart/existing\_wallet** directory.
2. In your **override-values.yaml** file for **oc-cn-helm-chart**, set the **ocbrm.existing\_rootkey\_wallet** key to **true**.
3. Deploy **oc-cn-helm-chart**.

Alternatively, you could deploy BRM with an existing schema by doing this:

1. Create images for each BRM Server component in the installed BRM staging area (using the same staging area that initialized the database).
2. In your **override-values.yaml** file for **oc-cn-helm-chart**, update the keys with the existing schema credentials and also set the following keys:
  - **ocbrm.existing\_rootkey\_wallet**: Set this to **true**.
  - **ocbrm.use\_oracle\_brm\_images**: Set this to **false**.
  - **ocbrm.db.queueName**: Set this to match the queue name in the existing database.
  - **ocbrm.db.roamqueueName**: If you want a database queue for Roaming Pipeline, create another queue in the Oracle database by following the instructions in "Creating Additional Queues for Multischema BRM Systems" in *BRM Installation Guide*. Then, set the **roamqueueName** key to the name of the queue you created.
3. Deploy **oc-cn-helm-chart**.

# 7

## Configuring BRM Server, PDC, and PCC Services

Learn how to configure the Oracle Communications Billing and Revenue Management (BRM) server, Pricing Design Center (PDC), and Pipeline Configuration Center (PCC) in your cloud native environment.

Topics in this document:

- [About Configuring BRM Cloud Native Services](#)
- [Creating Secrets for Docker Registry Authorization](#)
- [Configuring Global Values](#)
- [Specifying the BRM Services to Deploy](#)
- [Configuring the BRM Server](#)
- [Configuring BRM for a Multischema Database](#)
- [Configuring Pricing Design Center](#)
- [Configuring Pipeline Configuration Center](#)

### About Configuring BRM Cloud Native Services

You configure and deploy BRM cloud native services, such as BRM, PDC, and PCC, by using the BRM Helm chart (**oc-cn-helm-chart**). YAML descriptors in the **oc-cn-helm-chart/templates** directory use the **oc-cn-helm-chart/values.yaml** file for most of the values. The **values.yaml** file itself includes comments that describe each key. You can override the values by creating an **override-values.yaml** file.

Oracle recommends that you read the **values.yaml** file at least once to become familiar with all of the options available.

### Creating Secrets for Docker Registry Authorization

You can automatically pull images from your private container registry by creating an **ImagePullSecrets**, which contains a list of authorization tokens (or Secrets) for accessing a private container registry. You then add references to the **ImagePullSecrets** in your BRM Helm chart's **override-values.yaml** file. This allows pods to submit the Secret to the private container registry whenever they want to pull images.

Automatically pulling images from a private container registry involves these high-level steps:

1. Create a Secret outside of the Helm chart by entering this command:

```
kubectl create secret docker-registry SecretName --docker-server=RegistryServer --  
docker-username=UserName --docker-password=Password --namespace NameSpace
```

where:

- *SecretName* is the name of your Kubernetes Secret

- *RegistryServer* is your private container registry's FQDN (*repoHost:repoPort*)
- *UserName* and *Password* are your private container registry's user name and password
- *Namespace* is the namespace you will use for installing BRM Helm charts

For example:

```
kubectl create secret docker-registry cgbu-docker-registry --docker-
server=mydockerimages.com:2660/ --docker-username=xyz --docker-password=password --
namespace oms
```

2. Add the **imagePullSecrets** key to your **override-values.yaml** file for **oc-cn-helm-chart**:

```
imagePullSecrets:
  -name "SecretName1"
  -name "SecretName2"
```

3. Add the **ocbrm.imageRepository** key to your **override-values.yaml** file:

```
imageRepository: "RegistryServer"
```

4. Deploy **oc-cn-helm-chart**.

## Configuring Global Values

[Table 7-1](#) lists the keys that apply to all BRM components. To set or change the values, add them to your **override\_values.yaml** file for **oc-cn-helm-chart**.

**Table 7-1 Global Keys in Values.yaml File**

Key	Description
<b>imageRepository</b>	The registry server where you have pushed images. Typically, in the format "RepoHost:RepoPort". The value is added as a prefix to all image names when you install or upgrade Helm charts. This key is empty by default.
<b>imagePullSecrets</b>	The name of the Secret that contains credentials for accessing images from your private image server. This is added to each pod to give it permission to pull the image from your private registry server. See " <a href="#">Creating Secrets for Docker Registry Authorization</a> " for more information. This key is empty by default.
<b>uniPass</b>	Use this key to apply a uniform password to all BRM cloud native services, including: <ul style="list-style-type: none"> <li>• Database Schemas</li> <li>• BRM Root Login</li> <li>• Oracle Wallets</li> <li>• WebLogic User</li> </ul> To override this password for a specific service, specify a different password in the service's key. <b>Note:</b> Use this key for test or demonstration systems only.

Table 7-1 (Cont.) Global Keys in Values.yaml File

Key	Description
<b>db.*</b>	<p>The details for connecting to a shared database. The keys in this section take precedence over other database connection keys.</p> <p>Add these keys only if your system uses a shared database:</p> <ul style="list-style-type: none"> <li>• <b>sslMode</b>: The type of SSL connection required for connecting to the database: <ul style="list-style-type: none"> <li>– <b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li>– <b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li>– <b>NO</b>: SSL authentication is not required.</li> </ul> </li> <li>• <b>host</b>: The host name or IP address of the database server.</li> <li>• <b>port</b>: The port number of the database server.</li> <li>• <b>user</b>: The user name of the database administrator.</li> <li>• <b>password</b>: The password of the database system administrator.</li> <li>• <b>serviceName</b>: The service name that identifies the database.</li> <li>• <b>role</b>: The role assigned to the DBA user.</li> <li>• <b>walletPassword</b>: The password for accessing the certificates from the TrustStore and KeyStore. This is required if <b>sslMode</b> is set to <b>ONE_WAY</b> or <b>TWO_WAY</b>.</li> <li>• <b>walletType</b>: The type of file specified as the TrustStore or KeyStore for SSL connections: <b>SSO</b> or <b>pkcs12</b>.</li> </ul>
<b>security.java.overrideSecurityProperties</b>	Whether to override the default Java security property ( <b>true</b> ) or no ( <b>false</b> ). The default is <b>true</b> .
<b>monitoring.prometheus.jmx_exporter.enable</b>	Whether to enable the JMX exporter for Prometheus ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>false</b> .  See "Monitoring BRM Cloud Native Services" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.
<b>monitoring.prometheus.operator.*</b>	<p>The details for monitoring BRM cloud native services using Prometheus:</p> <ul style="list-style-type: none"> <li>• <b>enable</b>: Whether to use Prometheus Operator (<b>true</b>) or standalone Prometheus (<b>false</b>). The default is <b>false</b>.</li> <li>• <b>namespace</b>: The namespace in which Prometheus Operator is deployed. The default is <b>prometheus</b>.</li> <li>• <b>release</b>: The release name for Prometheus Operator. The default is <b>prometheus</b>.</li> </ul> <p>See "Monitoring BRM Cloud Native Services" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.</p>

## Specifying the BRM Services to Deploy

Some BRM cloud native services are enabled by default, while others are disabled. Ensure that your **override-values.yaml** file is set up to deploy the services that you want to include in your BRM cloud native environment.

### BRM Cloud Native Services Enabled by Default

[Table 7-2](#) lists the BRM cloud native services that are deployed by default. To exclude them from your deployment, set the keys to **false** in your **override-values.yaml** file for **oc-cn-helm-chart**.

**Table 7-2 BRM Services Enabled By Default**

BRM Service	override-values.yaml Key
Batch Pipeline	<code>ocbrm.batchpipe.isEnabled</code>
BRM REST Services Manager	<code>ocrsm.rsm.isEnabled</code>
Connection Manager	<code>ocbrm.cm.isEnabled</code>
Oracle Data Manager	<code>ocbrm.dm_oracle.isEnabled</code>
Pipeline Configuration Center	<code>ocpcc.pcc.isEnabled</code>
Pricing Design Center	<code>ocpdc.isEnabled</code>
Rated Event Loader	<code>ocbrm.rel_daemon.isEnabled</code>
Realtime Pipeline	<code>ocbrm.realtimepipe.isEnabled</code>

**BRM Cloud Native Services Disabled By Default**

[Table 7-3](#) lists the BRM cloud native services that are *not* deployed by default. To include them in your BRM cloud native deployment, set the keys to **true** in your `override-values.yaml` file for `oc-cn-helm-chart`.

**Table 7-3 BRM Services Disabled By Default**

BRM Service	override-values.yaml Key
Batch Controller	<code>ocbrm.batch_controller.isEnabled</code>
BRM Apps Jobs	<code>ocbrm.brm_apps.job.isEnabled</code>
Email Data Manager	<code>ocbrm.dm_email.isEnabled</code>
Enterprise Application Integration Data Manager	<code>ocbrm.dm_eai.isEnabled</code>
Invoicing Formatter	<code>ocbrm.formatter.isEnabled</code>
Invoicing Data Manager	<code>ocbrm.dm_invoice.isEnabled</code>
Kafka Data Manager	<code>ocbrm.dm_kafka.isEnabled</code>
LDAP Data Manager	<code>ocbrm.dm_ldap.isEnabled</code>
Paymentech Data Manager	<code>ocbrm.dm_fusa.isEnabled</code>
PDC REST Services Manager	<code>ocpdcrsm.isEnabled</code>
Provisioning Data Manager	<code>ocbrm.dm_prov_telco.isEnabled</code>
RE Loader Daemon Job	<code>ocbrm.rel_daemon.job.isEnabled</code>
RE Loader Manager Job	<code>ocbrm.rel_manager.job.isEnabled</code>
Rated Event Manager	<code>ocbrm.rem.isEnabled</code>
Roaming Pipeline	<code>ocbrm.roampipe.isEnabled</code>
Vertex Data Manager	<code>ocbrm.dm_vertex.isEnabled</code>
Webhook	<code>webhook.isEnabled</code>
Web Services Manager with WebLogic	<code>ocbrm.wsm.deployment.weblogic.isEnabled</code>
Web Services Manager with TomCat	<code>ocbrm.wsm.deployment.tomcat.isEnabled</code>

## Configuring the BRM Server

To configure the BRM server to run in your cloud native environment, you override the BRM server-specific keys in the **values.yaml** file for **oc-cn-helm-chart**. [Table 7-4](#) lists the keys that directly impact BRM Server pods. Add these keys to your **override-values.yaml** file with the same path hierarchy.

### Note

You can optionally deploy a simple demonstration version of BRM cloud native by using the sample **override\_values.yaml** file that is packaged with **oc-cn-helm-chart**. This sample override file contains the bare minimum keys that you need to update to create a simple BRM cloud native system with the following services enabled by default: Account Synchronization DM, Batch Pipeline, CM, Oracle DM, RE Loader, Realtime Pipeline, Billing Care, Billing Care REST API, Business Operations Center, and PDC.

**Table 7-4 BRM Server Keys**

Key	Path in values.yaml File	Description
<b>isAmt</b>	<b>ocbrm</b>	Whether account migration is enabled in your BRM database ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>false</b> .
<b>isIPv6Enabled</b>	<b>ocbrm</b>	Whether IPv6 is enabled in your Kubernetes environment ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>false</b> .
<b>ece_deployed</b>	<b>ocbrm</b>	Whether ECE is going to be deployed in your BRM cloud native environment ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>true</b> .
<b>pdcc_deployed</b>	<b>ocbrm</b>	Whether PDC is going to be deployed: <ul style="list-style-type: none"> <li><b>true</b>: Configuration data is <i>not</i> loaded into the BRM database. Only mandatory configuration records are loaded into the BRM database for starting the realtime pipeline pod. For the batch pipeline and roaming pipeline pods, you must load the required configuration data using PDC before deploying the pods.</li> <li><b>false</b>: Configuration data is loaded into the BRM database during deployment. This is the default.</li> </ul>
<b>use_oracle_brm_images</b>	<b>ocbrm</b>	Whether to use the default BRM images ( <b>true</b> ) or not ( <b>false</b> ). Set this to <b>false</b> if you are building custom images. The default is <b>true</b> .

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
<b>existing_rootkey_wallet</b>	<b>ocbrm</b>	<p>Whether you are deploying with an existing database or using an existing root key wallet:</p> <ul style="list-style-type: none"> <li><b>true</b>: You are deploying with an existing database or are using an existing root key wallet.</li> <li><b>false</b>: You are deploying with a new database and are using a new root key wallet. This is the default.</li> </ul> <p>When set to <b>false</b>, the BRM Helm chart runs <b>lds-config-job</b> to load default strings into BRM during the deployment process. Manually delete <b>lds-config-job</b> after BRM is deployed successfully.</p> <p>See "Rotating the BRM Root Key" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.</p>
<b>is_upgrade</b>	<b>ocbrm</b>	<p>Whether to upgrade the Helm chart (<b>true</b>) or not (<b>false</b>).</p> <p>See "<a href="#">Upgrading Your BRM Cloud Native Services</a>".</p>
<b>isSSLEnabled</b>	<b>ocbrm</b>	For SSL-enabled deployment required in <b>Infranet.properties</b> .
<b>cmSSLTermination</b>	<b>ocbrm</b>	<p>Whether to make the CM the SSL endpoint for the BRM cloud native deployment.</p> <ul style="list-style-type: none"> <li><b>true</b>: The CM is the SSL endpoint. In this case, TLS can be enabled only between BRM client applications and the CM. TLS is disabled between CM and all downstream components such as DMs and EMs. Communication between external clients and the CM will still be encrypted by TLS 1.2. This setting can increase performance, because it eliminates the overhead needed to handle TLS before processing the PCP packets.</li> </ul> <p>When set to <b>true</b>, also set the <b>isSSLEnabled</b> key to <b>true</b>.</p> <ul style="list-style-type: none"> <li><b>false</b>: SSL is enabled across all components, from client to CM to DMs and EMs. This is the default.</li> </ul>
<b>customSSLWallet</b>	<b>ocbrm</b>	<p>Whether to use a custom TLS certificate for the CM:</p> <ul style="list-style-type: none"> <li><b>true</b>: A custom TLS certificate is used for the CM.</li> </ul> <p>When set to <b>true</b>, also set the <b>cmSSLTermination</b> key to <b>true</b> and move the Oracle wallet (<b>brm_custom_wallet</b>) containing the custom TLS certificate to the top level of the Helm chart.</p> <ul style="list-style-type: none"> <li><b>false</b>: The default TLS certificate is used for the CM. This is the default.</li> </ul> <p>See "Using a Custom TLS Certificate" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>EnableSecurityContext</b>	<b>ocbrm</b>	Whether to enable a security context in the cluster ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>false</b> .
<b>root_key_rotate</b>	<b>ocbrm</b>	<p>Whether to rotate the BRM root key (<b>true</b>) or not (<b>false</b>). The default is <b>false</b>.</p> <p>See "Rotating the BRM Root Key" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>brm_root_pass</b>	<b>ocbrm</b>	<p>The root password.</p> <p>See "Rotating the BRM Root Password" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
<b>rotate_password</b>	<b>ocbrm</b>	<p>Whether to rotate the BRM root password:</p> <ul style="list-style-type: none"> <li><b>true</b>: The BRM root password is replaced with the one specified in the <b>new_brm_root_password</b> key.</li> <li><b>false</b>: The BRM root password is not changed. This is the default.</li> </ul> <p>See "Rotating the BRM Root Password" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>new_brm_root_password</b>	<b>ocbrm</b>	<p>The new BRM root password. Use this key only when <b>ocbrm.rotate_password</b> is set to <b>true</b>.</p> <p>See "Rotating the BRM Root Password" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>rotate_brm_role_passwords</b>	<b>ocbrm</b>	<p>Whether to rotate the BRM role passwords. The default is <b>false</b>.</p> <p>See "Rotating BRM Role Passwords" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>brm_role_pass.*</b>	<b>ocbrm</b>	<p>The passwords for each BRM role. The roles grant users the permission to access different BRM components, such as Customer Center or Pipeline Configuration Center (PCC).</p> <p><b>Note:</b> You must set all of these passwords when the <a href="#">unipass</a> key is not set.</p> <p>You set passwords for the following roles:</p> <ul style="list-style-type: none"> <li><b>acct_recv.0.0.0.1</b>: Accounts receivable</li> <li><b>bc_client.0.0.0.1</b>: Billing Care</li> <li><b>bill_inv_pymt_sub.0.0.0.1</b>: Invoice payments</li> <li><b>billing.0.0.0.1</b>: Billing</li> <li><b>boc_client.0.0.0.1</b>: Business Operations Center</li> <li><b>collections.0.0.0.1</b>: Collections</li> <li><b>crypt_utils.0.0.0.1</b>: Encryption utilities</li> <li><b>cust_center.0.0.0.1</b>: Customer Center</li> <li><b>cust_mgnt.0.0.0.1</b>: Customer management</li> <li><b>invoicing.0.0.0.1</b>: Invoicing</li> <li><b>java_client.0.0.0.1</b>: Java clients</li> <li><b>load_utils.0.0.0.1</b>: Load utilities</li> <li><b>payments.0.0.0.1</b>: Payments</li> <li><b>pcc_client.0.0.0.1</b>: PCC</li> <li><b>rerating.0.0.0.1</b>: Rerating</li> <li><b>rsm.0.0.0.1</b>: BRM REST Services Manager</li> <li><b>super_user.0.0.0.1</b>: Super User</li> <li><b>ui_client.0.0.0.1</b>: All GUI clients</li> <li><b>ece.0.0.0.1</b>: ECE</li> </ul> <p>The passwords in this key must match the passwords in <b>oc-cn-init-db-helm-chart</b>.</p> <p>See "Rotating BRM Role Passwords" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
wallet.*	ocbrm	Specify the passwords for these wallets: <ul style="list-style-type: none"> <li><b>client</b>: The password for the client wallet.</li> <li><b>server</b>: The password for the server wallet.</li> <li><b>root</b>: The password for the root wallet.</li> </ul> You must generate Secrets for these keys.
deployment.*	ocbrm.cm	<ul style="list-style-type: none"> <li><b>enable_publish</b>: Whether to publish events (1) or not (0). The default is 0.</li> <li><b>enable_prefs_enrichment</b>: Whether to enrich notifications with subscriber preferences (true) or not (false). The default is false.</li> <li><b>prefs_enabled_publisher_list</b>: The list of publishers with enrichment enabled. The default is 0.0.9.6.</li> <li><b>prefs_phone_no_location</b>: Where to retrieve the phone numbers for subscribers. The default is 0.</li> <li><b>provisioning_enabled</b>: Whether to enable provisioning of service orders (true) or not (false). The default is false.</li> <li><b>simulate_agent</b>: Whether to publish service orders (0) or not (1). The default is 1.</li> </ul>
enable	ocbrm.cm.custom_files	Whether to expose the <b>oc-cn-helm-chart/cm_custom_files</b> directory as a ConfigMap (true) or not (false). The default is false. See "Exposing Directories as ConfigMaps" in <i>BRM Cloud Native System Administrator's Guide</i> .
perflib_enabled	ocbrm.component.deployment	Whether to enable monitoring of the BRM service using the performance library (Perlib), where <i>component</i> is <b>cm</b> or <b>dm_oracle</b> . <ul style="list-style-type: none"> <li><b>true</b>: PerfLib is enabled for the specified component. The PerfLib container will expose the metrics for the specified component, which can be scraped by Prometheus.</li> <li><b>false</b>: PerfLib is disabled for the specified component. This is the default.</li> </ul> See "Monitoring BRM Cloud Native Services" in <i>BRM Cloud Native System Administrator's Guide</i> .
isEnabled	ocbrm.dm_kafka	Whether to enable the Kafka DM (true) or not (false). The default is false. For more information about integrating BRM cloud native with a Kafka Server, see "Integrating with Kafka Servers" in <i>BRM Cloud Native System Administrator's Guide</i> .

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
<b>deployment.*</b>	<b>ocbrm.dm_kafka</b>	<p>The details for configuring the Kafka DM.</p> <ul style="list-style-type: none"> <li><b>kafka_bootstrap_server_list</b>: Set this to a comma-separated list of addresses for the Kafka brokers in this format: <i>hostname1:port1, hostname2:port2</i>.</li> <li><b>poolSize</b>: Set this to the number of threads that can run in the JS server to accept requests from the CM. Enter a number from 1 through 2000. The default is <b>64</b>.</li> <li><b>topicName</b>: Set this to the name of the default Kafka topic. The default name is <b>BRM</b>.</li> <li><b>topicFormat</b>: Set this to the format of the payload that is published to the default Kafka topic: <b>XML</b> or <b>JSON</b>.</li> <li><b>topicStyle</b>: The style of XML payloads: <b>ShortName</b>, <b>CamelCase</b>, <b>NewShortName</b>, or <b>OC3CNotification</b>.</li> <li><b>isSecurityEnabled</b>: Whether to enable SSL between Kafka DM and Kafka Server (<b>true</b>) or not (<b>false</b>).</li> <li><b>trustStorePassword</b>: The TrustStore password in Base64 format. Required only if SSL is enabled.</li> <li><b>keyStorePassword</b>: The KeyStore password in Base64 format. Required only if SSL is enabled.</li> <li><b>keyPassword</b>: The password of a key in the KeyStore in Base64 format. Required only if SSL is enabled.</li> <li><b>password</b>: The password in Base64 format. Required only if SSL is enabled.</li> </ul> <p>For more information about integrating BRM cloud native with a Kafka Server, see "Integrating with Kafka Servers" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>smtpServer</b>	<b>ocbrm.dm_email.deployment</b>	Set this to your SMTP server name, such as <code>ocbrm.us.example.com</code> .
<b>create</b>	<b>ocbrm.storage_class</b>	Whether to create a Kubernetes StorageClass ( <b>true</b> ) or not ( <b>false</b> ).
<b>virtual_time.*</b>	<b>ocbrm</b>	<p>The details for configuring the <b>pin_virtual_time</b> utility.</p> <ul style="list-style-type: none"> <li><b>enabled</b>: Set this to <b>true</b> to enable <b>pin_virtual_time</b>.</li> <li><b>sync_pvt_time</b>: Set this to the number of seconds between each synchronization of <b>pin_virtual_time</b> with all pods. The default is <b>0</b> seconds.</li> </ul>

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
db.*	ocbrm	<p>The details for connecting to the BRM database:</p> <ul style="list-style-type: none"> <li>• <b>host</b>: The host name or IP address of the database server.</li> <li>• <b>port</b>: The port number of the database server.</li> <li>• <b>service</b>: The service name that identifies the database.</li> <li>• <b>sslMode</b>: The type of SSL connection required for connecting to the database: <b>TWO_WAY</b>, <b>ONE_WAY</b>, or <b>NO</b>.</li> <li>• <b>walletPassword</b>: The password for accessing the certificates from the TrustStore and KeyStore. This is required if <b>sslMode</b> is set to <b>ONE_WAY</b> or <b>TWO_WAY</b>.</li> <li>• <b>walletType</b>: The type of file specified as the TrustStore or KeyStore for SSL connections: <b>SSO</b> or <b>pkcs12</b>. The default is <b>SSO</b>.</li> <li>• <b>enable_partition</b>: Whether partitioning is enabled at the database level (<b>Yes</b>) or disabled (<b>No</b>). The default is <b>Yes</b>.</li> <li>• <b>storage_model</b>: The size of the BRM database tablespaces: <ul style="list-style-type: none"> <li>– <b>Test</b>: Less than 700 MB.</li> <li>– <b>Small</b>: Less than 1.5 GB.</li> <li>– <b>Medium</b>: Less than 30 GB.</li> <li>– <b>Large</b>: Greater than 30 GB.</li> </ul> </li> <li>• <b>schemauser</b>: The user name of the primary BRM schema. The default is <b>pin</b>.</li> <li>• <b>schemapass</b>: The password for the BRM schema.</li> <li>• <b>schematablespace</b>: The name of the tablespace for the primary BRM schema. This field is case-sensitive. The default is <b>pin</b>.</li> <li>• <b>indextablespace</b>: The name of the index tablespace for the primary BRM schema. This field is case-sensitive. The default is <b>pinx</b>.</li> <li>• <b>nls_lang</b>: The language, territory, and character set. Set this to <b>American_America.AL32UTF8</b>. You must use <b>American_America</b> as the language and territory, regardless of your locale.</li> <li>• <b>pipelineschemauser</b>: The BRM pipeline schema user name, which should be pre-created with all of the required grants.</li> <li>• <b>pipelineschemapass</b>: The BRM pipeline schema password. You must generate a Secret. Add this key to your <b>override-values.yaml</b> file with the same path hierarchy.</li> <li>• <b>pipelineschematablespace</b>: The name of the tablespace for the BRM pipeline schema. This field is case-sensitive.</li> <li>• <b>pipelinesindextablespace</b>: The name of the index tablespace for the BRM pipeline schema. This field is case-sensitive.</li> <li>• <b>skipPrimary</b>: Whether to deploy the primary schema (<b>false</b>) or skip the deployment of the primary schema (<b>true</b>). Set it to <b>true</b> only if you are adding a schema to</li> </ul>

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
		<p>an existing BRM multischema system. The default is <b>false</b>.</p> <p>Ensure these values match the <b>ocbrm.db.*</b> keys from <b>oc-cn-init-db-helm-chart</b>. See <a href="#">Table 6-1</a> for more information.</p>
<b>secondaryN.*</b>	<b>ocbrm.db.multiSchemas</b>	<p>The details for connecting to your secondary database schemas, where <i>N</i> is <b>1</b> for the first secondary schema, <b>2</b> for the next secondary schema, and so on.</p> <p>Add this block only if your BRM database contains multiple schemas. This section will be commented out by default:</p> <pre>multiSchemas:   secondary1:     deploy:       host: localhost       port: 1521       service: pindb       schemauser:       schemapass:       schematablespace:       indextablespace:</pre> <p>See "<a href="#">Configuring BRM for a Multischema Database</a>".</p>
<b>mountOptions</b>	<b>ocbrm.storage_class</b>	Set this to the version of the external provisioner.
<b>provisioner</b>	<b>ocbrm.dynamic_provisioner</b>	Set this to the name of the external provisioner.
<b>enabled</b>	<b>ocbrm.cmt</b>	Set this to <b>true</b> to run the Conversion Manager <b>pin_cmt</b> utility.
<b>config_jobs.*</b>	<b>ocbrm</b>	<p>The details for running a configurator job, which allows you to run BRM load utilities on demand without entering into a pod.</p> <ul style="list-style-type: none"> <li><b>run_apps</b>: Set to <b>true</b> to enable a configurator job.</li> <li><b>isMultiSchema</b>: Specifies whether to run the commands in the <b>loadme.sh</b> script on the secondary schemas.</li> <li><b>restart_count</b>: Increment this count by 1 to restart the CM.</li> <li><b>script_name</b>: The name of the script that contains the load utilities you want to run. The default is <b>loadme.sh</b>.</li> </ul> <p>See "Running Load Utilities through Configurator Jobs" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
<b>brm_apps.job.*</b>	<b>ocbrm</b>	<p>The details for running a brm-apps job, which allows you to run BRM applications and utilities on demand without entering into a pod.</p> <ul style="list-style-type: none"> <li><b>isEnabled</b>: Set to <b>true</b> to enable a brm-apps job.</li> <li><b>script_name</b>: The name of the script that contains the utilities and applications you want to run. The default is <b>loadme.sh</b>.</li> <li><b>configmap_path</b>: The path to the ConfigMap file.</li> </ul> <p>See "Running Applications and Utilities through BRM-Apps Jobs" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>

Table 7-4 (Cont.) BRM Server Keys

Key	Path in values.yaml File	Description
<code>brm_apps.deployment.isEnabled</code>	<code>ocbrm</code>	Whether to enable the brm-apps pod.
<code>utilityName.*</code>	<code>ocbrm.brm_apps.deployment</code>	The configuration details for running BRM utilities and applications, such as <code>pin_billd</code> , <code>pin_export_price</code> , and <code>pin_rerate</code> .  For example, you can configure performance parameters for multithreaded applications (MTA). See "Configuring MTA Performance Parameters" in <i>BRM Cloud Native System Administrator's Guide</i> .

## Configuring BRM for a Multischema Database

Using a BRM multischema database lets you distribute customer accounts among several database schemas, providing increased storage capacity, higher performance, and easier maintenance. For more information, see "A BRM Multischema Production System" in *BRM Installation Guide*.

To configure your BRM cloud native environment to connect to a multischema database, do this:

1. Ensure that you deployed a multischema database in your BRM cloud native environment. See "[Deploying BRM with a New Database Schema](#)".
2. Using SQL\*Plus, grant each database schema the privilege to insert and update tables on the other schemas:

- a. Connect to the BRM database with SQL\*Plus as **sysdba**:

```
sqlplus system@databaseAlias as sysdba
Enter password: password
```

where:

- `databaseAlias` is the Oracle system database alias.
- `password` is the Oracle system database user password.

- b. From the primary database schema, enter the following:

```
SQL> GRANT INSERT ANY TABLE TO SecondarySchema;
SQL> GRANT UPDATE ANY TABLE TO SecondarySchema;
```

where `SecondarySchema` is the name of the secondary schema.

- c. From each secondary database schema, enter the following:

```
SQL> GRANT INSERT ANY TABLE TO PrimarySchema;
SQL> GRANT UPDATE ANY TABLE TO PrimarySchema;
```

where `PrimarySchema` is the name of the primary schema.

3. Connect the BRM server to each secondary schema:

- a. Open your **override-values.yaml** file for **oc-cn-helm-chart**.
- b. Enable account migration by setting the **ocbrm.isAmt** key to **true**.
- c. Set the **ocbrm.db.skipPrimary** key to **false**.
- d. For each secondary schema in your system, add a **ocbrm.db.multiSchemas.secondaryN** block, where *N* is **1** for the first secondary schema, **2** for the next secondary schema, and so on.
- e. In each **ocbrm.db.multiSchemas.secondaryN** block, set the following keys:
  - **deploy**: Set this to **true**.
  - **host**: Set this to the hostname of the secondary schema. This key is optional.
  - **port**: Set this to the port number for the secondary schema. This key is optional.
  - **service**: Set this to the service name for the secondary schema. This key is optional.
  - **schemauser**: Set this to the schema user name.
  - **schemapass**: Set this to the schema password.
  - **schematablespace**: Set this to the name of the schema tablespace, such as `pin01`.
  - **indextablespace**: Set this to the name of the index tablespace, such as `pinx01`.
- f. Deploy **oc-cn-helm-chart** by running this command from the **helmcharts** directory:

```
helm install BrmReleaseName oc-cn-helm-chart --namespace BrmNameSpace --values OverrideValuesFile
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance. It must be different from the one used for **oc-cn-init-db-helm-chart**.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

The BRM Helm chart deploys new `dm-oracle`, `amt`, and `rel-dameon` pods, Rated Event (RE) Loader PVCs, services, ConfigMaps, and secrets. It also updates their corresponding schema entries in the primary CM and Oracle DM and deploys multiple containers for the batch-wireless-pipe pod.

4. Set each database schema's status and priority. BRM cloud native assigns accounts to an open schema with the highest priority.
  - a. Open the **configmap\_pin\_conf\_testnap.yaml** file.
  - b. Under the **config\_dist.conf** section, add the following entries for each secondary schema in your database:

```
DB_NO = "schema_number" ;           # database config. block
PRIORITY = priority ;
MAX_ACCOUNT_SIZE = 100000 ;
STATUS = "status" ;
SCHEMA_NAME = "schema_name" ;
```

- c. Set the **STATUS** and **PRIORITY** entries for each primary and secondary schema:

```

DB_NO = "0.0.0.1" ;           # Primary schema configuration block
PRIORITY = priority;
MAX_ACCOUNT_SIZE = 100000 ;
STATUS = "status" ;
SCHEMA_NAME = "schema_name" ;

DB_NO = "0.0.0.2" ;           # Secondary schema configuration block
PRIORITY = priority;
MAX_ACCOUNT_SIZE = 50000 ;
STATUS = "status" ;
SCHEMA_NAME = "schema_name" ;

```

where:

- *priority* is a number representing the schema's priority, with the highest number having the most priority. For example, 5 indicates a greater priority than a value of 1. For more information, see "Modifying Database Schema Priorities" in *BRM Cloud Native System Administrator's Guide*.
  - *status* specifies whether the schema is **open**, **closed**, or **unavailable**. For more information, see "Modifying Database Schema Status" in *BRM Cloud Native System Administrator's Guide*.
- d. Set up the configurator job to run the **load\_config\_dist** utility by adding the following lines to the **oc-cn-helm-chart/config\_scripts/loadme.sh** script:

```

#!/bin/sh

#cp /oms/config_dist.conf /oms/sys/test/config_dist.conf
cd /oms/sys/test ; load_config_dist
exit 0;

```

- e. In the **override-values.yaml** file for **oc-cn-helm-chart**, set this key:  
**ocbrm.config\_jobs.run\_apps**: Set this to **true**.
- f. Run the **helm upgrade** command to update the Helm release:

```

helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --
namespace BrmNameSpace

```

The distribution information is loaded into the primary schema.

- g. Update these keys in the **override-values.yaml** file for **oc-cn-helm-chart**:
- **ocbrm.config\_jobs.restart\_count**: Increment the existing value by 1.
  - **ocbrm.config\_jobs.run\_apps**: Set this to **false**.
- h. Update the Helm release again:

