

The background of the slide is a photograph of the Prague skyline at sunset. The sky is a warm, golden-orange color. In the foreground, the dark silhouettes of various church spires and domes are visible against the bright sky. The city buildings below are also in silhouette, creating a layered effect.

Upgrade, Migration and Patching to Oracle AI Database 26ai Advanced Edition

Prague, September 2026

Photo by [Denis Poltoradnev](#) on [Unsplash](#)

Oracle

DBAs

run the world





Daniel Overby Hansen

Distinguished Product Manager

 dohdatabase

 @dohdatabase.com

 <https://dohdatabase.com>



Rodrigo Jorge

Product Management Architect

 [rodrigoaraujorge](#)

 [@dbarj.com.br](#)

 <https://www.dbarj.com.br>



Find Slides and Much More on Our Blogs



MikeDietrichDE.com

Mike.Dietrich@oracle.com



dohdatabase.com

Daniel.Overby.Hansen@oracle.com



DBArj.com.br

Rodrigo.R.Jorge@oracle.com



AlexZaballa.com

Alex.Zaballa@oracle.com



axdiaz.com

jorge.a.diaz@oracle.com

Virtual Classroom Seminars	Slides
Episode 16 (replaces Episode 1 from Feb 2021) <u>Oracle Database Release and Patching Strategy for 19c and 23c</u> 115 minutes – May 10, 2023	
Episode 17 <u>From SR to Patch – Insights into the Oracle Database Development process</u> 55 minutes – June 22, 2023	
Episode 18 <u>Cross Platform Migration – Transportable Tablespaces to the Extreme</u> 145 min – February 22, 2024	
Episode 19 <u>Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 1</u> 145 min – May 16, 2024	
Episode 20 <u>Move to Oracle Database 23ai – Everything you need to know about Multitenant PART 2</u> 100 min – June 28, 2024	
Episode 21 <u>One-Button Patching with AutoUpgrade – Easing every DBA's life</u> 55 min – October 24, 2024	

Recorded Web Seminars

<https://MikeDietrichDE.com/videos>

More than 45 hours of technical content
On-demand, anytime, anywhere

09:30

Welcome
Release Strategy
Patching

11:15

AutoUpgrade
Multitenant

13:30

Data Pump
Autonomous

15:15

Migrating
"The Beast"

11:00

Coffee break

12:45

Lunch

15:00

Coffee break

Get the Slides


<https://dohdatabase.com/slides>






Release Strategy


Lifetime Support Policy



Lifetime Support Policy



Lifetime Support Policy



Do you want to upgrade?

Oracle Database 11.2.0.4

Oracle Database 12.1.0.2

Oracle Database 12.2.0.1

Oracle Database 18c

Oracle Database 11.2.0.4
Oracle Database 12.1.0.2
Oracle Database 12.2.0.1
Oracle Database 18c



Oracle Database 19c



Oracle Database 23ai

Oracle Database 11.2.0.4
Oracle Database 12.1.0.2
Oracle Database 12.2.0.1
Oracle Database 18c



Oracle Database 19c



Oracle Database 23ai

Oracle AI Database 26ai

26^{ai}

When is a database upgrade required?

Oracle Database 19c ⇒ Oracle Database 23ai ⇒ Oracle AI Database 26ai

UPGRADE

Oracle Database 19c



Oracle AI Database 26ai

UPGRADE

Oracle Database 19c ⇒ Oracle Database 23ai ⇒ Oracle AI Database 26ai

UPDATE

July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

July 2026

Oracle AI Database 26ai

23.26.3.0.0



July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

July 2026

Oracle AI Database 26ai

23.26.3.0.0

July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

July 2026

Oracle AI Database 26ai

23.26.3.0.0

July 2025

Oracle Database 23ai

23.9.0.25.07

October 2025

Oracle AI Database 26ai

23.26.0.0.0

January 2026

Oracle AI Database 26ai

23.26.1.0.0

July 2026

Oracle AI Database 26ai

23.26.3.0.0

July 2025	Oracle Database 23ai	23.9.0.25.07
-----------	----------------------	--------------

October 2025	Oracle AI Database 26ai	23.26.0.0.0
--------------	-------------------------	-------------

January 2026	Oracle AI Database 26ai	23.26.1.0.0
--------------	-------------------------	-------------

July 2026	Oracle AI Database 26ai	23.26.3.0.0
-----------	-------------------------	-------------

July 2025	Oracle Database 23ai	23.9.0.25.07
-----------	----------------------	--------------

October 2025	Oracle AI Database 26ai	23.26.0.0.0
--------------	-------------------------	-------------

January 2026	Oracle AI Database 26ai	23.26.1.0.0
--------------	-------------------------	-------------

July 2026	Oracle AI Database 26ai	23.26.3.0.0
-----------	-------------------------	-------------



26^{ai}

When will it be available for
non-Oracle hardware on-premises?

On-Premises Server Releases (includes client)[\(Download here\)](#)

Linux x86	<i>Not Planned</i>	<i>Not Planned</i>	<i>Not planned</i>	<i>Not planned</i>	<i>Not planned</i>	<i>Not planned</i>	<i>Not planned</i>	28-Aug-2013		
Linux x86-64	January 2026 Release Update (23.26.1) (Instant Client and Full Client are available today)	13-Aug-2021	25-Apr-2019	23-Jul-2018	1-Mar-2017	22-Jul-2014	25-Jun-2013	27-Aug-2013		
Linux on Arm	<i>TBA</i>	<i>Not Planned</i>	28-June-2023 Client: May-2021 download	<i>Not Planned</i>	<i>Not Planned</i>	<i>Not Planned</i>	<i>Not Planned</i>	<i>Not Planned</i>		
Oracle Solaris SPARC (64-	<i>TBA</i>	<i>See KB126264</i>	26-Apr-2019	30-Jul-2018	1-Mar-2017	22-Jul-2014	25-Jun-2013	29-Aug-2013		

Further Information

Release Strategy



- [Release Schedule of Current Database Releases \(PNEWS1360\)](#)

Patching



How often do you apply patches?

<https://www.oracle.com/security-alerts>

Critical Patch Updates

Critical Patch Updates provide security patches for supported Oracle on-premises products. A Critical Patch Update is a collection of patches for multiple security vulnerabilities. These patches address vulnerabilities in Oracle code and in third party components included in Oracle products. These patches are usually cumulative. They are available to customers with valid support contracts. Critical Patch Updates are released on the third Tuesday of January, April, July, and October. The next four dates are:

- 20 October 2026
- 19 January 2027
- 20 April 2027
- 20 July 2027

A pre-release announcement will be published on the Thursday preceding each Critical Patch Update release.

The Critical Patch Updates released since 2021 are listed in the following table. Critical Patch Updates released before 2021 are available [here](#).

Critical Patch Update	Latest Version/Date
Critical Patch Update - July 2026	Rev 7, 20 August 2026
Critical Patch Update - April 2026	Rev 2, 24 April 2026
Critical Patch Update - January 2026	Rev 3, 02 February 2026
Critical Patch Update - October 2025	Rev 1, 21 October 2025

Critical Patch Updates

Critical Patch Updates provide security patches for supported Oracle on-premises products. A Critical Patch Update is a collection of patches for multiple security vulnerabilities. These patches address vulnerabilities in Oracle code and in third party components included in Oracle products. These patches are usually cumulative. They are available to customers with valid support contracts. Critical Patch Updates are released on the third Tuesday of January, April, July, and October. The next four dates are:

- 20 October 2026
- 19 January 2027
- 20 April 2027
- 20 July 2027

A pre-release announcement will be published on the Thursday preceding each Critical Patch Update release.

The Critical Patch Updates released since 2021 are listed in the following table. Critical Patch Updates released before 2021 are available [here](#).

Critical Patch Update	Latest Version/Date
Critical Patch Update - July 2026	Rev 7, 20 August 2026
Critical Patch Update - April 2026	Rev 2, 24 April 2026
Critical Patch Update - January 2026	Rev 3, 02 February 2026
Critical Patch Update - October 2025	Rev 1, 21 October 2025

Oracle Critical Patch Update Advisory - July 2026

Description

A Critical Patch Update is a collection of patches for multiple security vulnerabilities. These patches address vulnerabilities in Oracle code and in third party components included in Oracle products. These patches are usually cumulative, but each advisory describes only the security patches added since the previous Critical Patch Update Advisory. Thus, prior Critical Patch Update advisories should be reviewed for information regarding earlier published security patches. Refer to [Critical Patch Updates](#), [Critical Security Patch Updates](#), [Security Alerts and Bulletins](#) for information about Oracle Security advisories.

Oracle continues to periodically receive reports of attempts to maliciously exploit vulnerabilities for which Oracle has already released security patches. In some instances, it has been reported that attackers have been successful because targeted customers had failed to apply available Oracle patches. Oracle therefore strongly recommends that customers remain on actively-supported versions and apply security patches without delay.

This Critical Patch Update contains 1448 new security patches across the product families listed below. Please note that a My Oracle Support (MOS) note summarizing the content of this Critical Patch Update and other Oracle Software Security Assurance activities is located at [July 2026 Critical Patch Update: Executive Summary and Analysis](#).

Affected Products and Patch Information

Security vulnerabilities addressed by this Critical Patch Update affect the products listed below.

