Backup, Cloning and DR for Oracle Analytics Cloud

Jason Lester – Managing Architect, Capgemini UK Ltd

Future & Past TechCasts:

Nov 7th
Gimme a Vector, Victor: Leveraging Vector Datatypes for Practical Generative AI Applications
Presented by **Jim Czuprynski**

Nov 21st
Backup, Cloning and DR for Oracle Analytics Cloud
Presented by **Jason Lester**

Dec 12th
Unlocking Insights: Mastering Data Storytelling with Oracle Analytics
Presented by **Philip Godfrey**

TechCast Archive

	2024	2023	2022	2021	2020	2019

Date	Title	Presenter(s)	Replay	Download(s)
Nov 7	Gimme a Vector, Victor: Leveraging Vector Datatypes for Practical Generative AI Applications	Jim Czuprynski	Video	Slides
Oct 17	Our Favorite New Features in OAC	Dan Vlamis, Wayne Van Sluys, Cathye Pendley, Tim Vlamis, Mystery Guest: Gautam Pisharam	Video	Slides
Oct 3	Harnessing Oracle GoldenGate 23ai with AI Automation: A New Era of Intelligent Monitoring, Building, and Diagnostics	Bobby Curtis	Video	Slides
Sept 19	Discovering Oracle Fusion Data Intelligence	Peter Koutroubis & Jai Gangwani	Video	Slides
Sept 5	Oracle CloudWorld 2024 Session Analysis and Expert Agenda Recommendations	Roger Cressey, Dan Vlamis, Jim Czuprynski, Tim Vlamis, Cathye Pendley	Video	Slides
Aug 22	Our Favorite Features of OAC	Dan Vlamis, Tim Vlamis, Cathye Pendley, & Oracle Analytics Mystery Guest	Video	Slides

Submit a topic to share at **https://andouc.org/techcasts/**

www.andouc.org

Let's Connect

Website
http://andouc.org/

Chat with the Experts
https://bit.ly/Join-ANDOUC-Slack

Watch Previous TechCasts
https://bit.ly/3qmGgHN

@AnalyticAndData

https://www.facebook.com/AnDOracleUserCommunity

https://www.linkedin.com/company/analytics-and-data-oracle-user-community

Spatial + Graph SIG
bit.ly/Spatial-Graph-LinkedIn

![Analytics and Data Oracle User Community]

Save the Date

Analytics and Data
Summit 2025

Registration now **open**!

April 8-10, 2025
Oracle Conference Center
Redwood Shores, California

www.andouc.org/analytics-and-data-summit-2025/

AGENDA

1

OAC ARCHITECTURE OPTIONS

Oracle Cloud components

The standard OAC model usually includes 3 components:

- Oracle Analytics Cloud (OAC) – The visualisation and processing engine, can also act as a datastore

- Oracle Autonomous Database (ADB) – The warehouse underpinning OAC, and where a majority of your data is stored

- An ETL layer – Multiple production available, but in this presentation will reference Oracle Data Integrator (ODI). Tool is used to bring in data from multiple different sources (Extract), modify it (Transform), and store it in the ADB (Load)

- And obviously your source system(s)