```

helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --
namespace BrmNameSpace

```

The CM is restarted.

5. Configure the account-router Pipeline Manager to route CDRs to pipelines based on the database schema POID. To do so, edit the ConfigMap file **configmap\_acc\_router\_reg.yaml**.

Based on the configuration, the account router Pipeline Manager does the following:

- Moves input files to the **data** PVC directory. The input file names have a prefix of **router** and a suffix of **.edr**.
- Moves the rated output files to the input of the Rating pipeline.

- Replicates the Rating pipeline based on the multischema entry. The Range function is used to replicate the rating pipeline.
- Moves the output files from the Rating pipeline to the **outputcdr** PVC directory.

Your BRM cloud native environment is connected to your BRM multischema database. To manage your multischema environment, see "Managing a BRM Cloud Native Multischema System" in *BRM Cloud Native System Administrator's Guide*.

## Configuring Pricing Design Center

Pricing Design Center (PDC) is a Web-based client application that you use to create and manage the product offerings that you sell to your customers. A product offering represents the services available to your customers and the price of those services. For more information about PDC, see *Pricing Design Center Online Help*.

You can optionally deploy a simple demonstration version of Pricing Design Center cloud native by using the sample **PDC\_OverrideValues.yaml** file provided with **oc-cn-helm-chart**. This simple demonstration version has both SSL and ECE enabled, uploads a sample JKS certificate file, loads sample RUMs and balance elements, and starts the BRM-to-PDC synchronization process with **SyncPDC**.

To configure PDC to run in your BRM cloud native environment:

1. Override the PDC-specific keys in the **values.yaml** file for **oc-cn-helm-chart**. See "[Adding PDC Keys for oc-cn-helm-chart](#)".
2. Override the PDC-specific keys in the **values.yaml** file for **oc-cn-op-job-helm-chart**. See "[Adding PDC Keys for oc-cn-op-job-helm-chart](#)".
3. Set up SAML for SSO in PDC. See "[Setting Up SSO for PDC Cloud Native](#)".

After you deploy PDC in your cloud native environment, you can access the PDC GUI at one of the following URLs:

- `http://kubernetesHost:pdPort/pdc`

where:

- *kubernetesHost* is the host name of the machine on which Kubernetes is deployed.
- *pdPort* is the PDC service node port.

- `http://LoadBalancerHost:pdNodePort/pdc`

where:

- *LoadBalancerHost* is the host name of the machine on which the load balancer is deployed.
- *pdNodePort* is the number assigned to the PDC node port.

## Adding PDC Keys for oc-cn-helm-chart

[Table 7-5](#) describes the most common PDC keys that you need to override. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart** with the same path hierarchy.

For information about all PDC-specific keys, see the descriptions in the **oc-cn-helm-chart/values.yaml** file.

**⚠ Caution**

Keys with the path **ocpdc.secretValue** hold sensitive data. Handle them carefully with controlled access to the file containing their values. Encode all of these values in Base64. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 7-5 Pricing Design Center Keys for oc-cn-helm-chart**

Key	Path in values.yaml	Description
<b>isEnabled</b>	<b>ocpdc</b>	Whether to enable and deploy PDC: <ul style="list-style-type: none"> <li><b>true</b>: Enables PDC and deploys the PDC application. This is the default.</li> <li><b>false</b>: Disables the PDC application.</li> </ul>
<b>lang</b>	<b>ocpdc</b>	The Linux system locale. The default is <b>en_US.UTF-8</b> .
<b>tz</b>	<b>ocpdc</b>	The Linux time zone. The default is <b>UTC</b> .
<b>pdcbmVolHostPath</b>	<b>ocpdc</b>	The host path for RRE, Import-Export, BRE, or SyncPDC. To use a dynamic provisioner, leave it empty. The default is empty. <b>Note:</b> Provide the required permissions to the volume path by following the guidelines in " <a href="#">Persistent Volume Storage Locations</a> " in the WebLogic Kubernetes Operator documentation.
<b>storageSize</b>	<b>ocpdc</b>	The size of the storage path for <b>pdcbmHostPath</b> .
<b>enableSecurityContext</b>	<b>ocpdc</b>	Whether to enable PDC pod-level security attributes and common container settings. The default is <b>false</b> .
<b>deployment.*</b>	<b>ocpdc</b>	The details for deploying the PDC image: <ul style="list-style-type: none"> <li><b>imageName</b>: The name of the PDC image. The default is <b>oracle/pdc</b>.</li> <li><b>imageTag</b>: The tag name for the PDC image. The default is <b>15.0.x.0.0</b>.</li> <li><b>imagePullPolicy</b>: When to pull images: only when one is not present locally (<b>IfNotPresent</b>) or always (<b>Always</b>). The default is <b>IfNotPresent</b>.</li> </ul>
<b>fmw.*</b>	<b>ocpdc.deployment</b>	The details for pulling WebLogic images for PDC from a container registry: <ul style="list-style-type: none"> <li><b>imageRepository</b>: The name of the container registry from which to pull the WebLogic image. The default is <b>container-registry.oracle.com/</b>.</li> <li><b>imageName</b>: The name of the container repository from which to pull the WebLogic image. The default is <b>middleware/fmw-infrastructure_cpu</b>.</li> <li><b>imageTag</b>: The tag name for the WebLogic image. The default is <b>12.2.1.4-jdk8-ol8</b>.</li> </ul> See " <a href="#">Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center</a> " for more information.

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
<b>nodeSelector</b>	<b>ocpdc</b>	<p>The name of the node on which to run the following PDC pods:</p> <ul style="list-style-type: none"> <li>• pdc</li> <li>• RRE</li> <li>• BRE</li> <li>• SyncPDC</li> <li>• Import-Export</li> </ul> <p>Set this key if you want to constrain the PDC pods to run only on the node you specify.</p> <p>For more information, see "<a href="#">nodeSelector</a>" in the Kubernetes documentation.</p> <p><b>Note:</b> To override the rules for a specific PDC pod, specify a different value for the pod's <b>nodeSelector</b> key. For example, set the <b>ocpdc.configEnv.transformation.syncPDC.nodeSelector</b> key to apply rules specifically to the SyncPDC pod.</p>
<b>affinity</b>	<b>ocpdc</b>	<p>The rules for running the following PDC pods on specific nodes:</p> <ul style="list-style-type: none"> <li>• pdc</li> <li>• RRE</li> <li>• BRE</li> <li>• SyncPDC</li> <li>• Import-Export</li> </ul> <p>Set this key if you want to constrain the PDC pod to run only on the nodes that meet your criteria.</p> <p>For more information about this key, see "<a href="#">Node Affinity</a>" in the Kubernetes documentation.</p> <p><b>Note:</b> To override the rules for a specific PDC pod, specify a different value for the pod's <b>affinity</b> key. For example, set the <b>ocpdc.configEnv.transformation.importExport.affinity</b> key to apply rules specifically to the Import-Export pod.</p>
<b>domainUID</b>	<b>ocpdc.wop</b>	The name of this PDC WebLogic Server domain.
<b>isVPAEnabled</b>	<b>ocpdc.wop</b>	(Release 15.0.1 or later) Whether to enable resource limits for the PDC domain pod. The default is <b>false</b> .
<b>isVPAEnabled</b>	<b>ocpdc.configEnv</b>	Whether to enable Vertical Pod autoscaling for the RRE, BRE, SyncPDC, and JobIE pods. The default is <b>false</b> .
<b>dbHostName</b>	<b>ocpdc.configEnv</b>	The host name of the PDC and cross-reference database. The value must match that of <b>oc-cn-op-job-helm-chart</b> .
<b>dbPort</b>	<b>ocpdc.configEnv</b>	The port for the PDC and cross-reference database. The value must match that of <b>oc-cn-op-job-helm-chart</b> .
<b>dbService</b>	<b>ocpdc.configEnv</b>	The service name for the PDC and cross-reference database. The value must match that of <b>oc-cn-op-job-helm-chart</b> .

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
<b>dbSSLMode</b>	<b>ocpdc.configEnv</b>	The type of SSL connection required for connecting to the cross-reference database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required.</li> <li><b>NO</b>: SSL authentication is not required. This is the default.</li> </ul> <b>Note</b> : This value must match that of the <b>dbSSLMode</b> key for <b>oc-cn-op-job-helm-chart</b> .
<i>Realtime rating engine and batch rating engine log file rotation</i>	<b>ocpdc.configEnv.transformation</b>	The settings for rotating RRE and BRE log files: <ul style="list-style-type: none"> <li><b>logLevel</b>: The logging level, which can be <b>SEVERE</b>, <b>WARNING</b>, <b>INFO</b>, <b>CONFIG</b>, <b>FINE</b>, <b>FINER</b>, or <b>FINEST</b>. The default is <b>WARNING</b>.</li> <li><b>logFileSize</b>: The maximum file size, in bytes, of the log files. After the log file reaches its maximum size, PDC closes the current log file and creates a new one. The default is <b>500000</b>.</li> <li><b>logFileCount</b>: The maximum number of log files to retain for the application. The default is <b>10</b>.</li> <li><b>persistTransactionLogs</b>: Whether to persist log files in the database after they are closed. Possible values are <b>all</b>, <b>disabled</b>, and <b>failed</b>. The default is <b>failed</b>.</li> </ul> See "Rotating PDC Log Files" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.
<b>MEM_ARGS</b>	<b>ocpdc.configEnv.transformation</b>	The memory argument, surrounded by quotes. For example: <b>"-Xms1024m -Xmx2048m -XX:CompileThreshold=8000"</b> .
<b>persistOutFiles</b>	<b>ocpdc.configEnv.transformation</b>	Whether to persist the output files generated by the transformation engine: <ul style="list-style-type: none"> <li><b>enabled</b>: Keeps the output and payload files in the <b>out</b> directory. This is not recommended.</li> <li><b>disabled</b>: Removes the output and payload files from the <b>out</b> directory. This is the default.</li> </ul>
<b>upgrade</b>	<b>ocpdc.configEnv.transformation</b>	Whether to upgrade from a previous release to PDC 15.0.
<b>nodeSelector</b>	<b>ocpdc.configEnv.transformation</b>	The name of the node on which to run the RRE and BRE pods.
<b>affinity</b>	<b>ocpdc.configEnv.transformation</b>	The rules for running the RRE and BRE pods on specific nodes.
<b>BE</b>	<b>ocpdc.configEnv.seedData</b>	Whether to load sample balance elements into the PDC database when PDC is deployed: <ul style="list-style-type: none"> <li><b>true</b>: Loads the sample balance elements.</li> <li><b>false</b>: Does not load the sample balance elements. This is the default.</li> </ul> <b>Note</b> : If balance element data already exists in the PDC database, it is not overwritten.

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
RUM	ocpdc.configEnv.seedData	<p>Whether to load sample RUMs into the PDC database when PDC is deployed:</p> <ul style="list-style-type: none"> <li><b>true</b>: Loads the sample RUMs.</li> <li><b>false</b>: Does not load the sample RUMs. This is the default.</li> </ul> <p><b>Note</b>: If RUM data already exists in the PDC database, it is not overwritten.</p>
IE_Operation.*	ocpdc.configEnv.importExport	<p>The operation for the <b>ImportExportPricing</b> utility to perform:</p> <ul style="list-style-type: none"> <li>Empty value: No operation is performed. This is the default.</li> <li><b>export</b>: The utility exports data from the PDC database into an XML file.</li> <li><b>import</b>: The utility imports data from the XML file into the PDC database.</li> <li><b>publish</b>: The utility publishes components from PDC to the batch rating engine, real-time rating engine, or ECE.</li> <li><b>keep</b>: The utility retains the latest version of successfully promoted PDC components.</li> <li><b>deleteprofile</b>: The utility deletes pricing profiles from PDC.</li> <li><b>type</b>: The utility displays the pricing or setup components available in PDC.</li> </ul> <p>When you deploy PDC, ensure that this key has an empty value.</p> <p>Before doing a Helm upgrade to run <b>ImportExportPricing</b>, delete the <b>pdcc-import-export-job</b> Kubernetes job.</p> <p>Don't include the hyphen (-) prefix with the value.</p> <p>For more information, see "Running PDC Applications" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>
IE_Component.*	ocpdc.configEnv.importExport	<p>The type of components and objects to import or export using the <b>ImportExportPricing</b> utility:</p> <ul style="list-style-type: none"> <li><b>config</b>: Imports or exports setup components, such as tax codes, business profiles, and general ledger IDs.</li> <li><b>pricing</b>: Imports or exports pricing components, such as events, charges, and chargeshares.</li> <li><b>metadata</b>: Imports or exports the event, service, account, and profile attribute specifications.</li> <li><b>profile</b>: Imports or exports pricing profile data.</li> <li><b>customfields</b>: Imports or exports custom fields.</li> <li><b>brmobject</b>: Exports BRM-mastered configuration objects such as services, events, and G/L IDs.</li> <li><b>all</b>: Imports or exports all objects and components.</li> </ul> <p>Don't include the hyphen (-) prefix with the value.</p> <p>For more information, see "Running PDC Applications" in <i>BRM Cloud Native System Administrator's Guide</i>.</p>

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
<b>IE_File_OR_Dir_Name</b>	<b>ocpdc.configEnv.importExport</b>	<p>The name of the XML file or <b>ImportExport</b> directory that contains the list of components and objects to import into the PDC database. This XML file is used by the <b>ImportExportPricing</b> utility.</p> <p>If importing or deleting components, copy the XML file to the HostPath specified in <b>pdcBRMHostPath</b> or to pdc-brm-pvc. Don't include the hyphen (-) prefix with the value.</p> <p>Set the appropriate read and write permissions for files or directories, and set the ownership to <i>runAsUser</i> using <b>chown runAsUser:0</b>.</p>
<b>extraCmdLineArgs</b>	<b>ocpdc.configEnv.importExport</b>	<p>The extra command-line arguments for the <b>ImportExportPricing</b> utility, apart from operation, component, and file name.</p> <p>The value must be surrounded by quotes. For example, "<b>-n ObjectName</b>".</p> <p>See "ImportExportPricing" in <i>PDC Creating Product Offerings</i> for more information.</p>
<i>ImportExport Log Rotation</i>	<b>ocpdc.configEnv.importExport</b>	<p>The settings for rotating ImportExport log files:</p> <ul style="list-style-type: none"> <li><b>logLevel</b>: Sets the logging level, which can be <b>SEVERE</b>, <b>WARNING</b>, <b>INFO</b>, <b>CONFIG</b>, <b>FINE</b>, <b>FINER</b>, or <b>FINEST</b>. The default is <b>WARNING</b>.</li> <li><b>logSize</b>: Sets the maximum file size, in bytes, of the log files. After the log file reaches its maximum size, PDC closes the current log file and creates a new one.</li> <li><b>logCount</b>: Specifies the maximum number of log files to retain for the application.</li> <li><b>persistIELogs</b>: Specifies whether to persist log files in the database after they are closed. Possible values are <b>all</b>, <b>disabled</b>, and <b>failed</b>.</li> </ul> <p>See "Rotating PDC Log Files" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.</p>
<b>nodeSelector affinity</b>	<b>ocpdc.configEnv.importExport</b>	The rules for deploying the Import-Export pod on specific nodes.
<b>upgradeFromPS2</b>	<b>ocpdc.configEnv.syncPDC</b>	<p>Whether to upgrade <b>SyncPDC</b> during the PDC upgrade process. Set this to <b>true</b> only if you are upgrading from Patch Set 2 to Patch Set 4.</p> <p>The default is <b>false</b>.</p> <p>See "<a href="#">Upgrading Your PDC Cloud Native Services</a>" for more information.</p>
<b>skipBREMigration</b>	<b>ocpdc.configEnv.syncPDC</b>	<p>Skips the synchronization of pipeline configuration data. The default is <b>false</b>.</p> <p>This key is ignored when ECE is enabled in a PDC system.</p>

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
<b>syncPDCStartAt</b>	<b>ocpdc.configEnv.syncPDC</b>	The scheduled time for running the <b>SyncPDC</b> utility. This key is set at deployment time only. Valid values include: <ul style="list-style-type: none"> <li><b>startAt</b>: The utility runs at the time the job was submitted.</li> <li><b>"HH:MM"</b>: The utility runs at the specified time, where <i>HH</i> is an hour between 0 and 23, and <i>MM</i> is the minutes between 0 and 59. For example, enter <b>"12:00"</b> to schedule the utility to run at noon.</li> </ul> For more information, see "Running PDC Applications" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>syncPDCInterval</b>	<b>ocpdc.configEnv.syncPDC</b>	The scheduled frequency for running the <b>SyncPDC</b> utility. This key is set at deployment time only. Enter a value in the format <b>"N:U"</b> , where <i>N</i> is a valid number and <i>U</i> is one of these units: <b>D</b> (Daily), <b>H</b> (Hourly), or <b>M</b> (Minute). For example, enter <b>"2:D"</b> to run the utility every other day. <b>Note</b> : A value of <b>"24:H"</b> is not the same as <b>"1:D"</b> due to daylight savings time (DST).
<b>enrichmentFileName</b>	<b>ocpdc.configEnv.syncPDC</b>	Set this to <b>ECEEEventEnrichmentSpec.xml</b> . Store the enrichment file in the path specified in <b>pdcbRmHostPath</b> . This is applicable at both PDC deployment time and individual <b>SyncPDC</b> runtime.
<b>runSyncPDC</b>	<b>ocpdc.configEnv.syncPDC</b>	Whether to create the SyncPDC pod: <ul style="list-style-type: none"> <li><b>true</b>: Creates the SyncPDC pod and starts the BRM-to-PDC synchronization process. This is the default.</li> <li><b>false</b>: Deletes the SyncPDC pod and stops the synchronization process.</li> </ul>
<i>SyncPDC Log Rotation</i>	<b>ocpdc.configEnv.syncPDC</b>	The settings for rotating <b>SyncPDC</b> log files: <ul style="list-style-type: none"> <li><b>logLevel</b>: Sets the logging level, which can be <b>SEVERE</b>, <b>WARNING</b>, <b>INFO</b>, <b>CONFIG</b>, <b>FINE</b>, <b>FINER</b>, or <b>FINEST</b>. The default is <b>WARNING</b>.</li> <li><b>logFileSize</b>: Sets the maximum file size, in bytes, of the log files. After the log file reaches its maximum size, PDC closes the current log file and creates a new one. The default is <b>20000</b>.</li> <li><b>logFileCount</b>: Specifies the maximum number of log files to retain for the application. The default is <b>10</b>.</li> </ul> See "Rotating PDC Log Files" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.
<b>nodeSelector affinity</b>	<b>ocpdc.configEnv.syncPDC</b>	The rules for deploying the <b>SyncPDC</b> pod on specific nodes.

Table 7-5 (Cont.) Pricing Design Center Keys for oc-cn-helm-chart

Key	Path in values.yaml	Description
<b>operator.*</b>	<b>ocpdc.configEnv.monitoring.prometheus</b>	The details for monitoring PDC cloud native services using Prometheus: <ul style="list-style-type: none"> <li><b>isEnabled:</b> Whether to enable monitoring of PDC by using Prometheus and Grafana.</li> <li><b>namespace:</b> The namespace in which Prometheus Operator is deployed. The default is <b>prometheus</b>.</li> </ul> For more information, see "Monitoring PDC in a Cloud Native Environment" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>walletPassword</b>	<b>ocpdc.secretValue</b>	The passwords for the PDC application wallet and PDC BRM integration wallet.

## Adding PDC Keys for oc-cn-op-job-helm-chart

You must create an **override-values.yaml** file for **oc-cn-op-job-helm-chart** and then add the PDC-specific keys in [Table 7-6](#).

For information about all PDC-specific keys, see the descriptions in the **oc-cn-op-job-helm-chart/values.yaml** file.

### Caution

Keys with the path **ocpdc.secretValue** hold sensitive data. Handle them carefully with controlled access to the file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

Table 7-6 Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<b>isEnabled</b>	<b>ocpdc</b>	Whether to enable PDC jobs: <ul style="list-style-type: none"> <li><b>true:</b> Enables PDC jobs. This is the default.</li> <li><b>false:</b> Disables PDC jobs.</li> </ul>
<b>isClean</b>	<b>ocpdc</b>	Whether to clean old PDC deployment and instance logs: <ul style="list-style-type: none"> <li><b>true:</b> Removes all existing PDC deployment and instance logs. This is the default.</li> <li><b>false:</b> Keeps all existing PDC logs.</li> </ul>
<b>lang</b>	<b>ocpdc</b>	The Linux system locale. The default is <b>en_US.UTF-8</b> .
<b>tz</b>	<b>ocpdc</b>	The Linux time zone. The default is <b>UTC</b> .

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<b>pdAppVolHostPath</b>	<b>ocpdc</b>	<p>The host path for pdc-domain.</p> <p>To use a dynamic provisioner, leave it empty. The default is empty.</p> <p><b>Note:</b> For a dynamic provisioner, ensure that <b>oc-cn-helm-chart</b> is deployed and that the storage class is reused.</p> <p><b>Note:</b> Provide the required permissions to the volume path by following the guidelines in "<a href="#">Persistent Volume Storage Locations</a>" in the WebLogic Kubernetes Operator documentation.</p>
<b>storageSize</b>	<b>ocpdc</b>	The size of the storage path for <b>pdHostPath</b> .
<b>enableSecurityContext</b>	<b>ocpdc</b>	<p>Whether to enable PDC pod-level security attributes and common container settings.</p> <p>The default is <b>false</b>.</p>
<b>deployment.*</b>	<b>ocpdc</b>	<p>The details for deploying the PDC image:</p> <ul style="list-style-type: none"> <li><b>imageName:</b> The name of the PDC image. The default is <b>oracle/pdc</b>.</li> <li><b>imageTag:</b> The tag name for the PDC image. The default is <b>15.0.0.0.0</b>.</li> <li><b>imagePullPolicy:</b> When to pull images: only when one is not present locally (<b>IfNotPresent</b>) or always (<b>Always</b>). The default is <b>IfNotPresent</b>.</li> </ul>
<b>fmw.*</b>	<b>ocpdc.deployment</b>	<p>The details for pulling WebLogic images for PDC from a container registry:</p> <ul style="list-style-type: none"> <li><b>imageRepository:</b> The name of the container registry from which to pull the WebLogic image. The default is <b>container-registry.oracle.com/</b>.</li> <li><b>imageName:</b> The name of the container repository from which to pull the WebLogic image. The default is <b>middleware/fmw-infrastructure_cpu</b>.</li> <li><b>imageTag:</b> The tag name for the WebLogic image. The default is <b>12.2.1.4-jdk8-ol8</b>.</li> </ul> <p>See "<a href="#">Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center</a>" for more information.</p>
<b>nodeSelector</b>	<b>ocpdc</b>	<p>The rules for scheduling a PDC job pod on a particular node using <b>nodeSelector</b> or <b>affinity</b>.</p> <p><b>pdc-domain-job:</b> Set this key to constrain the PDC pods to run only on the node you specify. For more information, see "<a href="#">nodeSelector</a>" in the Kubernetes documentation.</p> <p><b>Note:</b> To override the rules for a specific PDC pod, specify a different value for the pod's <b>nodeSelector</b> key. For example, set the <b>ocpdc.configEnv.transformation.syncPDC.nodeSelector</b> key to apply rules specifically to the SyncPDC pod.</p>

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<b>affinity</b>	<b>ocpdc</b>	The rules for scheduling a PDC job pod on a particular node using <b>nodeSelector</b> or <b>affinity</b> .  <b>pdcdomainjob</b> : Set this key if you want to constrain the PDC pod to run only on nodes that meet your criteria. For more information, see " <a href="#">Node Affinity</a> " in the Kubernetes documentation.  <b>Note</b> : To override the rules for a specific PDC pod, specify a different value for the pod's affinity key. For example, set the <b>ocpdc.configEnv.transformation.importExport.affinity</b> key to apply rules specifically to the PDC Import Export pod.
<b>domainUID</b>	<b>ocpdc.wop</b>	The name of this PDC WebLogic Server domain.
<b>includeServerOutlnPodLog</b>	<b>ocpdc.wop</b>	Whether to include the server out file in the pod's stdout log: <ul style="list-style-type: none"> <li><b>true</b>: The server out file is included. This is the default.</li> <li><b>false</b>: The server out file is excluded.</li> </ul>
<b>jtaTimeoutSeconds</b>	<b>ocpdc.wop</b>	The maximum amount of time, in seconds, an active transaction is allowed to be in the first phase of a two-phase commit transaction. The default is <b>10000</b> .  If the time expires, the transaction is automatically rolled back.
<b>jtaAbandonTimeoutSeconds</b>	<b>ocpdc.wop</b>	The maximum amount of time, in seconds, a transaction manager continues to attempt completing the second phase of a two-phase commit transaction.  The default is <b>10000</b> .
<b>stuckThreadMaxTime</b>	<b>ocpdc.wop</b>	The number of seconds a thread must be continually working before the server considers the thread to be stuck.  The default is <b>20000</b> .
<b>idlePeriodsUntilTimeout</b>	<b>ocpdc.wop</b>	The number of idle periods until the peer is considered to be unreachable. The default is <b>40</b> .
<b>dataSourceXaTxnTimeout</b>	<b>ocpdc.wop</b>	The number of seconds until the data source transaction times out. The default is <b>0</b> .  When set to <b>0</b> , the WebLogic Server Transaction Manager passes the global WebLogic Server transaction timeout in seconds in the method.
<b>pdcdAppSessionTimeout</b>	<b>ocpdc.wop</b>	The PDC application ( <b>pricingui.ear</b> ) session time out, in seconds. The default is <b>36000</b> .
<b>pdcdAppSessionInvalidIntervalTimeout</b>	<b>ocpdc.wop</b>	The PDC application ( <b>pricingui.ear</b> ) session invalid interval time out, in seconds. The default is <b>3000</b> .
<b>maxMessageSize</b>	<b>ocpdc.wop</b>	(Release 15.0.1 or later) The maximum number of bytes allowed in messages that are received over supported protocols. The default is <b>10000000</b> .
<b>isVPAEnabled</b>	<b>ocpdc.wop</b>	(Release 15.0.1 or later) Whether to enable resource limits for the PDC domain job pod.  The default file is present in <b>oc-cn-op-job-helm-chart/pdc/vpa_values.yaml</b> .  For more information, see "Using Resource Limits in PDC Domain Pods" in <i>BRM Cloud Native System Administrator's Guide</i> .

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<b>users</b>	<b>ocpdc.wop</b>	(Release 15.0.1 or later) The users to add to the PDC domain. Add the following sub-fields for each: <ul style="list-style-type: none"> <li><b>name:</b> The name of the user.</li> <li><b>description:</b> A brief description of the user.</li> <li><b>password:</b> The Base64-encoded password for the user.</li> <li><b>groups:</b> The list of groups that the user belongs to. The available PDC groups are <b>PricingAnalyst</b>, <b>PricingDesignAdmin</b>, and <b>PricingReviewer</b>.</li> </ul> For more information, see "Creating PDC Users" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>exposePorts</b>	<b>ocpdc.configEnv</b>	Exposes the SSL HTTPS port, the HTTP port, or both: <ul style="list-style-type: none"> <li><b>yes:</b> Exposes the SSL HTTPS port only.</li> <li><b>no:</b> Exposes the HTTP port only.</li> <li><b>all:</b> Exposes both the HTTP port and the SSL HTTPS port. This is the default.</li> </ul>
<b>t3ChannelPort</b>	<b>ocpdc.configEnv</b>	The port number for the t3 channel. The default is <b>30799</b> . Use this key if PDC needs to use the t3 protocol to communicate with an external system, such as Elastic Charging Engine (ECE). Set this to a Kubernetes port number from 30000 through 32767 that is not in use. This key is mandatory.
<b>t3ChannelAddresses</b>	<b>ocpdc.configEnv</b>	The IP address for the primary node or load balancer. This key is optional.
<b>t3sChannelPort</b>	<b>ocpdc.configEnv</b>	The port number for the t3s channel. The default is <b>30800</b> . Use this key if PDC needs to use the t3s protocol to communicate with an external system such as ECE. Set this to a Kubernetes port number from 30000 through 32767 that is not in use. This key is mandatory.
<b>t3sChannelAddresses</b>	<b>ocpdc.configEnv</b>	If SSL is enabled in the WebLogic domain, set this to the IP address for the primary node or load balancer. This key is optional.
<b>USER_MEMORY_ARGS</b>	<b>ocpdc.configEnv</b>	The custom memory arguments for WebLogic Admin Server.
<b>USER_JAVA_OPTIONS</b>	<b>ocpdc.configEnv</b>	The custom Java options for WebLogic Admin Server.
<b>tlsVersions</b>	<b>ocpdc.configEnv</b>	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <b>TLSv1.2, TLSv1.3</b> .

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<p><b>pdcaAppLogLevel</b></p> <p><b>pdcaAppLogFileSi ze</b></p> <p><b>pdcaAppLogFileCo unt</b></p>	<b>ocpdca.configEnv</b>	<p>The settings for rotating Pricing Server log and tracing log files:</p> <ul style="list-style-type: none"> <li><b>pdcaAppLogLevel:</b> Sets the logging level, which can be <b>SEVERE</b>, <b>WARNING</b>, <b>INFO</b>, <b>CONFIG</b>, <b>FINE</b>, <b>FINER</b>, or <b>FINEST</b>. The default is <b>WARNING</b>.</li> <li><b>pdcaAppLogFileSize:</b> Sets the maximum file size, in bytes, of the log files. After the log file meets the maximum, PDC closes the log file and creates a new log file.</li> <li><b>pdcaAppLogFileCount:</b> Specifies the maximum number of log files to retain for the application.</li> </ul> <p>See "Rotating PDC Log Files" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.</p>
<b>rcuJdbcURL</b>	<b>ocpdca.configEnv</b>	The connection string for connecting to a database where schemas needed by Oracle Fusion Middleware products will be created, especially OPSS. Use the format " <i>host:port/service</i> ".
<b>rcuPrefix</b>	<b>ocpdca.configEnv</b>	The prefix for the PDC domain RCU schema. For example, if the prefix is <b>XYZ</b> and the schema name is <b>STB</b> , the PDC domain RCU schema name will be <b>XYZ_STB</b> .
<b>rcuRecreate</b>	<b>ocpdca.configEnv</b>	Whether to re-create the PDC domain. <ul style="list-style-type: none"> <li><b>true:</b> If the PDC domain is present, the RCU drops and re-creates the domain. This is the default.</li> <li><b>false:</b> The PDC domain is left alone.</li> </ul>
<b>isCustomWLSPython</b>	<b>ocpdca.configEnv</b>	Whether to run your custom WebLogic Python files: <ul style="list-style-type: none"> <li><b>true:</b> Run your custom WebLogic Python files located in <b>oc-cn-op-job-chart/pdc/customWLSPython</b>.</li> <li><b>false:</b> Do not run a custom WebLogic Python file.</li> </ul> Set the appropriate read and write permissions for each file, and set the ownership to <b>1000</b> using <b>chown 1000:0</b> .
<b>honorOMF</b>	<b>ocpdca.configEnv</b>	Whether the RDS database honors the Oracle-Managed Files (OMF) naming format: <ul style="list-style-type: none"> <li><b>true:</b> OMF format is used.</li> <li><b>false:</b> OMF format is not used. This is the default.</li> </ul>
<b>keyStoreType</b>	<b>ocpdca.configEnv</b>	The SSL KeyStore type for the PDC domain. The default is <b>JKS</b> .
<b>keyStoreAlias</b>	<b>ocpdca.configEnv</b>	The alias name for the PDC domain SSL KeyStore. The default is <b>WeblogicPDCTestAlias</b> .
<b>keyStoreIdentityFile Name</b>	<b>ocpdca.configEnv</b>	The name of the PDC domain SSL KeyStore Identity file. The default is <b>defaultserver.jks</b> . The <b>defaultserver.jks</b> file is created during PDC deployment if it does not already exist.
<b>keyStoreTrustFile Name</b>	<b>ocpdca.configEnv</b>	The name of the PDC domain SSL TrustStore file. The default is <b>defaultclient.jks</b> . The <b>defaultclient.jks</b> file is created during PDC deployment if it does not already exist.

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
<b>isSSOEnabled</b>	<b>ocpdc.configEnv</b>	Set to <b>true</b> to configure and use SAML 2.0 SSO service. The default is <b>false</b> .
<b>samlAsserterName</b>	<b>ocpdc.configEnv</b>	The name of the SAML Asserter. It should be the same as OEM or IDCS. The default is <b>pdcsAML2IdentityAsserter</b> .
<b>ssoPublishedSiteURL</b>	<b>ocpdc.configEnv</b>	The base URL used to construct endpoint URLs, typically, the load balancer host and port at which the server is visible externally. It must be appended with <b>/saml2</b> . For example: <b>https://LoadBalancerHost:LoadBalancerPort/saml2</b> .
<b>ssoDefaultURL</b>	<b>ocpdc.configEnv</b>	The URL to which unsolicited authentication responses are sent if they do not contain an accompanying target URL.
<b>ssoLogoutURL</b>	<b>ocpdc.configEnv</b>	The URL where users are redirected after they log out from the application (OEM or IDCS log out).
<b>dbHostName</b>	<b>ocpdc.configEnv</b>	The host name of the PDC and cross-reference database.
<b>dbPort</b>	<b>ocpdc.configEnv</b>	The port for the PDC and cross-reference database.
<b>dbService</b>	<b>ocpdc.configEnv</b>	The service name for the PDC and cross-reference database.
<b>dbSysDBAUser</b>	<b>ocpdc.configEnv</b>	The SYS, System, or Sys DBA user for the PDC and cross-reference database. If this key is not configured, PDC assumes that <b>pdcschemaUserName</b> and <b>crossRefSchemaUserName</b> are already present on the database with the required permissions.
<b>dbSysDBARole</b>	<b>ocpdc.configEnv</b>	The role of the PDC and cross-reference database SYS, System, or Sys DBA user.
<b>dbSSLMode</b>	<b>ocpdc.configEnv</b>	The type of SSL connection required for connecting to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity.</li> <li><b>NO</b>: SSL authentication is not required. This is the default.</li> </ul> If set to <b>ONE_WAY</b> or <b>TWO_WAY</b> , place the database wallet in the <b>oc-cn-helm-chart/pdc/pdc_db_wallet</b> directory. Create the directory structure if it is not present and do not change the directory name.
<b>dbWalletType</b>	<b>ocpdc.configEnv</b>	The type of file specified as the TrustStore for SSL connections: <b>SSO</b> or <b>pkcs12</b> . SSO is the recommended value.
<b>crossRefSchemaPDCTableSpace</b>	<b>ocpdc.configEnv</b>	The name of the PDC tablespace for the transformation cross-reference schema. This field is case-sensitive.
<b>crossRefSchemaTempTableSpace</b>	<b>ocpdc.configEnv</b>	The name of the temporary tablespace for the transformation cross-reference schema. This field is case-sensitive.
<b>crossRefSchemaUserName</b>	<b>ocpdc.configEnv</b>	The cross-reference database schema user name.

Table 7-6 (Cont.) Pricing Design Center Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml	Description
pdcschemaPDCTableSpace	ocpdc.configEnv	The tablespace name of the PDC schema. This field is case-sensitive.
pdcschemaTempTableSpace	ocpdc.configEnv	The tablespace name of the temporary schema. This field is case-sensitive.
pdcschemaUserName	ocpdc.configEnv	The PDC database schema user name.
rcuWalletSchemaUserName	ocpdc.configEnv	The RCU wallet schema user name. The default schema user name is <b>PDRCUWALLET</b> . <b>Note:</b> The OPSS wallet file created for the RCU schema is stored in the RCU_WALLET_DETAILS table during the first run. If the wallet file is available for the given RCU prefix, it is reused in subsequent runs and the RCU schema is not re-created. If the OPSS wallet file is present in <b>oc-cn-op-job-helm-chart/pdc/opss_wallet</b> , it takes precedence.
pdAdminUser	ocpdc.configEnv	The PDC admin user name, which includes the Pricing Design Admin role.
supportECE	ocpdc.configEnv	The charging engine to use: Elastic Charging Engine ( <b>true</b> ) or the real-time and batch rating engine ( <b>false</b> ). The default is <b>true</b> .
adminPassword	ocpdc.secretValue	The password for the WebLogic domain's administrative user, which is used for accessing the WebLogic Console for administrative operations.
rcuSchemaPassword	ocpdc.secretValue	The password for the Oracle Fusion Middleware product schemas that will be created by RCU and used by OPSS.
keyStoreIdentityKeyPass	ocpdc.secretValue	The password for the PDC domain SSL identity key.
keyStoreIdentityStorePass	ocpdc.secretValue	The password for the PDC domain SSL identity store.
keyStoreTrustStorePass	ocpdc.secretValue	The password for the PDC domain SSL TrustStore.
dbPassword	ocpdc.secretValue	The Sys or System user password for the PDC and Cross Reference schema.
pdcschemaPassword	ocpdc.secretValue	The password for the PDC database schema user.
crossRefSchemaPassword	ocpdc.secretValue	The password for the transformation cross-reference database schema user.
rcuWalletSchemaPassword	ocpdc.secretValue	The password for the PDC RCU OPSS wallet schema.
dbWalletPassword	ocpdc.secretValue	The password for the database SSL wallet. This key is required if <b>dbWalletType</b> is set to <b>pkcs12</b> .
walletPassword	ocpdc.secretValue	The passwords for the PDC application wallet and PDC BRM integration wallet.
pdAdminUserPassword	ocpdc.secretValue	The password for the PDC admin user, which includes the Pricing Design Admin role.

## Setting Up SSO for PDC Cloud Native

SSO allows users to log in to applications using a single user name and password combination. You set up SSO for PDC cloud native services by using SAML 2.0.

To set up SSO for PDC:

1. Export the SAML 2.0 metadata XML file from your identity and access management (IAM) system.

For example, if you are using Oracle Access Management, you can export the file by following the instructions in "[Exporting Metadata](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

2. Rename the metadata XML file to **metadata.xml**, and then move **metadata.xml** to the **oc-cn-op-job-helm-chart/pdc/idp** directory.
3. Configure the KeyStores needed by SAML 2.0:
  - a. Generate Identity and Trust KeyStores.
  - b. Move your KeyStore files, such as **identity.p12** and **trust.p12**, to the **oc-cn-op-job-helm-chart/pdc/pdc\_keystore** directory.
4. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, set the following keys:
  - **ocpdc.configEnv.isSSOEnabled**: Set this to **true**.
  - **ocpdc.configEnv.keyStoreAlias**: Set this to the private key alias of the KeyStore.
  - **ocpdc.configEnv.keyStoreType**: Set this to the file type of the SSL Identity and Trust KeyStore, which is either **PKCS12** or **JKS**. The default is **PKCS12**.
  - **ocpdc.configEnv.keyStoreIdentityFileName**: Set this to the name of the Identity KeyStore file.
  - **ocpdc.configEnv.keyStoreTrustFileName**: Set this to the name of the Trust KeyStore file.
  - **ocpdc.configEnv.samlAsserterName**: Set this to the name of the SAML Asserter. The default is **pdcSAML2IdentityAsserter**.
  - **ocpdc.configEnv.ssoPublishedSiteURL**: Set this to the base URL used to construct endpoint URLs. This is typically the load balancer host and port where the server is visible externally. It must be appended with **/saml2**. For example: **https://LoadBalancerHost:LoadBalancerPort/saml2**.
  - **ocpdc.configEnv.ssoDefaultURL**: Set this to the URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
  - **ocpdc.secretValue.keyStoreIdentityStorePass**: Set this to the StorePass for the Identity KeyStore.
  - **ocpdc.secretValue.keyStoreIdentityKeyPass**: Set this to the KeyPass for the Identity KeyStore.
  - **ocpdc.secretValue.keyStoreTrustStorePass**: Set this to the StorePass for the Trust KeyStore.
5. Configure your load balancer's rules to send responses to the PDC WebLogic domain with **/saml2** appended to the URL path.

**Note**

Add this rule to your existing load balancer rules for routing responses to PDC (`/ pdc`), the load balancer host name, and so on.

See "[Installing an Ingress Controller](#)".

6. Deploy your PDC cloud native services by following the instructions in "[Deploying BRM Cloud Native Services](#)".
7. After PDC is deployed, retrieve the `sp-metadata-admin-server.xml` file from the `/shared/domains/domainUID` directory in your container, where `domainUID` is the name of your PDC domain specified in the `ocpdc.wop.domainUID` key.

The XML file configures the Web SSO Provider Partner. It contains the partner's KeyStore certificates, SAML assertion details, and the URLs where the SAML Identity Provider redirects to provide access to PDC.

8. Create a profile for your identity provider partner by loading the `sp-metadata-admin-server.xml` file into your IAM system.

For example, if you are using Oracle Access Management, you can load the file by following the instructions in "[Creating Remote Identity Provider Partners](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

## Configuring Pipeline Configuration Center

To configure Pipeline Configuration Center (PCC) to run in your BRM cloud native environment:

1. Override the PCC-specific keys in the `values.yaml` file for `oc-cn-op-job-helm-chart`. See "[Adding Pipeline Configuration Center Keys for oc-cn-op-job-helm-chart](#)".
2. Override the PCC-specific keys in the `values.yaml` file for `oc-cn-helm-chart`. See "[Adding Pipeline Configuration Center Keys for oc-cn-helm-chart](#)".
3. Set up volume mounts. See "[About PCC Volume Mounts](#)".
4. Create a WebLogic domain and install the PCC application. See "[Creating a WebLogic Domain and Installing the PCC Application](#)".
5. Set up SAML for SSO in PCC. See "[Setting Up SSO for PCC](#)".
6. Set up local users and groups for PCC. See "[Setting Up Local Users and Groups for PCC](#)".
7. Start and stop your WebLogic servers. See "[Starting and Stopping WebLogic Servers](#)".
8. Enable SSL in PCC. See "[Configuring SSL in PCC](#)".

## Adding Pipeline Configuration Center Keys for oc-cn-op-job-helm-chart

[Table 7-7](#) lists the keys that directly impact PCC deployment. Add these keys to your `override-values.yaml` file for `oc-cn-op-job-helm-chart`.

Table 7-7 PCC Keys for oc-cn-op-job-helm-chart

Key	Paths in values.yaml File	Description
<b>isEnabled</b>	<b>ocpcc.pcc</b>	Whether to deploy, configure, and start PCC services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using PCC.</li> <li><b>true</b>: Creates the Kubernetes resources for using PCC. This is the default.</li> </ul>
<b>imageName</b>	<b>ocpcc.pcc.deploy ment.app</b>	The name of the PCC image, such as <b>oracle/pcc</b>
<b>imageTag</b>	<b>ocpcc.pcc.deploy ment.app</b>	The tag associated with the image. This is generally the patch set number prefixed with a colon (:). For example, <b>:15.0.1.0.0</b>
<b>dbSSLMode</b>	<b>ocpcc.pcc.config Env</b>	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
<b>dbWalletType</b>	<b>ocpcc.pcc.config Env</b>	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
<b>rcuJdbcURL</b>	<b>ocpcc.pcc.config Env</b>	The connection string for connecting to the database where schemas needed by Oracle Fusion Middleware products will be created, especially OPSS.
<b>rcuDBARole</b>	<b>ocpcc.pcc.config Env</b>	The role of the database administrator user.
<b>rcuArgs</b>	<b>ocpcc.pcc.config Env</b>	The additional arguments for creating the RCU.
<b>ldapHost</b>	<b>ocpcc.pcc.config Env</b>	The host name or IP address of the LDAP Server (for example, OUD) where users and groups will be configured for access to PCC.
<b>ldapPort</b>	<b>ocpcc.pcc.config Env</b>	The port number on which the LDAP server is listening.
<b>ldapGroupBase</b>	<b>ocpcc.pcc.config Env</b>	The LDAP base DN that contains groups.
<b>ldapUserBase</b>	<b>ocpcc.pcc.config Env</b>	The LDAP base DN that contains users.
<b>keystoreAlias</b>	<b>ocpcc.pcc.config Env</b>	The private key alias of the KeyStore.
<b>keystoreType</b>	<b>ocpcc.pcc.config Env</b>	The file type of the SSL Identity and TrustStore, which is either <b>PKCS12</b> or <b>JKS</b> . The default is <b>PKCS12</b> .
<b>keystoreIdentityFile</b> <b>fileName</b>	<b>ocpcc.pcc.config Env</b>	The file name of the Identity KeyStore.
<b>keystoreTrustFile</b> <b>Name</b>	<b>ocpcc.pcc.config Env</b>	The file name of the Trust KeyStore.
<b>isSSOEnabled</b>	<b>ocpcc.pcc.config Env</b>	Whether to enable single sign-on (SSO) for PCC cloud native services through SAML 2.0: <ul style="list-style-type: none"> <li><b>true</b>: <b>SSO</b> is enabled for PCC cloud native services.</li> <li><b>false</b>: <b>SSO</b> is disabled. This is the default.</li> </ul>

Table 7-7 (Cont.) PCC Keys for oc-cn-op-job-helm-chart

Key	Paths in values.yaml File	Description
samlAsserterName	ocpcc.pcc.config Env	The name of the SAML Asserter. The default is <b>samlPCCAsserter</b> .
ssoPublishedSiteURL	ocpcc.pcc.config Env	The base URL that is used to construct endpoint URLs. This is typically the Load Balancer host and port at which the server is visible externally. It must be appended with <b>/saml2</b> . For example: <b>https://LoadBalancerHost:LoadBalancerPort/saml2</b> .
ssoDefaultURL	ocpcc.pcc.config Env	The URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
reloadVersion	ocpcc.pcc.config Env	Update this value with any value different from the current value to force a restart of the deployer.
adminPassword	ocpcc.pcc.secret Val	The password of the WebLogic domain's administrative user, which is used for accessing the WebLogic Console for administrative operations.
ldapPassword	ocpcc.pcc.secret Val	The password of the LDAP Server admin user.
rcuSysDBAPassword	ocpcc.pcc.secret Val	The password for the rcuJdbcURL database administrator.
rcuSchemaPassword	ocpcc.pcc.secret Val	The passwords for the schemas of Oracle Fusion Middleware products that will be created by RCU, which is used by OPSS.
dbWalletPassword	ocpcc.pcc.secret Val	The password for accessing the certificates from the TrustStore and KeyStore.
keystoreIdentityPassword	ocpcc.pcc.secret Val	The StorePass for the Identity KeyStore.
keystoreKeyPassword	ocpcc.pcc.secret Val	The KeyPass for the Identity KeyStore.
keystoreTrustPassword	ocpcc.pcc.secret Val	The StorePass for the Trust KeyStore.
domainUID	ocpcc.pcc.wop	The name of the domain. The default is <b>pcc-domain</b> .
adminChannelPort	ocpcc.pcc.wop	The NodePort where the admin-server's HTTP service will be accessible.
serverStartPolicy	ocpcc.pcc.wop	The WebLogic servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li><b>NEVER</b>: Does not start any server in the domain.</li> <li><b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li><b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
nodeSelector	ocpcc.pcc	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
affinity	ocpcc.pcc	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

## Adding Pipeline Configuration Center Keys for oc-cn-helm-chart

[Table 7-8](#) lists the keys that directly impact PCC deployment. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart**.

**Table 7-8 Pipeline Configuration Center Keys**

Key	Path in Values.yaml File	Description
appLogLevel	ocpcc	The logging level at which application logs must be captured in log files: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , <b>FINEST</b> , and <b>ALL</b> .
isEnabled	ocpcc.pcc	Whether to deploy, configure, and start PCC services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using PCC.</li> <li><b>true</b>: Creates the Kubernetes resources for using PCC. This is the default.</li> </ul>
imageName	ocpcc.pcc.deployment.app	The name of the PCC image, such as <b>oracle/pcc</b> .
imageTag	ocpcc.pcc.deployment.app	The tag associated with the image. This is generally the patch set number, prefixed with a colon (:). For example, <b>:15.0.1.0.0</b>
keystoreAlias	ocpcc.pcc.configEnv	The private key alias of the KeyStore.
dbSSLMode	ocpcc.pcc.configEnv	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each other's identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
dbWalletType	ocpcc.pcc.configEnv	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
tlsVersions	ocpcc.pcc.configEnv	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <b>TLSv1.2, TLSv1.3</b> .
login	ocpcc.pcc.infranet.user	The username of the service that has permission to access BRM.
serviceType	ocpcc.pcc.infranet.user	The POID type of the service that has permission to access BRM.
serviceID	ocpcc.pcc.infranet.user	The POID ID of the service that has permission to access BRM.
minSize	ocpcc.pcc.infranet.connectionpool	The minimum size of the connection pool.
maxSize	ocpcc.pcc.infranet.connectionpool	The maximum size of the connection pool.
loglevel	ocpcc.pcc.infranet	The log level for the infranet properties.

Table 7-8 (Cont.) Pipeline Configuration Center Keys

Key	Path in Values.yaml File	Description
<b>addOnProperties</b>	<b>ocpcc.pcc.infranet</b>	Empty by default, you can use this key to specify custom infranet properties.
<b>domainUID</b>	<b>ocpcc.pcc.wop</b>	The name of the domain. The default is pcc-domain.
<b>adminChannelPort</b>	<b>ocpcc.pcc.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.
<b>serverStartPolicy</b>	<b>ocpcc.pcc.wop</b>	The WebLogic servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li><b>NEVER</b>: Does not start any server in the domain.</li> <li><b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li><b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
<b>isEnabled</b>	<b>ocpcc.pcc.monitoring</b>	Whether to enable monitoring of PCC.
<b>nodeSelector</b>	<b>ocpcc.pcc</b>	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
<b>affinity</b>	<b>ocpcc.pcc</b>	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

[Table 7-9](#) lists the secret keys that directly impact PCC deployment. These keys hold sensitive data and must be handled carefully with controlled access to the file containing its values. See "[Secrets](#)" in *Kubernetes Concepts*.

Add these secret keys to your **override-values.yaml** file, and encode all of their values in Base64.

**Note**

- You can encode strings in Linux by using this command:

```
echo -n 'string' | base64
```

- You can decode strings in Linux by using this command:

```
echo 'encoded_string' | base64 --decode
```

Table 7-9 Pipeline Configuration Center Secret Keys

Key	Description
<b>ocpcc.pcc.secretVal.adminPassword</b>	The WebLogic Server administrative password encoded in Base64.
<b>ocpcc.pcc.secretVal.walletPassword</b>	The PCC wallet password encoded in Base64.
<b>ocpcc.pcc.secretVal.rcuSysDBAPassword</b>	The Database Administrator password encoded in Base64.

Table 7-9 (Cont.) Pipeline Configuration Center Secret Keys

Key	Description
<b>ocpcc.pcc.secretVal.rcuSchemaPassword</b>	The password for schemas of Oracle Fusion Middleware products that will be created by RCU, which is used by OPSS. The value must be Base64-encoded.
<b>ocpcc.pcc.secretVal.keystoreIdentityPassword</b>	The KeyPass of Identity Keystore, which is used for setting up the SSL-enabled domain. The value must be Base64-encoded.
<b>ocpcc.pcc.secretVal.keystoreKeyPassword</b>	The StorePass of Identity Keystore, which is used for setting up the SSL-enabled domain. This value must be Base64-encoded.
<b>ocpcc.pcc.secretVal.keystoreTrustPassword</b>	The StorePass of the Trust Keystore, which is used for setting up the SSL-enabled domain. This value must be Base64-encoded.
<b>ocpcc.pcc.secretVal.pccUserPassword</b>	The PCC user password encoded in Base64.

## About PCC Volume Mounts

The PCC container requires Kubernetes volume mounts for sharing the domain and application file system between the WebLogic Cluster servers. There is one volume for the domain. By default, these are created dynamically, using the provisioner defined in BRM, in the **storage-class** key in **oc-cn-op-job-helm-chart**.

To change the volume type or provider, modify the following keys in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

- **ocpcc.pcc.volume.domain.createOption** for the domain file system for PCC.

## Creating a WebLogic Domain and Installing the PCC Application

The WebLogic domain is created by a Kubernetes Deployment when **oc-cn-op-job-helm-chart** is installed. The same job also installs the PCC application and deploys the application EAR file onto the WebLogic Cluster.

The **oc-cn-op-job-helm-chart** chart also:

- Creates a Kubernetes ConfigMap and Secrets, which are used throughout the life-cycle of the WebLogic domain.
- Initializes the **PersistentVolumeClaim** for the domain and application file system as well as third-party libraries.

### Note

The **override-values.yaml** file that you use for this chart must include BRM override values.

After you install **oc-cn-op-job-helm-chart**, wait until the Kubernetes deployment has reached the **1/1 Running** status. Then, you can install or upgrade **oc-cn-helm-chart** for PCC services.

After the deployment is running, don't delete the chart. Its resources will be used for starting and stopping the servers through **oc-cn-helm-chart**.

## Setting Up SSO for PCC

SSO allows users to log in to applications using a single user name and password combination. You set up SSO for PCC cloud native services by using SAML 2.0.

To set up SSO for PCC:

1. Export the SAML 2.0 metadata XML file from your identity and access management (IAM) system.

For example, if you are using Oracle Access Management, you can export the file by following the instructions in "[Exporting Metadata](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

2. Rename the metadata XML file to **metadata.xml**, and then move **metadata.xml** to the **oc-cn-op-job-helm-chart/pcc/idp** directory.
3. Configure the KeyStores needed by SAML 2.0:
  - a. Generate identity and trust KeyStores.
  - b. Move your KeyStore files, such as **identity.p12** and **trust.p12**, to the **oc-cn-op-job-helm-chart/pcc/keystore** directory.
4. In your **override-values.yaml** file for **oc-cn-helm-chart**, set the **isSSOEnabled** key to **true**.
5. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, set the following keys:
  - **ocpcc.pcc.configEnv.isSSOEnabled**: Set this to **true**.
  - **ocpcc.pcc.configEnv.keystoreAlias**: Set this to the private key alias of the KeyStore.
  - **ocpcc.pcc.configEnv.keystoreType**: Set this to the file type of the SSL Identity and Trust store, which is either **PKCS12** or **JKS**. The default is **PKCS12**.
  - **ocpcc.pcc.configEnv.keystoreIdentityFileName**: Set this to the name of the Identity KeyStore file.
  - **ocpcc.pcc.configEnv.keystoreTrustFileName**: Set this to the name of the Trust KeyStore file.
  - **ocpcc.pcc.configEnv.samlAsserterName**: Set this to the name of the SAML Asserter. The default is **samlPCCAsserter**.
  - **ocpcc.pcc.configEnv.ssoPublishedSiteURL**: Set this to the base URL that is used to construct endpoint URLs. This is typically the load balancer host and port at which the server is visible externally. It must be appended with **/saml2**. For example: **https://LoadBalancerHost:LoadBalancerPort/saml2**.
  - **ocpcc.pcc.configEnv.ssoDefaultURL**: Set this to the URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
  - **ocpcc.pcc.secretVal.keystoreIdentityPassword**: Set this to the StorePass for the Identity KeyStore.
  - **ocpcc.pcc.secretVal.keystoreKeyPassword**: Set this to the KeyPass for the Identity KeyStore.
  - **ocpcc.pcc.secretVal.keystoreTrustPassword**: Set this to the StorePass for the Trust KeyStore.

6. Configure your load balancer's rules to send responses to the PCC WebLogic domain with `/saml2` appended to the URL path.

### Note

Add this rule to your existing load balancer rules for routing responses to PCC (`/pcc`), the load balancer host name, and so on.

See "[Installing an Ingress Controller](#)".

7. Deploy your PCC cloud native services by following the instructions in "[Deploying BRM Cloud Native Services](#)".
8. After PCC is deployed, retrieve the `sp-metadata-admin-server.xml` file from the `/shared/domains/domainUID` directory in your container, where `domainUID` is the name of your PCC domain specified in the `ocpcc.pcc.wop.domainUID` key.

The XML file configures the Web SSO Provider Partner. It contains the partner's KeyStore certificates, SAML assertion details, and the URLs where the SAML Identity Provider redirects to provide access to PCC.

9. Create a profile for your identity provider partner by loading the `sp-metadata-admin-server.xml` file into your IAM system.

For example, if you are using Oracle Access Management, you can load the file by following the instructions in "[Creating Remote Identity Provider Partners](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

## Setting Up Local Users and Groups for PCC

You have the option to customize the values for `oc-cn-op-job-helm-chart` to create users and groups locally in Oracle WebLogic Server. This would be especially useful for test environments where you might not have Identity Providers or LDAPs available. The groups for the admin user for WebLogic Server cannot be modified using this procedure.

Any passwords must be encoded using Base64. You can leave the password blank, but then the user will not be able to log in to the application directly.

To set up local users and groups for PCC, define the keys under `ocpcc.pcc.wlsUserGroups` in the `override-values.yaml` file for `oc-cn-op-job-helm-chart`.

A group has to be the Config Admin to access the PCC UI. Only a user associated with the Config Admin group has full access to the PCC user interface. For example:

```
Add users and groups to domain's DefaultAuthenticator (local)
wlsUserGroups:
  groups:
    - name: Config Admin
      description: PCC Admin
      # Each element for this takes "name", "description", "password" (base64
encoded) and list of "groups" that he is part of, like:
      # - name:
      # description:
      # password:
      # groups:
      # - "Regular CSR"
  users:
    - name: pccuser
      description: pccuser
```

```
password: QzFnMmIzdTQj
groups:
- "Config Admin"
```

## Starting and Stopping WebLogic Servers

When you install **oc-cn-op-job-helm-chart**, the default configuration sets up a WebLogic Cluster with five Managed Servers. When you install or upgrade **oc-cn-helm-chart** for the PCC service, two of the Managed Servers and one Admin Server are started.

By modifying the **override-values.yaml** file for **oc-cn-helm-chart**, you can control:

- The total number of Managed Servers and the initial server start up by using the **totalManagedServers** and **initialServerCount** keys.
- Whether the servers are started or stopped by using the **serverStartPolicy** key. To start the Admin Servers and the Managed Servers in a Cluster, set the key to **IF\_NEEDED**. To stop all servers, set the key to **NEVER**.

### Note

The keys in the **override-values.yaml** file should be the same as the ones used in **oc-cn-op-job-helm-chart** for keys that are common in both charts.

After you modify the **override-values.yaml** file, update the Helm release for the changes to take effect:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

## Configuring SSL in PCC

To access PCC over the HTTPS port, SSL must be enabled in the WebLogic domain where PCC is deployed. The BRM cloud native deployment package takes care of the configuration necessary to equip the WebLogic domain with SSL access.

To complete the configuration for SSL setup:

1. Copy PKCS12 files with valid certificates to the **oc-cn-op-job-helm-chart/pcc/keystore** directory:
  - **identity.p12**: Provides the certificate to identify the server.
  - **trust.p12**: Establishes trust for the certificate.

If your KeyStore files have different file names or file types, such as JKS, override the **keyStoreIdentityFileName**, **keyStoreTrustFileName**, and **keyStoreType** keys in the **override-values.yaml** file for **oc-cn-helm-chart**.

The **keystoreAlias** key is also mandatory along with **keyStoreIdentityFileName**, **keystoreTrustFileName** to enable SSL.

# 8

## Configuring REST Services

Learn how to integrate external applications with your Oracle Communications Billing and Revenue Management (BRM) cloud native environment by using the BRM and PDC REST services.

Topics in this document:

- [Configuring BRM REST Services Manager](#)
- [Configuring PDC REST Services Manager](#)

### Configuring BRM REST Services Manager

You use BRM REST Services Manager to integrate an external customer experience application with BRM. This allows you to manage billing and rating in BRM and then view your customers' account balances and bills in your external client. For more information, see *REST Services Manager API for Billing and Revenue Management*.

To configure BRM REST Services Manager in BRM cloud native:

1. Generate an SSL certificate. See "[Generating an SSL Certificate for BRM REST Services Manager](#)".
2. Optionally, configure the BRM REST Services Manager SDK. See "[Configuring the SDK \(Optional\)](#)".
3. Configure the HTTPS port for Oracle Unified Directory. See "[Configuring the Oracle Unified Directory HTTPS Port](#)".
4. If BRM and REST Services Manager are located in separate clusters, connect BRM REST Services Manager to BRM. See "[Connecting to a Separate BRM Cluster](#)".
5. Override the BRM REST Services Manager-specific keys in the **values.yaml** file. See "[Adding BRM REST Services Manager Keys](#)".

### Generating an SSL Certificate for BRM REST Services Manager

The following shows the steps for generating a sample SSL certificate:

1. Create a directory named **rsm\_keystore** under the **oc-cn-helm-chart/rsm** directory.
2. Generate an SSL certificate. For example:

```
openssl req -x509 -newkey rsa:4096 -keyout openSSLKey.pem -out cert.pem -days 365 -nodes
```

3. Generate a **PKCS12** KeyStore file. For example, this creates a KeyStore file named **keystore.p12**:

```
openssl pkcs12 -export -out keystore.p12 -inkey openSSLKey.pem -in cert.pem
```

4. Copy your SSL certificate file to the **oc-cn-helm-chart/rsm/rsm\_keystore** directory.

## Configuring the SDK (Optional)

To integrate the SDK with BRM REST Services Manager, generate an SDK image as follows:

1. Copy your extended SDK JAR `oc-cn-docker-files-15.0.x.0.0/oc-cn-docker-files/ocrsm/brm_rest_services_manager/SDK/libs` to the `oc-cn-docker-files-15.0.x.0.0/oc-cn-docker-files/ocrsm/brm_rest_services_manager/SDK` directory.

### Note

The SDK JAR can be used directly from `oc-cn-docker-files-15.0.x.0.0/oc-cn-docker-files/ocrsm/brm_rest_services_manager/SDK/libs` if no changes are required. If you need to make further customizations, follow the instructions in *REST Services Manager API for Billing and Revenue Management* and then copy the updated SDK JAR to the `oc-cn-docker-files-15.0.x.0.0/oc-cn-docker-files/ocrsm/brm_rest_services_manager/SDK` directory.

2. In your `override-values.yaml` file for `oc-cn-helm-chart`, set the `ocrsm.rsm.configEnv.rsmExtensionJar` key to the name of your extended SDK JAR file, such as `BRMRESTExtension.jar`.

3. Go to the `oc-cn-docker-files-15.0.x.0.0/oc-cn-docker-files/ocrsm/brm_rest_services_manager/SDK` directory.

4. Build the Podman image by running this command:

```
podman build --format docker --tag imagerepo/brm-rest-services-manager-extension:1 .
```

5. Push the SDK image to the repository by running this command:

```
podman login --username user --password password imagerepo
podman push imagerepo/brm-rest-services-manager-extension:1
```

## Configuring the Oracle Unified Directory HTTPS Port

If an HTTPS port is used for Oracle Unified Directory, do the following:

1. Create a directory named `rsm_oud_keystore` under the `oc-cn-helm-chart/rsm` directory.
2. Copy the Oracle Unified Directory certificate to the `oc-cn-helm-chart/rsm/rsm_oud_keystore` directory.

## Connecting to a Separate BRM Cluster

If BRM is located in a separate cluster from BRM REST Services Manager, do the following to connect BRM REST Services Manager to BRM:

1. Open the `configmap_env_brmrsm.yaml` file in a text editor.
2. Set `BRM_HOST_NAME` to the host name of the cluster on which BRM is located. The default value is `cm`.
3. Save and close the file.

## Adding BRM REST Services Manager Keys

[Table 8-1](#) lists the keys that directly impact BRM REST Services Manager. Add these keys to your `override-values.yaml` file with the same path hierarchy.

### Caution

Keys with the path `ocrsm.rsm.secretVal` hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 8-1 BRM REST Services Manager Keys**

Key	Path in Values.yaml File	Description
<code>isEnabled</code>	<code>ocrsm.rsm</code>	The flag to indicate if BRM REST Services Manager should be deployed with BRM cloud native.
<code>imageName</code>	<code>ocrsm.rsm.deploy ment</code>	The name of the BRM REST Services Manager image, such as <code>oracle/brm-rest-services-manager</code> .
<code>imageTag</code>	<code>ocrsm.rsm.deploy ment</code>	The tag associated with the image, such as 15.0.0.0.0.
<code>imageName</code>	<code>ocrsm.rsm.deploy ment.sdk</code>	The name of the BRM REST Services Manager SDK image, such as <code>brm-rest-services-manager-extension</code> .
<code>imageTag</code>	<code>ocrsm.rsm.deploy ment.sdk</code>	The tag associated with the BRM Services Manager SDK image, such as 1.
<code>httpPort</code>	<code>ocrsm.rsm.config Env</code>	The HTTP port in the container on which to deploy BRM REST Services Manager. <b>Note:</b> Set this to a port number from 30000 through 32767 that is not in use.
<code>httpsPort</code>	<code>ocrsm.rsm.config Env</code>	The HTTPS port in the container on which to deploy BRM REST Services Manager. <b>Note:</b> Set this to a port number from 30000 through 32767 that is not in use.
<code>tlsVersions</code>	<code>ocrsm.rsm.config Env</code>	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <code>TLSv1.2, TLSv1.3</code> .
<code>rsmCertificateFile Name</code>	<code>ocrsm.rsm.config Env</code>	The SSL certificate file name for BRM REST Services Manager.
<code>baseURL</code>	<code>ocrsm.rsm.config Env</code>	The base URL with resource details to return in the response of BRM REST Services Manager requests. <b>Note:</b> After deployment, you can update this value by editing your <code>override-values.yaml</code> file and then doing a Helm upgrade.
<code>brmLogin</code>	<code>ocrsm.rsm.config Env</code>	The user name of the service with permission to access BRM, such as <code>rsm.0.0.0.1</code> .
<code>brmServiceType</code>	<code>ocrsm.rsm.config Env</code>	The BRM service type, such as <code>/service/admin_client</code> .

Table 8-1 (Cont.) BRM REST Services Manager Keys

Key	Path in Values.yaml File	Description
<b>brmServicePoidId</b>	<b>ocrsm.rsm.config Env</b>	The BRM service POID, such as <b>1</b> .
<b>brmSSLWalletFileName</b>	<b>ocrsm.rsm.config Env</b>	The BRM SSL wallet file name.
<b>rsmExtensionJar</b>	<b>ocrsm.rsm.config Env</b>	The file name of the BRM REST Service Manager SDK JAR, such as <b>BRMRESTExtension.jar</b> .
<b>bipURL</b>	<b>ocrsm.rsm.config Env</b>	The Oracle Analytics Publisher URL.
<b>bipUserId</b>	<b>ocrsm.rsm.config Env</b>	The Oracle Analytics Publisher user ID.
<b>securityEnabled</b>	<b>ocrsm.rsm.config Env</b>	The flag to indicate if token-based authentication is enabled for BRM REST Services Manager.
<b>idcsURI</b>	<b>ocrsm.rsm.config Env.idcs</b>	The Oracle Identity Cloud Service (IDCS) URL.
<b>clientID</b>	<b>ocrsm.rsm.config Env.idcs</b>	The IDCS client ID.
<b>proxyHost</b>	<b>ocrsm.rsm.config Env.idcs</b>	The IDCS proxy host.
<b>scopeAudience</b>	<b>ocrsm.rsm.config Env.idcs</b>	The primary audience configured in IDCS.
<b>audience</b>	<b>ocrsm.rsm.config Env.idcs</b>	The secondary audience configured in IDCS. If a secondary audience is not configured, enter the primary audience.
<b>domainName</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Access Manager domain name.
<b>audience</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Access Manager OAuth server name.
<b>endpointURL</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Access Manager OAuth token endpoint URL.
<b>oudHostName</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory host name.
<b>oudRootUserDN</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory root user domain name.
<b>oudHttpPort</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory HTTP port.
<b>oudHttpsPort</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory HTTPS port.
<b>oudUserBaseDN</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory user domain name.
<b>oudGroupDN</b>	<b>ocrsm.rsm.config Env.oam</b>	The Oracle Unified Directory group domain name.
<b>logLevel</b>	<b>ocrsm.rsm.config Env</b>	The application log level: <b>SEVERE, WARNING, INFO, CONFIG, FINE, FINER, FINEST</b> .
<b>helidonSecurityLogLevel</b>	<b>ocrsm.rsm.config Env</b>	The security log level: <b>SEVERE, WARNING, INFO, CONFIG, FINE, FINER, or FINEST</b> .
<b>helidonWebServerLogLevel</b>	<b>ocrsm.rsm.config Env</b>	The server log level: <b>SEVERE, WARNING, INFO, CONFIG, FINE, FINER, or FINEST</b> .

Table 8-1 (Cont.) BRM REST Services Manager Keys

Key	Path in Values.yaml File	Description
helidonConfigLogLevel	ocrsm.rsm.configEnv	The Helidon configuration log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
helidonMicroProfileLogLevel	ocrsm.rsm.configEnv	Helidon MP log level <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
helidonCommonLogLevel	ocrsm.rsm.configEnv	The Helidon common log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
nettyServerLogLevel	ocrsm.rsm.configEnv	The embedded netty server log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
jerseyLogLevel	ocrsm.rsm.configEnv	The jersey log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
jbossWeldLogLevel	ocrsm.rsm.configEnv	The Helidon JBossWeld log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
auditLogLevel	ocrsm.rsm.configEnv	The audit log level: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , or <b>FINEST</b> .
rsmCertificatePassword	ocrsm.rsm.secretVal	The Base64-encoded certificate password for BRM REST Services Manager.
brmInfranetWalletPassword	ocrsm.rsm.secretVal	The Base64-encoded wallet password. You can use any password.  This password will be used to store the Oracle Analytics Publisher and Infranet connections in the wallet and can be used to access the same.
bipPassword	ocrsm.rsm.secretVal	The Base64-encoded Oracle Analytics Publisher password.
clientSecret	ocrsm.rsm.secretVal	The Base64-encoded IDCS client secret.
oudCertificateAliases	ocrsm.rsm.secretVal	The certificate alias is any unique name that is Base64-encoded for JDK KeyStore. This key is required for Oracle Unified Directory HTTPS protocols.
oudRootUserPassword	ocrsm.rsm.secretVal	The Base64-encoded Oracle Unified Directory root password.

You can use the following commands to encode and decode passwords in Base64 format:

- To encode strings in Linux:  
`echo -n 'password' | base64`
- To decode strings in Linux:  
`echo 'encoded_password' | base64 --decode`

## Sample override-values for IDCS Security Type

This shows sample content in the **override-values.yaml** for BRM REST Services Manager when the security type is Oracle Identity Cloud Service (IDCS):

```
ocrsm:
  rsm:
    configEnv:
      securityEnabled: true
      bipUrl: http://xxxxxxx:xxxx/xmlpserver/services/PublicReportService_v11
```

```

bipUserId: weblogic
baseURL: xxxxx.xxx.xxxxx.xxx
idcs:
  idcsURI: "https://xxxxx.xxxx.xxxxx.xxxxx.xxxx"
  clientID: xxxxxx
  scopeAudience: "https://xxxxx:xxxxx:xxxx:xxxxx:/"
  audience: "https://xxxxx:xxx:xxxxx:xxxxx:/"
rsmExtensionJar: brm-rest-extension.jar
rsmCertificateFileName: keyStore.pl2
secretVal:
  rsmCertificatePassword: xxxxxx
  brmInfranetWalletPassword: xxxxxx
  bipPassword: xxxxxx
  clientSecret: xxxxxx

```

## Sample override-values.yaml for OAM Security Type

This shows sample content in the **override-values.yaml** for BRM REST Services Manager when the security type is Oracle Access Manager:

```

ocrsm:
  rsm:
    configEnv:
      securityEnabled: true
      bipUrl: http://xxxxx:xxxxx/xmlpserver/services/PublicReportService_v11
      bipUserId: weblogic
      securityType: OAM
      baseURL: xxxxx.xxx.xxxxx.xxx
      oam:
        domainName: TMFOAuthOIDCDomain
        audience: TMFResourceServer
        endpointURL: "http://xxxx.xxxx.xxxx.xxx:14100/oauth2/rest/token/info"
        oudHostName: xxxx.xxxx.xxxx.xxx
        oudRootUserDN: cn=MyRootUser
        oudHttpPort: 9090
        oudUserBaseDN: ou=people,dc=tmf,dc=com
        oudGroupDN: ou=Groups,dc=tmf,dc=com
      rsmExtensionJar: brm-rest-extension.jar
      rsmCertificateFileName: keyStore.pl2
    secretVal:
      rsmCertificatePassword: xxxxxx
      brmInfranetWalletPassword: xxxxxx
      bipPassword: xxxxxx
      oudRootUserPassword: xxxxxx

```

## Sample BRM RSM override-values for Separate BRM Cluster

This shows sample content in the **override-values.yaml** for BRM REST Services Manager when BRM is deployed in a separate cluster:

**Note**

- Pass the values for **idcsURI**, **scopeAudience**, **audience**, and **OAM endpointURL** in quotes when the URL ends with characters such as colon colon (::).
- Ensure that you provide the hostname for connecting to the BRM cluster in the **configmap\_env\_brmrsm.yaml** file. See "[Connecting to a Separate BRM Cluster](#)".

```

ocbrm:
  brm_root_pass: xxxxx
  isSSLEnabled: true

ocrsm:
  rsm:
    configEnv:
      securityEnabled: true
      bipUrl: http://xxxxx:xxxx/xmlpserver/services/PublicReportService_v11
      bipUserId: weblogic
      baseURL: xxxxx.xxx.xxxxx.xxx
      idcs:
        idcsURI: "https://xxxxx.xxxx.xxxxx.xxxxx.xxxx"
        clientID: xxxx
        scopeAudience: "https://xxxxx:xxxxx:xxxx:xxxxx::"
        audience: "https://xxxxx:xxx:xxxxx:xxxxx::"
      rsmExtensionJar: brm-rest-extension.jar
      rsmCertificateFileName: keyStore.p12
    secretVal:
      rsmCertificatePassword: xxxxx
      brmInfranetWalletPassword: xxxxx
      bipPassword: xxxxx
      clientSecret: xxxxx

```

## Configuring PDC REST Services Manager

You use PDC REST Services Manager to integrate an enterprise product catalog, such as Oracle Digital Experience for Communications Launch Experience, with PDC. This enables you to create a variety of product offerings in your enterprise product catalog and then have all of the rating and billing performed by PDC and BRM. For more information, see "About PDC REST Services Manager" in *PDC REST Services Manager Integration Guide*.

To configure PDC REST Services Manager in BRM cloud native:

1. Override the PDC REST Services Manager-specific keys in the **values.yaml** file. See "[Adding PDC REST Services Manager Keys](#)".
2. Configure OAuth authentication:
  - a. If you are using Oracle Access Management for OAuth, create an identity domain, resource server, and OAuth client for PDC REST Services Manager in Oracle Access Management as described in "Setting Up OAuth for PDC REST Services Manager with Oracle Access Management" in *BRM Security Guide*.
  - b. Configure the keys in the **override-values.yaml** file for OAuth with either Oracle Identity Cloud Service or Oracle Access Management as described in "[Configuring OAuth Authentication in PDC REST Services Manager](#)".

3. Configure outbound communication to the enterprise product catalog. See "[Configuring Requests to the Enterprise Product Catalog](#)".
4. Enable TLS encryption in PDC REST Services Manager to secure the communications it receives from your enterprise product catalog. See "[Enabling TLS in PDC REST Services Manager](#)".
5. Enable the T3S protocol in PDC REST Services Manager to secure its communications to PDC. See "[Enabling T3S in PDC REST Services Manager](#)".
6. Map TMF620 priceType values to BRM events to ensure that PDC REST Services Manager triggers the correct charging events for your pricing components. See "[Configuring Mapping of TMF620 priceType to BRM Events](#)".

## Adding PDC REST Services Manager Keys

[Table 8-2](#) lists the keys that directly impact PDC REST Services Manager. Add these keys to your `override-values.yaml` file with the same path hierarchy.

### Caution

Keys with the path `ocpdcrrsm.secretValue` hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 8-2 PDC REST Services Manager Keys**

Key	Path in Values.yaml File	Description
<code>isEnabled</code>	<code>ocpdcrrsm.labels</code>	Whether to enable and deploy PDC REST Services Manager with BRM cloud native: <ul style="list-style-type: none"> <li>• <b>true</b>: Enables and deploys PDC REST Services Manager.</li> <li>• <b>false</b>: Does not deploy PDC REST Services Manager. This is the default.</li> </ul>
<code>imageName</code>	<code>ocpdcrrsm.deploy ment</code>	The name of the PDC REST Services Manager image, such as <code>oracle/pdcrrsm</code> .
<code>imageTag</code>	<code>ocpdcrrsm.deploy ment</code>	The tag associated with the image.
<code>rsmListenerPort</code>	<code>ocpdcrrsm.config Env</code>	The HTTPS port number assigned to listen for API requests from the enterprise product catalog.
<code>baseUrl</code>	<code>ocpdcrrsm.config Env</code>	The base URL with resource details to return in the response of PDC REST Services Manager requests. <b>Note:</b> After deployment, you can update this value by editing your <code>override-values.yaml</code> file and then doing a Helm upgrade.
<code>securityEnabled</code>	<code>ocpdcrrsm.config Env</code>	Whether to enable token-based authentication for PDC REST Services Manager.
<code>securityType</code>	<code>ocpdcrrsm.config Env</code>	Which OAuth provider to use for token-based authentication. Set this to <code>oam</code> for Oracle Access Management or <code>idcs</code> for Oracle Identity Cloud Service.

## Sample PDC REST Services Manager override-values.yaml Entries

The following shows sample content in the **override-values.yaml** for PDC REST Services Manager, when Oracle Access Management is used for OAuth authentication:

```

ocpdcrrsm:
  labels:
    name: "pdc-rsm"
    version: "15.0.0.0.0"
    isEnabled: true
  deployment:
    deadlineSeconds: 60
    revisionHistLimit: 10
    imageName: "oracle/pdcrrsm"
    imageTag: ":15.0.0.0.0"
    imagePullPolicy: IfNotPresent
  configEnv:
    name: "pdcrrsm-configmap-env"
    rsmListenerPort:
    baseUrl: xxxxx.xxx.xxxxx.xxx
    securityEnabled: true
    securityType: oam
    oam:
      domainName: PDCRSMDomain
      audience: PDCRSMResourceServer
      endpointURL: http://oam_host:oam_port/oauth2/rest/token
      introspectendpointuri: http://oam_host:oam_port/oauth2/rest/token/info
      scopeaudience: http://oam_host:oam_port/
      authorizationendpointuri: http://oam_host:oam_port/oauth2/authorize
      proxyhost: http://proxyhost:proxyport/
      frontenduri: http://oam_host:oam_port
  secretValue:
    name: "pdcrrsm-secret-env"
  service:
    name: "pdcrrsm"
    type: "NodePort"
    nodePort: 31000

```

## Configuring OAuth Authentication in PDC REST Services Manager

PDC REST Services Manager uses the OAuth 2.0 protocol to authenticate an enterprise product catalog's identity and to authorize the enterprise product catalog to access the PDC REST Services Manager API. It does this by validating an OAuth access token that is passed in the header of every HTTP/HTTPS request to the PDC REST Services Manager API.

To configure OAuth authentication in PDC REST Services Manager:

1. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart**:
  - If you are using Oracle Identity Cloud Service (IDCS) for OAuth:
    - **ocpdcrrsm.configEnv.isInboundOauthEnabled**: Set this to **true** to enable OAuth authentication.
    - **ocpdcrrsm.configEnv.inboundOauthUri**: Set this to the base URL of your Oracle Identity Cloud Service (IDCS) instance in this format:  
  
`https://idcs-TenantID.identity.oraclecloud.com`
    - **ocpdcrrsm.configEnv.inboundOauthClientId**: Set this to the client ID of your confidential application.

- **ocpdcrrsm.secretValue.inboundOauthClientSecret**: Set this to the Base64-encrypted client secret obtained from your IDCS application.
  - **ocpdcrrsm.configEnv.inboundOauthFrontendUri**: Set this to the base URL of your confidential application when run, such as `http://myapp.example.com:8080`.
  - **ocpdcrrsm.configEnv.inboundOauthAudience**: Set this to the primary audience as provisioned for the PDC REST Services Manager application in IDCS.
  - **ocpdcrrsm.configEnv.inboundOauthProxyHost**: Set this to the host name of your proxy server, if required.
  - **ocpdcrrsm.configEnv.inboundOauthPubEventScope**: Set this to the name of the scope for accessing the TMF620 Publish Event endpoint for inbound OAuth authentication, such as `pubevent`.
  - **ocpdcrrsm.configEnv.inboundOauthMetricsScope**: Set this to the name of the scope for accessing the metrics endpoint for inbound OAuth authentication, such as `metrics`.
- If you are using Oracle Access Management for OAuth:
    - **ocpdcrrsm.configEnv.oam.domainName**: Set this to the name of the OAuth identity domain created in Oracle Access Management for PDC REST Services Manager.
    - **ocpdcrrsm.configEnv.oam.audience**: Set this to the name of the OAuth resource server created in Oracle Access Management for PDC REST Services Manager.
    - **ocpdcrrsm.configEnv.oam.endpointURL**: Set this to the URL for requesting an OAuth token from Oracle Access Management.
    - **ocpdcrrsm.configEnv.oam.introspectendpointuri**: Set this to the URL for validating an OAuth token from Oracle Access Management.
    - **ocpdcrrsm.configEnv.oam.scopeaudience**: Set this to the primary audience for PDC REST Services Manager in the Oracle Access Management resource, used for error handling. This is the same as **ocpdcrrsm.configEnv.oam.frontenduri**, ending with `/`.
    - **ocpdcrrsm.configEnv.oam.authorizationendpointuri**: The URL for authorizing role-based access. PDC REST Services Manager does not support role-based access, so this will not be used.
    - **ocpdcrrsm.configEnv.oam.proxyhost**: Set this to the URL for your Oracle Access Management proxy server, if needed.
    - **ocpdcrrsm.configEnv.oam.frontenduri**: Set this to the URL for of the OAuth client created in Oracle Access Management for PDC REST Services Manager.

2. Run the **helm upgrade** command to update the Helm release:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

3. Restart the PDC REST Services Manager pods. If downtime is not a concern, both pods can be deleted and re-created by running the following command. Otherwise, delete one pod at a time, waiting for its replacement pod to become "Running" before deleting the next one.

```
kubectl -n BrmNameSpace delete pods --selector=app.kubernetes.io/name=pcrrsm
```

## Configuring Requests to the Enterprise Product Catalog

PDC REST Services Manager sends requests to the enterprise product catalog when calling the enterprise product catalog's REST API and when publishing acknowledgment notifications.

To configure PDC REST Services Manager to send requests to the enterprise product catalog:

1. Open the **override-values.yaml** file for **oc-cn-helm-chart**.
2. Edit the keys in the file based on the type of authentication required by your enterprise product catalog:
  - For OAuth 2.0 authentication, edit the keys in [Table 8-3](#).

**Table 8-3 OAuth 2.0 Keys**

Key	Path in Values.yaml file	Description
<b>tokenEndpoint</b>	<b>ocpdcrsm.configEnv.httpClients.security.oauth2</b>	The endpoint used to retrieve a token from.
<b>clientId</b>	<b>ocpdcrsm.configEnv.httpClients.security.oauth2</b>	The client ID used to authenticate the request from PDC REST Services Manager.
<b>username</b>	<b>ocpdcrsm.configEnv.httpClients.security.oauth2</b>	The user name required for accessing the enterprise product catalog.
<b>scope</b>	<b>ocpdcrsm.configEnv.httpClients.security.oauth2</b>	The scopes required by the enterprise product catalog.
<b>grantType</b>	<b>ocpdcrsm.configEnv.httpClients.security.oauth2</b>	The grant type to be used for the OAuth flow: <b>client_credentials</b> or <b>password</b> .
<b>clientsecret</b>	<b>ocpdcrsm.secretValue.htpClients.security.oauth2</b>	The encrypted client secret used to authenticate the request from PDC REST Services Manager.
<b>password</b>	<b>ocpdcrsm.secretValue.htpClients.security.oauth2</b>	The encrypted password required for accessing the enterprise product catalog.

- For basic authentication, edit the keys in [Table 8-4](#).

**Table 8-4 basicAuth Keys**

Key	Path in Values.yaml file	Description
<b>username</b>	<b>ocpdcrsm.configEnv.httpClients.security.basicAuth</b>	The user name required for accessing the enterprise product catalog.
<b>password</b>	<b>ocpdcrsm.secretValue.htpClients.security.basicAuth</b>	The password required for accessing the enterprise product catalog.

3. Run the **helm upgrade** command to update the Helm release:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

4. Restart the PDC REST Services Manager pods. If downtime is not a concern, both pods can be deleted and re-created by running the following command. Otherwise, delete one pod at a time, waiting for its replacement pod to have a "Running" status before deleting the next one.

```
kubectl --namespace BrmNameSpace delete pods --selector=app.kubernetes.io/name=pdcrsm
```

The following shows an example configuration for OAuth 2.0 authentication.

### Note

All `urlRegex` values in the file must be properly escaped with `\\`. The characters that must be escaped are: `\.[]{}()<>*+~!/? ^$|`.

```
configEnv:
  httpClients:
    - urlRegex: "http://hostname:port/mobile/custom/catalogManagement/.*"
      security:
        oauth2:
          tokenEndpoint: "https://idcs_hostname/oauth2/v1/token"
          clientId: "fcb3443f6c504ed789ba38a78341b88a"
          username: "user"
          scope: "https://hostnameurn:opc:resource:consumer::all"
          grantType: "password"
secretValue:
  httpClients:
    - urlRegex: "http://hostname:port/mobile/custom/catalogManagement/.*"
      security:
        oauth2:
          clientSecret: client_secret
          password: password
```

The following shows an example configuration for Basic authentication:

### Note

All `urlRegex` values in the file must be properly escaped with `\\`. The characters that must be escaped are: `\.[]{}()<>*+~!/? ^$|`.

```
configEnv:
  httpClients:
    - urlRegex: "http://hostname:port/mobile/custom/PublishingAPI.*"
      security:
        basicAuth:
          username: eccUser
secretValue:
  httpClients:
    - urlRegex: "http://hostname:port/mobile/custom/PublishingAPI.*"
      security:
        basicAuth:
          password: password
```

## Enabling TLS in PDC REST Services Manager

You can enable TLS encryption in PDC REST Services Manager to secure the communications it receives from your enterprise product catalog.

To enable TLS in PDC REST Services Manager:

1. Generate a self-signed SSL certificate:

- a. Create a directory for storing your SSL certificate that is accessible by the BRM Helm chart, such as **oc-cn-helm-chart/rsm\_cert**.
  - b. Generate an SSL certificate. For example, this creates a certificate file named **cert.pem**:
 

```
openssl req -x509 -newkey rsa:4096 -keyout openSSLKey.pem -out cert.pem -days 365 -nodes
```
  - c. Generate a PKCS12 KeyStore file. For example, this creates a KeyStore file named **keystore.p12**:
 

```
openssl pkcs12 -export -out keyStore.p12 -inkey openSSLKey.pem -in cert.pem
```
2. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart**:
    - **security.tlsVersions**: Set this to the list of TLS versions supported by PDC REST Services Manager, such as **TLSv1.2,TLSv1.3**. This is the default global value for PDC REST Services Manager.
    - **ocpdcrrsm.configEnv.isTlsEnabled**: Set this to **true** to enable TLS encryption for PDC REST Services Manager.
    - **ocpdcrrsm.configEnv.tlsVersions**: Set this to the list of supported TLS versions, such as **TLSv1.2,TLSv1.3**.

**Note**

Set this key if you want to override the value set it **security.tlsVersions** for communication with PDC or an enterprise product catalog.

- **ocpdcrrsm.configEnv.tlsCertificatePath**: Set this to the path of the TLS certificate bundle relative to this Helm chart, such as **rsm\_cert/keyStore.p12**.
  - **ocpdcrrsm.secretValue.tlsCertificatePassphrase**: Set this to the Base64-encoded passphrase for the TLS certificate.
3. Run the **helm upgrade** command to update the Helm release:
 

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```
  4. To apply the changes, re-create any previously existing PDC REST Services Manager pods:

```
kubectl --namespace BrmNameSpace delete pods --selector=app.kubernetes.io/name=pdcrrsm
```

After you enable TLS, connect to PDC REST Services Manager services using HTTPS only.

## Enabling T3S in PDC REST Services Manager

Enable the T3S protocol in PDC REST Services Manager to secure its communications to PDC.

To enable T3S in PDC REST Services Manager:

1. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart**:
  - **ocpdcrrsm.configEnv.useT3s**: Set this to **true**.
  - **ocpdcrrsm.configEnv.jksTrustStorePath**: Set this to the path of the JKS TrustStore for the PDC T3S connection, such as **rsm/truststore.jks**.
2. Run the **helm upgrade** command to update the Helm release:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

3. To apply the changes, re-create any previously existing PDC REST Services Manager pods:

```
kubectl --namespace BrmNameSpace delete pods --selector=app.kubernetes.io/name=pdcrsm
```

## Configuring Mapping of TMF620 priceType to BRM Events

If you are using PDC REST Services Manager, you must configure the mappings of BRM event names to the values your enterprise product catalog sends in the **priceType** property of the **ProductOfferingPrice** element of the TMF620 payload.

The mappings are configured in **configmap\_pdcrsm\_appeventCfg.yaml**. You can add mappings as needed for your deployment or use the default mappings provided at installation.

To add or edit mappings:

1. Open the **configmap\_pdcrsm\_appeventCfg.yaml** file.
2. Edit the existing mappings, or use them as templates to add new ones. Use the following format:

```
pricetype : "eventname"
```

where:

- *pricetype* is the value sent in the **priceType** property of the **ProductOfferingPrice** element of the TMF620 payload.
- *eventname* is the name of the BRM event the price type should be mapped to.

For example, the default mappings for one-time fees and usage events are:

```
ONE_TIME : "EventBillingProductFeePurchase"
ONE_TIME_PRICE_PLAN : "EventBillingProductFeePurchase"
USAGE : "EventSession"
USAGE_PRICE_PLAN : "EventSession"
```

3. Run the **helm upgrade** command to update the Helm release:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

4. Restart the PDC REST Services Manager pods. If downtime is not a concern, both pods can be deleted and re-created by running the following command. Otherwise, delete one pod at a time, waiting for its replacement pod to become "Running" before deleting the next one.

```
kubectl --namespace BrmNameSpace delete pods --selector=app.kubernetes.io/name=pdcrsm
```

# 9

## Configuring the Billing Care, Billing Care REST API, and Business Operations Center Services

Learn how to configure Billing Care, Billing Care REST API, and Business Operations Center to run in your Oracle Communications Billing and Revenue Management (BRM) cloud native environment.

Topics in this document:


- [About Configuring Business Operations Center, Billing Care, and Billing Care REST API](#)
- [Configuring Business Operations Center](#)
- [Configuring Billing Care](#)
- [Configuring the Billing Care REST API](#)

### About Configuring Business Operations Center, Billing Care, and Billing Care REST API

Business Operations Center, Billing Care, and the Billing Care REST API all of them share a similar image stack.

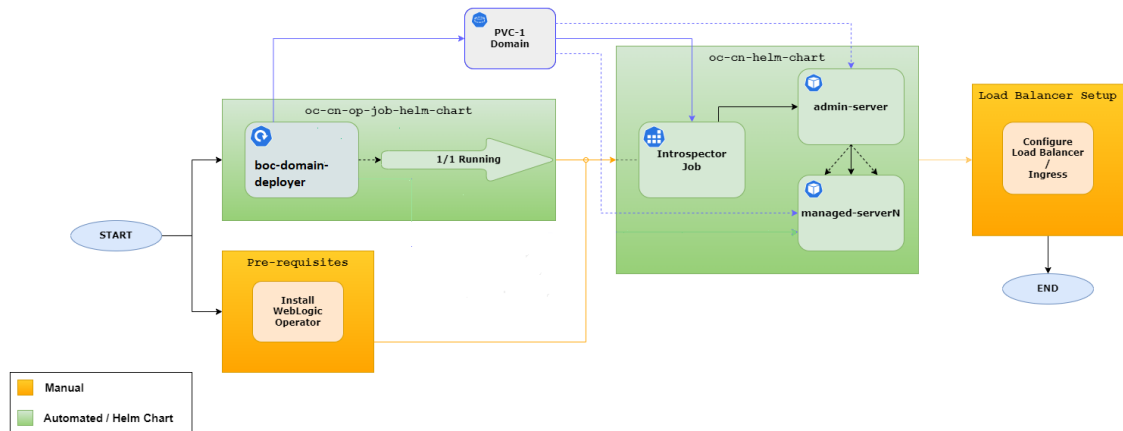
[Figure 9-1](#) shows the process for deploying Billing Care using WebLogic Operator. The same process is used for the Billing Care REST API. The only difference is the name of the deployer: **bcws-domain-deployer**.

**Figure 9-1 Billing Care Deployment Flow**



[Figure 9-2](#) shows the process for deploying Business Operations Center using WebLogic Operator. It is similar to the Billing Care process.

Figure 9-2 Business Operations Center Deployment Flow



### Note

It is important to wait until the *component-domain-deployer* process is in the **1/1 Running** status before running **oc-cn-helm-chart**.

You deploy these services by using the following Helm charts:

- **oc-cn-op-job-helm-chart:** This chart creates and configures the WebLogic domain, deploys the application, deploys and links the SDK (for Billing Care and Billing Care REST API), and loads the authorization policies.
- **oc-cn-helm-chart:** This chart starts the rolling restart of the WebLogic servers and the application update.
- **WebLogic Operator chart:** This chart manages the application domain, controlling the service availability when managed server pods are scaled up or down.

## Configuring Business Operations Center

Business Operations Center is a Web-based client application that you use to run business operations such as billing, invoicing, and payment collections. For more information, see "Using Business Operations Center" in *BRM System Administrator's Guide*.

To configure Business Operations Center to run in your BRM cloud native environment:

1. Override the Business Operations Center-specific keys in the **values.yaml** file for **oc-cn-op-job-helm-chart**. See "[Adding Business Operations Center Keys for oc-cn-op-job-helm-chart](#)".
2. Override the Business Operations Center-specific keys in the **values.yaml** file for **oc-cn-helm-chart**. See "[Adding Business Operations Center Keys for oc-cn-helm-chart](#)".
3. Set up volume mounts. See "[About Business Operations Center Volume Mounts](#)".
4. Create a WebLogic domain and install the Business Operations Center application. See "[Creating a WebLogic Domain and Installing the Business Operations Center Application](#)".
5. Set up SAML for SSO in Business Operations Center. See "[Setting Up SSO for Business Operations Center](#)".

6. Start and stop your WebLogic servers. See "[Starting and Stopping WebLogic Servers](#)".

**Note**

To set up Business Operations Center, ensure that you successfully complete the installation of **oc-cn-op-job-helm-chart** before you install or upgrade **oc-cn-helm-chart**.

## Adding Business Operations Center Keys for oc-cn-op-job-helm-chart

[Table 9-1](#) lists the keys that directly impact Business Operations Center. Add these keys to your **override-values.yaml** file for **oc-cn-op-job-helm-chart** with the same path hierarchy.

For a complete set of keys to personalize Business Operations Center deployment, see the keys with the path **ocboc.boc** in the **oc-cn-op-job-helm-chart/values.yaml** file.

**Caution**

Keys with the path **ocboc.boc.secretVal** hold sensitive data. Handle them carefully with controlled access to the file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 9-1** Keys for oc-cn-op-job-helm-chart

Key	Path in Values.yaml file	Description
<b>isEnabled</b>	<b>ocboc.boc</b>	Whether to deploy, configure, and start Business Operation Center services. <ul style="list-style-type: none"> <li><b>false</b>: Kubernetes resources meant for the Business Operation Center application will not be created.</li> <li><b>true</b>: Creates the necessary Kubernetes resources for using Business Operation Center. This is the default.</li> </ul>
<b>imageTag</b>	<b>ocboc.boc.deployment.app</b>	The tag associated with the image. This is generally the release number prefixed with a colon (:). For example: <b>:15.0.0.0.0</b> .
<b>rcuDBARole</b>	<b>ocboc.boc.configEnv</b>	The role of the database administrator user.
<b>rcuArgs</b>	<b>ocboc.boc.configEnv</b>	The additional arguments for creating the RCU.
<b>dbURL</b>	<b>ocboc.boc.configEnv</b>	Used to create the WebLogic data source for connecting to the Business Operations Center schema.  This is also the connection string for the database where schemas needed by Oracle Fusion Middleware products are created, especially OPSS.  Use one of these formats: <ul style="list-style-type: none"> <li><i>DatabaseHost:DatabasePort/ServiceName</i></li> <li><i>DatabaseHost:DatabasePort:ServiceID</i></li> </ul>

Table 9-1 (Cont.) Keys for oc-cn-op-job-helm-chart

Key	Path in Values.yaml file	Description
dbSSLMode	ocboc.boc.configEnv	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>Yes-Two Way</b>: Two-way SSL authentication is required.</li> <li><b>Yes-One Way</b>: One-way SSL authentication is required. This is the default.</li> <li><b>No</b>: SSL authentication is not required.</li> </ul>
dbWalletType	ocboc.boc.configEnv	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
dbWalletPassword	ocboc.boc.secretVal	The password for accessing the certificates from the TrustStore and KeyStore.
ldapHost	ocboc.boc.configEnv	The host name or IP address of the LDAP Server (for example, OUD) where users and groups are configured for access to Business Operations Center.
ldapPort	ocboc.boc.configEnv	The port number on which the LDAP server is listening.
ldapGroupBase	ocboc.boc.configEnv	The LDAP base DN that contains groups.
ldapUserBase	ocboc.boc.configEnv	The LDAP base DN that contains users.
bocSchemaUserName	ocboc.boc.configEnv	The Business Operations Center database schema user name.
bocSchemaBocTables pace	ocboc.boc.configEnv	The default tablespace for the Business Operations Center database administrator.
bocSchemaTempTable space	ocboc.boc.configEnv	The temp tablespace for the Business Operations Center database administrator.
billingCareUrl	ocboc.boc.configEnv	The URL of the Billing Care instance that is used with your BRM Server. Leave this blank if Billing Care isn't installed in your environment.
isSSOEnabled	ocboc.boc.configEnv	Whether to enable single sign-on (SSO) for Business Operations Center cloud native services using SAML 2.0: <ul style="list-style-type: none"> <li><b>true</b>: SSO is enabled for Business Operations Center cloud native services.</li> <li><b>false</b>: SSO is disabled. This is the default.</li> </ul>
keystoreAlias	ocboc.boc.configEnv	The private key alias of the KeyStore.
keystoreType	ocboc.boc.configEnv	The file type of the SSL Identity and Trust store, which is either <b>PKCS12</b> or <b>JKS</b> . The default is <b>PKCS12</b> .
keystoreIdentityFileName	ocboc.boc.configEnv	The file name of the Identity KeyStore.
keystoreTrustFileName	ocboc.boc.configEnv	The file name of the Trust KeyStore.
tlsVersions	ocboc.boc.configEnv	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <b>TLSv1.2, TLSv1.3</b> .

Table 9-1 (Cont.) Keys for oc-cn-op-job-helm-chart

Key	Path in Values.yaml file	Description
samlAsserterName	ocboc.boc.configEnv	The name of the SAML Asserter. The default is <b>samlBOCAsserter</b> .
ssoPublishedSiteURL	ocboc.boc.configEnv	The base URL that is used to construct endpoint URLs. This is typically the Load Balancer host and port at which the server is visible externally. It must be appended with <b>/saml2</b> . For example: <b>https://LoadBalancerHost:LoadBalancerPort/saml2</b> .
ssoDefaultURL	ocboc.boc.configEnv	The URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
reloadVersion	ocboc.boc.configEnv	Update this value with any value different from the current value to force a restart of the deployer.
adminPassword	ocboc.boc.secretVal	The Base64-encoded password for the WebLogic domain's administrative user. This is used for accessing the WebLogic Server Administration Console for administrative operations.
ldapPassword	ocboc.boc.secretVal	The Base64-encoded password of the LDAP Server admin user.
rcuSysDBAPassword	ocboc.boc.secretVal	The Base64-encoded database administrator's password.
bocSchemaPassword	ocboc.boc.secretVal	The Base64-encoded Business Operations Center database schema password.
rcuSchemaPassword	ocboc.boc.secretVal	The Base64-encoded password for schemas of Oracle Fusion Middleware products that will be created by RCU, which is used by OPSS.
keystoreIdentityPassword	ocboc.boc.secretVal	The StorePass for the Identity KeyStore.
keystoreKeyPassword	ocboc.boc.secretVal	The KeyPass for the Identity KeyStore.
keystoreTrustPassword	ocboc.boc.secretVal	The StorePass for the Trust KeyStore.
domainUID	ocboc.boc.wop	The name of the domain. The default is <b>boc-domain</b> .
adminChannelPort	ocboc.boc.wop	The NodePort where the admin-server's HTTP service will be accessible.
serverStartPolicy	ocboc.boc.wop	The WebLogic servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li><b>NEVER</b>: Does not start any server in the domain.</li> <li><b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li><b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
nodeSelector	ocboc.boc	The rules for scheduling WebLogic Server pods on particular nodes using simple selectors using Node Selector rules.
affinity	ocboc.boc	The rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors using affinity rules.

## Adding Business Operations Center Keys for oc-cn-helm-chart

[Table 9-2](#) lists the keys that directly impact Business Operations Center. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart** with the same path hierarchy.

For a complete set of keys to personalize Business Operations Center deployment, see the keys with the path **ocboc.boc** in the **oc-cn-helm-chart/values.yaml** file.

### Caution

Keys with the path **ocboc.boc.secretVal** hold sensitive data. Handle them carefully with controlled access to the file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 9-2** Keys for oc-cn-helm-chart

Key	Path in Values.yaml file	Description
<b>isEnabled</b>	<b>ocboc.boc</b>	Whether to deploy, configure, and start Business Operation Center services. <ul style="list-style-type: none"> <li><b>false</b>: Kubernetes resources meant for the Business Operation Center application are not created.</li> <li><b>true</b>: Creates the necessary Kubernetes resources for using Business Operation Center. This is the default.</li> </ul>
<b>imageTag</b>	<b>ocboc.boc.deployment.app</b>	The tag associated with the image. This is generally the release number prefixed with a colon (:). For example, <b>:15.0.0.0.0</b> .
<b>login</b>	<b>ocboc.boc.infranet.user</b>	The user name of the service with permission to access BRM, such as <b>boc_client.0.0.0.1</b> .
<b>serviceType</b>	<b>ocboc.boc.infranet.user</b>	The POID type of the service that has permission to access BRM.
<b>serviceID</b>	<b>ocboc.boc.infranet.user</b>	The POID ID of the service that has permission to access BRM.
<b>minSize</b>	<b>ocboc.boc.infranet.connectionpool</b>	The minimum size of the connection pool. The default is <b>25</b> .
<b>maxSize</b>	<b>ocboc.boc.infranet.connectionpool</b>	The maximum size of the connection pool. The default is <b>50</b> .
<b>loglevel</b>	<b>ocboc.boc.infranet</b>	The log level for the <b>Infranet.properties</b> file.
<b>addOnProperties</b>	<b>ocboc.boc.infranet</b>	This field is empty by default. You can use this key to specify custom <b>Infranet.properties</b> values.
<b>domainUID</b>	<b>ocboc.boc.wop</b>	The name of the domain. The default is <b>boc-domain</b> .
<b>adminChannelPort</b>	<b>ocboc.boc.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.

Table 9-2 (Cont.) Keys for oc-cn-helm-chart

Key	Path in Values.yaml file	Description
<b>serverStartPolicy</b>	<b>ocboc.boc.wop</b>	The WebLogic Servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li>• <b>NEVER</b>: Does not start any server in the domain.</li> <li>• <b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers are started).</li> <li>• <b>IF_NEEDED</b>: Starts the administration server and cluster servers up to the replica count.</li> </ul>
<b>isEnabled</b>	<b>ocboc.boc.monitoring</b>	Whether to enable monitoring of Business Operations Center. See "Monitoring and Autoscaling Business Operations Center Cloud Native" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>nodeSelector</b>	<b>ocboc.boc</b>	The rules for scheduling WebLogic Server pods on particular nodes using simple selectors using Node Selector rules.
<b>affinity</b>	<b>ocboc.boc</b>	The rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors using affinity rules.

## Updating Infranet.properties for Business Operations Center

The **Infranet.properties** file entries are located in the **values.yaml** file. This makes it easier to update them.

Following is a sample configuration block (located in the **ocboc.boc** path in **oc-cn-helm-chart**) for the **Infranet.properties** entries:

```
infranet:
  user:
    login: 'boc_client.0.0.0.1'
    serviceType: '/service/admin_client'
    serviceId: 2
  connectionpool:
    minSize: 25
    maxSize: 50
  logLevel: 3
  addOnProperties: ""
```

If you have custom properties, they should be defined here using the **addOnProperties** key. For example:

```
addOnProperties: |-
  infranet.connectionpool.timeout=90000
  infranet.pcp.debug.flags=0x3FFF
  infranet.pcp.debug.enabled=true
```

To update these properties, update the values in **oc-cn-helm-chart** and change the value of **ocboc.boc.wop.restartVersion** in **oc-cn-helm-chart** to any new value. This will force a pod restart and the new values will be used.

## Adding Custom Configuration to Deployment Workflow for Business Operations Center

You can provide additional configuration to be applied at particular checkpoints in the Business Operations Center deployment workflow. These checkpoints are:

- **ext\_deployer\_pre\_exit**: Called after the standard configuration in **deployer.sh** in **oc-cn-op-job-helm-chart**
- **ext\_init\_app\_pre\_exit**: Called after the standard configuration in the **init-app** **initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**
- **ext\_init\_config\_pre\_exit**: Called after the standard configuration in the **init-config** **initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**
- **ext\_init\_upgrade\_pre\_exit**: Called after the standard configuration in the **upgrade** container

Create a ConfigMap with your configuration scripts, including a shell script named **run\_hooks.sh** that calls your other scripts. For example:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: ext-scripts
data:
  run_hooks.sh: |+
    #!/bin/bash
    echo "executing extension for: $@"
    CURRENT_CHECKPOINT=$1
    if [ "$CURRENT_CHECKPOINT" == "ext_deployer_pre_exit" ] ; then
      sh my_deployer_extension.sh
    fi
  my_deployer_extension.sh: |+
    #!/bin/bash
    echo "executing my_deployer_extension"
  ...
```

Specify the name of your ConfigMap in the **ocboc.boc.extensions.scriptsConfigName** key in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

## About Business Operations Center Volume Mounts

The Business Operations Center container requires Kubernetes volume mounts for sharing the domain and application file system between the WebLogic Cluster servers. Business Operations Center requires a volume for the domain. By default, this is created dynamically, using the provisioner defined in BRM, in the **storage-class** key in **oc-cn-op-job-helm-chart**.

To change the volume type or provider, modify the **ocboc.boc.volume.domain.createOption** key in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

## Creating a WebLogic Domain and Installing the Business Operations Center Application

The WebLogic domain is created by a Kubernetes Deployment when **oc-cn-op-job-helm-chart** is installed. The same job also installs the Business Operations Center application and deploys the application WAR file onto the WebLogic Cluster.

The **oc-cn-op-job-helm-chart** chart also:

- Creates a Kubernetes ConfigMap and Secrets, which are used throughout the life-cycle of the WebLogic domain.
- Initializes the **PersistentVolumeClaim** for the domain and application file system as well as third-party libraries.

**Note**

The **override-values.yaml** file that you use for this chart must include BRM override values.

After you install **oc-cn-op-job-helm-chart**, wait until the Kubernetes deployment has reached the **1/1 Running** status. Then, you can install or upgrade **oc-cn-helm-chart** for Business Operations Center services.

After the deployment is running, don't delete the chart. Its resources will be used for starting and stopping the servers through **oc-cn-helm-chart**.

## Setting Up SSO for Business Operations Center

SSO allows users to log in to applications using a single user name and password combination. You set up SSO for Business Operations Center cloud native services by using SAML 2.0.

To set up SSO for Business Operations Center:

1. Export the SAML 2.0 metadata XML file from your identity and access management (IAM) system.

For example, if you are using Oracle Access Management, you can export the file by following the instructions in "[Exporting Metadata](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

2. Rename the metadata XML file to **metadata.xml**, and then move **metadata.xml** to the **oc-cn-op-job-helm-chart/boc/idp** directory.
3. Configure the KeyStores needed by SAML:
  - a. Generate identity and trust KeyStores.
  - b. Move your KeyStore files, such as **identity.p12** and **trust.p12**, under the **oc-cn-op-job-helm-chart/boc/keystore** directory.
4. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, set the following keys:
  - **ocboc.boc.configEnv.isSSOEnabled**: Set this to **true**.
  - **ocboc.boc.configEnv.keystoreAlias**: Set this to the private key alias of the KeyStore.
  - **ocboc.boc.configEnv.keystoreType**: Set this to the file type of the SSL Identity and Trust store, which is either **PKCS12** or **JKS**. The default is **PKCS12**.
  - **ocboc.boc.configEnv.keystoreIdentityFileName**: Set this to the name of the Identity KeyStore file.
  - **ocboc.boc.configEnv.keystoreTrustFileName**: Set this to the name of the Trust KeyStore file.

- **ocboc.boc.configEnv.samlAsserterName:** Set this to the name of the SAML Asserter. The default is **samlBOCAsserter**.
  - **ocboc.boc.configEnv.ssoPublishedSiteURL:** Set this to the base URL that is used to construct endpoint URLs. This is typically the load balancer host and port at which the server is visible externally. It must be appended with **/saml2**. For example: **https://LoadBalancerHost:LoadBalancerPort/saml2**.
  - **ocboc.boc.configEnv.ssoDefaultURL:** Set this to the URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
  - **ocboc.boc.secretVal.keystoreIdentityPassword:** Set this to the StorePass for the Identity KeyStore.
  - **ocboc.boc.secretVal.keystoreKeyPassword:** Set this to the KeyPass for the Identity KeyStore.
  - **ocboc.boc.secretVal.keystoreTrustPassword:** Set this to the StorePass for the Trust KeyStore.
5. Configure your load balancer's rules to send responses to the Business Operations Center WebLogic domain with **/saml2** appended to the URL path.

#### Note

Add this rule to your existing load balancer rules for routing responses to Business Operations Center (**opsdashboard**), the host name, and so on.

See "[Installing an Ingress Controller](#)".

6. Deploy your Business Operations Center cloud native services by following the instructions in "[Deploying BRM Cloud Native Services](#)".
7. After Business Operations Center is deployed, retrieve the **sp-metadata-admin-server.xml** file from the **/shared/domains/domainUID** directory in your container, where **domainUID** is the name of your Business Operations Center domain specified in the **ocboc.boc.wop.domainUID** key.

The XML file configures the Web SSO Provider Partner. It contains the partner's KeyStore certificates, SAML assertion details, and the URLs where the SAML Identity Provider redirects to provide access to Business Operations Center.

8. Create a profile for your identity provider partner by loading the **sp-metadata-admin-server.xml** file into your IAM system.

For example, if you are using Oracle Access Management, you can load the file by following the instructions in "[Creating Remote Identity Provider Partners](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

## Setting Up Local Users and Groups for Business Operations Center

You have the option to customize the values for **oc-cn-op-job-helm-chart** to create users and groups locally in Oracle WebLogic Server. This would be especially useful for test environments where you might not have Identity Providers or LDAPs available. The groups for the admin user for WebLogic Server cannot be modified using this procedure.

Any passwords must be encoded using Base64. You can leave the password blank, but then the user will not be able to log in to the application directly.

To set up local users and groups for Billing Care, define the keys under **ocboc.boc.wlsUserGroups** in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**. For example:

```
ocboc:
  boc:
    wlsUserGroups:
      groups:
        - name: "GroupA"
          description: "GroupA Description"
        - name: "GroupB"
          description: "GroupB Description"
      users:
        - name: csr1
          description: "csr1 description"
          password: "Base64_password"
          groups:
            - "GroupA"
            - "GroupB"
        - name: csr2
          description: "csr2 description"
          password: "Base64_password"
          groups:
            - "GroupB"
```

## Starting and Stopping WebLogic Servers

When you install **oc-cn-op-job-helm-chart**, the default configuration sets up a WebLogic Cluster with five Managed Servers. When you install or upgrade **oc-cn-helm-chart** for the Business Operations Center service, two of the managed servers and one Admin Server are started.

By modifying the **override-values.yaml** file for **oc-cn-helm-chart**, you can control:

- The total number of Managed Servers and the initial server start up by using the **totalManagedServers** and **initialServerCount** keys.
- Whether the servers are started or stopped by using the **serverStartPolicy** key. To start the Admin Servers and the Managed Servers in a Cluster, set the key to **IF\_NEEDED**. To stop all servers, set the key to **NEVER**.

### Note

The keys in the **override-values.yaml** file should be the same as the ones used in **oc-cn-op-job-helm-chart** for keys that are common in both charts.

Before installing or upgrading **oc-cn-helm-chart** for Business Operations Center, ensure that the **brm\_apps** values are configured correctly. If there is a change in any **brm\_apps** values, use **serverStartPolicy** to restart and have the changes take effect.