Oracle Communications Performance Intelligence Center, versions 10.5.0.1.0, 10.5.0.2.0	Oracle Communications Performance Intelligence Center
Oracle Communications Policy Management, version 15.0.0.0	Oracle Communications Policy Management
Oracle Communications Pricing Design Center, versions 15.0.0.0.0, 15.0.1.0.0, 15.1.0.0.0, 15.2.0.0.0	Oracle Communications Pricing Design Center
Oracle Communications Service Catalog and Design, versions 8.0.0.7.0-8.3.0.3.0	Oracle Communications Service Catalog and Design
Oracle Communications Session Border Controller, versions 9.3.0, 10.0.0, 10.1.0	Oracle Communications Session Border Controller
Oracle Communications Unified Assurance, versions 6.1.1-7.0.0	Oracle Communications Unified Assurance
Oracle Communications Unified Inventory Management, versions 7.5.0, 7.5.1, 7.6.0, 7.7.0, 7.8.0, 8.0.1	Oracle Communications Unified Inventory Management
Oracle Communications User Data Repository, version 15.0	Oracle Communications User Data Repository
Oracle Data Integrator, versions 12.3.1.4.0, 14.1.2.0.0	Fusion Middleware
Oracle Database Server, versions 19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	Database
Oracle Demanda Demand Management, versions 12.2.3-12.2.15	Oracle Supply Chain Products
Oracle E-Business Suite, versions 12.2.3-12.2.15, V16	Oracle E-Business Suite
Oracle Enterprise Communications Broker, versions 4.2.0, 5.0.0, 5.1.0	Oracle Enterprise Communications Broker
Oracle Enterprise Data Quality, versions 12.2.1.4.0, 14.1.2.0.0	Fusion Middleware
Oracle Enterprise Manager Base Platform, versions 13.5, 24.1	Oracle Enterprise Manager
Oracle Enterprise Manager for Fusion Middleware, version 13.5	Oracle Enterprise Manager
Oracle Enterprise Manager for MySQL Database, versions 13.5.4.0.0-13.5.5.0.0	Oracle Enterprise Manager

Oracle Critical Patch Update Advisory - July 2026

CVE ID	Component	Package and/ or Privilege Required	Protocol	Remote Exploit without Auth.?	CVSS VERSION 3.1 RISK (see Risk Matrix Definitions)									Supported Versions Affected	Notes
					Base Score	Attack Vector	Attack Complex	Privs Req'd	User Interact	Scope	Confid- entiality	Inte- grity	Avail- ability		
CVE-2026-61211	RDBMS	Execute DBMS_CLOUD	Oracle Net	No	9.9	Network	Low	Low	None	Changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-47040	Oracle Net Services	Connection Manager	Oracle Net	Yes	9.1	Network	Low	None	None	Un- changed	High	None	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-4738	Oracle Spatial and Graph (gdal)	None	HTTPS	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-60175	RDBMS	Authenticated User	Oracle Net	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47046	RDBMS	None	Oracle Net	Yes	8.2	Network	Low	None	None	Un- changed	None	Low	High	23.4.0-23.26.2	
CVE-2026-7383	Database (OpenSSL)	None	TLS	Yes	8.1	Network	High	None	None	Un- changed	High	High	High	23.4.0-23.26.2	
CVE-2026-54518	Fleet Patching and Provisioning (jackson- databind)	None	Multiple	No	7.5	Network	High	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47045	JDBC	None	Oracle Net	No	6.8	Network	Low	High	Required	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	

Oracle Critical Patch Update Advisory - July 2026

CVE ID	Component	Package and/ or Privilege Required	Protocol	Remote Exploit without Auth.?		CVSS VERSION 3.1 RISK (see Risk Matrix Definitions)								Supported Versions Affected	Notes
					Base Score	Attack Vector	Attack Complex	Privs Req'd	User Interact	Scope	Confid- entiality	Inte- grity	Avail- ability		
CVE-2026-61211	RDBMS	Execute DBMS_CLOUD	Oracle Net	No	9.9	Network	Low	Low	None	Changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-47040	Oracle Net Services	Connection Manager	Oracle Net	Yes	9.1	Network	Low	None	None	Un- changed	High	None	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-4738	Oracle Spatial and Graph (gdal)	None	HTTPS	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-60175	RDBMS	Authenticated User	Oracle Net	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47046	RDBMS	None	Oracle Net	Yes	8.2	Network	Low	None	None	Un- changed	None	Low	High	23.4.0-23.26.2	
CVE-2026-7383	Database (OpenSSL)	None	TLS	Yes	8.1	Network	High	None	None	Un- changed	High	High	High	23.4.0-23.26.2	
CVE-2026-54518	Fleet Patching and Provisioning (jackson- databind)	None	Multiple	No	7.5	Network	High	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47045	JDBC	None	Oracle Net	No	6.8	Network	Low	High	Required	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	

Oracle Critical Patch Update Advisory - July 2026

CVE ID	Component	Package and/ or Privilege Required	Protocol	Remote Exploit without Auth.?	CVSS VERSION 3.1 RISK (see Risk Matrix Definitions)									Supported Versions Affected	Notes
					Base Score	Attack Vector	Attack Complex	Privs Req'd	User Interact	Scope	Confid- entiality	Inte- grity	Avail- ability		
CVE-2026-61211	RDBMS	Execute DBMS_CLOUD	Oracle Net	No	9.9	Network	Low	Low	None	Changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-47040	Oracle Net Services	Connection Manager	Oracle Net	Yes	9.1	Network	Low	None	None	Un- changed	High	None	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-4738	Oracle Spatial and Graph (gdal)	None	HTTPS	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 23.4.0-23.26.2	
CVE-2026-60175	RDBMS	Authenticated User	Oracle Net	No	8.8	Network	Low	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47046	RDBMS	None	Oracle Net	Yes	8.2	Network	Low	None	None	Un- changed	None	Low	High	23.4.0-23.26.2	
CVE-2026-7383	Database (OpenSSL)	None	TLS	Yes	8.1	Network	High	None	None	Un- changed	High	High	High	23.4.0-23.26.2	
CVE-2026-54518	Fleet Patching and Provisioning (jackson- databind)	None	Multiple	No	7.5	Network	High	Low	None	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	
CVE-2026-47045	JDBC	None	Oracle Net	No	6.8	Network	Low	High	Required	Un- changed	High	High	High	19.3-19.31, 21.3-21.22, 23.4.0-23.26.2	



Be sure to [subscribe](#) to email notifications

- Critical Patch Updates and Security Alerts

Patching News

Security Updates

Accelerating Vulnerability Detection and Response at Oracle

April 29, 2026 | 3 minute read



Integrated Cyber Center (ICC)

The latest generation of AI is transforming how software vulnerabilities are identified and fixed, increasing the speed and scale of discovery and remediation.

At Oracle, we have long applied AI across our cloud and software environments to support security testing, vulnerability detection, and code analysis. These capabilities run on Oracle Cloud Infrastructure (OCI), leveraging OCI AI services, infrastructure, and development platforms to operate continuously at scale.

Oracle has access to leading frontier AI models, including Anthropic's Claude Mythos Preview and OpenAI's most capable models through Trusted Access for Cyber, and we are extending our capabilities with these models to improve how quickly and effectively vulnerabilities are identified. Combined with our AI-enabled security operations, these capabilities are applied across Oracle-developed software and services, Oracle Health, and the open-source components we build and use in our products.

The result is stronger code, earlier identification of risk and mitigations, and better protection for Oracle and our





The security practices of yesterday
are not sufficient today anymore



You must upgrade to Oracle Database 19c
or Oracle AI Database 26ai



You must patch your database
with the most recent Release Update

Older Release Updates?

File Download

1 Click each file name to download the selected files.

Tip: Use a Download Manager [Learn More...](#)

DATABASE RELEASE UPDATE 19.28.0.0.0(Patch:Linux x86)

 **p37960098_190000_LINUX.zip** (299.3 MB)

Why a Password?



Total Size 299.3 MB



Oracle also recommends
applying monthly patches

Update: Monthly Critical Security Patch Updates (CSPUs) Begin May 28, 2026




CSPUs complement Oracle's existing quarterly Critical Patch Updates (CPUs):

May 28, 2026	CSPU
June 16, 2026	CSPU
July 21, 2026	CPU
August 18, 2026	CSPU

Critical Security Patch Update

- **Monthly Critical Security Patch Updates**
Deliver timely, high-priority fixes for critical issues
- **Quarterly Release Updates**
Remain cumulative, including all fixes released in prior CSPUs

Critical Security Patch Update




Blog Posts



- [Accelerating Vulnerability Detection and Response at Oracle](#)
- [Update: Monthly Critical Security Patch Updates \(CSPUs\) Begin May 28, 2026](#)
- [Take Action Today: Protect Your Oracle Database Against AI-Enabled Cybersecurity Threats](#)
- [Recommendations to Help Protect Oracle Databases from Emerging AI-enabled Security Threats \(PNEWS3015\)](#)


Release Update Contents – Oracle Database 19c



[Database 19 Release Updates and Revisions Bugs Fixed Lists \(KB850150\)](#)



Release Update Contents – Oracle AI Database 26ai



Source: oradiff.oracle.com





If you don't apply a recent Release Update, you will miss thousands of fixes

- More than 16k fixes with 19.32.0
 - Almost 1400 security fixes
- More than 16k fixes with 23.26.3
 - More than 2300 security fixes

Just as easy as patching your smart phone

AutoUpgrade's mission for patching Oracle Database

AutoUpgrade and Patching



Download

- Find the right patch numbers
- For the right platform
- Get latest OPatch



Install

- Install brand-new Oracle home
- Update OPatch
- Apply all patches



Patch

- Datapatch Sanity Check
- Move instances and files
- Datapatch
- Recompilation
- Post-tasks



```
$ cat just_patch.cfg
```

```
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0  
patch1.sid=DB19
```

AutoUpgrade and Patching



Download

- Find the right patch numbers
- For the right platform
- Get latest OPatch



Install

- Install brand-new Oracle home
- Update OPatch
- Apply all patches



Patch

- Datapatch Sanity Check
- Move instances and files
- Datapatch
- Recompilation
- Post-tasks



```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RU,OPATCH,OJVM,DPBP
```




```
$ cat do_it_all.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RU,OPATCH,OJVM,DPBP
```

```
$ java -jar autoupgrade.jar -config do_it_all.cfg -patch -mode deploy
```


DEMO

One-Button Patching


- Download patches
- Install Oracle home
- Patch database



AutoUpgrade Patching




AutoUpgrade Patching




- Finds and downloads patches
- Uses your MOS credentials
- Chooses the right platform
- Stores patches in local repository


AutoUpgrade Patching



AutoUpgrade Patching




AutoUpgrade Patching




- Out-of-place patching
- Creates a brand-new Oracle home
- Uses *runInstaller settings* and binary options from source Oracle home
- Execute **root.sh** via sudo

AutoUpgrade Patching




- Moves database instance to new Oracle home
- Moves configuration files
- Executes Datapatch
- Updates system directories (`utlfixdirs.sql`)
- Updates `/etc/oratab`
- Supports read-only Oracle home

AutoUpgrade Patching



Patching Modes






Always download
the latest version of AutoUpgrade

- My Oracle Support [KB123450](#)

--Download the latest version of AutoUpgrade directly from oracle.com
--No authentication needed

wget https://download.oracle.com/otn-pub/otn_software/autoupgrade.jar



Patching

In-Place Patching

Oracle Home, 19.32.0



```
$ORACLE_HOME/OPatch/opatch rollback -id ...
```

Out-of-Place Patching

Oracle Home, 19.31.0



```
SQL> SHUTDOWN IMMEDIATE
```



```
[oracle]$ $ORACLE_HOME/OPatch/datapatch
```

New Oracle Home
Oracle Home, 19.32.0



Tim Hall ∞ 🤖 +∞ 📺

@oraclebase



When patching your production Oracle GI/DB installations, which method do you use?

In-Place = Current ORACLE_HOME

Out-Of-Place = New ORACLE_HOME

If you don't look after have production kit, then don't answer.