Autonomous Database

Source System

ODI
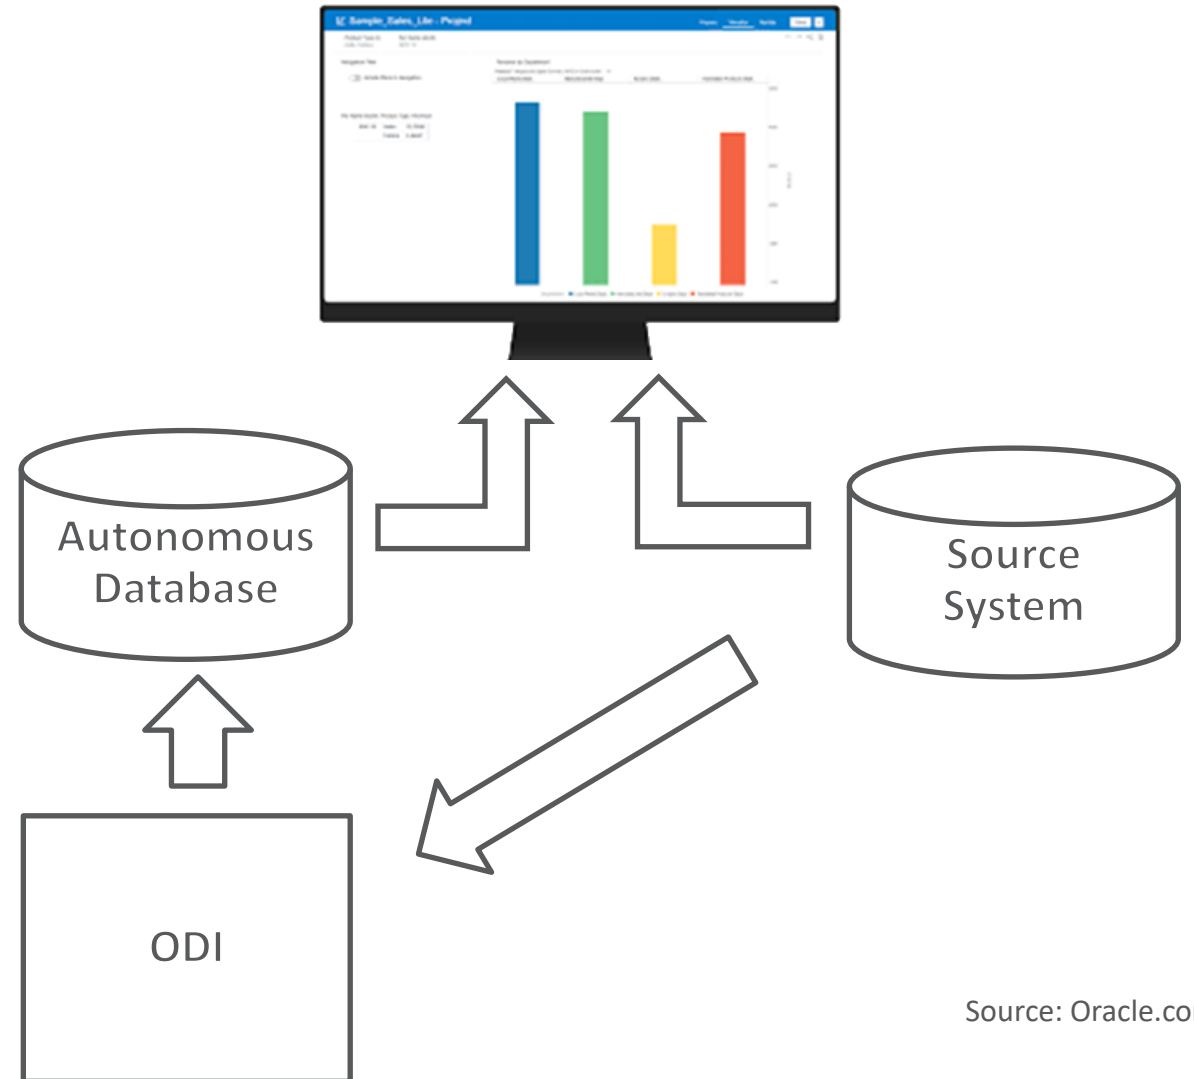
Source: Oracle.com

Warehousing Data

Why a warehouse ?

Most environments will report on data from source OLTP systems. These are optimised to manage (select, insert, update, delete) individual records very quickly.

Emp_id	Dept_id	Firstname	Lastname
1	1	Jason	Lester
2	2	Fred	Smith
3	2	Jane	Bloggs

Dept_id	Dept_name	Location
1	IT	London
2	Finance	Bristol

A warehouse wants to work with very large datasets, and so the underpinning data structure is different, and so the data is "de-normalised". This makes large queries quicker, but updates a lot slower