After you modify the **override-values.yaml** file, update the Helm release for the changes to take effect:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

## Adding Authorization for a Custom Job Resource

To grant user access to a custom job resource in Business Operations Center:

1. Add the following entry in **oc-cn-op-job-helm-chart/templates/configmap\_boc\_domain\_properties.yaml**:

```
bocauth-config.properties: |
                                JOB_BOC_ADMIN_HAS_ACCESS_TO_RESOURCE =
category_customName_resource
                                ACTIONS_GRANTED_FOR_category_customName_resource =
View,Create,Modify,Delete,Timeline,History
```

where *customName* is the name of the custom category that has been created or is yet to be created, displayed in Business Operations Center.

## Configuring Billing Care

Billing Care is a Web-based client application that CSRs use to manage billing, payments, and accounts receivable for your customers. For more information about using Billing Care, see *Billing Care Online Help*.

To configure Billing Care to run in your BRM cloud native environment:

1. Override the Billing Care-specific keys from the **values.yaml** file for **oc-cn-op-job-helm-chart**. See "[Adding Billing Care Keys for oc-cn-op-job-helm-chart](#)".
2. Override the Billing Care-specific keys from the **values.yaml** file for **oc-cn-helm-chart**. See "[Adding Billing Care Keys for oc-cn-helm-chart](#)".
3. Set up volume mounts for Billing Care. See "[About Billing Care Volume Mounts](#)".
4. Create a WebLogic domain and install Billing Care. See "[Creating a WebLogic Domain and Installing the Billing Care Application](#)".
5. Set up SAML for SSO in Billing Care. See "[Setting Up SSO for Billing Care](#)".
6. Start and stop your WebLogic servers. See "[Starting and Stopping WebLogic Servers](#)".

### Note

To set up Billing Care, ensure that you successfully complete the installation of **oc-cn-op-job-helm-chart** before you install or upgrade **oc-cn-helm-chart**.

## Adding Billing Care Keys for oc-cn-op-job-helm-chart

[Table 9-3](#) lists a few important keys that directly impact Billing Care. Add these keys to your `override-values.yaml` file for `oc-cn-op-job-helm-chart` with the same path hierarchy.

For the complete set of keys to personalize your Billing Care deployment, see the keys with the path `ocbc.bc` in the `oc-cn-op-job-helm-chart/values.yaml` file.

### Caution

Keys with the path `ocbc.bc.secretVal` hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 9-3** Keys for `oc-cn-op-job-helm-chart`

Key	Path in values.yaml File	Description
<code>isEnabled</code>	<code>ocbc.bc</code>	Whether to deploy, configure, and start Billing Care services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using Billing Care.</li> <li><b>true</b>: Creates the Kubernetes resources for using Billing Care. This is the default.</li> </ul>
<code>imageName</code>	<code>ocbc.bc.deployment.app</code>	The name of the Billing Care image, such as <b>oracle/billingcare</b> .
<code>imageTag</code>	<code>ocbc.bc.deployment.app</code>	The tag associated with the image. This is generally the release number prefixed with a colon (:). For example, <b>:15.0.x.0.0</b> .
<code>dbSSLMode</code>	<code>ocbc.bc.configEnv</code>	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
<code>dbWalletType</code>	<code>ocbc.bc.configEnv</code>	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
<code>rcuJdbcURL</code>	<code>ocbc.bc.configEnv</code>	The connection string for connecting to the database where schemas needed by Oracle Fusion Middleware products will be created, especially OPSS.
<code>rcuDBARole</code>	<code>ocbc.bc.configEnv</code>	The role of the database administrator user.
<code>rcuArgs</code>	<code>ocbc.bc.configEnv</code>	The additional arguments for creating the RCU.
<code>ldapHost</code>	<code>ocbc.bc.configEnv</code>	The host name or IP address of the LDAP Server (for example, OUD) where users and groups will be configured for access to Billing Care.

Table 9-3 (Cont.) Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml File	Description
<b>ldapPort</b>	<b>ocbc.bc.configEnv</b>	The port number on which the LDAP server is listening.
<b>ldapGroupBase</b>	<b>ocbc.bc.configEnv</b>	The LDAP base DN that contains groups.
<b>ldapUserBase</b>	<b>ocbc.bc.configEnv</b>	The LDAP base DN that contains users.
<b>keystoreAlias</b>	<b>ocbc.bc.configEnv</b>	The private key alias of the KeyStore.
<b>keystoreType</b>	<b>ocbc.bc.configEnv</b>	The file type of the SSL Identity and Trust store, which is either <b>PKCS12</b> or <b>JKS</b> . The default is <b>PKCS12</b> .
<b>keystoreIdentityFileName</b>	<b>ocbc.bc.configEnv</b>	The file name of the Identity KeyStore.
<b>keystoreTrustFileName</b>	<b>ocbc.bc.configEnv</b>	The file name of the Trust KeyStore.
<b>tlsVersions</b>	<b>ocbc.bc.configEnv</b>	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <b>TLSv1.2, TLSv1.3</b> .
<b>isSSOEnabled</b>	<b>ocbc.bc.configEnv</b>	Whether to enable single sign-on (SSO) for Billing Care cloud native services through SAML 2.0: <ul style="list-style-type: none"> <li><b>true</b>: SSO is enabled for Billing Care cloud native services.</li> <li><b>false</b>: SSO is disabled. This is the default.</li> </ul>
<b>samlAsserterName</b>	<b>ocbc.bc.configEnv</b>	The name of the SAML Asserter. The default is <b>samlBCAsserter</b> .
<b>ssoPublishedSiteURL</b>	<b>ocbc.bc.configEnv</b>	The base URL that is used to construct endpoint URLs. This is typically the Load Balancer host and port at which the server is visible externally. It must be appended with <b>/saml2</b> . For example: <b>https://LoadBalancerHost:LoadBalancerPort/saml2</b> .
<b>ssoDefaultURL</b>	<b>ocbc.bc.configEnv</b>	The URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
<b>reloadVersion</b>	<b>ocbc.bc.configEnv</b>	Update this value with any value different from the current value to force a restart of the deployer.
<b>adminPassword</b>	<b>ocbc.bc.secretVal</b>	The password of the WebLogic domain's administrative user, which is used for accessing the WebLogic Console for administrative operations.
<b>ldapPassword</b>	<b>ocbc.bc.secretVal</b>	The password of the LDAP Server admin user.
<b>rcuSysDBAPassword</b>	<b>ocbc.bc.secretVal</b>	The password for the rcuJdbcURL database administrator.
<b>rcuSchemaPassword</b>	<b>ocbc.bc.secretVal</b>	The passwords for the schemas of Oracle Fusion Middleware products that will be created by RCU, which is used by OPSS.
<b>dbWalletPassword</b>	<b>ocbc.bc.secretVal</b>	The password for accessing the certificates from the TrustStore and KeyStore.
<b>keystoreIdentityPassword</b>	<b>ocbc.bc.secretVal</b>	The StorePass for the Identity KeyStore.

Table 9-3 (Cont.) Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml File	Description
<b>keystoreKeyPassword</b>	<b>ocbc.bc.secretVal</b>	The KeyPass for the Identity KeyStore.
<b>keystoreTrustPassword</b>	<b>ocbc.bc.secretVal</b>	The StorePass for the Trust KeyStore.
<b>domainUID</b>	<b>ocbc.bc.wop</b>	The name of the domain. The default is <b>billingcare-domain</b> .
<b>adminChannelPort</b>	<b>ocbc.bc.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.
<b>serverStartPolicy</b>	<b>ocbc.bc.wop</b>	The WebLogic servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li>• <b>NEVER</b>: Does not start any server in the domain.</li> <li>• <b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li>• <b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
<b>nodeSelector</b>	<b>ocbc.bc</b>	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
<b>affinity</b>	<b>ocbc.bc</b>	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

## Adding Billing Care Keys for oc-cn-helm-chart

[Table 9-4](#) lists a few important keys that directly impact Billing Care. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart** with the same path hierarchy.

For the complete set of keys to personalize your Billing Care deployment, see the keys with the path **ocbc.bc** in the **oc-cn-helm-chart/values.yaml** file.

### Caution

Keys with the path **ocbc.bc.secretVal** hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

Table 9-4 Keys for oc-cn-helm-chart

Key	Path in values.yaml File	Description
<b>appLogLevel</b>	<b>ocbc</b>	The logging level at which application logs must be captured in log files: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , <b>FINEST</b> , and <b>ALL</b> .

Table 9-4 (Cont.) Keys for oc-cn-helm-chart

Key	Path in values.yaml File	Description
<b>isEnabled</b>	<b>ocbc.bc</b>	Whether to deploy, configure, and start Billing Care services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using Billing Care.</li> <li><b>true</b>: Creates the Kubernetes resources for using Billing Care. This is the default.</li> </ul>
<b>imageName</b>	<b>ocbc.bc.deployment.app</b>	The name of the Billing Care image, such as <b>oracle/billingcare</b> .
<b>imageTag</b>	<b>ocbc.bc.deployment.app</b>	The tag associated with the image. This is generally the release number, prefixed with a colon (:). For example, <b>:15.0.0.0.0</b> .
<b>keystoreAlias</b>	<b>ocbc.bc.configEnv</b>	The private key alias of the KeyStore.
<b>dbSSLMode</b>	<b>ocbc.bc.configEnv</b>	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
<b>dbWalletType</b>	<b>ocbc.bc.configEnv</b>	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
<b>login</b>	<b>ocbc.bc.infranet.user</b>	The user name of the service that has permission to access BRM, such as <b>bc_client.0.0.0.1</b> .
<b>serviceType</b>	<b>ocbc.bc.infranet.user</b>	The POID type of the service that has permission to access BRM.
<b>serviceID</b>	<b>ocbc.bc.infranet.user</b>	The POID ID of the service that has permission to access BRM.
<b>minSize</b>	<b>ocbc.bc.infranet.connectionpool</b>	The minimum size of the connection pool. The default is <b>25</b> .
<b>maxSize</b>	<b>ocbc.bc.infranet.connectionpool</b>	The maximum size of the connection pool. The default is <b>50</b> .
<b>loglevel</b>	<b>ocbc.bc.infranet</b>	The log level for the <b>Infranet.properties</b> file.
<b>addOnProperties</b>	<b>ocbc.bc.infranet</b>	This field is empty by default. You can use this key to specify custom <b>Infranet.properties</b> values.
<b>domainUID</b>	<b>ocbc.bc.wop</b>	The name of the domain. The default is <b>billingcare-domain</b> .
<b>adminChannelPort</b>	<b>ocbc.bc.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.

Table 9-4 (Cont.) Keys for oc-cn-helm-chart

Key	Path in values.yaml File	Description
<b>serverStartPolicy</b>	<b>ocbc.bc.wop</b>	The WebLogic Servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li><b>NEVER</b>: Does not start any server in the domain.</li> <li><b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers are started).</li> <li><b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
<b>isEnabled</b>	<b>ocbc.bc.monitoring</b>	Whether to enable monitoring of Billing Care. See "Monitoring and Autoscaling Billing Care Cloud Native" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>nodeSelector</b>	<b>ocbc.bc</b>	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
<b>affinity</b>	<b>ocbc.bc</b>	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

## Updating Infranet.properties for Billing Care

The **Infranet.properties** file entries are located in the **values.yaml** file. This makes it easier to update them.

Following is a sample configuration block (located in the **ocbc.bc** path in **oc-cn-helm-chart**) for the **Infranet.properties** entries:

```
infranet:
  user:
    login: 'boc_client.0.0.0.1'
    serviceType: '/service/admin_client'
    serviceId: 2
  connectionpool:
    minSize: 25
    maxSize: 50
  logLevel: 3
  ssoLogoutURL:
  addOnProperties: ""
```

If you have custom field classes, they should be provided through the SDK **.war** file and defined here using the **addOnProperties** key. For example:

```
addOnProperties:|-
  infranet.custom.field.package=com.portal.custom
  infranet.custom.field.100011=PIN_FLD_ABC
```

To update these properties, update the values in **override-values.yaml** file for **oc-cn-helm-chart**. If this is an upgrade, also update the **ocbc.bc.wop.restartVersion** key in the same file. This will force a pod restart and the new values will be used.

## Adding Custom Configuration to Deployment Workflow for Billing Care

You can provide additional configuration to be applied at particular checkpoints in the Billing Care deployment workflow. These checkpoints are:

- **ext\_deployer\_pre\_exit**: Called after the standard configuration in **deployer.sh** in **oc-cn-op-job-helm-chart**
- **ext\_init\_app\_pre\_exit**: Called after the standard configuration in the **init-app initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**
- **ext\_init\_config\_pre\_exit**: Called after the standard configuration in the **init-config initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**

Create a ConfigMap with your configuration scripts, including a shell script named **run\_hooks.sh** that calls your other scripts. For example:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: ext-scripts
data:
  run_hooks.sh: |+
    #!/bin/bash
    echo "executing extension for: $@"
    CURRENT_CHECKPOINT=$1
    if [ "$CURRENT_CHECKPOINT" == "ext_deployer_pre_exit" ] ; then
      sh my_deployer_extension.sh
    fi
  my_deployer_extension.sh: |+
    #!/bin/bash
    echo "executing my_deployer_extension"
  ...
```

Specify the name of your ConfigMap in the **ocbc.bc.extensions.scriptsConfigName** key in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

Since Billing Care is a web application that is deployed on WebLogic Server, refer to the WebLogic Server documentation for information about overriding timeouts, cookie attributes, and so on. See "[web.xml Deployment Descriptor Elements](#)" and "[weblogic.xml Deployment Descriptor Elements](#)" in *Developing Web Applications, Servlets, and JSPs for Oracle WebLogic Server* for more information about these configurations. You can find files to help you with this configuration in the **oc-cn-op-job-helm-chart/templates** directory.

## About Billing Care Volume Mounts

The Billing Care container requires Kubernetes volume mounts for sharing the domain and application file system between the WebLogic Cluster servers. There is one volume for the domain and one for batch payments. By default, these are created dynamically, using the provisioner defined in BRM, in the **storage-class** key in **oc-cn-op-job-helm-chart**.

To change the volume type or provider, modify the following keys in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

- **ocbc.bc.volume.domain.createOption** for the domain file system for Billing Care.
- **ocbc.bc.volume.batchPayment.createOption** for the batch payments file system.

## Creating a WebLogic Domain and Installing the Billing Care Application

The WebLogic domain is created by a Kubernetes Deployment when **oc-cn-op-job-helm-chart** is installed. The same job also installs the Billing Care application and deploys the application WAR file onto the WebLogic Cluster.

The **oc-cn-op-job-helm-chart** chart also:

- Creates a Kubernetes ConfigMap and Secrets, which are used throughout the life-cycle of the WebLogic domain.
- Initializes the **PersistentVolumeClaim** for the domain and application file system as well as third-party libraries.

### Note

The **override-values.yaml** file that you use for this chart must include BRM override values.

After you install **oc-cn-op-job-helm-chart**, wait until the Kubernetes deployment has reached the **1/1 Running** status. Then, you can install or upgrade **oc-cn-helm-chart** for Billing Care services.

After the deployment is running, don't delete the chart. Its resources will be used for starting and stopping the servers through **oc-cn-helm-chart**.

## Setting Up SSO for Billing Care

SSO allows users to log in to applications using a single user name and password combination. You set up SSO for Billing Care cloud native services by using SAML 2.0.

To set up SSO for Billing Care:

1. Export the SAML 2.0 metadata XML file from your identity and access management (IAM) system.

For example, if you are using Oracle Access Management, you can export the file by following the instructions in "[Exporting Metadata](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

2. Rename the metadata XML file to **metadata.xml**, and then move **metadata.xml** to the **oc-cn-op-job-helm-chart/billingcare/idp** directory.
3. Configure the KeyStores needed by SAML 2.0:
  - a. Generate identity and trust KeyStores.
  - b. Move your KeyStore files, such as **identity.p12** and **trust.p12**, to the **oc-cn-op-job-helm-chart/billingcare/keystore** directory.
4. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, set the following keys:
  - **ocbc.bc.configEnv.isSSOEnabled**: Set this to **true**.
  - **ocbc.bc.configEnv.keystoreAlias**: Set this to the private key alias of the KeyStore.
  - **ocbc.bc.configEnv.keystoreType**: Set this to the file type of the SSL Identity and Trust store, which is either **PKCS12** or **JKS**. The default is **PKCS12**.

- **ocbc.bc.configEnv.keystoreIdentityFileName**: Set this to the name of the Identity KeyStore file.
  - **ocbc.bc.configEnv.keystoreTrustFileName**: Set this to the name of the Trust KeyStore file.
  - **ocbc.bc.configEnv.samlAsserterName**: Set this to the name of the SAML Asserter. The default is **samlBCAsserter**.
  - **ocbc.bc.configEnv.ssoPublishedSiteURL**: Set this to the base URL that is used to construct endpoint URLs. This is typically the load balancer host and port at which the server is visible externally. It must be appended with **/saml2**. For example: **https://LoadBalancerHost:LoadBalancerPort/saml2**.
  - **ocbc.bc.configEnv.ssoDefaultURL**: Set this to the URL where unsolicited authentication responses are sent if they do not contain an accompanying target URL.
  - **ocbc.bc.secretVal.keystoreIdentityPassword**: Set this to the StorePass for the Identity KeyStore.
  - **ocbc.bc.secretVal.keystoreKeyPassword**: Set this to the KeyPass for the Identity KeyStore.
  - **ocbc.bc.secretVal.keystoreTrustPassword**: Set this to the StorePass for the Trust KeyStore.
5. Configure your load balancer's rules to send responses to the Billing Care WebLogic domain with **/saml2** appended to the URL path.

#### Note

Add this rule to your existing load balancer rules for routing responses to Billing Care (**/bc**), the load balancer host name, and so on.

See "[Installing an Ingress Controller](#)".

6. Deploy your Billing Care cloud native services by following the instructions in "[Deploying BRM Cloud Native Services](#)".
7. After Billing Care is deployed, retrieve the **sp-metadata-admin-server.xml** file from the **/shared/domains/domainUID** directory in your container, where *domainUID* is the name of your Billing Care domain specified in the **ocbc.bc.wop.domainUID** key.

The XML file configures the Web SSO Provider Partner. It contains the partner's KeyStore certificates, SAML assertion details, and the URLs where the SAML Identity Provider redirects to provide access to Billing Care.

8. Create a profile for your identity provider partner by loading the **sp-metadata-admin-server.xml** file into your IAM system.

For example, if you are using Oracle Access Management, you can load the file by following the instructions in "[Creating Remote Identity Provider Partners](#)" in *Oracle Fusion Middleware Administering Oracle Access Management*.

## Setting Up Local Users and Groups for Billing Care

You have the option to customize the values for **oc-cn-op-job-helm-chart** to create users and groups locally in Oracle WebLogic Server. This would be especially useful for test environments where you might not have Identity Providers or LDAPs available. The groups for the admin user for WebLogic Server cannot be modified using this procedure.

Any passwords must be encoded using Base64. You can leave the password blank, but then the user will not be able to log in to the application directly.

To set up local users and groups for Billing Care, define the keys under **ocbc.bc.wlsUserGroups** in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**. For example:

```
ocbc:
  bc:
    wlsUserGroups:
      groups:
        - name: "GroupA"
          description: "GroupA Description"
        - name: "GroupB"
          description: "GroupB Description"
      users:
        - name: csr1
          description: "csr1 description"
          password: "Base64_password"
          groups:
            - "GroupA"
            - "GroupB"
        - name: csr2
          description: "csr2 description"
          password: "Base64_password"
          groups:
            - "GroupB"
```

## Starting and Stopping WebLogic Servers

When you install **oc-cn-op-job-helm-chart**, the default configuration sets up a WebLogic Cluster with five Managed Servers. When you install or upgrade **oc-cn-helm-chart** for the Billing Care service, two of the Managed Servers and one Admin Server are started.

By modifying the **override-values.yaml** file for **oc-cn-helm-chart**, you can control:

- The total number of Managed Servers and the initial server start up by using the **totalManagedServers** and **initialServerCount** keys.
- Whether the servers are started or stopped by using the **serverStartPolicy** key. To start the Admin Servers and the Managed Servers in a Cluster, set the key to **IF\_NEEDED**. To stop all servers, set the key to **NEVER**.

### Note

The keys in the **override-values.yaml** file should be the same as the ones used in **oc-cn-op-job-helm-chart** for keys that are common in both charts.

After you modify the **override-values.yaml** file, update the Helm release for the changes to take effect:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.

- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

## Configuring the Billing Care REST API

You use the Billing Care REST API to integrate an external customer management application with BRM. This allows you to manage billing and rating in BRM and then manage your customers' accounts and bills in your external application. For more information, see *REST API Reference for Billing Care*.

To configure the Billing Care REST API to work with BRM cloud native:

1. Override the Billing Care REST API-specific keys from the **values.yaml** file for **oc-cn-op-job-helm-chart**. See "[Adding Billing Care REST API Keys for oc-cn-op-job-helm-chart](#)".
2. Override the Billing Care REST API-specific keys from the **values.yaml** file for **oc-cn-helm-chart**. See "[Adding Billing Care REST API Keys for oc-cn-helm-chart](#)".
3. Set up volume mounts for the Billing Care REST API. See "[About Billing Care REST API Volume Mounts](#)".
4. Create a WebLogic domain and install the Billing Care REST API. See "[Creating a WebLogic Domain and Installing the Billing Care REST API](#)".
5. Start and stop your WebLogic servers. See "[Starting and Stopping WebLogic Servers](#)".

### Note

To set up the Billing Care REST API, ensure that you successfully complete the installation of **oc-cn-op-job-helm-chart** before you install or upgrade **oc-cn-helm-chart**.

## Adding Billing Care REST API Keys for oc-cn-op-job-helm-chart

[Table 9-5](#) lists a few important keys that directly impact the Billing Care REST API. Add these keys to your **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

For the complete set of keys to personalize your Billing Care REST API deployment, see the keys with the path **ocbc.bcws** in the **oc-cn-op-job-helm-chart/values.yaml** file.

### Caution

Keys with the path **ocbc.bcws.secretVal** hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

Table 9-5 Billing Care REST API Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml File	Description
<b>isEnabled</b>	<b>ocbc.bcws</b>	Whether to deploy, configure, and start Billing Care REST API services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using the Billing Care REST API.</li> <li><b>true</b>: Creates the Kubernetes resources for using the Billing Care REST API. This is the default.</li> </ul>
<b>imageName</b>	<b>ocbc.bcws.deploy.app</b>	The name of the Billing Care REST API image, such as <b>oracle/bcws</b> .
<b>imageTag</b>	<b>ocbc.bcws.deploy.app</b>	The tag associated with the image. This is generally the release number. Prefix the value with a colon (:). For example, <b>:15.0.x.0.0</b> .
<b>dbSSLMode</b>	<b>ocbc.bcws.config Env</b>	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
<b>dbWalletType</b>	<b>ocbc.bcws.config Env</b>	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
<b>dbWalletPassword</b>	<b>ocbc.bcws.config Env</b>	The password for accessing the certificates from the TrustStore and KeyStore.
<b>rcuJdbcURL</b>	<b>ocbc.bcws.config Env</b>	The connection string for connecting to the database where schemas needed by Oracle Fusion Middleware products will be created, especially OPSS.
<b>rcuDBARole</b>	<b>ocbc.bcws.config Env</b>	The role of the database administrator user.
<b>rcuArgs</b>	<b>ocbc.bcws.config Env</b>	The additional arguments for creating the RCU.
<b>ldapHost</b>	<b>ocbc.bcws.config Env</b>	The host name or IP address of the LDAP Server (for example, OUD) where users and groups will be configured for access to the Billing Care REST API.
<b>ldapPort</b>	<b>ocbc.bcws.config Env</b>	The port number on which the LDAP server is listening.
<b>ldapGroupBase</b>	<b>ocbc.bcws.config Env</b>	The LDAP base DN that contains groups.
<b>ldapUserBase</b>	<b>ocbc.bcws.config Env</b>	The LDAP base DN that contains users.
<b>keystoreAlias</b>	<b>ocbc.bcws.config Env</b>	The private key alias of the KeyStore.
<b>keystoreType</b>	<b>ocbc.bcws.config Env</b>	The file type of SSL Identity and Trust store, either <b>PKCS12</b> or <b>JKS</b> .
<b>keystoreIdentityFileName</b>	<b>ocbc.bcws.config Env</b>	The file name of the Identity KeyStore.
<b>keystoreTrustFileName</b>	<b>ocbc.bcws.config Env</b>	The file name of the Trust KeyStore.

Table 9-5 (Cont.) Billing Care REST API Keys for oc-cn-op-job-helm-chart

Key	Path in values.yaml File	Description
<b>tlsVersions</b>	<b>ocbc.bcws.config Env</b>	(Release 15.0.1 or later) The list of TLS versions to support for connection with the WebLogic domain. List the version numbers in order, from lowest to highest, separated by a comma. For example: <b>TLSv1.2, TLSv1.3</b> .
<b>reloadVersion</b>	<b>ocbc.bcws.config Env</b>	Update this value with any value different from the current value to force a restart of the deployer.
<b>adminPassword</b>	<b>ocbc.bcws.secret Val</b>	The password of the WebLogic domain's administrative user, which is used for accessing the WebLogic Console for administrative operations.
<b>ldapPassword</b>	<b>ocbc.bcws.secret Val</b>	The password of the LDAP Server admin user.
<b>rcuSysDBAPassw ord</b>	<b>ocbc.bcws.secret Val</b>	The password for the rcuJdbcURL database administrator.
<b>rcuSchemaPassw ord</b>	<b>ocbc.bcws.secret Val</b>	The passwords for the schemas of Oracle Fusion Middleware products that will be created by RCU, which is used by OPSS.
<b>dbWalletPasswor d</b>	<b>ocbc.bcws.secret Val</b>	The password for accessing the certificates from the TrustStore and KeyStore.
<b>keystoreIdentityP assword</b>	<b>ocbc.bcws.secret Val</b>	The storepass of the Identity KeyStore.
<b>keystoreKeyPass word</b>	<b>ocbc.bcws.secret Val</b>	The KeyPass of the Identity KeyStore.
<b>keystoreTrustPas sword</b>	<b>ocbc.bcws.secret Val</b>	The storepass of Trust KeyStore.
<b>domainUID</b>	<b>ocbc.bcws.wop</b>	The name of the domain. The default is <b>bcws-domain</b> .
<b>adminChannelPor t</b>	<b>ocbc.bcws.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.
<b>serverStartPolicy</b>	<b>ocbc.bcws.wop</b>	The WebLogic servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li>• <b>NEVER</b>: Does not start any server in the domain.</li> <li>• <b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li>• <b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
<b>nodeSelector</b>	<b>ocbc.bcws</b>	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
<b>affinity</b>	<b>ocbc.bcws</b>	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

## Adding Billing Care REST API Keys for oc-cn-helm-chart

[Table 9-6](#) lists a few important keys that directly impact the Billing Care REST API. Add these keys to your **override-values.yaml** file for **oc-cn-helm-chart**.

For the complete set of keys to personalize your Billing Care REST API deployment, see the keys with the path **ocbc.bcws** in the **oc-cn-helm-chart/values.yaml** file.

**⚠ Caution**

Keys with the path `ocbc.bcws.secretVal` hold sensitive data. Handle them carefully with controlled access to the override file containing their values. Encode all of these values in Base64 format. See "[Secrets](#)" in *Kubernetes Concepts*.

**Table 9-6 Billing Care REST API Keys for oc-cn-helm-chart**

Key	Path in values.yaml File	Description
<code>appLogLevel</code>	<code>ocbc</code>	The logging level at which application logs must be captured in log files: <b>SEVERE</b> , <b>WARNING</b> , <b>INFO</b> , <b>CONFIG</b> , <b>FINE</b> , <b>FINER</b> , <b>FINEST</b> , and <b>ALL</b> .
<code>isEnabled</code>	<code>ocbc.bcws</code>	Whether to deploy, configure, and start Billing Care REST API services: <ul style="list-style-type: none"> <li><b>false</b>: Does not create the Kubernetes resources for using the Billing Care REST API.</li> <li><b>true</b>: Creates the Kubernetes resources for using the Billing Care REST API. This is the default.</li> </ul>
<code>imageName</code>	<code>ocbc.bcws.deploy ment.app</code>	The name of the Billing Care REST API image, such as <b>oracle/bcws</b> .
<code>imageTag</code>	<code>ocbc.bcws.deploy ment.app</code>	The tag associated with the image. This is generally the release number. Prefix the value with a colon (:). For example, <b>:15.0.0.0.0</b> .
<code>keystoreAlias</code>	<code>ocbc.bcws.config Env</code>	The private key alias of the KeyStore.
<code>dbSSLMode</code>	<code>ocbc.bcws.config Env</code>	The type of connection required to connect to the database: <ul style="list-style-type: none"> <li><b>TWO_WAY</b>: Two-way SSL authentication is required. In this case, both the client and server must authenticate each others identity.</li> <li><b>ONE_WAY</b>: One-way SSL authentication is required. In this case, the client must authenticate the server's identity. This is the default.</li> <li><b>NO</b>: SSL authentication is not required.</li> </ul>
<code>dbWalletType</code>	<code>ocbc.bcws.config Env</code>	The type of TrustStore and KeyStore file that is used for the SSL connection: <b>SSO</b> or <b>PKCS12</b> .
<code>login</code>	<code>ocbc.bcws.infran et.user</code>	The user name of the service that has permission to access BRM.
<code>serviceType</code>	<code>ocbc.bcws.infran et.user</code>	The POID type of the service that has permission to access BRM.
<code>serviceID</code>	<code>ocbc.bcws.infran et.user</code>	The POID ID of the service that has permission to access BRM.
<code>minSize</code>	<code>ocbc.bcws.infran et.connectionpool</code>	The minimum size of the connection pool. The default is <b>25</b> .
<code>maxSize</code>	<code>ocbc.bcws.infran et.connectionpool</code>	The maximum size of the connection pool. The default is <b>50</b> .
<code>loglevel</code>	<code>ocbc.bcws.infran et</code>	The log level for the <b>Infranet.properties</b> file.
<code>addOnProperties</code>	<code>ocbc.bcws.infran et</code>	This key is empty by default. You can use this key to specify custom <b>Infranet.properties</b> values.

Table 9-6 (Cont.) Billing Care REST API Keys for oc-cn-helm-chart

Key	Path in values.yaml File	Description
<b>domainUID</b>	<b>ocbc.bcws.wop</b>	The name of the domain. The default is <b>bcws-domain</b> .
<b>adminChannelPort</b>	<b>ocbc.bcws.wop</b>	The NodePort where the admin-server's HTTP service will be accessible.
<b>serverStartPolicy</b>	<b>ocbc.bcws.wop</b>	The WebLogic Servers that the Operator starts when it discovers the domain: <ul style="list-style-type: none"> <li><b>NEVER</b>: Does not start any server in the domain.</li> <li><b>ADMIN_ONLY</b>: Starts only the administration server (no managed servers will be started).</li> <li><b>IF_NEEDED</b>: Starts the administration server and clustered servers up to the replica count.</li> </ul>
<b>isEnabled</b>	<b>ocbc.bc.monitoring</b>	Whether to enable monitoring of Billing Care REST API. See "Monitoring and Autoscaling Billing Care Cloud Native" in <i>BRM Cloud Native System Administrator's Guide</i> .
<b>nodeSelector</b>	<b>ocbc.bcws</b>	The node selector rules for scheduling WebLogic Server pods on particular nodes using simple selectors.
<b>affinity</b>	<b>ocbc.bcws</b>	The affinity rules for scheduling WebLogic Server pods on particular nodes using more powerful selectors.

## Updating Infranet Properties for the Billing Care REST API

The **Infranet.properties** file entries are located in the **values.yaml** file. This makes it easier to update them.

Following is a sample configuration block (located in the **ocbc.bcws** path in **oc-cn-helm-chart**) for the **Infranet.properties** entries:

```
infranet:
  user:
    login: 'root.0.0.0.1'
    serviceType: '/service/admin_client'
    serviceId: 2
  connectionpool:
    minSize: 25
    maxSize: 50
  logLevel: 3
  addOnProperties: ""
```

If you have custom field classes, they should be provided through the SDK **.war** file and defined here using the **addOnProperties** key. For example:

```
addOnProperties:|-
  infranet.custom.field.package=com.portal.custom
  infranet.custom.field.100011=PIN_FLD_ABC
```

To update any of these properties after an install or upgrade, update the values in **override-values.yaml** file for **oc-cn-helm-chart**. If this is an upgrade, also update the **ocbc.bcws.wop.restartVersion** key in the same file. This will force a pod restart and the new values will be used.

## Adding Custom Configuration to Deployment Workflow for Billing Care REST API

You can provide additional configuration to be applied at particular checkpoints in the Billing Care REST API deployment workflow. These checkpoints are:

- **ext\_deployer\_pre\_exit**: Called after the standard configuration in **deployer.sh** in **oc-cn-op-job-helm-chart**
- **ext\_init\_app\_pre\_exit**: Called after the standard configuration in the **init-app initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**
- **ext\_init\_config\_pre\_exit**: Called after the standard configuration in the **init-config initContainer** container in both **oc-cn-op-job-helm-chart** and **oc-cn-helm-chart**

Create a ConfigMap with your configuration scripts, including a shell script named **run\_hooks.sh** that calls your other scripts. For example:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: ext-scripts
data:
  run_hooks.sh: |+
    #!/bin/bash
    echo "executing extension for: $@"
    CURRENT_CHECKPOINT=$1
    if [ "$CURRENT_CHECKPOINT" == "ext_deployer_pre_exit" ] ; then
      sh my_deployer_extension.sh
    fi
  my_deployer_extension.sh: |+
    #!/bin/bash
    echo "executing my_deployer_extension"
  ...
```

Specify the name of your ConfigMap in the **ocbc.bcws.extensions.scriptsConfigName** key in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

## About Billing Care REST API Volume Mounts

The Billing Care REST API container requires Kubernetes volume mounts for sharing the domain and application file system between the WebLogic Cluster servers. There is one volume for the domain and one for batch payments. By default, these are created dynamically, using the provisioner defined in BRM, in the **storage-class** key in **oc-cn-op-job-helm-chart**.

### Note

The selected location must be accessible on all worker nodes across which WebLogic Servers will be distributed based on defined nodeSelector or affinity rules.

To change the volume type or provider, modify the following keys in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.

- **ocbc.bcws.volume.domain.createOption** for the domain file system for Billing Care.
- **ocbc.bcws.volume.batchPayment.createOption** for the batch payments file system.

## Creating a WebLogic Domain and Installing the Billing Care REST API

The WebLogic domain is created by a Kubernetes Deployment when **oc-cn-op-job-helm-chart** is installed. The same job also installs the Billing Care REST API and deploys the application WAR file onto the WebLogic Cluster.

The **oc-cn-op-job-helm-chart** chart also:

- Creates a Kubernetes ConfigMap and Secrets, which are used throughout the life-cycle of the WebLogic domain.
- Initializes the **PersistentVolumeClaim** for the domain and application file system as well as third-party libraries.

### Note

The **override-values.yaml** file that you use for this chart must include BRM override values.

After you install **oc-cn-op-job-helm-chart**, wait until the Kubernetes deployment has reached the **1/1 Running** status. Then, you can install or upgrade **oc-cn-helm-chart** for Billing Care REST API services.

After the deployment is running, don't delete the chart. Its resources will be used for starting and stopping the servers through **oc-cn-helm-chart**.

## Setting Up Local Users and Groups for Billing Care REST API

You have the option to customize the values for **oc-cn-op-job-helm-chart** to create users and groups locally in Oracle WebLogic Server. This would be especially useful for test environments where you might not have Identity Providers or LDAPs available. The groups for the admin user for WebLogic Server cannot be modified using this procedure.

Any passwords must be encoded using Base64. You can leave the password blank, but then the user will not be able to log in to the application directly.

To set up local users and groups for Billing Care, define the keys under **ocbc.bcws.wlsUserGroups** in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**. For example:

```
ocbc:
  bcws:
    wlsUserGroups:
      groups:
        - name: "GroupA"
          description: "GroupA Description"
        - name: "GroupB"
          description: "GroupB Description"
      users:
        - name: csr1
          description: "csr1 description"
          password: "Base64_password"
          groups:
            - "GroupA"
            - "GroupB"
        - name: csr2
```

```
description: "csr2 description"
password: "Base64_password"
groups:
- "GroupB"
```

## Starting and Stopping WebLogic Servers

When you install **oc-cn-op-job-helm-chart**, the default configuration sets up a WebLogic Cluster with five Managed Servers. When you install or upgrade **oc-cn-helm-chart** for the Billing Care REST API service, two of the Managed Servers and one Admin Server are started.

By modifying the **override-values.yaml** file for **oc-cn-helm-chart**, you can control:

- The total number of Managed Servers and the initial server start up by using the **totalManagedServers** and **initialServerCount** keys.
- Whether the servers are started or stopped by using the **serverStartPolicy** key. To start the Admin Servers and the Managed Servers in a Cluster, set the key to **IF\_NEEDED**. To stop all servers, set the key to **NEVER**.

### Note

The keys in the **override-values.yaml** file should be the same as the ones used in **oc-cn-op-job-helm-chart** for keys that are common in both charts.

After you modify the **override-values.yaml** file, update the Helm release for the changes to take effect:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

# 10

## Configuring ECE Services

Learn how to configure Oracle Communications Elastic Charging Engine (ECE) services by configuring and deploying the ECE Helm chart.

Topics in this document:

- [Adding Elastic Charging Engine Keys](#)
- [Enabling SSL in Elastic Charging Engine](#)
- [Connecting ECE Cloud Native to an SSL-Enabled Database](#)
- [About Elastic Charging Engine Volume Mounts](#)
- [Loading Custom Diameter AVP](#)
- [Generating CDRs for Unrated Events](#)
- [Configuring ECE to Support Prepaid Usage Overage](#)
- [Recording Failed ECE Usage Requests](#)
- [Loading BRM Configuration XML Files](#)
- [Setting Up Notification Handling in ECE](#)
- [Configuring ECE for a Multischema BRM Environment](#)

For information about performing administrative tasks on your ECE cloud native services, see "Administering ECE Cloud Native Services" in *BRM Cloud Native System Administrator's Guide*.

Before installing the ECE Helm chart, you must first publish the metadata, config, and pricing data from the PDC pod.

### Note

Kubernetes looks for the CPU limit setting for pods. If it's not set, Kubernetes allocates a default value of 1 CPU per pod, which causes CPU overhead and Coherence scalability issues. To prevent this from happening, override each ECE pod's CPU limit to be the maximum CPU available on the node.

## Adding Elastic Charging Engine Keys

[Table 10-1](#) lists the keys that directly impact ECE deployment. Add these keys to your **override-values.yaml** file for **oc-cn-ece-helm-chart**. In the table, *component-name* should be replaced with the name of the ECE component, such as **emgateway**, **radiusgateway**, **diametergateway**, **httpgateway**, and **ratedeventformatter**.

Table 10-1 Elastic Charging Engine Keys

Key	Path in values.yaml File	Description
<code>imagePullPolicy</code>	<code>container</code>	The default value is <b>IfNotPresent</b> , which specifies to not pull the image if it's already present. Applicable values are <b>IfNotPresent</b> and <b>Always</b> .
<code>containerPort</code>	<code>container</code>	The port number that is exposed by this container.
<code>chargingSettingManagementPath</code>	<code>volume</code>	The location of the management folder, which contains the <b>charging-settings.xml</b> , <b>test-tools.xml</b> , and <b>migration-configuration.xml</b> files. The default is <b>/home/charging/opt/ECE/occeserver/config/management</b> .
<code>chargingSettingPath</code>	<code>volume</code>	The location of the configuration folder for ECE. The default is <b>/home/charging/opt/ECE/occeserver/config</b> .
<code>walletPassword</code>	<code>secretEnv</code>	The string password for opening the wallet.
<code>JMSQUEUEPASSWORD</code>	<code>secretEnv</code>	The password for the JMS queue, which is stored under the key <b>jms.queue.notif.pwd</b> in the wallet.
<code>RADIUSHAREDSECRET</code>	<code>secretEnv</code>	The RADIUS secret password, which is stored as <b>radius.secret.pwd</b> in the wallet.
<code>BRMGATEWAYPASSWORD</code>	<code>secretEnv</code>	The BRM Gateway password. This is the root user password for CM. <b>Note:</b> The password is not always of the root user and is dependent on its configuration in the <code>loginName</code> provided in the <code>BRMConnectionConfiguration</code> .
<code>PDCPASSWORD</code>	<code>secretEnv</code>	The PDC password, which is stored as <b>pdcc.pwd</b> in the wallet. <b>Note:</b> This key must match the <b>pdccAdminUserPassword</b> key in the <b>override-values.yaml</b> file for <b>oc-cn-helm-chart</b> .
<code>PDCKEYSTOREPASSWORD</code>	<code>secretEnv</code>	The PDC KeyStore password, which is stored as <b>pdcc.keystore.pwd</b> in the wallet. <b>Note:</b> This key must match the <b>keyStoreIdentityStorePass</b> key in the <b>override-values.yaml</b> file for <b>oc-cn-helm-chart</b> .
<code>PERSISTENCEDATABASEPASSWORD</code>	<code>secretEnv</code>	The database schema user password. This user is created using <code>ece-persistence-job</code> if it doesn't exist in the database.
<code>ECEHTTPGATEWAYSERVERSSLKEYSTOREPASSWORD</code>	<code>secretEnv</code>	The server SSL KeyStore password for the HTTP Gateway.
<code>BRM_SERVER_WALLET_PASSWORD</code>	<code>secretEnv</code>	The password to open the BRM server wallet.
<code>BRM_ROOT_WALLET_PASSWORD</code>	<code>secretEnv</code>	The root wallet password of the BRM wallet.

Table 10-1 (Cont.) Elastic Charging Engine Keys

Key	Path in values.yaml File	Description
<b>BRMDATABASEPASSWORD</b>	<b>secretEnv</b>	<p>The password for the BRM database.</p> <p>If you are connecting ECE to a BRM multischema database, use these entries instead:</p> <pre>BRMDATABASEPASSWORD:   - schema: 1     PASSWORD: Password   - schema: 2     PASSWORD: Password</pre> <p>where:</p> <ul style="list-style-type: none"> <li><b>schema</b> is the schema number. Enter <b>1</b> for the primary schema, <b>2</b> for the secondary schema, and so on.</li> <li><b>PASSWORD</b> is the schema password.</li> </ul>
<b>SSEENABLED</b>	<b>sslconnectioncertificates</b>	Whether SSL is enabled in ECE ( <b>true</b> ) or not ( <b>false</b> ).
<b>DNAME</b>	<b>sslconnectioncertificates</b>	The domain name. For example: "CN=Admin, OU=Oracle Communication Application, O=Oracle Corporation, L=Redwood Shores, S=California, C=US"
<b>SSLKEYSTOREVALIDITY</b>	<b>sslconnectioncertificates</b>	The validity of the KeyStore, in days. A value of <b>200</b> indicates that the validity is 200 days.
<b>runjob</b>	<b>job.sdk</b>	<p>Whether the SDK job needs to be run as part of the deployment (<b>true</b>) or not (<b>false</b>). The default value is <b>false</b>.</p> <p>If set to <b>true</b>, a default SDK job is run as part of the Helm installation or upgrade.</p>
<b>serviceFqdn</b>	<b>emgateway</b>	The default is <b>ece-emg</b> .
<b>tlsVersion</b>	<b>customerUpdater.customerUpdaterList.oracleQueueConnectionConfiguration</b>	(Release 15.0.1 or later) The TLS version to support, such as 1.2 or 1.3
<b>replicas</b>	<i>component-name.component-nameList</i>	The number of replicas to be created while deploying the chart. The default replica count is 3 for ecs server, and 1 for all other components.
<b>coherenceMemberName</b>	<i>component-name.component-nameList</i>	The Coherence member name under which this component will be added to the Coherence cluster.
<b>jmxEnabled</b>	<i>component-name.component-nameList</i>	Whether the component is JMX-enabled ( <b>true</b> ) or not ( <b>false</b> ).
<b>coherencePort</b>	<i>component-name.component-nameList</i>	The optional value indicating the Coherence port used by the component.
<b>jvmGCOpts</b>	<i>component-name.component-nameList</i>	This field helps to provide the Java JVM options such as GC details, max memory, and min memory.

Table 10-1 (Cont.) Elastic Charging Engine Keys

Key	Path in values.yaml File	Description
<b>jvmJMXOpts</b>	<i>component-name.component-nameList</i>	This field helps to provide the JMX-related option.
<b>jvmCoherenceOpts</b>	<i>component-name.component-nameList</i>	This field helps to provide the Coherence-related options such as the override file and cache config file.
<b>jvmOpts</b>	<i>component-name.component-nameList</i>	This field is empty by default, and any additional JVM arguments can be provided here.
<b>labels</b>	<b>charging</b>	The label for all pods in the deployment. The default value is <b>ece</b> .
<b>jmxport</b>	<b>charging</b>	The JMX port exposed by ece, which can be used to log in to JConsole. The default is <b>31022</b> .
<b>terminationGracePeriodSeconds</b>	<b>charging</b>	Used for graceful shutdown of the pods. The default value is 180 seconds.
<b>persistenceEnabled</b>	<b>charging</b>	Whether to persist the ECE cache data into the Oracle database. The default is <b>true</b> . See "Enabling Persistence in ECE" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.
<b>hpaEnabled</b>	<b>charging</b>	Whether to enable autoscaling using Kubernetes Horizontal Pod Autoscaler. See "Setting Up Autoscaling of ECE Pods" in <i>BRM Cloud Native System Administrator's Guide</i> for more information.
<b>timeoutSurvivorQuorum</b>	<b>charging</b>	The minimum number of cluster members that must remain in the cluster when the cluster service is terminating suspect members, without data loss. The default is <b>3</b> . To calculate the minimum number, use this formula: <b>(chargingServerWorkerNodes – 1) * (sum of all ecs pods/chargingServerWorkerNodes)</b>
<b>chargingServerWorkerNodes</b>	<b>charging</b>	The number of charging server worker nodes. The default is <b>3</b> .
<b>primary.*</b>	<b>charging.cluster</b>	The details about your primary cluster: <ul style="list-style-type: none"> <li><b>clusterName</b>: The name of the primary cluster.</li> <li><b>eceServiceName</b>: The ECE service name that creates the Kubernetes cluster with all of the ECE components in the primary cluster. The default is <b>ece-server</b>.</li> <li><b>eceServicefqdnOrExternalIP</b>: The fully qualified domain name (FQDN) or external IP address of the ECE service running in the primary cluster. For example: <b>ece-server.NameSpace.svc.cluster.local</b>.</li> </ul>

Table 10-1 (Cont.) Elastic Charging Engine Keys

Key	Path in values.yaml File	Description
<b>secondary.*</b>	<b>charging.cluster</b>	The details about your secondary cluster: <ul style="list-style-type: none"> <li><b>clusterName</b>: The name of the secondary cluster.</li> <li><b>eceServiceName</b>: The ECE service name that creates the Kubernetes cluster with all of the ECE components in the secondary cluster. The default is <b>ece-server</b>.</li> <li><b>eceServicefqdnOrExternalIP</b>: The fully qualified domain name (FQDN) or external IP address of the ECE service running in the secondary cluster. For example: <b>ece-server.NameSpace.svc.cluster.local</b>.</li> </ul>
<b>&lt;tags&gt;</b>	<b>migration</b>	The different tags indicating the values that will be stored under <b>migration-configuration.xml</b> . The tag names are the same as the ones used in the <b>migration-configuration.xml</b> file for ease of mapping.
<b>&lt;tags&gt;</b>	<b>testtools</b>	The different tags indicating the values that will be stored under <b>test-tools.xml</b> . The tag names are the same as the ones used in the <b>test-tools.xml</b> file for ease of mapping.
<b>&lt;module&gt;</b>	<b>log4j2.logger</b>	The different log levels for each module represents the logging level for the corresponding module.
<b>&lt;tags&gt;</b>	<b>eceproperties</b>	The different tags indicating the values that will be stored under <b>ece.properties</b> . The tag names are the same as the ones used in the <b>ece.properties</b> file for ease of mapping.
<b>&lt;tags&gt;</b>	<b>JMSConfiguration</b>	The different tags indicating the values that will be stored under <b>JMSConfiguration.xml</b> . The tag names are the same as the ones used in the <b>JMSConfiguration.xml</b> file for ease of mapping.
<b>name</b>	<b>secretEnv</b>	The user-defined name to give for the Secrets. The default is <b>secret-env</b> .
<b>SSEENABLED</b>	<b>sslconnectioncertificates</b>	Whether to install ECE under SSL mode ( <b>true</b> ) or not ( <b>false</b> ). The default is <b>true</b> .
<b>name</b>	<b>pv.external</b>	The name of the external PV. The default is <b>external-pv</b> .
<b>hostpath</b>	<b>pv.external</b>	The location on the host system of the external PV. The default is <b>/scratch/qa/ece_config/</b> .
<b>accessModes</b>	<b>pv.external</b>	The access mode for the PV. The default is <b>ReadWriteMany</b> .
<b>capacity</b>	<b>pv.external</b>	The maximum capacity of the external PV.
<b>name</b>	<b>pvc.logs</b>	The name for the ECE log files. The default is <b>logs-pv</b> .
<b>hostPath</b>	<b>pvc.logs</b>	The location on the host system for the ECE log files.