In-Place

55.4%

Out-Of-Place

44.6%



AutoUpgrade patches out-of-place



Which patches should you install?



```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RECOMMENDED
```

Recommended Patches

patch1.patch=RECOMMENDED

OPATCH

The latest OPatch

RU

The latest Release Update

OJVM

OJVM bundle matching Release Update

DPBP

Data Pump bundle patch matching RU



OJVM is embedded in Release Updates

- No separate download
- Complete RAC Rolling patching support



Stay secure with monthly Critical Security Patch Updates

- CSPUs complement quarterly Release Updates



AutoUpgrade will support
Critical Security Patch Updates (CSPU)

Monthly Recommended Patches

A collection of recommended one-off fixes
provided at monthly intervals
via a single downloadable patch

Quarterly Release Updates

	2024				2025				2026				2027		
	January	April	July	October	January	April	July	October	January	April	July	October	January	April	July
19c	19.22.0	19.23.0	19.24.0	19.25.0	19.26.0	19.27.0	19.28.0	19.29.0	19.30.0	19.31.0	19.32.0	19.33.0	19.34.0	19.35.0	19.36.0
21c	21.13.0	21.14.0	21.15.0	21.16.0	21.17.0	21.18.0	21.19.0	21.20.0	21.21.0	21.22.0	21.23.0	21.24.0	21.25.0	21.26.0	21.27.0
26ai			23.5.0	23.6.0	23.7.0	23.8.0	23.9.0	23.26.0	23.26.1	23.26.2	23.26.3	23.26.4	23.26.5	23.26.6	23.26.7



Monthly Recommended Patches

	2026												2027	
	January	February	March	April	May	June	July	August	September	October	November	December	January	February
19.30.0	19.30.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6							
19.31.0				19.31.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6				
19.32.0							19.32.0	MRP1	MRP2	MRP3	MRP4	MRP5	MRP6	
19.33.0										19.33.0	MRP1	MRP2	MRP3	MRP4
19.34.0													19.34.0	MRP1





MRPs and CSPUs are cumulative
but only within one Release Update



CSPU and MRPs will co-exist



Customers on Linux may adopt the MRP

- CSPU content is part of MRP

Content

MRP – Monthly Recommended Patch

- [KB188772](#)
 - Critical fixes
 - Regression fixes
 - RAC rolling fixes
- Security fixes

CSPU – Critical Security Patch Update

- Security fixes



Primary Note for Database Quarterly Release Updates

KB106822

Last Updated
Aug 25, 2026

Service
Oracle Database - Enterprise
Edition



2.5

Applies To

All Users

Summary

The purpose of this document is to list Database patches for both proactive and reactive maintenance.

This was MOS Document ID: 888.1 in Legacy MOS.

Solution

1.0 Overview

2.0 Oracle Recommendation

3.0 Latest Release Updates

- 3.1 Oracle Database Long Term Supported Releases
 - 3.1.1 Database 19c
 - 3.1.2 Oracle AI Database 26ai
- 3.2 Oracle Database Innovation Release
- 3.3 Oracle Database Releases Under Sustaining Support

4.0 Information On Database Patches

- 4.1 Types of Database Patches
- 4.2 Rolling Install of OJVM PSU
- 4.3 Interim Patches are generated Proactively for fixes that conflict with the latest RU Patches

Patch Type	Date Published	Version	Patch Number	Name	Known Issues	Bugs Fixed List	Patch Availability	Comments
MRP	Aug 18, 2026	19.32.0.0.260818 - Database MRP	39834034	Database MRP 19.32.0.0.260818	KB886381	KB425801	Available	Prerequisite is Database RU 19.32
	Aug 18, 2026	19.32.0.0.260818 - GI MRP	39834042	GI MRP 19.32.0.0.260818	KB886381	KB425801	Available	Prerequisite is GI RU 19.32
	Aug 18, 2026	19.31.0.0.260818 - Database MRP	39834048	Database MRP 19.31.0.0.260818	KB869205	KB425801	Available	Prerequisite is Database RU 19.31
	Aug 18, 2026	19.31.0.0.260818 - GI MRP	39834056	GI MRP 19.31.0.0.260818	KB869205	KB425801	Available	Prerequisite is GI RU 19.31
	Aug 18, 2026	19.28.0.0.260818 - Database MRP	39833833	Database MRP 19.28.0.0.260818	KB858306	KB425801	Available	Prerequisite is Database RU 19.28
	Aug 18, 2026	19.28.0.0.260818 - GI MRP	39833998	GI MRP 19.28.0.0.260818	KB858306	KB425801	Available	Prerequisite is GI RU 19.28
	Aug 18, 2026	19.32.0.0.260818	39835323	DATABASE CSPU 19.32.0.0.260818	KB886381	KB425801	AIX: ETA: 28-Aug-2026 All other ports Available	Prerequisite is Database RU 19.32



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,39834034
```



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RECOMMENDED,39834034
patch1.patch=RECOMMENDED,MRP
```



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0
patch1.sid=DB19
patch1.folder=/home/oracle/autoupgrade-patching/patch
patch1.patch=RECOMMENDED,39834034
patch1.patch=RECOMMENDED,MRP
```



You can also add specific one-off fixes

- Oracle Database 19c and Oracle AI Database 26ai
Important Recommended One-off Patches ([KB188772](#))

Oracle Database 19c and Oracle AI Database 26ai Important Recommended One- off Patches

KB188772

Last Updated
Aug 24, 2026

Service
Oracle Database - Enterprise
Edition



2.3

Summary

This note provides a list of important recommended patches that resolve issues affecting Database and Grid Infrastructure 19c and 23ai software homes. These patches resolve issues that have a high likelihood of occurring or that can cause significant impact if they do occur.

This document was MOS Document ID: 555.1 in Legacy MOS.

Solution

NOTE: Unless there is an urgent fix, this document will be updated once a week on Tuesdays at 9AM EST if any new patch is identified.

Through our review of service requests, we often find that issues encountered are the result of customers being on an older Release Update (RU) or Monthly Recommended Patch (MRP). Many issues will have been fixed in the latest updates, which we always recommend.

As noted in the RU-specific tables below, fixes for known issues are targeted for inclusion in the next available MRP and RU. You can always find the latest RUs, MRPs, other patches, lists of fixed bugs and known issues in [KB106822](#) Identifying and installing the latest updates (patches) helps ensure you are using the most current content for security, functional, regression and bug fixes, as well as minor enhancements and any emergency one-offs.

Note:

Oracle encourages customers to install the latest MRP whenever possible. MRPs deliver the safest, fully-tested path to receive the fixes in this document. Refer to [FAQ2283](#) for further details.

In addition to the relevant patches listed below, you should apply patches based on the specific RU after reviewing the following My Oracle Support knowledge documents:

1. [Database PSU/BP/Update/Revision - Known Issues Primary Note](#) [Database PSU/BP/Update/Revision - Known](#)

Recommended Patches for 19.32 DB Home

Below is the list of important patches to consider applying on top of 19.32.



Bug	KM Article	Fixed in RU	Fixed in MRP	Description	Patches	RAC Rolling Installable	Database Online Installable	Added
39890510	KI55777			[TYPE CHECK] ORA-00932 After Upgrading to 19.32 or 23.26.3	PATCH 39890510	YES	YES	24-AUG-2026
38238416 (replaces 37690446)	KI40882			[SQL EXECUTION] Bug 38238416 - ORA-600 [ktatmkref-rs] with RAC roll option for 19c	PATCH 38238416	YES	NO	30-APR-2026
38723830 (replaces 34672698)	KI55221	Not Eligible	Not Eligible	[VOS] DB50: ORA-800: soft external error, arguments: [set priority failed], [vktm] , dism(16) or ORA-7445 [__intel_sse2_strncmp()+379]	PATCH 38723830	YES	NO	31-MAR-2026
34774667	KI36203	Not Eligible	Not Eligible	[AQ] ORA-7445 in Purge Queue Table / High CPU usage in SVCB Service	PATCH 34774667	YES	NO	11-JUL-2023
29213893	KI24589	Not Eligible	Not Eligible	[QRY OPTIMIZER] DBMS_STATS Failing With Error Ora-1422 When Gathering Stats for User's Table	PATCH 29213893	NO	NO	01-SEP-2021

Version DB 19.32_555.1: 38238416,38723830,34774667,29213893

Recommended Patches for 19.32 GI Home

Below is the list of important patches to consider applying on top of 19.32.

Only one OCW (Oracle Clusterware) patch should be applied to GI homes.

Bug	KM Article	Fixed in RU	Fixed in MRP	Description	Patches	RAC Rolling Installable	Database Online Installable	Added
38723830 (replaces 34672698)	KI55221	Not Eligible	Not Eligible	[VOS] DB50: ORA-800: soft external error, arguments: [set priority failed], [vktm] , dism(16) or ORA-7445 [__intel_sse2_strncmp()+379]	PATCH 38723830	YES	NO	31-MAR-2026

Version GI 19.32_555.1: 38723830



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,38238416,38723830,34774667,29213893
```



```
$ cat DB19.cfg
```

```
global.keystore=/home/oracle/autoupgrade-patching/keystore  
patch1.source_home=/u01/app/oracle/product/19/dbhome_19_31_0  
patch1.target_home=/u01/app/oracle/product/19/dbhome_19_32_0  
patch1.sid=DB19  
patch1.folder=/home/oracle/autoupgrade-patching/patch  
patch1.patch=RECOMMENDED,OCW
```


DEMO

Applying recommended patches

- Specifying patches
- Patch database

 Watch on [YouTube](#)



How to download patches

Network Connectivity

The download stage requires:

1. Internet access
2. My Oracle Support credentials

Internet Access

AutoUpgrade connects to:

- `https://updates.oracle.com`
- `https://login-ext.identity.oraclecloud.com`
- `https://aru-akam.oracle.com`

URLs are part of a CDN, so expect changing IP addresses

- Use DNS names instead of IP addresses in your firewall

For connections via proxy use environment variables

- `https_proxy`

MOS Credentials

You must have:

- A valid My Oracle Support credential
- Connected to a Customer Support Identifier (CSI)
- Privilege to download patches using that CSI

Failure to meet the requirements:

- `*Connection Failed - You entered an incorrect user name or password.*`

MOS Credentials


For *download patches* privilege contact an MOS administrator in your organization.

- In *User Details*
- Select the appropriate *Support Identifier*
- Set the *Patches* drop-down to *Download*

Keystore

AutoUpgrade stores MOS credentials in a keystore

- Config file parameter: `global.keystore`
- Governs directory of AutoUpgrade keystore
- Password protected software keystore
- Optionally, an auto-open keystore
- No additional license needed



```
$ java -jar autoupgrade.jar ... -patch -load_password
```




```
$ java -jar autoupgrade.jar ... -patch -load_password
```

```
...
```

```
MOS> add -user <MOS username>
```


```
Enter your secret/Password:
```

```
Re-enter your secret/Password:
```



Your database host
doesn't have internet access?