Data Warehouse
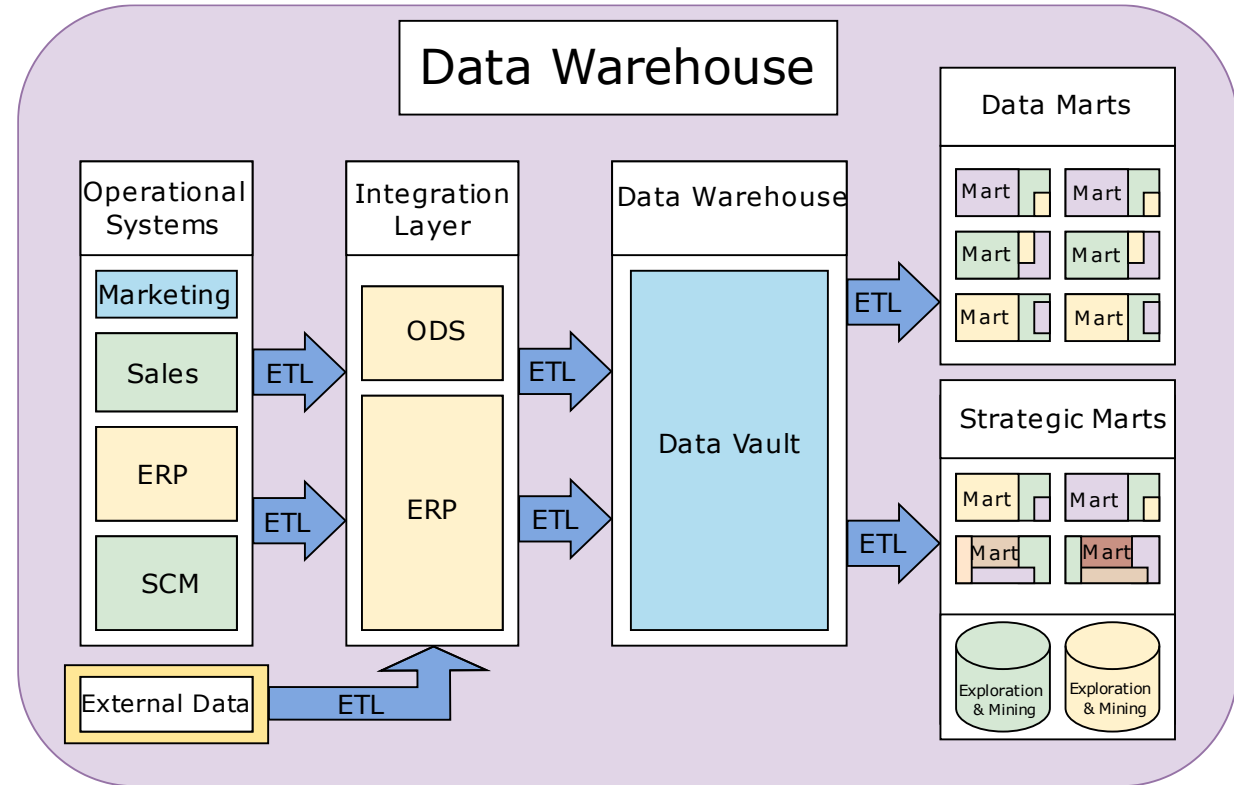
Source: Wikipedia

Emp_id	Dept_id	Firstname	Lastname	Dept_name	Location
1	1	Jason	Lester	IT	London
2	2	Fred	Smith	Finance	Bristol
3	2	Jane	Bloggs	Finance	Bristol

2

SCENARIOS

Scenarios

There are a number of different scenarios an OAC administrator will want to consider:

- Cloning (creating a new complete environment or rebuilding an existing one from another environment e.g. Production to UAT)

- Backups to revert a change if a deployment has not gone as planned

- User error – A user has done something wrong (deployed the wrong code into production, deleted something they shouldn't have etc)

- Disaster recovery to another region in OCI

Each of these scenarios is slightly different, but the processes for handling them are very similar with the execution just varying. We'll cover each in turn later

3

HANDLING EACH COMPONENT

Handling each component – The autonomous database

- Oracle's "innovative" database product.

- Based on the existing Oracle technology product.

- Includes automation for **backups**, performance, capacity management, monitoring.

- Automated cross-region disaster recovery.

- Automated cloning, restores, flashback etc. All controllable via API's so can be scripted / scheduled.

- Comes with a pre-build, managed APEX layer.

- Used to store both warehouse data and ODI config.

- Backup / recovery / cloning

  1. Using the OCI native capability – Creates a complete copy of everything, although you may need to repoint some functionality (like OAC) because it will have a new name, IP etc).

  2. Schema export / import – Selective copy of tables and structures.

- After cloning, you may need to re-run FULL ETL to build schema objects.

N.B. Store archive data separate from "live" data, to allow for replication.

Fully automated database service

Supports all modern data types, workloads including transaction processing, AI and analytics, reducing the need for multiple speciality databases.

Maximise data security and availability

With Oracle's automation services, you can keep your mission critical applications running by leveraging 99.995% availability.

Bring AI to your data

Use LLMs with built-in AI Vector Search and your proprietary data to get more accurate answers. There's no need to duplicate data to a separate vector database.

Backup

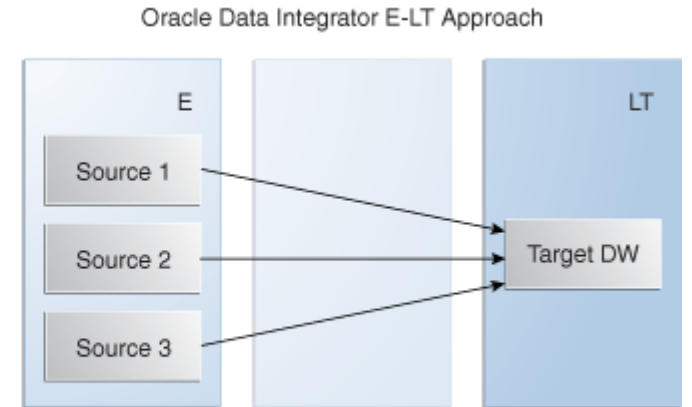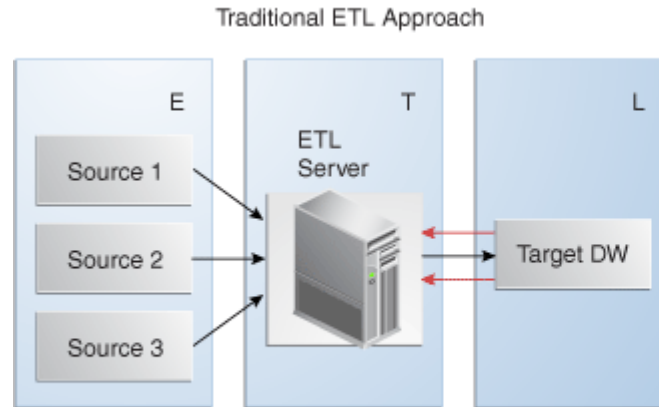**Automatic backup retention period:** 60 days Edit

