

Oracle Utilities Analytics

Administration Guide

Release 2.5.2.0.1

E48998-07

June 2016

Copyright © 2016, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

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

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

Contents

Preface	i-i
Audience	i-i
Prerequisite Knowledge.....	i-i
How This Guide is Organized.....	i-ii
Related Documents	i-ii
Conventions.....	i-iii
Acronyms	i-iii
Terminology	i-iv
Oracle Utilities Analytics.....	i-iv
Oracle Data Integrator	i-v
Oracle GoldenGate.....	i-vii
Oracle Business Intelligence Enterprise Edition.....	i-viii
Chapter 1	
Overview of Oracle Utilities Analytics	1-1
Introduction.....	1-1
Architecture	1-2
Product Overview	1-2
Deployment.....	1-3
Chapter 2	
Oracle Utilities Extractors and Schema	2-1
Data Flow	2-1
ELT Architecture.....	2-2
Replication	2-3
Replication Patterns	2-4
Structure of the Replicated Table	2-5
DDL Support.....	2-6
Default Keys in Dimensions.....	2-9
Referencing Foreign Keys.....	2-10
Metadata Model	2-10
Products	2-11
Instances	2-11
Global Configuration.....	2-11
Entities	2-12
Job Configuration.....	2-12
Objects.....	2-12
Mapped Objects	2-12
Source Tables.....	2-13
Job Executions.....	2-13
Dependencies.....	2-13
DSI Mapping.....	2-13
GoldenGate Configuration.....	2-13
GoldenGate Checkpoint.....	2-13
Extract Parameters	2-13

Storage Configuration.....	2-13
Configuration Views	2-13
Context Configuration.....	2-14
Replication Configuration.....	2-14
Job Configuration.....	2-14
Storage Configuration.....	2-14
Dependencies.....	2-14
Jobs.....	2-14
Ready for Load	2-15
Wait Reasons.....	2-15
Oracle Data Integrator Components.....	2-16
Project	2-16
Model Folders.....	2-16
Scenario Folders	2-17
Logical Schemas	2-17
Connections	2-18
Contexts.....	2-18
Chapter 3	
Oracle Utilities Analytics Dashboards.....	3-1
Dashboard Catalogs	3-1
Oracle Business Intelligence Repository.....	3-2
Physical Layer	3-2
Business Model and Mapping Layer	3-2
Presentation Layer.....	3-3
Chapter 4	
Oracle Utilities Analytics Administration Tool.....	4-1
Configuring Global Settings	4-2
Configuring Oracle GoldenGate	4-3
Configuring Product Instance.....	4-4
Configuring Storage.....	4-5
Configuring Source Table	4-6
Configuring Target Entities.....	4-8
Configuring Jobs	4-10
Customizing the Column Properties of Type 2 Slowly Changing Dimension.....	4-12
ETL Replication View.....	4-13
NMS Configuration	4-14
ETL Job Execution.....	4-14
Chapter 5	
Configuring Oracle Utilities Analytics	5-1
Configuring Source.....	5-1
Oracle Utilities Operational Device Management (ODM).....	5-1
Oracle Utilities Customer Care & Billing (CCB).....	5-6
Oracle Utilities Network Management System (NMS).....	5-13
Oracle Utilities Mobile Workforce Management (MWM).....	5-16
Oracle Utilities Meter Data Management (MDM).....	5-22
Configuring ELT	5-35
Oracle Utilities Analytics Basic Setup	5-35
Oracle Utilities Analytics Global Configurations	5-35
Initial Setup for Star Schema	5-37
Adding Sources	5-37
Configuring the Socks Proxy Setup.....	5-38
Invoking Source Configuration User Interface	5-38
Using Source Configuration User Interface to Register or Upgrade Sources	5-40
Deploying Oracle GoldenGate Scripts	5-44
Initiating the ELT Jobs.....	5-46

Setting up the Scheduler.....	5-46
Enabling the ELT Jobs	5-47
Configuring the Security.....	5-47
Configuring the Analytics.....	5-47
Maintaining the Administration Dashboards.....	5-47
Configuring the Dashboards	5-48
Enabling Weather Information.....	5-54

Chapter 6

Managing Oracle Utilities Analytics	6-1
Managing ELT Processes	6-1
Recommended Configuration	6-2
Weblogic	6-3
Scheduling	6-3
Monitoring the Jobs	6-4
Debugging	6-5
Handling the Errors	6-5
Reloading the Data.....	6-6
Performance Recommendations	6-10
Partitioning.....	6-10
Parallelism.....	6-11
Optimizing the Top N Answers	6-12
Archiving the Historical Data	6-12

Preface

This guide provides instructions for configuring and administering Oracle Utilities Analytics (OUA), including:

- [Audience](#)
- [Prerequisite Knowledge](#)
- [How This Guide is Organized](#)
- [Related Documents](#)
- [Conventions](#)
- [Acronyms](#)
- [Terminology](#)

Audience

This administration guide is intended for system and data warehouse administrators who need to configure and administer Oracle Utilities Analytics. However, it is a good reference material for implementers and developers interested in understanding how the application works.

Prerequisite Knowledge

Oracle Utilities Extractors and Schema and Oracle Utilities Analytics Dashboards use several technologies. You should have knowledge of the following before configuring and administering Oracle Utilities Analytics:

- Oracle Data Warehouse:
http://docs.oracle.com/cd/E11882_01/server.112/e25554/toc.htm
- Oracle Data Integrator:
http://docs.oracle.com/cd/E28280_01/nav/odi.htm
- Oracle GoldenGate:
<http://docs.oracle.com/goldengate/1212/gg-winux/index.html>
- Oracle WebLogic Server:
http://docs.oracle.com/cd/E23943_01/nav/wls.htm
- Oracle Business Intelligence Enterprise Edition:
http://docs.oracle.com/cd/E28280_01/nav/bi.htm

How This Guide is Organized

This guide helps you configure and administer Oracle Utilities Extractors and Schema and Oracle Utilities Analytics Dashboards. This guide refers to these two products together as Oracle Utilities Analytics. If any topic is specific to only one of the products, then it is specifically mentioned.

[Chapter 1: Overview of Oracle Utilities Analytics](#) introduces the reader to Oracle Utilities Analytics product, its fundamental architecture and how the product works.

[Chapter 2: Oracle Utilities Extractors and Schema](#) lists the various components related to the extractors and schemas within the product and how they work together in the product.

[Chapter 3: Oracle Utilities Analytics Dashboards](#) lists the various components related to the dashboards within the product and how they work together in the product.

[Chapter 4: Oracle Utilities Analytics Administration Tool](#) describes how to maintain the out-of-the-box ELT jobs using the Oracle Utilities Analytics Administration Tool.

[Chapter 5: Configuring Oracle Utilities Analytics](#) goes into detail on how to configure the Oracle Utilities Analytics product and how to manage and maintain an implementation of Oracle Utilities Analytics.

[Chapter 6: Managing Oracle Utilities Analytics](#) delves into the important configurations related to the Oracle Utilities Analytics product and additional recommendations on utilizing the partitioning feature of the Oracle database.

Related Documents

The following documentation is included with this release:

- *Oracle Utilities Analytics Getting Started Guide*
- *Oracle Utilities Analytics License Information User Guide*
- *Oracle Utilities Analytics Developer's Guide*
- *Oracle Utilities Analytics Installation Guide*
- *Oracle Utilities Analytics Quick Install Guide*
- *Oracle Utilities Analytics Release Notes*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Meter Data Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Customer Analytics, Revenue Analytics and Credit & Collections Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Exception Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Mobile Workforce Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Distribution Analytics and Outage Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Work and Asset Analytics Metric Reference Guide*
- *Oracle Utilities Analytics Dashboards for Oracle Utilities Operational Device Analytics Metric Reference Guide*
- *Oracle Utilities Extractors and Schema for Oracle Utilities Customer Care and Billing Data Mapping Guide*
- *Oracle Utilities Extractors and Schema for Oracle Utilities Meter Data Management Data Mapping Guide*

- *Oracle Utilities Extractors and Schema for Oracle Utilities Mobile Workforce Management Data Mapping Guide*
- *Oracle Utilities Extractors and Schema for Oracle Utilities Network Management System Data Mapping Guide*
- *Oracle Utilities Extractors and Schema for Oracle Utilities Operational Device Management Data Mapping Guide*
- *Oracle Utilities Extractors and Schema for Oracle Utilities Work & Asset Management Data Mapping Guide*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Acronyms

The list of acronyms used in this guide is as explained below:

- **APEX**: Oracle Application Express
- **CC&B**: Oracle Utilities Customer Care and Billing
- **CDC**: Changed Data Capture
- **ELT**: Extraction, Loading and Transformation
- **ETL**: Extraction, Transformation and Loading
- **MDM**: Oracle Utilities Meter Data Management
- **MWM**: Oracle Utilities Mobile Workforce Management
- **NMS**: Oracle Utilities Network Management System
- **OBIEE**: Oracle Business Intelligence Enterprise Edition
- **ODI**: Oracle Data Integrator
- **ODM**: Oracle Utilities Operational Device Management
- **OGG**: Oracle GoldenGate
- **OUA**: Oracle Utilities Analytics
- **OWB**: Oracle Warehouse Builder
- **WAM**: Oracle Utilities Work and Asset Management

Terminology

This section covers basic terminology used in the document:

- [Oracle Utilities Analytics](#)
- [Oracle Data Integrator](#)
- [Oracle GoldenGate](#)
- [Oracle Business Intelligence Enterprise Edition](#)

Oracle Utilities Analytics

This sub-section includes the following:

- [Product](#)
- [Instance](#)
- [Entity](#)
- [Late Arriving Dimension](#)
- [Data Source Indicator](#)

Product

In this context, 'Product' is referred to the source systems or edge applications, which act as source for the transactional data being loaded in the data warehouse. The following scenarios are expected depending on your global presence and implementation of your products:

- The edge applications may be deployed in different locations across the globe, and thus the transactions can occur in multiple time zones.
- There can be multiple instances of each of the source systems that you have implemented.
- Oracle Utilities products support cross product integration and hence it is possible that some of the master data is integrated across the multiple systems.

Instance

You may have multiple instances of the same product that you want to integrate with the Oracle Utilities Analytics solution. Each source product deployment is considered as an instance of the product. A product instance may include multiple servers in a clustered deployment.

Instances differ by the source database connections and possibly by the differing configurations. The Data Source Indicator (DSI) is unique across instances and products as this allows the ability to identify the source of the data in the data warehouse.

Entity

The data warehouse is built of star schemas, and the star schemas consist of facts and dimensions. Each fact and dimension is considered as an entity. We also use materialized views to aggregate data in the facts for improving the analytics performance. These are also considered entities.

Late Arriving Dimension

The late arriving dimensions are the dimensions where the fact (measurable quantities) table records come early when compared to the dimension table records.

Data Source Indicator

The data warehouse is an integrated database receiving data from the multiple source applications. The Data Source Indicator (DSI) in Oracle Utilities Analytics is used as an identifier for the

originating source application or system-of-record for the data. Both fact and dimension tables have a DSI value for each record.

Oracle Data Integrator

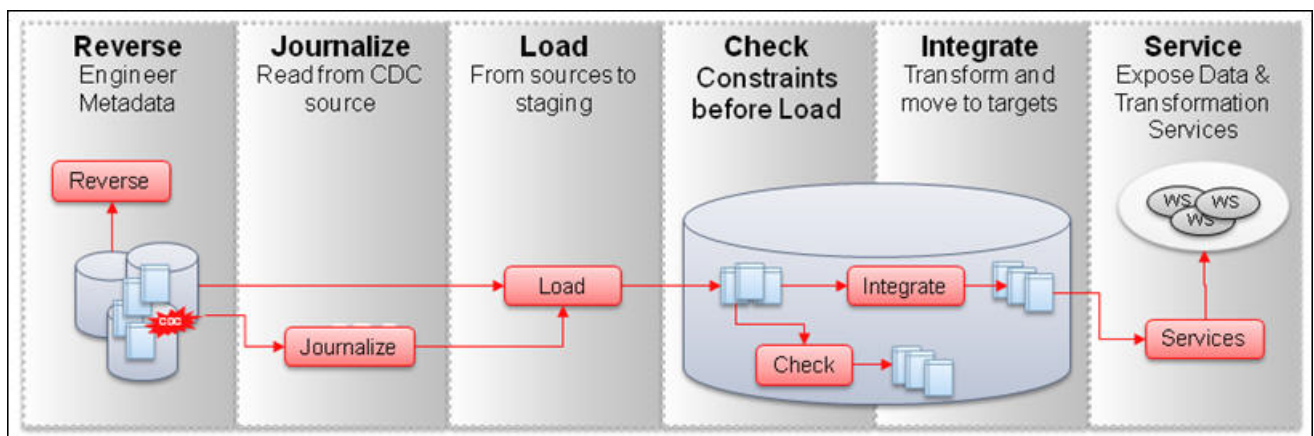
This sub-section includes the following:

- [Knowledge Modules](#)
- [Model](#)
- [Interface](#)
- [Package](#)
- [Context](#)
- [Logical Schema](#)
- [Physical Schema](#)
- [Agent](#)
- [Reverse Engineering](#)
- [Journalization](#)
- [Change Data Capture](#)
- [Scenario](#)

Knowledge Modules

The Knowledge Modules (KMs) in Oracle Data Integrator are the components that implement reusable transformation and ELT (Extraction, Loading and Transformation) strategies across different technologies.

- The knowledge modules are code templates. Each KM is dedicated to an individual task in the overall data integration process.
- A KM is reused across several interfaces or models. The benefit of Knowledge Modules is that you make a change once and it is instantly propagated to hundreds of transformations.
- The knowledge modules are based on logical tasks that are performed. They do not contain references to the physical objects (data stores, columns, physical paths, etc.).
- The knowledge modules can be written with a mix and match of different programming languages, types, and styles (native RDBMS SQL, scripting languages such as Jython or JavaScript, or even Java).



Types of Knowledge Modules

The types of knowledge modules are:

1. Reverse-engineering Knowledge Modules (RKM)
2. Check Knowledge Modules (CKM)
3. Loading Knowledge Modules (LKM)
4. Integration Knowledge Modules (IKM)
5. Journalizing Knowledge Modules (JKM)
6. Service Knowledge Modules (SKM)

Note: Refer to Oracle Data Integrator documents for more details.

Model

A model is a set of data stores corresponding to data structures contained in a physical schema. The tables, files, JMS messages, elements from an XML file are represented as data stores.

Interface

An interface is a declarative mapping between the source to the target, and it utilizes Loading Knowledge Modules (LKM), Integration Knowledge Module (IKM) and Check Knowledge Module (CKM) as required for execution. The joins, aggregations, filters and data transformations are performed in an interface.

Package

A package is made up of a sequence of steps organized into an execution diagram.

Context

A context is a collection of logical schemas mapped to physical schemas. All logical schemas may not be mapped to the physical schemas in a context. A context identifies the physical schemas and connections that are utilized for the data access. The interfaces can be executed in any context as this allows reusing the same code for multiple instances.

Logical Schema

A logical schema is a logical representation of a storage location. The same physical schemas can be associated with different logical schemas.

Physical Schema

A physical schema is the actual schema where data is retrieved or stored.

Agent

An agent is a component that executes the Oracle Data Integrator jobs. Multiple agents can be created and utilized for load balancing.

Reverse Engineering

It is the process of identifying the data structures, constraints and relations from an existing data store. The data store can be any database management system, files or any other Oracle Data Integrator supported data store.

Journalization

It is the process of keeping track of data changes, enabling incremental data access, transport and transformation.

Change Data Capture

It is the process of identifying and tracking data that has changed since the last known state. The change data capture mechanisms are implemented by the Journalizing Knowledge Modules (JKM).

Scenario

A scenario is a self-contained prepackaged executable version of a package, interface, or procedure in Oracle Data Integrator. The logic within a scenario cannot be viewed using the Oracle Data Integrator client.

Oracle GoldenGate

This sub-section includes the following:

- [Manager](#)
- [Extract](#)
- [Data Pump](#)
- [Replicat](#)
- [Parameter files](#)
- [Trail files](#)

Manager

The Manager process performs a number of functions such as starting other Oracle GoldenGate processes, trail log file management and reporting. The Manager process needs to be configured on both source as well as target systems and configuration is carried out via a parameter file. This is the mgr.prm file.

Extract

The Extract process is the extraction (capture) mechanism of Oracle GoldenGate. The extract process runs on a source system or on a downstream database, or both, depending on the database and the implementation requirements.

When configured for change synchronization, extract process captures the DML and DDL operations that are performed on objects in the Extract configuration. The Extract process stores these operations until it receives commit records or rollbacks for the transactions that contain them. When a rollback is received, the Extract discards the operations for that transaction. When a commit is received, the Extract persists the transaction to disk in a series of files called a trail, where it is queued for propagation to the target system. All of the operations in each transaction are written to the trail as a sequentially organized transaction unit. This design ensures both speed and data integrity.

Data Pump

A data pump is a secondary extract group within the source Oracle GoldenGate configuration. If a data pump is not used, the Extract must send the captured data operations to a remote trail on the target. In a typical configuration with a data pump, however, the primary Extract group writes to a trail on the source system. The data pump reads this trail and sends the data operations over the network to a remote trail on the target. The data pump adds storage flexibility and also serves to isolate the primary extract process from the TCP/IP activity.

In general, a data pump can perform data filtering, mapping, and conversion, or it can be configured in pass-through mode, where data is passively transferred as-is, without manipulation. The pass-through mode increases the throughput of the data pump. This bypasses all object definition lookups.

Replicat

The Replicat process runs on the target system, reads the trail on that system, and then reconstructs the DML or DDL operations and applies them to the target database. The Replicat uses the dynamic SQL to compile a SQL statement once, and then execute it many times with different bind variables.

Parameter files

Most Oracle GoldenGate functionalities are controlled by means of parameters specified in the parameter files. A parameter file is a plain text file that is read by an associated Oracle GoldenGate process.

Trail files

To support the continuous extraction and replication of the database changes, Oracle GoldenGate stores records of the captured changes temporarily on the disk in a series of files called a trail. A trail can exist on the source system, an intermediary system, the target system, or any combination of those systems, depending on how you configure Oracle GoldenGate. On the local system, it is known as an extract trail (or local trail). On a remote system, it is known as a remote trail.

Oracle Business Intelligence Enterprise Edition

This section defines the certain key terms for Oracle Business Intelligence Enterprise Edition Tool used in the Oracle Utilities Analytics product, including:

- [Oracle Business Intelligence Repository](#)
- [Physical Layer](#)
- [Business Model and Mapping layer](#)
- [Presentation Layer](#)
- [Subject Area](#)
- [Oracle Business Intelligence Presentation Catalog](#)
- [Dashboard](#)
- [Analysis/Report](#)

Note: For an exhaustive list, refer to the product documentation of Oracle Business Intelligence Enterprise Edition.

Oracle Business Intelligence Repository

It is a file that stores Oracle Business Intelligence metadata. The metadata defines logical schemas, physical schemas, physical-to-logical mappings, aggregate table navigation, and other constructs. The repository file has an extension of .rpd. Oracle Business Intelligence repositories can be edited using the Oracle Utilities Analytics Administration Tool.

Physical Layer

It is a layer of the Oracle Business Intelligence repository that contains objects that represent physical data constructs from back-end data sources. The physical layer defines the objects and relationships available for writing physical queries. This layer encapsulates source dependencies to enable portability and federation.

Business Model and Mapping layer

It is a layer of the Oracle Business Intelligence repository that defines the business, or logical, model of the data and specifies the mapping between the business model and the physical layer schemas. This layer can contain one or more business models. The business model and mapping layer determines the analytic behavior that is seen by users, and defines the superset of objects

available to you. It also hides the complexity of the source data models. Sometimes, this layer is also referred to as the logical layer

Presentation Layer

This layer provides a way to present customized, secure, role-based views of a business model to you. It adds a level of abstraction over the business model and the mapping layer in the Oracle Business Intelligence repository. The presentation layer provides the view of the data seen by you who build analyses in **Presentation Services** and other client tools and applications.

Subject Area

In an Oracle Business Intelligence repository, an object in the presentation layer that organizes and presents data about a business model. It is the highest-level object in the presentation layer and represents the view of the data that you see in **Presentation Services**. Oracle Business Intelligence repository subject areas contain presentation tables, presentation columns, and presentation hierarchies. In **Presentation Services**, subject areas contain folders, measure columns, attribute columns, hierarchical columns, and levels.

Oracle Business Intelligence Presentation Catalog

It stores business intelligence objects, such as analyses and dashboards, and provides an interface where you create, access, and manage objects, and perform specific object-based tasks (for example, export, print, and edit). The catalog is organized into folders that are either shared or personal.

Dashboard

An object that provides personalized views of corporate and external information. A dashboard consists of one or more pages. Pages can display anything that you can access or open with a web browser, such as results of analyses, images, alerts from agents, and so on.

Analysis/Report

This is a query that you create on the **Criteria** tab in **Presentation Services**. An analysis can optionally contain one or more filters or selection steps to restrict the results.

Chapter 1

Overview of Oracle Utilities Analytics

This chapter provides overview of Oracle Utilities Analytics (OUA), including:

- [Introduction](#)
- [Architecture](#)

Introduction

The purpose of the data warehouse is to publish the organization's data for supporting the decision making. Oracle Utilities Analytics consists of the pre-built analytics applications as well as a collection of the extractors and schema.

Oracle Utilities Analytics is a suite of pre-built analytics applications for the utilities industry. It is a single data repository with the pre-built data integration to the Oracle Utilities Source Applications listed below:

- Oracle Utilities Customer Care and Billing (CC&B)
- Oracle Utilities Meter Data Management (MDM)
- Oracle Utilities Network Management System (NMS)
- Oracle Utilities Work and Asset Management (WAM)
- Oracle Utilities Mobile Workforce Management (MWM)
- Oracle Utilities Operational Device Management (ODM)

Built on world class Oracle Business Intelligence Enterprise Edition (OBIEE) platform with integrated spatial features, Oracle Utilities Analytics supports end to end analytic workflows including the ability to drill back to the source applications. The extractors and schema are designed with the pre-built mapping between the source and the target. They support schema extensibility with built-in user-defined fields, dimensions and measures.

Oracle Utilities Analytics data warehouse is a separate database from the operational database. The data warehousing involves large volumes of data used primarily for analysis. The data warehouse has the following features:

- Includes pre-built star schemas, pre-built metrics, pre-built reports and pre-built dashboards.
- Makes the data structures easily accessible for the reporting needs of the end users.
- Retrieves the large volumes of data quickly. This in turn allows fast rendering of the graphics that showcase the Key Performance Indicators (KPIs).
- Contains star schemas and graphics suited for data retrieved from the various Oracle Utilities edge applications.

- Provides you with the ability to add additional star schemas and graphics as per your requirement using the required development tools.

Architecture

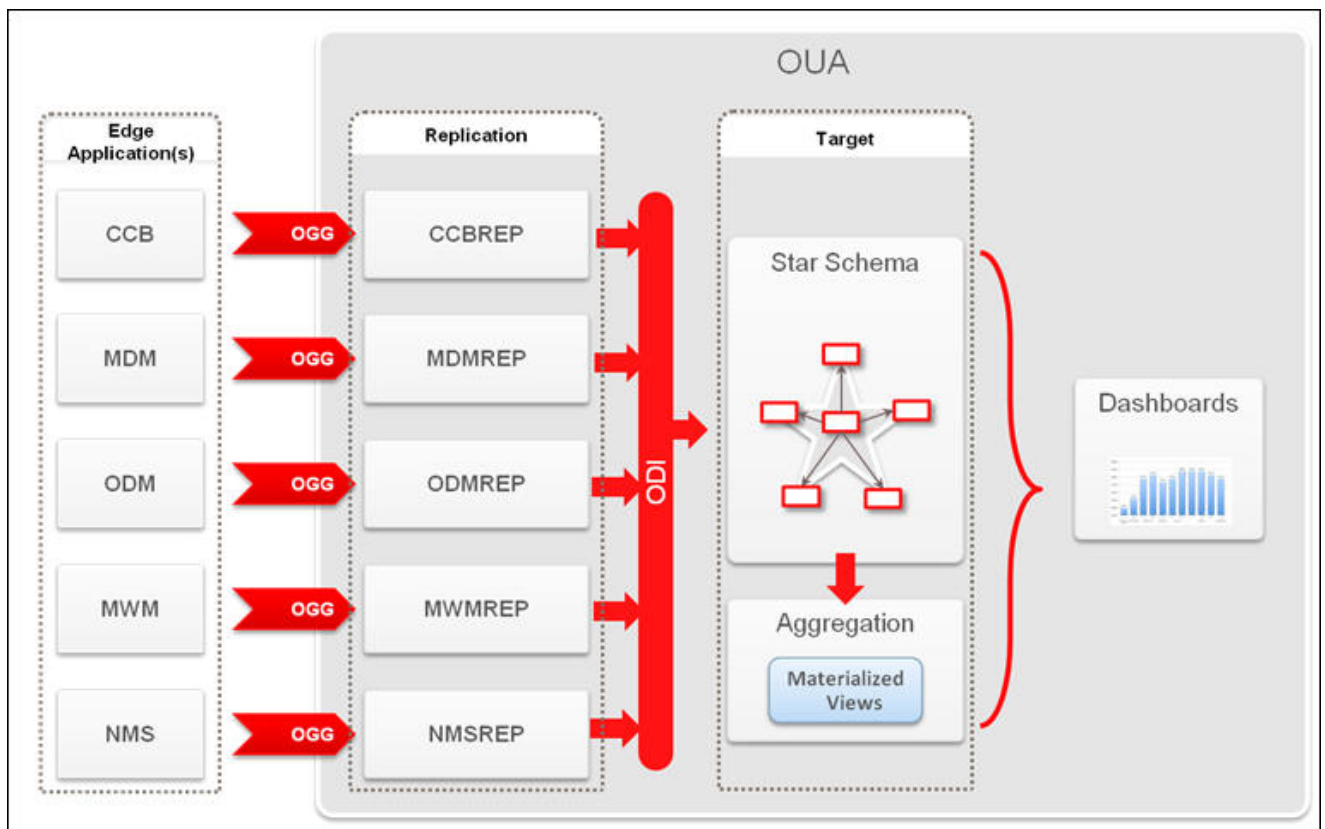
This section provides details on the product architecture and how the different components work together. The objective of this section is to provide you a high level understanding of the key components of the Oracle Utilities Analytics product and how these components work together, including:

- [Product Overview](#)
- [Deployment](#)

Product Overview

The selected tables from the edge applications are replicated to the corresponding replication schema using Oracle GoldenGate. The data in the replication area is transformed into appropriate star schemas using Oracle Data Integrator. Oracle Data Integrator is utilized to generate the aggregated materialized views based on the star schemas. Oracle Business Intelligence Enterprise Edition is used for presenting the data in the star schemas and the associated materialized views on dashboards for each subject area.

This diagram depicts the conceptual components involved in the Oracle Utilities Analytics product.




Oracle Utilities Analytics Product Components

Deployment

A standard deployment of the Oracle Utilities Analytics product consists of a database and an application server. Two Weblogic installations are performed on the application server, one for Oracle Business Intelligence Enterprise Edition and another for Oracle Data Integrator. And, Oracle GoldenGate is installed on the source database server and on the target database server.

The diagram below depicts a standard deployment of the Oracle Utilities Analytics product.



Standard Deployment

Chapter 2

Oracle Utilities Extractors and Schema

This chapter lists the various components related to the extractors and schemas within the product and how they work together in the product, including:

- [Data Flow](#)
- [ELT Architecture](#)
- [Replication](#)
- [Default Keys in Dimensions](#)
- [Metadata Model](#)
- [Configuration Views](#)
- [Oracle Data Integrator Components](#)

Data Flow


A data warehouse typically pulls data from multiple sources and transforms data stored in a relational model into a dimensional model. Oracle Utilities Analytics product is a data warehouse built on data pulled in from multiple Oracle Utilities products.

The core of the Oracle Utilities Analytics product is the ELT (Extraction, Loading and Transformation) processes, which extract data from source system tables and transform them into dimensional structure. These processes have been designed to be metadata-driven to ensure scalability and customizations without the need to change the out of the box components.

The diagram provided below shows the following four stages that all the data passes through during the ELT process:

- **Source:** This represents one of the source systems attached to the Oracle Utilities Analytics product. A predefined set of tables supports the load of the star schemas provided in the product. Oracle Utilities Analytics supports multiple source products and up to nine instances of the same source product can be attached.
- **Replicated:** Each instance of the source product gets a dedicated schema on the target database. This separates the objects from each source system and avoids any naming conflicts across multiple sources. The replicated tables are created without foreign key constraints to reduce overheads. This schema is created dynamically when a source product instance is attached to Oracle Utilities Analytics. The data is synchronized using Oracle GoldenGate.
- **Staging:** Oracle Data Integrator processes use staging schema for storing data temporarily in between transformations. The data in this schema is semi-persistent and is retained purely for debugging and traceability. The retention period is configurable.