Using Download Mode





You can download patches manually and place them in the *patch* folder


- Use a shared folder accessible to all hosts

Webinar

One-Button Patching

Makes life easier for every Oracle DBA

Recording on [YouTube](#)
Get the [slides](#)



Patch smarter, not harder

Makes life easier for every Oracle DBA

Recording on [YouTube](#)
Get the [slides](#)



Hands-on Lab

Patch Me If You Can

It's better to fail in our lab, than in production



[Access lab on Oracle Live Labs](#)

Datapatch

What is Datapatch?

`$ORACLE_HOME/OPatch/datapatch`

↳ `$ORACLE_HOME/sqlpatch/sqlpatch`

↳ `$ORACLE_HOME/sqlpatch/sqlpatch.pl`

```
use strict;  
use Getopt::Long;  
  
use sqlpatch.pm
```




Datapatch uses OPatch to determine patches installed in Oracle home

- From within the database via `DBMS_QOPATCH`

ORA-20001: Latest xml inventory is not loaded into table

--Avoid the use of Queryable Inventory by using OPatch directly
--Safe for single instance databases. In RAC, it's vital you ensure
--all nodes are patched and the inventories on each node is similar.

./datapatch -local_inventory

Patch Apply Sequence

datapatch



1



Java patches

2




Release Update

3



One-off patches

Datapatch | Patch Rollback and Apply Queue



Datapatch | Patch Rollback and Apply Queue

Binary Registry after opatch:

OJVM 19.31

Release Update 19.31

One-off Patch 39142906

SQL Registry before datapatch:

OJVM 19.30

Release Update 19.30

One-off Patch 29213893

Rollback:

Apply:

Rollback:

Cumulative:

Apply:

datapatch queue

Release Update 19.30 to 19.31



Purge old rollback scripts
from the data dictionary

Rollback Scripts



Apply/rollback scripts

```
$ORACLE_HOME/sqlpatch/.../nnn_apply.sql  
$ORACLE_HOME/sqlpatch/.../nnn_rollback.sql
```



Rollback scripts:

```
INSERT INTO REGISTRY$SQLPATCH  
VALUES ... ;
```

```
INSERT INTO REGISTRY$SQLPATCH_RU_INFO  
VALUES ... ;
```


PATCH_ID	RU_VERSION	RU_BUILD_TS		LOB_SIZE_MB
29517242	19.3.0.0.0	04/10/2019	12:27:20	4
30125133	19.5.0.0.0	09/09/2019	18:05:49	18
30557433	19.6.0.0.0	12/17/2019	15:50:04	24
30869156	19.7.0.0.0	04/04/2020	03:50:18	30
31281355	19.8.0.0.0	07/03/2020	03:15:01	37
31771877	19.9.0.0.0	09/30/2020	18:32:49	43
32218454	19.10.0.0.0	01/08/2021	18:50:17	51
32545013	19.11.0.0.0	04/13/2021	00:40:09	59
32904851	19.12.0.0.0	07/16/2021	14:18:10	67
33192793	19.13.0.0.0	10/04/2021	16:50:50	75
33515361	19.14.0.0.0	12/25/2021	12:21:23	83
34133642	19.16.0.0.0	07/03/2022	02:22:23	99
34419443	19.17.0.0.0	09/24/2022	22:40:51	108
34765931	19.18.0.0.0	01/11/2023	17:17:38	116
35643107	19.21.0.0.0	09/30/2023	15:19:51	141

15 rows selected.

CON_ID	LOB_SIZE_MB
1	953
2	953
3	953
...	...
4098	953

-- Available since 19.28.0. Purges no longer needed zip files.
-- Run in off-peak period after patching


`./datapatch -purge_old_metadata`

Major retailer in the US

Saved **20 TB/year** in total:

- *5000 DBs*
- *5 PDBs (3 PDBs + Root + Seed)*
- *200 Mb in each per patch*
- *4 patches annualy*

Multitenant



- Datapatch starts with CDB\$ROOT
- Datapatch sorts PDBs by *priority* and *con_id*
 - Set priority using `ALTER PLUGGABLE DATABASE ... PRIORITY`
- Datapatch determines parallel degree based on `CPU_COUNT`



Datapatch patches *PDB\$SEED* automatically

- New PDBs are ready to go
- No need to execute Datapatch on new PDBs
- Always set: `_exclude_seed_cdb_view=FALSE`



Datapatch only patches open PDBs

- `READ WRITE`, or `UPGRADE`



Unpatched PDBs will open in **RESTRICTED** mode

- Applies to failed or incomplete Datapatch runs
- Reported as plug-in violation

--Prevent PDBs from opening in restricted mode after a failed/incomplete
--Datapatch run. Quickly resolve the patching issue and re-run Datapatch.
--Use with caution.

```
alter system set "_pdb_datapatch_violation_restricted"=false;
```

Advanced Patching




Reduce downtime to the time it takes
to perform a switchover



Safely test and verify patches with
Standby-First Patch Apply

Standby-First Patching





Patch must Standby-First installable

- Release Updates (RU) are always standby-first installable
- Always check the patches' readme



Execute Datapatch on the primary database

- Only execute Datapatch when all homes are on the new patch

Standby-First Patching




- [Introduction to Patching Oracle Data Guard](#)
- [KB137118 Oracle Patch Assurance - Data Guard Standby-First Patch Apply](#)



Avoid database downtime with
RAC Rolling Patch Apply

RAC Rolling Patching



- New DB home

```
autoupgrade.jar ... -mode create_home
```

- Move to new home

```
autoupgrade.jar ... -mode deploy
```



Keep GI and DB patch levels in sync

- This is what we test and run in our Cloud

GI and DB Patch Levels

Method	Status	Example
In sync	Recommended, best practice	GI 19.32 + DB 19.32
Within two Release Updates	Supported, good practice	GI 19.32 + DB 19.30 GI 19.30 + DB 19.32
Within three or more Release Updates	Supported, not recommended	GI 19.32 + DB 19.29 GI 19.29 + DB 19.32
Different releases	Supported, use mostly for upgrades	GI 26.3 + DB 19.32

GI and DB Patch Levels

Method	Status	Example
In sync	Recommended, best practice	GI 19.32 + DB 19.32
Within two Release Updates	Supported, good practice	GI 19.32 + DB 19.30 GI 19.30 + DB 19.32
Within three or more Release Updates	Supported, not recommended	GI 19.32 + DB 19.29 GI 19.29 + DB 19.32
Different releases	Supported, use mostly for upgrades	GI 26.3 + DB 19.32



GI and DB Patch Levels

Method	Status	Example
In sync	Recommended, best practice	GI 19.32 + DB 19.32
Within two Release Updates	Supported, good practice	GI 19.32 + DB 19.30 GI 19.30 + DB 19.32
Within three or more Release Updates	Supported, not recommended	GI 19.32 + DB 19.29 GI 19.29 + DB 19.32
Different releases	Supported, use mostly for upgrades	GI 26.3 + DB 19.32



GI and DB Patch Levels

Method	Status	Example
In sync	Recommended, best practice	GI 19.32 + DB 19.32
Within two Release Updates	Supported, good practice	GI 19.32 + DB 19.30 GI 19.30 + DB 19.32
Within three or more Release Updates	Supported, not recommended	GI 19.32 + DB 19.29 GI 19.29 + DB 19.32
Different releases	Supported, use mostly for upgrades	GI 26.3 + DB 19.32



Complete a rolling patching operation
always as quickly as possible

- [FAQ7365](#)
[RAC: Frequently Asked Questions](#)



Unusual combinations are supported, but we **strongly advise against it**

- Node 1 with GI 19.32.0, node 2 with GI 19.31.0
- Patching node 1 on Monday, node 2 on Tuesday ...
- Mix of GI and DB versions on various nodes

Release updates are **always**:



Standby-First installable



RAC Rolling installable



Always patch Out-of-Place



But if you insist ...



```
du -sh $ORACLE_HOME/.patch_storage
```

```
14.2G    /u01/app/oracle/product/19/.patch_storage
```

-- Clean up Oracle home if you use in-place patching
-- or use out-of-place patching with cloned Oracle homes.
-- Delete obsolete files from \$ORACLE_HOME/.patch_storage and trims OPatch XML files

./opatch util deleteinactivepatches

Key Learnings



- 1 Patch your Oracle AI Database regularly
- 2 Patch out-of-place
- 3 Use AutoUpgrade Patching

Break

We will start again at 11:15

AutoUpgrade

Overview and new features





AutoUpgrade is the **only supported tool** to upgrade your Oracle Database

- DBUA is **desupported**



```
$ java -jar autoupgrade.jar -version
```

```
build.version 26.5.260807
```

```
build.date 2026/08/07 17:06:01 +0000
```

```
build.supported_target_versions 12.2,18,19,21,23,26
```

```
build.type production
```

```
build.MOS_NOTE KB123450
```

```
build.MOS_LINK https://support.oracle.com/...
```

--Download the latest version of AutoUpgrade directly from oracle.com
--No authentication needed

wget https://download.oracle.com/otn-pub/otn_software/autoupgrade.jar

New Features

--Upgrade RMAN catalog after upgrade

--<https://dohdatabase.com/autoupgrade-new-features-upgrade-rman-catalog-schema>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_30
```

```
upg1.target_home=/u01/app/oracle/product/26/dbhome_26_1
```

```
upg1.sid=CDB1
```

```
upg1.rman_catalog_connect_string=catalogdb
```


--Update OEM configuration after upgrade
--<https://dohdatabase.com/autoupgrade-new-features-update-enterprise-manager-configuration>

```
$ cat CDB1.cfg
```

```
upg1.source_home=/u01/app/oracle/product/19/dbhome_19_30  
upg1.target_home=/u01/app/oracle/product/26/dbhome_26_1  
upg1.sid=CDB1  
upg1.upg1.emcli_path=/u01/app/oracle/oem  
upg1.em_target_name=ORCL_myhost.domain.int
```




Upgrading and patching on Windows





```
$ cat patch_on_win.cfg
```

```
global.keystore=c:\oracle\autoupgrade\keystore  
patch1.source_home=c:\oracle\product\dbhome_19_31_0  
patch1.target_home=c:\oracle\product\dbhome_19_32_0  
patch1.sid=DB19  
patch1.folder=c:\oracle\patches  
patch1.patch=RECOMMENDED  
patch1.wincredential=c:\oracle\autoupgrade\credential
```



```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "DB19"
```



```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-loadwincredential "DB19"
```