Total backup storage: 85 GB

Last automatic backup: Fri, Nov 15, 2024, 10:34:10 UTC
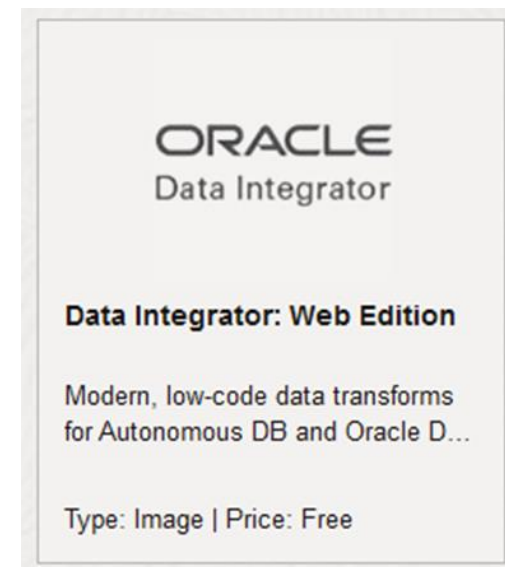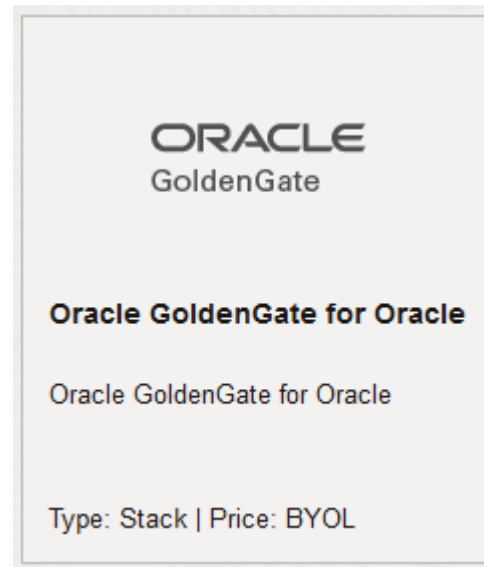
Next long-term backup: -