Table 10-1 (Cont.) Elastic Charging Engine Keys

Key	Path in values.yaml File	Description
<b>accessModes</b>	<b>pvc.logs</b>	The access mode for the PVC. The default is <b>ReadWriteMany</b> , since all of the different component pods will be writing their respective log files into a single logs directory.
<b>storage</b>	<b>pvc.logs</b>	The storage space required initially to create this PVC. If the storage specified here is not available in the machine, ensure that the PVC is not created and that the pods do not get initialized.
<b>name</b>	<b>pvc.brmconfig</b>	The name of the BRM Config PVC, in which all BRM configuration files such as the payload config file are exposed outside of the pod.
<b>accessModes</b>	<b>pvc.brmconfig</b>	The access mode for the PVC. The default is <b>ReadWriteMany</b> .
<b>storage</b>	<b>pvc.brmconfig</b>	The storage space required initially to create this PVC. If the storage specified here is not available in the machine, ensure that the PVC is not created and that the pods do not get initialized.
<b>name</b>	<b>pvc.sdk</b>	The name for the SDK PVC, in which all of the SDK files such as the config, sample scripts, and source files are exposed to the user.
<b>accessModes</b>	<b>pvc.sdk</b>	The access mode for the PVC. The default is <b>ReadWriteMany</b> .
<b>storage</b>	<b>pvc.sdk</b>	The storage space required initially to create this PVC. If the storage specified here is not available in the machine, ensure that the PVC is not created and that the pods do not get initialized.
<b>name</b>	<b>pvc.wallet</b>	The name for the wallet PVC, in which the wallet directory will be stored and shared by all of the ecs pods. The default is <b>ece-wallet-pvc</b> .
<b>accessModes</b>	<b>pvc.wallet</b>	The access mode for the PVC. The default is <b>ReadWriteMany</b> .
<b>storage</b>	<b>pvc.wallet</b>	The storage space required initially to create this PVC. If the storage specified here is not available in the machine, ensure that the PVC is not created and that the pods do not get initialized.
<b>name</b>	<b>pvc.external</b>	The name for the external PVC, in which the third-party JARs can be placed to share with the pods. The default is <b>external-pvc</b> .
<b>accessModes</b>	<b>pvc.external</b>	The access mode for the PVC. The default is <b>ReadWriteMany</b> .
<b>storage</b>	<b>pvc.external</b>	The storage space required initially to create this PVC. If the storage specified here is not available in the machine, ensure that the PVC is not created and that the pods do not get initialized.
<b>name</b>	<b>pvc.rel</b>	The name of the RE Loader PVC as created in the BRM deployment.
<b>name</b>	<b>storageClass</b>	The name of the storage class.

## Enabling SSL in Elastic Charging Engine

To complete the configuration for SSL setup in ECE:

1. Set these keys in the **override-values.yaml** file for **oc-cn-ece-helm-chart**:
  - **sslconnectioncertificates.SSLENABLED**: Set this to **true**.
  - **sslEnabled**: Set this to **true** in **emGatewayConfigurations**, **httpGatewayConfigurations**, and **BRMConnectionConfiguration**.
  - **migration.pricingUpdater.keyStoreLocation**: Set this to **/home/charging/opt/ECE/oceceserver/config/client.jks**.
  - **charging.brmWalletServerLocation**: Set this to **/home/charging/wallet/brmwallet/server/cwallet.sso**.
  - **charging.brmWalletClientLocation**: Set this to **/home/charging/wallet/brmwallet/client/cwallet.sso**.
  - **charging.brmWalletLocation**: Set this to **/home/charging/wallet/brmwallet**.
  - **charging.emGatewayConfigurations.emGatewayConfigurationList.emGateway1Config.wallet**: Set this to the BRM wallet location.
  - **charging.emGatewayConfigurations.emGatewayConfigurationList.emGateway2Config.wallet**: Set this to the BRM wallet location.
  - **charging.radiusGatewayConfigurations.wallet**: Set this to the BRM wallet location.
  - **charging.connectionConfigurations.BRMConnectionConfiguration.brmwallet**: Set this to the BRM wallet location.
2. Copy the SSL certificates, such as **client.jks** and **public-admin.cer**, generated from PDC to the **pdc\_ssl\_keystore** directory in the external PVC.
3. Configure the **connectionURL**, **port**, and **protocol** as per the PDC-configured t3s channel.

## Connecting ECE Cloud Native to an SSL-Enabled Database

To connect your ECE cloud native services to an SSL-enabled Oracle database:

1. Prepare for persistence schema creation.
  - a. Go to the **oc-cn-ece-helm-chart** directory, and then create a directory named **ece\_ssl\_db\_wallet/schema1**:

```
cd oc-cn-ece-helm-chart
mkdir -p ece_ssl_db_wallet/schema1
```
  - b. Save the contents of the ECE SSL database wallet to the **schema1** directory.
  - c. Grant the necessary permissions to the **ece\_ssl\_db\_wallet** directory:

```
chmod -R 775 ece_ssl_db_wallet
```
  - d. For multischema systems only, create a directory named **schema2** in the **ece\_ssl\_db\_wallet** directory and then copy the ECE SSL database wallet to the **schema2** directory.
2. Configure the SSL database wallets in the external volume mount.
  - a. Go to the external volume mount location (**external-pvc**).

- b. Create a directory named **ece\_ssl\_db\_wallet/schema1**:  

```
mkdir -p ece_ssl_db_wallet/schema1
```
  - c. Save the contents of the ECE SSL database wallet to the **ece\_ssl\_db\_wallet/schema1** directory.
  - d. Create a directory named **brm\_ssl\_db\_wallet/schema1**:  

```
mkdir -p brm_ssl_db_wallet/schema1
```
  - e. Save the contents of the BRM SSL database to the **brm\_ssl\_db\_wallet/schema1** directory.
  - f. Grant the necessary permissions to both new directories:  

```
chmod -R 775 ece_ssl_db_wallet brm_ssl_db_wallet
```
  - g. For multischema systems only, create a **schema2** directory inside both **ece\_ssl\_db\_wallet** and **brm\_ssl\_db\_wallet** directories. Then, copy the contents of the ECE SSL database to the **ece\_ssl\_db\_wallet/schema2** directory, and copy the contents of the BRM SSL database to the **brm\_ssl\_db\_wallet/schema2** directory.
3. Configure ECE for an SSL-enabled Oracle persistence database.  
Under the **charging.connectionConfigurations.OraclePersistenceConnectionConfigurations** section, set the following keys:
    - **dbSSLEnabled**: Set this to **true**.
    - **dbSSLType**: Set this to the type of SSL connection required for connecting to the database: **oneway**, **twoway**, or **none**.
    - **sslServerCertDN**: Set this to the SSL server certificate distinguished name (DN). The default is **DC=local,DC=oracle,CN=pindb**.
    - **trustStoreLocation**: Set this to **/home/charging/ext/ece\_ssl\_db\_wallet/schema1/cwallet.sso**.
    - **trustStoreType**: Set this to the type of file specified as the TrustStore for SSL connections: **SSO** or **pkcs12**.
  4. Configure customerUpdater for an SSL-enabled Oracle AQ database queue.  
Under the **customerUpdater.customerUpdaterList.oracleQueueConnectionConfiguration** section, set the following keys:
    - **dbSSLEnabled**: Set this to **true**.
    - **dbSSLType**: Set this to the type of SSL connection required for connecting to the database: **oneway**, **twoway**, or **none**.
    - **sslServerCertDN**: Set this to the SSL server certificate distinguished name (DN). The default is **DC=local,DC=oracle,CN=pindb**.
    - **trustStoreLocation**: Set this to **/home/charging/ext/brm\_ssl\_db\_wallet/schema1/cwallet.sso**.
    - **trustStoreType**: Set this to the type of file specified as the TrustStore for SSL connections: **SSO** or **pkcs12**.

**Note**

For database connectivity, ECE supports only the database service name and not the database service ID. Therefore, set the following keys to the database service name:

- **charging.connectionConfigurations.OraclePersistenceConnectionConfigurations.sid**
- **customerUpdater.customerUpdaterList.oracleQueueConnectionConfiguration.sid**

5. Configure your Oracle database configuration files to connect to the SSL-enabled BRM and ECE databases:

- a. Copy the **tnsnames.ora** and **sqlnet.ora** files from your SSL database host to your ECE cloud native instance.

- b. On your ECE cloud native instance, go to the ECE Helm chart directory:

```
cd oc-cn-ece-helm-chart
```

- c. Create the **ora\_files/ece** and **ora\_files/brm** directories:

```
mkdir -p ora_files/ece/
mkdir -p ora_files/brm/
```

- d. Copy the ECE database **tnsnames.ora** and **sqlnet.ora** files to the **oc-cn-ece-helm-chart/ora\_files/ece/** directory.

- e. In the **oc-cn-ece-helm-chart/ora\_files/ece/sqlnet.ora** file, set the wallet location to **/home/charging/opt/ECE/occeserver/config/ece\_ssl\_db\_wallet/schema1**:

```
WALLET_LOCATION =
  (SOURCE =
    (METHOD = FILE)
    (METHOD_DATA =
      (DIRECTORY = /home/charging/opt/ECE/occeserver/config/ece_ssl_db_wallet/
schema1)
    )
  )
```

- f. Copy the BRM database **tnsnames.ora** and **sqlnet.ora** files to the **oc-cn-ece-helm-chart/ora\_files/brm**.

- g. In the **oc-cn-ece-helm-chart/ora\_files/brm/sqlnet.ora** file, set the wallet location to **/home/charging/opt/ECE/occeserver/config/brm\_ssl\_db\_wallet/schema1**:

```
WALLET_LOCATION =
  (SOURCE =
    (METHOD = FILE)
    (METHOD_DATA =
      (DIRECTORY = /home/charging/opt/ECE/occeserver/config/brm_ssl_db_wallet/
schema1)
    )
  )
```

- h. Set the permissions for the **ora\_files** directory:

```
chmod -R 775 ora_files/
```

- i. Copy the **ora\_files** directory to the external PV mount location:

```
cp -rf ora_files externalPvMountLocation
```

## About Elastic Charging Engine Volume Mounts

### Note

You must use a provisioner that has ReadWriteMany access and sharing between pods.

The ECE container requires Kubernetes volume mounts for third-party libraries. The third-party volume mount shares the third-party libraries required by ECE from the host system with the container file system. For the list of third-party libraries to download, see "ECE Software Compatibility" in *BRM Compatibility Matrix*. Place the library files under the third-party volume mount.

The default configuration comes with a hostPath PersistentVolume. For more information, see "[Configure a Pod to Use a PersistentVolume for Storage](#)" in *Kubernetes Tasks*.

To use a different type of PersistentVolume, modify the `oc-cn-ece-helm-chart/templates/ece-pvc.yaml` file.

## Loading Custom Diameter AVP

To load custom Diameter AVPs into your ECE cloud native environment:

1. Create a **diameter** directory inside **external-pvc**.
2. Move the custom AVP file, such as **dictionary\_custom.xml**, to the **diameter** directory.
3. If you need to load a custom AVP after ECE is set up, restart the diametergateway pod by doing the following:
  - a. Increment the **diametergateway.diametergatewayList..restartCount** key by 1.
  - b. Run the **helm upgrade** command to update the release.

## Generating CDRs for Unrated Events

By default, the httpgateway pod sends all 5G usage requests to the ecs pod for online and offline charging.

You can configure httpgateway to convert some 5G usage requests into call detail record (CDR) files based on the charging type. You can then send the CDR files to roaming partners, a data warehousing system, or legacy billing systems for rating. For more information, see "About Generating CDRs" in *ECE Implementing Charging*.

You use the following to generate CDRs:

- httpgateway pod
- cdrgateway pod
- cdrFormatter pod
- CDR database

The cdrgateway and cdrFormatter pods can be scaled together, with one each per schema, or independently of the schemas. For more information, see "[Scaling the cdrgateway and cdrFormatter Pods](#)".

For details about the CDR format, see "CHF-CDR Format" in *ECE 5G CHF Protocol Implementation Conformance Statement*.

To set up ECE cloud native to generate CDRs:

1. Configure your httpgateway pod to do the following:
  - Generate CDRs (set **cdrGenerationEnabled** to **true**).
  - Route offline charging requests to the ecs pod for rating (set **rateOfflineCDRinRealtime** to **true**) or to the cdrgateway pod for generating CDRs (set **rateOfflineCDRinRealtime** to **false**).
  - Route online charging requests to the ecs pod for rating (set **generateCDRsForOnlineRequests** to **false**) or to the cdrgateway pod for generating CDRs (set **generateCDRsForOnlineRequests** to **true**).
2. Configure the cdrgateway pod to connect to the CDR database and do the following:
  - Generate individual CDR records for each request (set **individualCdr** to **true**) or aggregate multiple requests into a CDR record based on trigger criteria (set **individualCdr** to **false**). For information about the trigger criteria, see "About Trigger Types" in *ECE Implementing Charging*.
  - Store CDR records in an Oracle NoSQL database (set **isNoSQLConnection** to **true**) or in an Oracle database (set **isNoSQLConnection** to **false**).
3. Configure the cdrFormatter pod to do the following:
  - Retrieve batches of CDR records from the CDR database and pass them to a specified cdrFormatter plug-in for processing.
  - Purge processed CDR records from the CDR database older than a specified number of days (configured in **retainDuration**).
  - Purge orphan CDR records from the CDR database.  
  
Orphan CDR records are incomplete ones that are older than a specified number of seconds (configured in **cdrOrphanRecordCleanupAgeInSec**). Orphan CDR records can be created when your ECE system goes down due to maintenance or failure.
4. Configure the cdrFormatter plug-in to do the following:
  - Write a specified number of CDR records to each CDR file (set **maxCdrCount** to the maximum number).
  - Create JSON-formatted CDR files and then store them in your file system (set **enableDiskPersistence** to **true**) or send them to your Kafka messaging service (set **enableKafkaIntegration** to **true**).

To generate CDRs in ECE cloud native, you configure the following entries in your **override-values.yaml** file. This example configures:

- httpgateway to route both online and offline charging requests to cdrgateway.
- cdrgateway to aggregate multiple requests into a CDR record and then store it in an Oracle NoSQL database.
- cdrFormatter to retrieve CDR records in batches of 2500 from the Oracle NoSQL database and then send them to the default plug-in module. Immediately after CDR records are retrieved, cdrFormatter purges them from the database. It would also purge orphan records older than 200 seconds from the database.

- The `cdrFormatter` plug-in to create CDR files with a maximum of 20000 CDR records and an `.out` file name extension. It would store them in your file system in the path `/home/charging/cdr_input`.

```

cdrFormatter:
  cdrFormatterList:
    - schemaNumber: "1"
      replicas: 1
      coherenceMemberName: "cdrformatter1"
      jmxEnabled: true
      jvmGCOpts: "-XX:+UnlockExperimentalVMOptions -XX:+AlwaysPreTouch -
XX:G1RSetRegionEntries=2048 -XX:ParallelGCThreads=10 -XX:+ParallelRefProcEnabled -
XX:MetaspaceSize=100M -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintGCTimeStamps -
XX:+PrintTenuringDistribution -XX:+PrintAdaptiveSizePolicy -XX:-UseGCLogFileRotation -
XX:+UseG1GC -XX:NumberOfGCLogFile=99"
      jvmOpts: "-Xms16g -Xmx20g -Dece.metrics.http.service.enabled=true"
  cdrFormatterConfiguration:
    name: "cdrformatter1"
    clusterName: "BRM"
    primaryInstanceName: "cdrformatter1"
    partition: "1"
    isNoSQLConnection: "true"
    noSQLConnectionName: "noSQLConnection"
    connectionName: "oraclePersistenceelbrm"
    threadPoolSize: "6"
    retainDuration: "0"
    ripeDuration: "60"
    checkPointInterval: "6"
    maxPersistenceCatchupTime: "0"
    pluginPath: "ece-cdrformatter.jar"
    pluginType:
      "oracle.communication.brm.charging.cdr.formatterplugin.internal.SampleCdrFormatterCustomP
ugin"
    pluginName: "cdrFormatterPlugin1"
    noSQLBatchSize: "2500"
    cdrStoreFetchSize: "2500"
    cdrOrphanRecordCleanupAgeInSec: "200"
    cdrOrphanRecordCleanupSleepIntervalInSec: "200"
    enableIncompleteCdrDetection: "false"

cdrgateway:
  cdrgatewayList:
    - coherenceMemberName: "cdrgateway1"
      replicas: 6
      jmxEnabled: true
      jvmGCOpts: "-XX:+UnlockExperimentalVMOptions -XX:+AlwaysPreTouch -
XX:G1RSetRegionEntries=2048 -XX:ParallelGCThreads=10 -XX:+ParallelRefProcEnabled -
XX:MetaspaceSize=100M -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintGCTimeStamps -
XX:+PrintTenuringDistribution -XX:+PrintAdaptiveSizePolicy -XX:-UseGCLogFileRotation -
XX:+UseG1GC -XX:NumberOfGCLogFile=99"
      jvmJMXOpts: "-Dcom.sun.management.jmxremote -
Dcom.sun.management.jmxremote.authenticate=false -
Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.local.only=false"
      jvmCoherenceOpts: "-Dpof.config=charging-pof-config.xml -
Dcoherence.override=charging-coherence-override-dev.xml -Dcoherence.security=false -
Dsecure.access.name=admin"
      jvmOpts: "-Xms6g -Xmx8g -Dece.metrics.http.service.enabled=true -
DcdrServerCorePoolSize=64 -Dserver.sockets.metrics.bind-address=0.0.0.0 -
Dece.metrics.http.port=19612"
      restartCount: "0"
      cdrGatewayConfiguration:
        name: "cdrgateway1"

```

```

clusterName: "BRM"
primaryInstanceName: "cdrgateway1"
schemaNumber: "1"
isNoSQLConnection: "true"
noSQLConnectionName: "noSQLConnection"
connectionName: "oraclePersistence1"
cdrPort: "8084"
cdrHost: "ece-cdrgatewayservice"
individualCdr: "false"
cdrServerCorePoolSize: "32"
cdrServerMaxPoolSize: "256"
enableIncompleteCdrDetection: "false"
retransmissionDuplicateDetectionEnabled: "false"

httpgateway:
  cdrGenerationEnabled: "true"
  cdrGenerationStandaloneMode: "true"
  rateOfflineCDRinRealtime: "false"
  generateCDRsForOnlineRequests: "true"
  httpgatewayList:
    - coherenceMemberName: "httpgateway1"
      replicas: 8
      maxreplicas: 8
      jvmGCOpts: "-XX:+AlwaysPreTouch -XX:G1RSetRegionEntries=2048 -
XX:ParallelGCThreads=10 -XX:+ParallelRefProcEnabled -XX:MetaspaceSize=100M -
XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintGCTimeStamps -
XX:+PrintTenuringDistribution -XX:+PrintAdaptiveSizePolicy -XX:-UseGCLogFileRotation -
XX:+UseG1GC -XX:NumberOfGCLogFiles=99"
      jvmOpts: "-Xms10g -Xmx14g -Djava.net.preferIPv4Addresses=true -
Dece.metrics.http.service.enabled=true -Dserver.sockets.metrics.bind-address=0.0.0.0 -
Dece.metrics.http.port=19612"
      httpGatewayConfiguration:
        name: "httpgateway1"
        processingThreadPoolSize: "200"
        processingQueueSize: "32768"
        kafkaBatchSize: "10"

connectionConfigurations:
  OraclePersistenceConnectionConfigurations:
    retryCount: "1"
    retryInterval: "1"
    maxStmtCacheSize: "100"
    connectionWaitTimeout: "3000"
    timeoutConnectionCheckInterval: "3000"
    inactiveConnectionTimeout: "3000"
    databaseConnectionTimeout: "6000"
    persistenceInitialPoolSize: "4"
    persistenceMinPoolSize: "4"
    persistenceMaxPoolSize: "20"
    reloadInitialPoolSize: "0"
    reloadMinPoolSize: "0"
    reloadMaxPoolSize: "20"
    ratedEventFormatterInitialPoolSize: "6"
    ratedEventFormatterMinPoolSize: "6"
    ratedEventFormatterMaxPoolSize: "24"

charging:
  cdrFormatterPlugins:
    cdrFormatterPluginConfigurationList:
      cdrFormatterPluginConfiguration:
        name: "cdrFormatterPlugin1"
        tempDirectoryPath: "/tmp/tmp"

```

```

doneDirectoryPath: "/home/charging/cdr_input"
doneFileExtension: ".out"
enableKafkaIntegration: "false"
enableDiskPersistence: "true"
maxCdrCount: "20000"
staleSessionCauseForRecordClosingString: "PARTIAL_RECORD"
enableStaleSessionCleanupCustomField: "false"

```

## Scaling the cdrgateway and cdrFormatter Pods

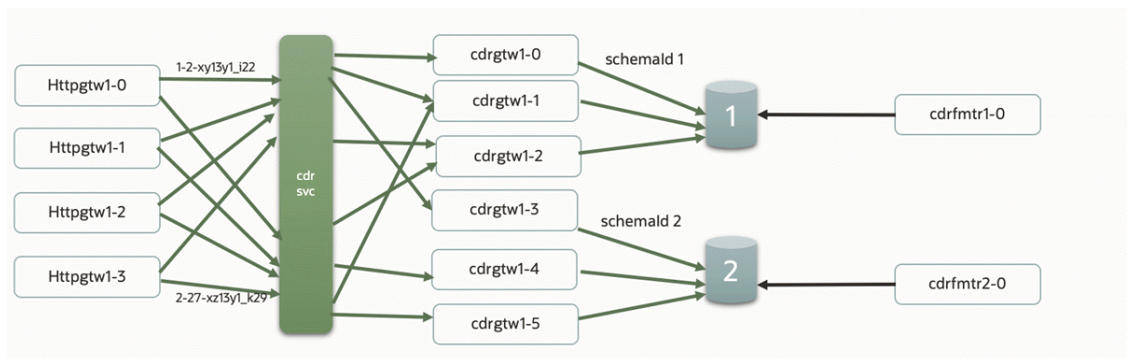
To increase performance and throughput, you can scale the cdrgateway and cdrFormatter pods together, with one each per schema, or scale them independently of the schemas.

[Figure 10-1](#) shows an example of scaled cdrgateway and cdrFormatter pods that have CDR storage in an Oracle Database. This example contains:

- One cdrgateway multi-replica deployment for all ECE schemas. All cdrgateway replicas have a single CDR Gateway service acting as a front end to httpgateway.
- One cdrFormatter single-replica deployment for each ECE schema. Each cdrFormatter reads persisted CDRs from its associated ECE schema.

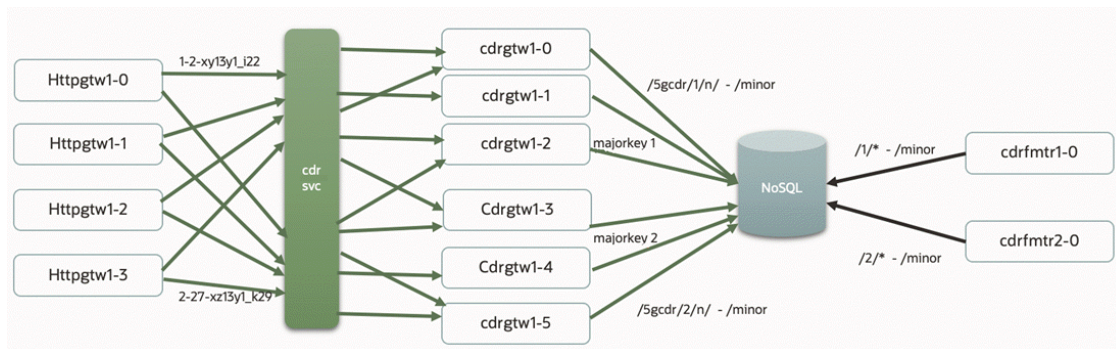
httpgateway forwards CDR requests to cdrgateway replicas in round-robin fashion. In this example, cdrgateway replicas 1-0, 1-1, and 1-2 persist CDRs in schema 1 tables, and replicas 1-3, 1-4, and 1-5 persist CDRs in schema 2 tables.

**Figure 10-1 Scaled Architecture with an Oracle Database**



[Figure 10-2](#) shows an example of scaled cdrgateway and cdrFormatter pods that have CDR storage in an Oracle NoSQL Database. This example contains:

- One cdrgateway multi-replica deployment for all ECE schemas. All cdrgateway replicas have a single CDR Gateway service acting as a front end to the httpgateway.
- One cdrFormatter single-replica deployment for each major key partition in the ECE schema. Each cdrFormatter reads persisted CDRs from its associated partition.

**Figure 10-2 Scaled Architecture with a NoSQL Database**

## Configuring ECE to Support Prepaid Usage Overage

You can configure ECE cloud native to capture any overage amounts by prepaid customers during an active session, which can help you prevent revenue leakage. If the network reports that the number of used units during a session is greater than a customer's available allowance, ECE cloud native charges the customer up to the available allowance. It then creates an overage record with information about the overage amount and sends it to the ECE Overage topic. You can create a custom solution for reprocessing the overage amount later on.

For example, assume a customer has a prepaid balance of 100 minutes, but uses 130 minutes during a session. ECE cloud native would charge the customer for 100 minutes, create an overage record for the remaining 30 minutes of usage, and then write the overage topic to the ECE Overage Kafka topic.

When the prepaid usage overage is disabled, ECE cloud native charges the customer for the full amount regardless of the amount of funds in the customer's balance.

To configure ECE cloud native to support prepaid usage overage, do the following:

- Ensure that ECE cloud native is connected to your Kafka Server
- Enable ECE cloud native to support prepaid usage overage
- Create an ECE Overage topic in your Kafka Server

To do so, set the following keys in your **override-values.yaml** file for **oc-cn-helm-chart**:

- **charging.kafkaConfigurations.kafkaConfigurationList.\***: Specify how to connect ECE to your Kafka Server.
- **charging.server.checkReservationOverImpact**: Set this to **true**.
- **charging.kafkaConfigurations.kafkaConfigurationList.overageTopicName**: Set this to the name of the Kafka topic where ECE will publish overage records.

## Recording Failed ECE Usage Requests

ECE cloud native may occasionally fail to process usage requests. For example, a data usage request could fail because a customer has insufficient funds. You can configure ECE cloud native to publish details about failed usage requests, such as the user ID and request payload, to the ECE failure topic in your Kafka server. Later on, you can reprocess the usage requests or view the failure details for analysis and reporting.

To configure ECE cloud native to record failed ECE usage requests:

- Ensure that ECE cloud native is connected your Kafka Server
- Enable the recording of failed ECE usage requests
- Create an ECE failure topic in your Kafka Server

To do so, set the following keys in your **override-values.yaml** file for **oc-cn-helm-chart**:

- **charging.kafkaConfigurations.kafkaConfigurationList.\***: Specify how to connect ECE to your Kafka Server.
- **charging.kafkaConfigurations.kafkaConfigurationList.persistFailedRequestsToKafkaTopic**: Set this to **true**.
- **charging.kafkaConfigurations.kafkaConfigurationList.failureTopicName**: Set this to the name of the topic that stores information about failed ECE usage requests.

## Loading BRM Configuration XML Files

BRM is configured by using the **pin\_notify** and **payloadconfig\_ece\_sync.xml** files. To ensure that the BRM pod can access these files for configuring the EAI Java Server (eai\_js), they are exposed through the brm\_config PVC within the pricingupdater pod. When new metadata is synchronized with ECE, if there are updates to the payload configuration file, it will create a new file in the location which can be accessed and configured in BRM.

For more information, see "Enabling Real-Time Synchronization of BRM and ECE Customer Data Updates" in *ECE Implementing Charging*.

## Setting Up Notification Handling in ECE

You can configure ECE cloud native to send notifications to a client application or an external application during an online charging session. For example, ECE cloud native could send a notification when a customer has breached a credit threshold or when a customer needs to request reauthorization.

You can set up ECE cloud native to send notifications by using either Apache Kafka topics or Oracle WebLogic queues:

- [Creating an Apache Kafka Notification Topic](#)
- [Creating an Oracle WebLogic Notification Queue](#)

## Creating an Apache Kafka Notification Topic

To create notification topics in Apache Kafka:

1. Create these Kafka topics either in the Kafka **entrypoint.sh** script or after the Kafka pod is ready:
  - **kafka.topicName**: ECENotifications
  - **kafka.suspenseTopicName**: ECESuspenseQueue
2. In the ZooKeeper runtime ConfigMap, set the **ece-zookeeper-0.ece-zookeeper.ECENamespace.svc.cluster.local** key to the name of the Kafka Cluster.
3. Set these Kafka and ZooKeeper-related environment variables appropriately:
  - **KAFKA\_PORT**: Set this to the port number in which Apache Kafka is up and running.

- **KAFKA\_HOST\_NAME**: Set this to the host name of the machine in which Apache Kafka is up and running. If it contains multiple Kafka brokers, create a comma-separated list.
  - **REPLICATION\_FACTOR**: Set this to the number of topic replications to create.
  - **PARTITIONS**: Set this to the total number of Kafka partitions to create in your topics. The recommended number to create is calculated as follows:  

$$[(\text{Max Diameter Gateways} * \text{Max Peers Per Gateway}) + (1 \text{ for BRM Gateway}) + \text{Internal Notifications}]$$
  - **TOPIC\_NAME**: Set this to **ECENotifications**. This is the name of the Kafka topic where ECE will publish notifications.
  - **SUSPENSE\_TOPIC\_NAME**: Set this to **ECESuspenseQueue**. This is the name of the Kafka topic where BRM will publish failed notifications and will retry later.
  - **ZK\_CLUSTER**: Set this to the name of your ZooKeeper cluster. This should match the value you set in step 2.
  - **ZK\_CLIENT\_PORT**: Set this to the port number in which ZooKeeper listens for client connections.
  - **ZK\_SERVER\_PORT**: Set this to the port number of the ZooKeeper server.
4. Ensure that the Kafka and ZooKeeper pods are in a READY state.
  5. Set these keys in your **override-values.yaml** file for **oc-cn-ece-helm-chart**:
    - **charging.server.kafkaEnabledForNotifications**: Set this to **true**.
    - **charging.server.kafkaConfigurations.name**: Set this to the name of your ECE cluster.
    - **charging.server.kafkaConfigurations.hostname**: Set this to the host name of the machine on which Kafka is up and running.
    - **charging.server.kafkaConfigurations.topicName**: Set this to **ECENotifications**.
    - **charging.server.kafkaConfigurations.suspenseTopicName**: Set this to **ECESuspenseQueue**.

6. Install the ECE cloud native service by entering this command from the **helmcharts** directory:

```
helm install EceReleaseName oc-cn-ece-helm-chart --namespace BrmNameSpace --values
OverrideValuesFile
```

The notification topics are created in Apache Kafka.

## Creating an Oracle WebLogic Notification Queue

To create notification queues and topics in Oracle WebLogic:

1. Ensure the following:
  - Oracle WebLogic is running in your Kubernetes cluster.
  - A separate WebLogic domain for the ECE Notification queues has been created.

**Note**

Do not create your ECE notification queues in an existing WebLogic domain. For example, do not use the Billing Care, Business Operations Center, PDC, or Billing Care REST API domains.

- The following third-party libraries are in the **3rdparty\_jars** directory inside **external-pvc**:
    - **external-pvc: com.oracle.weblogic.beangen.general.api.jar**
    - **wlthint3client.jar**
  - For SSL-enabled WebLogic in a disaster recovery environment, move a common JKS certificate file for all sites to the **ece\_ssl\_keystore** directory inside **external-pvc**.
2. Create an **override-values.yaml** file for **oc-cn-ece-helm-chart**.
  3. Set the following keys in your **override-values.yaml** file:
    - Set the **secretEnv.JMSQUEUEPASSWORD** key to the WebLogic user password.
    - If WebLogic SSL is enabled, set the **secretEnv.NOTIFYEVENTKEYPASS** key to the KeyStore password.
    - Set the **job.jmsconfig.runjob** key to **true**.
    - If the job needs to create the ECE JMS module and subdeployment, set the **job.jmsconfig.preCreateJmsServerAndModule** key to **true**.
    - Set the **charging.server.weblogic.jmsmodule** key to **ECE**.
    - Set the **charging.server.weblogic.subdeployment** key to **ECEQueue**.
    - Set the **charging.server.kafkaEnabledForNotifications** key to **false**.
    - In the **JMSConfiguration** section, set the **HostName**, **Port**, **Protocol**, **ConnectionURL**, and **KeyStoreLocation** keys to the appropriate values for your system.

For more information about these keys, see [Table 10-1](#).

4. Copy the SSL certificate file (**client.jks**) to the **ece\_ssl\_keystore** directory in the external PVC.
5. Install the ECE cloud native service by entering this command from the **helmcharts** directory:

```
helm install EceReleaseName oc-cn-ece-helm-chart --namespace BrmNameSpace --values OverrideValuesFile
```

The following are created in the ECE domain of your WebLogic Server:

- A WebLogic notification topic named **NotificationTopic**.
- A WebLogic notification queue named **SuspenseQueue**.
- A WebLogic connection factory named **NotificationFactory**.

Next, configure the connection factory resource so your clients can connect to the ECE notification queues and topics in Oracle WebLogic.

To configure the connection factory resource:

1. On the WebLogic Server in which the JMS ECE notification queue resides, sign in to WebLogic Server Administration Console.

2. In the **Domain Structure** tree, expand **Services**, expand **Messaging**, and then click **JMS Modules**.  
The Summary of JMS Modules page appears.
3. In the **JMS Modules** table, click on the name **ECE**.  
The Settings for ECE page appears.
4. In the **Summary of Resources** table, click on the name **NotificationFactory**.  
The Settings for NotificationFactory page appears.
5. Click the **Configuration** tab, and then click the **Client** tab.
6. On the **Client** page, do the following:
  - a. In **Client ID Policy**, select **Unrestricted**.
  - b. In **Subscription Sharing Policy**, select **Sharable**.
  - c. In **Reconnect Policy**, select **None**.
  - d. Click **Save**.
7. Click the **Transactions** tab.
8. On the **Transactions** page, do the following:
  - a. In **Transaction Timeout**, enter **2147483647** which is the maximum timeout value.
  - b. Click **Save**.

For more information, see [Oracle WebLogic Administration Console Online Help](#).

## Configuring ECE for a Multischema BRM Environment

If your BRM database contains multiple schemas, you must configure ECE to connect to each schema.

To configure ECE for a BRM multischema database:

1. Open your **override-values.yaml** file for the **oc-cn-ece-helm-chart** chart.
2. Specify the password for accessing each schema in the BRM database. To do so, configure these keys for each schema:
  - **secretEnv.BRMDATABASEPASSWORD.schema**: Set this to the schema number. Enter **1** for the primary schema, **2** for the secondary schema, and so on.
  - **secretEnv.BRMDATABASEPASSWORD.PASSWORD**: Set this to the schema password.

This shows example settings for two schemas:

```
secretEnv:
  BRMDATABASEPASSWORD:
    - schema: 1
      PASSWORD: Password
    - schema: 2
      PASSWORD: Password
```

3. Configure a customerUpdater pod for each schema. To do so, add a **-schemaNumber** list for each schema. In the list:
  - Set the **SchemaNumber** key to **1** for the primary schema, **2** for the secondary schema, and so on.

- Set the **amtAckQueueName** key to the fully qualified name of the acknowledgment queue to which the **pin\_amt** utility listens to Account Migration Manager (AMM)-related acknowledgment events. The value is in the format *primarySchema.ECE\_AMT\_ACK\_QUEUE*, where *primarySchema* is the name of the primary schema.
- Set the **hostName** and **jdbcUrl** keys to their corresponding values for each schema.

This shows example settings for two schemas:

```
customerUpdater:
  customerUpdaterList:
    - schemaNumber: "1"
      coherenceMemberName: "customerupdater1"
      replicas: 1
      jmxEnabled: true
      coherencePort: ""
      jvmGCOpts: ""
      jvmJMXOpts: ""
      jvmCoherenceOpts: ""
      jvmOpts: ""
      jmxport: ""
      restartCount: "0"
      oracleQueueConnectionConfiguration:
        name: "customerupdater1"
        gatewayName: "customerupdater1"
        hostName: ""
        port: "1521"
        sid: "pindb"
        userName: "pin"
        jdbcUrl: ""
        queueName: "IFW_SYNC_QUEUE"
        suspenseQueueName: "ECE_SUSPENSE_QUEUE"
        ackQueueName: "ECE_ACK_QUEUE"
        amtAckQueueName: "pin0101.ECE_AMT_ACK_QUEUE"
        batchSize: "1"
        dbTimeout: "900"
        retryCount: "10"
        retryInterval: "60"
        walletLocation: "/home/charging/wallet/ecewallet/"

    - schemaNumber: "2"
      coherenceMemberName: "customerupdater2"
      replicas: 1
      jmxEnabled: true
      coherencePort: ""
      jvmGCOpts: ""
      jvmJMXOpts: ""
      jvmCoherenceOpts: ""
      jvmOpts: ""
      jmxport: ""
      oracleQueueConnectionConfiguration:
        name: "customerupdater2"
        gatewayName: "customerupdater2"
        hostName: ""
        port: "1521"
        sid: "pindb"
        userName: "pin"
        jdbcUrl: ""
        queueName: "IFW_SYNC_QUEUE"
        suspenseQueueName: "ECE_SUSPENSE_QUEUE"
        ackQueueName: "ECE_ACK_QUEUE"
        amtAckQueueName: "pin0101.ECE_AMT_ACK_QUEUE"
```

```

batchSize: "1"
dbTimeout: "900"
retryCount: "10"
retryInterval: "60"
walletLocation: "/home/charging/wallet/ecewallet/"

```

4. Configure a `ratedEventFormatter` pod for processing rated events belonging to each BRM schema. To do so, add a **-schemaNumber** list for each schema. In the list, set the **schemaNumber** and **partition** keys to **1** for the primary schema, **2** for the secondary schema, and so on.

This shows example settings for two schemas:

```

ratedEventFormatter:
  ratedEventFormatterList:
    - schemaNumber: "1"
      replicas: 1
      coherenceMemberName: "ratedeventformatter1"
      jmxEnabled: true
      coherencePort:
      jvmGCOpts: ""
      jvmJMXOpts: ""
      jvmCoherenceOpts: ""
      jvmOpts: ""
      jmxport: ""
      restartCount: "0"
      ratedEventFormatterConfiguration:
        name: "ratedeventformatter1"
        primaryInstanceName: "ratedeventformatter1"
        partition: "1"
        noSQLConnectionName: "noSQLConnection"
        connectionName: "oraclePersistencel"
        threadPoolSize: "6"
        retainDuration: "0"
        ripeDuration: "600"
        checkPointInterval: "6"
        maxPersistenceCatchupTime: "60"
        siteName: ""
        pluginPath: "ece-ratedeventformatter.jar"
        pluginType:
"oracle.communication.brm.charging.ratedevent.formatterplugin.internal.BrmCdrPluginDi
rect"
        pluginName: "brmCdrPlugin1"
        noSQLBatchSize: "25"

    - schemaNumber: "2"
      replicas: 1
      coherenceMemberName: "ratedeventformatter2"
      jmxEnabled: true
      coherencePort:
      jvmGCOpts: ""
      jvmJMXOpts: ""
      jvmCoherenceOpts: ""
      jvmOpts: ""
      jmxport: ""
      ratedEventFormatterConfiguration:
        name: "ratedeventformatter2"
        primaryInstanceName: "ratedeventformatter2"
        partition: "2"
        noSQLConnectionName: "noSQLConnection"
        connectionName: "oraclePersistencel"
        threadPoolSize: "6"
        retainDuration: "0"

```

```

    ripeDuration: "600"
    checkPointInterval: "6"
    maxPersistenceCatchupTime: "60"
    siteName: ""
    pluginPath: "ece-ratedeventformatter.jar"
    pluginType:
"oracle.communication.brm.charging.ratedevent.formatterplugin.internal.BrmCdrPluginDirect"
    pluginName: "brmCdrPlugin1"
    noSQLBatchSize: "25"

```

5. Save and close your **override-values.yaml** file for **oc-cn-ece-helm-chart**.
6. In the **oc-cn-ece-helm-chart/templates/charging-settings.yaml** ConfigMap, add **poidIdConfiguration** in **itemAssignmentConfig** for each schema.

This shows example settings for three schemas:

```

<itemAssignmentConfigconfig-
class="oracle.communication.brm.charging.appconfiguration.beans.item.ItemAssignmentCo
nfig" itemAssignmentEnabled="true" delayToleranceIntervalInDays="0"
poidPersistenceSafeCount="12000">
  <schemaConfigurationGroup config-class="java.util.ArrayList">
    <poidIdConfigurationconfig-
class="oracle.communication.brm.charging.appconfiguration.beans.item.PoidIdConfigurat
ion" schemaName="1" poidQuantity="2000000">
      </poidIdConfiguration>
    <poidIdConfigurationconfig-
class="oracle.communication.brm.charging.appconfiguration.beans.item.PoidIdConfigurat
ion" schemaName="2" poidQuantity="2000000">
      </poidIdConfiguration>
    <poidIdConfigurationconfig-
class="oracle.communication.brm.charging.appconfiguration.beans.item.PoidIdConfigurat
ion" schemaName="3" poidQuantity="2000000">
      </poidIdConfiguration>
    </schemaConfigurationGroup>
  </itemAssignmentConfig>

```

After you deploy **oc-cn-ece-helm-chart** in "[Deploying BRM Cloud Native Services](#)", the ECE pods will be connected to your BRM database schemas.

# 11

## Deploying BRM Cloud Native Services

Learn how to deploy Oracle Communications Billing and Revenue Management (BRM) cloud native services by running the Helm install command.

Topics in this document:

- [Deploying BRM Cloud Native Services](#)

## Deploying BRM Cloud Native Services

### Note

The **oc-cn-init-db-helm-chart** and **oc-cn-helm-chart** charts must be deployed in different namespaces.

To deploy BRM cloud native services, do this:

1. Create a namespace for the BRM Helm chart.

```
kubectl create namespace BrmNameSpace
```

where *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.

2. Validate the content of your Helm charts by using the **Helm lint** command.

- For Helm 3.6.0 and later releases, enter these commands from the **helmcharts** directory:

```
helm lint --strict oc-cn-helm-chart --values oc-cn-helm-chart/values.yaml --values OverrideValuesFile
helm lint --strict oc-cn-ece-helm-chart --values oc-cn-ece-helm-chart/values.yaml --values OverrideValuesFile
helm lint --strict oc-cn-op-job-helm-chart --values oc-cn-op-job-helm-chart/values.yaml --values OverrideValuesFile
```

- For previous Helm releases, enter these commands from the **helmcharts** directory:

```
helm lint --strict oc-cn-helm-chart
helm lint --strict oc-cn-ece-helm-chart
helm lint --strict oc-cn-op-job-helm-chart
```

You'll see this if the commands complete successfully:

```
3 chart(s) listed, no failures
```

3. If you are using Pricing Design Center (PDC), Billing Care, the Billing Care REST API, Web Services Manager, or Business Operations Center, do this:
  - a. Ensure BRM images are available and BRM is deployed successfully, for PDC pods to deploy successfully
  - b. Direct WebLogic Kubernetes Operator to monitor the BRM namespace:

```
helm upgrade weblogic-operator weblogic-operator/weblogic-operator \
--namespace Operator \
--reuse-values \
--set "domainNamespaces={BrmNameSpace}" \
--wait
```

where *Operator* is the namespace you created for WebLogic Kubernetes Operator as part of the prerequisite tasks.

- c. Create WebLogic domains by entering this command from the **helmcharts** directory:

```
helm install OpJobReleaseName oc-cn-op-job-helm-chart --namespace BrmNameSpace --
values OverrideValuesFile
```

where *OpJobReleaseName* is the release name for **oc-cn-op-job-helm-chart** and is used to track this installation instance. It must be different from the one used for the BRM Helm chart.

4. Install BRM cloud native services by entering this command from the **helmcharts** directory:

```
helm install BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace
BrmNameSpace
```

where *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance. It must be different from the one used for **oc-cn-init-db-helm-chart**.

5. To install the ECE cloud native service, enter this command from the **helmcharts** directory:

```
helm install EceReleaseName oc-cn-ece-helm-chart --namespace BrmNameSpace --values
OverrideValuesFile
```

where *EceReleaseName* is the release name for **oc-cn-ece-helm-chart** and is used to track this installation instance. It must be different from the one used for the BRM Helm chart.

# Deploying into Oracle Cloud Infrastructure

Learn how to deploy Oracle Communications Billing and Revenue Management (BRM) cloud native services into Oracle Cloud Infrastructure.

Topics in this document:

- [Deploying into Oracle Cloud Infrastructure](#)

## Deploying into Oracle Cloud Infrastructure

Oracle Cloud Infrastructure is a set of complementary cloud services that enable you to run a wide range of applications and services in a highly available hosted environment. It offers high-performance computing capabilities (as physical hardware instances) and storage capacity in a flexible overlay virtual network that is securely accessible from your on-premise network. BRM cloud native deployment is tested in Oracle Cloud Infrastructure for the following services both on Virtual Machine and Bare Metal:

- BRM cloud native application and database running on IaaS
- BRM cloud native application managed by Oracle Kubernetes Engine and database on IaaS
- BRM cloud native application managed by Oracle Kubernetes Engine and database on DBaaS

Deploying the BRM cloud native services into Oracle Cloud Infrastructure involves these high-level steps:

### Note

These are the bare minimum tasks for deploying BRM cloud native services in Oracle Cloud Infrastructure. Your steps may vary from the ones listed below.

1. Sign up for Oracle Cloud Infrastructure.
2. Create a database system on a bare metal or virtual machine instance.  
Select a database version that is compatible with the BRM cloud native software requirements. See "BRM Software Compatibility" in *BRM Compatibility Matrix*.
3. Create a Kubernetes cluster and deselect the **Tiller (Helm) Enabled** option. The version of Helm used by Oracle Cloud Infrastructure isn't compatible with the BRM cloud native software requirements.
4. Install and configure the Oracle Cloud Infrastructure Command Line Interface (CLI).  
CLI is a small footprint tool that you can use on its own or with the Console to complete OCI tasks. It's needed here to download the **kubeconfig** file.
5. Install and configure **kubectl** on your system to perform operations on your cluster in Oracle Cloud Infrastructure.

6. The **kubeconfig** file (by default named **config** and stored in the **\$HOME/.kube** directory) provides the necessary details to access the cluster using **kubectl** and the Kubernetes Dashboard.

Download **kubeconfig** to access your cluster on Oracle Cloud Infrastructure by entering this command:

```
oci ce cluster create-kubeconfig --cluster-id ClusterId --file $HOME/.kube/  
config --region RegionId --token-version 2.0.0
```

where *ClusterId* is the Oracle Cloud Identifier (OCID) of the cluster, and *RegionId* is the region identifier such as us-phoenix-1 and us-ashburn-1.

7. Set the **\$KUBECONFIG** environment variable to the downloaded **kubeconfig** file by entering this command:

```
export KUBECONFIG=$HOME/.kube/config
```

8. Verify access to your cluster. You can enter this command and then match the output Internal IP Addresses and External IP Addresses against the nodes in your cluster in the Oracle Cloud Infrastructure Console.

```
kubectl get node -o wide
```

9. Download and configure Helm in your local system.
10. Place the BRM cloud native Helm chart on your system where you have downloaded and configured **kubectl** and Helm. Then, follow the instructions in "Configuring and Deploying BRM Cloud Native" in *BRM Cloud Native Deployment Guide*.

# 13

## Uninstalling Your BRM Cloud Native Deployment

Learn how to uninstall the Oracle Communications Billing and Revenue Management (BRM) cloud native deployment from your system.

Topics in this document:

- [Uninstalling Your BRM Cloud Native Deployment](#)
- [Uninstalling Selected BRM Cloud Native Services](#)

### Uninstalling Your BRM Cloud Native Deployment

When you uninstall a Helm chart from your BRM cloud native deployment, it removes only the Kubernetes objects that it created during installation.

To uninstall, enter this command:

```
helm delete ReleaseName --namespace Namespace
```

where:

- *ReleaseName* is the name you assigned to this installation instance.
- *Namespace* is the namespace in which the BRM Kubernetes objects reside.

### Uninstalling Selected BRM Cloud Native Services

Depending on the scenario, you might need to temporarily or permanently uninstall BRM cloud native services, such as Billing Care, the Billing Care REST API, or Business Operations Center, while retaining other BRM services. To do this, you upgrade your **oc-cn-helm-chart** release by disabling the service you intend to remove.