```
C:\> java -jar autoupgrade.jar  
-config patch_on_win.cfg  
-patch  
-mode deploy
```

DEMO

Patch on Windows



[Watch on YouTube](#)

--AutoUpgrade deletes the credential file after use for security reasons
--For automation it is desirable to keep the credential file

patch1.delete_credential_file=false



AutoUpgrade fully supports upgrading
Oracle RAC Database on Windows

Refreshable Clone PDBs

Non-CDB to PDB conversion is irreversible

What are your rollback options?

ROLLBACK



Backup / restore

Ensure you have a recent backup and requires time to restore and recover



Copy data files

Requires time and disk space to hold a copy of the data files

ROLLBACK



Backup / restore

Ensure you have a recent backup and requires time to restore and recover



Copy data files

Requires time and disk space to hold a copy of the data files



Refreshable clone PDB

Requires ~~time and~~ disk space to hold a copy of the data files

Requires Oracle Database 12.2 or newer

Refreshable Clone PDB



CREATE

Create PDB from non-CDB over a database link



REFRESH

Apply redo from non-CDB to keep PDB up-to-date



OUTAGE

Disconnect users and refresh PDB for the last time




CONVERT


To become a proper PDB, it must be converted

Refreshable Clone PDB

system01.dbf
sysaux01.dbf
users01.dbf
undo01.dbf
...



Source
non-CDB



Target
CDB


Refreshable Clone PDB

system01.dbf
sysaux01.dbf
users01.dbf
undo01.dbf
...



*Could be same
server as well*

Refreshable Clone



Refreshable Clone PDB

Source non-CDB

Target CDB



```
CREATE USER dblinkuser  
  IDENTIFIED BY ... ;  
  
GRANT CREATE SESSION,  
  CREATE PLUGGABLE DATABASE,  
  SELECT_CATALOG_ROLE TO dblinkuser;  
  
GRANT READ ON sys.enc$ TO dblinkuser;
```

```
CREATE DATABASE LINK CLONEPDB  
  CONNECT TO dblinkuser  
  IDENTIFIED BY ...  
  USING 'noncdb-alias';
```

Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_32
upg1.target_home=/u01/app/oracle/product/dbhome_26_3
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=26
upg1.source_dblink.NONCDB1=CLONEPDB
upg1.target_pdb_name.NONCDB1=PDB1
```



You can drop user and database link after migration

- Config file parameter `drop_dblink`

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_32
upg1.target_home=/u01/app/oracle/product/dbhome_26_3
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=26
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_32
upg1.target_home=/u01/app/oracle/product/dbhome_26_3
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=26
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
```



Rename your PDB to [avoid name collision](#)

- If CDB is on same host,
it also registers for the default service

Refreshable Clone

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_32
upg1.target_home=/u01/app/oracle/product/dbhome_26_3
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=26
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=19/01/2038 03:14:07
```

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> proceed -job 101

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> proceed -job 101

--Or postpone it
upg> proceed -job 101 -newstarttime +2h30m

--When a job is in REFRESHPDB stage,
--you can force it to start immediately
upg> proceed -job 101

--Or postpone it
upg> proceed -job 101 -newstarttime +30m

--Or reschedule it
upg> proceed -job 101 -newstarttime 13/09/2026 02:30:00

Refreshable Clone PDB

Source non-CDB

Target CDB



```
upg1.source_home=/u01/app/oracle/product/dbhome_19_32
upg1.target_home=/u01/app/oracle/product/dbhome_26_3
upg1.sid=NONCDB1
upg1.target_cdb=CDB1
upg1.target_version=26
upg1.source_dblink.NONCDB1=CLONEPDB 300
upg1.target_pdb_name.NONCDB1=PDB1
upg1.start_time=19/01/2038 03:14:07
upg1.parallel_pdb_creation_clause=4
```

Refreshable Clone

1

Run on source


```
autoupgrade.jar ... -mode analyze  
autoupgrade.jar ... -mode fixups
```

2

Run on target

```
autoupgrade.jar ... -mode deploy
```

Refreshable Clone PDB



`autoupgrade.jar ... -mode deploy`

`upg1.start_time=19/01/2038 03:14:07`

`upg> proceed -job n`



Works for **unplug-plug** upgrades as well



The source non-CDB stays **intact**
to allow rollback



**Techniker
Krankenkasse**

Move hundreds of Oracle Database 19c PDBs to
Oracle Database 23ai with ExaScale-only



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Large German Public Health Insurance

- **Founded:** 1884
- **Legal Form:** Public Corporation
- **Headquarters:** Hamburg, Germany
- **Members:** 11.9 million
- **Employees:** aprx. 19,000
- **Customer Service:** 174 service centers
- **Budget 2024:** €62.5 billion

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Move to Oracle Database 23ai with ExaScale

- No ASM, no Oracle Database 19c anymore



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

COMPATIBLE=23.0.0

No ASM, no Oracle Database 19c anymore

Data Pump as rollback

Customer Case | Techniker Krankenkasse

Customer

Project


Constraints

Preparation

Upgrade

Success?

Remarks



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation


Upgrade

Success?

Remarks

AutoUpgrade:

```
global.autoupg_log_dir=/home/oracle/autoupgrade_logs
upg1.source_home=/u01/app/oracle/product/19.0.0/dbhome_1
upg1.target_home=/u01/app/oracle/product/23.0.0/dbhome_1
upg1.sid=CDB19
upg1.pdbs=PDBSOURCE
upg1.target_cdb=CDB23
upg1.target_pdb_name,PDBSOURCE=PDB23
upg1.source_dbLink,PDBSOURCE=COPYPDB_PDB23_PREPARE 60
upg1.timezone_upg=NO
upg1.close_source=NO
upg1.keep_source_pdb=yes
upg1.target_pdb_copy_option,PDBSOURCE=file_name_convert=NONE
```



Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation


Upgrade

Success?

Remarks



Oracle 19.28



Standby

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Yes!!

- Migrated 150 PDBs from TEST & DEV within less than 4 weeks
- Migration of all Oracle Database 19c instances completed by end of 2025
- Exascale snapshots are super-fast

Customer Case | Techniker Krankenkasse

Customer

Project

Constraints

Preparation

Upgrade

Success?

Remarks

Standby building

- There were some issues
- Refreshable Clones don't propagate
- PDB recovery isolation doesn't work as expected

A number of SRs had to be opened for various areas

Optimizer works very well

In the cloud, upgrading Oracle databases
is as **easy** as pressing a button



The **easy** approach



Upgrade entire CDB, including all PDBs



The **fast** approach



Upgrade single PDB

Upgrading in the cloud



Blog post:

[Upgrade Encrypted PDB in Cloud to Oracle AI Database 26ai](#)

Blog post:

[Upgrade Oracle Base Database Service to Oracle AI Database 26ai](#)

Blog post:

[Upgrade from 19c to 26ai using AutoUpgrade –ExaDB-D or ExaC@C – Part 1](#)

Blog post:

[When A Refreshable Clone Takes Over The Service](#)

Blog post:

[Upgrade to Oracle AI Database 26ai](#)

Documentation: [Proceed command](#)



Refreshable clone works only with deferred recovery on standby database

- You must restore the PDB on standby database after disconnect from non-CDB

Multitenant

Create Container Database



1 Character set

2 Components

3 COMPATIBLE

Create Container Database

1 Character set

- Always choose AL32UTF8
- Allows PDBs with any character set

2 Components

3 COMPATIBLE

Database Configuration Assistant - Create 'orcl' database - Step 9 of 15

Specify Configuration Options

ORACLE 26^{ai}
AI Database

Memory Sizing Character sets Connection mode

The database character set determines how character data is stored in the database.

☒ Use Unicode (AL32UTF8)
Setting character set to Unicode (AL32UTF8) enables you to store multiple language groups.

☐ Use OS character set (WE8MSWIN1252)
Character set is based on the language setting of this operating system.

☐ Choose from the list of character sets

Database character set: AL32UTF8 - Unicode UTF-8 Universal character set

☒ Show recommended character sets only

National character set: AL16UTF16 - Unicode UTF-16 Universal character set

Default language: American

Default territory: United States

Create Container Database

1 Character set

2 Components

- Install as many as you need
- No more than that

3 COMPATIBLE

Database Configuration Assistant - Create 'orcl' database - Step 8 of 15

Select Database Options

ORACLE 26^{ai}
AI Database

Database components

Select the standard database components you want to configure for use in your database. Oracle recommends that you always install these components in your database. Deselecting these components may cause you to no longer be able to choose some components on the subsequent page.

Select Component	Tablespace	Include in PDBs
☐ Oracle JVM	SYSTEM	☐
☐ Oracle Text	SYSAUX	☐
☐ Oracle OLAP	SYSAUX	☐
☐ Oracle Spatial	SYSAUX	☐
☐ Oracle Label Security	SYSTEM	☐
☐ Oracle Database Vault	SYSAUX	☐

Create Container Database

1 Character set

2 Components

3 COMPATIBLE

- Keep at **23.6.0** to benefit from AI Vector Search features
- Unless you want the option of downgrade

All initialization parameters

⚠ Update the initialization parameters only when it is required. Refer to the Oracle documentation to learn more about each initialization parameter and its valid set of values.


(Storage related parameter(s) value is shown in MB) ☐ Show advanced parameters

Name	Value	Include in spfile	Category
undo_tablespace	UNDOTBS1	☒	Cluster Database
sga_target	19202	☒	SGA Memory
db_block_size (bytes)	8192	☒	Cache and I/O
nls_language	AMERICAN	☒	NLS
control_files	{*{ORACLE_BASE}/oradata/...	☒	File Configuration
remote_login_passwordfile	EXCLUSIVE	☒	Security and Auditing
processes	640	☒	Processes and Sessions
pga_aggregate_target	6401	☒	Sort, Hash Joins, Bitmap Indexes
nls_territory	AMERICA	☒	NLS
open_cursors	300	☒	Cursors and Library Cache
db_domain	livelabs.oraclevcn.com	☒	Database Identification
compatible	23.6.0	☒	Miscellaneous
db_name	orcl	☒	Database Identification
cluster_database	FALSE	☐	Cluster Database

Description:

compatible: Allows you to use a new release, while at the same time guaranteeing backward compatibility with an earlier release.
Range of Values: Default to current release. Default Value: Release dependent

Compatible



- You can now update the COMPATIBLE parameter without a restart
- Works within a release, e.g., from 23.6.0 to 23.26.0
- Cross-release changes, e.g., 19.0.0 to 23.6.0, still requires a restart

alter database set downgrade compatibility to '23.26.0';

Further Information

Container Database

YouTube:

[Create CDB in Oracle AI Database](#)

Blog post:

[How I Think You Should Create Your Container Databases](#)



[COMPATIBLE Parameter - Documentation](#)

Blog post:

[Online COMPATIBLE raise starting with 23.9](#)

```
SQL> select con$name, comp_id, status from cdb_registry;
```

CON\$NAME	COMP_ID	STATUS
-----------	---------	--------

-----	-----	-----
-------	-------	-------

CDB\$ROOT	CATALOG	VALID
-----------	---------	-------

CDB\$ROOT	CATPROC	VALID
-----------	---------	-------

CDB\$ROOT	XDB	VALID
-----------	-----	-------

RED	CATALOG	VALID
-----	---------	-------

RED	CATPROC	VALID
-----	---------	-------

RED	XDB	VALID
-----	-----	-------

No information on
PDB\$SEED



```
alter system set "_exclude_seed_cdb_view"=false;
```

```
alter system set "_exclude_seed_cdb_view"=false;
```

```
SQL> select con$name, comp_id, status from cdb_registry;
```

CON\$NAME	COMP_ID	STATUS
-----	-----	-----
CDB\$ROOT	CATALOG	VALID
CDB\$ROOT	CATPROC	VALID
CDB\$ROOT	XDB	VALID
PDB\$SEED	CATALOG	VALID
PDB\$SEED	CATPROC	VALID
PDB\$SEED	XDB	VALID
RED	CATALOG	VALID
RED	CATPROC	VALID
RED	XDB	VALID



How does a standby database handle a plug-in operation?