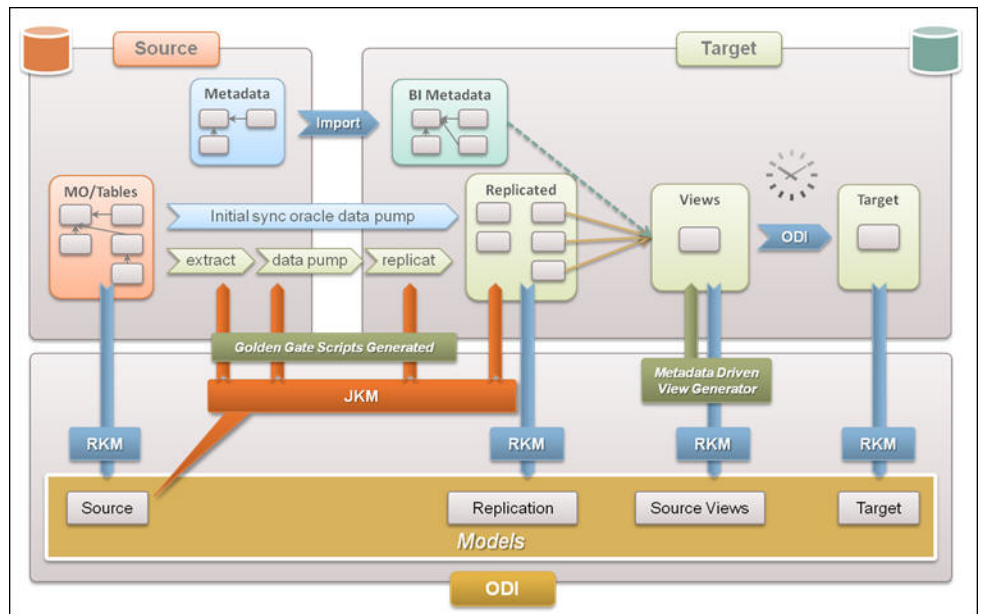
- **Target:** This is the schema, which stores the facts, dimensions, and materialized views used by the analytic dashboards.



Data Flow

ELT Architecture

The diagram below shows the various components and how they are involved in the ELT architecture.



ELT Architecture

In this diagram, the ELT architecture has been divided into three sections.

- **Source:** This section contains the components of the source product instance and components that are added to the source product environment to facilitate the real-time synchronization.
- **Target:** This section covers the components on the target Oracle Utilities Analytics database and application server to facilitate the initial and ongoing data loads into the star schemas.

-
- **Oracle Data Integrator:** This section contains the Oracle Data Integrator components that are leveraged for the configuration and the ongoing data loads.

After the installation, the product contains the metadata and the Oracle Data Integrator components designed to load data into the target star schemas using the replication schema objects. However, no source product instance is associated with Oracle Utilities Analytics at this point of time. Attaching a source product instance is the first step after installation. This involves the following actions.

1. **Importing Source System Metadata:** The product imports source product metadata into the Oracle Utilities Analytics metadata structure. This involves creation of product instance entry and detecting some of the configurations that have been implemented on the source product.
2. **Creating Job Configuration Entries:** The job configurations with predefined defaults are created for all entities associated with the source instance. All jobs are disabled by default.
3. **Reverse Engineering Models:** This step creates one or more models (Oracle Data Integrator) and executes a custom Reverse-Engineering Knowledge Module (RKM) to import specific tables into each model. The association of tables to models is pre-configured. The tables are reverse-engineered at the time of implementation to ensure that any changes in the table structures are automatically captured.
4. **Enabling Ongoing Sync:** Once the source models have been created, the Journalizing Knowledge Module (JKM) is executed. This enables journalization on the tables of the model. The replication schema is created, and the corresponding replica tables are created for each model. This step also generates “extract”, “data pump” and “replicat” parameter files for each model. These parameter files control the ongoing synchronization. The “extract” and “data pump” parameter files are deployed on the source and the “replicat” is deployed on the target.
5. **Performing Initial Sync:** Once these parameter files have been deployed successfully, the initial synchronization is performed. The data volume to be synced up depends on how long the source product instance has been active at your location. This sync up is performed by executing an Oracle Data Integrator scenario, which internally utilizes Oracle Data Pump to copy data from the source database to the replica database.
6. **Generating Metadata Driven Views:** Now, that the replicated tables and the source instance specific metadata has been imported, the metadata-driven views are generated in the replication schema. These views are the sources for the Oracle Data Integrator interfaces that load the dimensions and facts.
7. **Enabling Jobs:** At this point of time, the Oracle Utilities Analytics has been configured to work with the source instance to load the associated star schemas. An administrator can now utilize the Oracle Utilities Analytics Administration user interface to enable jobs, proceed with loading the star schemas and viewing the loaded data in the Dashboards.

Replication

This section describes about replication and DDL support, including:

- [Replication Patterns](#)
- [Structure of the Replicated Table](#)
- [DDL Support](#)

Replication Patterns

The following patterns are utilized when replicating tables from the source. These patterns are based on how these tables are being utilized by the source system and which target entities are loaded from these tables.

Pattern	Description	Approach
Overwrite	The historical changes need not to be tracked for these entities. These are typically the transactional or the master data used to populate the Type 1 dimensions.	<ul style="list-style-type: none">Any changes in the source are updated on the target replication table.The type of transaction is marked on the replicated table.The transaction timestamp is marked on the replicated table.The rows deleted on the source are not marked as deleted on the replication layer.
History	The source does not maintain history of changes, but in business intelligence, these tables are required to be tracked for populating type 2 dimensions, where each change needs to be treated as a new historical item.	<ul style="list-style-type: none">Every change is treated as a new historical item.The previous row for the same key must be marked with an end date.The latest record is marked as current.
Effective-Dated	The source maintains history by using an effective date column, or both the effective start and end dates. If the end date is not defined, it is calculated for better responsiveness of the data load processes.	<ul style="list-style-type: none">A change in the source is updated as is.The effective end date is calculated.The latest record is marked as current.
Key Tables	Oracle Utilities Analytics Framework product stores the keys for the primary tables of a maintenance object. This data is not deleted when data from the base tables is archived or purged. This is required to differentiate between a deletion on the source system v/s an archival on the source system.	The table consists of key columns only and is replicated as is with the timestamp and the source system change number.

Pattern	Description	Approach
SCD2	The SCD2 behavior requires the usage of columns to store the effective start and the effective end date of a historical record. This may also include the optional column to identify the current row.	In Oracle Utilities Network Management System source system, some of the tables follow the SCD2 behavior utilizing columns birth and death to store the effective start and effective end dates.

Structure of the Replicated Table

A replicated table has all the columns of the source table and the following additional attributes depending on the categorization. The column order of the replicated table matches the column order of the source table and any additional columns added will be after these columns.

The table mentioned below shows the different patterns of replication and the additional columns, which are present in each of the patterns.

Column	Overwrite	Effective Dated	SCD2	History	Key
All Source Columns	Yes	Yes	Yes	Yes	Yes
Slicing Timestamp	Yes	Yes	Yes	Yes	Yes
System Change Number	Yes	Yes	Yes	Yes	Yes
Last Change Timestamp	Yes	Yes	Yes	Yes	
Operation Indicator	Yes	Yes	Yes	Yes	
Cleansed Indicator	Yes	Yes	Yes	Yes	
Journal Effective Start Timestamp				Yes	
Journal Effective end Timestamp				Yes	
Effective End Date		Yes			
Current Indicator		Yes	Yes	Yes	

DDL Support

Any product can undergo data structure changes as a result of upgrades or customizations. The listed below are some possible edge application upgrade scenarios and the expected behavior on Analytics. The DDL replication feature of Oracle GoldenGate is utilized to support structural changes.

Category	Change	Expected Behavior
Structure change	The table/column is renamed or dropped.	<p>The DDL replication propagates the DDL change on the source to the target.</p> <ul style="list-style-type: none">• The renaming of a table is not supported in the DDL replication. However, rename of a column is supported in the DDL replication.• The dropping of a table is also not supported, but dropping of a column is handled by the DDL replication.• In case a table is renamed or a column is renamed, then views referring to the original tables and column are invalidated and target fact dimension loads fail. This requires manual steps to resolve the table/column rename.

Category	Change	Expected Behavior
Structure change	Column(s) added	<p>The DDL replication adds the new column(s) before the JRN columns in the replication table and no components are impacted. Oracle GoldenGate processes continues to run and the new columns store values updated or inserted at the source. The restrictions are given below:</p> <ul style="list-style-type: none"> • If a columns is added to a HIST type table, then any subsequent updates on the added column are not captured if the update does not include a key column filter in the where clause. • The XMLTYPE columns are not supported as DBMS_REDEFINITION do not support online redefinition of tables with binary XML columns. • Every DDL replication that is applied on the replication layer via Oracle GoldenGate causes the replicat process to be stopped. This is required to prevent any data loss during the time that the replicated table is redefined. After every DDL replication, the replicat process needs to be started again.
Structure change	Table(s) added.	No impact on Analytics.
Structure change	Column data type changed.	The corresponding change is applied on the replication layer and the new data is processed without interruptions.

Category	Change	Expected Behavior
Content change	CLOB xml structure change.	<p>Assuming that this node <code><isResidentialCustomer>true</isResidentialCustomer></code> is changed to <code><customerType>residential</customerType></code>.</p> <p>Post this change:</p> <ul style="list-style-type: none"> • None of the interfaces fails. • The logic to identify customer type as residential now returns null and all customers are now tagged as non-residential. • The data prior to this change used to have old xml structure and new data had the new structure. • The data correction is required on the replication layer to sync the structures to match new structure. This may require a change on the Analytics interfaces.
Content change	New/modified characteristic types	No impact on Analytics. The new characteristic types might need to be configured in the job parameters.

Category	Change	Expected Behavior
Content change	Data correction patch	<p>Post an upgrade where table structures have been changed, it is possible that a data correction patch is provided to convert the older version data to the newer version.</p> <ul style="list-style-type: none"> For the tables where you do not maintain history in the replication layer, this data correction may automatically propagate to the replication layer depending on the type of structural change. For the tables where you maintain history, this causes a new row to be created on the target for each of the impacted tables. The older data still have the old data and require a cleanup on the replication layer to fix the data appropriately. This potentially creates a spike in data volume to be reprocessed for the dimensions and the facts. Depending on the functionality impact of the data correction, facts/dimensions may need to be reprocessed completely.
Version	Multiple instances on different versions.	Multiple instances of the same edge application product on the different versions are not supported.

Default Keys in Dimensions

All the dimensions have two default rows:

Primary Key	Default Values	Data Source Indicator	Remarks
0	***	NULL	<p>This dimensional record is used as a reference for empty foreign key relation on a fact. If a particular fact record does not have any relation to a particular dimension, then the fact record is mapped to the 0th key in the dimension.</p> <p>For example, if no meter is installed at a service point in Oracle Utilities Meter Data Management, the service point fact record in the Service Point Fact points to the 0th key in the meter device dimension.</p>

Primary Key	Default Values	Data Source Indicator	Remarks
-99	-N/A-	NULL	<p>This dimensional record is used as a reference for missing foreign key relation on a fact. This scenario arises when a particular fact record does have reference to a particular dimension as per the source data, but the referenced value itself is not yet available in the dimension. In such cases, the fact record is mapped to the -99 key in the dimension.</p> <p>For example, consider that a meter has been installed at a service point in Oracle Utilities Meter Data Management, but a record for the meter has not yet been loaded in the meter device dimension. The service point fact record in the Service Point Fact points to the -99 key in the meter device dimension.</p> <p>This means that the ELT administrator should look out for these -99 records in the fact and double check why the dimension record is not loaded. There can be two reasons:</p> <ol style="list-style-type: none"> 1. This can be a classic case of late arriving dimension, wherein the fact record is attempted to be loaded even before a referenced dimension record was not yet loaded. In such cases, the ELT framework automatically tries to map the fact record again. Whenever the dimension record gets loaded, the fact record would get automatically corrected. 2. There can be other cases such as the source data reference for the dimension being incorrect/corrupt data.

Referencing Foreign Keys

- The Foreign Key (FK) value is set to 0 for any rows with NULL values for “Ref Natural Key”.
- The FK value is set to -99 for any rows with some value for “Ref Natural Key”, but no corresponding record in the dimension. The error reprocessing automatically binds these rows with the correct rows in the dimension, once the dimension value has arrived.

Metadata Model

The objectives of the metadata model are:

- To maintain configuration details for each source system.
- To maintain configuration of each instance of the source system, where multiple instances of the same source system have been implemented.
- To maintain configuration for each of the target entities, reusing interfaces and packages for multiple instances.

-
- To manage dependencies.
 - To manage multiple Data Source Indicators (DSIs) allowing external references at the instance level and also at granularity of individual rows of the source data.

You should not change any of the metadata entries that are delivered out of box in these metadata tables. Wherever your inputs are needed/allowed, the Oracle Utilities Analytics Administration Tool supplied with the Oracle Utilities Analytics product should be used to provide inputs. Other than this method, any other means to change the metadata is strictly prohibited.

This section describes the metadata model utilized by the implementation, including:

- [Products](#)
- [Instances](#)
- [Global Configuration](#)
- [Entities](#)
- [Job Configuration](#)
- [Objects](#)
- [Mapped Objects](#)
- [Source Tables](#)
- [Job Executions](#)
- [Dependencies](#)
- [DSI Mapping](#)
- [GoldenGate Configuration](#)
- [GoldenGate Checkpoint](#)
- [Extract Parameters](#)
- [Storage Configuration](#)

Products

This table is utilized to store a list of products that are supported by the implementations. This table is used to identify which products are supported, and to identify the Oracle Data Integrator scenario to be used to import source metadata and configurations.

Instances

You may have multiple instances of the same product that you want to integrate with the BI solution. The instance object represents each instance of the same product that can be utilized as a source. The objective is to enable development of a single interface that can be utilized across multiple instances.

The instances differ by the source database connections and possibly by the differing configurations. The Data Source Indicator (DSI) is unique across instances and products as this allows ability to identify the source of the data in the data warehouse.

Global Configuration

This table is utilized to store the configuration name value pairs that are applicable on the product. These configurations are utilized to control the execution modes for interfaces.

Entities

This table holds the configurations for the target entities in the data warehouse. The target entities can be dimensions, facts or materialized views. The same target entity may be loaded from multiple source instances. Configuration that is common across multiple instances is stored in this table.

Job Configuration

This table is used to provide configuration for the package (the executable logic) for populating the target entities for each instance. This table stores the current state information and an override package for use when the logic for one instance differs from the logic used for the other instances. The override package name can be used to provide a custom scenario in case you want to deviate from the logic provided.

Objects

This table is utilized to store all objects and configurations that are common across all objects. The objects can be the source tables that need to be replicated, the database objects that need to be generated or source Maintenance Objects (MOs) or Business Objects (BOs). Any object that is created during installation or when a source is attached is listed in this table along with configuration associated with the object.

Mapped Objects

This table contains a list of all source objects mapped to a target objects. This is the primary data entry point mapping a source MO/Table with the target entity. This is used to map views to be generated for the target entities. The entries are identified based on object type flag. Below are the possible values and their usage.

Object Type flag	Description
MO	Maintenance Object: The value in the source object field is identified as OUAF Maintenance object based on this flag.
TBL	Table: The value in the source object field is identified as NMS source table based on this flag
PRVW	Primary View: The value in the source object field is identified as primary view to load the data into the target object based on this flag
MVDP	Materialized View Dependency: Based on this flag, it is assumed that the value in the source object column is a dimension or fact entities and the value in the target object is a materialized view. This also indicates that the materialized view load should depend on the load of the entity specified in the source object column.
DMDP	Dimension Dependency: Based on this flag, it is assumed that the value in the source object column is a dimension entity and the value in the target object is a fact entity. This also indicates that the fact load should depend on the dimension load.

Source Tables

This table contains configurations controlling which tables are replicated and the mode of replication.

Job Executions

This table is used for tracking the execution of the ELT processes. An entry is created for each execution. Some attributes are populated from the SNP_SESSION table which is used by Oracle Data Integrator to track sessions.

Dependencies

This table is used to map a job to the jobs it is dependent on. The dependent objects are only executed up to the minimum sync timestamp of all the dependencies.

DSI Mapping

A source system may be integrated with other source systems. When pulling information from multiple source instances, it is possible that the data in one instance refers to the master data from another instance. This table allows such cross references to be configurable. By default, the configuration is at an instance level that is a product and instance number level. However, more fine grained control is possible by specifying the entity for which the configuration should be applied.

GoldenGate Configuration

This table is used for generating the Oracle GoldenGate scripts for initial as well as ongoing sync. This stores the database specific information of the target and source system.

GoldenGate Checkpoint

This table is used by Oracle GoldenGate replication process to track its processing activities. This table is used by the scheduler process to ensure that the warehouse loading tasks only process data that has been synced. The structure is controlled by Oracle GoldenGate.

Extract Parameters

This table stores configurations that control the extraction of data. Each source system may have different codes used for the same purpose. This table allows you to configure the codes so that the extraction routine can utilize it appropriately.

Storage Configuration

This table is utilized to store configuration for allocation of table spaces while creating objects within the knowledge models. This table has the default configurations at the pre-loaded “Global” context level.

Configuration Views

This section covers the following views:

- [Context Configuration](#)
- [Replication Configuration](#)
- [Job Configuration](#)

-
- [Storage Configuration](#)
 - [Dependencies](#)
 - [Jobs](#)
 - [Ready for Load](#)
 - [Wait Reasons](#)

Context Configuration

The view associates the product instance with the logical and physical schemas defined in Oracle Data Integrator.

Replication Configuration

The replication configuration view is a summary of the settings that are used to control the tables to be replicated.

Job Configuration

The storage configuration view shows the tablespaces configured for different object types for each logical schema. The configuration can be common or can be granular at logical schema level. This view breaks down the configurations at the lowest granularity and provides easy to access all attributes of storage configuration handling the complexities of identifying context and logical schema specific configurations that override the global configurations.

Storage Configuration

All the attributes of storage configuration handling the complexities of identifying context and logical schema specific configurations that override the global configurations.

Dependencies

The Oracle GoldenGate processes can be tracked via the Oracle GoldenGate Checkpoint table. A source instance can have multiple Oracle GoldenGate processes and each process consists of multiple tables. A target entity can be loaded only up to the maximum sync timestamp of all the Oracle GoldenGate processes associated with the context in which it is executing.

A fact load is dependent on the load of type 2 dimensions associated with the fact. The fact associated with a type 2 dimension cannot be loaded beyond the time that the dimension has been synced.

A materialized view becomes stale if any of the underlying tables is modified. Hence, the materialized view depends on the execution of any of its underlying tables.

The view identifies the base object and the timestamp associated with its dependencies so that appropriate actions can be taken to avoid any data loss.

Jobs

This is a consolidated way of viewing all jobs, which are being executed or have completed. The view shows the slicing units, duration of execution and volumes being processed. This can be utilized for health checks and for monitoring the load processes.

Ready for Load

This view governs the selection of jobs for execution and controlling the pattern of execution. This is governed by the following rules.

- The entity should not be disabled.
- The configured scenario should exist.
- All dependencies should have run.
- The number of running/error instances of the job should be less than the maximum parallel executions allowed.
- If the instance has been configured as run once, then it should not execute once it has been successfully executed.
- If a job fails, it should be retried again until the maximum retries per day is reached. The interval between successive retries should be based on the configuration.
- The Oracle GoldenGate models comprising of the source tables used in the entity should have been synced up. In case the sync timestamps vary across multiple models, then the least common sync timestamp is used.
- The snapshot entities are executed on or after the snapshot end period.
- The schedule time can be used to stagger loads and distribute processing. A job is not executed until the current time crosses the scheduled time.

Wait Reasons

This view can be used to diagnose the reasons why a job is not being executed. This view applies the rules for scheduling and reverses the logic to obtain the reasons for not executing. The issues that can be identified by using this view are given below:

- An entity is deactivated.
- The Oracle GoldenGate replication is not configured for model {model name}.
- The configured scenario does not exist.
- This job is scheduled to run at {timestamp}.
- This one time load interface has already executed successfully.
- The reached max retries allowed per day.
- The approximately {X} minutes till next retry.
- The maximum number of parallel materialized views execution reached.
- Waiting for a change to occur in dependent entity.
- Waiting on a dependent entity to be loaded.
- Waiting for at least one minute of incremental data to arrive.
- The cut off time has been enabled. All the entities are loaded only up to {timestamp}.
- The temporary table is not reversed in the staging model.
- The staging table is not reversed in the staging model.
- The flexfields are not set for table.

Oracle Data Integrator Components

This section lists out the high level components in Oracle Data Integrator:

- [Project](#)
- [Model Folders](#)
- [Scenario Folders](#)
- [Logical Schemas](#)
- [Connections](#)
- [Contexts](#)

Project

The out of the box implementation installs the project “Oracle Utilities BP”. This project contains all the out of the box code for loading facts, dimensions and materialized views. This project and all its underlying objects may be modified in subsequent releases.

All the customizations should be done in a separate project.

Note: Any custom objects if created within this project may be deleted during a subsequent release.

Model Folders

The Oracle Utilities Analytics product comes with a set of predefined model folders and models.

Name	Purpose
Source	The model folder is created when the first source product instance is attached to Oracle Utilities Analytics. This folder contains all models associated with a source instance. These models are also created when a source instance is attached to Oracle Utilities Analytics.
Pre-Staging	This folder contains models for all the replicated table objects and the views generated in the replicated schema.
Target	This folder contains the models for the star schema facts and dimensions. These are further grouped together based on the source system from which it is populated. The shared dimensions, which are populated from multiple edge applications, are grouped together.
Metadata	This contains the models for all metadata tables.

The objects within these models should not be modified. All custom models should be created under a separate custom model folder.

Note: Any custom objects if created within the Oracle Utilities Analytics models may be deleted during a subsequent release.

Scenario Folders

All scenarios that are delivered are grouped into scenario folders based on the functionality and usage. The following scenario folders are available

Name	Purpose
Framework	All the framework components are included in this folder.
Accelerators	Implementation accelerators for each source product framework are available in this folder. These scenarios are executed automatically during the steps for attaching a source instance.
Materialized Views	All the scenarios that are used to create or refresh materialized views are organized under this folder. These are automatically executed based on the configured schedule.
Replication	All scenarios that are used to create views in the replication schemas are organized under this folder. These are executed when attaching a new source instance to Oracle Utilities Analytics or when the view generator process is executed.
Star Schema	All the scenarios that are used to load the star schema tables are organized within this folder. These are automatically executed based on the configured schedule.

Note: Any custom objects if created within these folders may be deleted during a subsequent release.

Logical Schemas

The logical connections as defined in Oracle Data Integrator are:

Logical Name	Purpose
Source	The logical connection associated with the source data. This contains all the tables that will be replicated.
Journal	The journal schema is utilized as the work schema where objects can be created for processing. This is also the user connecting to the source system and managing the GoldenGate replication.
Replication	The logical connection associated with the replicated area where the data is replicated and history is maintained before the warehouse load process is initiated.
Staging	The logical connection associated with the staging area where the semi-persistent staging tables reside.
Target	The logical connection for the final target star schemas.
Repository	The logical connection to the work repository. This is used for internal objects associated with the Oracle Utilities Analytics processes.
Master	The Oracle Data Integrator master repository is maintained in a separate schema associated with the master logical schema.
Metadata	The metadata configurations utilized to manage the solution are stored in this data area.

Logical Name	Purpose
Configuration	This is a flat file logical schema used for import and export of the primary metadata.

Connections

During installation a single connection named “B1” is created to the target database. This connection is used for all data load processes and it connects using the metadata user. Subsequently when attaching a source product instance, the additional connections are created. These connections are named based on the source product code and instance number.

Contexts

The default context “Global” is the only context created during install and this is automatically mapped to the appropriate logical schemas during installation. Additional contexts are created when adding source instances to Oracle Utilities Analytics. The contexts map the physical schemas to the logical schemas based on the source connection and the mappings defined in the “Global” context for the target schemas. The jobs for a particular source are always executed in the context defined for that source instance.

Note: The contextual mapping should not be modified manually.

Chapter 3

Oracle Utilities Analytics Dashboards

A number of pre-built analytics are delivered with Oracle Utilities Analytics. They cater to a rich set of KPIs for each source product whose data can be extracted into the data warehouse. The Dashboards components in Oracle Utilities Analytics are described in the following sections:

- [Dashboard Catalogs](#)
- [Oracle Business Intelligence Repository](#)

Dashboard Catalogs

For each source product multiple dashboards are delivered with Oracle Utilities Analytics. The dashboards are grouped in logical divisions and are a part of a Presentation Catalog. Below is mapping between the various source products and the catalogs available for them.

Source Product	Presentation Catalog in Oracle Utilities Analytics
Oracle Utilities Customer Care & Billing	Credit & Collection Analytics Customer Analytics Revenue Analytics Exception Analytics
Oracle Utilities Meter Data Management	Meter Data Analytics
Oracle Utilities Mobile Workforce Management	Mobile Workforce Analytics
Oracle Utilities Network Management System	Outage Analytics Distribution Analytics
Oracle Utilities Operational Device Management	Operational Device Analytics
Oracle Utilities Work and Asset Management	Work & Asset Analytics

The details about each dashboard are described in detail in the Metric Reference Guide for each of the Analytics above. The details about the star schema tables underneath the dashboards are covered in the Data Mapping Guide for each source product. All dashboards delivered along with base product are supported via a dedicated materialized view layer. These materialized views ensure that data is pre-aggregated and pre-joined as per the report requirements so that data is instantly available for the reports. Data in the materialized views is refreshed via separate ELT jobs which would get triggered after an ELT job loads data in a star schema table.

Apart from these dashboards for various source products, a special Administration catalog is also provided. The dashboards in this catalog cater to the administration of the labels to be used in report and some map related configuration. These dashboards are covered in detail in [Chapter 5](#):

[Configuring Oracle Utilities Analytics](#), under the section [Configuring the Analytics](#). You can add more dashboards/reports to this or can choose to customize any of the existing ones. The details of the same are described in the *Oracle Utilities Analytics Developer's Guide*.

Oracle Business Intelligence Repository

Just like in any other Oracle Business Intelligence Repository file, you can find all three layers in the repository file for Oracle Utilities Analytics. This section describes what is available in each of these layers in Oracle Utilities Analytics.

- [Physical Layer](#)
- [Business Model and Mapping Layer](#)
- [Presentation Layer](#)

Physical Layer

The physical layer in the repository file for Oracle Utilities Analytics has the following physical data construct groups

- **Initialization Blocks** - This data construct is used as the source for the initialization block queries in the repository which are used to initialize the session variables used in various dashboards. There are two connection pools defined, one for the data warehouse and the other for the source database of Oracle Utilities Meter Data Management system, which hosts the aggregator materialized views for some of the dashboards in Meter Data Analytics
- **Materialized Views** - This data construct is used as the source for fetching the materialized view data from the data warehouse.
- **MDM Database** - This data construct is used to connect to the source database of Oracle Utilities Meter Data Management system, to fetch the data from the aggregator materialized views for some of the dashboards in Meter Data Analytics
- **ODIDEV** - This data construct defines a special connection pool with write access to certain set of tables in the data warehouse. These are the tables used in the Administration group of dashboards, where Oracle Business Intelligence Enterprise Edition's writeback feature is used to update data back into the warehouse.
- **Oracle Data Warehouse** - This data construct defines the connection pool to fetch the data from the star schema tables from the data warehouse.

Business Model and Mapping Layer

The BMM layer in the repository file for Oracle Utilities Analytics has the following groups

- **Core** - This group defines the data model of all star schemas in the data warehouse. The dimensional hierarchies are also defined here. Logical columns needed for calculations on top of the fact measures or the dimension attributes would be defined here. The materialized view mapping is also done here. Each fact and dimension has the logical mapping to one or more materialized views, which were designed for it. This ensures that the reports based on the respective fact or dimension goes to the materialized view first in the same order as the logical sources are defined
- **OUA** - This group defines the data model of the special set of tables where writeback capabilities are enabled.

Presentation Layer

The subject areas in the **Presentation Layer** are defined as per the star schema in the data warehouse. In most of the subject areas, there is one fact along with all of its associated dimensions. In certain cases multiple facts might be present to support certain complex report requirements. Each subject area name starts with a three letter prefix of the source product it's based upon. This is the layer which inserts the label mechanism in the product, which provides the multi language support on the Oracle Utilities Analytics dashboards. Each presentation layer field is associated with a session variable name. These session variables are dynamically mapped to label descriptions of an appropriate language based on the logged in your preference.