For example, to remove only the Billing Care REST API service, you would set the **ocbc.bcws.isEnabled** key to **false** in your **override-values.yaml** file and then upgrade your release of **oc-cn-helm-chart**:

```
helm upgrade --namespace Namespace ReleaseName oc-cn-helm-chart --values override-values.yaml
```

This would bring down the WebLogic servers that are hosting the Billing Care REST API and remove all resources created for this service through **oc-cn-helm-chart**.

# Part IV

## Customizing BRM Cloud Native

This part provides information about customizing Oracle Communications Billing and Revenue Management (BRM) cloud native. It contains the following chapters:

- [Customizing BRM Cloud Native Services](#)
- [Building Your Own Images](#)

# 14

## Customizing BRM Cloud Native Services

Learn how to customize the Oracle Communications Billing and Revenue Management (BRM) server and clients in a cloud native environment to meet your business needs.

The Podman build commands in this chapter reference Dockerfile and related scripts as is from the **oc-cn-docker-files-15.0.x.0.0.tgz** package. Ensure you use your own version of Dockerfile and related scripts before running the build command.

Topics in this document:

- [Customizing BRM Server](#)
- [Customizing Billing Care](#)
- [Customizing ECE](#)

### **Caution**

The Dockerfiles and related scripts are provided for reference only. You can refer to them to build or extend your own images. Support is restricted to core product issues only and no support will be provided for custom Dockerfiles and scripts.

## Customizing BRM Server

You can customize BRM Server by layering the BRM cloud native image with a customized library file.

For example, you could extend the **fm\_subscription\_pol\_custom.so** library file and layer it with the BRM cloud native image by doing this:

1. Customize your **lib/fm\_subscription\_pol\_custom.so** library file as follows:
  - a. Enable the BRM SDK by setting the following keys in your **override-values.yaml** file for **oc-cn-helm-chart**:

```
brm_sdk:
  isEnabled: true
  deployment:
    imageName: brm_sdk
    imageTag: 15.0.x.0.0
  pvc:
    storage: 50Mi
```

- b. Run the **helm upgrade** command to deploy the brm-sdk pod:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values
OverrideValuesFile --namespace BrmNameSpace
```

- c. Run the following **kubectl** command to retrieve the brm-sdk pod name:

```
kubectl get pods --namespace BrmNameSpace | grep brm-sdk
```

If successful, you should see something similar to this:

NAME	READY	STATUS	RESTARTS	AGE
brm-sdk-f67b95777-bf8j5	1/1	Running	0	18m

- d. Run the following **kubectl** command to retrieve the name of the PVC volume for brm-sdk:

```
kubectl get pvc --namespace BrmNameSpace | grep brm-sdk
```

If successful, you should see something similar to this:

NAME	STATUS	VOLUME	CAPACITY
brm-sdk	Bound	pvc-094feae0-4d11-4887-83a0-b47a0fc6a3f4	50Mi
		myclass	23h

- e. List the files and folders in **/mnt/oke\_test/brm** to verify that the PVC volume is present:

```
ls /mnt/oke_test/brm/ | grep pvc-094feae0-4d11-4887-83a0-b47a0fc6a3f4
```

If successful, you should see something similar to this:

```
brm-sdk-pvc-094feae0-4d11-4887-83a0-b47a0fc6a3f4
```

- f. Do one of the following:

- Copy the custom C file to the PVC:

```
cp customFile nfsMountPath/BrmNameSpace/pvcVolumePath/
```

For example:

```
cp fm_cust_pol_valid_billinfo.c /mnt/oke_test/brm/brm-sdk-  
pvc-094feae0-4d11-4887-83a0-b47a0fc6a3f4/
```

- Copy the custom C file to the **oc-cn-helm-chart/brm\_sdk\_scripts/** directory and run the **helm upgrade** command:

```
cp customFile oc-cn-helm-chart/brm_sdk_scripts/  
helm upgrade BrmReleaseName oc-cn-helm-chart --namespace  
BrmNameSpace --values oc-cn-helm-chart/override_values.yaml
```

For example:

```
cp fm_cust_pol_valid_billinfo.c oc-cn-helm-chart/brm_sdk_scripts/  
helm upgrade release oc-cn-helm-chart --namespace brm --values oc-  
cn-helm-chart/override_values.yaml
```

The files from **oc-cn-helm-chart/brm\_sdk\_scripts/** will be present at **/oms/load** in the brm-sdk pod.

- g. Run the **kubectl** command to get a shell to a running container:

```
kubectl exec --namespace BrmNameSpace -it brmSDKPod bash
```

For example:

```
kubectl exec --namespace brm -it brm-sdk-f67b95777-bf8j5 bash
```

- h. Build your custom library file in one of these ways:

- If you copied your custom C file to the PVC in step [1.f](#), do this:

```
cd source/sys/libraryName  
cp /oms/ext/fileName .  
make
```

For example:

```
cd source/sys/fm_cust_pol/  
cp /oms/ext/fm_cust_pol_valid_billinfo.c .  
make
```

- If you copied your custom C file to **oc-cn-helm-chart/brm\_sdk\_scripts/** in step [1.f](#), do this:

```
cd source/sys/libraryName  
cp /oms/load/fileName .  
make
```

For example:

```
cd source/sys/fm_cust_pol/  
cp /oms/load/fm_cust_pol_valid_billinfo.c .  
make
```

- i. Copy the generated library file to the PVC:

```
cp customLibrary.so /oms/ext/
```

For example:

```
cp fm_cust_pol_custom.so /oms/ext/
```

- j. Copy the library file from the PVC to the **\$PIN\_HOME/lib** directory:

```
cp nfsMountPath/BrmNameSpace/brmSDKPod/customLibrary.so $PIN_HOME/lib
```

For example:

```
cp /mnt/oke_test/brm/brm-sdk-pvc-094feae0-4d11-4887-83a0-b47a0fc6a3f4/
fm_cust_pol_custom.so $PIN_HOME/lib
```

2. Build the custom Connection Manager (CM) image using the Dockerfile:

```
FROM cm:15.0.x.0.0
USER root
COPY lib/fm_subscription_pol_custom.so $PIN_HOME/lib/fm_subscription_pol_custom.so
RUN chown -R omsuser:root $PIN_HOME/lib/fm_subscription_pol_custom.so && \
  chmod 755 ${PIN_HOME}/lib/fm_subscription_pol_custom.so
USER omsuser
```

3. Build the BRM Server image by entering this command:

```
podman build --format docker --tag cm:imageTag --file Dockerfile_cm .
```

4. Push the image to the image repository:

```
podman tag cm:imageTag imageRepository/cm:imageTag
podman push imageRepository/cm:imageTag
```

5. Update the custom image name in the **override-values.yaml** file. For example:

```
cm:
  isEnabled: true
  deployment:
    replicaCount: 1
    imageName: cm
    imageTag: imageTag
```

where *imageTag* must match the value used in step 4.

6. Add the custom configuration for the CM **pin.conf** to the **configmap\_pin\_conf\_cm.yaml** file.

For example, for the **fm\_cust\_pol\_custom** library:

```
- cm fm_module ${PIN_HOME}/lib/fm_cust_pol_custom${LIBRARYEXTENSION}
  fm_cust_pol_custom_config fm_cust_pol_init pin
```

7. Run the **helm upgrade** command to update the release with the new CM image:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --namespace BrmNameSpace --
values oc-cn-helm-chart/override_values.yaml
```

## Customizing Billing Care

Extensibility is one of the biggest features of on-premise Billing Care, and this same extensibility is available in the Billing Care cloud native deployment. You can override the

existing Billing Care behavior, such as changing labels and icons, add new flows and screens, and so on, by using the Billing Care SDK.

To use the Billing Care SDK in a cloud native environment, do this:

1. Build the Billing Care SDK WAR the same way as described in "Packaging and Deploying Customizations" in *Billing Care SDK Guide*.
2. Create a Billing Care SDK image by using the Linux image as a base.
3. Update the **override-values.yaml** file to direct **oc-cn-op-job-helm-chart** to deploy the SDK WAR file and link it with Billing Care or the Billing Care REST API WAR after deploying them.

The cloud native package includes all of the scripts necessary to prepare and run an SDK image. For example, if your SDK WAR is named **BillingCareCustomizations.war**, you would build the Billing Care SDK image like this:

1. Go to the **oc-cn-docker-files/ocbc/billing\_care\_sdk** directory.
2. Copy the **BillingCareCustomizations.war** file to your current working directory (**oc-cn-docker-files/ocbc/billing\_care\_sdk**).
3. Build the SDK image by entering this command:
 

```
podman build --format docker --build-arg SDK_WAR=BillingCareCustomizations.war --tag oracle/billingcare_sdk:15.0.x.0.0 .
```
4. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, edit the keys shown in [Table 14-1](#). This directs **oc-cn-op-job-helm-chart** to deploy the Billing Care SDK image rather than the Billing Care image and to include additional files that are needed for successful deployment of SDK.

**Table 14-1 Billing Care SDK Keys**

Key	Path	Description
<b>imageName</b>	<b>ocbc.bc.deployment.sdk</b> <b>ocbc.bcws.deployment.sdk</b>	Set this to <b>oracle/billingcare_sdk</b> . This is the name of the image, which must be used for the billingcare pod.
<b>imageTag</b>	<b>ocbc.bc.deployment.sdk</b> <b>ocbc.bcws.deployment.sdk</b>	Set this to <b>15.0.x.0.0</b> . This tags the image used for the billingcare pod.
<b>isEnabled</b>	<b>ocbc.bc.sdk</b> <b>ocbc.bcws.sdk</b>	Set this key to <b>true</b> if you want to deploy SDK.
<b>deployName</b>	<b>ocbc.bc.sdk</b> <b>ocbc.bcws.sdk</b>	The name of the SDK Library in the <b>Manifest.MF</b> file. The default is <b>BillingCareCustomizations</b> .

5. Install **oc-cn-op-job-helm-chart** followed by **oc-cn-helm-chart** to customize Billing Care or the Billing Care REST API with SDK.

## Customizing ECE

You can customize the ECE image by layering the native image with the customized code.

For example:

```
> cat Dockerfile_custom_ece
FROM oc-cn-ece:15.0.x.0.0
USER root
#commands that need to be run
USER eceuser
```

To build the image, run this Podman command:

```
podman build --format docker --tag customECE:15.0.x.0.0 --file customECEDockerfile .
```

where *customECE* is the name of your custom ECE Helm chart, and *customECEDockerfile* is the name of your custom Dockerfile.

For the Helm chart to take the new custom image for installation, set these keys in your **override-values.yaml** file for the ECE Helm chart:

```
imageRepository: "imageRepo:imagePort"
container:
  image: "customECEImageName"
```

# 15

## Building Your Own Images

Learn how to build your own images of the Oracle Communications Billing and Revenue Management (BRM), Elastic Charging Engine (ECE), Pipeline Configuration Center, Pricing Design Center (PDC), Billing Care, and Business Operations Center applications.

The Podman build commands in this chapter reference Dockerfile and related scripts as is from the **oc-cn-docker-files-15.0.x.0.0.tgz** package. Ensure you use your own version of Dockerfile and related scripts before running the build command.

Topics in this document:

- [Building BRM Server Images](#)
- [Building BRM REST Services Manager Images](#)
- [Building PDC REST Services Manager Images](#)
- [Building PDC Images](#)
- [Building Pipeline Configuration Center Images](#)
- [Building Billing Care Images](#)
- [Building Business Operations Center Images](#)

Sample Dockerfiles included in the BRM cloud native deployment package (**oc-cn-docker-files-15.0.x.0.0.tgz**) are examples that depict how default images are built for BRM. If you want to build your own images, refer to the sample Dockerfiles shipped with the product as a reference. Create your own Dockerfiles and then build your images.

### **Caution**

The Dockerfiles and related scripts are provided for reference only. You can refer to them to build or extend your own images. Support is restricted to core product issues only and no support will be provided for custom Dockerfiles and scripts.

## Building BRM Server Images

To build images for BRM Server, your staging area (\$PIN\_HOME) must be available from where the images are built. After you unpack **oc-cn-docker-files-15.0.x.0.0.tgz**, the BRM Server directory structure will be **oc-cn-docker-files/ocbrm**.

### **Note**

If you are using Podman to build your images, pass the **--format docker** flag with the **podman build** command.

Building your own BRM Server images involves these high-level steps:

1. You build the BRM Server base image. See "[Building Your BRM Server Base Image](#)".
2. You build images for each BRM Server component. See "[Building Images of BRM Server Components](#)".
3. You build the Web Services Manager image. See "[Building Web Services Manager Images](#)".
4. You build the BRM REST Services Manager image. See "[Building BRM REST Services Manager Images](#)".
5. You containerize the Email Data Manager. See "[Containerization of Email Data Manager](#)".
6. You containerize the roaming pipeline. See "[Containerization of Roaming Pipeline](#)".
7. You build and deploy Vertex Manager. See "[Building and Deploying Vertex Manager](#)".

## Building Your BRM Server Base Image

To make your directory structure ready for building base images:

1. Edit the **\$PIN\_HOME/bin/orapki** binary to replace the staging Java path with `{JAVA_HOME}`.
2. Create the **\$PIN\_HOME/installer** directory.
3. If you're behind a proxy server, set the **\$PROXY** variable:
 

```
export PROXY=ProxyHost:Port
```
4. Download the Java binary and then copy it to **\$PIN\_HOME**. See "BRM Software Compatibility" for the latest supported version of Java.
5. Download the Perl binary and then copy it to **\$PIN\_HOME**. See "BRM Software Compatibility" for the latest supported version of Perl.
6. For your database client:
  - a. Copy **oracle\_client\_response\_file.rsp** (64 bit), **downloadOracleClient.sh**, and **waitForOracleClientInst.sh** from **oc-cn-docker-files/ocbrm/base\_images** to **\$PIN\_HOME**.
  - b. Modify these parameters in the **downloadOracleClient.sh** file:
    - **ORACLE\_CLIENT\_ZIP**: Enter the binary name.
    - **REPOSITORY\_URL**: Enter the location to fetch the database client binary.
  - c. If the **db\_client** binary is already downloaded, copy the binary to the **\$PIN\_HOME/installer** directory.

After preparing your directory structure, build your BRM Server base image:

- For database client 12CR2 (64 Bit) + Java + Perl, enter this command:
 

```
podman build --format docker --build-arg PROXY=$PROXY --tag db_client_and_java_perl:15.0.x.0.0 --file DockerFileLocation/Dockerfile_db_client_and_java_perl .
```
- For database client 12CR2 (64 Bit) + Java, enter this command:
 

```
podman build --format docker --build-arg PROXY=$PROXY --tag db_client_and_java:15.0.x.0.0 --file DockerFileLocation/Dockerfile_db_client_and_java .
```
- For Java, enter this command:

```
podman build --format docker --build-arg PROXY=$PROXY --tag java:15.0.x.0.0 --file
DockerFileLocation/Dockerfile_java .
```

- For Java + Perl, enter this command:

```
podman build --format docker --build-arg PROXY=$PROXY --tag java_perl:15.0.x.0.0 --
file DockerFileLocation/Dockerfile_java_perl .
```

### Note

If the existing database is used with custom build images, do this:

- Override the **ocbrm.use\_oracle\_brm\_images** key in the Helm chart with a value of **false**.
- Set the **ocbrm.existing\_rootkey\_wallet** key to **true**.
- Copy your client wallet files to the **oc-cn-helm-chart/existing\_wallet** directory.

## Building Images of BRM Server Components

The **oc-cn-docker-files-15.0.x.0.0.tgz** package includes references to all of the Dockerfiles and scripts needed to build images of BRM Server components (except for **oraclelinux:8**).

To build an image of a BRM Server component:

1. Copy these scripts from the **oc-cn-docker-files/ocbrm** directory to your staging area at **\$PIN\_HOME**:
  - **entrypoint.sh**
  - **createWallet.sh**
  - **cm/preStopHook.sh\_cm**
  - **cm/postStartHook.sh**
  - **cm/updatePassword.sh**
  - **eai\_js/preStopHook.sh\_eai**
2. Do one of these:
  - For the batch pipeline, roaming pipeline, and real-time pipeline, copy **entrypoint.sh** and **createWallet.sh** to **\$PIN\_HOME/..**, and copy **\$PIN\_HOME/./setup/BRMActions.jar** to the **\$PIN\_HOME/jars** directory for building the images.
  - For all other components, copy the **\$PIN\_HOME/./setup/BRMActions.jar** file to **\$PIN\_HOME**.
3. Set these environment variables:
  - **\$PIN\_HOME**: Set this to your staging area.
  - **\$PERL\_HOME**: Set this to the path of Perl. See "BRM Software Compatibility" for the latest supported version of Perl.
  - **\$JAVA\_HOME**: Set this to the Java path. See "BRM Software Compatibility" for the latest supported version of Java.
4. Build the image for your BRM component.

For example, to build a CM image, you'd enter this:

```
podman build --format docker --tag cm:15.0.x.0.0 --build-arg
STAGE_PIN_HOME=$PIN_HOME --build-arg STAGE_JAVA_HOME=$JAVA_HOME --build-arg
STAGE_PERL_HOME=$PERL_HOME --file DockerfileLocation/Dockerfile .
```

To build a roaming pipeline image, you'd enter this:

```
podman build --format docker --tag roam_pipeline:$BRM_VERSION --build-arg
STAGE_PERL_HOME=StagePerlPath .
```

where *StagePerlPath* is the path to the Perl files in your staging area at \$PIN\_HOME.

To build a dm-oracle image, you'd enter this:

```
podman build --format docker --force-rm=true --no-cache=true --tag
dm_oracle:15.0.x.0.0 --file DockerfileLocation/Dockerfile .
```

where *DockerfileLocation* is the path to the Dockerfiles for your BRM component.

### Note

Build batch and realtime pipeline images from the **\$PIN\_HOME/..** directories.

## Building Web Services Manager Images

To containerize images for Web Services Manager, your staging area (\$PIN\_HOME) must be available from where the Docker images are built.

You can create one of these Web Services Manager containers:

- [Building and Deploying Web Services Manager for Apache Tomcat Image](#)
- [Building and Deploying Web Services Manager for WebLogic Server Image](#)

## Building and Deploying Web Services Manager for Apache Tomcat Image

The Web Services Manager Dockerfile is based on the official Apache Tomcat image. The sample Web Services Manager Dockerfile includes both the XML element-based and XML string-based SOAP Web Services implementation. Use this Dockerfile to build an image that can call any standard BRM opcode that is exposed as a SOAP Web service.

The Web Services Manager **Infranet.properties** configuration is available as a Kubernetes ConfigMap. To expose a custom opcode as a Web service, place your customized WAR filepath in the Dockerfile. When multiple pod replicas are configured, each pod runs its own copy of Apache Tomcat. By default, Web Services Manager is exposed as a Kubernetes NodePort service running on port 30080.

Containerizing the Web Services Manager for Tomcat image involves these high-level steps:

1. [Building the Web Services Manager Tomcat Image](#)
2. [Deploying the Web Services Manager Tomcat Image in Kubernetes](#)

### Building the Web Services Manager Tomcat Image

To build the Web Services Manager for Apache Tomcat image:

1. Download the JAX-WS reference implementation JARs from *JAX-WS Java API for XML Web Services* (<https://javaee.github.io/metro-jax-ws/>).
2. Copy the **jaxws-ri-2.3.x.zip** file to your staging area at \$PIN\_HOME.

3. Unzip the **jaxws-ri-2.3.x.zip** file.
4. Download Apache Tomcat 9 from the Apache Tomcat website:  
<https://tomcat.apache.org/download-90.cgi>  
See "Additional BRM Software Requirements" in *BRM Compatibility Matrix* for information about compatible versions of Apache Tomcat.
5. Copy **apache-tomcat-9.x.tar.gz** to your staging area at \$PIN\_HOME.
6. Copy these files from the **oc-cn-docker-files** directory to your staging area at \$PIN\_HOME.
  - **wsm\_entrypoint.sh**
  - **Dockerfile**
  - **context.xml**
  - **BRMActions.jar**
7. Update Tomcat in the Dockerfile to the latest version.
8. Build the Web Services Manager image by entering this command:  

```
podman build --format docker --tag brm_wsm:$BRM_VERSION .
```

### Deploying the Web Services Manager Tomcat Image in Kubernetes

To deploy the Web Services Manager for Tomcat image in Kubernetes:

1. Configure your Web services by updating the **configmap\_infranet\_properties\_wsm.yaml** file.
2. In the **override-values.yaml** file for **oc-cn-helm-chart**, set the following values:
  - **ocbrm.wsm.deployment.tomcat.isEnabled**: Set this to **true**.
  - **ocbrm.wsm.deployment.tomcat.walletPassword**: Set this to the Base64-encoded wallet password for the Web Services Manager image.
  - **ocbrm.wsm.deployment.tomcat.basicAuth**: Optionally, set this to **true** to enable BASIC authentication.
3. Optionally, for BASIC authentication, configure users in the **wsm\_config/tomcat-users.xml** file for **oc-cn-helm-chart**:
  - a. Open **tomcat-users.xml** in a text editor.
  - b. Locate the following lines and specify the login details of the user:
 

```
<role rolename="role"/>
<user username="username" password="password" roles="role"/>
```

where:

    - *role* is the role with permissions to access Web services, for example, **brmws**.
    - *username* is the user name for accessing Web services.
    - *password* is the password for accessing Web services.
  - c. Save and close the file.

See "[User File Format](#)" under *MemoryRealm* in the Apache Tomcat documentation for more information about the format of **tomcat-users.xml**.
4. Deploy the BRM Helm chart:

```
helm install ReleaseName oc-cn-helm-chart --namespace NameSpace --values
OverrideValuesFile
```

where:

- *ReleaseName* is the release name, which is used to track this installation instance.
- *NameSpace* is the namespace in which to create BRM Kubernetes objects.
- *OverrideValuesFile* is the path to the YAML file that overrides the default configurations in the BRM helm chart's **values.yaml** file.

## Building and Deploying Web Services Manager for WebLogic Server Image

To deploy and use Web Services Manager on WebLogic Server, you should be familiar with:

- Oracle WebLogic Server 12.2.1.3. See the Oracle WebLogic Server 12.2.1.3 documentation (<https://docs.oracle.com/middleware/12213/wls/index.html>).
- Oracle WebLogic Kubernetes Operator. See the WebLogic Kubernetes Operator documentation (<https://oracle.github.io/weblogic-kubernetes-operator/>).

The image for deploying BRM Web Services Manager on Oracle WebLogic Server 12.2.1.3 uses the domain in image approach. The image includes a WebLogic domain named **brmdomain**. When you build the image, the BRM SOAP Web Services application WAR files get deployed in this domain.

Containerizing the Web Services Manager for WebLogic Server image involves these high-level steps:

1. [Building the Web Services Manager WebLogic Image](#)
2. [Deploying the Web Services Manager WebLogic Image in Kubernetes](#)
3. [Updating the BRM Web Services Manager Configuration](#)
4. [Restarting the WebLogic Server Pods](#)
5. [Scaling Your WebLogic Managed Server](#)

### Building the Web Services Manager WebLogic Image

The BRM Web Services Manager on WebLogic Server image uses two images that run two containers inside each WebLogic Server pod.

To build the **brm\_wsm\_wls15.0.x.0.0** image:

1. Copy the contents of the **oc-cn-docker-files/ocbrm/brm\_soap\_wsm/weblogic/dockerfiles** directory to your staging area at \$PIN\_HOME.
2. Customize the WebLogic domain-related properties by editing the **dockerfiles/properties/docker-build/domain.properties** file. For example:

```
DOMAIN_NAME=brmdomain
ADMIN_PORT=7111
ADMIN_NAME=admin-server
ADMIN_HOST=wlsadmin
MANAGED_SERVER_PORT=8111
MANAGED_SERVER_NAME_BASE=managed-server
CONFIGURED_MANAGED_SERVER_COUNT=3
CLUSTER_NAME=cluster-1
DEBUG_PORT=8453
DB_PORT=1527
DEBUG_FLAG=true
PRODUCTION_MODE_ENABLED=true
```

```
CLUSTER_TYPE=DYNAMIC
JAVA_OPTIONS=-Dweblogic.StdoutDebugEnabled=false
T3_CHANNEL_PORT=30012
T3_PUBLIC_ADDRESS=kubernetes
IMAGE_TAG=brm_wsm_wls:$BRM_VERSION
```

3. Set the WebLogic domain user name and password by editing the **dockerfiles/properties/docker-build/domain\_security.properties** file. For example:

```
username=UserName
password=Password
```

#### Note

It is strongly recommended that you set a new user name and password when building the image.

For details about securing the **domain\_security.properties** file, see <https://github.com/oracle/docker-images/tree/master/OracleWebLogic/samples/12213-domain-home-in-image>.

4. Build the **brm\_wsm\_wls:15.0.x.0.0** image by running the **build.sh** script.

The script creates an image based on the custom tag defined in **dockerfiles/properties/docker-build/domain.properties**. By default, it creates the **brm\_wsm\_wls:15.0.x.0.0** image and then deploys the **BRMWebServices.war** and **infarnetwebsvc.war** files.

#### Note

If you don't want to deploy either **BRMWebServices.war** or **infarnetwebsvc.war**, modify the **dockerfiles/container-scripts/app-deploy.py** script.

5. Build the **brm\_wsm\_wl\_init:15.0.x.0.0** image by running this command:

```
podman build --format docker --tag brm_wsm_wl_init:15.0.x.0.0 --file
Dockerfile_init_wsm .
```

This image runs an init container, which populates the Oracle wallet that is used by Web Services Manager to connect to the CM.

## Deploying the Web Services Manager WebLogic Image in Kubernetes

You deploy the WebLogic Operator Helm chart so that Web Services Manager can work in a Kubernetes environment.

To deploy the Web Services Manager for WebLogic Server image in Kubernetes:

1. Clone the Oracle WebLogic Kubernetes Operator Git project:

```
git clone https://github.com/oracle/weblogic-kubernetes-operator
```

2. Modify these keys in the **override-values.yaml** file for **oc-cn-helm-chart**:

#### Note

Ensure that you set the **wsm.deployment.weblogic.enabled** key to **true**.

```
wsm:
  deployment:
    weblogic:
      enabled:true
      imageName:brm_wsm_wls
      initImageName:brm_wsm_wl_init
      imageTag:$BRM_VERSION
      username:d2VibG9naWM=
      password:password
      replicaCount:1
      adminServerNodePort:30611
      log_enabled:false
      minPoolSize:1
      maxPoolSize:8
      poolTimeout:30000
```

3. If the WebLogic user name and password was updated when building the **brm\_wsm\_wls:15.0.x.0.0** image, also update the base64-encoded WebLogic user name and password in these keys:

```
.Values.ocbrm.wsm.deployment.weblogic.username
.Values.ocbrm.wsm.deployment.weblogic.password
```

4. Add the BRM WebLogic Server namespace in the **kubernetes/charts/weblogic-operator/values.yaml** file:

```
domainNamespaces:
  - "default"
  - "NameSpace"
```

5. Deploy the WebLogic Operator Helm chart:

```
helm install weblogic-operator kubernetes/charts/weblogic-operator --namespace
WebOperatorNameSpace --values WebOperatorOverrideValuesFile --wait
```

where:

- *WebOperatorNameSpace* is the namespace in which to create WebLogic Operator Kubernetes objects.
- *WebOperatorOverrideValuesFile* is the path to a YAML file that overrides the default configurations in the WebLogic Operator Helm chart's **values.yaml** file.

6. Deploy the BRM helm chart:

```
helm install ReleaseName oc-cn-helm-chart --namespace NameSpace --values
OverrideValuesFile
```

where:

- *ReleaseName* is the release name, which is used to track this installation instance.
- *NameSpace* is the namespace in which **oc-cn-helm-chart** will be installed.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the BRM Helm chart's **values.yaml** file.

## Updating the BRM Web Services Manager Configuration

Update the basic configurations for BRM Web Services Manager by editing the Kubernetes ConfigMap (**configmap\_infranet\_properties\_wsm\_wl.yaml**). After updating the configuration, restart your WebLogic Server pods.

## Restarting the WebLogic Server Pods

To restart your WebLogic Server pods:

1. Stop the WebLogic Server pods by doing this:
  - a. In the `domain_brm_wsm.yaml` file, set the `serverStartPolicy` key to **NEVER**.
  - b. Update your Helm release.

```
helm upgrade ReleaseName oc-cn-helm-chart --namespace NameSpace --values
OverrideValuesFile
```

where `NameSpace` is the namespace in which `oc-cn-helm-chart` will be installed.

2. Start the WebLogic Server pods by doing this:
  - a. In the `domain_brm_wsm.yaml` file, set the `serverStartPolicy` key to **IF\_NEEDED**.
  - b. Update your Helm release:

```
helm upgrade ReleaseName oc-cn-helm-chart --namespace NameSpace --values
OverrideValuesFile
```

### Scaling Your WebLogic Managed Server

The default configuration starts one WebLogic Managed Server pod. To modify the configuration to start up to three pods, do this:

1. In the `oc-cn-helm-chart/values.yaml` file, set the `.Values.ocbrm.wsm.deployment.weblogic.replicaCount` key to **1**, **2**, or **3** WebLogic Managed Server pods.
2. Update your Helm release:

```
helm upgrade ReleaseName oc-cn-helm-chart --namespace NameSpace --values
OverrideValuesFile
```

You set the maximum number of managed servers in the BRM Web Services Manager image by modifying the `CONFIGURED_MANAGED_SERVER_COUNT` property in the `dockerfiles/properties/docker-build/domain.properties` file.

## Containerization of Email Data Manager

The Email Data Manager (DM) enables you to send customer notifications and invoices to your customers through email automatically. The Email DM uses the Sendmail client to forward emails to Postfix, which is the SMTP server. In-turn, Postfix sends the emails to your customers.

The Email DM will have the Sendmail client, and the Kubernetes host will have Postfix running. You must install and configure Postfix on your Kubernetes host.

To configure your cm pod to point to the Email DM, add this key to the `oc-cn-helm-chart/values.yaml` file:

```
ocbrm.dm_email.deployment.smtp: EmailHostName
```

where `EmailHostName` is the hostname of the server on which the Email DM is deployed. For example: `em389.us.example.com`.

To configure the Kubernetes host or SMTP server to accept data from the Email DM, do this:

1. Log in as the root user to the Kubernetes host.
2. Add the IP address for the Kubernetes host to the `/etc/postfix/main.cf` file:

```
inet_interfaces=localhost, HostIPAddress
```

For example, if the Kubernetes host is `10.242.155.149`.

```
inet_interfaces=localhost, 10.242.155.149
```

3. Retrieve the container network configuration by running this command on the Kubernetes host:

```
/sbin/ifconfig cni0 | grep netmask | awk '{print$2"\n"$4}'
```

The output will be similar to this:

```
10.244.0.1 ← The Kubernetes host IP, which is in the container network.
255.255.255.0
```

4. Edit the **mynetworks** field in the **/etc/postfix/main.cf** file to include the Kubernetes network in the list of trusted SMTP clients. If the Kubernetes host IP and Email DM container IP are in different networks, add both networks to the **mynetworks** field:

```
mynetworks = TrustedNetworks
```

where *TrustedNetworks* is the IP addresses for the SMTP clients that are allowed to relay mail through Postfix.

For example:

```
mynetworks = 168.100.189.0/28, 127.0.0.0/8, 10.244.0.0/24
```

5. Do one of these:
  - If Postfix is already running in the host, run this command:

```
systemctl restart postfix.service
```

- If Postfix isn't running in the host, run this command:

```
systemctl start postfix.service
```

### Note

In case of a multi-node environment, you can configure Postfix on the primary node (or any one node).

## Containerization of Roaming Pipeline

Roaming allows a wireless network operator to provide services to mobile customers from another wireless network. For example, when a mobile customer makes a phone call from outside the home network, roaming allows the customer to access the same wireless services that he has with his home network provider through a visited wireless network operator.

You feed the input files for the roaming pipeline through a Kubernetes PersistentVolumeClaim (PVC). The EDR output files will be available in a PVC for consumption of the rel-daemon pod. When building the roaming pipeline image, pass the Perl path in these files as part of **build-arg**.

To containerize the roaming pipeline, update the **configmap\_infranet\_properties\_rel\_daemon.yaml** file to specify how to load your rated CDR output files. For example:

```
batch.random.events TEL, ROAM
ROAM.max.at.highload.time 4
ROAM.max.at.lowload.time 2
ROAM.file.location /oms/ifw/data/roamout
```

```
ROAM.file.pattern test*.out
ROAM.file.type STANDARD
```

**Note**

The input file to the splitter pipeline must start with **Roam\_**.

## Building and Deploying Vertex Manager

To deploy Vertex Manager (dm-vertex), you layer the dm-vertex image with the libraries for Vertex Communications Tax Q Series (CTQ) or Vertex Sales Tax Q Series (STQ). For the list of supported library versions, see "Additional BRM Software Requirements" in *BRM Compatibility Matrix*.

### Deploying with Vertex Communications Tax Q Series

You deploy Vertex Manager with Vertex CTQ by doing the following:

1. Building the new Vertex Manager image by layering it with Vertex CTQ libraries.
  - a. Copy the entire Vertex CTQ installation directory to the **\$PIN\_HOME** directory, where **\$PIN\_HOME** is set to the path of your staging area.
  - b. Update the paths in the **64bit/bin/ctqcfg.xml**, **64bit/cfg/ctqcfg.xml**, and other Vertex CTQ files present in the Vertex CTQ installation directory. For example:

```
<configuration name="CTQ Test">
<fileControl>
  <updatePath>/oms/vertex/64bit/dat</updatePath>
  <archivePath>/oms/vertex/64bit/dat</archivePath>
  <callFilePath>/oms/vertex/64bit/dat</callFilePath>
  <reportPath>/oms/vertex/64bit/rpt</reportPath>
  <logPath>/oms/vertex/64bit/log</logPath>
</fileControl>
```

- c. In your copied Vertex CTQ installation directory, update the **64bit/bin/odbc/odbc.ini** file. For example:

**Note**

Set the **Driver** and **TNSNamesFile** entries to the file system path inside the pod.

```
[CtqTestOracle]
Description=Vertex, Inc. 8.0 Oracle Wire Protocol
Driver=/oms/vertex/64bit/bin/odbc/lib/VXor827.so
...
HostName=DBhostname
LogonID=DBuser
PortNumber=1521
Password=DBpassword
ServerName=//IPaddress:1521/DBalias
SID=DBalias
TNSNamesFile=/oms/ora_k8/tnsnames.ora
```

where:

- *DBhostname* is the host name of the machine on which the Vertex tax calculation database is installed.
- *DBuser* is the Vertex database schema user name.
- *DBpassword* is the password for the Vertex database schema user.
- *IPaddress* is the IP address of the machine on which the Vertex tax calculation database is installed.
- *DBalias* is the Vertex database alias name, which is defined in your **tnsnames.ora** file.

d. Layer the default images provided by Oracle.

For example, to layer dm-vertex with Vertex CTQ, you could add these sample commands to its Dockerfile. In this example, **\$PIN\_HOME** is set to **/oms** inside the pod.

```
FROM dm_vertex:15.0.x.0.0

USER root
RUN mkdir -p /oms/vertex/64bit/cfg
RUN chown -R omsuser:root /oms/vertex/64bit/cfg
COPY ./Vertex_CTQ_30206/ /oms/vertex
COPY Vertex_CTQ_30206/64bit/lib/libctq.so /oms/lib/
COPY Vertex_CTQ_30206/64bit/bin/odbc/lib/libodbc.so /oms/lib/libodbc.so

RUN chown -R omsuser:root /oms/vertex
RUN chown -R omsuser:root /oms/lib/libctq.so
RUN chown -R omsuser:root /oms/lib/libodbc.so
USER omsuser
```

e. Build your new Vertex Manager image. For example:

```
podman build --format docker --tag dm_vertex_ctq:15.0.x.0.0 --file
Dockerfile_vertex_ctq .
```

2. Enabling and configuring Vertex Manager in your BRM cloud native deployment.

a. Set these environment variables in your **oc-cn-helm-chart/templates/dm\_vertex.yaml** file:

```
- name: LD_LIBRARY_PATH
  value: "/oms/vertex/64bit/bin/odbc:/oms/lib:/oms/sys/dm_vertex:/oms/vertex/64bit/lib"
- name: CTQ_CFG_HOME
  value: "/oms/vertex/64bit/bin"
- name: ODBCINI
  value: "/oms/vertex/64bit/bin/odbc/odbc.ini"
```

b. Uncomment these entries in your **oc-cn-helm-chart/templates/configmap\_pin\_conf\_dm\_vertex.yaml** file:

```
- dm_vertex commtax_sm_obj ${DM_VERTEX_CTQ_SM}
- dm_vertex commtax_config_name ${DM_VERTEX_CTQ_CFG_NAME}
- dm_vertex commtax_config_path ${DM_VERTEX_CTQ_CFG_PATH}
```

c. Update these key in your **override-values.yaml** file for **oc-cn-helm-chart**:

```
dm_vertex:
  isEnabled: true
  deployment:
    replicaCount: 1
    imageName: dm_vertex_ctq
```

```

imageTag: 15.0.x.0.0
quantum_db_password: password
ctqCfg: /oms/vertex/64bit/cfg
ctqCfgName: CTQ Test
ctqSmObj: ./dm_vertex_ctq30206.so

```

- d. Run the **helm upgrade** command to update your BRM Helm release:

```

helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --
namespace BrmNameSpace

```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

## Deploying with Vertex Sales Tax Q Series

You deploy Vertex Manager with Vertex STQ by doing the following:

1. Copying the required libraries from the Vertex STQ installation directory to your **\$PIN\_HOME/req\_libs** directory.
2. Layer the default images provided by Oracle. For example, to layer dm-vertex with Vertex STQ, you could add these sample commands to its Dockerfile:

```

FROM dm_vertex:15.0.x.0.0

USER root
COPY ["req_libs/libvst*.so", "req_libs/libqutil*.so", "req_libs/libloc*.so",
"/oms/lib/"]
RUN chown omsuser:root -R /oms/lib/ /lib64
USER omsuser

```

3. Build your new Vertex Manager image. For example:

```

podman build --format docker --tag dm_vertex_stq:15.0.x.0.0 --file
Dockerfile_vertex_stq .

```

4. Update these key in your **override-values.yaml** file for **oc-cn-helm-chart**:

```

dm_vertex:
  isEnabled: true
  deployment:
    replicaCount: 1
    imageName: dm_vertex
    imageTag: 15.0.x.0.0
    quantum_db_password: password

```

5. Update these entries in your **oc-cn-helm-chart/templates/configmap\_env\_dm\_vertex.yaml** file:

```

SERVICE_FQDN: dm-vertex
QUANTUM_DB_SOURCE: quantum
QUANTUM_DB_SERVER: qsul22a
QUANTUM_DB_USER: quantum

```

6. Update these entries in your **oc-cn-helm-chart/templates/configmap\_odbc\_ini\_dm\_vertex.yaml** file:

```

data:
  odbc.ini: |
    [ODBC Data Sources]
    Server = Oracle Server v12.2
    [Server]
    Description = Oracle Server v12.2
    Driver = /usr/lib/oracle/19.20/client64/lib/libsqora.so.19.1
    Servername = PINDB
    UserID = DBuser
    Password = DBpassword
    Port = 1521
    Trace = yes
    TraceFile = /oms_logs/odbc.log
    Database = //DBhostname:DBport

```

where:

- *Server* is the name of the server on which the Vertex database is installed.
- *DBuser* is the Vertex database schema user name.
- *DBpassword* is the password for the Vertex database schema user.
- *DBhostname* is the host name of the machine on which the Vertex tax calculation database is installed.
- *DBport* is the port number of the Vertex tax calculation database.

7. Set these entries in your **oc-cn-helm-chart/templates/configmap\_pin\_conf\_dm\_vertex.yaml** file:

```

- dm_vertex quantum_sm_obj ./dm_vertex_stq100.so
- dm_vertex quantumdb_source ${QUANTUM_DB_SOURCE}
- dm_vertex quantumdb_server ${QUANTUM_DB_SERVER}
- dm_vertex quantumdb_user ${QUANTUM_DB_USER}

```

8. Run the **helm upgrade** command to update the BRM Helm release:

```

helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace BrmNameSpace

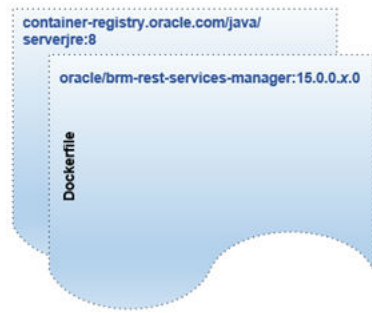
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.
- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **values.yaml** file for **oc-cn-helm-chart**.

## Building BRM REST Services Manager Images

In a production deployment, containers for BRM REST Services Manager will run in their own pods on a Kubernetes node. [Figure 15-1](#) shows how to stack images for BRM REST Services Manager.

**Figure 15-1 Image Stack for BRM REST Services Manager**

In this figure:

- **container-registry.oracle.com/java/serverjre:8**: The base image on which BRM REST Services Manager will be deployed. The official image is available at <https://container-registry.oracle.com/>.
- **oracle/brm-rest-services-manager:15.0.x.0.0**: The sample Dockerfile and related scripts used for creating the BRM REST Services Manager image (**oracle/brm-rest-services-manager:15.0.x.0.0**).

The **oc-cn-docker-files/ocrsm/brm\_rest\_services\_manager** directory in the **oc-cn-docker-files-15.0.x.0.0.tgz** package contains a Dockerfile, container scripts, and an API JAR file.

You can load or build the BRM REST Services Manager image in the following ways:

- The **oc-cn-brm-rest-services-manager-15.0.x.0.0.tar** image is included in the package. Apply the image in your machine by running this command:

```
podman load < oc-cn-brm-rest-services-manager-15.0.x.0.0.tar
```

- If the image needs customization, modify the Dockerfile and then deploy it using this command:

```
podman build --format docker --tag oracle/brm-rest-services-manager:15.0.x.0.0 .
```

## Building PDC REST Services Manager Images

In a production deployment, containers for PDC REST Services Manager will run in their own pods on a Kubernetes node. You create PDC REST Services Manager images by stacking these Dockerfiles in the following order:

1. **container-registry.oracle.com/java/serverjre:8**: The base image on which PDC REST Services Manager will be deployed. The official image is available at <https://container-registry.oracle.com/>.
2. **oracle/pdcrsm:15.0.x.0.0**: The sample Dockerfile and related scripts used for creating the PDC REST Services Manager image (**oracle/pdcrsm:15.0.x.0.0**).

To build PDC REST Services Manager images:

1. Copy the Dockerfile and the **oc-cn-pdc-rsm-jars-15.0.x.0.0** file into the current working directory.
2. Run the following commands:

```
tar xvf oc-cn-pdc-rsm-jars-15.0.x.0.0.tar
podman build --format docker --tag oracle/pdcrsm:15.0.x.0.0 .
```

## Building PDC Images

To build the PDC image:

1. (Release 15.0.0 only) Layer the brm-apps and realtimepipe images by doing the following:

- a. Download the brm-apps and realtimepipe images from the repository by entering this command:

```
podman pull RepoHost:RepoPort/ImageName
```

where:

- *RepoHost* is the IP address or host name of the repository.
- *RepoPort* is the port number for the repository.
- *ImageName* is either **brm\_apps:15.0.0.0** or **realtimepipe:15.0.0.0**.

- b. Tag the images by entering these commands:

```
podman tag RepoHost:RepoPort/brm_apps:15.0.0.0 brm_apps:15.0.0.0
podman tag RepoHost:RepoPort/realtimepipe:15.0.0.0 realtimepipe:15.0.0.0
```

2. Download **PricingDesignCenter-15.0.x.0.0.zip** to the *ParentFolder/Docker\_files/PDCImage/other-files* directory.
3. Pull the Java Image from the Oracle Container Registry (<https://container-registry.oracle.com>). This image is regularly updated with the latest security fixes. You can pull this image to your local system, where you will build other images, with the name **container-registry.oracle.com/java/serverjre:JavaVersion**.

where *JavaVersion* is the Oracle Java version number. See "Additional BRM Software Requirements" in *BRM Compatibility Matrix* for supported versions.

4. Set the following environment variables:

- HTTP\_PROXY: Set this to the host name or IP address of your proxy server
- JAVA\_VERSION: Set this to **container-registry.oracle.com/java/serverjre:JavaVersion**
- BRM\_VERSION: Set this to **15.0.x.0.0**

5. Build your Oracle PDC BRM integration image by entering this command from the *ParentFolder/Docker\_files/PDCImage* directory:

- For release 15.0.0 only, run this command:

```
podman build --format docker --force-rm=true --no-cache=true --build-arg
DB_VERSION=DBRelease --build-arg HTTP_PROXY=$HTTP_PROXY --build-arg
JAVA_VERSION=$JAVA_VERSION --build-arg BRM_VERSION=$BRM_VERSION --
tag $IMAGE_NAME --file Dockerfile .
```

where *DBRelease* is the Oracle database version number.

- For release 15.0.1 or later, run this command:

```
podman build --format docker --force-rm=true --no-cache=true --build-arg
HTTP_PROXY=$HTTP_PROXY --build-arg JAVA_VERSION=$JAVA_VERSION --tag $IMAGE_NAME
--file Dockerfile .
```

6. (Optional) To use custom fields in your PDC RUM expressions:

- a. Create a **custom\_flds.h** file that contains your custom fields.

For information about the syntax to use in a header file, view the *BRM\_home/include/pin\_flds.h* file in the brm-sdk pod.

- b. Parse the **custom\_flds.h** file and generate a **custom\_flds.bat** file.
- c. Layer the *BRM\_home/lib/custom\_flds.bat* file in the following images: brm-apps, pdc, and cm.

For example, to layer the file in the pdc image:

```
FROM pdc:15.0.x.0.0

USER root
COPY custom_flds.bat ${PIN_HOME}/lib/custom_flds.bat
RUN chown=oracle:root ${PIN_HOME}/lib/custom_flds.bat
```

- d. In your brm-apps-2 ConfigMap (**configmap\_pin\_conf\_brm\_apps\_2.yaml**), add the following entry under the **load\_config.conf** section:

```
load_config.conf: |
# Making custom fields entry
- - ops_fields_extension_file ${PIN_HOME}/lib/custom_flds.h
```

- e. In your CM ConfigMap (**configmap\_pin\_conf\_cm.yaml**), add the following entry under the **pin.conf** section:

```
data:
  pin.conf:
    # Making custom fields entry
    - - ops_fields_extension_file ${PIN_HOME}/lib/custom_flds.h
```

- f. In your testnap ConfigMap (**configmap\_pin\_conf\_testnap**), add the following entry under the **pin.conf** section:

```
data:
  pin.conf: |
    # Making custom fields entry
    - - ops_fields_extension_file ${PIN_HOME}/lib/custom_flds.h
```

- g. In your **override-values.yaml** file for **oc-cn-helm-chart**, update the **imageTag** keys to point to the new cm, brm-apps and pdc images.

### Note

Skip this step if your **override-values.yaml** file does not already contain **imageTag** keys for the cm, brm-apps, and pdc images.

- h. Run the **helm upgrade** command to update your Helm release:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --
namespace BrmNameSpace
```

## Building Pipeline Configuration Center Images

The Pipeline Configuration Center image extends the Fusion Middleware Infrastructure image by packaging its own installer **PipelineConfigurationCenter\_15.0.x.0.0\_generic.jar** file along with scripts and configurations.

To build your own image of Pipeline Configuration Center, you must have these base images ready. The **oc-cn-docker-files-15.0.x.0.0.tgz** package includes references to all Dockerfiles

and scripts that are needed to build images of Pipeline Configuration Center. You can refer to them when building a Pipeline Configuration Center image in your own environment.

## Pulling the Fusion Middleware Infrastructure Image

The Fusion Middleware Infrastructure Image is available on the Oracle Container Registry (<https://container-registry.oracle.com/>). This image is regularly updated with the latest security fixes. You can pull this image to your local system, where you will build other images, with the name **container-registry.oracle.com/middleware/fmw-infrastructure\_cpu:12.2.1.4-jdk8-ol7**.

## Building the Pipeline Configuration Center Image

To build the Pipeline Configuration Center image, do this:

1. Go to the **oc-cn-docker-files/ocpcc/pcc** directory.
2. Download the Oracle Communications Pipeline Configuration Center installation JAR file.
3. Copy **PipelineConfigurationCenter\_15.0.x.0.0\_generic.jar** to the current working directory (**oc-cn-docker-files/ocpcc/pcc**).
4. Build the Pipeline Configuration Center image by entering this command:

```
podman build --format docker --tag oracle/pcc:15.0.x.0.0 .
```

## Building Billing Care Images

The Billing Care image extends the Linux image by packaging the application archive along with scripts and configurations.