- Take care of your Data Guard during migration

Data Guard



*Plug-in on primary propagates
to standby database via redo*



- 1 Enabled recovery
- 2 Deferred recovery

Enabled Recovery

1

Enabled recovery

```
create pluggable database ... standbys=all
```

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB


PDB is immediately protected

Default

2

Deferred recovery

Enabled Recovery





All data files on primary and standby
must be at the same SCN

Enabled Recovery

- The plug-in happens on the primary database
- The plug-in uses the *manifest file*
- The manifest file contains information on data files from the primary database only

```
<?xml version="1.0" encoding="UTF-8"?>
<PDB>
  <xmlversion>1</xmlversion>
  <pdbrname>DB12</pdbrname>
  <cid>0</cid>
  <byteorder>1</byteorder>
  <vsn>203424000</vsn>
  <vsns>
    <vsnum>12.2.0.1.0</vsnum>
    <cdbcompt>12.2.0.0.0</cdbcompt>
    <pdbrcompt>12.2.0.0.0</pdbrcompt>
    <vslibnum>0.0.0.0.24</vslibnum>
    <vsnsql>24</vsnsql>
    <vsnsbv>8.0.0.0.0</vsnsbv>
  </vsns>
  <dbid>1852833295</dbid>
  <ncdb2pdb>1</ncdb2pdb>
  <cdbid>1852833295</cdbid>
  <guid>8605DC2587337002E0532AB2A8C0A57C</guid>
  <uscnbas>4437941</uscnbas>
  <uscnwrp>0</uscnwrp>
  <undosc>8</undosc>
  <rdba>4194824</rdba>
  <tablespace>
    <name>SYSTEM</name>
    <type>0</type>
    <tsn>0</tsn>
    <status>1</status>
    <issft>0</issft>
    <isnft>0</isnft>
    <encts>0</encts>
    <flags>0</flags>
    <bmbunitsize>8</bmbunitsize>
    <file>
      <path>/u02/oradata/DB12/system01.dbf</path>
      <afn>1</afn>
      <rfn>1</rfn>
```

Enabled Recovery

- The plug-in happens on the primary database
- The plug-in uses the *manifest file*
- The manifest file contains information on data files from the primary database only

How does the standby database know which files to plug in?

Enabled Recovery

How does the standby database know which files to plug in?

- 1 Regular files
- 2 OMF in regular file system
- 3 ASM

Enabled Recovery

1 Regular files

- Standby search for data files at the same location as the primary
- Override with `DB_FILE_NAME_CONVERT`
- Or, override with `STANDBY_PDB_SOURCE_FILE_DIRECTORY`

Enabled Recovery

2 OMF in regular file system

- Standby search for data files at the OMF location (**DB_CREATE_FILE_DEST**)
- Move data files from non-CDB location into OMF location
- Or, create soft links in OMF location pointing to data file location

Enabled Recovery

3 ASM

- Standby search for data files at the OMF location (**DB_CREATE_FILE_DEST**)
- Use ASM aliases to find data files at non-CDB OMF location

Enabled Recovery | ASM

Primary



Standby



```
SQL> select name from v$datafile;  
  
NAME  
-----  
+DATA/DB_BOSTON/DATAFILE/system.269.1103046537  
+DATA/DB_BOSTON/DATAFILE/sysaux.270.1103046537  
+DATA/DB_BOSTON/DATAFILE/users.273.1103046827
```

```
SQL> select name from v$datafile;  
  
NAME  
-----  
+DATA/DB_CHICAGO/DATAFILE/system.265.1103050007  
+DATA/DB_CHICAGO/DATAFILE/sysaux.266.1103050007  
+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009
```

Same file,
but different name

Enabled Recovery | ASM

19c
Non-CDB
Primary



19c
Non-CDB
Standby




The manifest file contains

```
SQL> select dbms_pdb.describe('step1/manifest_DB.xml');
```

- Not standby database

```
<?xml version="1.0" encoding="UTF-8"?>
<PDB>
  <xmlversion>1</xmlversion>
  <pdbname>PDB1</pdbname>
  ...
  <guid>DDB49CFEFD8ED4FCE053E801000A078C</guid>
  ...
  <tablespace>
    <name>USERS</name>
    ...
  <file>
    <path>+DATA/DB_BOSTON/DATAFILE/users.273.1103046827</path>
```

Enabled Recovery | ASM



Enabled Recovery | ASM

26ai
CDB
Primary



26ai
CDB
Standby



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827

Redo record says:
Plug in this data file

No good, data file
has a different name

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

Enabled Recovery | ASM

26ai
CDB
Primary



+DATA/DB_BOSTON/DATAFILE/users.273.1103046827



26ai
CDB
Standby



OK, let's check the OMF directory

+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009

+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE

It's empty



Could I just move the file in ASM??



There's no **move** command in ASM.
But you can create *aliases*

- Similar to file system soft links



```
SQL> alter diskgroup data add alias
      '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'
for
      '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```

Non-CDB standby data file



```
SQL> alter diskgroup data add alias  
      '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'  
for  
      '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```



```
SQL> alter diskgroup data add alias  
      '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'  
for  
      '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```



Standby PDB OMF location

```
SQL> alter diskgroup data add alias  
      '+DATA/DB_CHICAGO/DATAFILE/users.269.1103050009'  
for  
      '+DATA/CDB1_CHICAGO/<PDB_GUID>/DATAFILE/users.269.1103050009':
```



Name does not matter.
Standby scans all files in OMF directory

Data Guard | Re-use Data Files

26ai
CDB
Primary



26ai
CDB
Standby



- Standby database scans its own OMF directory for data files
- Standby ignores file names and look at file headers
- Standby will find aliases and find the real file locations

Data Guard | Re-use Data Files

Primary



Standby



Looking for file like on primary



```
Recovery scanning directory +DATA/DB_BOSTON/... for any matching files
Deleted Oracle managed file +DATA/DB_BOSTON/...
Successfully added datafile 37 to media recovery
Datafile #37: +DATA/DB_CHICAGO/DATAFILE/users.269.1103050009
```



Follows alias and finds the real file



What happens with enabled recovery
if the standby fails to find the data files?

Enabled Recovery | Missing Data Files

What if a standby database fails to find data files?

- If Active Data Guard and PDB Recovery Isolation is turned on
 - New feature in Oracle Database 21c
 - Recovery disabled for PDB
 - Recovery proceeds in the entire CDB, except in specific PDB
 - Standby automatically restores data files from primary and re-enables recovery afterward
 - PDB protected after auto-restore
- If not, recovery halts in the **entire** CDB
 - **This is a critical situation**



What about AutoUpgrade
and enabled recovery?

----- Continue with the manual steps -----

There is a job with manual steps pending.

The checkpoint change number is 2723057 for database DB19.

For the standby database <DB19_AARHUS>, use the checkpoint number <2723057> to recover the database.

You can find the SCN information in:

/home/oracle/autoupgrade/upgrade26/DB19/101/drain/scn.json

Once these manual steps are completed, you can resume job 101

Deferred Recovery

1

Enabled recovery

```
create pluggable database ... standbys=all
```

Standby records PDB creation

Standby locates data files

MRP applies redo to PDB

PDB is immediately protected

2

Deferred recovery

```
create pluggable database ... standbys=none
```


Standby records PDB creation

Standby ignores data files



MRP skips redo

PDB protected after restore


Deferred Recovery



Deferred Recovery



```
SQL> create pluggable database ...  
standby=None;
```



Deferred Recovery



```
SQL> create pluggable database ...  
standbys=none;
```



PDB created
Data files missing

Deferred Recovery



```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	READ	WRITE

```
SQL> show pdbs
```

CON_NAME	OPEN	MODE
PDB1	MOUNTED	

Deferred Recovery



```
SQL> select name, recovery_status  
       from v$pdb;
```

NAME	RECOVERY_STATUS
PDB1	DISABLED

Deferred Recovery



```
RMAN> restore pluggable database  
... from service ... ;
```

```
SQL> alter pluggable database  
enable recovery;
```

```
SQL> alter database datafile  
... online;
```

Deferred Recovery



```
RMAN> restore pluggable database  
... from service ... ;
```

```
SQL> alter pluggable database  
enable recovery;
```

```
SQL> alter database datafile  
... online;
```

- Automated process in Oracle Database 21c
- PDB Recovery Isolation
- Requires Active Data Guard

Further Details

Multitenant Conversion



- [Move to Oracle Database 23ai – Everything you need to know about Oracle Multitenant – Part 1](#)



Don't jeopardize your Data Guard

- Test the procedure and verify your environment



Operations



Share resources between PDBs

Resource Consolidation

Non-CDB
database



Memory




Background processes




Files



Resource Consolidation




Resource Consolidation



Consolidation Strategies?

There is no "*best*" strategy



Using a Swingbench benchmark,
a single-core machine could host **nine non-CDBs**
before reaching 75 % CPU utilization

By going multitenant the number of databases reached **123 PDBs**

A US Health Care provider managed to

- Reduce the number of database instances by 7x
- Reduce the number of physical servers by 50%



You can run multiple CDBs on the same host and out of the same Oracle home

Consolidation



Schema consolidation



Virtual Private Database



PDB consolidation

- Less complexity
- Better isolation
- Operational benefits
- Easier cloning

A global provider of financial services states

*The multitenant architecture gives us **complete client separation out of the box**, without having to maintain a Virtual Private Database setup.*

We went away from Virtual Private Database and consolidated our different clients in individual PDBs.