Chapter 4

Oracle Utilities Analytics Administration Tool

The Oracle Data Integrator based Extract, Load and Transform (ELT) architecture is a metadata driven framework that allows you to configure the out of the box ELT jobs, as well as making it easy to integrate custom ELT jobs to Oracle Utilities Analytics. The out-of-the-box ELT jobs are set up during the installation of the product. However, you can maintain some of the parameters using Oracle Utilities Analytics Administration Tool. This tool is based on Oracle Application Express (APEX) platform. Below are the tasks that can be performed using Oracle Utilities Analytics Administration Tool:

- [Configuring Global Settings](#)
- [Configuring Oracle GoldenGate](#)
- [Configuring Product Instance](#)
- [Configuring Storage](#)
- [Configuring Source Table](#)
- [Configuring Target Entities](#)
- [Configuring Jobs](#)
- [Customizing the Column Properties of Type 2 Slowly Changing Dimension](#)
- [ETL Replication View](#)
- [NMS Configuration](#)
- [ETL Job Execution](#)

Configuring Global Settings

This table is used to capture the global settings needed for Oracle Utilities Analytics:

Attribute	Purpose
Product	A reference to the product code, identifying the product for this instance. This could be one of the configured source products of the base Oracle Utilities Analytics product itself.
Instance Number	The unique number of the configured source product instance. For the configuration settings of the base Oracle Utilities Analytics product, this value is set to null.
Description	This column describes the purpose of the specific global configuration setting.
Value	The configuration value that you can update based on their setup.
Data Type	The data type of the configuration value that a user is expected to enter.
Data Format	The specific format of the data that a user is expected to enter the configuration value in.

The **Global Configuration** page in the Administration Tool allows you to enter a configuration value for each of the configuration setting. As a part of the Oracle Utilities Analytics installation process, the global configuration records are generated for the configured source application and the target Oracle Utilities Analytics application. You can then use this page to change the configurations.

You are expected to enter the configuration value strictly in the specified data format only. Entering value in other formats results in errors during the ETL processing.

Global Configuration			
Q		Go	Actions
			Edit
Product	Instance Number	Description	Value
Customer Care and Billing	1	Product Extract Date - The date from which all ETL jobs for the specific product instance will be configured to start the initial load	01-Jan-2001
B1	(None)	Number of days for which the logs would be retained	11
B1	(None)	Maximum number of MV interfaces which can run in parallel	6
B1	(None)	Database Edition Type - Mention the type of the Oracle Database, whether it is Standard Edition(SE) or Enterprise Edition(EE)	EE
B1	(None)	Enable Partitioning - If a partitioning license has been purchased from Oracle, set this to Y	N
B1	(None)	Cut Off Time - The time of the day upto which all changes from the source will be loaded by the ETL jobs. Any changes after this will not be loaded on the same day. Set this to default value (i.e. 23:59:59) to keep loading data continuously.	23:59:59
B1	(None)	Global Extract Date - The date from which all ETL jobs will be configured to start the initial load	01-Jan-2001

Enable Partitioning: *Partitioning* feature is an additional license on top of the Oracle Enterprise Edition. Set this property to 'Y' if you have licensed the partitioning features. The ETL processes utilize the partitioning functionality. Setting this flag to 'Y' causes the intermediate staging tables to be partitioned by the job number and significantly improve the parallel execution performance of individual entities.

Some features such as database parallelism are not available in the Standard Edition. Setting this property ensures that the ELT processes appropriately utilize the underlying database features for processing the loads.

Configuring Oracle GoldenGate

Oracle Utilities Analytics uses Oracle GoldenGate to capture the changed data. This table is used by the Oracle GoldenGate replication process to track its processing activities. This information is used by the scheduler process to ensure that the warehouse loading tasks only the processed data that has been synced.

Attribute	Purpose
Configuration Code	The code is either be the context code identifying the instance uniquely or B1 for the target data warehouse.
Manager Port	The port that the Oracle GoldenGate Manager Port utilizes.
Dynamic Port Range	A port range that the Oracle GoldenGate process utilizes.
Algorithm	The encryption algorithm to be used. For example, BLOWFISH.
Encryption Key	The name of the key used for encryption/decryption.
GoldenGate Home	The home folder for Oracle GoldenGate.
Shared Secret	The Shared Secret is the shared public key. The Encryption mechanism utilizes a public key private key pair.
Database Host	The host name of the database.
Database Port	The port name for the database.
Database ID	The system identifier for the database.
Database Home	The database home folder where the database software is installed.

Oracle GoldenGate Configuration page in the Administration Tool allows you to add, edit, and delete records. As part of the Oracle Utilities Analytics installation process, Oracle GoldenGate configuration records are generated for the configured source application and the target Oracle Utilities Analytics application. You can then use this page to change the configurations.

Maintain Golden Gate Configuration

Cancel Delete Save

Configuration Code * ODM1

Manager Port *

Dynamic Port Range Start *

Dynamic Port Range End *

Algorithm *

Encryption Key *

Shared Secret *

Golden Gate Home *

Database Connection Details

Database Host *

Database Port *

Database SID *

Database Home *

Configuring Product Instance

The Product Instance represents a specific instance of a source application that can be configured as a source for Oracle Utilities Analytics solution. A record is created for every instance of the product for which ELT is setup to extract data. Oracle Utilities Analytics delivers a record for the product which uses Oracle Data Integrator for the ELT. You may have multiple instances of the same source application to integrate with the Oracle Utilities Analytics data warehouse.

The following attributes are available on the Product Instance metadata table.

Attribute	Purpose
Product Code	A reference to the product code, identifying the product for this instance.
Instance Number	A unique number starting with 1 to uniquely identify the instance.
Context Code	A unique code comprised of the product code and instance number which is used to identify the connections in Oracle Data Integrator. Due to limitations imposed by Oracle GoldenGate (used for replication), the context code cannot be more than 5 characters in length.
Change Data Capture	A flag identifying the methodology to be used for replicating source tables. The default is Oracle GoldenGate.
Drill Back URL	A URL to be used to allow you to drill back to the source system from the analytics.
Time Zone Code	The time zone of the source product instance. This allows the interfaces to be built so that dependencies on job execution for multiple time zones are handled correctly.
Currency Code	The currency used in the source product instance. Multiple currencies are not supported.
Language Code	The primary language supported by the data warehouse. This is used to filter language specific data in the data warehouse.
Data Source Indicator	A unique value representing the instance. For Oracle Utilities Application Framework products, this is the environment ID of the source instance.
Source Instance Version	The current version number of the source product instance. This can be used to validate whether the version is supported by the data warehouse, or provide suggestions in case upgrades, or patches need to be applied to enable support for the source product.

The **Product Instance** page in the Administration Tool allows you to maintain the existing records. To add a new product instance, use the script delivered as a part of the installation package

Maintain Product Instance

Cancel Delete Save

Source Product * Oracle Operational Device Management

Instance Number * 1

Context Code ODM1

Change Data Capture * Oracle Golden Gate ▼

Drillback URL *
Note : Enter the URL to access the source application in the following format http://(host-name):(port)/ouaf/cis.jsp

Information Retrieved from Source

Time Zone Offset
Note : Specify in the following format 'HH24:MI:SS'

Currency

Language

Data Source Indicator 273498

Product Version V2.0.1.1.0

For more details, refer to the **Configuring the Source** sub-section under the **Upgrading Oracle Utilities Analytics Based ELT Component** section in the **Chapter 5: Upgrading Oracle Utilities Analytics** of *Oracle Utilities Analytics Installation Guide*.

Configuring Storage

In this release, an additional storage configuration is provided that enables you to manage the tablespaces utilized for the database objects that are created during the execution of the Oracle Data Integrator jobs.

The following attributes are available:

Attribute	Purpose
Context Code	Identifies the context for which this is applicable. A value of “GLOBAL” is allocable across all the contexts unless overridden by an entry for the specific context.
Logical Schema	The logical Schema for which the configuration is applicable. A empty value indicates that this is applicable to all logical schemas.
Object Type	The type of object for which the configuration is applicable. This can be “Table, Index, Partition”.
Sequence	The default is 0. This is applicable only for partitions. This can be used to map multiple table spaces in a round-robin scheduling while creating the new partitions within the same table.
Tablespace Name	The name of the table space to be associated with the object.

The **Storage Configuration** page in the Administration Tool allows you to maintain the existing records related to configuration. Records are delivered out of the box for the star schemas tables delivered with Oracle Utilities Analytics. You cannot delete base product owned records. Create new records for any entity that you want to add to the ELT. For owned records, you have full access to edit and delete them.

Storage Configuration							
Storage Config ID	Context Code	Logical Schema	Object Type	Sequence	Tablespace Name	Owner Flg	
303	CCB1	Master	Index	4	TEST_01	B1	
198	CCB1	Master	Index	5	ts_003	CM	
149	CCB1	Master	Index	22	20	B1	
146	CCB1	Master	Index	99	20	B1	
209	CCB1	Master	Index	2	tablespace1	CM	
152	CCB1	Master	Index	77	ss	B1	
141	CCB1	Master	Index	11	20	B1	
32	NMS1	Staging	Index	0	CISTS_04	B1	
31	CCB2	Staging	Index	0	STAGE_03	B1	
30	CCB1	Staging	Index	0	STAGE_02	B1	
26	NMS1	Metadata	Index	0	CISTS_04	B1	
25	CCB2	Metadata	Index	0	CISTS_03	B1	
24	CCB1	Metadata	Index	0	CISTS_02	B1	
10	NMS1	Replication	Index	0	REPOS_04	B1	
9	CCB2	Replication	Index	0	REPOS_03	B1	
















Configuring Source Table

This table contains configurations controlling, such as the source tables to be replicated and the mode of replication.

Attribute	Purpose
Source Product	A reference to the product code, identifying the product for this instance.
Table Name	The name of the source table. There is a corresponding entry in the objects table also.
History Type	The following scenarios exist: <ul style="list-style-type: none"> Source system tracks history using an effective date column No history in source but needs history for the type 2 dimensions No history required
Effective Date Column Name	The column name used for storing the effective dates in the source.
Characteristic Entity	Identifies the char entity flag in source Oracle Utilities Application Framework application.

Attribute	Purpose
Base Replication	Controls whether the table is required to be replicated for target entity load.
Group Number	Defines the grouping of various tables for a source product. These group numbers is used to define the various GoldenGate models that are setup to extract data from the source system.
Custom Replication	Extension for customizations. The additional tables to be marked for replication.
Purge Enabled	Controls whether the replicated table should be purged or not.
Replication Retention Days	The number of days the data should be retained in the replication layer.
Owner Flag	Indicates whether the record is owned by the base product (B1) or you (CM).

The **Source Table** page in the Administration Tool allows you to add, edit, and delete records. For this table, records are delivered out of the box for the source tables that are configured to be part of the out-of-the-box ELT. For these records, you can edit only the selected fields. You cannot delete base product owned records. However, for owned records, you have full access to edit and delete them.

Source Table							
Source Table Id	Source Product	Table Name	History Type	Base Replication	Custom Replication	Owner	
 111	Customer Care and Billing	CI_DD_EVT_TYPE	Override Values	Y	N	B1	
 108	Customer Care and Billing	CI_CC_LOG	Override Values	N	N	B1	
 100	Customer Care and Billing	CI_COLL_EVT_TYP	Override Values	Y	N	B1	
 99	Customer Care and Billing	CI_COLL_EVT_CC	Override Values	N	N	B1	
 93	Customer Care and Billing	CI_BILL	Override Values	Y	N	B1	
 90	Customer Care and Billing	CI_WO_EVT_TYP_L	Override Values	Y	N	B1	
 89	Customer Care and Billing	CI_WO_EVT_CC	Override Values	N	N	B1	
 87	Customer Care and Billing	CI_SEV_EVT_DEP	Override Values	N	N	B1	
 86	Customer Care and Billing	CI_SQI	Override Values	N	N	B1	
 85	Customer Care and Billing	CI_CUT_EVT_DEP	Override Values	N	N	B1	
 84	Customer Care and Billing	CI_CASE_CHAR	Override Values	Y	N	B1	
 83	Customer Care and Billing	CI_CASE_LOG	Override Values	Y	N	B1	
 82	Customer Care and Billing	CI_PP	Override Values	Y	N	B1	
 80	Customer Care and Billing	CI_TNDR_SRCE_L	Override Values	Y	N	B1	
 79	Customer Care and Billing	CI_TNDR_SRCE	Override Values	Y	N	B1	

Configuring Target Entities

This table holds the configurations for target entities and their attributes in the data warehouse. The target entities can be dimensions, or facts. The same target entity may be loaded from multiple source instances. Configuration that is common across multiple instances is stored in this table.

Attribute	Purpose
Entity Name	The name of the entity, which needs to be scheduled for loading into the target.
Entity Type	The type of the entity supported by Oracle Utilities Analytics: <ul style="list-style-type: none">• Slowly changing dimension type 1• Slowly changing dimension type 2• Accumulation facts• Snapshot facts• Materialized views
Maximum Parallel Executions	To efficiently load data, it may be necessary to execute multiple instances, each instance working on a different data set. This attribute controls how many parallel executions can be spawned for a single entity load.
Maximum Retries	This attribute controls the maximum number of retry attempts in a day. Once this limit is reached and the load is still failing, it is retried next day.
Retry Interval	The base architecture has been designed for automatic retries. In case of failures, jobs are retried and this attribute controls the interval between successive retries. Default is 30 minutes, but can be configured as per the requirement.
Scheduling Type	Three different modes of schedules are supported: <ul style="list-style-type: none">• Daily Incremental Load: Loads are executed as soon as data is available for load.• Near Real Time Load: Loads are executed within a configured interval. This is for future use.• One Time Load: A load is executed only once.
Scheduling Interval	This is applicable for NRT loads only and the value controls the frequency of execution of the load process.
Scheduling Time	A job is executed only after the scheduled time each day. The default is 00:00. This can be changed as per the requirement.

Attribute	Purpose
Slice Duration Type	<p>In any data warehouse, the basic challenge is to get the data loaded quickly and efficiently whether it is the initial load or the incremental load. Data volumes are high during initial load and during incremental loads, the volumes are considerably smaller. A slice is a volume of data bound within duration of time. Different objects have differing data distribution and load processing requirements. This attribute controls the duration between two slices. The following slicing intervals are supported:</p> <ul style="list-style-type: none"> • Day(s) • Week(s) • Month(s) • Quarter(s) • Year(s) • Hour(s) • Minute(s)
Slice Duration	The number identifying the slice duration based on the slice duration type.
ODI Package Name	The name of the Oracle Data Integrator scenario that should be executed.
Staging Retention Days	The number of days to retain data in the staging tables.
Owner Flag	Indicates whether the record is owned by the base product (B1) or you (CM).

The **Target Entity** page in the Administration Tool allows you to add, edit, and delete records. Records are delivered out of the box for the star schemas tables delivered with Oracle Utilities Analytics. For these records, you can edit only select fields. You cannot delete base product owned records. Create new records for any entity that you want to add to the ELT. For owned records, you have full access to edit and delete them.

Target Entity							
Target Entity Id	Entity Name	Entity Type	Scheduling Time (HH:MI:SS)	Scheduling Interval (mins)	ODI Package Name	Staging Retention Days	Owner
103	CD_UOM	Dimension SCD1	00:00:00	1	B1_PKG_CD_UOM	1	B1
101	CD_UCOLPROC_STATUS	Dimension SCD1	00:00:00	0	B1_PKG_CD_UCOLPROC_STATUS	1	B1
100	CD_UCOLEVT_TYPE	Dimension SCD1	00:00:00	0	B1_PKG_CD_UCOLEVT_TYPE	1	B1
99	CD_TOU	Dimension SCD1	00:00:00	0	B1_PKG_CD_TOU	1	B1
98	CD_TNDR_TYPE	Dimension SCD1	00:00:00	0	B1_PKG_CD_TNDR_TYPE	1	B1
97	CD_TNDR_STATUS	Dimension SCD1	00:00:00	0	B1_PKG_CD_TNDR_STATUS	1	B1
96	CD_TNDR_SRCE	Dimension SCD1	00:00:00	0	B1_PKG_CD_TNDR_SRCE	1	B1
94	CD_TD_STATUS	Dimension SCD1	00:00:00	0	B1_PKG_CD_TD_STATUS	1	B1
93	CD_TD_SKILL	Dimension SCD1	00:00:00	0	B1_PKG_CD_TD_SKILL	1	B1
92	CD_TD_ROLE	Dimension SCD1	00:00:00	0	B1_PKG_CD_TD_ROLE	1	B1
91	CD_TD_PRIORITY	Dimension SCD1	00:00:00	0	B1_PKG_CD_TD_PRIORITY	1	B1
90	CD_TD	Dimension SCD1	00:00:00	0	B1_PKG_CD_TD	1	B1
89	CD_SQI	Dimension SCD1	00:00:00	0	B1_PKG_CD_SQI	1	B1
88	CD_SERVICE_HIST_TYPE	Dimension SCD1	00:00:00	0	B1_PKG_CD_SERVICE_HIST_TYPE	1	B1
87	CD_SA_STATUS	Dimension SCD1	00:00:00	0	B1_PKG_CD_SA_STATUS	1	B1

Configuring Jobs

The **Job Configuration** page in Oracle Utilities Analytics Administration user interface is used to provide configuration for the package (the executable logic) for populating the target entities for each instance.

Attribute	Purpose
Source Product	A reference to the product code, identifying the product for this instance.
Instance Number	To be populated for instance specific deviations. Ideally most data should be reusable across instances of the same product.
Target Entity	A reference to the entity.
Entity Active Flag	Indicates whether this ELT job should be picked up for execution by the Oracle Data Integrator Scheduler.
Slice Start Date/ Time	The starting timestamp for the next execution.
Initialize Flag	Controls the initialization of the entity for slicing.
Execution Sequence	Sequence of execution within the same group.
Last sync timestamp	Stores the timestamp of the most recent successful slice.
Overrides ODI Package Name	Overrides the package name defined at entity level. This is to be used in case the interface implementation for a specific instance differs from the common interface
User Exit Procedure	The procedure name, which can be used to load the user-defined columns in the entity.

The **Configuring Jobs** page in the Administration Tool allows you to add, edit, and delete records. Records are delivered out of the box for the star schemas tables delivered with Oracle Utilities Analytics. You cannot delete base product owned records. Create new records for any entity that you want to add to the ELT. For owned records, you have full access to edit and delete them.

Job Configuration Id	Source Product	Instance Number	Target Entity Name	Entity Active Flag	User Exit Procedure	Last Sync Date/Time
663	Operational Device Management	1	CF_SERVICE_HIST	N	-	-
662	Operational Device Management	1	CF_OPR_DEVICE_SNAP	N	-	-
661	Operational Device Management	1	CF_OPR_DEVICE	N	-	-
660	Operational Device Management	1	CF_ASSET_LOC	N	-	-
659	Operational Device Management	1	CD_UTIL_ASSET	N	-	-
658	Operational Device Management	1	CD_SERVICE_HIST_TYPE	N	-	-
657	Operational Device Management	1	CD_OPR_DEVICE	N	-	-
656	Operational Device Management	1	CD_LOCATION	N	-	-
655	Operational Device Management	1	CD_ASSET_INSTORE_AGE	N	-	-
654	Operational Device Management	1	CD_ASSET_INSTALL_AGE	N	-	-
653	Operational Device Management	1	CD_ASSET_INSP_STATUS	N	-	-
652	Operational Device Management	1	CD_ASSET_DISP	N	-	-
651	Operational Device Management	1	CD_ASSET_AGE	N	-	-
650	Operational Device Management	1	CD_ADDOR	N	-	-
649	Operational Device Management	1	B1_SERVICE_HIST_MOH_MV1	N	-	-

Maintaining Job Configuration

The **Job Configuration** page in the Administration Tool allows you to add, edit and delete records. As a part of the Oracle Utilities Analytics installation process, job configuration records are generated for the target entities available out of the box.

You can then use this page to change the configurations. When you click the **Add** button or click the **Edit** button for an existing job record, the **Maintain Job Configuration** page pops up where you can specify the desired configuration for the job.

Cancel Delete Save

Job Configuration Id: 663

Source Product * Operational Device Management

Instance Number * 1

Target Entity * 127

Scheduling Parameters

Entity Active Flag * No

Slice Start Date/Time * 01-JAN-2000 00:00:00
Note : Specify in the following format 'DD-MON-YYYY HH24:MI:SS'

Initialize Flag * Yes

Execution Sequence

Last Sync Date/Time

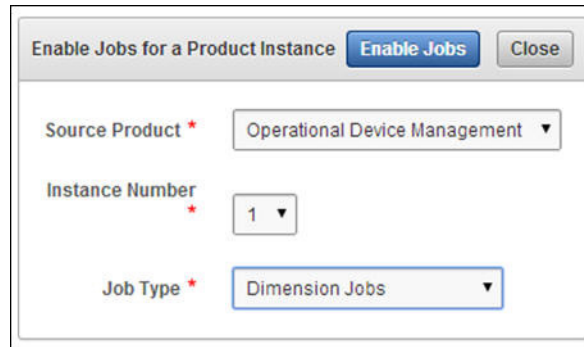
Customization Attributes

Override ODI Package Name

User Exit Procedure

Enabling Jobs

The jobs can be enabled from the **Job configuration** page on Oracle Utilities Analytics Administration user interface. The job type can be dimension jobs, dimension and fact jobs and materialized view jobs. Each job can be individually enabled by using the edit button. Alternately, jobs of a group can be enabled together using the “Enable Jobs” button. There are three options (Dimension Jobs, Facts, Facts and Materialized Views) for the job type.



The screenshot shows a dialog box titled "Enable Jobs for a Product Instance". It contains three dropdown menus: "Source Product" (Operational Device Management), "Instance Number" (1), and "Job Type" (Dimension Jobs). There are "Enable Jobs" and "Close" buttons at the top right.

Customizing the Column Properties of Type 2 Slowly Changing Dimension

A type 2 slowly changing dimension has the following column types:

- Natural Key
- SCD Start
- SCD End
- SCD Flag

Any change in any of the remaining columns is treated as a new history item for the same natural key. In the next release of Oracle Utilities Analytics, you will be able to decide which of the columns should not be considered for creating a new history item.

Using Oracle Data Integrator, this is done by changing the corresponding property of the column in the model of the target dimension. However, this change requires regeneration of the dimension load scenario. The customizations done by you will also get overridden by any future upgrades provided by Oracle.

You can control the columns where a change causes a history to be created. You can change the properties by accessing the Oracle Utilities Analytics Administration user interface. Once the configurations have been changed, the following script needs to be executed to ensure that the changes have been applied to the appropriate scenarios:

For UNIX:

Perform the following steps:

1. Navigate to <Install_Dir>/bin directory.
2. Initialize the environment with the `./splenviron.sh -e <envname>` command.
3. `cd $SPLEBASE/bin.`
4. Run `ksh ./updateODIMetadata.sh`

For Windows:

Perform the following steps:

1. Navigate to <Install_Dir>/bin directory.
2. Initialize the environment with the `./splenviron.cmd -e <envname>` command.

3. cd %SPLEBASE%\bin.
4. Run updateODIMetadata.cmd.

Important Note: Once the behavior of the SCD2 dimension is customized, the above-mentioned script needs to be executed after every upgrade to retain the customized behavior.

ETL Replication View

Every ETL job in the data warehouse is based on one or more Oracle Data Integrator interfaces to populate the target fact/dimension table. The starting point is a view defined for the ETL job in replication layer of the specific source product. This ETL replication view fetches the base values from the replication tables as a starting point for further transformations to be done by the Oracle Data Integrator interfaces.

In certain business scenarios, it might be necessary for you to override the definition of this ETL replication view provided out-of-the-box. For example, in Oracle Utilities Network Management System, a particular type of data may be stored in different tables based on the network model. The base view definitions supplied for Oracle Utilities Network Management System ETL are designed based on the OPAL network model. When you use a different model, you might have to change the definition of some of these views in the replication layer as per the network model you use.

To take a more specific example, the B1_D_ACCT_VW for the account dimension in the Oracle Utilities Network Management System star schema, tries to fetch the account related attributes from the CU_SERVICE_LOCATIONS table. If the model used by you, stores the account attributes in a different table, then the view definition needs to be changed. This page in the Oracle Utilities Analytics Administration tool allows you to do this. You need to define a custom Oracle Data Integrator interface to generate the view.

The view name and the column names in the view should be exactly as it was delivered out-of-the-box. However, the view definition can be changed as per your requirements.

A scenario should be generated for this custom Oracle Data Integrator interface and should be mentioned in this Oracle Utilities Analytics Administration Tool page against the appropriate view name, whose definition needs to be overridden.

Next time, when views are generated using the 'Run View Generator' script, the ETL framework generates the view based Oracle Data Integrator scenario provided by you.

Maintain ETL Replication Views

Maintain ETL Replication View Cancel Save


Object Id	961
View Name	B1_D_ACCT_VW
Source Product	Network Management System
View Type	Primary View

Maintain ETL Replication Views

Custom ODI Scenario Name

NMS Configuration


Before attaching a specific source system to Oracle Utilities Analytics, there are some configurations that need to be done for every source system for its ETL to function properly. These configurations including setting up of certain parameters to be referred to during the extraction or setting up bucket ranges for classifying the measures in a fact. Set up the configurations needed for the Oracle Utilities Network Management System ETL processes via the Oracle Utilities Analytics Administration Tool using the two pages grouped under the **NMS Configuration**.



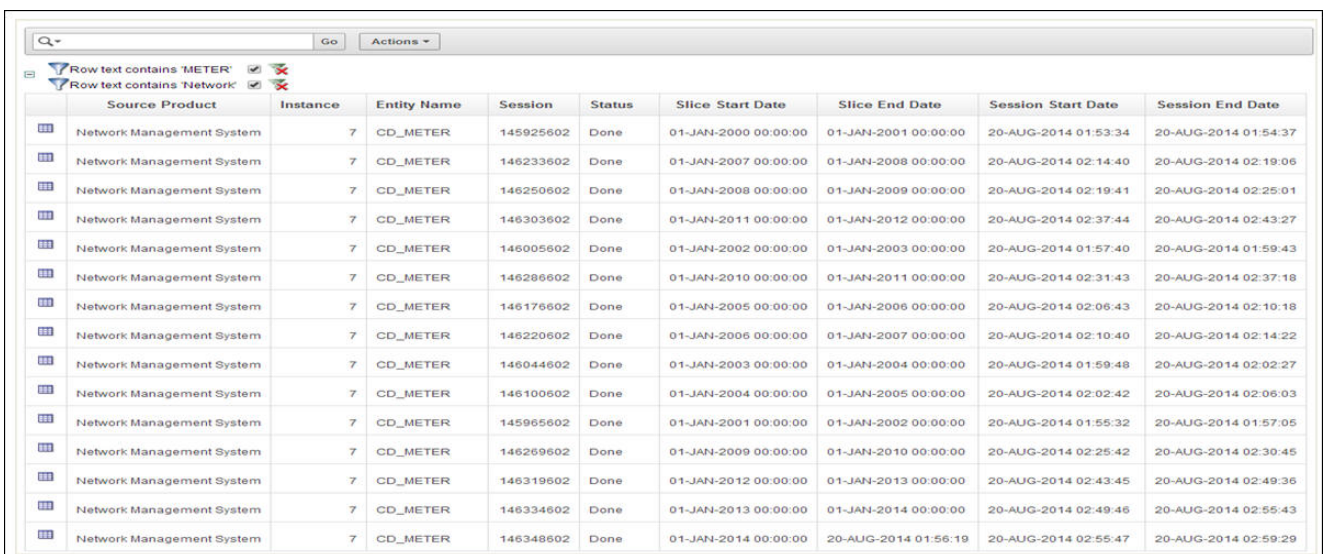
Note: Details of this are covered under the [Configuring Source > Oracle Utilities Network Management System \(NMS\)](#) sub section under the **Chapter 5: Configuring Utilities Analytics**, where the configurations needed for every source system are discussed at length with all necessary details.

ETL Job Execution

The **Job Execution** page in Oracle Utilities Analytics Administration user interface is used for tracking the execution of the ETL processes.



An entry is created for each execution. Here is a screenshot for the job execution log table.