Long-term backup schedule: Schedule

Handling each component – ODI

- Different methodologies

  - ETL – Extract, transform (in ODI) and Load

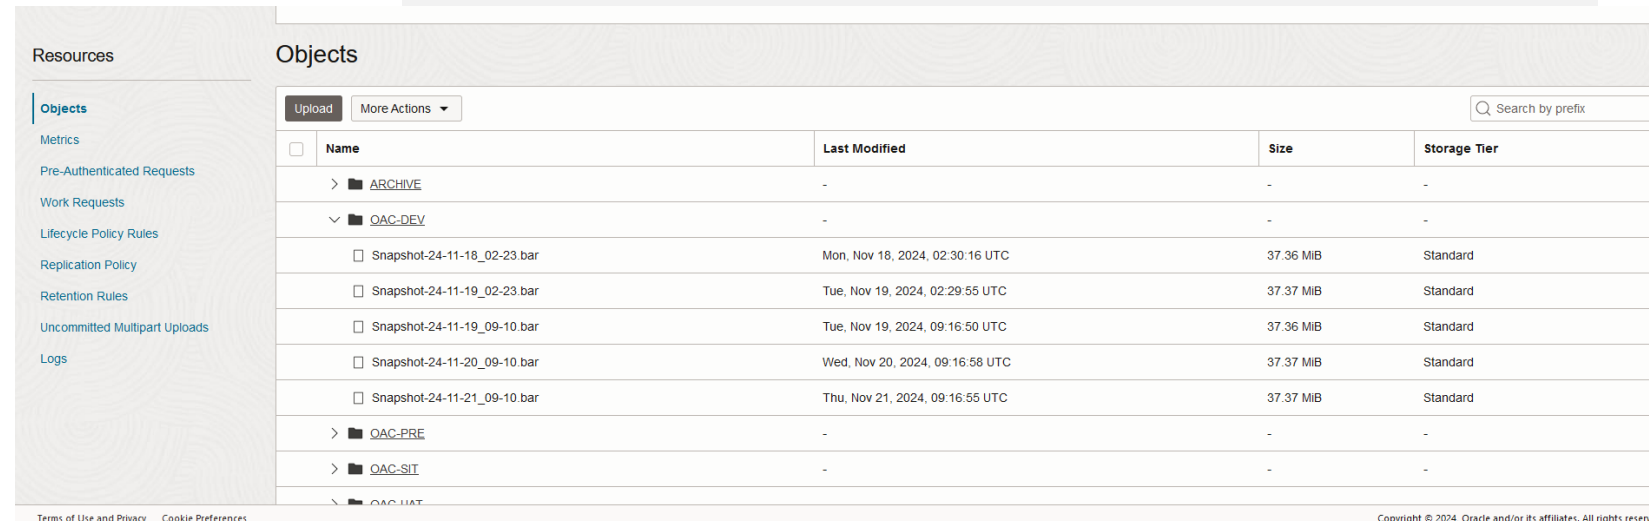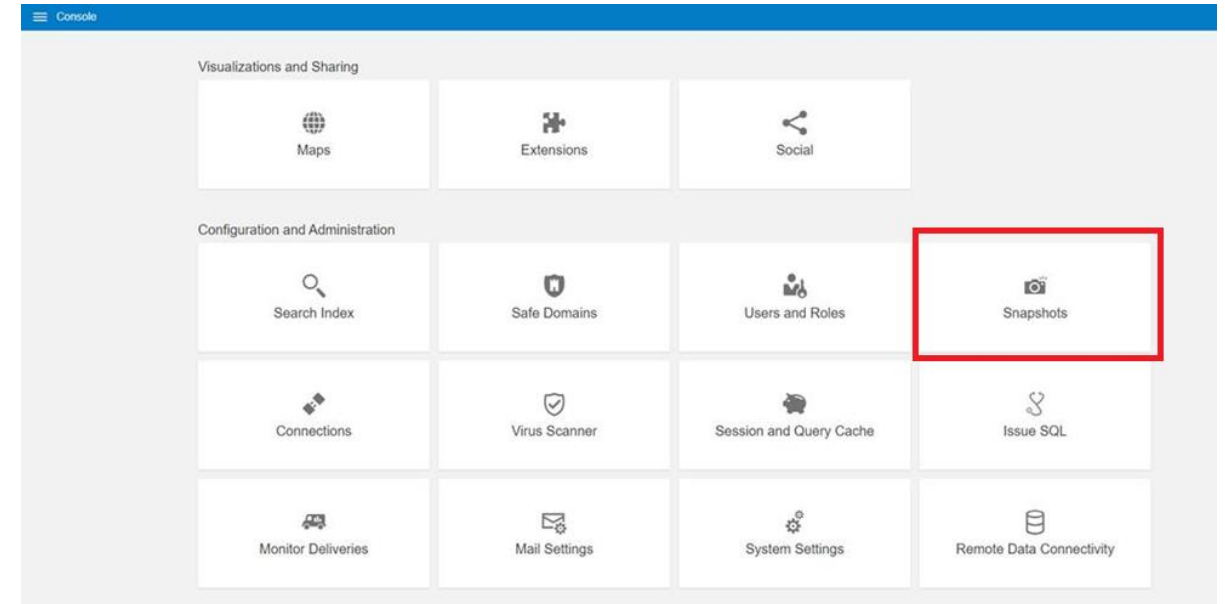 - ELT – Extract then load into target, and transform in-situ

- Deployed from the Oracle Marketplace

- Configuration store in ADB, so can be easily replicated

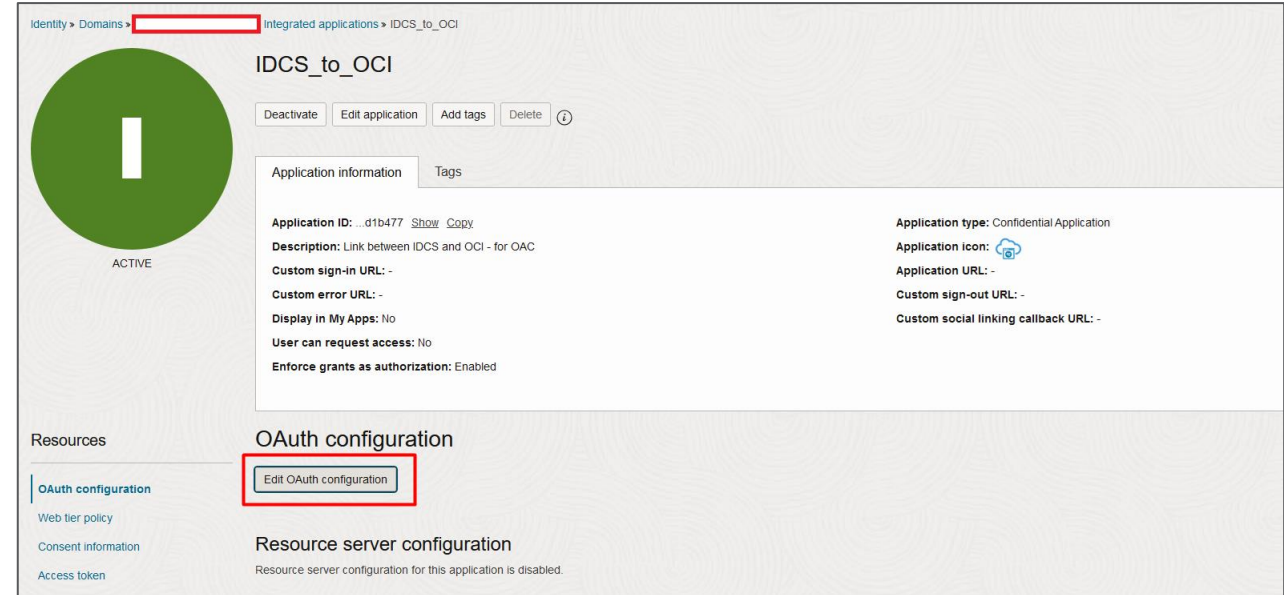- No option for incrementally updating the config when "cloning" using DB method