*This reduced the complexity of our database implementation and **made operations much easier**.*



Avoid *noisy neighbors*

- Allow sharing resources
but everyone must get a fair share

Method 1



Instance caging

- Most simple
- Define `CPU_COUNT` for each PDB
- Hard limit

Method 1



8 CPUs



CPU_COUNT=3



CPU_COUNT=2



CPU_COUNT=2

←
All non-CDBs
share 7 CPU

Method 1



8 CPUs



CPU_COUNT=7

CDB never uses more than 7 CPUs,
despite sum of PDBs
At peak, use more resources,
but never deplete the CDB



CPU_COUNT=4



CPU_COUNT=4



CPU_COUNT=4

PDBs might fight over CPUs,
but each process gets a fair share

Method 2



Memory allocation

- Simple
- Define **SGA_TARGET** for each PDB
- Hard limit

Method 2



8 GB memory



SGA_TARGET=7G



SGA_TARGET=4G



SGA_TARGET=4G



SGA_TARGET=4G

PDB may never use more than 4G
of shared memory



If all PDBs are active,
cache management comes into play



Requires use of Automatic Shared Memory Management

- Both in CDB and PDB



Optionally, allocate minimum shared pool and buffer cache for a PDB

- Use `SHARED_POOL_SIZE` and `DB_CACHE_SIZE`



You can combine method 1 and 2

- Instance caging and memory allocation

Method 3



Simple Resource Manager

- Elaborate, yet simple to implement
- Enable CDB resource manager
- Allocate minimum shares instead of hard limits
- For advanced use cases

Method 3



8 CPUs



CPU_COUNT=7



CPU_MIN_COUNT=2



CPU_MIN_COUNT=1



CPU_MIN_COUNT=1

At peak, may use up to 5 CPUs



4 CPUs are reserved,
3 are free for all

Method 3



8 GB memory



SGA_TARGET=7G



SGA_MIN_SIZE=2G



SGA_MIN_SIZE=1G



SGA_MIN_SIZE=1G

At peak, may use
up to 4G shared memory





Requires Resource Manager at root level

```
alter session set container=cdb$root;
```

```
-- Create an empty resource manager plan with no directives
```

```
exec dbms_resource_manager.clear_pending_area;
```

```
exec dbms_resource_manager.create_pending_area;
```

```
exec dbms_resource_manager.create_cdb_plan('CDB_PLAN');
```

```
exec dbms_resource_manager.validate_pending_area;
```

```
exec dbms_resource_manager.submit_pending_area;
```

```
-- Make plan active in root to enable CDB resource manager
```

```
alter system set resource_manager_plan=CDB_PLAN;
```

Method 4



Advanced Resource Manager

- Requires additional configuration, but much greater control
- Use directives instead of shares



You can still control resources inside a
PDB with Resource Manager



What about I/O?

- Exadata I/O Resource Management
- Or `MAX_MBPS` and `MAX_IOPS`



You can run multiple CDBs on the same host and out of the same Oracle home

Inter-instance Resource Management

Shares resources like with non-CDBs:

- CPU_COUNT
- SGA_MAX_SIZE

Inter-instance CPU resource manager:


- Controls resource sharing using Linux c-groups
- Check [RESOURCE MANAGER CPU SCOPE](#)
- Exadata Database Machine and Autonomous AI Database






A word about automated maintenance tasks

Automated Maintenance Tasks




Automated Maintenance Tasks



--Change the amount of PDBs that can run maintenance tasks at the same time
--Default value 2

```
alter system set autotask_max_active_pdb=3;
```


Automated Maintenance Tasks



--Selectively disable maintenance tasks in a PDB
--For instance, test databases or databases that are rebuilt frequently

```
alter session set container=PDB2;  
alter system set enable_automatic_maintenance_pdb=false;
```

Automated Maintenance Tasks







Shift maintenance windows

- Optionally, shorten maintenance windows

Automated Maintenance Tasks



Automated Maintenance Tasks





Selectively disable individual maintenance tasks using `DBMS_AUTO_TASK_ADMIN`

- Does a test database need Automatic Segment Advisor?
- Or Evolve Advisor?



Resource Manager prevents maintenance tasks from *stealing* resources from users

- Consumer group `ORA$AUTOTASK`

Key Learnings



- 1 Get started with Multitenant
- 2 Understand plug-in on standby databases
- 3 Finetune your container database

Lunch Break

We will start again at 13:30

Data Pump Top Tips

Supercharge data loading/unloading



Always use the Data Pump Bundle Patch



More than 230 functional
and performance fixes

*Importing a complete application with data
drops from almost 2.5 hours to 48 minutes –
by just applying the Data Pump bundle patch*

A global provider of financial services



Apply the Data Pump Bundle Patch
without downtime

DEMO

Apply Data Pump Bundle Patch



[Watch on YouTube](#)

Further Information | Data Pump Bundle Patch



[\(KB107134\)](#)

[Data Pump Recommended Proactive Patches For 19.10 and Above](#)

The patch is not RAC rolling installable

- The patch is non-binary online installable
- Apply while the database instance is running
- Don't use Data Pump or DBMS_METADATA



Use PARALLEL and multiple dump files

--Apply parallelism by simply specifying a degree

```
expdp ... parallel=8
```

--Use different parallel degree on import

```
impdp ... parallel=32
```



Oracle Cloud Infrastructure

Number of ECPUs / 4

Number of OCPUs




On-prem (x86-64)

2 x physical cores



On-prem (other)

Depends



Use at least 32 ECPUs when importing into Autonomous AI Database

- Use the elasticity in Autonomous AI Database and scale even higher

Parallel Architecture

expdp ... parallel=4

Parallel Architecture

expdp ... `parallel=4`



`select * from t1`

`select * from t2`

`select * from t3`

`select * from t4`

Control process

Worker processes

Parallel Architecture

expdp ... parallel=4



Control process

Worker processes

select /*+ parallel(2) */ * from t1

select * from t2

select * from t3

select * from t4

Worker 4 goes idle

--Use %L to allow multiple dump files
expdp ... parallel=8 dumpfile=exp%L.dmp



--Split dump files into minor files for easier transport
expdp ... parallel=8 dumpfile=exp%L.dmp **filesize=10000M**

- After export, store a checksum in the dump file.
- Detects in-flight corruption or alteration.
- Specify other algorithms using `checksum_algorithm` parameter.

```
expdp ... checksum=yes
```

```
impdp ... verify_checksum=yes  
         verify_only=yes
```



For **best protection** against dump file tampering, use encrypted dump files

- Requires Advanced Security Option

- Protect your dump files from alteration by using encryption
- Creating an encrypted dump file requires a license for ASO

```
expdp ... encryption=all encryption_algorithm=AES256
```

- Protect your dump files from alteration by using encryption
- Creating an encrypted dump file requires a license for ASO

```
expdp ... encryption=all encryption_algorithm=AES256
```

DEMO

Encrypted exports




[Watch on YouTube](#)



Transportable jobs **can** use parallel metadata from Oracle Database 21c

- But you need to specify **PARALLEL**



- Any transportable jobs can now run in parallel
- Parallel unload/load of metadata provide a significant
- performance boost

```
expdp ... full=y transportable=always parallel=16
```

```
impdp ... parallel=16
```


Parallel Transportable Benchmark

Oracle E-Business Suite database

600.000+ objects

Export parallel 1	2h 2m	Export parallel 16	1h 8m
Import parallel 1	6h 44m	Import parallel 16	1h 23m
Total	8h 46m	Total	2h 31m





Speed up imports by using
NOVALIDATE constraints

A Constraint Can Be

VALIDATED

All data in the table obeys the constraint.
The database guarantees that data is good.

NOT VALIDATED

All data in the table **may** obey the constraint.
The database **does not know** if data is good.



Most constraints are **VALIDATED**



On import, Data Pump creates constraints
in the same state as in the source

--Example of which commands Data Pump import might execute as part of an import

```
create table sales ( .... );
```

```
insert into sales as select ... ;
```

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable validate;
```

```
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable validate;
```

```
alter table sales add constraint c_sales_3 check (c3 > 0) enable validate;
```

Recursive full table scan

Recursive full table scan

Recursive full table scan



```
-- Add constraints with NOVALIDATE keyword regardless of state in source database  
-- Significantly speeds up add constraints for larger tables
```

```
impdp ... transform=constraint_novalidate:y
```

--Transforming constraints to NOVALIDATE to speed up import

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```

No full table scan



--Transforming constraints to NOVALIDATE to speed up import

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;  
alter table sales add constraint c_sales_2 check (c2 in ('A','B')) enable novalidate;  
alter table sales add constraint c_sales_3 check (c3 > 0) enable novalidate;
```

Database validates new rows

Benchmark, 1 billion rows

Importing VALIDATE constraints

```
10-AUG-24 00:32:28.716: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:36:42.762: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:45:41.226: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:55:35.787: W-1      Completed 7 CONSTRAINT objects in 594 seconds
```

Importing NOVALIDATE constraints

```
10-AUG-24 00:14:56.050: W-1 Processing object type TABLE_EXPORT/TABLE/TABLE_DATA
10-AUG-24 00:19:10.311: W-1 . . imported "FUSION"."hwr_topic_t1" 151.2 GB 1044625000 rows in 254 seconds using external_table
10-AUG-24 00:29:20.841: W-1 Processing object type TABLE_EXPORT/TABLE/CONSTRAINT/CONSTRAINT
10-AUG-24 00:29:21.101: W-1      Completed 7 CONSTRAINT objects in 1 seconds
```

Real-World Example, Importing 180G Schema

Starting point	44m 50s
Adding Data Pump Bundle Patch	36m 53s
Using NOVALIDATE constraints	7m 25s



NOVALIDATE constraints prevent the optimizer from certain **query rewrites**

Validate constraints after import, or even after go-live

- Still requires a full scan of the table
- But can use parallel query
- And no table lock!

```
alter table sales add constraint c_sales_1 check (c1 in (0,1)) enable novalidate;
```

```
-----
```

```
----- GO LIVE -----
```

```
-----
```

```
-- Validate constraints
```

```
-- Optionally, use parallel query
```

```
alter session force parallel query;
```

```
alter table sales modify constraint c_sales_1 enable validate;
```



Also available in Oracle Database 19c
via 19.27 Data Pump Bundle Patch



Also available in
Oracle Autonomous AI Database 19c



```
--Add the following to your ZDM response file to set the parameter during  
--ADB migrations using ZDM
```

```
DATAPUMPSETTINGS_METADATATRANSFORMS-1=name:CONSTRAINT_NOVALIDATE,value:1
```

Further Information | NOVALIDATE Constraints



- 19.26 and earlier also requires patch 37280692
- Allow optimizer to perform query rewrites even with NOVALIDATE constraints using [QUERY REWRITE INTEGRITY](#)
- Data Pump always validates certain constraints:
 - On DEFAULT ON NULL columns
 - Used by a reference partitioned table
 - Used by a reference partitioned child table
 - Table with Primary key OID
 - Used as clustering key on a clustered table



Use index size to determine
parallel degree on index creation


Index Creation

```
impdp ... parallel=16
```

Before 12.1

Worker 1 CREATE INDEX PARALLEL 16

Really good for few big indexes



Index Creation

```
impdp ... parallel=16
```

From 12.1

Worker 1	CREATE INDEX PARALLEL 1
Worker 2	CREATE INDEX PARALLEL 1
...	CREATE INDEX PARALLEL 1
Worker 16	CREATE INDEX PARALLEL 1

Really good for many small indexes



Index Creation

```
impdp ... parallel=16
```

From 26

Worker 1	CREATE INDEX PARALLEL 1
Worker 2	CREATE INDEX PARALLEL 8
Worker 3	CREATE INDEX PARALLEL 4
Worker 4	CREATE INDEX PARALLEL 3

The best of both worlds



How Data Pump Create Indexes

- 1 Calculate the optimal parallel degree
- 2 Create indexes

How Data Pump Create Indexes

1 Calculate the optimal parallel degree

- Always parallel 1 when a table is less than 150 MB
- Customizable via `INDEX_THRESHOLD`
- Get optimal parallel degree using `EXPLAIN PLAN`



```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> explain plan for create index i1 on t1(c1) parallel;
```

Explained.

```
SQL> select * from table(dbms_xplan.display(format => 'ALL'));
```

...