Source Product	Instance	Entity Name	Session	Status	Slice Start Date	Slice End Date	Session Start Date	Session End Date
Network Management System	7	CD_METER	145925602	Done	01-JAN-2000 00:00:00	01-JAN-2001 00:00:00	20-AUG-2014 01:53:34	20-AUG-2014 01:54:37
Network Management System	7	CD_METER	146233602	Done	01-JAN-2007 00:00:00	01-JAN-2008 00:00:00	20-AUG-2014 02:14:40	20-AUG-2014 02:19:06
Network Management System	7	CD_METER	146250602	Done	01-JAN-2008 00:00:00	01-JAN-2009 00:00:00	20-AUG-2014 02:19:41	20-AUG-2014 02:25:01
Network Management System	7	CD_METER	146303602	Done	01-JAN-2011 00:00:00	01-JAN-2012 00:00:00	20-AUG-2014 02:37:44	20-AUG-2014 02:43:27
Network Management System	7	CD_METER	146005602	Done	01-JAN-2002 00:00:00	01-JAN-2003 00:00:00	20-AUG-2014 01:57:40	20-AUG-2014 01:59:43
Network Management System	7	CD_METER	146286602	Done	01-JAN-2010 00:00:00	01-JAN-2011 00:00:00	20-AUG-2014 02:31:43	20-AUG-2014 02:37:18
Network Management System	7	CD_METER	146176602	Done	01-JAN-2005 00:00:00	01-JAN-2006 00:00:00	20-AUG-2014 02:06:43	20-AUG-2014 02:10:18
Network Management System	7	CD_METER	146220602	Done	01-JAN-2006 00:00:00	01-JAN-2007 00:00:00	20-AUG-2014 02:10:40	20-AUG-2014 02:14:22
Network Management System	7	CD_METER	146044602	Done	01-JAN-2003 00:00:00	01-JAN-2004 00:00:00	20-AUG-2014 01:59:48	20-AUG-2014 02:02:27
Network Management System	7	CD_METER	146100602	Done	01-JAN-2004 00:00:00	01-JAN-2005 00:00:00	20-AUG-2014 02:02:42	20-AUG-2014 02:06:03
Network Management System	7	CD_METER	145965602	Done	01-JAN-2001 00:00:00	01-JAN-2002 00:00:00	20-AUG-2014 01:55:32	20-AUG-2014 01:57:05
Network Management System	7	CD_METER	146269602	Done	01-JAN-2009 00:00:00	01-JAN-2010 00:00:00	20-AUG-2014 02:25:42	20-AUG-2014 02:30:45
Network Management System	7	CD_METER	146319602	Done	01-JAN-2012 00:00:00	01-JAN-2013 00:00:00	20-AUG-2014 02:43:45	20-AUG-2014 02:49:36
Network Management System	7	CD_METER	146334602	Done	01-JAN-2013 00:00:00	01-JAN-2014 00:00:00	20-AUG-2014 02:49:46	20-AUG-2014 02:55:43
Network Management System	7	CD_METER	146348602	Done	01-JAN-2014 00:00:00	20-AUG-2014 01:56:19	20-AUG-2014 02:55:47	20-AUG-2014 02:59:29

Some attributes are populated from the SNP_SESSION table, which is used by Oracle Data Integrator to track sessions.

Attribute	Purpose
Job Configuration	A reference to the job configuration.
Session Number	Reference to the Oracle Data Integrator session number.
Scheduled Start Time	Start time as set in the schedule.
Slice Start	Starting timestamp of the slice.
Slice End	Ending timestamp of the slice.
Status	A composite status for the job. Primarily, the status is derived from SNP_SESSION, however additional statuses are tracked in job executions: <ul style="list-style-type: none">• Pending• Submitted• Reprocessed• Error• Done• Running
Actual Start time	The timestamp when the job actually started.
Actual End time	The timestamp when the job actually ended.
Duration	The total execution time in seconds.
Insert Count	The number of records inserted.
Update Count	The number of records updated.
Delete Count	The number of deletions.
Error Count	The number of rows identified as error.
Total Count	The number of rows processed. Sum of all the counts mentioned above.

Chapter 5

Configuring Oracle Utilities Analytics

This chapter describes how to configure the Oracle Utilities Analytics (OUA) product and how to manage and maintain an implementation of Oracle Utilities Analytics, including:

- [Configuring Source](#)
- [Configuring ELT](#)
- [Adding Sources](#)
- [Initiating the ELT Jobs](#)
- [Configuring the Analytics](#)

Configuring Source

Before attaching a specific source system to Oracle Utilities Analytics, there are some configurations that need to be done on the source application. This section provides instructions to configure the source system prior to attaching the corresponding source instance to Oracle Utilities Analytics, including:

- [Oracle Utilities Operational Device Management \(ODM\)](#)
- [Oracle Utilities Customer Care & Billing \(CCB\)](#)
- [Oracle Utilities Network Management System \(NMS\)](#)
- [Oracle Utilities Mobile Workforce Management \(MWM\)](#)
- [Oracle Utilities Meter Data Management \(MDM\)](#)

Oracle Utilities Operational Device Management (ODM)

To enable proper data extracts for Oracle Utilities Extractors and Schema, we need to define some parameters in the Oracle Utilities Operational Device Management application. This section provides information on the steps to be taken to enable this configuration.

- [BI Configuration Portal](#)

BI Configuration Portal

The **BI Configuration** portal holds information on all the BI-oriented configuration tasks. It is a display-only portal that gives a bird's eye view of how much configuration has been set up, and also provides links and guidelines for the areas that need configuration, at the minimum, to successfully run the Oracle Utilities Analytics ELT processes from Oracle Utilities Analytics.

To access the **BI Configuration** portal in Oracle Utilities Operational Device Management, navigate to **Menu > Admin Menu > B > BI Configuration**.

Use the **BI Configuration** portal to perform the following tasks:

- [BI-Oriented Master Configuration](#)
- [Bucket Configuration](#)
- [BI-Oriented Extendable Lookups](#)

Note: These configurations must be done before starting the ELT (Extract, Load, and Transform) processes on the Oracle Utilities Analytics data warehouse.

BI-Oriented Master Configuration

This section lists the master configuration BO that was created for Oracle Utilities Extractors and Schema and guides you during the configuration. Click the link in the **Master Configuration** zone to navigate to the **Master Configuration Maintenance** portal where the configuration is maintained.

Extract Parameters

During extraction of source data to BI data warehouse, there are certain parameters that need to be configured by implementation to identify or filter the data per business requirements. Once these parameters have been set up (by the end user), the ELT process can then use these information to selectively extract data and/or transform data from the source application and populate them into the warehouse.

Note: The ELT job that loads these parameters into the warehouse is configured to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing parameters can cause inconsistency to the loaded data. However, if a need arises to reconfigure the parameters, the parameters can be changed on the source system, and certain additional steps need to be performed to reload them into the warehouse. The star schema tables would also need to be truncated and reloaded.

For the detailed steps involved in reloading the parameters, see [Reloading the Extract Parameters](#) section.

The **BI Master Configuration BO** is delivered to contain the list of parameter values that need to be configured for use by the ELT process:

- **Firmware Parameters:** Holds the firmware identifier types that are used in retrieving and extracting the appropriate firmware versions of the asset. The Operational Device dimension in the Business Intelligence data warehouse has allocated 10 slots for storing an asset's firmware versions.

Use this firmware parameter configuration to specify the type of firmware versions that should be extracted into each of the columns.

- **Operation Device BO Parameters:** Holds the business object names that are used to identify the operational devices that should be extracted. The Operational Device Snapshot and Operation Device Accumulation facts, and the Operational Device dimension in the Business Intelligence data warehouse are designed to hold only the operational device data.

Use this operational device parameter configuration to specify the business objects that are used to identify the assets that should be extracted into the above mentioned star schema tables.

- **Asset Failure Parameters:** Holds the failure condition codes to be considered for identifying asset failures during the extraction process. The Operational Device Snapshot and Asset Location facts in the Business Intelligence data warehouse have specific columns that indicate an asset's failure.

Use this asset failure parameter configuration to specify the conditions that should be considered for marking an asset's failure.

Bucket Configuration

Oracle Utilities Analytics allows you to define a set of ranges, each representing a bucket, for which extracted measures can be grouped and classified under the relevant bucket. This allows for reports that show the data as a group that matches a particular classification, or bucket. For example, in Oracle Utilities Operational Device Management, an asset's age can be classified according to an age grouping (such as Less than 6 Months Old, Between 6 Months to a Year Old, Between One to 5 Years Old, and Older than 5 Years).

- [Defining Age Buckets](#)
- [Asset Age Buckets](#)
- [Asset Installation Age Buckets](#)
- [Asset In Storage Age Buckets](#)
- [Bucket Definition Considerations](#)

Defining Age Buckets

Use the **Bucket Configuration** portal in Oracle Utilities Operational Device Management to define any one of these types of age buckets:

- Asset Age
- Asset Installation Age
- Asset In Storage Age

These buckets can be defined for a specific asset type or as a default set of buckets that is applicable to all assets that do not have any specific buckets defined on the asset type.

Click the link on the **Bucket Configuration** zone to navigate to the **Bucket Configuration Maintenance** portal where the bucket values can be configured.

Note: All the ELT jobs that load these bucket configurations into the warehouse are set to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing bucket configuration can cause inconsistency to the loaded fact data. However, if there is a need to reconfigure the buckets, certain additional steps need to be carried out to pull in the new configuration into the warehouse and ensure that this is reflected in the star schema tables.

For the detailed steps involved in reloading the bucket configuration, see [Reloading the Bucket Configuration](#) section.

Asset Age Buckets

The Asset Age bucket configuration defines various bucket ranges that are used to classify how long an asset has been in the system. See the [Bucket Definition Considerations](#) section for various rules to be followed while defining the buckets.

No values are delivered as a part of the base product. Here are some sample values:

Default: Yes

Sequence	Start Range	End Range	Description
10	0	90	Less Than 3 Months
20	90	180	3 - 6 Months
30	180	270	6 - 9 Months
40	270	365	9 - 12 Months

Sequence	Start Range	End Range	Description
50	365	730	1 - 2 Years
60	730	1095	2 - 3 Years
70	1095	99999	Older Than 3 Years

This data is extracted onto the Asset Age dimension in the Business Intelligence data warehouse. This is referenced in the Operational Device Snapshot fact to categorize an asset's age into one of these buckets.

Asset Installation Age Buckets

The Asset Installation Age bucket configuration defines various bucket ranges that are used to classify how long an asset has been installed on the field. See the [Bucket Definition Considerations](#) section for various rules to be followed while defining the buckets.

No values are delivered as a part of the base product. Here are some sample values:

Default: Yes

Sequence	Start Range	End Range	Description
10	0	90	Less Than 3 Months
20	90	180	3 - 6 Months
30	180	270	6 - 9 Months
40	270	365	9 - 12 Months
50	365	730	1 - 2 Years
60	730	1095	2 - 3 Years
70	1095	99999	Older Than 3 Years

This data is extracted onto the **Asset Install Age** dimension in the Business Intelligence data warehouse. This is referenced in the Operational Device Snapshot fact to categorize an asset's install age into one of these buckets.

Asset In Storage Age Buckets

The Asset In Storage Age Buckets configuration defines various bucket ranges that are used to classify how long an asset has been available in the storerooms. See the [Bucket Definition Considerations](#) section for various rules to be followed while defining the buckets.

No values are delivered as a part of the base product. Here are some sample values:

Default: Yes

Sequence	Start Range	End Range	Description
10	0	30	Less Than 1 Month
20	30	180	2 - 6 Months
30	180	365	7 - 12 Months
40	365	99999	Older Than a Year

This data is extracted onto the Asset Instore Age dimension in the Business Intelligence data warehouse. This is referenced in the Operational Device Snapshot fact to categorize an asset's instore age into one of these buckets.

Bucket Definition Considerations

To define the age bucket ranges follow these rules:

- At least one bucket has the start range as 0.
- At least one bucket has the end range as 99999.
- No overlapping is allowed between buckets. For example: 0-10, 5-20 is an incorrect way to define buckets.
- No gap in allowed between buckets. For example: 0-10, 15-20 or 0-10, 11-15 is an incorrect way to define buckets.

BI-Oriented Extendable Lookups

This section lists all the extendable lookup BOs that were created for Business Intelligence, and guides you to set up the lookup values. Click the link on the **Extendable Lookup** zone to navigate to the **Extendable Lookup Maintenance** portal where the lookup values are configured.

DSI Identifier Mapping

The data source indicator (DSI) identifier mapping defines the integration points of Oracle Utilities Operational Device Management with other Oracle Utilities products (such as Oracle Utilities Meter Data Management). This information is pulled into the BI data warehouse to be used by the ELT processes. During the actual ELT, Oracle Utilities Operational Device Management facts are joined with Oracle Utilities Meter Data Management dimensions using the DSI identifier mapping information.

The Business Intelligence data warehouse uses Service Point integration between Oracle Utilities Operational Device Management with Oracle Utilities Meter Data Management to join Oracle Utilities Operational Device Management facts and Oracle Utilities Meter Data Management / Oracle Utilities Customer Care and Billing dimensions. Use the DSI identifier mapping to configure the following:

- **Maintenance Object:** The maintenance object of the entity being integrated between Oracle Utilities Operational Device Management and the external system.

Note: In the OUASA 2.4.1 release, the ELT processes extract DSI mapping for 'Node' maintenance object from the source Oracle Utilities Operational Device Management system.

- **External System Mapping:** Configure the following items under this group:
 - **Identifier Type:** The identifier type code used to identify the external identifier number from the integrating system.
 - **Product:** The product to which Oracle Utilities Operational Device Management is being integrated with.
 - **Data Source Indicator:** The unique identifier of the integrating system. This is the environment ID of the Oracle Utilities Meter Data Management system. It can be identified from the Oracle Utilities Meter Data Management system using the below query.

```
select ENV_ID from F1_INSTALLATION;
```

Oracle Utilities Customer Care & Billing (CCB)

To enable proper data extracts for Oracle Utilities Extractors and Schema, certain parameters have to be defined in the Oracle Utilities Customer Care and Billing application. This section provides information on the steps to be performed to enable this configuration.

- [BI Configuration Portal](#)

BI Configuration Portal

To enable successful data extracts for Oracle Utilities Extractors and Schema, certain configurations are required in the Oracle Utilities Customer Care and Billing application. This section provides information on what are the needed configuration in Oracle Utilities Customer Care and Billing and how to properly set them up.

To access the **BI Configuration** portal in Oracle Utilities Customer Care and Billing, select **Menu > Admin Menu > B > BI Configuration**.

Use the **BI Configuration** portal to perform the following tasks:

- [BI-Oriented Master Configuration](#)
- [Bucket Configuration](#)

Note: These configurations must be done before starting the ELT processes on the Oracle Utilities Analytics data warehouse.

BI-Oriented Master Configuration

This section lists the master configuration BO that was created for Oracle Utilities Extractors and Schema, and guides you during the configuration. The **BI Configuration** Portal provides a link that allows you to navigate to the **Master Configuration** portal, where master configurations are maintained.

Extract Parameters

During extraction of source data to BI data warehouse, there are certain parameters that need to be configured by implementation to identify or filter the data per business requirements. Once these parameters have been set up (by the end user), the ELT process can then use these information to selectively extract data and/or transform data from the source application and populate them into the warehouse.

Note: The ELT job that loads these parameters into the warehouse is configured to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing parameters can cause inconsistency to the loaded data. However, if a need arises to reconfigure the parameters, the parameters can be changed on the source system, and certain additional steps need to be performed to reload them into the warehouse. The star schema tables would also need to be truncated and reloaded.

For the detailed steps involved in reloading the parameters, see [Reloading the Extract Parameters](#) section.

The BI Extract Parameters Configuration BO is delivered to contain the list of parameter values that need to be configured for use by the ELT process:

- **To Do Characteristic Types:** Used for configuring the Account, Service Agreement, Person, and Premise characteristic types that are used to determine the corresponding entities related to a To Do. This is used for the To Do facts.
- **Recent To Do Miscellaneous:** Used for configuring a numeric value that denotes the 'Completed in Days' parameter. Completed in Days, if specified, is used to retrieve the To Do entries completed within the last 'x' days, for inclusion on the Recent To Do Fact.

- **Financial Transaction Characteristic Values:** Holds the Revenue and Tax characteristic values that are used to determine the Revenue and Tax amount, respectively, of the FT. These values are used by the Financial Transaction, Financial Transaction General Ledger, and SA Billing extracts.
- **Financial Transaction Characteristic Types:** Holds the GL Account characteristic types that are used to determine the GL Accounts that are extracted for the Financial Transactions, Financial Transaction General Ledger, and SA Billing extracts.
- **Exclude Archive Adjustment Types:** Holds all the archive adjustment types that should be excluded by the FT and FT GL extracts.
- **Broken PA Characteristic Types:** Holds the Broken characteristic type that is used to identify if a payment arrangement is broken. This is used by the Payment Arrangement extract.
- **Broken PA Characteristic Values:** Holds the Broken characteristic value that is used to identify if a payment arrangement is broken. This is used by the Payment Arrangement extract.
- **Broken PA Adjustment Cancel Reasons:** Holds the adjustment cancel reason that is used to identify if a payment arrangement is broken. This is used by the Payment Arrangement extract.
- **Overdue Bill Characteristic Types:** Holds the overdue bill characteristic type that is used to identify the bill that is associated with the Overdue Process. This is used by the Overdue Process and Event extracts.
- **High Bill Complaint Characteristic Types:** Holds the characteristic type that is used to identify the bill associated with the High Bill Complaint case. This is used by the SA Billing extract.
- **Exclude High Bill Complaint Case States:** Holds the case statuses for which High Bill Complaint cases should not be considered. This is used by the SA Billing extract to identify if a bill is associated to a high bill complaint.
- **High Bill Complaint Case Types:** Holds the case types for high bill complaint cases. This is used by the SA Billing extract.
- **SA Snapshot Miscellaneous:** Holds a numeric value to denote the Exclude SAs Closed in Days parameter. Exclude SAs Closed in Days, if specified, is used to retrieve SAs closed within the last 'x' days, for inclusion on the SA Snapshot and Payment Arrangement Snapshot extracts.

Bucket Configuration

Oracle Utilities Analytics allows you to define a set of ranges, each representing a bucket, for which extracted measures can be grouped and classified under the relevant bucket. This allows for reports that show data as a group that matches a particular classification, or bucket. For example, in Oracle Utilities Customer Care and Billing, arrears can be classified according to their arrears age (For example, Less Than 30 Days, 30-60 Days, 6-90 Days, Older Than 90 Days).

Defining Age Buckets

The **BI Configuration** portal in Oracle Utilities Customer Care and Billing has a **Bucket Configuration List** zone that lists all the buckets that need to be configured for business intelligence:

- [Billing Day In Window Buckets](#)
- [Days Before Bill Window Closes Buckets](#)
- [Days of Unbilled Usage Buckets](#)
- [Days Since Last Frozen Bill Segment Buckets](#)

- [PA Future Payment Age Buckets](#)
- [PA Number of Installments Buckets](#)
- [PA Recurring Charge Amount Buckets](#)[PA Recurring Charge Amount Buckets](#)
- [PP Future Payment Age Buckets](#)
- [SA Arrears Buckets](#)

The **Bucket Configuration List** zone provides an overview of which bucket configurations have been set up so far, and which are yet to be set up. Those that have not been set up yet will have a corresponding **Add** button which, when clicked, will navigate you to the **Bucket Configuration Maintenance** portal in **Add** mode. Those that have been previously set up will have a foreign key reference link which, when clicked, will navigate you to the **Bucket Configuration Maintenance** portal where you can view or edit the existing bucket values.

Note: All the ELT jobs that load these bucket configurations into the warehouse are set to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing bucket configuration can cause inconsistency to the loaded fact data. However, if there is a need to reconfigure the buckets, certain additional steps need to be carried out to pull in the new configuration into the warehouse and ensure that this is reflected in the star schema tables.

For the detailed steps involved in reloading the bucket configuration, see [Reloading the Bucket Configuration](#) section.

Billing Day In Window Buckets

The Billing Day In Window bucket configuration defines the bucket ranges that identify the day in the billing window when the bill segment was frozen. It also indicates whether the range is within the window or outside the window. Two sets of ranges can be defined for this bucket configuration: one set for day ranges that fall inside the billing window (bill window status of In Window), and another set for day ranges that fall outside the billing window (bill window status of Outside Window).

No values are delivered as a part of the base product. Here are some sample values:

Sample 1:

Bill Window Status: In Window

Sequence	Bill Window Category	Start Range	End Range	Description
10	First Day	0	1	Day 1
20		1	2	Day 2
30		2	3	Day 3
40		3	4	Day 4
50		4	5	Day 5
60		5	10	Day 5+

Sample 2:**Bill Window Status: Outside Window**

Sequence	Bill Window Category	Start Range	End Range	Description
10		0	99,999	Outside Billing Window

This data is extracted onto the Billing Day In Window dimension in the Business Intelligence data warehouse. This is referenced in the SA Billing fact to determine when in the billing window was the bill segment frozen.

The billing day in window buckets enforce the following rules:

- Only one set of ranges can be defined for each bill window status (In Window and Outside Window).
- No gaps or overlaps are allowed within the bucket ranges for a bill window status.
- One initial bucket is required (bucket range starts at 0).
- One catch all bucket is required for the bucket configuration with the bill window status of Outside Window (bucket end range has the maxim value of 99999).
- One and only one bucket range should be categorized as “First Day” for the bucket configuration with the bill window status of In Window.

Days Before Bill Window Closes Buckets

The Days Before Bill Window Closes bucket configuration defines the bucket ranges that identify the number of days left before the bill window closes. Two sets of ranges can be configured for this bucket configuration: one set for day ranges when the bill window is still open (window closure status is Open), and another set for days when the bill window is closed (window closure status is Closed).

No values are delivered as a part of the base product. Here are some sample values.

Sample 1:**Window Closure Status: Open**

Sequence	Start Range	End Range	Description
10	0	1	1 Day Before Bill Window Closes
20	1	2	2 Days Before Bill Window Closes
30	2	3	3 Days Before Bill Window Closes
40	3	4	4 Days Before Bill Window Closes
50	4	99,999	5+ Days Before Bill Window Closes

Sample 2:**Window Closure Status: Closed**

Sequence	Start Range	End Range	Description
10	0	99,999	Bill Window is Closed

This data is extracted onto the Days Before Bill Window Closes dimension in the Business Intelligence data warehouse. This is referenced in the SA Billing fact to determine, for bill segments in Error state, the number of days that are left for correcting the bill segment before the bill window closes.

- The days before bill window closes bucket configuration enforces the following rules:
- Only one set of bucket configuration can be defined for each window closure status (Open and Closed).
- No gaps or overlaps are allowed within the bucket ranges of a window closure status.
- Only one range is allowed for the closed window status and it has to catch all the measures (end range has the maximum value of 99999).

Days of Unbilled Usage Buckets

The Days of Unbilled Usage bucket configuration defines the bucket ranges used to identify how many days the service agreement has usage but has not been billed for.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	0 - 30 Days
20	30	60	30 - 60 Days
30	60	90	60 - 90 Days
40	90	99,999	90+ Days

This data is extracted onto the Days of Unbilled Usage dimension in the Business Intelligence data warehouse. This is referenced in the SA Arrears Snapshot fact to determine how many days that the service agreement has usage but has not been billed.

The following rules are enforced for days of unbilled usage buckets:

- Only one set of the days of unbilled usage bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (the bucket end range has the maximum value of 99999).
- One and only one initial bucket is allowed and required (bucket range that starts at 0).

Days Since Last Frozen Bill Segment Buckets

The Days Since Last Frozen Bill Segment bucket configuration defines the bucket ranges used to identify the number of days since the service agreement has a frozen bill segment.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	0 - 30 Days
20	30	60	30 - 60 Days
30	60	90	60 - 90 Days
40	90	180	90 - 180 Days
50	180	99,999	180+ Days

This data is extracted onto the Days Since Last Frozen Bill Segment dimension in the Business Intelligence data warehouse. This is referenced in the SA Arrears Snapshot fact to determine how long it has been since the service agreement has frozen bill segment.

The following rules are enforced for days since last frozen bill segment buckets:

- Only one set of the days since last frozen bill segment bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (the bucket end range has the maximum value of 99999).
- One and only one initial bucket is allowed and required (bucket range starts at 0).

PA Future Payment Age Buckets

The PA Future Payment Age defines the configurable day ranges indicating the age of future scheduled payments for payment arrangements. The future payments for a payment arrangement are approximated using the number of installments left and the frequency that the payment arrangement SA is billed (via the bill cycle schedule of the account).

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	30 Days
20	30	60	30 - 60 Days
30	60	90	60 - 90 Days
40	90	99,999	90+ Days

This data is replicated into the data warehouse and is referenced in the PA Accumulation/Snapshot fact to categorize a payment arrangement's future payments into one of these age buckets.

The following rules are enforced for PA future payment age buckets:

- Only one set of future payment age bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (bucket end range has maximum value of 99999).

PA Number of Installments Buckets

The PA Number of Installments bucket defines the bucket ranges used to group payment arrangements based on the number of installments.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	3	< 3 Installations
20	3	5	3 - 5 Installations
30	5	8	6 - 8 Installations
40	8	99,999	8+ Installations

This data is extracted into the PA Number of Installments dimension in the Business Intelligence data warehouse. This is referenced in the PA Accumulation/Snapshot fact to determine which bucket a payment arrangement falls into based on its number of installments.

The following rules are enforced on the PA number of installments buckets:

- Only one set of number of installment bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (bucket end range has maximum value of 99999).
- One and only one initial bucket is allowed and required (bucket range starts at 0).

PA Recurring Charge Amount Buckets

The PA Recurring Charge Amount bucket defines the bucket ranges used to group payment arrangements based on their recurring charge amounts.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	25	Less than \$25
20	25	50	\$25 - \$50
30	50	100	\$50 - \$100
40	100	200	\$100 - \$200
50	200	99,999	Greater than \$200

This data is extracted onto the PA Recurring Charge Amount dimension in the Business Intelligence data warehouse. This is referenced in the PA Accumulation/Snapshot fact to determine which bucket a payment arrangement falls into based on its recurring charge amount.

The following rules are enforced for recurring charge amount buckets:

- Only one set of recurring charge amount bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (bucket end range has maximum value of 99999).
- One and only initial bucket is allowed and required (bucket range starts at 0).

PP Future Payment Age Buckets

The PP Future Payment Age bucket defines the age buckets for which future payments for pay plans are scheduled to be paid.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	0 - 30 Days
20	30	60	30 - 60 Days
30	60	90	60 - 90 Days
40	90	99,999	90+ Days

This data is replicated into the data warehouse and referenced in the Pay Plan Accumulation/Snapshot fact to categorize a payment plan's future payments into one of these buckets.

The following rules are enforced on the pay plan future payment age buckets:

- Only one set of future payment age bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.
- One and only one catch all bucket is allowed and required (bucket end range has maximum value of 99999).

SA Arrears Buckets

The SA Arrears bucket defines the bucket ranges for the service agreement's arrears age.

Sequence	Start Range	End Range	Description
10	0	15	0 - 15 Days
20	15	31	15 - 30 Days
30	31	61	31 - 60 Days
40	61	91	61 - 90 Days
50	91	99,999	90+ Days
60	31	99,999	31+ Days

This data is replicated into the data warehouse and is referenced in the SA Arrears Snapshot fact to categorize a service agreement's arrears into one of these buckets based on its age.

Note that the out-of-the-box ETL considers only the first 10 buckets and uses the same to populate UDM1 to UDM10 fields. To bucketize the arrears amount into more than 10 buckets, customize to populate the remaining arrears bucket columns.

The following rules are enforced for SA arrears buckets:

- Only one set of arrears age bucket configuration is allowed.
- At least one catch all bucket has to be defined (bucket end range has maximum value of 99999).

Oracle Utilities Network Management System (NMS)

To enable successful data extracts for Oracle Utilities Extractors and Schema, certain configurations are required in the Oracle Utilities Analytics Administration Tool and Oracle Utilities Network Management System (NMS).

The sections mentioned below provides information on how to setup this configuration:

- [Configuring Oracle Utilities Analytics](#)
- [Configuring Oracle Utilities Network Management System](#)

Configuring Oracle Utilities Analytics

This section provides information about the various configurations to be setup so that the Oracle Utilities Network Management System's ELT functions successfully. Unlike other source systems, these configurations have to be setup using the Oracle Utilities Analytics Administration Tool.

ELT Parameters

During extraction of source data to BI data warehouse, there are certain parameters that need to be configured by implementation to identify or filter the data per business requirements. For example, in order to identify a particular outage event as a momentary outage, implementation has to define the duration of an outage so that it is considered momentary. Once these parameters have been set up (by the end user), the ELT process can then use these information to selectively

extract data and/or transform data from the source application and populate them into the warehouse.

Note: The ELT job that loads these parameters into the warehouse is configured to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing parameters can cause inconsistency to the loaded data. However, if a need arises to reconfigure the parameters, the parameters can be changed on the source system, and certain additional steps need to be performed to reload them into the warehouse. The star schema tables would also need to be truncated and reloaded.

For the detailed steps involved in reloading the parameters, see [Reloading the Extract Parameters](#) section.