Traditional ETL Approach

Oracle Data Integrator E-LT Approach

ETL Server

Source 1
Source 2
Source 3

Target DW

Source 1
Source 2
Source 3

Target DW

ORACLE
GoldenGate

Oracle GoldenGate for Oracle

Oracle GoldenGate for Oracle

Type: Stack | Price: BYOL

ORACLE
Data Integrator

Data Integrator: Web Edition

Modern, low-code data transforms for Autonomous DB and Oracle D...

Type: Image | Price: Free

Handling each component – OAC

- Option inside the OAC console to take "snapshots", but these are manual, and only exist within the instance, but you can choose what to snapshot and what to restore

- REST API for OAC: https://docs.oracle.com/en/cloud/paas/analytics-cloud/acapi/quick-start.html

- When cloning OAC using the "external" snapshot, it's "all or nothing" ie. you cannot choose which components to clone.

- Saves the snapshot to OCI storage bucket, so can be used to restore to another instance (clone)

- The API can be called from a command line, meaning it can be scheduled (e.g. using crontab on Linux), giving you the ability to perform a schedule backup

Handling each component – OAC

- The API needs authenticating in the IAM domain that OAC is registered to

- Register OAC to a domain will create a new application automatically, but you need to enable it for API usage

- Click the "Edit Oauth Configuration" button

- Add a new scope to the resources (will show all available OAC instances)

- Ensure the correct Authorization options are selected

- Remember to SAVE the config before exiting IAM

- Allocate the application to an account

Handling each component – OAC

- The API:

  – Requires an "auth token" to be included in the URL

```
curl -s --request POST \
      --url https://${IDCS_URL}/oauth2/v1/token \
      --header 'authorization: Basic ${IDCS_CLIENT}' \
      --header 'content-type: application/x-www-form-urlencoded;charset=UTF-8' \
      -d \"grant_type=password&username=${API_USERNAME}&password=${API_PWD}&scope=https://${vURL}urn:opc:resource:consumer::all\"
```

  – Which generates something like this (about 2,500 characters)

```
eyJ4NXQjUzI1NiI6IktZZFZxMWlMVU1JOTJZdzB5YVZnUHo0Q1Q1M1dlTjB3T09hX3RxVVotU0kiLCJ4NXQiOiJUdjQ4TFhTlZMYUFidjdZVjhzdWZFdUtYMWciL
CJraWQiOiJTSUdOSU5HX0tFWSIsImFsZyI6IlJTMjU2In0.eyJjbGllbnRfb2NpZCI6Im9jaWQxLmRvbWFpbmFwcC5vYzQudWdWstZ292LWxvbmRvbi0xLmFtYWFhY
WFhYXF0cDViYWFlbWF5ajd0dmlua3EzbDU3N2RhNzQYmRvbmFzenB4ZXBmMnRndG53dGNscSIsIn.
.
.
spLGsWrzlAaTaKtBnk_fRtds7tn1IvFDhQfthTivwRtNKKKOVDD0jo27DLpWKXnXspDd7xI4GfGF-
5qlXZy5kXBbMoTOL_DnqGB9urc5U_FlpNDigwpAF80L29JqDGOMdYoBGUlJ7yK8bvzRr_C3DUfVon9GDs8u3Hq7Uv5s0FeLkVMPXgjU4bTKkF1XjCeNe56hKN
gG-wAVGZKXg3cA
```

Handling each component – OAC

- The API:
  - Requires an "auth token" to be included in the subsequent API calls
  - The vURL value is the "Scope" from OAM and NOT the OAC url

```
curl -s --request POST \
      --url https://${IDCS_URL}/oauth2/v1/token \
      --header 'authorization: Basic ${IDCS_CLIENT}' \
      --header 'content-type: application/x-www-form-urlencoded;charset=UTF-8' \
      -d \"grant_type=password&username=${API_USERNAME}&password=${API_PWD}&scope=https://${vURL}urn:opc:resource:consumer::all\"
```