To build your own image of Billing Care, you need the Linux and JRE images, available on the Oracle Container Registry (<https://container-registry.oracle.com/>). These images are regularly updated with the latest security fixes. You can pull these images to your local system, where you will build other images, with the names:

- `container-registry.oracle.com/os/oraclelinux:8`
- `container-registry.oracle.com/java/serverjre:8-oraclelinux8`

The **oc-cn-docker-files-15.0.x.0.0.tgz** package includes references to all Dockerfiles and scripts that are needed to build images of Billing Care. You can refer to them when building a Billing Care image in your own environment.

## Building the Billing Care Image

To build the Billing Care image, do this:

1. Go to the **oc-cn-docker-files/ocbc/billing\_care** directory.
2. Download the Oracle Communications Billing Care installation JAR file.
3. Copy **BillingCare\_generic.jar** to the current working directory (**oc-cn-docker-files/ocbc/billing\_care**).
4. Build the Billing Care image by entering this command:

```
podman build --format docker --tag oracle/billingcare:15.0.x.0.0 .
```

## Building the Billing Care REST API Image

To build the Billing Care REST API image:

1. Go to the **oc-cn-docker-files/ocbc/bcws** directory.
2. Download the Oracle Communications Billing Care REST API installation JAR file.
3. Copy **BillingCare\_generic.jar** to the current working directory (**oc-cn-docker-files/ocbc/bcws**).
4. Build the Billing Care REST API image by entering this command:

```
podman build --format docker --tag oracle/bcws:15.0.x.0.0 .
```

## Building Business Operations Center Images

The Business Operations Center image extends the Linux image by packaging the application archive along with scripts and configurations.

To build your own image of Business Operations Center, you need the Linux and JRE images, available on the Oracle Container Registry (<https://container-registry.oracle.com>). These images are regularly updated with the latest security fixes. You can pull these images to your local system, where you will build other images, with the names:

- container-registry.oracle.com/os/oraclelinux:8
- container-registry.oracle.com/java/serverjre:8-oraclelinux8

The **oc-cn-docker-files-15.0.x.0.0.tgz** package includes references to all of the Dockerfiles and scripts needed to build images of Business Operations Center. You can refer to them when building a Business Operations Center image in your own environment.

To build the Business Operations Center image, do this:

1. Go to the **oc-cn-docker-files/ocboc/boc** directory.
2. Download the Oracle Communications Business Operations Center installation JAR file.
3. Copy **BusinessOperationsCenter\_generic.jar** to the current working directory (**oc-cn-docker-files/ocboc/boc**).
4. Build the Business Operations Center image by entering this command:

```
podman build --format docker --tag oracle/boc:15.0.x.0.0 .
```

# Part V

## Upgrading BRM Cloud Native

This part provides information about upgrading your Oracle Communications Billing and Revenue Management (BRM) cloud native environment to the latest patch set or interim patch release. It contains the following chapters:

- [Upgrading Your BRM Cloud Native Environment](#)
- [Performing Zero-Downtime Upgrades](#)
- [Performing Zero-Downtime Upgrades of Disaster Recovery Cloud Native Systems](#)
- [Migrating from On-Premise BRM to BRM Cloud Native](#)

# Upgrading Your BRM Cloud Native Environment

Learn how to upgrade your Oracle Communications Billing and Revenue Management (BRM) cloud native environment to the latest release.

Topics in this document:

- [Tasks for the BRM Cloud Native Upgrade](#)

In this document, the BRM release currently running on your production system is called the *old* release. The release you are upgrading to is called the *new* release. For example, if you upgrade from BRM 12.0.0.x.0 to BRM 15.0.x.0.0, 12.0.0.x.0 is the old release and 15.0.x.0.0 is the new one.

## Tasks for the BRM Cloud Native Upgrade

This section provides a list of tasks required to upgrade your BRM cloud native deployment to the latest release, patch set, or interim patch release. All patch sets and interim patches are cumulative, so they include the fixes from previous patch sets and interim patches. You can perform a direct upgrade from one patch set to another. For example, you can perform a direct upgrade of BRM cloud native from 12.0.0.x.0 to 15.0.x.0.0.

To upgrade your BRM cloud native deployment, complete these tasks in the specified order:

1. If you are upgrading from 12.0.0.2.0 to 12.0.0.3.0 or later, migrate your BRM cloud native Helm charts to the v3.x format by using the **helm2to3** utility. The Helm charts in BRM cloud native 12.0.0.2.0 use Helm v2.x, and later patch set releases use Helm v3.x. Helm v3.x doesn't readily understand the releases created by Helm v2.x.

For more information, see "[Migrating Helm v2 to v3](#)" in the Helm documentation. The documentation contains references to the migration plugin and to a blog with a comprehensive walk-through of steps using a sample chart.

2. Upgrade your BRM cloud native database schema. See "[Upgrading Your Database Schema](#)".
3. Upgrade your BRM cloud native services. See "[Upgrading Your BRM Cloud Native Services](#)".
4. Upgrade your Elastic Charging Engine (ECE) cloud native services. See "[Upgrading Your ECE Cloud Native Services](#)".
5. Upgrade your client application services in any order:
  - Upgrade your Pricing Design Center (PDC) service and database schema. See "[Upgrading Your PDC Cloud Native Services](#)".
  - Upgrade your BRM REST Services Manager service. See "[Upgrading BRM REST Services Manager](#)".
  - Upgrade your Business Operations Center service and database schema. See "[Upgrading Your Business Operations Center Cloud Native Services](#)".

- Upgrade your Pipeline Configuration Center (PCC) service. See "[Upgrading Your PCC Cloud Native Services](#)".
- Upgrade your Billing Care and Billing Care REST API services. See "[Upgrading Your Billing Care and Billing Care REST API Cloud Native Services](#)".

## Upgrading Your Database Schema

To upgrade your BRM cloud native database schema to the 15.0.x.0.0 release:

1. Download Oracle Communications Cloud Native Database Initializer Helm Chart 15.0.x.0.0 from the Oracle Software Delivery Cloud website. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and load the BRM 15.0.x.0.0 cloud native component images in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".

3. Extract the BRM database initializer Helm chart from the archive. For example:

```
tar xvzf oc-cn-init-db-helm-chart-15.0.x.0.0.tgz
```

where *x* is **0** for the 15.0.0.0.0 release, **1** for BRM 15.0.1.0.0 release, and so on. If you are extracting an interim patch, the file name will also have the interim patch number appended to it, such as **oc-cn-init-db-helm-chart-15.0.x.0.0-12345678.tgz**.

4. Copy the **/oms/wallet/client** files (**ewallet.p12** and **cwallet.sso**) from the BRM 12.0.0.x.0 version of the dm-oracle pod to the BRM 15.0.x.0.0 **oc-cn-init-db-helm-chart/existing\_wallet/** directory.

### Note

This step is only required if you are upgrading from 12.0.0.x.0 to 15.0.x.0.0. Skip this step if you are upgrading from 15.0.0.0.0 to 15.0.x.0.0.

5. Create an **override-init-db-15.yaml** file for the 15.0.x.0.0 version of **oc-cn-init-db-helm-chart**.
6. In the **override-init-db-15.yaml** file, do the following:
  - Set the **ocbrm.is\_upgrade** key to **true**.
  - Set the **existing\_rootkey\_wallet** key to **true**.
  - Set the other keys in [Table 6-1](#) as needed.

### Note

The BRM root password, wallet passwords, and database details should be the same as in your old release.

7. Validate the **oc-cn-init-db-helm-chart** 15.0.x.0.0 chart's content by using the **helm lint** command.

- For Helm 3.6.0 and later releases, enter this command from the **helmcharts** directory:

```
helm lint --strict oc-cn-init-db-helm-chart --values oc-cn-init-db-helm-chart/
values.yaml --values override-init-db-15.yaml
```

- For previous Helm releases, enter this command from the **helmcharts** directory:

```
helm lint --strict oc-cn-init-db-helm-chart
```

You'll see this if the command completes successfully:

```
1 chart(s) listed, no failures
```

8. Upgrade the database schema by entering this command from the **helmcharts** directory. Ensure that you run the Helm chart with a new release name and namespace.

```
helm install newRelease oc-cn-init-db-helm-chart --namespace newNameSpace --values
override-init-db-15.yaml
```

where:

- *newRelease* is the release name for your new 15.0.x.0.0 release. This release name must be different from that of your old release.
- *newNameSpace* is the namespace in which to create BRM Kubernetes objects for the new 15.0.x.0.0 release. This namespace must be different from that of your old release.

Your BRM cloud native database schema is upgraded to the new 15.0.x.0.0 release.

To determine if the upgrade was successful, enter the following:

```
kubectl --namespace newNameSpace get pods
```

If successful, you will see something similar to this:

NAME	READY	STATUS	RESTARTS	AGE
upgrade-wc6sx	0/1	Completed	0	22h

## Upgrading Your BRM Cloud Native Services

### Note

The steps for upgrading your BRM cloud native services are the same for old and new schemas.

When you upgrade your BRM cloud native services, it upgrades all BRM core services in your BRM cloud native environment.

To upgrade your BRM cloud native services to the 15.0.x.0.0 release:

1. Download and install Oracle Communications Cloud Native Helm Chart 15.0.x.0.0 from the Oracle Software Delivery Cloud website. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and load the BRM 15.0.x.0.0 cloud native images in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".

3. Extract the BRM 15.0.x.0.0 Helm chart from the archive:

```
tar xvzf oc-cn-helm-chart-15.0.x.0.0.tgz
```

where *x* is **0** for the 15.0.0.0.0 release, **1** for BRM 15.0.1.0.0 release, **2** for BRM 15.0.2.0.0 release, and so on. If you are extracting an interim patch, the file name will also have the interim patch number appended to it, such as **oc-cn-helm-chart-15.0.x.0.0-12345678.tgz**.

4. Create an **upgrade-brm-15.yaml** file for the 15.0.x.0.0 version of **oc-cn-helm-chart**.
5. In the **upgrade-brm-15.yaml** file, set the following keys:
  - Set the **ocbrm.is\_upgrade** key to **true**.
  - Set the **ocbrm.existing\_rootkey\_wallet** key to **false**.
  - Set the other keys in [Table 7-4](#) as needed.
6. Run the **helm upgrade** command for the 15.0.x.0.0 version of **oc-cn-helm-chart** using the same release name and namespace that you used for your old release:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --values  
upgrade-brm-15.yaml --namespace oldBrmNamespace
```

where:

- *oldBrmReleaseName* is the release name assigned to your old **oc-cn-helm-chart** installation.
  - *oldOverrideValues* is the file name and path to the **override-values.yaml** file for your old BRM installation.
  - *oldBrmNamespace* is the same namespace as for your old BRM deployment.
7. If you are upgrading a multischema system, do the following:
    - a. Add the following lines to the 15.0.x.0.0 version of the **oc-cn-helm-chart/brmapps\_scripts/loadme.sh** script:

```
#!/bin/sh

cd /oms/setup/scripts; perl pin_multidb.pl -i
cd /oms/setup/scripts; perl pin_multidb.pl -f
cd /oms/setup/scripts; perl pin_amt_install.pl
exit 0;
```

- b. Enable the brm-apps job by setting these keys in your **override-values.yaml** file for the 15.0.x.0.0 version of **oc-cn-helm-chart**:
    - **ocbrm.brm\_apps.job.isEnabled**: Set this to **true**
    - **ocbrm.brm\_apps.job.isMultiSchema**: Set this to **false**
  - c. Run the **helm upgrade** command for **oc-cn-helm-chart**:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --  
namespace BrmNameSpace
```

Your BRM core services have been upgraded to the latest 15.0.x.0.0 release.

**Note**

The first time you run `pin_virtual_time` after upgrading the BRM core services to the 15.0.x.0.0 release, it generates a new 64-bit `pin_virtual_time_file` utility. You must restart the CM after the 64-bit `pin_virtual_time_file` is created. To do so:

1. In the `override-values.yaml` for `oc-cn-helm-chart`, set the following keys:
  - `ocbrm.config_jobs.restart_count`: Increment the value by 1.
  - `ocbrm.config_jobs.run_apps`: Set this to `true`.
2. Run the `helm upgrade` command for `oc-cn-helm-chart` using the same release name and namespace that you used for your old release:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --
values upgrade-brm.yaml --namespace oldBrmNamespace
```

## Upgrading Your ECE Cloud Native Services

To upgrade your ECE cloud native services to the 15.0.x.0.0 release:

1. Download the ECE cloud native Helm chart. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Extract the ECE Helm chart from the archive into a separate staging area. For example:
 

```
tar xvzf oc-cn-ece-helm-chart-15.0.x.0.0.tgz StagingArea
```
3. In your staging area, create a diff between the new `values.yaml` file and your old `oc-cn-ece-helm-chart/values.yaml` file. For example, if you are upgrading from 12.0.0.5.0 to 15.0.x.0.0, do a diff between the 12.0.0.5.0 and 15.0.x.0.0 versions of the `values.yaml` file.
4. Using the diff, make a list of the keys that were added, changed, and removed in the new release.
5. Open your old release's `override-values.yaml` file for `oc-cn-ece-helm-chart`. This file contains all of the customizations that you made in previous releases.
6. Do the following:
  - Add and configure any new keys that you want to use.
  - Delete the keys that were removed.
  - If a key's default value changed, determine whether you want to modify the key's value.
7. Ensure that you have at least three `ecs` and `ecs1` pod replicas configured in the file:
  - `charging.ecs1.replicas`: Set this to a value of 3 or greater.
  - `charging.ecs.replicas`: Set this to a value of 3 or greater.
8. If your current ECE cloud native deployment does not support `cdrstore` tablespaces and you are upgrading to a version that requires them, do one of the following:
  - Continue without separate tablespaces for `cdrstore`. To do so, in your `override-values.yaml` file, set the `cdrstoretablespace` and `cdrstoreindexspace` keys to an empty value:

```
cdrstoretablespace: ""
cdrstoreindexspace: ""
```

- Use separate tablespaces for cdrstore. To do so, in **override-values.yaml** file, set the **cdrstoretablespace** and **cdrstoreindexspace** keys to the tablespace names.

Also, grant quota on the new tablespaces to the ECE schema user. Connect to your database as the **system** user using SQL\*Plus, and enter these commands:

```
SQL> ALTER USER EceSchemaUser quota unlimited on ECECDRTABLESPACE;
SQL> ALTER USER EceSchemaUser quota unlimited on ECECDRINDEXSPACE;
```

9. Save and close your **override-values.yaml** file to your staging area.
10. Delete all old ECE template files from your staging area's **oc-cn-ece-helm-chart/templates** directory.
11. Copy the new ECE template files to your staging area's **oc-cn-ece-helm-chart/templates** directory.
12. Upgrade your ECE cloud native services by running these commands:

```
cd StagingArea/oc-cn-ece-helm-chart/
sh upgradeECE_15.0.x.0.0.sh -o OverrideValuesFile -n BrmNameSpace -r EceReleaseName -s y
```

where:

- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **oc-cn-ece-helm-chart/values.yaml** file.
- *BrmNameSpace* is the namespace in which BRM Kubernetes objects reside for the BRM Helm chart.
- *EceReleaseName* is the release name for **oc-cn-ece-helm-chart** and is used to track this installation instance. It must be different from the one used for the BRM Helm chart.

## Upgrading ECE Cloud Native to the Latest Interim Patch

To upgrade your ECE cloud native services from 15.0 to the latest 15.0 interim patch:

1. Delete any existing ECE Kubernetes jobs:

```
kubectl --namespace BrmNameSpace get job
kubectl --namespace BrmNameSpace delete job JobName
```

where *BrmNameSpace* is the namespace in which BRM Kubernetes objects reside for the BRM Helm chart, and *JobName* is the name of the Kubernetes job.

2. Download the latest 15.0 interim patch release from the My Oracle Support website (<https://support.oracle.com>).
3. Extract the interim patch's ECE Helm chart from the archive into a separate staging area. For example:

```
tar xvzf oc-cn-ece-helm-chart-15.0.0.0.0.xyz.tgz StagingArea
```

where *xyz* is the interim patch number.

4. In your staging area, create a diff between the 15.0 interim patch **values.yaml** file and your old 15.0 **oc-cn-ece-helm-chart/values.yaml** file.
5. Using the diff, make a list of the keys that were added, changed, and removed in the 15.0 interim patch release.
6. Open your 15.0 **override-values.yaml** file for **oc-cn-ece-helm-chart**.
7. Do the following:

- Add and configure any new keys that you want to use.
  - Delete the keys that were removed.
  - If a key's default value changed, determine whether you want to override the key's value.
8. Ensure that you have at least three ecs and ecs1 pod replicas configured in the file:
    - **charging.ecs1.replicas**: Set this to a value of 3 or greater.
    - **charging.ecs.replicas**: Set this to a value of 3 or greater.
  9. Save and close your **override-values.yaml** file.
  10. Upgrade your ECE cloud native services to the latest 15.0 interim patch release by running these commands:

```
cd StagingArea/oc-cn-ece-helm-chart/
sh upgradeECE_15.0.0.0.sh --override OverrideValuesFile --namespace BrmNameSpace --release EceReleaseName
```

where:

- *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **oc-cn-ece-helm-chart/values.yaml** file.
- *EceReleaseName* is the release name for **oc-cn-ece-helm-chart** and is used to track this installation instance. It must be different from the one used for the BRM Helm chart.

## Upgrading Your PDC Cloud Native Services

### ① Note

- When you upgrade your PDC cloud native services, your PDC database is also upgraded.
- If you reuse an old PDC schema, you must choose the same rating engine as the old PDC installation. That is, if your old PDC installation uses Elastic Charging Engine (ECE) for usage rating, you cannot switch to the Real-time Rating and Batch Rating Engines. Likewise, you cannot switch to ECE if your old PDC installation uses the Real-time Rating and Batch Rating Engines for usage rating. If you attempt to switch rating engines, the PDC upgrade fails and generates an error message.
- You cannot migrate pricing data to PDC cloud native systems, because the PDC **MigrateBRMPricing** utility is not supported in cloud native environments.

To upgrade your PDC cloud native services and the PDC database to the latest 15.0.x.0.0 release:

1. Download and extract the BRM 15.0.x.0.0 cloud native Helm charts. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the PDC 15.0.x.0.0 component images into your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".

- From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Remove all PDC pods, services, and volume mounts from your PDC 12.0.0.x.0 release by doing the following:
    - a. In your **override-values.yaml** file for the 12.0.0.x.0 version of **oc-cn-helm-chart**, set the **ocpdc.isEnabled** key to **false**.
    - b. Run the **helm upgrade** command for the 12.0.0.x.0 release of **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values
oldBrmOverrideValuesFile --namespace oldBrmNameSpace
```

where:

- *oldBrmReleaseName* is the release name assigned to your 12.0.0.x.0 version of the **oc-cn-helm-chart** installation.
- *oldBrmOverrideValuesFile* is the file name and path of your 12.0.0.x.0 version of the **override-values.yaml** file for **oc-cn-helm-chart**.
- *oldBrmNameSpace* is the namespace for your 12.0.0.x.0 version of the BRM deployment.

**Note**

Ensure that all PDC pods, services, and volume mounts have been deleted.

4. If you are upgrading from PDC 12.0.0.8.0 or later to PDC 15.0.x.0.0, do the following:
  - a. In your **override-values.yaml** file for your *old* release of **oc-cn-op-job-helm-chart**, set the **ocpdc.isEnabled** key to **false**.
  - b. Run the **helm upgrade** command for your *old* release of **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
oldOpJobOverrideValuesFile --namespace oldBrmNameSpace
```

where:

- *oldOpJobReleaseName* is the release name assigned to your old release of the **oc-cn-op-job-helm-chart** installation.
- *oldOpJobOverrideValuesFile* is the file name and path of your old release of the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.
- *oldBrmNameSpace* is the namespace for your old BRM deployment.

Ensure all PDC jobs are deleted.

- c. Compare your old release's **oc-cn-op-job-helm-chart/values.yaml** file with the 15.0.x.0.0 version of that file.
  - Create a diff between your old release's **values.yaml** file and the 15.0.x.0.0 release's **values.yaml** file.
  - Using the diff, make a list of the keys that were added, changed, and removed in the 15.0.x.0.0 release.
  - Open *oldOpJobOverrideValuesFile*. This file contains all of the customizations that you made in previous releases.
  - Do the following:

- Add and configure any new keys that you want to use.
  - Delete the keys that were removed.
  - If a key's default value changed, determine whether you want to modify the key's value.
- Close and save the file as **new-op-job-override-values.yaml**.
5. Do one of the following:
    - If you are upgrading from 12.0 Patch Set 8 or later, open your **new-op-job-override-values.yaml** file.
    - If you are upgrading from 12.0 Patch Set 7 or earlier, create an **override-values.yaml** file for the 15.0.x.0.0 version of **oc-cn-op-job-helm-chart** named **new-op-job-override-values.yaml**.
  6. In **new-op-job-override-values.yaml**, set the following keys:
    - **ocpdc.isEnabled**: Set this to **true**.
    - **ocpdc.configEnv.upgrade**: Set this to **true**.
    - **ocpdc.configEnv.deployAndUpgradeSite2**: Set this to **false**.
  7. Run the **helm upgrade** command for the 15.0.x.0.0 version of **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values new-op-job-override-values.yaml --namespace oldBrmNameSpace
```

Wait for the PDC domain job to complete.
  8. Compare your 12.0.0.x.0 release's **oc-cn-helm-chart/values.yaml** file with the 15.0.x.0.0 version of the file.
    - Create a diff between your old release's **values.yaml** file and the 15.0.x.0.0 release's **values.yaml** file.
    - Using the diff, make a list of the keys that were added, changed, and removed in the 15.0.x.0.0 release.
    - Open your 12.0.0.x.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
    - Do the following:
      - Add and configure any new keys that you want to use.
      - Delete the keys that were removed.
      - If a key's default value changed, determine whether you want to modify the key's value.
    - Close and save the file as **new-brm-override-values.yaml**.
  9. In **new-brm-override-values.yaml**, set the following keys:
    - **ocbrm.pdc\_deployed**: Set this to **true**.
    - **ocpdc.configEnv.upgrade**: Set this to **true**.

This value must match the one set in your **override-values.yaml** file for **oc-cn-op-job-helm-chart** (in [Step 6](#)).
    - **ocpdc.isEnabled**: Set this to **true**.

10. Run the **helm upgrade** command for the 15.0.x.0.0 version of **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values new-brm-override-values.yaml --namespace oldBrmNameSpace
```

## Upgrading BRM REST Services Manager

To upgrade your BRM REST Services Manager cloud native services to the 15.0.x.0.0 release:

1. Download and extract the BRM cloud native Helm charts. See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the BRM REST Services Manager component image to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling WebLogic Images for PDC, Billing Care, Billing Care REST API, and Business Operations Center](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Disable the **brm-rest-services-manager** service in your BRM cloud native environment.
  - a. Create an **upgrade-brm-rsm.yaml** file. This file will be used by **oc-cn-helm-chart**.
  - b. In your **upgrade-brm-rsm.yaml** file, set the **ocrsm.rsm.isEnabled** key to **false**.
  - c. Stop the **brm-rest-services-manager** pod by running the Helm upgrade command for the **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --values upgrade-brm-rsm.yaml --namespace oldBrmNameSpace
```

where:

- *oldBrmReleaseName* is the BRM release name for your old release.
  - *oldOverrideValues* is the file name and path to the **override-values.yaml** file for your old **brm-rest-services-manager** installation.
  - *oldBrmNameSpace* is the BRM namespace for your old release.
4. Wait for the **brm-rest-services-manager** pod to stop.
  5. In your **upgrade-brm-rsm.yaml** file, set the these keys:
    - **ocrsm.rsm.isEnabled**: Set this to **true**.
    - **ocrsm.rsm.deployment.imageTag**: Set this to the new release number in the format 15.0.x.0.0 or 15.0.x.0.0-*nnnnnnnn* for interim patches.
  6. Copy the SSL Certificate for BRM REST Services Manager.
    - a. Create a directory named **rsm\_keystore** under the newly extracted **oc-cn-helm-chart/rsm** directory.
    - b. Copy the files created in the step "[Generating an SSL Certificate for BRM REST Services Manager](#)" to the newly created **oc-cn-helm-chart/rsm** directory.
    - c. Start your **brm-rest-services-manager** services by running the Helm upgrade command for **oc-cn-helm-chart**:

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values oldOverrideValues --values upgrade-brm-rsm.yaml --namespace oldBrmNameSpace
```

## Upgrading Your Business Operations Center Cloud Native Services

The instructions to upgrade your Business Operations Center services differ depending on the patch set you are upgrading to or from.

- To upgrade the Business Operations Center service from 12.0.0.7.0 or earlier to 15.0.x.0.0, follow the instructions in "[Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.7.0 or Earlier to 15.0.x.0.0](#)".
- To upgrade the Business Operations Center service from 12.0.0.8.0 to 15.0.x.0.0, follow the instructions in "[Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.8.0 to 15.0.x.0.0](#)".

### Note

When you upgrade your Business Operations Center cloud native service, you can also upgrade your Business Operations Center database schema.

## Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.7.0 or Earlier to 15.0.x.0.0

To upgrade your Business Operations Center cloud native service and database schema from 12.0.0.7.0 or earlier to 15.0.x.0.0:

1. Download and extract the 15.0.x.0.0 versions of the BRM cloud native Helm chart (**oc-cn-helm-chart**) and the cloud native operator job chart (**oc-cn-op-job-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the Business Operations Center component image (**boc**) to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Disable all Business Operations Center 12.0.0.x.0 services in your BRM cloud native environment.
  - a. Create an **upgrade-boc.yaml** file and then set the **ocboc.boc.isEnabled** key to **false**.  
The **upgrade-boc.yaml** file will be used with both **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.
  - b. Stop the WebLogic domain by running the **helm upgrade** command for **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --values upgrade-boc.yaml --namespace oldBrmNamespace
```

where:

- *oldBrmReleaseName* is the BRM release name for your old release.
- *oldOverrideValues* is the **override-values.yaml** file for your 12.0.0.x.0 release.
- *oldBrmNamespace* is the BRM namespace for your old release.

- c. Remove the WebLogic domain by running the **helm upgrade** command for **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
oldOverrideValues --values upgrade-boc.yaml --namespace oldBrmNameSpace
```

where *oldOpJobReleaseName* is the **oc-cn-op-job-helm-chart** release name for the old release.

4. Clean up the data in your Business Operations Center 12.0.0.x.0 persistent volumes (PVs).
- a. Clean up the domain home from the PV for Business Operations Center 12.0.0.x.0:

```
rm -rf Domain_home/domains/domainUID
```

where:

- *Domain\_home* is the location specified in the **ocboc.boc.wop.domainVolHostPath** key.
- *domainUID* is the domain name specified in the **ocboc.boc.wop.domainUID** key. The default is **boc-domain**.

See [Table 9-1](#) for more information.

- b. Clean up the application home from the PV for Business Operations Center:

```
rm -rf Application_home/BOC
```

where *Application\_home* is the path specified in the **ocboc.boc.wop.appVolHostPath** key.

5. Compare your old 12.0.0.x.0 versions of the **oc-cn-helm-chart/values.yaml** and **oc-cn-op-job-helm-chart/values.yaml** files with the 15.0.x.0.0 versions of those files.
- Create a diff between the 15.0.x.0.0 **values.yaml** file and your old 12.0.0.x.0 **values.yaml** file.
  - Using the diff, make a list of the keys that were added, changed, and removed in the new release.
  - Open your old 12.0.0.x.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
  - Do the following:
    - Add and configure any new keys that you want to use.
    - Delete the keys that were removed.
    - If a key's default value changed, determine whether you want to modify the key's value.
  - Close and save the files as *updatedoldOverrideValues*.
6. Deploy Business Operations Center 15.0.x.0.0 with the latest changes and upgrade the Business Operations Center database schema.
- a. In your **upgrade-boc.yaml** file, set these Business Operations Center keys:
- **ocboc.boc.isEnabled**: Set this to **true**.
  - **ocboc.boc.deployment.imageTag**: Set this to the new release number in the format 15.0.x.0.0 for patch sets and 15.0.x.0.0-*nnnnnnnn* for interim patches.
- b. Run the **helm upgrade** command for **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
updatedOldOverrideValues --values upgrade-boc.yaml --namespace oldBrmNamespace
```

Wait for the jobs to complete their tasks.

7. Start your Business Operations Center 15.0.x.0.0 services by running the **helm upgrade** command for **oc-cn-helm-chart**:

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values updatedOldOverrideValues --
values upgrade-boc.yaml --namespace oldBrmNamespace
```

## Upgrading Your Business Operations Center Cloud Native Service from 12.0.0.8.0 to 15.0.x.0.0

To upgrade your Business Operations Center cloud native service and database schema from 12.0.0.8.0 to 15.0.x.0.0:

1. Download and extract the 15.0.x.0.0 versions of the BRM cloud native Helm chart (**oc-cn-helm-chart**) and the cloud native operator job chart (**oc-cn-op-job-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the Business Operations Center 15.0.x.0.0 image (**boc**) to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Compare your old 12.0.0.8.0 versions of the **oc-cn-helm-chart/values.yaml** and **oc-cn-op-job-helm-chart/values.yaml** files with the 15.0.x.0.0 versions of those files.
  - Create a diff between the 15.0.x.0.0 **values.yaml** file and your old 12.0.0.8.0 **values.yaml** file.
  - Using the diff, make a list of the keys that were added, changed, and removed in the new release.
  - Open your old 12.0.0.8.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
  - Do the following:
    - Add and configure any new keys that you want to use.
    - Delete the keys that were removed.
    - If a key's default value changed, determine whether you want to modify the key's value.
  - Close and save the files as *updatedOldOverrideValues*.
4. Deploy Business Operations Center 15.0.x.0.0 with the latest changes and upgrade the Business Operations Center database schema.
  - a. In your **upgrade-boc.yaml** file, set the **ocboc.boc.deployment.app.imageTag** key to the new release number in the format 15.0.x.0.0 or 15.0.x.0.0-*nnnnnnnn* for interim patches.  
The **upgrade-boc.yaml** file will be used with both **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.
  - b. Run the **helm upgrade** command for **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
updatedOldOverrideValues --values upgrade-boc.yaml --namespace oldBrmNameSpace
```

Wait for the jobs to complete their tasks.

5. Start your Business Operations Center 15.0.x.0.0 services by running the **helm upgrade** command for **oc-cn-helm-chart**:

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values updatedOldOverrideValues --
values upgrade-boc.yaml --namespace oldBrmNameSpace
```

## Upgrading Your PCC Cloud Native Services

The instructions to upgrade your PCC cloud native services differ depending on the release you are upgrading from.

- To upgrade your PCC cloud native services from 12.0.0.x.0 or earlier to 15.0.0.0.0, follow the instructions in "[Upgrading Your PCC Cloud Native Services from 12.0.0.x.0 or Earlier to 15.0.0.0.0](#)".
- To upgrade your PCC cloud native services from 15.0.0.0.0 to 15.0.x.0.0, follow the instructions in "[Upgrading Your PCC Cloud Native Services from 15.0.0.0.0 to 15.0.x.0.0](#)".

## Upgrading Your PCC Cloud Native Services from 12.0.0.x.0 or Earlier to 15.0.0.0.0

To upgrade your PCC cloud native services from 12.0.0.x.0 or earlier to the 15.0.0.0.0 release:

1. Download and extract the BRM 15.0.0.0.0 cloud native Helm chart (**oc-cn-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the PCC 15.0.0.0.0 component image (**oracle/pcc:15.0.0.0.0**) to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Disable the PCC 12.0.0.x.0 cloud native services in your BRM cloud native environment:
  - a. Create an **upgrade-pcc.yaml** file for **oc-cn-helm-chart**.
  - b. In your **upgrade-pcc.yaml** file, set the **ocpcc.pcc.isEnabled** key to **false**.
  - c. Stop the running pcc pod by running the **helm upgrade** command for the 12.0.0.x.0 version of **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --
values upgrade-pcc.yaml --namespace oldBrmNameSpace
```

where:

- *oldBrmReleaseName* is the BRM release name for your 12.0.0.x.0 release.
  - *oldOverrideValues* is the file name and path to the **override-values.yaml** file for your 12.0.0.x.0 BRM installation.
  - *oldBrmNameSpace* is the BRM namespace for your 12.0.0.x.0 release.
4. Wait for the pcc pod to stop.
  5. Copy the SSL Certificate for PCC:
    - a. Create a **keystore\_pcc** directory under **oc-cn-helm-chart/pcc**.

- b. Copy the default PKCS12 certificate and KeyStore files to the **oc-cn-helm-chart/pcc** directory.  
During deployment, Helm uses the KeyStore files to create a Secret, which will be mounted as a volume inside the pcc pod.
  - c. If your KeyStore files have file names different from what is specified in the **values.yaml** file, update the **keyStoreType**, **keyStoreIdentityFileName**, and **keyStoreTrustFileName** keys in your **override-values.yaml** file.
6. Compare your 12.0.0.x.0 release's **oc-cn-helm-chart/values.yaml** file with the 15.0.0.0.0 version of the file.
    - Create a diff between your 12.0.0.x.0 release's **values.yaml** file and the 15.0.0.0.0 release's **values.yaml** file.
    - Using the diff, make a list of the keys that were added, changed, and removed in the 15.0.0.0.0 release.
    - Open your 12.0.0.x.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
    - Do the following:
      - Add and configure any new keys that you want to use.
      - Delete the keys that were removed.
      - If a key's default value changed, determine whether you want to modify the key's value.
    - Close and save the file as **new-brm-override-values.yaml**.
  7. In your **upgrade-pcc.yaml** file, set the following keys:
    - **ocpcc.pcc.isEnabled**: Set this to **true**.
    - **ocpcc.pcc.deployment.imageTag**: Set this to the new release number in the format 15.0.0.0.0 or 15.0.0.0.0-*nnnnnnnn* for interim patches.
  8. Start your PCC cloud native services by running the **helm upgrade** command for the 15.0.0.0.0 version of **oc-cn-helm-chart**:

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values new-brm-override-values.yaml --values upgrade-pcc.yaml --namespace oldBrmNamespace
```

## Upgrading Your PCC Cloud Native Services from 15.0.0.0.0 to 15.0.x.0.0

To upgrade your PCC cloud native services from 15.0.0.0.0 to 15.0.x.0.0:

1. Download and extract the BRM 15.0.x.0.0 cloud native Helm charts (**oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the PCC 15.0.x.0.0 component image (**pcc**) to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Disable all PCC 15.0.0.0.0 cloud native services in your BRM cloud native environment:
  - a. Create an **upgrade-pcc.yaml** file and then set the **ocpcc.pcc.isEnabled** key to **false**.

The **upgrade-pcc.yaml** file will be used with both **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.

- b. Stop the WebLogic domain by running the **helm upgrade** command for the 15.0.0.0.0 version of **oc-cn-helm-chart**:

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values
oldOverrideValues --values upgrade-pcc.yaml --namespace oldBrmNamespace
```

where:

- *oldBrmReleaseName* is the BRM release name for your 15.0.0.0.0 release.
  - *oldOverrideValues* is the **override-values.yaml** file for your 15.0.0.0.0 release.
  - *oldBrmNamespace* is the BRM namespace for your 15.0.0.0.0 release.
- c. Remove the WebLogic domain by running the **helm upgrade** command for the 15.0.0.0.0 version of **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
oldOverrideValues --values upgrade-pcc.yaml --namespace oldBrmNameSpace
```

where *oldOpJobReleaseName* is the **oc-cn-op-job-helm-chart** release name for the 15.0.0.0.0 release.

4. Clean up the data in your PCC 15.0.0.0.0 persistent volumes (PVs).
  - a. Clean up the domain home from the PV for PCC 15.0.0.0.0:

```
rm -rf Domain_home/domains/domainUID
```

where:

- *Domain\_home* is the location specified in the **ocpcc.pcc.wop.domainVolHostPath** key.
  - *domainUID* is the domain name specified in the **ocpcc.pcc.wop.domainUID** key. The default is **pcc-domain**.
- b. Clean up the application home from the PV for PCC 15.0.0.0.0:

```
rm -rf Application_home/PCC
```

where *Application\_home* is the path specified in the **ocpcc.pcc.wop.appVolHostPath** key.

5. Compare your 15.0.0.0.0 version of the **oc-cn-helm-chart/values.yaml** file with the 15.0.x.0.0 version of that file.
  - Create a diff between the 15.0.x.0.0 **values.yaml** file and the 15.0.0.0.0 **values.yaml** file.
  - Using the diff, make a list of the keys that were added, changed, and removed in the 15.0.x.0.0 release.
  - Open your 15.0.0.0.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
  - Do the following:

- Add and configure any new keys that you want to use.
  - Delete the keys that were removed.
  - If a key's default value changed, determine whether you want to modify the key's value.
  - Close and save the file as **new-brm-override-values.yaml**.
6. Configure the 15.0.x.0.0 version of **oc-cn-op-job-helm-chart**:
    - a. Create an **op-job-override-values.yaml** file.
    - b. In the file, set the PCC-related keys for **oc-cn-op-job-helm-chart**. See "[Adding Pipeline Configuration Center Keys for oc-cn-op-job-helm-chart](#)".
  7. In your **upgrade-pcc.yaml** file, set these PCC keys:
    - **ocpcc.pcc.isEnabled**: Set this to **true**.
    - **ocpcc.pcc.deployment.imageTag**: Set this to the new release number in the format 15.0.x.0.0 for patch sets and 15.0.x.0.0-*nnnnnnnn* for interim patches.
  8. Deploy PCC 15.0.x.0.0 with the latest changes by running the **helm upgrade** command for **oc-cn-op-job-helm-chart**:
 

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values op-job-override-values.yaml --values upgrade-pcc.yaml --namespace oldBrmNameSpace
```
  9. Start your PCC 15.0.x.0.0 cloud native services by running the **helm upgrade** command for **oc-cn-helm-chart**:
 

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values new-brm-override-values.yaml --values upgrade-pcc.yaml --namespace oldBrmNameSpace
```

## Upgrading Your Billing Care and Billing Care REST API Cloud Native Services

The instructions to use to upgrade your Billing Care and Billing Care REST API services are different, depending on the patch set you are upgrading from.

- To upgrade Billing Care and Billing Care REST API services from 12.0 Patch Set 7 or earlier, follow the instructions in "[Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.7.0 or Earlier to 15.0.x.0.0](#)".
- To upgrade the Billing Care and Billing Care REST API services from 12.0 Patch Set 8 or later, follow the instructions in "[Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.8.0 to 15.0.x.0.0](#)".

## Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.7.0 or Earlier to 15.0.x.0.0

To upgrade Billing Care and Billing Care REST API from 12.0.0.7.0 or earlier to 15.0.x.0.0:

1. Download and extract the 15.0.x.0.0 versions of the BRM cloud native Helm chart (**oc-cn-helm-chart**) and the cloud native operator job chart (**oc-cn-op-job-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the 15.0.x.0.0 versions of the Billing Care image (**billingcare**) and the Billing Care REST API image (**bcws**) to your repository in one of these ways:

- From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".
  - From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Disable all Billing Care and Billing Care REST API services in your BRM 12.0.0.x.0 cloud native environment.
    - a. Create an **upgrade-billing.yaml** file.  
This file will be used with both **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.
    - b. In your **upgrade-billing.yaml** file, set these keys:
      - **ocbc.bc.isEnabled**: Set this to **false**.
      - **ocbc.bcws.isEnabled**: Set this to **false**.
    - c. Stop the WebLogic domain by running the Helm upgrade command for **oc-cn-helm-chart**:
 

```
helm upgrade oldBrmReleaseName oc-cn-helm-chart --values oldOverrideValues --values upgrade-billing.yaml --namespace oldBrmNamespace
```

 where:
      - *oldBrmReleaseName* is the BRM release name for your old release.
      - *oldOverrideValues* is the file name and path to the **override-values.yaml** file for your old Billing Care installation.
      - *oldBrmNamespace* is the BRM namespace for your old release.
    - d. Remove the WebLogic domain by running the Helm upgrade command for **oc-cn-op-job-helm-chart**:
 

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values oldOverrideValues --values upgrade-billing.yaml --namespace oldBrmNameSpace
```

 where *oldOpJobReleaseName* is the **oc-cn-op-job-helm-chart** release name for the old release.
  4. Clean up the data in the 12.0.0.x.0 versions of the Billing Care and Billing Care REST API persistent volumes (PVs).
    - a. Clean up the domain home from the PV for Billing Care and Billing Care REST API:
 

```
rm -rf Domain_home/domains/domainUID
```

 where:
      - *Domain\_home* is the location specified in the **ocbc.bc.wop.domainVolHostPath** and **ocbc.bcws.wop.domainVolHostPath** keys.
      - *domainUID* is the domain name specified in the **ocbc.bc.wop.domainUID** and **ocbc.bcws.wop.domainUID** keys. The defaults are **billingcare-domain** and **bcws-domain**.
 See [Table 9-3](#) and [Table 9-5](#).
    - b. Clean up the application home from the PV for Billing Care and Billing Care REST API:
 

```
rm -rf Application_home/billingcare
```

 where *Application\_home* is the path specified in the **ocbc.bc.wop.appVolHostPath** and **ocbc.bcws.wop.appVolHostPath** keys.

5. Compare your old 12.0.0.x.0 versions of the **oc-cn-helm-chart/values.yaml** and **oc-cn-op-job-helm-chart/values.yaml** files with the 15.0.x.0.0 versions of those files.
  - Create a diff between the 15.0.x.0.0 **values.yaml** file and your old 12.0.0.x.0 **values.yaml** file.
  - Using the diff, make a list of the keys that were added, changed, and removed in the new release.
  - Open your old 12.0.0.x.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
  - Do the following:
    - Add and configure any new keys that you want to use.
    - Delete the keys that were removed.
    - If a key's default value changed, determine whether you want to modify the key's value.
  - Close and save the files as *updatedOldOverrideValues*.
6. In your **upgrade-billing.yaml** file, set these Billing Care and Billing Care REST API keys:
  - **ocbc.bc.isEnabled**: Set this to **true**.
  - **ocbc.bcws.isEnabled**: Set this to **true**.
  - **ocbc.bc.deployment.imageTag**: Set this to the new release number in the format 15.0.x.0.0 or 15.0.x.0.0-*nnnnnnnn* for interim patches.
  - **ocbc.bcws.deployment.imageTag**: Set this to the new release number in the format 15.0.x.0.0 or 15.0.x.0.0-*nnnnnnnn* for interim patches.
7. Deploy the 15.0.x.0.0 versions of the Billing Care and Billing Care REST API with the latest changes by running the Helm upgrade command for **oc-cn-op-job-helm-chart**:
 

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
updatedOldOverrideValues --values upgrade-billing.yaml --namespace oldBrmNameSpace
```

Wait for the jobs to complete their tasks.
8. Start your 15.0.x.0.0 versions of the Billing Care and Billing Care REST API services by running the Helm upgrade command for **oc-cn-helm-chart**:
 

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values updatedOldOverrideValues --
values upgrade-billing.yaml --namespace oldBrmNameSpace
```

## Upgrading Your Billing Care and Billing Care REST API Cloud Native Services from 12.0.0.8.0 to 15.0.x.0.0

To upgrade Billing Care and Billing Care REST API cloud native from the 12.0.0.8.0 release to 15.0.x.0.0:

1. Download and extract the 15.0.x.0.0 versions of the BRM cloud native Helm chart (**oc-cn-helm-chart**) and the cloud native operator job chart (**oc-cn-op-job-helm-chart**).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)".
2. Download and push the 15.0.x.0.0 versions of the Billing Care image (**billingcare**) and the Billing Care REST API image (**bcws**) to your repository in one of these ways:
  - From the Oracle Container Registry. To do so, see "[Pulling BRM Images from the Oracle Container Registry](#)".

- From the Oracle Software Delivery website. To do so, see "[Downloading BRM Images from Oracle Software Delivery Website](#)".
3. Compare your old 12.0.0.8.0 versions of the **oc-cn-helm-chart/values.yaml** and **oc-cn-op-job-helm-chart/values.yaml** files with the 15.0.x.0.0 versions of those files.
    - Create a diff between the 15.0.x.0.0 **values.yaml** file and your old 12.0.0.8.0 **values.yaml** file.
    - Using the diff, make a list of the keys that were added, changed, and removed in the new release.
    - Open your old 12.0.0.8.0 release's **override-values.yaml** file. This file contains all of the customizations that you made in previous releases.
    - Do the following:
      - Add and configure any new keys that you want to use.
      - Delete the keys that were removed.
      - If a key's default value changed, determine whether you want to modify the key's value.
    - Close and save the files as *updatedOldOverrideValues*.
  4. Create an **upgrade-billing.yaml** file and set these Billing Care and Billing Care REST API keys:
    - **ocbc.bc.deployment.app.imageTag**: Set this to the new release number in the format 15.0.x.0.0 for patch sets and 15.0.x.0.0-*nnnnnnnn* for interim patches.
    - **ocbc.bcws.deployment.app.imageTag**: Set this to the new release number in the format 15.0.x.0.0 for patch sets and 15.0.x.0.0-*nnnnnnnn* for interim patches.

The **upgrade-billing.yaml** file will be used with both **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.

5. Deploy the 15.0.x.0.0 versions of Billing Care and Billing Care REST API with the latest changes by running the Helm upgrade command for **oc-cn-op-job-helm-chart**:

```
helm upgrade oldOpJobReleaseName oc-cn-op-job-helm-chart --values
updatedOldOverrideValues --values upgrade-billing.yaml --namespace oldBrmNameSpace
```

Wait for the jobs to complete their tasks.

6. Start your 15.0.x.0.0 versions of the Billing Care and Billing Care REST API services by running the Helm upgrade command for **oc-cn-helm-chart**:

```
helm upgrade oldBRMReleaseName oc-cn-helm-chart --values updatedOldOverrideValues --
values upgrade-billing.yaml --namespace oldBrmNameSpace
```

# Performing Zero-Downtime Upgrades

Learn how to upgrade an Oracle Communications Billing and Revenue Management (BRM) cloud native deployment without having to take the environment offline. A zero-downtime upgrade allows your customers to continue using BRM's major services during the upgrade process.

Topics in this document:

- [Performing a Zero-Downtime Upgrade of BRM](#)
- [Performing a Zero Downtime Upgrade of PDC](#)

In this document, the BRM release running on your production system is called the *existing* release. The release you are upgrading to is called the *new* release. For example, if you upgrade from BRM 12.0 Patch Set 5 to BRM 15.0, 12.0 Patch Set 5 is the existing release and 15.0 is the new one.

## Performing a Zero-Downtime Upgrade of BRM

You can perform a zero-downtime upgrade of your BRM cloud native services and the BRM database schema from 12.0.0.x.0 to 15.0.x.0.0.

To perform a zero-downtime upgrade of BRM cloud native:

1. Download the BRM 15.0.x.0.0 cloud native package from the Oracle Software Delivery website (<https://support.oracle.com>) or the Oracle Support website (<https://support.oracle.com>).
2. Configure and deploy the BRM 12.0.0.x.0 **oc-cn-helm-chart** Helm chart on your cloud native environment by doing the following:
  - a. In your **override-values.yaml** file for the BRM 12.0.0.x.0 **oc-cn-helm-chart**, set the following keys:

```
ocbrm:
  refreshInterval: 10
  terminationGracePeriodSeconds: 120
  pcprReconnectDelayOnSocketError: 10
  pcpcConnectRetryDelayOnError: 10
```

- b. In your **oc-cn-helm-chart/templates/configmap\_pin\_conf\_cm.yaml** file, set the following key:

```
- cm pcm_connect_max_retries 10
```

- c. Ensure that at least two replica pods of cm, dm-oracle, dm-ifw-sync, and realtime-pipeline are up and running.
- d. (BRM 12.0.0.7.0 only) Apply BRM 12.0.0.7.0 Interim Patch 34939558 to your cloud native environment.
- e. (BRM 12.0.0.7.0 or later) In your BRM 12.0.0.x.0 **oc-cn-helm-chart/templates/configmap\_pin\_conf\_dm\_oracle.yaml** file, set the following key:

```
- dm dm_ignore fld_mismatch_err 1
```

- f. Run the **helm upgrade** command for the 12.0.0.x.0 **oc-cn-helm-chart**:

```
helm upgrade Brm_12_ReleaseName oc-cn-helm-chart --values OverrideValuesFile --
namespace Brm_12_NameSpace
```

where:

- *Brm\_12\_ReleaseName* is the release name assigned to your existing 12.0.0.x.0 **oc-cn-helm-chart** installation.
  - *OverrideValuesFile* is the file name and path to the file that overrides the **oc-cn-helm-chart/values.yaml** file.
  - *Brm\_12\_NameSpace* is the namespace for your existing 12.0.0.x.0 BRM deployment.
- g. Back up your existing BRM 12.0.0.x.0 Helm charts.
- h. Copy the 15.0.x.0.0 **oc-init-db-helm-chart** and **oc-cn-helm-chart** Helm charts to your BRM cloud native environment.
- i. (12.0.0.6.0 or earlier) Upgrade only the dm-oracle pod from 12.0.0.x.0 to 15.0.x.0.0 by doing the following:
- i. In your **override-values.yaml** file for the 15.0.x.0.0 **oc-cn-helm-chart**, set the **ocbrm.dm\_oracle.deployment.imageTag** key to **15.0.x.0.0**.
  - ii. Run the **helm upgrade** command for the 15.0.x.0.0 **oc-cn-helm-chart**:

```
helm upgrade Brm_12_ReleaseName oc-cn-helm-chart --values OverrideValuesFile
--namespace Brm_12_NameSpace
```

Verify that only the dm-oracle pod is running with the 15.0.x.0.0 image. The remaining pods run with 12.0.0.x.0 images.