The various types of parameters that need to be configured for Oracle Utilities Network Management System ELT are as follows:

- **Country:** The country code where the addresses are located.
- **Interruptions Count for CEMI:** Number of interruptions that you should have experienced to calculate the CEMI index. This is required by the Control Zone Outage and City Outage extracts.
- **Momentary Outage Duration:** The duration in minutes to be used for determining if an outage is a momentary outage or a sustained outage. This is required by the Control Zone Outage and City Outage extracts.
- **Number of Retention Days:** The number of days in which restored outage data is retained in the Recent Call, Recent Crew, Recent Job and Recent Customer Outage facts.
- **Priority Indicator Bit:** The Nth bit, which indicates the priority of the Complaint registered in the Incidents table required by the Call fact extracts.
- **Switch Plan Aborted Step Status Key(s):** The codes used to identify the aborted steps status of a Switch Plan. This is required by the Switch Plan extracts.
- **Switch Plan Begin State Key(s):** The codes used to identify the initial states of Switch Plans.
- **Switch Plan End State Key(s):** The codes used to identify the final states of Switch Plans.
- **Switch Plan Failed Step Status Key(s):** The codes used to identify the failed steps states of a Switch Plan. This is required by the Switch Plan extracts.
- **Switch Sheet Planned Category:** The Planned Category code for Switch Sheets.

Bucket Configuration

Oracle Utilities Analytics allows you to define a set of ranges, each representing a bucket, for which extracted measurement values can be grouped and placed against the relevant bucket. This allows analysts to view the data as a group that matches a particular classification, or bucket. For example, in Oracle Utilities Network Management System, analysts may want to classify the outages based on certain ranges of duration.

Note: All the ELT jobs that load these bucket configurations into the warehouse are set to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing bucket configuration can cause inconsistency to the loaded fact data. However, if there is a need to reconfigure the buckets, certain additional steps need to be carried out to pull in the new configuration into the warehouse and ensure that this is reflected in the star schema tables.

For the detailed steps involved in reloading the bucket configuration, see [Reloading the Bucket Configuration](#) section.

There are three types of buckets that need to be configured for the source system Oracle Utilities Network Management System.

- [Outage Duration Buckets](#)
- [Estimated Restoration Duration Buckets](#)
- [Duration Deviation Buckets](#)

Outage Duration Buckets

The outage duration bucket configuration defines the bucket ranges used for classifying each outage based on how long it took before the outage was addressed and closed. The restored customer outage fact uses this bucket configuration to map the outage duration measure to the corresponding bucket range. The configured data is extracted into the Outage Duration dimension via a separate ELT job.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	Restored within 30 minutes
20	30	60	30 minutes to 1 hour
30	60	9999999	More than an hour for restoration

The following rules are enforced for outage duration buckets:

- No gaps or overlaps are allowed within the bucket ranges.
- One catch all bucket should be defined i.e., the bucket end range has maximum value of 9999999.

Estimated Restoration Duration Buckets

The estimated restoration duration bucket configuration defines the bucket ranges used for classifying outages based on its estimated restoration time. The restored customer outage fact uses this bucket configuration to map the estimated restoration duration measure to the corresponding bucket range. The configured data is extracted into the Estimated Restoration Duration dimension via a separate ELT job.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	To be restored within 30 minutes
20	30	60	30 minutes to 1 hour
30	60	9999999	More than an hour for restoration

The following rules are enforced for estimated restoration duration buckets:

- No gaps or overlaps are allowed within the bucket ranges.
- One catch all bucket should be defined i.e., the bucket end range has maximum value of 9999999.

Duration Deviation Buckets

The duration deviation bucket configuration defines the bucket ranges used for classifying an outage based on how much its actual restoration time deviated from its estimated restoration time. The restored customer outage fact uses this bucket configuration to map the duration deviation measure to the corresponding bucket range. The configured data is extracted into the Duration Deviation dimension via a separate ELT job.

No values are delivered as a part of the base product. Here are some sample values:

Sequence	Start Range	End Range	Description
10	0	30	0 to 30 minutes delay
20	30	60	30 minutes to 1 hour delay
30	60	120	1 to 1.5 hours delay
40	120	9999999	More than 2 hours delay

The following rules are enforced for the duration deviation buckets:

- No gaps or overlaps are allowed within the bucket ranges.
- One catch all bucket should be defined i.e., the bucket end range has maximum value of 99999999.

Configuring Oracle Utilities Network Management System

An entry for 'Data Source Indicator', a unique identifier for each Oracle Utilities Network Management System instance has to be configured in Oracle Utilities Network Management System. The data source indicator is used to uniquely identify the exact source of the data in the warehouse. Each instance should have a different configuration value for Data Source Indicator. Below is a sample Insert statement:

```
INSERT INTO CES_PARAMETERS (APP, ATTRIB, VALUE) VALUES ('BI',  
'DATA_SOURCE_INDICATOR', '000004');
```

Oracle Utilities Mobile Workforce Management (MWM)

To enable successful data extracts for Oracle Utilities Extractors and Schema, certain configurations are required in the Oracle Utilities Mobile Workforce Management (MWM) application. This chapter provides information on what are the needed configurations in Oracle Utilities Mobile Workforce Management and how to properly set them up.

- [BI Configuration Portal](#)

BI Configuration Portal

Oracle Utilities Mobile Workforce Management provides a **BI Configuration** portal that holds all the BI-oriented configuration tasks. There are two configurations that are currently supported in Oracle Utilities Mobile Workforce Management: one for the new **Oracle Data Integrator -based ELT** (via the **Main** tab) and another one for the Oracle Warehouse Builder-based ETL (via the **Oracle Warehouse Builder-Based ETL** tab). Starting with Oracle Utilities Analytics v2.5.2, the ELT methodology uses Oracle Data Integrator therefore configuration should be made on the Main tab.

The **BI Configuration** portal is a display-only portal that gives a bird's eye view of the configuration set up for Oracle Utilities Extractors and Schema, and also provides links and guidelines for the areas that need configuration, at the minimum, to successfully run the ELT processes from Oracle Utilities Analytics.

To access the **BI Configuration** portal in Oracle Utilities Mobile Workforce Management, navigate to **Menu > Admin Menu > B > BI Configuration**.

Use the **BI Configuration** portal to perform the following tasks:

- [Bucket Configuration](#)
- [BI-Oriented Business Object List](#)

Note: These configurations must be done before starting the ELT processes on the Oracle Utilities Analytics data warehouse.

Bucket Configuration

Oracle Utilities Analytics allows you to define a set of ranges, each representing a bucket, for which extracted measurement values can be grouped and classified under the relevant bucket. This allows for reports that show the data as a group that matches a particular classification, or bucket. For example, in Oracle Utilities Mobile Workforce Management, appointments can be classified if the crew missed the appointment, and if so, by how much time.

The **BI Configuration** portal in Oracle Utilities Mobile Workforce Management has a Bucket Configuration List zone that lists all the buckets that need to be configured for BI:

- [Appointment Time Buckets](#)
- [Appointment Time of Day Buckets](#)
- [Late Logon Buckets](#)
- [Early Logoff Buckets](#)
- [Work Duration Deviation Buckets](#)
- [Travel Duration Deviation Buckets](#)
- [Travel Distance Deviation Buckets](#)
- [Response Time Deviation Buckets](#)

The **Bucket Configuration List** zone provides an overview of which bucket configurations have been set up so far, and which are yet to be set up. Those that have not been set up yet will have a corresponding **Add** button which, when clicked, will navigate you to the **Bucket Configuration Maintenance** portal in **Add** mode. Those that have been previously set up will have a foreign key reference link which, when clicked, will navigate you to the **Bucket Configuration Maintenance** portal where you can view or edit the existing bucket values.

Note: All the ELT jobs that load these bucket configurations into the warehouse are set to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing bucket configuration can cause inconsistency to the loaded fact data. However, if there is a need to reconfigure the buckets, certain additional steps need to be carried out to pull in the new configuration into the warehouse and ensure that this is reflected in the star schema tables.

For the detailed steps involved in reloading the bucket configuration, see [Reloading the Bucket Configuration](#) section.

Appointment Time Buckets

The appointment time bucket configuration defines the time ranges (in minutes) indicating whether the crew missed an appointment and if so, by how much time. Each bucket range has an appointment state, which is used to indicate whether the appointment was met on time or if it was missed.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	Appointment Status	Description
-99999	0	On Time	On Time
0	5	On Time	1 – 5 minutes late
5	15	Late	5 – 15 minutes late
15	30	Late	15 – 30 minutes late
30	99999	Late	More than 30 minutes late

The configured data is extracted into the Appointment Time dimension via a separate ELT job. The Crew Task fact and Field Activity Fact use these bucket configurations to determine if appointments are missed and by how much.

The following rules are enforced for appointment time buckets:

- Only one set of the appointment time bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Appointment Time of Day Buckets

The appointment time of day bucket configuration defines the time ranges (in minutes) classifying an appointment time of day.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	Description
420	720	7:00am – 12:00pm
720	1080	12:00pm – 6:00pm
420	1080	All Day 7:00am – 6:00pm
0	99999	Other

The configured data is extracted into the Appointment Time of Day dimension via a separate ELT job. The Crew Task Fact and Field Activity Fact use these bucket configurations to map the appointment time to the corresponding bucket range.

The following rules are enforced for appointment time of day buckets:

- Only one set of the appointment time of day bucket configuration is allowed.

Late Logon Buckets

The late logon bucket configuration defines the time ranges (in minutes) indicating by how much time the crew is late to logon for a shift. Each bucket range has an indicator used for determining if the logon is on time or late.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	On Time Indicator	Description
-99999	0	Yes	On Time
0	4	Yes	On Time
4	15	No	15 Min Late

Start Range	End Range	On Time Indicator	Description
15	30	No	30 Min Late
30	99999	No	Really Late

The configured data is extracted into the Late Logon dimension via a separate ELT job. The completed shift fact uses these bucket configurations to determine if the crew logged in late for the shift and by how much time.

The following rules are enforced for late logon buckets:

- Only one set of the late logon bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Early Logoff Buckets

The early logoff bucket configuration defines the time ranges (in minutes) indicating by how much time the crew is logging off earlier than planned for a shift. Each bucket range has an indicator used for determining if the logoff is on time or early.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	On Time Indicator	Description
-99999	0	Yes	On Time
0	4	Yes	On Time
4	15	No	15 Min Early
15	30	No	30 Min Early
30	99999	No	Really Early

The configured data is extracted into the Early Logoff dimension via a separate ELT job. The completed shift fact use these bucket configurations to map the early logoff of a crew to the corresponding bucket range.

The following rules are enforced for early logoff buckets:

- Only one set of the early logoff bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Work Duration Deviation Buckets

The work duration deviation bucket configuration defines the percentage ranges to indicate the deviation of actual work duration relative to estimated work duration.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	Description
-99999	-25	> 25% Under
-25	-10	10 – 25% Under
-10	-1	1-10% Under
-1	0	As Estimated
0	10	1 – 10% Over

Start Range	End Range	Description
10	25	10 – 25% Over
25	99999	> 25 % Over

The configured data is extracted into the Work Duration Deviation dimension via a separate ELT job. The Field Activity fact uses these bucket configurations determine the percentage of how much the actual work duration is equal, under, or over the estimated duration.

The following rules are enforced for work duration deviation buckets:

- Only one set of the work duration deviation bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Travel Duration Deviation Buckets

The travel duration deviation bucket configuration defines the percentage ranges to indicate the deviation of actual travel duration relative to estimated travel duration.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	Description
-99999	-25	> 25% Under
-25	-10	10 – 25% Under
-10	-1	1-10% Under
-1	0	As Estimated
0	10	1 – 10% Over
10	25	10 – 25% Over
25	99999	> 25 % Over

The configured data is extracted into the Travel Duration Deviation dimension via a separate ELT job. The Crew Task fact uses these bucket configurations to determine the percentage of how much the actual travel duration is equal, under or over the estimated duration.

The following rules are enforced for travel duration deviation buckets:

- Only one set of travel duration deviation bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Travel Distance Deviation Buckets

The travel distance deviation bucket configuration defines the percentage ranges to indicate the deviation of actual travel distance relative to estimated travel distance.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	Description
-99999	-25	> 25% Under
-25	-10	10 – 25% Under
-10	-1	1-10% Under
-1	0	As Estimated
0	10	1 – 10% Over

Start Range	End Range	Description
10	25	10 – 25% Over
25	99999	> 25 % Over

The configured data is extracted into the Travel Distance Deviation dimension via a separate ELT job. The Completed Shift fact uses these bucket configurations to determine how much the actual distance traveled by the crew deviated from the estimated distance.

The following rules are enforced for travel distance deviation buckets:

- Only one offset of travel distance deviation bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

Response Time Deviation Buckets

The response time deviation bucket configuration defines ranges (in minutes) to indicate the deviation of actual response time relative to response time SLA. Each bucket has an indicator used for determining if the SLA is met or missed.

No values are delivered as a part of the base product. Here are some sample values:

Start Range	End Range	SLA Status	Description
-99999	-25	Met	> 25% Under
-25	-10	Met	10 – 25% Under
-10	-1	Met	1-10% Under
-1	0	Met	Meets SLA
0	10	Missed	1 – 10% Over
10	25	Missed	10 – 25% Over
25	99999	Missed	> 25 % Over

The configured data is extracted into the Response Time Deviation dimension via a separate ELT job. The Field Activity fact uses these bucket configurations to determine how much the actual response time deviated from the SLA duration.

The following rules are enforced for response time deviation buckets:

- Only one set of response time deviation bucket configuration is allowed.
- No gaps or overlaps are allowed within the bucket ranges.

BI-Oriented Business Object List

The BI-Oriented Business Object List zone provides a list of records that need to be created to support a successful BI extract:

- [Crew Time Usage](#)
- [Crew Hierarchy](#)
- [Service Area Hierarchy](#)

Crew Time Usage

Crew Time Usage is an extendable lookup defining crew time usage types (Work, Travel, POU, Break, Depot, Idle, Logon, Logoff, Out of Service, Wait), where each crew time usage type is further classified as either productive or not. In the following examples, Travel is defined as "Not Productive" and Work is defined as "Productive".

The Crew Time Usage is used to support KPIs that show how much time crews spend on productive versus non-productive tasks, as well as KPIs that show how much time crews spend on performing logon and logoff procedures at the start and end of their shift, traveling to work, waiting to start activities, performing activities, taking breaks and handling events that cause them to be out of service.

Crew Hierarchy

Crew hierarchies organize crews into logical structures of up to three levels that represent their position relative to other crews. For example, your business practices may require crew hierarchy levels of Business Unit, Department and Office with different crews belonging to each level.

The hierarchies where the crew belongs to is included in the Crew dimension extract so that reports involving crews can be sliced and diced using this information.

Note: For more information about crew hierarchies and how to set them up, refer to the **Chapter 8: Hierarchies** in *Oracle Utilities Analytics Mobile Workforce Management Configuration Guide*.

Service Area Hierarchy

Service Area hierarchies organize service areas into logical structures of up to three levels that represent their position relative to other service areas. For example, your business practices may require service area hierarchy levels of Division, District and County.

The hierarchies where the service area belongs to is included in the Service Area dimension extract so that reports involving service areas can be sliced and diced using this information.

Note: For more information about service area hierarchies and how to set them up, refer to the **Chapter 8: Hierarchies** in *Oracle Utilities Analytics Mobile Workforce Management Configuration Guide*.

Oracle Utilities Meter Data Management (MDM)

To enable successful data extracts for Oracle Utilities Extractors and Schema, certain configurations are required in Oracle Utilities Meter Data Management (MDM) application. This chapter provides information on what are the needed configurations in Oracle Utilities Meter Data Management and how to properly set them up.

This section provides information about the following:

- [BI Configuration Portal](#)
- [BI Aggregators](#)
- [Setting Up the Dynamic Usage Group Extracts](#)
- [Setting Up Snapshot Fact Extracts](#)

BI Configuration Portal

Oracle Utilities Meter Data Management provides a **BI Configuration** portal that holds all the BI-oriented configuration tasks. There are two configurations that are currently supported in Oracle Utilities Meter Data Management: one for the Oracle Data Integrator-based ELT (via the Main tab) and another one for the Oracle Warehouse Builder-based ETL (via the Oracle Warehouse Builder-Based ETL tab). Starting with Oracle Utilities Analytics v2.5.2, the ELT methodology uses Oracle Data Integrator therefore configuration should be made on the Main tab.

The **BI Configuration** portal is a display-only portal that gives a bird's eye view of how much configuration has been set up for Oracle Utilities Extractors and Schema, and also provides links and guidelines for the areas that need configuration, at the minimum, to successfully run the data extraction process in Oracle Utilities Meter Data Management.

To access the **BI Configuration** portal in Oracle Utilities Meter Data Management, navigate to **Menu > Admin Menu > B > BI Configuration**.

Use the **BI Configuration** portal to set up the following:

- [BI-Oriented Master Configuration](#)
- [BI-Oriented Extendable Lookups](#)
- [Service Point Configuration](#)

BI-Oriented Master Configuration

During extraction of source data to BI data warehouse, there are certain parameters that need to be configured by implementation to identify or filter the data per business requirements. Once these parameters have been set up (by the end user), the ELT process can then use these information to selectively extract data and/or transform data from the source application and populate them into the warehouse.

Note: The ELT job that loads these parameters into the warehouse is configured to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing parameters can cause inconsistency to the loaded data. However, if a need arises to reconfigure the parameters, the parameters can be changed on the source system, and certain additional steps need to be performed to reload them into the warehouse. The star schema tables would also need to be truncated and reloaded.

For the detailed steps involved in reloading the parameters, see [Reloading the Extract Parameters](#) section.

The BI Master Configuration - Oracle Data Integrator-Based master configuration BO is delivered to contain the list of parameter values that need to be configured for use by the ELT process:

- [Market Relationship Types](#)
- [Subscription Types](#)
- [Activity Category Types](#)
- [Device Event BOs to Exclude](#)

Market Relationship Types

A service point may have several service providers (for example: distributor, retailer, etc.) where each is defined with a specific market relationship type on either the service point directly or on the service point's market indirectly. The service point-oriented facts can extract up to two of these service provider records and these two service providers are determined via the specified Market Relationship Type. If a specific service point does not have a value for the given market relationship type, the service provider is taken from the market referenced on the service point. It is not an error if no such service provider(s) is found.

Subscription Types

A service point may have several usage subscriptions. The service point-oriented facts can extract up to two of these usage subscription records and these two usage subscriptions are determined via the Subscription Type. The third subscription type, the Oracle Utilities Customer Care and Billing (CCB) Subscription Type, is used to determine the usage subscription that holds the external ID to the Customer Care and Billing service agreement so that references to the Customer Care and Billing dimensions can be populated on the facts. It is not an error if no such usage subscription(s) is found.

Activity Category Types

The Activity accumulation fact can be limited to activities of specific activity type categories. Only activities, whose Activity Type Category is specified on this master configuration BO, are included in the extract.

Device Event BOs to Exclude

The Device Event accumulation fact can be filtered to exclude certain device events. The device events, whose Device Event BO is specified on this master configuration BO, are excluded in the extract.

BI-Oriented Extendable Lookups

Every Extendable Lookup BO that was created for BI is listed in this section and it should serve as a guide/list for you when setting up the lookup values. Click the link on the **Configuration** portal to navigate to the **extendable Lookup Maintenance** portal where the lookup values can be configured. This section describes these lookups:

- [Measurement Condition](#)
- [Days Since Last Normal Measurement](#)
- [Usage Snapshot Type](#)
- [Days Since Last Usage Transaction](#)
- [Unreported Usage Analysis Snapshot Type](#)
- [External System Mapping](#)
- [External System Entity Name](#)
- [External System ID Mapping](#)

Measurement Condition

This extendable lookup is used to define anything interesting about a measurement, such as the source/type of the measurement (for example: a system estimate vs. a normal read vs. a human override). For more details on how condition codes are used on measurements, refer to the “Measurements” section in the *Oracle Utilities Meter Data Management User Guide*.

The measurement condition is a dimension on the SP Usage Snapshot star schema, which means it has dimensional attributes that can be used to slice and dice, and filter consumption.

If you want to quantify consumption by lateness and/or quality, you need to create finer grains of condition codes than those used in the base. For example, base has delivered Regular Measurement Condition as 500100. If you want to further classify a regular measurement that arrived on time, or late, or very, very late, you will have to create additional condition codes to represent them.

500101 - Regular / On Time

500102 - Regular / Late

500103 - Regular / Very Late

500104 - Regular / Very, Very Late

This means that the extendable lookup needs to be extended for additional attributes, such as lateness category, that needs to be sent to the data warehouse during extract.

Days Since Last Normal Measurement

This extendable lookup is used to define the age ranges for days since the last normal measurement was received. Each active instance in this extendable lookup is a bucket definition, where you describe what it is and what is the bucket's upper threshold. If a bucket is meant to have no upper limit (for example: 90+ days), it should be defined with an empty threshold (there should ideally only be one of this).

The lookup value codes should be defined in such a way such that when all the instances for the lookup BO are read in ascending order of the lookup value code, their corresponding upper thresholds are also in ascending order. This is important because the extract logic retrieves all the buckets in ascending order of lookup code, will compare the actual number of days on the ordered threshold, and will also apply whichever matches first (less than the upper threshold or the empty threshold). What this means is that the catch-all bucket (empty threshold) should ideally be defined so that it will be the last bucket retrieved, otherwise, any other bucket definition after the empty threshold will be of no use.

For example, the desired bucket definition is as follows: 0-30 Days, 30-45 Days, 45-60 Days and 60+ Days. This means that the following extendable lookups should be created:

Lookup Value Code	Upper Threshold	Description
10	30	0-30 Days
20	45	30-45 Days
30	60	45-60 Days
40		60+ Days

The SP Snapshot fact has a measure for the number of days since the service point's last normal measurement, and this configuration is used to find the age bucket that corresponds with that number of days since the normal measurement.

Usage Snapshot Type

This extendable lookup is used to define the granularity of the aggregated consumptions of a service point. It defines the TOU map that is applied to the service point's consumptions, where every resultant TOU and condition results in a row on the SP Usage Snapshot fact. It is also used to define the target UOM that is used to convert the source UOM prior to TOU mapping (for example: convert KW to KWH).

The Usage Snapshot Type allows an implementation to have different snapshots, such as:

- On/Off/Sh for CCF
- Day of Week for Therm
- Seasonal On/Off/Sh for Loss Adjusted kWh

This means that a given service point can have several consumption snapshots (but most implementations will have just one).

Days Since Last Usage Transaction

This extendable lookup is used to define the bucket definitions for the number of days since the last usage transaction was created for the service point. These age bucket definitions are used while extracting data for the Unreported Usage Snapshot Fact.

Unlike the Days Since Last Normal Measurement, we do not define the upper threshold here because of the possibility of differing bill cycles for different customer classes. For example, residential customers that bill quarterly vs. commercial/industrial customers that bill monthly. The thresholds are defined on the service point type configuration instead. For more details, see the [Service Point Configuration](#) section.

Unreported Usage Analysis Snapshot Type

This extendable lookup is used to define the different aging snapshots that can be taken for a service point for different types of usage subscriptions. Multiple snapshots of a single service point are allowed as implementations could have multiple systems it is sending consumption to and it may need a snapshot for each.

This allows an implementation to have different snapshots, such as:

- Oracle Utilities Customer Care and Billing
- Distribution
- ...

This means that a given service point can have different Unreported Usage Analysis snapshots (for different types of usage subscriptions).

External System Mapping

This extendable lookup provides a mapping between an external system and its data source indicator. If the external system is one of Oracle Utilities edge applications, the data source indicator should be the environment id on the installation options of that external system. The data source indicator is used to uniquely identify the source of the system data.

External System Entity Name

This extendable lookup defines the entity names of entity information that can be received from external systems.

External System ID Mapping

This extendable lookup provides a mapping between the external identifiers captured in Oracle Utilities Meter Data Management and the external system where that identifier is from. Since it is possible that different sets of records for a particular entity may be synchronized from multiple external systems, it is necessary to specify the identifier type to use for each external system. The key information that is captured in this mapping is the maintenance object (which identifies the MO that stores the identifier collection), the identifier type, the external system, and the entity name (which identifies the type of entity that uses this identifier as its external primary key). For example, the Service Point MO can contain identifiers for the Oracle Utilities Customer Care and Billing Service Point ID and Oracle Utilities Customer Care and Billing Premise ID, as well as the Oracle Utilities Operational Device Management Service Point ID.

Service Point Configuration

This section lists all the service point types in the system and indicates whether the BI configuration has been setup for each of them or not. It provides a navigation link to the service point type where the necessary configuration can be setup or modified.

The service point type allows two configurations:

- [Usage Snapshot Configuration](#)
- [Unreported Usage Snapshot Configuration](#)

Usage Snapshot Configuration

This section of the service point type defines the configuration(s) to be used to take the weekly or monthly usage snapshot(s). The configuration here is used when extracting data for the Usage Snapshot fact.

The Usage Snapshot Type defines the type of usage snapshot. Its extendable lookup definition contains the TOU map (used to map the consumption), and the target unit of measure (used if it is necessary to convert the source UOM to a target UOM prior to TOU mapping).

The UOM, TOU, and SQI are used to define the source MCs value identifier that will be TOU mapped.

A given service point type can have many usage snapshot types if there are different ways to look at the monthly consumption. This is not limited to just different TOU maps, but could also be used to create snapshots of different measured values. For example, if a measurement contains two values, actual and normally used, this can be used to create a snapshot of normal usage so it can be compared to a separate snapshot of actual usage.

If the service point type does not have at least one configuration type, service points of this type do not have their usage snapshot taken.

Unreported Usage Snapshot Configuration

This section of the service point type defines the configuration(s) to be used to take the weekly or monthly unreported usage snapshot(s). The configuration here is used while extracting data for the Unreported Usage Snapshot fact.

The Unreported Usage Analysis Snapshot Type defines the type of unreported usage snapshot.

The UOM, TOU, and SQI are used to define the source MC's value identifier that will be used to calculate the amount of unreported usage in various age buckets.

The Subscription Type is the type of subscription that the analysis will be performed on for this type of snapshot. If the service point is not linked to such a subscription, the fact will be linked to a “none” usage subscription, so analysis of consumption with no usage subscription can be performed.

The Days Since UT Buckets and their corresponding descriptions are used to categorize into different age buckets the amount of consumption that has not been billed. For example: if Bucket 1 is defined as 30, Bucket 2 is 45, and Bucket 3 is 60, any unbilled consumption that is less than or equal to 30 days old will fall into the bucket 1. Any unbilled consumption that is older than 30 days but is less than or equal to 45 days old will fall into the bucket 2. Any unbilled consumption that is older than 45 days but is less than or equal to 60 days old will fall into the bucket 3. Any unbilled consumption that is older than 60 days will fall into the bucket 4.

A separate snapshot can be taken for different subscription types, so that a given service point can have multiple snapshots for a given month/week. If the service point type does not have at least one unreported usage configuration type, service points of this type do not have their unreported usage snapshot calculated.

BI Aggregators

The following types of groupings are used in Oracle Utilities Business Intelligence to calculate the totals by interval for every combination of the dimensions.

- **Postal Code:** This is retrieved from the service point's address information.
- **City:** This is retrieved from the service point's address information.
- **Head-End System:** This is retrieved either from the override head-end system defined on the device or from the fallback head-end system on the device type, if there's no override on the device.
- **Usage Group:** This is retrieved either the override usage group on the usage subscription or the fallback usage group on the usage subscription type, if none on the usage subscription. The Usage Group can be set up as an optional dimension, and if so, service points without primary usage subscription can be included.
- **Market:** This is retrieved from the service point's market. The Market can be set up as an optional dimension, and if so, service points not participating in a market can be included.
- **Service Provider of Role 'X':** If the service point participates in a deregulated market, the service provider is for the role specified. The Service Provider can be set up as an optional dimension, and if so, service points that do not participate in a market and do not have any service provider for a given role can be included.
- **Service Type:** This is retrieved from the service point's details.
- **Device Type:** This is retrieved from the device currently installed on the service point.
- **Manufacturer and Model:** This is retrieved from the device. The Manufacturer and Model can be set up as optional dimensions, and if so, devices with no manufacturer/model definition can be included.

-
- **Geographic Code:** This is retrieved from the service point. The Geographic Code can be set up as an optional dimension, and if so, service points without geographic information can be included.

This section covers the following:

- [Aggregator Measuring Components](#)
- [Setting Up Aggregation Parameters](#)
- [Creating and Aggregating BI Aggregators](#)
- [Refreshing Materialized Views](#)

Aggregator Measuring Components

For every combination of the dimensions listed above and for each type of aggregation there is a distinct aggregator measuring component. These measurements include the aggregated totals for their constituent measuring components.

The four types of aggregators are:

- **Measured Quantity:** In this type of aggregation, the aggregated measurements of the constituent measuring components are spread across buckets as per their measurement conditions.
- **Quality Count:** In the quality count aggregation type, a count for each interval related to a constituent measuring component is placed into one of the quality buckets.
- **Timeliness Count:** In this aggregation type, a count for each interval related to a constituent measurement component is placed into one of the late buckets.
- **Timeliness Quantity:** In this aggregation type, the aggregated measurements of the constituent measuring components are spread across late buckets as per their measurement conditions.