  - Which generates something like this (about 2,500 characters). Auth tokens have a default lifetime of 100 second (configurable)

eyJ4NXQjUzI1NiI6IktZZFZxMWlMVU1JOTJZdzB5YVZnUHo0Q1Q1M1dITjB3T09hX3RxVVotU0kiLCJ4NXQiOiJUdjQ4TFhTSlZMYUFidjdZVjhzdWZFdUtYMWciL
CJraWQiOiJTSUdOSU5HX0tFWSIsImFsZyI6IlJTMjU2In0.eyJjbGllbnRfb2NpZCI6Im9jaWQxLmRvbWFpbmFwcC5vYzQudWStZ292LWxxvbmRvbi0xLmFtYWFhY
WFhYXF0cDViYWFlbWF5ajd0dmlua3EzbDU3N2RoNzN0YmRvbmFzenB4ZXBmMnRwdG53dGNscISsIn.
.
.
spLGsWrzlAaTaKtBnk_fRtds7tn1IvFDhQfthTivwRtNKKKOVDD0jo27DLpWKXnXspDd7xI4GfGF-
5qlXZy5kXBbMoTOL_DnqGB9urc5U_FlpNDigwpAF80L29JqDGOMdYoBGUlJ7yK8bvzRr_C3DUfVon9GDs8u3Hq7Uv5s0FeLkVMPXgjU4bTKkF1XjCeNe56hKN
gG-wAVGZKXg3cA
```

# Handling each component – OAC

- The API:

  - To create a snapshot, you need a JSON file

  - Then invoke the JSON file using the API

```
curl -s -i -x "http://1.2.3.4:8080" \
    --header "Authorization: Bearer $TOKEN" \
    --header "Content-Type: application/json" \
    --request POST https://${vURL}/api/20210901/snapshots \
    -d @$DIR/create_snapshot_${SOURCE}.json
```

  - This returns a WorkRequestID (or an error)

  - You can monitor the job using another API

```
curl -s -i -x "http://1.2.3.4:8080" \
    --header "Authorization: Bearer $TOKEN" \
    --request GET https://${vURL}/api/20210901/workRequests/$WorkRequestID
```

```
{
    "type": "CREATE",
    "name": "ProductionSnapshot",
    "storage": {
        "type": "OCI_NATIVE",
        "bucket": "SnapshotBUCKET",
        "auth": {
            "type": "OSS_AUTH_OCI_USER_ID",
            "ociRegion": ""uk-london-1",
            "ociTenancyId": "ocid1.tenancy.oc1..aaaaaaaalunq2y...wj5wga",
            "ociUserId": "ocid1.user.oc1..aaaaaaaado7hwuv...pou7z3oqlq",
            "ociKeyFingerprint": "be:b3:...:3b:c4:59",
            "ociPrivateKeyWrapped": "LS0tLS1CRUdJT...0VZLS0tLS0K"
        }
    },
    "bar": {
        "uri": "file:///PROD/Snapshot-24-11-17_02-23.bar",
        "password": "snapshotPWD"
    }
}
```

Handling each component – OAC

- The API:

  - Performing a restore is very similar, but you need to first register the snapshot with the OAC instance

```
curl -s -i -x "http://1.2.3.4:8080" \
    --header "Authorization: Bearer $TOKEN" \
    --header "Content-Type: application/json" \
    --request POST https://${vURL}/api/20210901/snapshots \
    -d @$DIR/create_snapshot_${SOURCE}.json
```

```
71989f79....3e558:OAC-SIT/Snapshot-24-11-17_04-23.bar
da1c174e....3741e:OAC-SIT/Snapshot-24-11-15_04-23.bar
d7e03425....41b02:OAC-SIT/Snapshot-24-11-16_04-23.bar
cf135ada....79616:OAC-SIT/Snapshot-24-11-14_04-23.bar
65222e15....21b01:OAC-SIT/Snapshot-24-11-13_04-23.bar
006e1cca....14496:OAC-SIT/Snapshot-24-11-12_04-23.bar
```

  - If the snapshot you want to restore doesn't exist, you need to register it (which requires another JSON file)