3. Upgrade the BRM cloud native database schema to 15.0.x.0.0 by doing the following:

- a. Ensure that the **ConfigCacheRefreshInterval** business parameter in **bus\_params\_system.xml** is set to 0.

For information about how to set business parameters, see "Running Load Utilities through Configurator Jobs" in *BRM Cloud Native System Administrator's Guide*.

- b. In your **override-values.yaml** file for the 15.0.x.0.0 **oc-cn-init-db-helm-chart**, set the following keys:

```
ocbrm:
  is_upgrade: true
  existing_rootkey_wallet: true
```

- c. Copy the wallet files (**ewallet.p12** and **cwallet.sso**) from the **/oms/wallet/client/** directory of the 12.0.0.x.0 primary dm-oracle pod to the 15.0.x.0.0 **oc-cn-init-db-helm-chart/existing\_wallet** directory.
- d. Upgrade the cloud native database to 15.0.x.0.0 using the **helm install** command:

```
helm install NewInitDbReleaseName oc-cn-init-db-helm-chart --values
OverrideValuesFile --namespace NewInitDbNameSpace
```

where:

- *NewInitDbReleaseName* is the release name assigned to your new 15.0.x.0.0 **oc-cn-init-db-helm-chart** installation.
- *NewInitDbNameSpace* is the namespace for your new 15.0.x.0.0 **oc-cn-init-db-helm-chart** deployment.

- Verify that the database has upgraded successfully.
- e. (12.0.0.6.0 or earlier) Restart the 15.0.x.0.0 dm-oracle pod.
4. Upgrade the BRM cloud native server to 15.0.x.0.0 by doing the following:

### ① Note

To run billing during the BRM cloud native server upgrade for non-production systems, set the next billing date in the **oc-cn-helm-chart/templates/configmap\_env\_common.yaml** file to the following:

```
VIRTUAL_TIME_SETTING: "-m 2 billingDate"
```

For example, to set the next billing date to Feb 12, 2025, set *billingDate* to 021210002025.

- a. Ensure the **ConfigCacheRefreshInterval** business parameter in **bus\_params\_system.xml** is set to **0**.  
For information about how to set business parameters, see "Running Load Utilities through Configurator Jobs" in *BRM Cloud Native System Administrator's Guide*.
- b. In your **override-values.yaml** file for the 15.0.x.0.0 **oc-cn-helm-chart**, set the following keys:
 

```
ocbrm:
  refreshInterval: 10
  terminationGracePeriodSeconds: 120
  pcpReconnectDelayOnSocketError: 10
  pcpConnectRetryDelayOnError: 10
  virtual_time:
    enabled: true
    sync_pvt_time: 5
```
- c. In the 15.0.x.0.0 **oc-cn-helm-chart/templates/configmap\_pin\_conf\_cm.yaml** file, set the following:
 

```
- cm pcm_connect_max_retries 10
- cm pcm_em_proto_vers 0
```
- d. In the 15.0.x.0.0 **oc-cn-helm-chart/templates/configmap\_pin\_conf\_rtp\_pipeline.yaml** file, set the following:
 

```
- cm-em pcm_em_proto_vers 0
```
- e. In the 15.0.x.0.0 **oc-cn-helm-chart/templates/cm.yaml** file, set the following:
 

```
spec.template.spec.containers(- name: cm).livenessProbe.initialDelaySeconds: 60
```
- f. In the 15.0.x.0.0 **oc-cn-helm-chart/templates/configmap\_pin\_conf\_dm\_oracle.yaml** file, set the following:
 

```
- dm dm_ignore fld_mismatch_err 1
```
- g. Upgrade the BRM cloud native server to 15.0.x.0.0 using the **helm upgrade** command:
 

```
helm upgrade Brm_12_ReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace Brm_12_NameSpace
```

Verify that all 12.0.0.x.0 pods terminate and all 15.0.x.0.0 pods come up and run with 15.0.x.0.0 images.

5. Revert the configuration values in your **override-values.yaml** file for the 15.0.x.0.0 **oc-cn-helm-chart**:

- a. In non-production systems, set the **VIRTUAL\_TIME\_SETTING** parameter to the default value in your **oc-cn-helm-chart/templates/configmap\_env\_common.yaml** file:

```
VIRTUAL_TIME_SETTING: "-m 0"
```

- b. In your **oc-cn-helm-chart/templates/configmap\_pin\_conf\_rtp\_pipeline.yaml** file, remove or comment out the following entry:

```
- cm-em pcm_em_proto_vers 0
```

- c. In your **oc-cn-helm-chart/templates/configmap\_pin\_conf\_cm.yaml** file, remove or comment out the following entry:

```
- cm pcm_em_proto_vers 0
```

- d. In your **override-values.yaml** file for **oc-cn-helm-chart**, set the **ocbrm.refreshInterval** key to its original value.

- e. Set the **ConfigCacheRefreshInterval** business parameter in **bus\_params\_system.xml** to its original value.

For information about how to set business parameters, see "Running Load Utilities through Configurator Jobs" in *BRM Cloud Native System Administrator's Guide*.

- f. Run the **helm upgrade** command for the 15.0.x.0.0 **oc-cn-helm-chart**:

```
helm upgrade Brm_12_ReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace Brm_12_NameSpace
```

## Performing a Zero Downtime Upgrade of PDC

You can perform a zero downtime upgrade of your PDC cloud native services and the PDC database schema from the 12.0 or 12.0 Patch Set release to a 15.0.x release.

You do so using a two-namespace approach in which you create an instance of PDC cloud native in a standby namespace, redirect PDC traffic to services in the standby namespace, upgrade PDC cloud native to release 15.0.x in your original namespace, and then redirect PDC traffic back to your original namespace.

To upgrade PDC in zero downtime upgrade mode:

1. Create a temporary namespace, such as **BrmStandbyNameSpace**.
2. Clone your PDC *OverrideValuesFile* file to *StandbyOverrideValuesFile*.
3. In your *StandbyOverrideValuesFile* file for **oc-cn-op-job-helm-chart**, set the following keys:

```
ocpdc:
  configEnv:
    rcuPrefix: NewPrefix
    crossRefSchemaUserName: XrefSchema
    pdcSchemaUserName: PdcSchema
```

```
deployAndUpgradeSite2: true
upgrade: true
```

where:

- *NewPrefix* is the new prefix for the PDC domain RCU schema.
  - *XrefSchema* is the same XREF schema name used for deploying PDC in **BrmNameSpace**.
  - *PdcSchema* is the same PDC schema name used for deploying PDC in **BrmNameSpace**.
4. In your *StandbyOverrideValuesFile* file for **oc-cn-helm-chart**, set the following keys:

```
ocpdc:
  configEnv:
    upgrade: true
```

These settings will upgrade the PDC and XREF schema.

5. Copy the following templates from your **BrmNameSpace** Helm chart **template** directory to the **BrmStandbyNameSpace** Helm chart **template** directory:
- **secret\_env\_brm.yaml**
  - **configmap\_pin\_conf\_brm\_apps\_2.yaml**
  - **configmap\_loadifwconfig\_reg.yaml**
  - **configmap\_env\_common.yaml**
  - **virtual\_time\_pvc.yaml**
  - **configmap\_infranet\_properties\_brm\_apps.yaml**
  - **config\_jobs.yaml**
  - **storage\_class\_green.yaml**
  - **realtime\_pipeline\_common\_pvc.yaml**
  - **configmap\_tns\_admin.yaml**
  - **secret\_wallet\_db.yaml**
  - **\_helpers\_utils.tpl**
  - **configmap\_pdc\_aux\_engines.yaml**
  - **configmap\_log\_properties\_pdc.yaml**
  - **configmap\_env\_pdc.yaml**
  - **configmap\_env\_pdc\_rre.yaml**
  - **job\_ie\_pdc.yaml**
  - **domain\_pdc.yaml**
  - **\_pdchelpers.tpl**
  - **pdc\_domain\_monitoring\_role.yaml**
  - **pdc\_domain\_monitoring\_rbac.yaml**
  - **service\_monitor\_pdc\_domain.yaml**
  - **secret\_pdc.yaml**
  - **volume\_pdc\_brm.yaml**

- **deployment\_pdc\_rre.yaml**
  - **deployment\_pdc\_bre.yaml**
  - **deployment\_pdc\_syncpdc.yaml**
6. Create a configuration file named **cm-service-external-name.yaml** and add the following content:

```
apiVersion: v1
kind: Service
metadata:
  name: cm
  namespace: BrmStandbyNameSpace
spec:
  externalName: cm.BrmNameSpace.svc.cluster.local
  internalTrafficPolicy: Cluster
  ports:
    - port: 11960
      protocol: TCP
      targetPort: 11960
  sessionAffinity: None
  type: ExternalName
```

7. Apply the configuration file to a resource:
 

```
kubectl apply -f cm-service-external-name.yaml
```
8. Deploy the PDC 12.0 Patch Set 7 or 8 Helm charts in your standby namespace:
 

```
helm install OpJobStandbyReleaseName oc-cn-op-job-helm-chart --values
StandbyOverrideValuesFile --namespace BrmStandbyNameSpace

helm install BrmStandbyReleaseName oc-cn-helm-chart --values
StandbyOverrideValuesFile --namespace BrmStandbyNameSpace
```
9. Redirect PDC traffic to services in **BrmStandbyNameSpace**.
10. Upgrade your PDC cloud native services to release 15.0.x in **BrmNameSpace** while requests are temporarily routed to **BrmStandbyNameSpace**.
  - a. Set the following keys in your **override-values.yaml** file for **oc-cn-op-job-helm-chart**:

```
ocpdc:
  configEnv:
    deployAndUpgradeSite2: false
    upgrade: true
```

- b. Follow the instructions in "[Upgrading Your PDC Cloud Native Services](#)" to upgrade your original namespace to 15.0.x.
11. Redirect PDC traffic back to services in your original namespace (**BrmNameSpace**).

# Performing Zero-Downtime Upgrades of Disaster Recovery Cloud Native Systems

You can perform a zero-downtime upgrade of an Oracle Communications Billing and Revenue Management (BRM) cloud native deployment in an active-active disaster recovery system.

Topics in this document:

- [About the Zero-Downtime Upgrade of an Active-Active Disaster Recovery System](#)
- [Tasks for Upgrading a BRM Cloud Native Active-Active System](#)

## About the Zero-Downtime Upgrade of an Active-Active Disaster Recovery System

The steps for performing a zero-downtime cloud native upgrade of an active-active disaster recovery system assume that you are upgrading from the BRM 12.0 Patch Set 3 release to a BRM 15.0.x release. It also assumes that your system contains a secondary BRM and ECE instance, which can be on a disaster recovery site or the primary site.

### **Caution**

Be aware that during the upgrade process:

- Usage processing is done in a single instance only.
- PDC is not available.
- ECE cache federation is supported from the old to the new version but not from the new to the old version.
- If an Oracle Data Guard role reversal is required for the BRM database, a short (few minutes) downtime occurs for provisioning.

For information about the supported software versions in BRM 15.0.x, see "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix*.

## Tasks for Upgrading a BRM Cloud Native Active-Active System

To upgrade your BRM cloud native active-active disaster recovery system using the zero-downtime upgrade process:

1. Turn off Site 2. See "[Switching Off Site 2](#)".
2. From Site 2, uninstall the 12.0 Patch Set 3 version of BRM and ECE. See "[Uninstalling BRM and ECE from Site 2](#)".

3. On Site 2, upgrade the Kubernetes platform and all prerequisite software to the versions required by the 15.0.x release. See "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix* for the list of supported software versions.  
For information, see "[Upgrade a Cluster](#)" in the Kubernetes documentation.
4. On Site 2, upgrade the BRM database schema to the 15.0.x release. See "[Upgrading Your BRM Database Schema in Site 2](#)".
5. On Site 2, install the BRM 15.0.x cloud native software. See "[Installing BRM 15.0.x Cloud Native on Site 2](#)".
6. From Site 2, drop the ECE 12.0 Patch Set 3 persistence database schema. See "[Dropping the ECE Persistence Database Schema from Site 2](#)".
7. On Site 2, install ECE 15.0.x. See "[Installing ECE 15.0.x Cloud Native on Site 2](#)".
8. Transfer usage processing from Site 1 to Site 2. See "[Failing Over Site 1 to Site 2](#)".
9. From Site 1, uninstall the 12.0 Patch Set 3 version of BRM and ECE. See "[Uninstalling BRM and ECE from Site 1](#)".
10. On Site 1, upgrade the Kubernetes platform and all prerequisite software to the versions required by the 15.0.x release. See "BRM Cloud Native Deployment Software Compatibility" in *BRM Compatibility Matrix* for the list of supported software versions.  
For information, see "[Upgrade a Cluster](#)" in the Kubernetes documentation.
11. From Site 1, drop the ECE 12.0 Patch Set 3 persistence database schema. See "[Dropping the ECE Persistence Database Schema from Site 1](#)".
12. On Site 1, install ECE 15.0.x. See "[Installing ECE 15.0.x Cloud Native on Site 1](#)".
13. Restart the federation process between Site 1 and Site 2. See "[Federating ECE Cache Data Between Site 1 and Site 2](#)".
14. If required, move the provisioning flow back to Site 1 and do an Oracle Data Guard role reversal.

**Note**

The Data Guard role reversal may cause a few minutes of downtime in the provisioning flow, but it does not impact the usage flow.

## Switching Off Site 2

During the first phase of the upgrade process, Site 1 processes usage requests while you upgrade the software in Site 2.

To switch off Site 2:

1. Stop all usage requests to Site 2 and direct all usage requests to Site 1.
2. Stop the connection from BRM on Site 1 to ECE on Site 2. To do so:
  - On ECE Site 2, stop the EM Gateway.
  - On BRM Site 1, remove any connection to the EM Gateway on Site 2.
3. On ECE Site 1, mark Site 2 as failed.

This stops ECE from forwarding rating requests to Site 2 for subscribers with Site 2 as their preferred site. Usage requests are now rated on Site 1.

4. Disable the federation from ECE Site 1 to Site 2.

**Note**

Keep the Site 2 to Site 1 federation active to drain any remaining data from Site 2.

5. Check that the federation backlog from Site 2 to Site 1 has been cleared.

All traffic to Site 2 is stopped. Ensure all data from Site 2 has been synchronized with Site 1. The Coherence Federation Metrics should show as IDLE instead of YIELDING.

6. On Site 2, check that all rated events in the ECE cache have been extracted by running the following with the **query.sh** utility:

```
./query.sh
Coherence Command Line Tool
CohQ1> select value() from AggregateObjectUsage
```

If successful, the query displays zero entries. See "query" in *ECE Implementing Charging* for information about the utility's syntax and parameters.

7. On Site 2, use SQL\*Plus to check that all Site 2 rated events have been extracted from the persistence database and are present in BRM:

```
sqlplus pin@databaseName
Enter password: password

SQL> select count(*) from ratedevent_site2Name
```

where *databaseName* is the service name or database alias of the BRM database, and *password* is the password for the **pin** user.

**Note**

The Site 2 persistence database might contain some Site 1 events. After the Site 2 to Site 1 federation process is stopped, these events are not extracted or purged. However, they will be processed on Site 1. You can ignore these events because they get purged when you re-create the Site 2 persistence database later.

8. Stop the federation process from Site 2 to Site 1.

All data is now synchronized with Site 1. Site 2 is isolated and ready for the upgrade.

## Uninstalling BRM and ECE from Site 2

To uninstall BRM and ECE from Site 2:

1. Uninstall BRM 12.0 Patch Set 3 cloud native from Site 2:

```
helm uninstall BrmReleaseName --namespace BrmNameSpace
```

2. Uninstall ECE 12.0 Patch Set 3 cloud native from Site 2:

```
helm uninstall EceReleaseName --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *EceReleaseName* is the release name for **oc-cn-ece-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.

These commands delete all the resources associated with the chart's last release and the release history.

## Upgrading Your BRM Database Schema in Site 2

To upgrade your BRM database schema in Site 2 to release BRM 15.0.x:

1. Download and extract the BRM 15.0.x cloud native database initializer Helm chart (**oc-cn-op-job-helm-chart**) from Oracle Software Delivery Cloud (<https://edelivery.oracle.com>). See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)" for more information.
2. Copy your BRM 12.0 Patch Set 3 wallet files (**ewallet.p12** and **cwallet.sso**) from the dm-oracle pod's **/oms/wallet/client** directory to the BRM 15.0.x Helm chart's **oc-cn-init-db-helm-chart/existing\_wallet/** directory.
3. Create an **override-init-db-15.yaml** file for the 15.0.x version of **oc-cn-init-db-helm-chart**.
4. In your **override-init-db-15.yaml** file, set the following keys:
  - **ocbrm.is\_upgrade**: Set this to **true**.
  - **ocbrm.existing\_rootkey\_wallet**: Set this to **true**.
  - **ocbrm.wallet.client**: Set this to the password for the client wallet. This value must match that of your 12.0 Patch Set 3 release.
  - **ocbrm.wallet.server**: Set this to the password for the server wallet. This value must match that of your 12.0 Patch Set 3 release.
  - **ocbrm.wallet.root**: Set this to the password for the root wallet. This value must match that of your 12.0 Patch Set 3 release.
5. Enter this command from the **helmcharts** directory to upgrade the database schema. Ensure you run the 15.0.x **oc-cn-init-db-helm-chart** Helm chart with a new release name and namespace.

```
helm install newInitDbReleaseName oc-cn-init-db-helm-chart --namespace  
newInitDbNameSpace --values oldOverrideValues --values override-init-  
db-15.yaml
```

where:

- *newInitDbReleaseName* is the new release name for the 15.0.x version of **oc-cn-init-db-helm-chart**.

- *newInitDbNameSpace* is the new namespace for the 15.0.x version of **oc-cn-init-db-helm-chart**.
- *oldOverrideValues* is the name and path to your 12.0 Patch Set 3 **override-values.yaml** file for **oc-cn-init-db-helm-chart**.

## Installing BRM 15.0.x Cloud Native on Site 2

To configure and deploy BRM 15.0.x on your cloud native system on Site 2:

1. Download and extract the BRM 15.0.x cloud native (**oc-cn-helm-chart**) and BRM 15.0.x Operator Job Helm Chart (**oc-cn-op-job-helm-chart**) from Oracle Software Delivery Cloud (<https://edelivery.oracle.com>).

See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)" for more information.

2. To reuse your 12.0 Patch Set 3 SSL KeyStore with the new release, copy the PDC KeyStore files from the 12.0 Patch Set 3 Helm chart to the 15.0.x **oc-cn-op-job-helm-chart/pdc/pdc\_keystore/** directory.

3. Create an **override-values-15.yaml** file.

You will use this file with the 15.0.x version of **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.

4. In your **override-values-15.yaml** file, set the following keys:

- In the BRM section:
  - **ocbrm.is\_upgrade**: Set this to **true**.
  - **ocbrm.existing\_rootkey\_wallet**: Set this to **false**.
  - **ocbrm.db.\***: Set the BRM database schema details to the same values as your 12.0 Patch Set 3 release.
- In the PDC section:
  - **ocpdc.configEnv.pdcSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
  - **ocpdc.configEnv.crossRefSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
  - **ocpdc.configEnv.rcuPrefix**: Set this key to a new prefix to create a new RCU schema.
  - **ocpdc.configEnv.transformation.upgrade**: Set this to **true**. (For upgrades to 15.0.0 only)
  - **ocpdc.secretValue.walletPassword**: Set this to the same value as your 12.0 Patch Set 3 release.
  - **ocpdc.configEnv.deployAndUpgradeSite2**: Set this to **true**. (For upgrades to 15.0.1 only)
  - **ocpdc.configEnv.upgrade**: Set this to **true**. (For upgrades to 15.0.1 only)
- In the Business Operations Center section:
  - **ocboc.boc.configEnv.bocSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
  - **ocboc.boc.configEnv.runUpgrade**: Set this to **true**.

- **ocboc.boc.configEnv.rcuPrefix**: Set this to a new prefix to create a new RCU schema.
  - **ocboc.boc.secretVal.\***: Set the Business Operations Center passwords to the same values as your 12.0 Patch Set 3 release.
  - In the Billing Care section:
    - **ocbc.bc.configEnv.rcuPrefix**: Set this to a new prefix to create a new RCU schema.
    - **ocbc.bc.secretVal.\***: Set the Billing Care passwords to the same values as your 12.0 Patch Set 3 release.
  - In the Billing Care REST API section:
    - **ocbc.bcws.configEnv.rcuPrefix**: Set this to a new prefix to create a new RCU schema.
    - **ocbc.bcws.secretVal.\***: Set the Billing Care REST API passwords to the same values as your 12.0 Patch Set 3 release.
5. Create WebLogic domains by running the 15.0.x version of **oc-cn-op-job-helm-chart** from the **helmcharts** directory:

```
helm install oldOpJobReleaseName oc-cn-op-job-helm-chart --namespace
oldBrmNameSpace --namespace oldOverrideValuesFile --values override-
values-15.yaml
```

6. Install BRM cloud native services by running the 15.0.x version of **oc-cn-helm-chart** from the **helmcharts** directory:

```
helm install oldBrmReleaseName oc-cn-helm-chart --namespace
oldBrmNameSpace --values oldOverrideValuesFile --values override-
values-15.yaml
```

## Dropping the ECE Persistence Database Schema from Site 2

To drop the ECE persistence database schema from Site 2:

1. Uninstall the ECE 12.0 Patch Set 3 Helm chart:

```
helm uninstall oldEceReleaseName
```

where *oldEceReleaseName* is the release name for your 12.0 Patch Set 3 version of **oc-cn-ece-helm-chart**.

2. Delete the 12.0 Patch Set 3 version of **ece-persistence-job** from your system by running this command:

```
kubectl delete job ece-persistence-job --namespace oldBrmNameSpace
```

where *oldBrmNameSpace* is the namespace for your 12.0 Patch Set 3 version of your BRM Helm release.

**Note**

The ECE 15.0.x Helm chart re-creates the persistence database schema when you install it later.

## Installing ECE 15.0.x Cloud Native on Site 2

To install ECE 15.0.x cloud native on Site 2:

1. Download and extract the ECE 15.0.x cloud native (**oc-cn-ece-helm-chart**) from Oracle Software Delivery Cloud (<https://edelivery.oracle.com>).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)" for more information.
2. Create an **override-values-ece-15.yaml** file.
3. In the file, configure your ECE 15.0.x cloud native services by following the instructions in "[Configuring ECE Services](#)".
4. Deploy the ECE 15.0.x cloud native services by entering this command from the **helmcharts** directory:

```
helm install oldEceReleaseName oc-cn-ece-helm-chart --namespace  
oldBrmNameSpace --values oldOverrideValuesFile --values override-values-  
ece-15.yaml
```

where:

- *oldEceReleaseName* is the release name for your 12.0 Patch Set 3 version of **oc-cn-ece-helm-chart**.
- *oldBrmNameSpace* is the namespace for your 12.0 Patch Set 3 **oc-cn-helm-chart** deployment.
- *oldOverrideValuesFile* is the file name and path to the 12.0 Patch Set 3 version of your **override-values.yaml** file for **oc-cn-ece-helm-chart**.

## Failing Over Site 1 to Site 2

Site 2 has been upgraded at this stage, but it is not handling any usage requests. You now switch usage request processing from Site 1 to Site 2.

To fail over from Site 1 to Site 2:

**Note**

Delete the **NFProfileConfiguration** and **NFServiceConfiguration** ECE configuration MBeans from the site running the ECE 12.0 patch set (production site) before federating cache data to the backup site from the production site. See "Changing the ECE Configuration During Runtime" in *BRM Cloud Native System Administrator's Guide* for information about changing ECE configuration MBeans.

1. Start the Coherence federation process from Site 1 to Site 2. This provisions the empty ECE Site 2 cache with the latest data from Site 1.
2. On Site 2, mark Site 1 inactive to ensure no preferred site routing occurs from Site 2 to Site 1.

**Note**

Site 1 also has Site 2 marked as inactive.

3. Check that the federation process is up to date. After the federation process completes successfully:
  - The ECE pods in Site 2 transition to the **Running** state.
  - Site 2 transitions to the **Usage Processing** state and spawns Monitoring Agent pods.

**Note**

Follow the steps in "Reloading ECE Application Configuration Changes" in the *BRM Cloud Native System Administrator's Guide* to update the values:

- Under **brmGatewayConfigurations** set the **kafkaPartition** value to **3**.
- Under **httpGatewayConfigurations** set the **kafkaPartition** value to **5,6,7,8,9**.

Stop all usage requests to Site 1 and then redirect them to Site 2.

5. Switch provisioning to BRM on Site 2. To do so:
  - a. On BRM Site 2, remove any connections to EM Gateway on Site 1.
  - b. On the client side, switch provisioning to BRM on Site 2.
  - c. On ECE Site 1, stop the EM Gateway.

**Note**

If latency between Site 1 and Site 2 is too high to give you acceptable performance for the provisioning flow, reverse the Oracle Data Guard roles and make the BRM database on Site 2 active. Be aware that this can cause a service interruption of a few minutes for provisioning.

6. Check that the federation process from Site 1 to Site 2 has completed.

Now that all traffic to Site 1 has stopped, ensure all data from Site 1 has been synchronized with Site 2. The Coherence Federation Metrics should display **IDLE** instead of **YIELDING**.
7. Check that all rated events in the ECE Site 1 cache have been extracted by running the **query** utility:

```
./query.sh  
Coherence Command Line Tool  
CohQ1> select value() from AggregateObjectUsage
```

If successful, this command returns zero entries.

8. On Site 1, check that all Site 1 rated events have been extracted from the persistence database and are present in BRM by using SQL\*Plus:

```
sqlplus pin@databaseName
Enter password: password
```

```
SQL> select count(*) from ratedevent_site1Name
```

where *databaseName* is the service name or database alias of the BRM database, and *password* is the password for the **pin** user.

#### Note

The Site 1 persistence database might contain some Site 2 events. After the Site 2 to Site 1 federation process is stopped, these events are not extracted or purged. However, they will be processed on Site 2. You can ignore these events because they get purged when you re-create the Site 1 persistence database later.

9. Stop the federation process from Site 1 to Site 2.  
Site 1 is isolated and ready for the upgrade.

## Uninstalling BRM and ECE from Site 1

Uninstall the 12.0 Patch Set 3 version of BRM and ECE from Site 1.

To do so, perform the following steps on Site 1:

1. Uninstall BRM 12.0 Patch Set 3 cloud native:

```
helm uninstall BrmReleaseName --namespace BrmNameSpace
```

2. Uninstall ECE 12.0 Patch Set 3 cloud native

```
helm uninstall EceReleaseName --namespace BrmNameSpace
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance.
- *EceReleaseName* is the release name for **oc-cn-ece-helm-chart** and is used to track this installation instance.
- *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM Helm chart.

These commands delete all of the resources associated with the chart's last release and the release history.

## Installing BRM Cloud Native on Site 1

To configure and deploy BRM 15.0.x cloud native on Site 1:

1. Download and extract the BRM 15.0.x cloud native Helm chart (**oc-cn-helm-chart**) and BRM Operator Job Helm Chart (**oc-cn-op-job-helm-chart**) from Oracle Software Delivery Cloud (<https://edelivery.oracle.com>).  
  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)" for more information.
2. To reuse your 12.0 Patch Set 3 SSL KeyStore, copy the PDC KeyStore files from the 12.0 Patch Set 3 Helm chart to the 15.0.x **oc-cn-op-job-helm-chart/pdc/pdc\_keystore/** directory.
3. Create an **override-values-15.yaml** file.  
  
You will use this file with the 15.0.x version of **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.
4. In your **override-values-15.yaml** file, set the following keys:
  - In the BRM section:
    - **ocbrm.is\_upgrade**: Set this to **true**.
    - **ocbrm.existing\_rootkey\_wallet**: Set this to **false**.
    - **ocbrm.db.\***: Set the BRM database schema details to the same values as your 12.0 Patch Set 3 release.
  - In the PDC section:
    - **ocpdc.configEnv.pdcSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
    - **ocpdc.configEnv.crossRefSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
    - **ocpdc.configEnv.rcuPrefix**: Set this key to a new prefix to create a new RCU schema. The value must be different from the one used in Site 2.
    - **ocpdc.configEnv.transformation.upgrade**: Set this to **true**. (For upgrades to 15.0.0 only)
    - **ocpdc.secretValue.walletPassword**: Set this to the same value as your 12.0 Patch Set 3 release.
    - **ocpdc.configEnv.deployAndUpgradeSite2**: Set this to **false**. (For upgrades to 15.0.1 only)
    - **ocpdc.configEnv.upgrade**: Set this to **true**. (For upgrades to 15.0.1 only)
  - In the Business Operations Center section:
    - **ocboc.boc.configEnv.bocSchemaUserName**: Set this to the same value as your 12.0 Patch Set 3 release.
    - **ocboc.boc.configEnv.runUpgrade**: Set this to **true**.
    - **ocboc.boc.configEnv.rcuPrefix**: Set this to a new prefix for the RCU schema. The value must be different from the one used in Site 2.
    - **ocboc.boc.secretVal.\***: Set the Business Operations Center passwords to the same values as your 12.0 Patch Set 3 release.
  - In the Billing Care section:
    - **ocbc.bc.configEnv.rcuPrefix**: Set this to a new prefix for the RCU schema. This value must be different from the one used in Site 2.
    - **ocbc.bc.secretVal.\***: Set the passwords to the same values as your 12.0 Patch Set 3 release.

- In the Billing Care REST API section:
  - **ocbc.bcws.configEnv.rcuPrefix**: Set this to a new prefix for the RCU schema. This value must be different from the one used in Site 2.
  - **ocbc.bcws.secretVal.\***: Set the passwords to the same values as your 12.0 Patch Set 3 release.
- 5. Save and close the file.
- 6. Create WebLogic domains by running the 15.0.x version of **oc-cn-op-job-helm-chart** from the **helmcharts** directory:

```
helm install oldOpJobReleaseName oc-cn-op-job-helm-chart --namespace
oldBrmNameSpace --namespace oldOverrideValuesFile --values override-
values-15.yaml
```

where:

- *oldOpJobReleaseName* is the release name assigned to your 12.0 Patch Set 3 release of the **oc-cn-op-job-helm-chart** installation.
  - *oldBrmNameSpace* is the namespace for your 12.0 Patch Set 3 version of the BRM deployment.
  - *oldOverrideValuesFile* is the file name and path of your 12.0 Patch Set 3 version of the **override-values.yaml** file for **oc-cn-op-job-helm-chart**.
7. Install BRM cloud native services by running the 15.0.x version of **oc-cn-helm-chart** from the **helmcharts** directory:

```
helm install oldBrmReleaseName oc-cn-helm-chart --namespace
oldBrmNameSpace --values oldOverrideValuesFile --values override-
values-15.yaml
```

where *oldBrmReleaseName* is the release name assigned to your 12.0 Patch Set 3 version of the **oc-cn-helm-chart** installation.

## Dropping the ECE Persistence Database Schema from Site 1

To drop the ECE 12.0 Patch Set 3 persistence database schema from Site 1:

1. Uninstall the ECE 12.0 Patch Set 3 Helm chart:

```
helm uninstall oldEceReleaseName
```

*oldEceReleaseName* is the release name for the 12.0 Patch Set 3 version of **oc-cn-ece-helm-chart**.

2. Uninstall the 12.0 Patch Set 3 version of **ece-persistence-job** from your system by running this command:

```
kubectl delete job ece-persistence-job --namespace oldBrmNameSpace
```

where *BrmNameSpace* is the namespace for the 12.0 Patch Set 3 version of the BRM Helm chart.

**Note**

The ECE Helm chart re-creates the database schema when you install it later.

## Installing ECE 15.0.x Cloud Native on Site 1

To install ECE 15.0.x cloud native on Site 1:

1. Download and extract the ECE 15.0.x cloud native (**oc-cn-ece-helm-chart**) from Oracle Software Delivery Cloud (<https://edelivery.oracle.com>).  
See "[Downloading Packages for the BRM Cloud Native Helm Charts and Docker Files](#)" for more information.
2. Create an **override-values-ece-15.yaml** file.
3. In the file, configure your ECE 15.0.x cloud native services by following the instructions in "[Configuring ECE Services](#)".
4. Deploy the ECE 15.0.x cloud native services by entering this command from the **helmcharts** directory:

```
helm install oldEceReleaseName oc-cn-ece-helm-chart --namespace  
oldBrmNameSpace --values oldOverrideValuesFile --values override-values-  
ece-15.yaml
```

where:

- *oldEceReleaseName* is the release name for your 12.0 Patch Set 3 version of **oc-cn-ece-helm-chart**.
- *oldBrmNameSpace* is the release name for your 12.0 Patch Set 3 **oc-cn-helm-chart** deployment.
- *oldOverrideValuesFile* is the file name and path to the 12.0 Patch Set 3 version of your **override-values.yaml** file for **oc-cn-ece-helm-chart**.

Site 1 is now upgraded, but it is idle.

## Federating ECE Cache Data Between Site 1 and Site 2

Now that Site 1 and Site 2 have been upgraded to release 15.0.x, you can restart the federation process between them.

To start the federation process of ECE cache data between Site 1 and Site 2:

1. Enable the two-way federation process between Site 1 and Site 2.
2. Check that the federation backlog is processed successfully.

After the federation process completes, ECE in Site 1 transitions to the **Usage Processing** state and spawns the Monitoring Agent pods.

3. In ECE Site 1, mark ECE in Site 2 as active. Likewise, in ECE Site 2, mark ECE in Site 1 as active.

This enables usage rating requests to be forwarded to the subscriber's preferred site.

4. Enable usage requests to Site 2.

5. Enable failover connection to the EM Gateway. On both sides, ensure the CM has a failover EM connection to the EM Gateway on the other site.

# 19

## Migrating from On-Premise BRM to BRM Cloud Native

Learn how to migrate from an on-premise release of Oracle Communications Billing and Revenue Management (BRM) 7.5.x or 12.x to a BRM 15.0 cloud native release.

Topics in this document:

- [Migrating to BRM Cloud Native](#)
- [Migrating from PDC On Premises to PDC Cloud Native](#)

### Migrating to BRM Cloud Native

To migrate from an on-premises release to BRM cloud native:

1. Upgrade your BRM and PDC database schemas to release 15.0:
  - If you are upgrading from a 12.0 or 12.0 Patch Set 1 database schema, follow the instructions in "[Installing BRM 12.0 Patch Sets](#)" in *BRM Patch Set Installation Guide*.
  - If you are upgrading from a 12.0 Patch Set 2 or later database schema, follow the instructions in "[Upgrading Your Database Schema](#)".

#### Note

- To use an existing BRM database schema with your cloud native deployment, follow the instructions in "[Deploying BRM with an Existing Schema](#)".
- To use an existing Business Operations Center database schema, point to your existing Business Operations Center schema in the **override-values.yaml** file for **oc-cn-op-job-helm-chart**. See "[Adding Business Operations Center Keys for oc-cn-op-job-helm-chart](#)".

2. If you customized BRM or Billing Care, layer your customizations on top of the images provided with this release before deploying the images. For guidelines about customization, see "[Customizing BRM Cloud Native Services](#)".

### Migrating from PDC On Premises to PDC Cloud Native

Before migrating your system, perform these steps:

- Ensure you have upgraded your PDC on-premises release to version 12.0 Patch Set 3 or later. Follow the instructions in "[Upgrading Pricing Design Center Software](#)".
- Back up your existing PDC and cross-reference database schemas.
- Download and load the PDC 15.0.x cloud native images to the image repository or to the virtual machine (VM) where PDC is deployed.
- Download and extract the BRM 15.0.x Helm charts.

- Verify that BRM has been upgraded to the same 15.0.x cloud native version as PDC cloud native.

The following are the high-level steps for migrating PDC from an on-premises release to cloud native:

1. Do the following on your PDC on-premises system:
  - a. Stop your on-premises PDC domain.
  - b. (Release 15.0.0 only) Connect to your PDC database schema and rename the following tables:
    - Rename the PDCJD\_WLSTORE table to PDCJD\_WLSTORE\_BACKUP
    - Rename the WL\_LL\_R\_ADMINSERVER table to WL\_LL\_R\_ADMINSERVER\_BACKUP
2. In your **override-values.yaml** file for **oc-cn-op-job-helm-chart**, set the following keys to match the PDC and cross-reference database schema details for your existing on-premises release:

```
ocpdc:
  configEnv:
    crossRefSchemaPDCTableSpace:
    crossRefSchemaTempTableSpace:
    crossRefSchemaUserName: UserName
    pdcSchemaPDCTableSpace:
    pdcSchemaTempTableSpace:
    pdcSchemaUserName: UserName
  secretValue:
    crossRefSchemaPassword: Password
    pdcSchemaPassword: Password
```

3. Set any other mandatory keys in your **override-values.yaml** file. For more information, refer to "[Adding PDC Keys for oc-cn-op-job-helm-chart](#)".
4. Deploy the PDC 15.0.x domain to your cloud native environment:

```
helm install OpJobReleaseName oc-cn-op-job-helm-chart --namespace
BrmNameSpace --values OverrideValuesFile
```

where:

- *OpJobReleaseName* is the release name for **oc-cn-op-job-helm-chart** and is used to track this installation instance. It must be different from the one used for the BRM Helm chart.
  - *BrmNameSpace* is the namespace in which to create BRM Kubernetes objects for the BRM and Job Helm charts.
  - *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **oc-cn-op-job-helm-chart/values.yaml** file.
5. In your **override-values.yaml** file for **oc-cn-helm-chart**, set the **ocpdc.configEnv.upgrade** key to **true**.

6. Deploy the PDC 15.0.x cloud native services to your cloud native environment, and ensure that the **pd-domain** pod is in completed status:

```
helm install BrmReleaseName oc-cn-helm-chart --namespace BrmNameSpace --  
values OverrideValuesFile
```

where:

- *BrmReleaseName* is the release name for **oc-cn-helm-chart** and is used to track this installation instance. It must be different from the one used for **oc-cn-op-job-helm-chart**.
  - *OverrideValuesFile* is the path to a YAML file that overrides the default configurations in the **oc-cn-helm-chart/values.yaml** file.
7. Ensure all files are upgraded successfully. To do so, check the status of the SQL upgrade log files in `pd-brm-pvc`. If there is no failure message, all files upgraded successfully.

# Part VI

## Troubleshooting BRM Cloud Native Deployments

This part provides information about troubleshooting issues that may occur while deploying Oracle Communications Billing and Revenue Management (BRM) cloud native in your system. It contains the following chapters:

- [Troubleshooting Your BRM Cloud Native Deployment](#)

# Troubleshooting Your BRM Cloud Native Deployment

Learn how to solve problems that may occur after the installation or upgrade of your Oracle Communications Billing and Revenue Management (BRM) cloud native system.

Topics in this document:

- [Problems with the Helm Installation](#)
- [Helm Installation Fails with Time-Out Error](#)
- [BRM Cloud Native Deployment Out of Memory Errors](#)
- [PDC Messages Stuck in Rating Engine Queues](#)
- [PDC Interceptor Pod is Started But Went to Error State](#)
- [eceTopology.conf Errors While Restarting Pods](#)
- [SyncPDC Pod Fails to Start and Enters CrashLoopBackOff State](#)

## Problems with the Helm Installation

If a Helm installation encounters errors, such as an incorrect namespace, follow these steps to get back to a state where you can fix the issue and do a new installation.

### Note

For more information about Kubernetes commands, see "[kubectl Cheat Sheet](#)" in the Kubernetes documentation.

1. Check the state of the deployment:

```
kubect1 get pods --output wide --namespace NameSpace
```

To see information about a specific pod:

```
kubect1 describe pod PodName --namespace NameSpace
```

2. Use the **helm rollback** command to go back to a previous revision of the chart, or use the **helm uninstall** command to uninstall the chart. See "Rolling Back A Release To A Previous Revision" in *BRM Cloud Native System Administrator's Guide*, or see "[Helm Uninstall](#)" in the Helm documentation.
3. If neither rolling back nor uninstalling the chart are successful, do the following to identify Kubernetes resources that did not install correctly and then delete them:

- Check and delete all other stateful set components from the cluster:

```
kubect1 get sts
```

If you identify a stateful set that you want to delete, scale the number of replicas:

```
kubectl scale statefulsets StatefulSetName --replicas=n
```

where *StatefulSetName* is the name of a stateful set, and *n* is the number of replicas you are scaling to. For more information, see "[Scale a StatefulSet](#)" in the Kubernetes documentation.

Then, delete the stateful set:

```
kubectl delete StatefulSetName
```

You can run **kubectl get sts** again to verify the deletions.

- If you need to clean up Apache Kafka and Apache ZooKeeper, scale to 0 and then delete:

```
kubectl scale sts/kafka_pod --replicas=0
kubectl scale sts/zookeeper_pod --replicas=0
kubectl get pods
kubectl get sts
kubectl delete sts kafka_pod zookeeper_pod
```

- If necessary, check any PVC, Secret, ConfigMap, or service that was created by the deployment. If the output from any of these commands shows something that you want to clean up, you can use **kubectl delete** to remove it.

For example:

```
kubectl get pvc --all-namespaces
kubectl delete pvc PVCName
```

```
kubectl get secrets --all-namespaces
kubectl delete secret SecretName
```

```
kubectl get configmap --all-namespaces
kubectl delete configmap ConfigMapName
```

```
kubectl get svc --all-namespaces
kubectl delete svc SVC1 SVC2
```

## Helm Installation Fails with Time-Out Error

After you deploy a Helm chart, you may receive the following error message indicating that the Helm chart installation failed:

```
Error: failed post-install: timed out waiting for the condition
```

This occurs because a post-installation job took longer than five minutes to complete.

To resolve the issue:

1. Purge your Helm release:

```
helm delete BrmReleaseName --purge
```

This removes and purges all resources associated with the last revision of the release.

2. Run the **Helm install** command again.

If that does not fix the problem, increase the amount of time Kubernetes waits for a command to complete by including the **--timeout** *duration* argument with the **helm install** command. For example, to set the timeout duration to 10 minutes, you would enter this command:

```
helm install BrmReleaseName oc-cn-helm-chart --namespace BrmNameSpace --timeout 10m --values OverrideValuesFile
```

## BRM Cloud Native Deployment Out of Memory Errors

After you deploy BRM cloud native, you may receive an error message similar to the following:

```
ERROR: cm_cache_heap_malloc: name="fm_bparams_cache" - out of memory, size
requested=2216,high_val=960
cm_cache_flist: PIN_ERR_NO_MEM:requested=2216, used=121456, allocated=122880, chunk=30,
cache name="fm_bparams_cache"
```

To resolve the issue:

1. In your **oc-cn-helm-chart** directory, open your CM ConfigMap file **configmap\_pin\_conf\_cm.yaml**.
2. Add the following **fm\_bparams\_cache** entry to the file:
  - `cm_cache fm_bparams_cache 40,245760,23`
3. Save and close the file.
4. Run the **helm** upgrade command for **oc-cn-helm-chart**:

```
helm upgrade BrmReleaseName oc-cn-helm-chart --values OverrideValuesFile --namespace
BrmNameSpace
```

## PDC Messages Stuck in Rating Engine Queues

Occasionally, PDC messages and changesets may become stuck in the rating engine queues.

To resolve the issue, delete both the RRE and BRE pods by running the following command:

```
kubectl --namespace BrmNameSpace delete pod PdcPodName
```

where:

- *BrmNameSpace* is the namespace in which the BRM Kubernetes objects reside.
- *PdcPodName* is the name of the pod.

Kubernetes automatically restarts the deleted pod, which restarts the transformation engine. Messages should start flowing again.

## PDC Interceptor Pod is Started But Went to Error State

After you deploy PDC, the Interceptor pod may start but immediately transition to an error state.

This may occur because the RCU prefix is configured incorrectly. To find out if this is the case, run the following command:

```
kubectl describe domain DomainName --namespace NameSpace
```

If the issue is related to the RCU prefix, you will see something similar to the following:

```
WLSDPY-12409: createDomain failed to create the domain: Failed to get FMW
infrastructure database defaults from the service table : Got exception when auto
configuring the schema component(s) with data obtained from shadow table:
Failed to build JDBC Connection object:
```

To resolve the issue:

1. Make sure that the RCU prefix is configured successfully as part of **oc-cn-op-job-helm-chart**. To do so, run the following command:

```
kubectl get pod --namespace BrmNameSpace
```

If it is configured correctly, **pdcc-configure-rcu-xxxxx** will show a Completed status.

2. Make sure that the RCU prefix and password configured in the **override-values.yaml** file for **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart** matches, and that a valid host name, port, and service name have been configured in the **values.yaml** file.

If the values are not configured properly, do the following:

1. Uninstall PDC and then set the **ocpdc.isEnabled** key to **false** in your **override-values.yaml** file for **oc-cn-helm-chart**.
2. Run the **Helm upgrade** command for **oc-cn-helm-chart**.  
Wait until the PDC pods have stopped.
3. In the **override-values.yaml** file for **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**, configure the **ocpdc.configEnv.rcuPrefix** key and set the **ocpdc.isEnabled** key to **true**.
4. Run the **helm upgrade** command for **oc-cn-helm-chart** and **oc-cn-op-job-helm-chart**.  
Wait until the PDC pods are in Running status.

For more information about troubleshooting pod errors, see "[Troubleshooting](#)" in *Oracle WebLogic Kubernetes Operator Samples*.

## eceTopology.conf Errors While Restarting Pods

While restarting the pricingupdater and brmgateway pods in your ECE cloud native deployment, you may receive an error message similar to the following:

```
ERROR MonitorFrameworkMessagesBundle-31300: Failed to initialize grid manager based on topology file: eceTopology.conf property file:
ece.propertiesjava.lang.IllegalArgumentException: Nodes on a given host must be assigned unique JMX ports
(check nodes: '[PricingUpdater node pricingupdater-6d575bf75b-r5q2t on Host pricingupdater.ece-server.cluster, PricingUpdater node pricingupdater-fbb9d7fb7-kqmxv on Host pricingupdater.ece-server.cluster]')
```

This occurs because ECE cloud native has written invalid entries to the **eceTopology.conf** property file during the startup process.

To resolve the issue, do not restart the pricingupdater and brmgateway pods. Instead, scale down and then scale up those pods.

For example, to scale down and scale up the brmgateway pod:

1. Scale down the brmgateway pod to **0**:

```
kubectl --namespace BrmNameSpace scale deploy brmgateway1 --replicas=0
```

Wait for the brmgateway pod to stop.

2. Scale back up the brmgateway pod to **1**:

```
kubectl --namespace BrmNameSpace scale deploy brmgateway1 --replicas=1
```

## SyncPDC Pod Fails to Start and Enters CrashLoopBackOff State

After a fresh deployment of PDC 15.x, the **SyncPDC** pod deploys and the **SyncPDC** job runs, but the pod fails and enters a **CrashLoopBackOff** state.

This may occur because the deployment fails to create the **WorkEntity** table. The **cdc\_domain** logs show that the table is missing.

If this issue occurs, the **SyncPDC** log file shows an error similar to the following:

```
oracle.communications.brm.pdc.syncpdc.SyncPDCJob execute
SEVERE: Exception executing SyncPDC job
javax.ejb.EJBException: Problem finding error class; nested exception is:
java.lang.ClassNotFoundException: Failed to load class
javax.persistence.PersistenceException
```

If you try to create the table manually and see the following error, you can see the reason that the deployment failed to create the table:

```
ORA-60019: Creating initial extent of size 14 in tablespace of extent size 5
```

To resolve the issue, update the PDC schema as follows:

1. Delete the existing PDC tablespace.
2. Create a new PDC tablespace with higher values for the **STORAGE** parameters.

The following is a sample SQL\*Plus command for creating a tablespace for PDC:

```
create tablespace pdctables datafile '/PINDB/pdctables.dbf'
size 600M reuse autoextend on default
STORAGE (initial 1M next 1M pctincrease 0);
```

### Note

Choose the datafile name, location, size and storage values (/PINDB/pdcfiles.dbf, 600M and 1M in the example above) based on your environment.

For more information about creating tablespaces, see "Creating Tablespace Storage" in *BRM Installation Guide*.

3. Redeploy PDC cloud native.