Setting Up Aggregation Parameters

This section describes the required and recommended aggregation parameters that should be set up while configuring the Oracle Utilities Meter Data Management application in BI, including:

- [Master vs. Sub Aggregator](#)
- [Measurement BOs](#)
- [Timeliness Master Configuration](#)
- [Determining Initial Measurement Data Timeliness](#)
- [Value Identifiers on Aggregator MC Type \(Recommended\)](#)

Master vs. Sub Aggregator

Required Configuration: Nominate one of the BI aggregation types as the master aggregator, particularly if the data analysis in BI for various aggregation types is for the same set of customers. Nominating a master aggregator makes the setup and the aggregation process easier, because the master aggregator controls how and when the aggregations are performed.

The first step is to define this type's aggregator measuring component type as the master (its MC Type BO must have an Aggregation Hierarchy Type Option of "Master"). The master aggregator measuring component type controls the aggregation parameters (horizon, lag, and cutoff time), the valid measuring component types to aggregate and its sub aggregator measuring component types.

The next step is to ensure that the measuring component BO of the master aggregator measuring component type has the appropriate algorithms plugged in on it - a BO system event Find Constituent Measuring Components algorithm (which contains the logic on how to find the

constituent measuring components) and an Enter algorithm on its Aggregate state (which contains the logic on how to aggregate the measurements).

The sub aggregator measuring component types can only define the value identifiers that are applicable to it. Its aggregation parameters and valid measuring component types are inherited from the master aggregator measuring component type. Similarly, the Find Constituent MC and Aggregate algorithms are defined on the master (defining these algorithms on the sub aggregators BO will be ineffectual as they will never get triggered).

Note: The usage of the each type of the BI aggregation depends on the implementation. In this case, the implementation should exclude defining this type of BI aggregation as master aggregator or in any of the sub aggregator types.

Measurement BOs

Required Configuration: Define, as appropriate, the special measurement BOs for the four types of BI aggregation (described above under [Aggregator Measuring Components](#)). These are used to differentiate the aggregated measurements from normal measurements created by initial measurement data.

These business objects should be defined on the corresponding aggregation type's measurement component type. The materialized views that were built to aggregate the individual intervals use these specific measurement BOs for performance reasons. These views are accessed directly by the BI analytics.

The measurement business objects are listed as follows:

- Measured Quantity - Measured Quantity Measurement (D2-MeasuredQuantityMsrmnt)
- Quality Count - Quality Count Measurement (D2-QualityCountMsrmnt)
- Timeliness Count - Timeliness Count Measurement (D2-TimelinessCountMsrmnt)
- Timeliness Quantity - Timeliness Quantity Measurement (D2-TimelinessQuantityMsrmnt)

Timeliness Master Configuration

Required Configuration: Add definitions for the timeliness buckets in the Timeliness Master Configuration. This is where definitions for when initial measurement data is considered late, the severity of that lateness, and also the top and bottom time range for that lateness are created.

The severity of lateness can be on time, late, very late, or missing (a measurement is considered missing if it does not exist or if it's condition is either System Estimated or No Read-Other).

In addition, this Master Configuration is where the heating and cooling degree day's factors are setup.

Determining Initial Measurement Data Timeliness

Required Configuration: Mark the initial measurement data with the number of hours that it is late. This definition is used to qualify whether or not a measurement arrived on time or late. This is done on the initial measurement data-level to allow dynamic configuration of what it means for initial measurement data (and measurements) to be late without the need to re-configure the aggregation logic. There is a measure on the initial measurement data's process data that should store the number of hours the initial measurement data is late.

The base application is delivered with an Enter algorithm that calculates the initial measurement data's timeliness as the difference between the initial measurement data's end date time and its actual creation date/time in the system (Determine initial measurement data's Timeliness D2-DET-TML). This algorithm can be plugged in on the Pending state of the initial measurement data life cycle.

Value Identifiers on Aggregator MC Type (*Recommended*)

The base Enter algorithm for aggregating BI measurements (Aggregate Measurement Counts and Quantity D2-AGG-MCQ) has allocated the measurement buckets for Measured Quantity and

Quality Count. While using this base algorithm, it is recommended that the following value identifiers are set up on the appropriate aggregator measuring component type, so that you looking at the aggregated data within the Oracle Utilities Meter Data Management system can see the appropriate identifier labels for each quantity:

Measured Quantity

- Measurement Value - Measured Quantity
- Value 1 - Regular Measurement Quantity
- Value 2 - Estimated Measurement Quantity
- Value 3 - User-Edited Measurement Quantity
- Value 4 - Misc Condition 1 Measurement Quantity
- Value 5 - Misc Condition 2 Measurement Quantity
- Value 6 - MC Count (per interval)
- Value 7 - Heating Degree Days
- Value 8 - Cooling Degree Days
- Value 9 - Average Consumption

Quality Count

- Value 1 - Regular Measurement Count
- Value 2 - Estimated Measurement Count
- Value 3 - User Edited Measurement Count
- Value 4 - No Measurement/No IMD Count
- Value 5 - No Measurement/IMD Exists Count
- Value 6 - No Read Outage Count
- Value 7 - No Read Other Count
- Value 8 - Missing Count
- Value 9 - Misc 1 Count
- Value 10 - Misc 2 Count

The late buckets for Timeliness Count and Quantity are configured via the Timeliness Master Configuration. Similarly, these buckets are recommended to be reflected on the corresponding aggregator measuring component type's Value Identifiers definition.

Creating and Aggregating BI Aggregators

The BI aggregators can be created and aggregated either manually or automatically, just like other the aggregators. For more information, refer to the “Aggregation” section in the *Oracle Utilities Meter Data Management Configuration Guide*.

Refreshing Materialized Views

The two types of materialized views built for each type of BI Aggregation are explained below.

- A materialized view that contains all the BI aggregator measuring components only, where the dimensional values are flattened on the view. The pseudo-dimensions in BI are listed below.

Measured Quantity - D2_MEASR_QTY_MV

Quality Count - D2_QUALITY_CNT_MV

Timeliness Count - D2_TIMELINESS_CNT_MV

Timeliness Quantity - D2_TIMELINESS_QTY_MV

- A materialized view that contains all the aggregated measurements for BI only. The pseudo-facts in BI are listed below.

Measured Quantity - D2_MEASR_QTY_AGR_MV

Quality Count - D2_QUALITY_CNT_AGR_MV

Timeliness Count - D2_TIMELINES_CNT_AGR_

Timeliness Quantity - MVD2_TIMELINES_QTY_AGR_MV

Whenever the dimension scanning or the aggregation is done, it is important that these materialized views are refreshed, so that these materialized views contain the latest data set. An idiosyncratic batch job is delivered to perform this refresh, called Materialized View Refresh (D2-MVREF). This batch job can refresh all the materialized views in one run, so long as the materialized view names are provided as batch parameters in the run.

Setting Up the Dynamic Usage Group Extracts

The Oracle Utilities Analytics Meter Data Analytics allows for usage groups to be used for slicing and dicing the meter data analytics facts. However, in Oracle Utilities Meter Data Management, usage groups are determined dynamically via an algorithm plugged in on the usage subscription's type. Due to this dynamic configuration of usage groups, it is difficult to determine and replicate the logic that will be involved in determining the usage group. A download table is instead created which will store the effective usage group(s) for usage subscriptions.

Extraction

An idiosyncratic batch process can be run on Initial Load or Incremental Load mode that will be responsible for determining the effective dynamic usage group of usage subscriptions and writing the results into the US usage snapshot download table. This batch job is recommended to run the incremental mode daily, to ensure that data in the download table will not be stale. Run the "Usage Subscription Usage Group Download (D2-DUGDL)" batch process to determine the effective usage group of usage subscriptions.

Setting Up Snapshot Fact Extracts

The Oracle Utilities Analytics Meter Data Analytics include four snapshot facts which capture several SP-based information:

- [Consumption Snapshot](#)
- [Service Point Snapshot](#)
- [Service Point Unreported Usage Snapshot](#)
- [Service Point VEE Exception Snapshot](#)

Taking these weekly or monthly snapshots involve massaging and/or transforming the measurement data by applying a few core Oracle Utilities Meter Data Management functionalities such as converting scalar consumptions to intervals, axis conversion and TOU mapping. The huge volume of measurement data and these core Oracle Utilities Meter Data Management functionalities cannot easily be replicated in the BI data warehouse. Hence, these snapshots will be taken in Oracle Utilities Meter Data Management via idiosyncratic batch jobs which will write the results into download tables. These download tables are the ones that the ELT process will replicate and load to the data warehouse.

Consumption Snapshot

The Consumption snapshot allows KPIs that report on every service point's measured consumption. Once a month (or once a week), the system applies a TOU map to every active service point's consumption; every resultant TOU code, quantity, and measurement condition results in a separate row on the Consumption Snapshot fact.

Setup

The following is a list of configuration steps required for a successful SP Usage Snapshot fact extract.

1. Navigate to the appropriate Service Point BO and ensure that the correct algorithm is plugged in on the system event Usage Snapshot (The system event Usage Snapshot on the Service Point BO controls if and how snapshots are taken for the service point. If there is no such algorithm plugged in on the service point's BO, the service point is skipped):
 - a. Deactivate the algorithm "Aggregate SP Usage Snapshot and Write to Flat File (D2-SP-CA)" if the algorithm is plugged in and has not been deactivated yet.
 - b. Plug in the algorithm "Aggregate SP Usage Snapshot and Write to Download Table (D2-SP-USG-DL)" for system event Usage Snapshot.

Note: If the Business Object Algorithm table's caching regime is configured to be "Cached for Batch", it is necessary to clear the cache in order for the algorithm changes to take effect. This can be done by either restarting the thread pool worker or running the Flush All Cache (F1-FLUSH) batch job.

2. Setup the Usage Snapshot Type(s). Refer to the BI Oriented Extendable Lookups section under **BI Configuration** Portal for details on how to setup this data.
3. Configure the Service Point Type's usage snapshot configuration. The Service Point Type holds the below information that controls the extract:
 - a. Usage Snapshot Type
 - b. UOM/TOU/SQI - These are used to define the source measuring component's value identifier that will be TOU mapped. If there are multiple measuring components linked to the service point with such a combination of UOM/TOU/SQI, all will be mapped. If the service point has no measuring components with this combination, it will be skipped.

A given Service Point Type can have many Usage Snapshot Types for different ways to look at the monthly consumption. This is not limited to just TOU maps; this feature can also be used to create snapshots of different measured values.

For example: If a measurement contains two values (actual and normally used), use this feature to create a snapshot of normal consumption, so it could be compared to a different snapshot of actual consumption.

If the Service Point Type does not have at least one entry in the Usage Snapshot Configuration list, it means service points of this type do not have their snapshot taken.

Extraction

The Usage Snapshot fact has its own idiosyncratic batch process that takes the monthly or weekly consumption snapshot of all active service points in the system. Run the "Service Point Usage Snapshot Usage Download (D2-USGDL)" batch process to retrieve the service point's consumptions and populate the SP Usage Snapshot download table.

Service Point Snapshot

The Service Point Snapshot fact allows KPIs that report on the number of service points and/or devices that are installed over time. Once a month (or once a week), the system takes a snapshot of every service point in the system, extracting service point-oriented information as well as information about the service point's install event as of the snapshot date.

Setup

To configure a Service Point Snapshot fact extract successfully, follow these steps:

1. Navigate to the appropriate Service Point BO and ensure that the correct algorithm is plugged in on the system event Service Point Snapshot (The system event Service Point Snapshot on the service point business object controls if and how snapshots are taken for the

service point. If there is no such algorithm plugged in on the service point's BO, the service point is skipped.)

- a. Deactivate the algorithm "Take SP Snapshot and Write to Flat File (D1-SPSNAP-SE)", if the algorithm is plugged and has not been deactivated yet.
- b. Plug in the algorithm "Take SP Snapshot and Insert to Download Table (D1-SPSNAP-DL)".

Note: If the Business Object Algorithm table's caching regime is configured to be "Cached for Batch", it is necessary to clear the cache in order for the algorithm changes to take effect. This can be done by either restarting the thread pool worker or running the Flush All Cache (F1-FLUSH) batch job.

2. Setup the Days Since Last Normal Measurement. Refer to the BI Oriented Extendable Lookups section under **BI Configuration** portal for details on how to setup this data. This will be used to reference the dimension containing the specific age bucket that corresponds with the number of days since the last normal measurement of the service point was received.

Extraction

The Service Point Snapshot fact has its own idiosyncratic batch process that takes the monthly or weekly snapshot of all active service points in the system. Run the "Service Point Snapshot Download (D1-SPSDL)" batch process to extract service point's information and populate the SP Snapshot download table.

Service Point Unreported Usage Snapshot

The Unreported Usage Analysis Snapshot fact allows KPIs that show consumption that has not appeared on usage transactions. Once a month (or once a week), the system reviews all active service points and determines the consumption that has taken place since the most recent usage transaction in the "sent" state.

Setup

To successfully configure the Unreported Usage Analysis Snapshot extract, follow these steps:

1. Navigate to the appropriate Service Point BO and ensure that the correct algorithm is plugged in on the system event Unreported Usage Analysis Snapshot (This system event on the service point BO controls if and how snapshots are taken for the service point. If there is no such algorithm plugged in on the service point's BO, the service point is skipped.)
 - a. Deactivate the algorithm "Analyze Unreported Usage Snapshot and Write to Flat File (D2-SP-UT-AGE)", if the algorithm is plugged and has not been deactivated yet.
 - b. Plug in the algorithm "Take SP Unreported Usage Snapshot and Insert to Download Table (D2-SPUT-DL)".

Note: If the Business Object Algorithm table's caching regime is configured to be "Cached for Batch", it is necessary to clear the cache in order for the algorithm changes to take effect. This can be done by either restarting the thread pool worker or running the Flush All Cache (F1-FLUSH) batch job.

2. Set up the Days Since Last Usage Transaction. This will be used to reference the dimension containing the specific age bucket that corresponds to the number of days since the most recent "sent" usage transaction of the service point. See Unreported Usage Analysis Snapshot Type dimension for setup instructions.
3. Set up the Unreported Usage Analysis Snapshot Type. See Unreported Usage Analysis Snapshot Type dimension for setup instructions.
4. Configure the Service Point Type's unreported usage analysis snapshot configuration. The Service Point Type holds the following information that controls the extract:
 - a. Unreported Usage Analysis Snapshot Type

-
- b. UOM/TOU/SQI - These are used to define the source measuring component's value identifier that will be used to calculate the "unbilled" consumption in the various age buckets. If there are multiple measuring components linked to the service point with such a combination of UOM/TOU/SQI, all will be mapped. If the service point has no measuring components with this combination, it will be skipped.
 - c. Subscription Type - This is the type of subscription that the analysis will be performed on for this type of snapshot; if the service point is not linked to such a subscription, the fact will be linked to the "none" US, so that the analysis of consumption with no US can be performed.
 - d. The Days Since UT Buckets and their corresponding descriptions - These are used to categorize the "unbilled" consumption into different age buckets. For example, if Bucket 1 is defined as 30; Bucket 2 is 45; and Bucket 3 is 60: Any unbilled consumption that is less than or equal to 30 days old will fall into bucket 1. Any unbilled consumption that is older than 30 days but is less than or equal to 45 days old will fall into bucket 2. Any unbilled consumption that is older than 45 days but is less than or equal to 60 days old will fall into bucket 3. Any unbilled consumption that is older than 60 days will fall into bucket.

A separate snapshot can be taken for different subscription types so that a given service point can have multiple snapshots for a given month/week. If the Service Point Type does not have at least one entry in the Unreported Usage Snapshot Configuration list, it means that the service points of this type do not have their snapshot taken.

Extraction

The Unreported Usage Analysis Snapshot fact has its own idiosyncratic batch process that takes the monthly or weekly snapshot of all active service points in the system. Run the "SP Unreported Usage Snapshot Download (D2-UUSDL)" to extract service point's information and populate the SP Unreported Usage download table.

Service Point VEE Exception Snapshot

The Service Point VEE Exception Snapshot fact allows KPIs that report on VEE Exceptions. Once a month (or once a week), the system looks at every service point and count the amount of initial measurement data with and without exceptions. For initial measurement data with exceptions, it further subdivides the count by VEE Exception Type, IMD Type, VEE Severity, VEE Group, and Rule; every resultant subtotal will have a row on the VEE Exception Snapshot fact.

Setup

To successfully configure a Service Point VEE Exception Snapshot fact extract, follow these steps:

Navigate to the appropriate Service Point BO and ensure that the correct algorithm is plugged in on the system event SP VEE Exception Snapshot (This system event on the service point BO controls if and how snapshots are taken for the service point.

If there is no such algorithm plugged in on the service point's BO, the service point is skipped.)

- a. Deactivate the algorithm "Aggregate SP VEE Exceptions and Write to Flat File (D2-SPVEEEXC)", if the algorithm is plugged and has not been deactivated yet.
- b. Plug in the algorithm "Aggregate SP VEE Exceptions and Insert to Download Table (D2-SPVEE-DL)".

Note: If the Business Object Algorithm table's caching regime is configured to be "Cached for Batch", it is necessary to clear the cache in order for the algorithm changes to take effect. This can be done by either restarting the thread pool worker or running the Flush All Cache (F1-FLUSH) batch job.

Extraction

The Service Point VEE Exception Snapshot fact has its own idiosyncratic batch process that takes the monthly or weekly snapshot of all active service points in the system. Run the “SP VEE Exception Snapshot Download (D2-VEEDL)” to extract service point's information and populate the SP VEE Exception download table.

Configuring ETL

The Oracle Utilities Analytics product is extensible and configuration driven. There are multiple configurations that need to be done on the out of the box product. Some of these configurations are applicable system wide. And, some are applicable to the source products and some are entity and instance specific.

This section covers the configuration steps that are part of the initial deployment:

- [Oracle Utilities Analytics Basic Setup](#)
- [Oracle Utilities Analytics Global Configurations](#)

Oracle Utilities Analytics Basic Setup


The global configurations apply to all entities or affect the behavior of the system as a whole. There are two components to the global configuration. Oracle Utilities Analytics global configurations apply to the system and are applicable to each entity. The source product specific global configurations apply to entities associated with the specified source product. Oracle Utilities Analytics basic setup is completed only when these configurations have been set up correctly.



Oracle Utilities Analytics Global Configurations

To access the **Global Configuration** page, perform the following steps:

1. Open the browser to go to the **Oracle Utilities Analytics Administration** user interface.
2. Click on the **ETL Configurations** tab.
3. On the sidebar, click on the item **Global Configurations** as shown below.



The following screenshot shows the available configuration items applicable to the entire product. The first column shows the product for which the configuration applies to.

Product	Instance Number	Description	Value
Oracle Utilities Analytics	(None)	Cut Off Time - The time of the day upto which all changes from the source will be loaded by the ETL jobs. Any changes after this will not be loaded on the same day. Set this to default value (i.e. 23:59:59) to keep loading data continuously.	23:59:59
Oracle Utilities Analytics	(None)	Number of days for which the logs would be retained	0
Oracle Utilities Analytics	(None)	Maximum number of MV interfaces which can run in parallel	4
Oracle Utilities Analytics	(None)	Threshold value to consider for parallel loading for high volume table	1000000
Oracle Utilities Analytics	(None)	Global Extract Date - The date from which all ETL jobs will be configured to start the initial load	01-Jan-2000
Oracle Utilities Analytics	(None)	Enable Partitioning - If a partitioning license has been purchased from Oracle, set this to Y	Y
Oracle Utilities Analytics	(None)	Database Edition Type - Mention the type of the Oracle Database, whether it is Standard Edition(SE) or Enterprise Edition(EE)	EE

The parameter descriptions are self explanatory. However, some additional detail needs to be mentioned about the partitioning parameter.

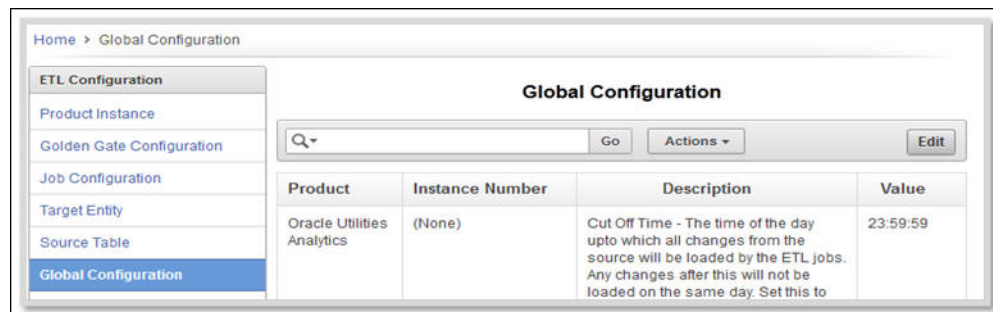
Note: The *Partitioning* feature is an additional license on top of the Oracle Enterprise Edition. Set this property to 'Y' if you have licensed the partitioning features. The Extraction, Loading and Transformation (ELT) processes utilize the partitioning functionality. Setting this flag to 'Y' causes the intermediate staging tables to be partitioned by the job number and significantly improve the parallel execution performance of individual entities.

Note: Some features such as *database parallelism* is not available in the Standard Edition. Setting this property ensures that the ELT processes appropriately utilize the underlying database features for processing the loads.

Editing Global Configuration

To edit the Global Configurations, perform the following steps:

1. Login to **Oracle Utilities Analytics Administration** user interface.
2. Navigate to **ETL Configuration > Global Configuration**.
3. Click **Edit** to open **Maintain Global Configuration** dialog box.



4. Set the **Partitioning** to 'Y' and the **Database Edition** to 'EE'.

Product	Instance Number	Description	Value	Data Type	Data Format
Customer Care and Billing	(None)	Product Extract Date - The date from which all ETL jobs for the specific product instance will be configured to start the initial load	01-Jan-2000	DATE	DD-MON-YYYY
Operational Device Management	(None)	Product Extract Date - The date from which all ETL jobs will be configured to start the initial load	01-Jan-2000	DATE	DD-MON-YYYY
Oracle Utilities Analytics	(None)	Database Edition Type - Mention the type of the Oracle Database, whether it is Standard Edition(SE) or Enterprise Edition(EE)	EE	VARCHAR2(2)	EE/SE
Oracle Utilities Analytics	(None)	Enable Partitioning - If a partitioning license has been purchased from Oracle, set this to Y	Y	VARCHAR2(1)	Y/N
Oracle Utilities Analytics	(None)	Global Extract Date - The date from which all ETL jobs will be configured to start the initial load	01-Jan-2000	DATE	DD-MON-YYYY
Oracle Utilities Analytics	(None)	Threshold value to consider for	1000000	NUMBER	99999

5. Click **Save** to save your changes.

Initial Setup for Star Schema

Perform the following steps to populate the default key records (See the section [Default Keys in Dimensions](#)) into all the dimensions in the data warehouse:


1. Log into **Oracle Data Integrator Studio** and navigate to **Designer > Load Plans and Scenarios > Framework**.
2. Right-click the **B1_INITIAL_SETUP_PKG Version 001** and execute.
3. Select **global** as context, **WLS_AGENT** as logical agent, and then click **OK**.
4. Provide the Minimum date and maximum date range to populate the CD_DATE dimension. Uncheck the option and provide the values in the option.

Note: Check the transactional data for the minimum and maximum date range before executing this package.

5. Go to the **Operator** and view the logs for successful execution of the package.

Adding Sources

The diagram below illustrates the basic steps involved in configuring an instance of a product as a source for Oracle Utilities Analytics.



This section covers the following:

- [Configuring the Socks Proxy Setup](#)
- [Invoking Source Configuration User Interface](#)
- [Using Source Configuration User Interface to Register or Upgrade Sources](#)
 - [Registering a Source](#)
 - [Upgrading a Source](#)
- [Deploying Oracle GoldenGate Scripts](#)

Configuring the Socks Proxy Setup

1. Login in to the source system and open the command prompt.
2. Ensure that the port which will act as the proxy port is not presently active or listening for any other connections on the source system. Use the netstat command to verify that the proxy port is free to use.
3. Run the following command

```
ssh -f -N -D 0.0.0.0:<<UNUSEDPORT> osusername@socks_proxy_server
```

where

- osusername: An account with SSH login capabilities
 - socks_proxy_server: The hostname of socks proxy server
4. Use the netstat command to verify whether the port is now listening.
 5. When configuring source i.e. when running config.sh GUI tool, configure SOXPROXY as sourcesystem:UNUSEDPORT

Invoking Source Configuration User Interface

Run the following commands in the specified sequence from bin folder of the Oracle Home where Oracle Utilities Analytics was installed (This is referred as OUA home in this document):

- **On UNIX Operating Systems:**

```
./splenviron.sh -e <OUA Environment Name>
```

```
cd bin
```

```
./sourceConfig.sh
```

Example: If Oracle Home chosen during installation of Oracle Utilities Analytics is /u01/OUA/OUAProdEnv, Navigate to /u01/OUA/OUAProdEnv/bin and execute:

```
./splenviron.sh -e OUAProdEnv
```

```
cd bin
```

```
./sourceConfig.sh
```

- **On Windows Operating Systems:**

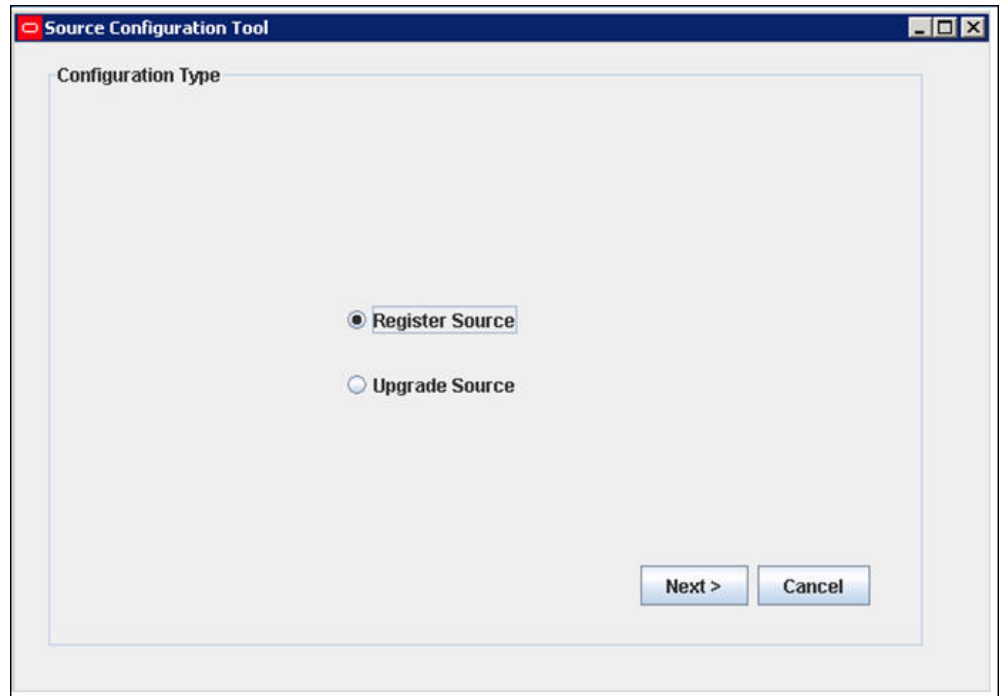
```
splenviron.cmd -e <OUA Environment Name>  
cd bin  
sourceConfig.cmd
```

Example: If Oracle Home chosen during installation of Oracle Utilities Analytics is C:\OUA\OUAProdEnv, Navigate to C:\OUA\OUAProdEnv\bin and execute:

```
splenviron.cmd -e OUAProdEnv  
cd bin  
sourceConfig.cmd
```

Using Source Configuration User Interface to Register or Upgrade Sources

The **Configuration Type** page appears as shown below:



The **Configuration Type** page has two options:

- **Register Source:** The Register Source feature allows you to register a source.
- **Upgrade Source:** The Upgrade Source feature allows you to upgrade a registered source.

Registering a Source

To register a source, perform the following steps:

1. On the **Configuration Type** page, select **Register Source** option and click **Next**.
2. The **Source Product** page appears.
3. Select the source product from the **Source Product** list.
4. Select the instance number from the **Instance Number** list and click **Next**.
5. The **Source Details** page appears. Enter the following details in the respective fields and Click **Next**:

Field Name	Description	Value
DB Host	This is the source database host name.	
DB Port	This is the source database port.	The default port is 1521
DB Service Name	This is the source database service name.	
DB Home Path	This is the source database home installed location.	
	In case GoldenGate for source is not installed on the source database server, provide oracle client home location on the server on which GoldenGate is installed.	