```
curl -s -i -x "http://10.210.52.164:8080" --header "Authorization: Bearer $TOKEN" --header "Content-Type: application/json" \
    --request POST https://${tURL}/api/20210901/snapshots -d @$DIR/register_snapshot_${TARGET}.json
```

  - Finally initiate the restore with the call /api/20210901/system/actions/restoreSnapshot

  - This API will also return a WorkRequestID. If the restore has been invoked successfully, this will go to 100% in about 2 minutes, but the restore may continue running for some time, depending on the size of the OAC instance, and there's no real indication of when the restore is completed ☹

  - During the restore, certain operations (like console access) will not be fully available

Warning! When you restore the snapshot, it restores everything.. Including permissions. Make sure the users you're doing this with has full admin privileges in both the target and destination system

4

PREPARING FOR THE WORST

Preparing for the worst

- Think about "What changes every day":
  - The data in the warehouse – back this up using ADB backups
  - The OAC reports, dashboards and configuration – backup this up using OAC external snapshots
  - Integration configuration – You should back this up before and after making any changes (but included in the ADB backup)
- All components can be replicated to a secondary location (another region ?)
  - ADW can be replicated real-time to a second region using standard OCI functionality
  - Integration doesn't change unless you're deploying a release, but included in ADW replication
  - OAC need to take snapshot, so should consider what is an acceptable loss of data (daily?)
    - RPD / logical layer
    - Dashboards
    - Reports
    - Custom builds

Preparing for the worst – Each Component

- Database replicated using standard OCI functionality
- ODI deployed from marketplace and connected to config in ADB
- Object storage replicated using standard OCI functionality
- OAC deployed from console as blank instance, and then loaded from replicated snapshot

5

DAY TO DAY

Day to Day

- Not just thinking about disasters.. What about human errors ?
  - Deployed code to the wrong environment
  - Deleted some data from production accidentally
    - For both these scenarios you could restore "yesterday's" backup to another environment, and extract just the small pieces of information you want, rather than restore the whole environment from a backup (reducing downtime)
  - Clones for development
    - The ability to allow your developers to always be working on a current copy of production
    - Ensure your production release (e.g. UAT) environment is as close to production as possible before doing the pre-production code deployment
    - Process can be automated, or "self-service", so developers can rebuild on their own timetable

Day to Day

- Saving money !!
  - Automate shutdown of environments overnight
- Putting the power in their hands:
  - DBA's / Admins are always a bottleneck to development
  - Giving the developers the ability to perform self-service ETL's, backups or restores (in case they want the latest data, or accidentally delete something they didn't mean to)
  - Ability to restart or change automated schedules, resize environments, and to understand the cost implications (and record who did it)

6

QUESTIONS

Helpful Links

ORACLE AUTONOMOUS CLOUD
https://cloud.oracle.com/tryit

ORACLE AUTONOMOUS HANDS ON LAB FOR DEVELOPERS
https://go.oracle.com/e/f2?LP=82486

ORACLE ANALYTICS CLOUD
Examples: https://www.oracle.com/solutions/business-analytics/data-visualization/examples.html

ORACLE ANALYTICS CLOUD (API Reference)
https://docs.oracle.com/en/cloud/paas/analytics-cloud/acapi/quick-start.html

Questions?

Or contact me on LinkedIn, Email or check out my blog

LinkedIn: https://www.linkedin.com/in/jasonrlester/
Email: Jason.Lester@Capgemini.com
Blog: https://jasonlesterdba.wordpress.com/

Capgemini

Helpful Links –

ORACLE ANALYTICS VIDEOS:

https://www.youtube.com/@OracleAnalytics/videos

OAC SEPTEMBER NEW FEATURES VIDEOS BY ORACLE: https://bit.ly/OACSept24Features

OAC NEW FEATURES DOCUMENTATION BY ORACLE:

https://docs.oracle.com/en/cloud/paas/analytics-cloud/acswn/index.html#GUID-CFF90F44-BCEB-49EE-B40B-8D040F02D476

ORACLE ANALYTICS COMMUNITY:

https://community.oracle.com/products/oracleanalytics

ORACLE ANALYTICS LIBRARY/EXAMPLES:

https://www.oracle.com/business-analytics/data-visualization/examples/

ORACLE ANALYTICS LIVE DEMOS:

https://www.oracle.com/business-analytics/data-visualization/demos/

ANALYTICS AND DATA
TechCasts

Future & Past TechCasts:

Nov 7th
Gimme a Vector, Victor: Leveraging Vector Datatypes for Practical Generative AI Applications

Presented by **Jim Czuprynski**

Nov 21st
Backup, Cloning and DR for Oracle Analytics Cloud

Presented by **Jason Lester**

Dec 12th
Unlocking Insights: Mastering Data Storytelling with Oracle Analytics

Presented by **Philip Godfrey**

TechCast Archive

| 2024 | 2023 | 2022 | 2021 | 2020 | 2019 |

Date	Title	Presenter(s)	Replay	Download(s)
Nov 7	Gimme a Vector, Victor: Leveraging Vector Datatypes for Practical Generative AI Applications	Jim Czuprynski	Video	Slides
Oct 17	Our Favorite New Features in OAC	Dan Vlamis, Wayne Van Sluys, Cathye Pendley, Tim Vlamis, Mystery Guest: Gautam Pisharam	Video	Slides
Oct 3	Harnessing Oracle GoldenGate 23ai with AI Automation: A New Era of Intelligent Monitoring, Building, and Diagnostics	Bobby Curtis	Video	Slides
Sept 19	Discovering Oracle Fusion Data Intelligence	Peter Koutroubis & Jai Gangwani	Video	Slides
Sept 5	Oracle CloudWorld 2024 Session Analysis and Expert Agenda Recommendations	Roger Cressey, Dan Vlamis, Jim Czuprynski, Tim Vlamis, Cathye Pendley	Video	Slides
Aug 22	Our Favorite Features of OAC	Dan Vlamis, Tim Vlamis, Cathye Pendley, & Oracle Analytics Mystery Guest	Video	Slides

Submit a topic to share at **https://andouc.org/techcasts/**

Analytics and Data ORACLE USER COMMUNITY

Oracle Spatial & Graph SIG

www.andouc.org

Save the Date

Analytics and Data Summit 2025

*Registration now **open**!*

April 8-10, 2025

Oracle Conference Center

Redwood Shores, California

www.andouc.org/analytics-and-data-summit-2025/