Field Name	Description	Value
DB Schema Name	This is the source schema name.	
Drill Back URL	This is the drill back URL for the source database.	
Extract Start Date	This is the date from which data should be extracted from the source.	Default value is 20000101 (YYYYMMDD)
Socks Proxy	This is the socks proxy host and port separated by a :	Provide the value only if a socks proxy has been setup. Leave the field blank otherwise.

6. The **GoldenGate Details** page appears. Enter the following details in the respective fields and Click **Next**:

Field Name	Description	Value
Host	This is the source GoldenGate server host.	
Home Path	This is the Oracle GoldenGate installed location on the source database server.	For example: opt/local/ggs_home
Manager Port	The port number on which Oracle GoldenGate Manager is running on the Oracle GoldenGate host.	The default dynamic min port is 7830. The default dynamic max port is 7880.
Dynamic Port Range	This is the dynamic port range configured in Oracle GoldenGate on the GoldenGate host.	The default port is 7830.
Encryption Algorithm	This is the algorithm configured in Oracle GoldenGate on the source server.	Enter value 'AES128'.
Encrypt Key	This is the Encrypt Key configured in Oracle GoldenGate on the source server.	Enter the key value generated in <i>Setting up Oracle GoldenGate on the Source Database</i> section of the <i>Installation Guide</i> .
Shared Secret	This is the shared secret key configured in Oracle GoldenGate on the source server.	Go to the Oracle GoldenGate prompt and run the command: encrypt password <password of source GoldenGate owner user>, encryptkey DEFAULT. Provide the result of above-mentioned command as the value for shared secret.
GoldenGate Owner User	This is the user name of GoldenGate Owner user.	

Field Name	Description	Value
GoldenGate Owner Password	This is the password of GoldenGate Owner user.	

- Click **Finish** to complete the registration of a source.
- Upon completion, the status of source registration is shown in a popup dialog. The detailed logs of the operation will be available in the logs/system/deployodi.log file in the Oracle Utilities Analytics home.

Upgrading a Source

To upgrade a source, perform the following steps:

- On the **Configuration Type** page, select **Upgrade Source** option and click **Next**.
- The **Source Product** page appears. All the source contexts that are already registered will show up in the **Source Product** page.
- Select the source product from the **Source Product** list. and click **Next**:
- The **Source Details** page appears. The previously configured values for the selected source are visible in this page.
- Modify the values in the page where required and Click **Next**. The following is a brief description of the fields in this page:

Field Name	Description	Value
DB Host	This is the source database host name.	
DB Port	This is the source database port.	The default port is 1521
DB Service Name	This is the source database service name.	
DB Home Path	This is the source database home installed location.	
DB Schema Name	This is the source schema name.	
Drill Back URL	This is the drill back URL for the source database.	
Extract Start Date	This is the date from which data should be extracted from the source.	Default value is 20000101 (YYYYMMDD)
Socks Proxy	This is the socks proxy host and port separated by a :	Provide the value only if a socks proxy has been setup. Leave the field blank otherwise.

Note: While upgrading a source that was registered using Oracle Utilities Analytics versions prior to Oracle Utilities Analytics 2.5.2, values for the database schema name and drillback URL will not appear by default. These parameters must be entered to proceed with upgrade of the source.

- The **GoldenGate Details** page appears. The previously configured values for the selected source are visible in this page.

- Modify the values where required and Click **Next**. The following is a brief description of the fields in this page:

Field Name	Description	Value
Host	This is the source GoldenGate server host.	
Home Path	This is the Oracle GoldenGate installed location on the source database server.	For example: opt/local/ggs_home
Manager Port	The port number on which Oracle GoldenGate Manager is running on the Oracle GoldenGate host.	The default dynamic min port is 7830. The default dynamic max port is 7880.
Dynamic Port Range	This is the dynamic port range configured in Oracle GoldenGate on the GoldenGate host.	The default port is 7830.
Encryption Algorithm	This is the algorithm configured in Oracle GoldenGate on the source server.	The default value is 1.
Encrypt Key	This is the Encrypt Key configured in Oracle GoldenGate on the source server.	The default value is "DEFAULT".
Shared Secret	This is the shared secret key configured in Oracle GoldenGate on the source server.	Go to the Oracle GoldenGate prompt and run the command: encrypt password <password of source GoldenGate owner user>, encryptkey DEFAULT. Provide the result of above-mentioned command as the value for shared secret.
GoldenGate Owner User	This is the user name of GoldenGate Owner user.	
GoldenGate Owner Password	This is the password of GoldenGate Owner user.	

Note: While upgrading a source that was registered in using Oracle Utilities Analytics versions prior to Oracle Utilities Analytics 2.5.2, the parameter values for Oracle GoldenGate Owner User and Oracle GoldenGate Owner Password are not populated by default. These parameters must be re-entered to proceed with upgrade of the source.

- Click **Finish** to complete the upgrade of the source.
- Upon completion, the status of source registration will be shown in a popup dialog. The detailed logs of the operation are available in logs/system/deployodi.log file in the Oracle Utilities Analytics home.

Deploying Oracle GoldenGate Scripts

The journalization step results in the replication schema being created and the generation of the GoldenGate scripts for each model. The scripts are generated in the directory “ODI Environment” mentioned in the section “Target Oracle Utilities Analytics”. The scripts are generated in the GGScriptsGen folder under the Oracle Data Integrator environment path. The GoldenGate scripts are generated for each model configured for replication. There will be a separate folder with the name of the model. Under each folder, there will be two folders “src” containing the scripts to be deployed on source and “stg” containing the scripts to be deployed on the target. You will find the *Readme.txt* file under each of the model folders (CCB1AA, CCB1AB...) under the path. The scripts are generated on the target server. These need to be copied over to the source server. Typically, this requires a remote file transfer (SCP/FTP/SFTP).

Once these files are copied over, view the *ReadMe.txt* and execute the steps as mentioned in the *ReadMe.txt*. Depending on the source system being configured, there may be one or more folders. Each folder has its own *ReadMe.txt* and the instructions need to be followed and executed for all the models.

For initial load from source to replication, the B1_SYNC_CONTEXT scenario is executed after logging into Oracle Data Integrator. B1_SYNC_CONTEXT uses the database data pump functionality.

The data can be copied to replication layer using one of these two methods:

- DB link
- Export and Import

DBLink

This is the default method used. The Sync method has to be updated in the Admin tool in Global Configuration page for the product and instance that is being configured.

Export and Import

The Sync method has to be updated in the Admin tool in Global Configuration page for the product and instance that is being configured. Update the parameter value to DPUMP_EXPORT to export the data to the file in the pre-configured Oracle directory B1_DATA_DUMP_DIR. Execute the B1_SYNC_CONTEXT for the respective context.

Note: The export code uses Compression feature of the database to compress the export data. This requires an additional license. To use export and import functionality, ensure the appropriate license is available.

Once the data is exported, move the files to the target server in the location pointing to B1_DATA_DUMP_DIR directory.

Change the parameter value to DPUMP_IMPORT in the Admin tool’s Global configuration page. After the parameter value is changed and the files are moved, trigger the B1_SYNC_CONTEXT for the same context to import the data.

Upgrading Oracle Utilities Analytics

If you are upgrading Oracle Utilities Analytics and using one of these source products - Oracle Utilities Meter Data Management, Oracle Utilities Mobile Workforce Management, Oracle Utilities Network Management System whose ELT was earlier based on Oracle Warehouse Builder, certain additional steps are needed. The generated *ReadMe.txt* file has a step to execute the Oracle Data Integrator scenario B1_SYNC_CONTEXT for the initial data load into the replication layer. After this particular step has been completed successfully, note the timestamp using the query given below on the target database:

```
alter session set nls_date_format = 'DD-MON-YYYY HH24:MI:SS';
select max(COMPLETED_ON) from <Context code>REP.CDC_SYNC_LOG;
```

where <Context code> is the context Code with which the source product instance was configured like MDM1, MWM1 etc.

This timestamp is important since the Oracle Data Integrator based ELT jobs should start loading data only after timestamp. Data prior to this timestamp would have already been loaded via the earlier Oracle Warehouse Builder based ETL processes. The next step is to update the configuration of the ELT jobs with this timestamp.

For the ELT jobs of Oracle Utilities Network Management System:

```
update MDADM.B1_JOB_CONFIG
set SLICE_START_DTTM = <timestamp_noted_earlier>
where PROD_FLG = 'NMS'
and TARGET_ENTITY_ID in (
select TARGET_ENTITY_ID from MDADM.B1_TARGET_ENTITY
where ENTITY_NAME in
('CD_ACCT', 'CD_ADDR', 'CD_CALL_INFO', 'CD_CITY', 'CD_CREW',
'CD_CTRL_ZONE', 'CD_DEVICE', 'CD_EVENT', 'CD_EVENT_STATUS',
'CD_FEEDER', 'CD_METER', 'CD_PER', 'CD_PHASE', 'CD_PREM', 'CD_SNL',
'CD_STORM', 'CD_STORM_OUTAGE_TYPE', 'CD_SW_PLAN', 'CD_SW_PLAN_STATE',
'CD_USER', 'CF_CUST_RECENT_OUTG', 'CF_CUST_RST_OUTG',
'CF_RECENT_CALL', 'CF_RECENT_CREW', 'CF_RECENT_JOB', 'CF_RST_CALL',
'CF_RST_CREW', 'CF_RST_JOB', 'CF_SW_PLAN', 'CF_SW_PLAN_STATE'));
```

For the snapshot fact ETL jobs alone, the timestamp should be updated for the first second of the next month. For example, if the last sync timestamp noted earlier is “12-NOV-2013 04:02:30”, then the arrears ETL job should have the timestamp as “01-DEC-2013 00:00:00”. Use the below query on the target database to update the timestamp on the ETL jobs. Derive the timestamp to be used based on the timestamp noted in the earlier.

```
update MDADM.B1_JOB_CONFIG
set SLICE_START_DTTM = <First second of next month to load>
where PROD_FLG = 'NMS'
and TARGET_ENTITY_ID in
(select TARGET_ENTITY_ID from MDADM.B1_TARGET_ENTITY
where ENTITY_NAME in
('CF_FEEDER_DLVRD_LOAD', 'CF_OUTG',
'CF_CTRL_ZONE_OUTG', 'CF_CITY_OUTG'));
```

For the ETL jobs of Oracle Utilities Meter Data Management:

```
update MDADM.B1_JOB_CONFIG
set SLICE_START_DTTM = <timestamp_noted_earlier>
where PROD_FLG = 'MDM'
and TARGET_ENTITY_ID in (
select TARGET_ENTITY_ID from MDADM.B1_TARGET_ENTITY
where ENTITY_NAME in
('CD_ADDR','CD_CONS_TYPE','CD_CONTACT','CD_DAYS_LASTTUT_TYPE',
'CD_DAYS_LAST_MSRMT','CD_DEVICE_ACTIVITY_STATUS',
'CD_DEVICE_ACTIVITY_TYPE','CD_DEVICE_EVT_STATUS',
'CD_DEVICE_EVT_TYPE','CD_EXCP_SEV','CD_EXCP_TYPE',
'CD_IE_STATUS','CD_IMD_TYPE','CD_MC','CD_MSRMT_COND',
'CD_MTR_DEVICE','CD_SP','CD_SPR','CD_SP_STATUS',
'CD_SP_UT_AGE_TYPE','CD_UOM_TOU','CD_UOM_TOU_SQI',
'CD_US','CD_USAGE_GROUP','CD_VEE_RULE',
'CF_DEVICE_EVT','CF_DEVICE_ACTIVITY','CF_INSTALL_EVT','CF_SP'));
```

For the Snapshot fact ETL jobs alone, the timestamp should be updated for the first second of the next month. For example, if the last sync timestamp noted earlier is “12-NOV-2013 04:02:30”, then the arrears ETL job should have the timestamp as “01-DEC-2013 00:00:00”. Use the below mentioned query on the target database to update the timestamp on the ETL jobs. Derive the timestamp to be used based on the timestamp noted in the earlier.

```
update MDADM.B1_JOB_CONFIG
set SLICE_START_DTTM = <First second of next month to load>
where PROD_FLG = 'MDM'
and TARGET_ENTITY_ID in
(select TARGET_ENTITY_ID from MDADM.B1_TARGET_ENTITY
where ENTITY_NAME in
('CF_CONSUMPTION','CF_SP_SNAP','CF_SP_UT_AGE','CF_VEE_EXCP'));
```

For the ETL jobs of Oracle Utilities Mobile Workforce Management:

```
update MDADM.B1_JOB_CONFIG
set SLICE_START_DTTM = <timestamp_noted_earlier>
where PROD_FLG = 'MDM'
and TARGET_ENTITY_ID in (
select TARGET_ENTITY_ID from MDADM.B1_TARGET_ENTITY
where ENTITY_NAME in
('CD_ADDR','CD_CREW','CD_APPT_TM','CD_APPT_TM_OF_DAY',
'CD_CREW_SHIFT','CD_CREW_TM_USG','CD_EARLY_LOGOFF_TM',
'CD_LATE_LOGON_TM','CD_RESP_TM_DEV','CD_SERVICE_AREA',
'CD_SHIFT_BO_STATUS','CD_TASK_BO_STATUS','CD_TASK_TYPE',
'CD_TRAVEL_DIST_DEV','CD_TRAVEL_DUR_DEV','CD_WORK_DUR_DEV',
'CF_CMP_SHIFT','CF_CREW_TASK','CF_FLD_ACTIVITY'));
```

Initiating the ELT Jobs

To start the ELT jobs, two things need to be done. The first task is to setup a scheduler in Oracle Data Integrator which would run continuously looking for jobs to execute. The next task is to enable the specific set of ELT jobs to actually load the target star schemas.

Setting up the Scheduler

Perform the following steps:

1. In **Oracle Data Integrator Studio**, navigate to **Designer > Load Plans and Scenarios > Framework> B1_RUN_ALL Version 001 > Scheduling**.
2. Right-click scheduling and select **New Scheduling** option.

-
3. Select the source context as global, agent as **WLS_AGENT** and the log level as 1.
 4. To specify how often the scheduler should run, navigate to the **Execution Cycle** tab of the Scheduler and select the '**Many Times**' option. Set the interval between repetitions.
 5. Navigate to **Topology > Agents > OracleDIAgent**.
 6. Right-click on **OracleDIAgent** and click **Update Schedule**.

Enabling the ELT Jobs

Using the Oracle Utilities Analytics Administration tool, navigate to **ETL Configuration > Job Configuration**. The jobs can be enabled individually one by one by updating each record in this page. To mass enable multiple jobs at the same time, use the **Enable Jobs** button on the top right side of the page. It is recommend that customers enable the jobs in a progressive manner such as enabling the dimensions first, the facts next and finally the materialized view jobs. This helps reduce the load on the server during the initial data load.

Configuring the Security

For configuring the security, perform the following steps:

1. In the Oracle Utilities Analytics database component installation package, navigate to the `..\BI252\DWADM\Security` folder
2. Edit the file **OraGenSec.bat**, which is in the same location, and replace the parameter `database-name` with the name of your database. This file is provided for your convenience and executes the **Oragensec.exe** utility based on the parameters passed into it.

Note: Ensure to run **OraGenSec.bat** from a Windows desktop that has the Oracle 12c 32 bit client installed. Your database should already be listed in the local file **tnsnames.ora**.

The script executes as the following:

```
oragensec -d DWADM,<DWADM password>,database_name -r  
DW_READ,DW_USER - a A -u DWUSER,DWREAD
```

3. Execute the edited **OraGenSec.bat** file from the command prompt.

Configuring the Analytics

This section describes the various Administration Dashboards and the details about configuration of dashboards.

- [Maintaining the Administration Dashboards](#)
- [Configuring the Dashboards](#)
- [Deploying the Custom MapViewer](#)

Maintaining the Administration Dashboards

The following dashboard pages are described in this section.

- [Maintaining the Base Fields](#)
- [Maintaining the Custom Fields](#)
- [Configuring the Source Drill Back](#)
- [Monitoring the ETL Jobs](#)

Maintaining the Base Fields

This dashboard allows you to provide an override description for the fields delivered along with the product. These fields contain descriptions, which appear on the report titles and the column titles for the dashboards delivered with the product.

Note: For the complete details on the field labels, refer to the [Configuring the Labels](#) in the section [Configuring the Dashboards](#).

Maintaining the Custom Fields

This dashboard allows you to define additional fields and your description, which you may want to add as a part of customization.

This dashboard has the following pages to allow you to add, edit, and delete your own fields:

- Update
- Insert
- Delete

Note: For the complete details on the field labels, refer to the [Configuring the Labels](#) in the section [Configuring the Dashboards](#).

Configuring the Source Drill Back

This dashboard page allows you to provide the source application details for Oracle Utilities Analytics integration with other application. The source application details, such as hostname, port number, and context are entered in this dashboard page. These details can be entered for Oracle Utilities Meter Data Management and Oracle Utilities Mobile Workforce Management edge applications depending on which application is being used.

This information is used by reports in the detail pages for providing a link back to a specific source application page.

Note: For more detailed information on drill back configuration, refer to the section [Configuring the Drill Back](#).

Monitoring the ETL Jobs

Note: This information is applicable **only** for Oracle Warehouse Builder ETL applications.

This dashboard page allows you to monitor the Oracle Warehouse Builder ETL jobs (**specific to Oracle Utilities Work and Asset Management**). You can track the status of the jobs along with its various attributes. During its lifecycle, when a job fails, you can resubmit the job for reprocessing using this page.

If you are using Oracle Data Integrator based source applications (**other than Oracle Utilities Work and Asset Management**), follow these steps to disable the ETL Job Control dashboard:

1. Login to the Oracle Utilities Analytics application.
2. On the **Home** page, click **Catalog**.
3. On the **Folders** pane, navigate to **Catalog Root > Shared Folders > Administration > Dashboards > ETL Job Control**.
4. On the **Tasks** pane, click **Properties**.
5. In the **Properties** window, select the **Hidden** check box under the **Attributes** section.
6. Click **OK**.

These steps allow you to hide the ETL Job Control dashboard.

Configuring the Dashboards

This section describes the configuration dashboards in Oracle Utilities Analytics, including:

- [Configuring the Drill Back](#)
- [Configuring the Oracle Utilities Meter Data Management Answers](#)
- [Configuring the Labels](#)
- [Configuring the Spatial Data](#)
- [Viewing the About Page](#)

Configuring the Drill Back

Oracle Utilities Analytics provides multiple **Drill Back** functionality from various reports in Oracle Utilities Analytics to the source applications, such as Oracle Utilities Meter Data Management (MDM), Oracle Utilities Mobile Workforce Management (MWM), and Oracle Utilities Customer Care and Billing (CC&B). You must configure **Drill Back** by providing the required information for this functionality to work.

The **Configuration Dashboard** under **Administration** group can be used for configuring various options in Oracle Utilities Analytics. The **Configuration** tab contains **Drill Back** settings for the source applications.

You can update the host name, port, and the context root folder for various edge applications with Oracle Utilities Analytics application through this page. After updating the values for the environment, the **Drill Back** links on various dashboard pages use these new values when an item is selected.

Currently, this drill back configuration supports the following edge applications:

- Oracle Utilities Meter Data Management
- Oracle Utilities Mobile Workforce Management

Note: You are required to configure only those source applications for which you have implemented Oracle Utilities Analytics.

Note: The **Drill Back Configuration** option for Oracle Utilities Network Management System, Oracle Utilities Customer Care and Billing is not available. Since the ELT methodology for these source products are based on Oracle Data Integrator, the source Drill Back URL is provided as a part of the **Product Instance** configuration in Oracle Utilities Analytics Administration Tool. For further information, refer to the [Configuring Product Instance](#) section in the chapter 4: Oracle Utilities Analytics Administration Tool.

Configuring the Oracle Utilities Meter Data Management Answers

The following Oracle Utilities Meter Data Management Answers should be configured before viewing the data:

- Tamper Events Answer (**Overview** dashboard)
- Usage Unreported for > 30 Days (**Overview** dashboard)
- Percent of Normal Intervals (**Overview** dashboard)
- Percent of On-Time Intervals (**Overview** dashboard)
- Degree Days (**Overview** dashboard)

After customizing the answers, save the reports in a separate CM catalog.

Note: For details, refer to *Oracle Utilities Analytics Dashboards for Oracle Utilities Meter Data Analytics Metric Reference Guide*.

Configuring the Labels

This section describes how to create and customize the labels that appear in answers and dashboards.

Note: You must have a full license of Oracle Utilities Analytics to use this feature.

This section includes the following topics:

- [Overview of Labels](#)
- [Overriding the Base Labels](#)
- [Supporting the Multiple Languages](#)

Overview of Labels

Oracle Utilities Analytics uses labels for the columns and the tables in the delivered Oracle Business Intelligence Enterprise Edition repository file when displaying the columns in the **Presentation** folders for you. These labels are displayed in the answers for the columns on the dashboards. In addition, the answers are also titled based on labels stored in the metadata displayed in report titles, sub-titles, and other dashboard strings.

The application uses labels instead of the hard coding the text values in the answers and the RPD file for supporting translation of the dashboards into the different languages and allowing easy overriding of the labels for you if you wish to customize the field label.

For example, within an answer, the labels can be referred to by Oracle Business Intelligence Enterprise Edition server variables. For example, the **Device Activity - Distributions** report uses this variable in the title section of the answer:

```
@{biServer.variables['NQ_SESSION.B1_RPT_DEV_ACTI_DISTRIBUTION']}
```

The **B1_RPT_DEV_ACTI_DISTRIBUTION** label is defined in the **B1_MD_FLD** table in **DWADM** schema.

For the columns in the fact and dimension tables, the labels exist for every field. For example, the **UDF1_DESCR** column in the **CD_ACCT** table has the description of the **Customer Class**, and the **Customer Class** label is displayed in the **Presentation** folder for this field.

Overriding the Base Labels

There are several reasons that an implementer may want to update an existing label:

- A field may contain data that does not match the default extracted data for that field. In the **CD_ACCT** example, described in the above section, you may elect to store information other than customer class in the **UDF1_DESCR** field. If an extract change is made to the default **CD_ACCT** extract, then an implementation change in the label for the **UDF1_DESCR** field of the **CD_ACCT** table at one place changes the label in all the dashboards and answers that display that field. This reason also applies if data is extracted to a User Defined Field (UDF) field that is not already having a default population.
- Even if you use the default extract code, you may choose to use some other name for the extracted data other than the default name. In the **CD_ACCT** example, if you execute the field extracted into the **UDF1_DESCR** field account class instead of customer class, you can make this change at one place and have it updated on all dashboards and answers.
- You may want to provide multilingual labels for your users. Oracle Utilities Analytics application provides the labels to a user based on the language selected when logging into Oracle Business Intelligence Enterprise Edition, assuming that the language is present in the **B1_MD_FLD** table. An implementation can add its own translated fields, or can download supported language packs from the Oracle Software Delivery Cloud.

Note: The multilingual support is only provided for labels and not for the underlying data in the data warehouse. The data displayed in all the database tables is not translatable from the extract language.

Supporting the Multiple Languages

Oracle Utilities Analytics is released with default support for English labels on all the dashboards and answers. Both Oracle Business Intelligence Enterprise Edition and Oracle Utilities Analytics support the multiple languages.

The default language on the **Login** page is English. However, you can select any of the supported language on the **Login** page or can change the preferred language under the **Administration** menu to view dashboards in a different language. If you have not purchased and applied the specific language pack and if you select a language other than English, the default Oracle Business Intelligence Enterprise Edition labels is still be translated in the selected language, but Oracle Utilities Analytics product specific labels appear in English.

Oracle Utilities Analytics may release various language packs depending on user demands. Hence, for the language that is already released, installing the language pack is sufficient for creating the labels needed by the dashboards.

To view the list of Oracle Utilities Analytics language pack applied on an environment, you can navigate to the **About Oracle Utilities Analytics** dashboard under the [Viewing the About Page](#) in the Oracle Business Intelligence Enterprise Edition dashboards menu.

Contact your Oracle support representative to purchase an Oracle Utilities Analytics language pack for additional language support.

To update a label for a base field, use the **Base Field Maintenance** dashboard in the **Administration** portal.

- **Table Labels:** For records that have the **Table Name** field populated, but not the **Field Name**, this label is shown in the **Presentation** folder for the fields available in this table. For example, the **CD_ACCT** table has the label '**Account Dimension**' displayed in the **Presentation** folder wherever it is used.
- **Field Labels:** For records that have both the **Table Name** and **Field Name** fields populated, this label is shown in the **Presentation** folder and on answers whenever that field is used. For example, the **UDF1_DESCR** field in the **CD_ACCT** table has the label '**Customer Class**' displayed whenever it is used in an answer or when you select it from the **Presentation** folder when creating a new answer.
- **Report Labels:** Records that have a field name, such as '**B1_RPT%**' and no table name value are used for the titles of answers in the dashboards. For example, the **B1_RPT_DEV_ACTI_DISTRIBUTION** label is defined to be '**Device Activity Distribution**', and this is displayed on the **Oracle Utilities Meter Data Management** dashboard when the answer is displayed.
- **Other Labels:** All other non-report labels that have a field name, but no table name is used for calculations that are computed in the RPD logical layer for display on answers. For example, the **B1_APPT_IND_SUM** label is defined to be '**Number of Appointments**', and is used in Oracle Utilities Mobile Workforce Management answers that compute the number of crew appointments based on the **Appointment Indicator** field in the **CF_CREW_TASK** fact table.

If a base field label should be changed, then the implementation team can query the required record on the **Base Field Maintenance** dashboard; populate a value in the **Override Description** field, and click **Update**. Once populated, the Oracle Business Intelligence Enterprise Edition Server must be restarted for the changes to take effect.

Configuring the Spatial Data

This section describes how to configure mapping for Oracle Utilities Analytics. It includes the following topics:

- [Loading the Geographical Data](#)
- [Setting Up the Network Model Spatial Data in Outage Analytics](#)
- [Configuring the Google Map Tile Layer](#)
- [Implementing the Maps](#)

Loading the Geographical Data

In order to place information on a geographic map, data in the data warehouse must match geographic data (themes) that are configured in Oracle MapViewer.

The standard map answers delivered with Oracle Utilities Analytics include maps that query state, city, county, postal codes, and network model summary data. As Oracle Utilities Analytics does not have access to this spatial data (and as each user may require different spatial data), you must set up the geographic themes used in the maps.

Note: For details regarding setting up these standard spatial themes, refer to *Oracle Utilities Analytics Installation Guide*.

The installation instructions refer to shape files downloaded from the US Census Bureau. However, the shape files can also be used for the state, city, county, and zip code boundaries. The only requirement is that the names of the geographic entities in the shape file should match the corresponding name in the **CD_ADDR** table. This is not usually a problem for postal code data, but can be an issue for city and county names, as different sources may use different names to refer to the geographic places. Make sure that after loading the **MapViewer shapefiles** that the names in the geographic tables match the names that are in the **CD_ADDR** table. If these do not match, then the maps may not display correctly.

Setting Up the Network Model Spatial Data in Outage Analytics

This section provides Oracle Utilities Network Management System (NMS) network model representation in Outage Analytics.

Oracle Utilities Network Management System provides model build process to generate geometry data for various components of network model. The geographic data is stored in the **DIAGRAM_OBJECTS** table.

Note: For details regarding the steps required to set up the Network Management System Model Build process, refer to *Oracle Utilities Network Management System Installation Guide*.

To plot these elements in Outage Analytics Dashboards, the **DIAGRAM_OBJECTS** table should be replicated as per the instructions available in the section **Installing US State Spatial Data** in *Oracle Utilities Analytics Install Guide*.

The proper mechanism should be in place to be up-to-date with the **DIAGRAM_OBJECTS** table in the Source database.


To ensure that the network model can be displayed without coordinate translations during runtime, one of the geometry columns should use the same projection coordinate system as the base map used by the Outage Analytics Outage Maps. Out of the box Oracle eLocation is used for the base map and the SRID (Spatial Reference System Identifier) is 54004. Note that the SRID is a unique value used to identify the coordinate system used in a Geographic Information System (GIS) application.

Configuring the Google Map Tile Layer

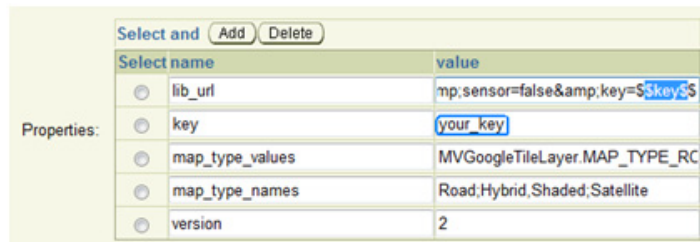
Out of the box Oracle Utilities Analytics is configured to fetch the map tiles from the eLocation. However, you can switch to Google as an alternative source for the map tiles.

Perform the following steps:

1. Log onto the **MapViewr** console.
Sample URL:
`http://<host-name>:<port>/mapviewer`
2. Click the **Admin** and enter the credentials.
3. On the **Management** tab, click the **Manage Map Tile Layers**.
4. Create a new tile layer by selecting **Google**.




5. Enter the key fetched from Google in the **Key** field, edit the default **lib_url** to include the key value and choose the appropriate data source.



6. Click **Submit** to save the information.



7. Click **View map/Manage tiles** to verify the new tile layer shows up properly.
8. Click **Show Map** to view the map.
9. Logon to Oracle Business Intelligence Enterprise Edition Analytics and navigate to **Administration > Manage Map Data** and click on the **Background Maps** tab.



10. Import the tile layer '**GOOGLE_MAP**' created on the **MapViewr** console.

11. Add layers to the map and save it.

Edit Background Map - GOOGLE_MAP


Name:

Description:

Interactive BI Layers
 For each layer, select the zoom levels at which associated BI data can display

	Zoom Level																		
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
City																			
City Name																			
Postal Code																			

12. Subject areas will now be associated to the map.
13. Create a new answer and add a new map-view to have the data displayed on the map.



Implementing the Maps

This section describes the method of implementing maps in Oracle Utilities Analytics. The implementation method is the default implementation method for Oracle Business Intelligence Enterprise Edition 11g. This form of map can be seen in various dashboards, such as Oracle Utilities Meter Data Analytics, Oracle Utilities Mobile Workforce Analytics, and Oracle Utilities Customer Care and Billing Analytics. Using this method, you can create new answers using the MapView. This view uses the configuration defined under the **Administration** menu in the **Manage Map Data**.

The layers, background maps, and the images being used in the map must be defined in this page. The key column and geographical columns are mapped for each subject area used in the analysis. This is a one-time setup unless new subject areas are added.

Note: You should not customize the map metadata until you import the Spatial Catalog file.

For customizations that involve map analysis, all the modifications must be done in a separate folder in order for those modifications to be preserved when upgrading Oracle Utilities Analytics.

Viewing the About Page

The **About** page shows information about the product name along with the current release version and patch number. It also lists all of the languages that are currently installed in the product.

Enabling Weather Information

Oracle Utilities Analytics product supports weather information display on some of its Oracle Business Intelligence Enterprise Edition map reports. Currently, this feature is limited only to the users of Outage Analytics. To enable this feature, perform these steps:

- [Deploying the Custom MapViewer](#)
- [Subscribing to the Weather Service Provider](#)
- [Configuring the Weather API Details](#)

Deploying the Custom MapViewer

This is an optional step in the installation process. Outage Analytics users who wish to enable the weather features should carry out all the steps, documented in the section **Configuring and Deploying the MapViewer** under the chapter **Oracle Utilities Analytics Initial Installation** in *Oracle Utilities Analytics Installation Guide*.

Subscribing to the Weather Service Provider

Weather information is not captured and stored anywhere in the Oracle Utilities Analytics product. This is a complex process and is beyond the scope of Oracle Utilities Analytics. There are third party vendors who provide weather information. Each vendor has their own set of APIs, which provide weather information in varied formats.

The previous step is to deploy the Oracle Utilities Analytics code, which would contact one of such weather vendors, process the data and render it on the maps. Out-of-box the Oracle Utilities Analytics code will work with Pulse APIs provided by Earth Networks. You need to subscribe these from their website (<http://weather.weatherbug.com/pulseapi.html>) and opt for one of their package.

Note: This is a third party vendor and you need to purchase the license separately from them.

Upon successful registration and purchase, you would be provided with a consumer key and secret key. Using this, you should generate an access token using the URL mentioned below:

```
https://thepulseapi.earthnetworks.com/oauth20/
token?client_id=<insert_consumer_key_here>&client_secret=<insert_secret_key_here>&grant
_type=client_credentials
```

This should provide a JSON response with the access token as shown below:

```
Response format: {"OAuth20":{"access_token":{"token":"<access token>
","refresh_token":"<refresh token>","token_type":"bearer","expires_in":86399}}
```

Configuring the Weather API Details

The final step is to configure couple of values for the Oracle Utilities Analytics product to make use of. More specifically, the weather API provider's URL domain name and the access token should be configured.


Log onto the Oracle Business Intelligence Enterprise Edition dashboard environment, navigate to **Administration dashboards > Configuration > Map Profile** page. Update the two records as highlighted below. Use the access token that was generated in the earlier step.

Map Profile

Type	Feature	Description	Value	Override Value
Map Attribute Profile	General Attributes	Copyright Text	©2014 Oracle Corp. ©2014 HERE Maps™	
		Profile Cache	false	
	Map Source Attributes	Base Map	b1_world_map	
		Bing Map API Key		
		Bing Map Source	http://ecn.dev.virtualearth.net/mapcontrol/mapcontrol.ashx?v=6.2	
		Default Data Source	MAPCONN	
		Google Map API Key		
		Google Map Source	http://maps.google.com/maps?file=api&v=2	
		User Third Party Map	none	
		User Third Party Map Type	hybrid	
	Weather API Attributes	Weather API Access Token	dummy_value	hKKaCeB32A4rHo0MoqZSicDIAqTn
		Weather API Domain	dummy_value	https://thepulseapi.earthnetworks.com
		Weather API Password Key	dummy_value	
		Weather API User Key	dummy_value	

Update Revert

Click **Update** to save the changes. The weather configuration is now complete. To see the weather information, navigate to Outage Analytics, go to **Current Outages** dashboard > **Overlay Map** page, you would see something similar to the below snapshot.



Note: The out-of-box weather integration code of Oracle Utilities Analytics works fine for Pulse APIs. If you want to opt for other weather information vendors, they would need to customize the code to work with other APIs. Every weather vendor will have their own set of APIs and the response format will vary from one to another. Hence, the weather integration code is different for every vendor.

Note: For additional details, refer to the white paper **Weather Integration using Third Party Weather APIs (Doc ID 1947857.1)** available at the My Oracle Support website (<https://support.oracle.com>).

Chapter 6

Managing Oracle Utilities Analytics

This chapter describes some of the important configurations that need to be set up to ensure that the Oracle Utilities Analytics product works smoothly. This chapter also provides additional recommendations on utilizing partitioning feature of Oracle database to improve the performance of the overall Oracle Utilities Analytics product, including:

- [Managing ELT Processes](#)
- [Performance Recommendations](#)

Managing ELT Processes

Oracle Data Integrator (ODI) based ELT uses Oracle GoldenGate and Oracle Data Integrator combination to extract, load and transform data from the source database to the target database. Oracle GoldenGate captures the changes in the selected tables of the source schema, and transfers these data to the replication schema. The tables in the replication schema are similar in the structure to the source tables with a few additional columns added for tracking history and audit. Data retention in the replication schema is controlled by configuration stored in the metadata.

Based on the scheduler configuration, data from replication schema is transferred by Oracle Data Integrator to the staging schema. The tables in the staging schema are similar to the tables in the target schema. It contains a few additional columns for data transformations. There are no constraints on the staging tables. The foreign key mapping and other transformations are performed in the staging area. The data retention in the staging schema is controlled by configuration stored in the metadata.

From the staging schema, Oracle Data Integrator transfers data to the facts and dimensions in the target schema. Oracle Data Integrator also loads and refreshes the materialized views in the target schema. This chapter describes the following:

- [Recommended Configuration](#)
- [Scheduling](#)
- [Monitoring the Jobs](#)
- [Debugging](#)
- [Handling the Errors](#)
- [Reloading the Data](#)

Recommended Configuration

This sub-section covers the following:

- [Database](#)
- [Estimating Connection Pool Size](#)

Database

Setting	Recommended	Reason
Open Cursors	1000	<p>Oracle Data Integrator based ELT processes utilize database connections to perform transformations to load data into the data warehouse. Depending on the number of processes running in parallel, the additional connections are created to the database.</p> <p>Oracle Business Intelligence Enterprise Edition also uses the database connections to query and fetch data for the answers on various dashboards. Depending on the number of concurrent connected users, Oracle Business Intelligence Enterprise Edition spawns the additional database connections.</p> <p>Considering these standard database accesses and other ad-hoc database access, the open cursors parameter should be set to a minimum of 1000. The higher values may be set based on the implementation usage.</p>
Partitioning	Enabled	<p>A data warehouse typically stores several years of transactional data. This can lead to a significantly large volume of data. Most analytics are monthly aggregated summaries, and hence partitioning the data by month and year will enable faster access to the smaller subset of data in a specific answer.</p> <p>The Oracle Data Integrator processes also leverage the partitioning feature to isolate the data processed across multiple slices and process significantly higher load volumes in the lesser time.</p>

Setting	Recommended	Reason
Connection Pool Size	200 (Default is 12)	<p>Each Oracle Data Integrator job utilizes a connection from the pool and the connection is released when not in use by the job.</p> <p>Oracle Utilities Analytics ELT executes multiple jobs in parallel. The default configuration enables 6 parallel instances of each entity. This is particularly helpful during initial load when a large volume of historical data needs to be loaded into the data warehouse.</p> <p>Depending on the source system being configured, this can result in spawning up to 300 jobs. In case more than one source system is attached to Oracle Utilities Analytics, this number can increase further.</p> <p>Ensure that the connection pool size is increased appropriately when increasing the max_parallel configuration for an entity or when attaching more than one source system.</p> <p>In case the connection pool is not sized appropriately, you will notice that many jobs go into the waiting state until a connection becomes available.</p>

Estimating Connection Pool Size

Use the query mentioned below to estimate the connection pool size. This should be executed after any changes in the max_parallel setting for the entities or after attaching a new source instance.

```
select round(
    greatest(
        sum(case when te.entity_type_flg in ('SCD1','SCD2')
            then max_parallel
            else 0
        end)
        ,sum(case when te.entity_type_flg in
('FACT','RACC','SNAP')
            then max_parallel
            else 0
        end)) *.6,-2)
estimated_connection_pool_size
from mdadm.bl_target_entity te
, mdadm.bl_job_config jc
where jc.target_entity_id = te.target_entity_id;
```

Scheduling

In any data warehouse, the basic challenge is to get the data loaded quickly and efficiently whether it is the initial load or the incremental load. Data volumes are high during initial load and during incremental loads the volumes are considerably smaller. The Oracle Data Integrator based ELT processes utilize a time based slicing mechanism to split the load volumes into more manageable slices to process the initial as well as incremental loads. The metadata configurations allow you to control the size of the slice, parallelism of the loads and much more. To ensure that the data

loaded is always consistent, the process executions are governed by a set of rules. The following criteria are considered for the job execution:

- There should be no errors, which need reprocessing.
- The maximum retries limit for the day should not be exceeded or reached.
- Tasks wait for the configured retry interval before submitting a retry for the job.
- Number of parallel jobs is always being limited to the maximum parallel configured.
- Jobs are not executed beyond the time of the most recent Oracle GoldenGate sync, or to a specified cut off time whichever is less.
- If a job is dependent on tables, which are being synced by separate Oracle GoldenGate processes, then the common sync time of both processes is considered.
- If a scenario does not exist, then the jobs are not executed.
- An interface should be active.
- All dependencies should be run.
- The number of running/error instances of the job should be less than the maximum parallel executions allowed.
- If the instance is configured as run once, then it should not execute once it is successfully executed.
- If a job fails, it should be retried again until the maximum retries per day is reached. The interval between successive retries should be based on configuration.
- Oracle GoldenGate models comprising of the source tables used in the entity should have been synced up. In case, the sync timestamps vary across multiple models, then the maximum sync timestamp is used.
- The snapshot entities are executed on or after the snapshot end period.
- The schedule time can be used to stagger loads and distribute processing. A job is not being executed until the current time crosses the scheduled time.

Note: Refer to the [Configuring Target Entities](#) section in the chapter 4: Oracle Utilities Analytics Administration Tool for where to schedule jobs.

All jobs executions are internally managed by the scenario named **B1_RUN_ALL**. This is scheduled using Oracle Data Integrator to run every 1-10 minutes as per your requirement.

Monitoring the Jobs

As jobs keep running on a regular basis, these are monitored or tracked to ensure that the required performance parameters are met. Jobs are created so that these are capable of automatic re-execution and retries. However, reasons for failure should be looked into and appropriate actions should be taken to resolve those issues. The following views are provided for achieving this:

View	Purpose
b1_jobs_vw	Provides a list of jobs executed with details of the slice start and end, scheduled execution time, actual start and end times, status, and record counts.
b1_config_vw	The consolidated view of the current configuration for all entities in the data warehouse.

View	Purpose
b1_wait_reasons_vw	There are possibilities of erroneous configuration or other reasons due to which entities may not be running. To figure out what may be preventing a job from executing, have a look at this view's results.

Debugging

Oracle Data Integrator jobs are designed to make it easy to look at the data that was processed to figure out issues in the processing pipeline. To do this, the staging tables are utilized that retain data for a configured duration. Each execution results in some data that has been processed from the replication layer into the staging area, and then further on to the target entity. Each such set is stored in the staging table with the session number to identify the executing session.

At any point of time if a job fails or to look at how data was before it was loaded into the target, you can query the associated staging tables to look at the data.

The staging tables' names are derived by prefixing the target entity name with "STG_".

Handling the Errors

There are three basic components utilized in the data processing pipeline:

- Oracle GoldenGate for replicating changes from the source to the target
- Oracle Data Integrator to load data from the replication layer to the target
- Oracle database to store and manage data in the different layers. Oracle data pump is utilized to perform the initial sync between the source database and the target database.

There is a possibility of error occurring in one or more stages of the process. This section covers some of the details to help you identify the root cause and provides pointers to take appropriate action to resolve the issues that may arise:

1. **Failure of Initial Sync Processes:** In case of an error, you can execute the initial sync process again and it starts up from where it was left off. If the Oracle GoldenGate replication process is running while the initial sync is performed it is possible that the process will fail due to the duplicate keys. Here are the steps to resolve the issue:
 - i) Stop the relevant replicat process.
 - ii) Delete the data from the tables, which are listed in the error.
 - iii) Execute the initial sync process again.
 - iv) After the initial sync has been successfully loaded start the replicat processes.
2. **Failure of Oracle GoldenGate Processes:** Oracle GoldenGate processes can fail if the source or the target database goes down or the network connectivity fails between the source and the target. To identify whether Oracle GoldenGate processes are running properly, you should login to the server and run the command `info all` on the Oracle GoldenGate command prompt. If you see that any jobs shows up as `ABENDED` or `STOPPED`, verify the reason for the error by looking at the Oracle GoldenGate logs and start the jobs. Oracle GoldenGate processes will restart from the point of failure and continue the replication processes.
3. **Failure of Oracle Data Integrator Processes:** The Oracle Data Integrator processes have been designed to automatically retry for a configurable number of attempts per day. If the issue is related to a table spaces not available or database connectivity issues, the jobs would automatically execute once the issue has been resolved.

If the issue resolution is anticipated to take few hours, it is advisable to disable the failing jobs for the duration of the fix.

Query the view B1_JOBS_VW in the metadata schema to identify if any jobs have failed. In case of a failure, the status_flg column will show 'E'. You can log into Oracle Data Integrator and filter by the session number (available in the jobs view) and look at the error encountered.

Any failures will rollback any changes to the target. A late arriving dimension will be tagged to a default key (-99) in the dimension and these will be reprocessed in subsequent loads. If the dimensional data has arrived into the warehouse, then the foreign key references will be corrected automatically.

Reloading the Data

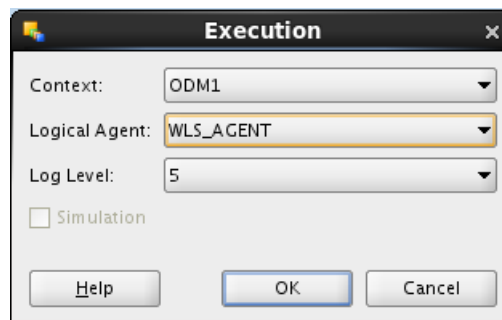
Data warehouses are usually designed based on the assumption that data is added only once to the warehouse. It is rare, although, sometimes to reset specific entities, the data is reloaded again. Oracle Utilities Analytics provides the functionality to reset and reload individual entities, or all data associated with a specific instance of a source system:

- [Resetting an Entity](#)
- [Resetting an Instance](#)
- [Reloading the Extract Parameters](#)
- [Reloading the Bucket Configuration](#)


Resetting an Entity

Perform the following steps to reset an [Entity](#):

1. Login to **Oracle Data Integrator** client.
2. Navigate to the **Load Plans & Scenarios** section. Select the scenario named **B1_RESET_ENTITY**. Right-click and execute.
3. In the popup menu, select the context.
For example, to reset the fact for Oracle Utilities Device Management product instance 1, select the context **ODM1**.



- Uncheck **Last Value** checkbox, specify the variable value, and click on the **Name** field before clicking **OK**.



This cleans up the metadata entries for the executions of the ELT job for this specific entity. The data loaded in the particular table for the entity will also be clear. Note that the ELT job for this entity will now be in disabled state.

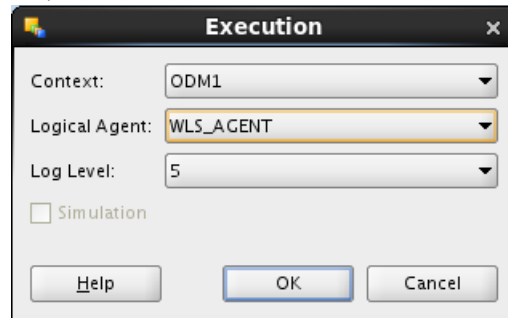
- To start loading data again, enable the ELT job for the entity using the Oracle Utilities Analytics Administration Tool.

Resetting an Instance

Perform the following steps to reset an **Instance**:

Note: The steps shown below to reset an instance should be used with care. The replication layer data should not be purged already, or else, the target does not load with the expected amount of data.

- Login to the **Oracle Data Integrator** client.
- Navigate to **Load Plans and Scenarios** section. Select the scenario named **B1_RESET_INSTANCE**. Right-click and execute.
- In the popup, select the context.
For example, to reset all entities for Oracle Utilities Operational Device Management (ODM) product instance 1, select the context **ODM1**.



- After verifying the configurations, enable all the entities to be load.

Note: When a reset instance is done on the data warehouse, the materialized view refresh jobs that are triggered subsequently can result into an error. It would result in a materialized view refresh path error sometimes. This is a known behavior of the Oracle database. To workaround this issue, do a complete refresh of all the materialized views manually.

Reloading the Extract Parameters

If you need to modify the extract parameters defined on the source, perform the following steps:

1. Modify the parameters on the source application. Refer to the section [Configuring Source](#) for additional details.
2. Regenerate the replication views for the source instance for which parameters have been updated.

For Unix:

Perform the following steps:

- a. Navigate to the <Install_Dir>/bin directory.
- b. Initialize the environment with the `./splenviron.sh -e <envname>` command.
- c. `cd $SPLEBASE/bin.`
- d. Run the `runviewGenerator.sh -c <Source Instance>`.

For Windows:

Perform the following steps:

- a. Navigate to the <Install_Dir>/bin directory.
- b. Initialize the environment with the `./splenviron.cmd -e <envname>` command.
- c. `cd %SPLEBASE%\bin.`
- d. Run the `runviewGenerator.cmd -c <Source instance>`.

Where Source instance = Instance name given to the source in the **Configure Source** step mentioned above. For example: ODM1

3. Reload all the facts for the source products using the steps mentioned in the section [Resetting an Entity](#).

Reloading the Bucket Configuration

Several key performance indicators in Oracle Utilities Analytics look at the measurement values and classify the value into an age range. The analysts can use these metrics to review the ages classified into different groups, such as 0-30 days, 30-90 days, or 90+ days. Such age ranges configuration is usually done before the implementation, but there may be scenarios where you would like to change these age ranges.

All the ELT jobs that load these bucket configurations into the warehouse are set to be initial load only. Any incremental changes to these buckets after the initial run is not captured in Oracle Utilities Analytics. Because, any changes made to the existing bucket configuration can cause inconsistency to the loaded fact data. However, if there is a need to reconfigure the buckets, certain additional steps need to be carried out to pull in the new configuration into the warehouse and ensure that this is reflected in the star schema tables.

The exact steps vary based on how the bucket configurations are used in the star schema ELT jobs. The below sections covers the two possible scenarios:

- [Buckets Referenced in Dimensions](#)
- [Buckets Referenced in Facts](#)

The exact detail of which bucket configuration falls under scenario is discussed in the section [Configuring Source](#).

Buckets Referenced in Dimensions

The typical scenario is that the configured buckets are loaded into dedicated bucket dimensions. The ELT jobs for the facts reference these buckets dimensions and map a fact record to an

appropriate bucket configuration record in the dimension. This is the typical scenario that most bucket configurations fall under.

To reload such type of bucket configurations, perform the following steps:

1. Make the necessary changes to the bucket configurations in the source application.
2. Identify the bucket dimensions for which the bucket configurations have been updated. Refer to the section [Configuring Source](#) for details.
3. The data that has been loaded into the bucket dimensions need to be cleared and reloaded with the updated configuration. Since the dimensions need to be reloaded, the related facts also need to be re-loaded. This list can be identified using the star schema mapping details provided in the *Oracle Utilities Data Mapping Guides* for the respective source product. Follow the steps mentioned in the section [Resetting an Entity](#) to reload the bucket dimensions and the related facts.
4. Reset the fact tables first, followed by resetting the bucket dimensions.
5. Enable it for the bucket dimensions first followed by the facts when enabling the ELT jobs back.

Buckets Referenced in Facts

The other scenario is that bucket configurations can be referenced by the fact ELT directly. This is a special use of the bucket configurations where a fact ELT would try to calculate a measure value based on the configured bucket ranges. Few bucket configurations fall under this scenario.

To reload such type of bucket configurations, perform the following steps:

1. Make the necessary changes to the bucket configurations in the source application.
For such type of bucket configurations where there is no dedicated bucket dimensions, the data is loaded in the metadata table B1_RANGE_LOOKUP in the MDADM schema.
2. Delete the existing rows in the MDADM.B1_RANGE_LOOKUP table for the specific bucket whose configurations are being updated.
Below is the table that can be used to identify the bucket business object name. This name is stored as the range code in the metadata table and can be used to identify the rows to be deleted.

Source Product	Bucket	Business Object Name	Fact Name
CC&B	PA Future Payment Age Configuration	C1-PAFuturePaymentAge	CF_PA, CF_PA_SNAP
CC&B	PP Future Payment Age Configuration	C1-PPFuturePaymentAge	CF_PAY_PLAN, CF_PAY_PLAN_SNAP
CC&B	SA Arrears Configuration	C1-SAArrearsBuckets	CF_ARREARS

3. Follow the steps mentioned below to reload the updated bucket configuration once the identified rows have been deleted from the metadata table:

For Unix:

- a. Navigate to the <Install_Dir>/bin directory.
- b. Initialize the environment with the command `./splenviron.sh -e <envname>`.
- c. `cd $SPLEBASE/bin.`
- d. Run the `runviewGenerator.sh -c <Source Instance>`.

For Windows:

- a. Navigate to the <Install_Dir>/bin directory.
 - b. Initialize the environment with the command `./splenviron.cmd -e <envname>`.
 - c. `cd %SPLEBASE%\bin`.
 - d. Run the `runviewGenerator.cmd -c <Source instance>`.
4. Identify the affected facts for which the bucket configurations have been updated using the above table. Refer to the section [Configuring Source](#) for details.
 5. The data that had been loaded into these facts would have been using the older bucket configuration, which gets invalidated when the configuration changes. Hence, the fact data needs to be re-loaded. Follow the steps mentioned in the section [Resetting an Entity](#) to reload the related facts.

Performance Recommendations

This section discusses how to configure the database for better system performance, including:

- [Partitioning](#)
 - [Partitioning of the Snapshot Facts](#)
 - [Partitioning of the Accumulation Facts](#)
- [Parallelism](#)
- [Optimizing the Top N Answers](#)
- [Archiving the Historical Data](#)

Partitioning

Partitioning helps to scale a data warehouse by dividing the database objects into smaller pieces, enabling access to smaller, more manageable objects. Having direct access to the smaller objects addresses the scalability requirements of the data warehouses.

It takes longer to scan a big table than it takes to scan a small table. Queries against partitioned tables may access one or more partitions that are small in contrast to the total size of the table. Similarly, queries may take advantage of partition elimination on indexes. It takes less time to read a smaller portion of an index from the disk than to read the entire index. The index structures that share the partitioning strategy with the table, such as local partitioned indexes, can be accessed and maintained on a partition-by-partition basis.

The partitioning of the facts is explained below in:

- [Partitioning of the Snapshot Facts](#)
- [Partitioning of the Accumulation Facts](#)

Partitioning of the Snapshot Facts

Most of the out of the box snapshot facts are delivered as Monthly snapshots. This means that the fact will have rows for each month from the beginning of the implementation.

All snapshot facts should be partitioned by the snapshot date. Listed below are the Snapshot facts and the column on which the partition should be created. When converting an existing fact into a partitioned fact ensure that the indexes, which do not have the partition key columns, are created as the local indexes. Ensure that the partitions are also created for the future periods so that the future snapshot loads do not fail while adding data.

For monthly snapshots primary partition based on the year and sub-partition based on the calendar month work well. For snapshots that are not monthly, different set can be used.

Fact	Column	Partition Type
CF_ARREARS	SNAPSHOT_DT	List
CF_PAY_PLAN_SNAP	SNAPSHOT_DT	List
CF_PA_SNAP	SNAPSHOT_DT	List
CF_CITY_OUTG	SNAPSHOT_DT	List
CF_CTRL_ZONE_OUTG	SNAPSHOT_DT	List
CF_FEEDER_DLVRD_LOAD	SRC_DTTM	List
CF_OUTG	SNAPSHOT_DTTM	List
CF_OPR_DEVICE_SNAP	SNAPSHOT_DT	List

Partitioning of the Accumulation Facts

The accumulation facts can also be partitioned based on a date key. However, the data in the accumulation facts will be spread across multiple dates and a list based partition will not work for the accumulation facts. A range based partitioning scheme should be employed on the primary date_key column in the specific fact.

Given below is a query which gets the start and end range values from the CD_DATE dimension. Using this range values for individual partitions can be defined.

```
select abs_month_nbr  part_name_suffix
       , min(date_key) part_range_start
       , max(date_key) part_range_end
from dwadm.cd_date
group by abs_month_nbr
order by 1;
```

Parallelism

The database can take advantage of the distinct data sets in separate partitions if you use parallel execution to speed up queries, Data Manipulation Language (DML), and Data Definition Language (DDL) statements. Individual parallel execution servers can work on their own data sets, identified by the partition boundaries.

The parallel execution enables the application of multiple CPU and I/O resources to the execution of a single database operation. It dramatically reduces response time for data-intensive operations on large databases typically associated with a Decision Support System (DSS) and data warehouses. You can also implement parallel execution on an Online Transaction Processing (OLTP) system for batch processing or schema maintenance operations, such as index creation. Parallel execution is also called 'Parallelism'. The parallelism involves breaking down a task so that instead of one process doing all of the work in a query; many processes do part of the work at the same time. For example, when four processes combine to calculate the total sales for a year, each process handles one quarter of the year instead of a single process handling all four quarters by it. The improvement in performance can be quite significant. The parallel execution improves processing for:

- Queries requiring large table scans, joins, or partitioned index scans
- Creation of large indexes
- Creation of large tables (including materialized views)
- Bulk insertions, updates, merges, and deletions

Note: For details on parallelism, partitioning, and other performance enhancement options, refer to *Oracle Database VLDB and Partitioning Guide 11g Release 2*.

Optimizing the Top N Answers

The Top N charts have to go through millions of records to find out the Top N objects that meet the criteria. Top N materialized views rearrange and partition the data so that the data reads are optimal. However, at times, depending on the amount of data, additional configuration may be required to reduce the size of the data set, which is being scanned in order to identify the Top N objects.

For example, a large sized utility with 6 million customers may have around 6 million service points in their Oracle Utilities Meter Data Management application. This means that the Consumption fact in the Oracle Utilities Meter Data Management star schema have several million records for every snapshot month. So, for each month level partition on the **Consumption Detail Level** materialized view, there will be several million records. The detail level pages (Top N Analysis and Unreported Usage Details) on the **Usage Summary** dashboard under Oracle Utilities Meter Data Analytics tries to access a huge volume of data, which can result in sluggish performance sometimes. The recommendation in such scenarios is to make a key prompt filter such as, 'City' as mandatory. This ensures the report looks at a smaller data set; thereby, improving the report performance. You can choose to make any set of filters mandatory as per your data requirement.

Archiving the Historical Data

Assuming that the facts are partitioned as per the recommendations, it is now possible to archive historical data by selecting specific partitions and archiving them to a backup. Archival rules should be based on your data retention policies. This can be set up as a scheduled activity.