Note

- automatic DOP: Computed Degree of Parallelism is 4 because of degree limit
- estimated index size: 655K bytes


How Data Pump Create Indexes

2 Create indexes

- One worker creates small indexes (parallel 1) in large batches
- The next worker starts with the biggest index (measured by optimal parallel degree)

How Data Pump Create Indexes

impdp ... parallel=16



Importing with former index method

```
10-MAY-25 16:18:55.130: W-12 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX  
10-MAY-25 16:36:46.902: W-30      Completed 480 INDEX objects in 1071 seconds
```

Importing with new index method

```
10-MAY-25 16:47:50.267: W-4 Processing object type SCHEMA_EXPORT/TABLE/INDEX/INDEX  
10-MAY-25 16:59:17.006: W-3      Completed 480 INDEX objects in 686 seconds
```

Further Information | Index Creation

- Requires:
 - Oracle AI Database 26ai
 - Oracle Database 23.8 plus Data Pump Bundle Patch
 - Oracle Database 19.31 plus Data Pump Bundle Patch



Also available in Oracle Database 19c
via 19.31 Data Pump Bundle Patch



Convert time zone data on import

Time Zone File Version Check



Source
Version 43

Target
Version 42

```
create table t1 (  
  ...  
  c1 timestamp with timezone  
  ...  
)
```

Import: Release 19.0.0.0.0 - Production on Sun Sep 30 06:17:06 2025
Version 19.27.0.0.0

ORA-39002: invalid operation

ORA-39405: Oracle Data Pump does not support importing from a source database with TSTZ version 43 into a target database with TSTZ version 42.

DEMO

Convert time zone data on import

 [Watch on YouTube](#)



Further Information | Time Zone Data Conversion



- If needed, Data Pump converts `TIMESTAMP WITH TIMEZONE` to the target time zone file version
- This works if the target time zone file is higher or lower than the source
- Expect a small overhead for the conversion
- Requires 19.27 including Data Pump Bundle Patch



Also available in Oracle Database 19c
via 19.27 Data Pump Bundle Patch




Copy the dump files
before the export completes

- Oracle AI Database 26ai

Copy Dump Files



Copy Dump Files





```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);
```

```
for i in l_status.job_status.files.first..l_status.job_status.files.last() loop
```

```
end loop;
```

```
dbms_datapump.get_status(l_handle, l_stsmask, NULL, l_job_state, l_status);

for i in l_status.job_status.files.first..l_status.job_status.files.last() loop

    l_dump_file := l_status.job_status.files(i);

    if (l_dump_file.file_bytes_written = l_dump_file.file_size) then
        dbms_output.put_line('DONE: ' || l_dump_file.file_name);
    end if;

end loop;
```

Further Information | Copy Dump Files

- Blog post:
[Copy Data Pump Files Before the End of the Export](#)





Use diagnostics views to assist in troubleshooting

- Oracle AI Database 26ai

```
SQL> select waiting_session, event, dp_state_in_wait
       from   v$datapump_sessionwait_info;
```

WAITING_SESSION	EVENT	DP_STATE_IN_WAIT
10	direct path sync	WAITING
77	log buffer space	WAITING
191	log buffer space	WAITING
428	enq: TT - contention	WAITING

Further Information | Diagnostic Views



New views available in Oracle AI Database 26ai:

- v\$datapump_process_info
- v\$datapump_sessionwait_info
- v\$datapump_processwait_info

Key Learnings



- 1 Use Data Pump Bundle Patch
- 2 Use NOVALIDATE constraints
- 3 Create indexes faster

Autonomous AI Database Migrations

Made easy – a sneak preview


KEYNOTE PRESENTATION

Oracle OpenWorld San Francisco 2017




Recap AutoUpgrade

End-to-end automation



Recap AutoUpgrade

End-to-end automation




- Small, flexible, agile, available on all platforms
- Prechecks, pre/post-fixups
- Used by thousands of customers
- Deployed in every Oracle home

Why shouldn't AutoUpgrade migrate databases
to Oracle Autonomous AI Database Serverless?



Overview



What if we combine these tools together?
And there it is, **AutoMigrate!!**

AutoMigrate

What we are going to do



ADB migrations for every DBA

Simple and well-known interface


Available everywhere

Works with Windows as source platform


Can be embedded in any automation framework

Can be embedded into other tools


Conceptual Overview



Architecture Overview




Architecture Overview



Load passwords into CPAT keystore (SEPS)

- ADMIN user
- Data Pump dump file encryption password (optional)
- Token for objectstore (if no pre-auth URL is used)

Architecture Overview



global.**keystore**=/home/oracle/AUTOMIGRATE_DEMO/KEYSTORE

global.**username**=ADMIN

demo26ai.**source_home**=/opt/oracle/product/26ai/dbhomeFree

demo26ai.**sid**=FREE

demo26ai.**src_pdb_name**=FREEPDB1

demo26ai.**src_pdb_service_name**=localhost:1521/FREEPDB1


demo26ai.**target_service_name**=amademo26_tpurgent

demo26ai.**target_wallet_location**=/home/oracle/AUTOMIGRATE_DEMO/WALLET


demo26ai.**curl_target_url**=https://swiftobjectstorage.eu-frankfurt-1.oraclecloud.com/v1/oradbcclouducm/AMADemo/

demo26ai.**curl_target_user_name**=ORACLE-SS0/klaus.gronau@oracle.com


Architecture Overview




Architecture Overview



Architecture Overview



Architecture Overview



AutoMigrate Sneak Preview



Download required

- **AutoUpgrade** from oracle.com or MOS
- **CPAT** from oracle.com or MOS
- **Wallet** for Autonomous Serverless from cloud.oracle.com
- **Token** for the Object Storage

DEMO

AutoMigrate

- Migrate a local database of 15GB into Oracle Autonomous AI Database



Demo

Regression testing DB with plenty of objects

- Size: 15GB
- Duration: 20 mins – end-to-end migration

```
[oracle@47544c5be592 100]$ more autoupgrade_20260310.log
2026-03-10 07:16:26.809 INFO [23]
build.MOS_LINK:https://support.oracle.com/support/?anchorId=&kmContentId=2485457&page=sptemplate&sptemplate=km-article
....
2026-03-10 07:32:44.057 INFO [27] Adding Job context key[dispatcher_stage] value[4] for job [100] - CommonBackBone.addJobContext#270
2026-03-10 07:32:44.057 INFO [27] Stopping background threads for Job - JobDefinition.stopJobThreads#132
2026-03-10 07:32:44.058 INFO [27] [0] threads have been stopped by JobDefinition - JobDefinition.stopJobThreads#134
2026-03-10 07:32:44.058 INFO [27] Dispatcher finished for 100 - RunJobDefinition.runJob#177
2026-03-10 07:32:44.058 INFO [27] End - RunJobDefinition.runJob#178
2026-03-10 07:32:44.059 INFO [27] -----Stages Summary----- - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27]      SETUP <1 min - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27]      MIGRATIONCPATANALYZE <1 min - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27]      MIGRATIONCPATFIXUPS <1 min - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27]      MIGRATIONAMA 16 min - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27]      COMPLETED <1 min - Utilities.writeStageSummary#1153
2026-03-10 07:32:44.059 INFO [27] End of dispatcher instance for N/A - CommonBackBone.finalJobLogging#436
```



Prototype is already fully operational

- If you want to [beta test](#) from July 1 on, send mail to: mike.dietrich@oracle.com

DBA vs. AI

Harvard Business Review

***“AI won’t replace humans -
but humans with AI will
replace humans without AI”***

Karim Lakhani





AI Database




Using AI


Vector Database



Vector Database



Vector Database



Vector Database





AI Database



Using AI

Let us create an
AutoMigrate Data Pump Log Analyzer
using Oracle AI Database 26ai

What do we need?

Oracle AI Database 26ai

- On-premises or cloud or [Oracle Autonomous AI Database](#)

Transforming model

- [Turning text into Vectors](#), e.g. [Hugging Face's all_MiniLM_L12_v2.onnx](#)

LLM


Table with VECTOR column and VECTOR index

Predefined rules for errors

Analyzing hundreds of Data Pump migration log files



Analyzing hundreds of Data Pump migration log files



It may take a while to load, classify and report

```
Processed 1000 rows; last severity=low
Processed 2000 rows; last severity=medium
Processed 3000 rows; last severity=low
Processed 4000 rows; last severity=medium
Processed 5000 rows; last severity=low
Processed 6000 rows; last severity=medium
Processed 7000 rows; last severity=low
Processed 8000 rows; last severity=medium
Processed 9000 rows; last severity=low
Processed 10000 rows; last severity=medium
Processed 11000 rows; last severity=low
Processed 12000 rows; last severity=medium
Processed 13000 rows; last severity=low
Processed 14000 rows; last severity=medium
Processed 15000 rows; last severity=low
Processed 16000 rows; last severity=low
Done. Total processed=16846
```

Elapsed: 00:32:01.43

Report uploaded successfully.
Object URI:
https://.../oracle_error_severity_report.html

Elapsed: 00:00:26.75

Example report from test migration with Fusion Apps (FA)

Oracle Error Severity Report

Status: WARN

Generated: 2026-05-06 13:10:53.042732000 +00:00

Files: 151 Errors: 37413 Critical: 0 Medium: 18087 Low: 19326 Unclassified: 0 AI: 12226 Exact: 0 Vector: 25187

Critical 0 Medium 18087 Low 19326 Unclassified 0



AI 12226 Exact 0 Vector 25187



Scope: all errors



Example report from test migration with Fusion Apps (FA)

IMPDP (125 file(s))				
	critical 0	medium 5	low 13126	errors 13131
00018_PDB1_IMPDP_SCHEMA_FUSION_260430_202439.log	critical 0	medium 3	low 12357	errors 12360
00034_PDB1_IMPDP_SCHEMA_AZ_MIG_260430_202439.log	critical 0	medium 2	low 2	errors 4
00022_PDB1_IMPDP_SCHEMA_FUSION_GRC_260430_202439.log	critical 0	medium 0	low 323	errors 323
00030_PDB1_IMPDP_SCHEMA_FUSION_SETUP_260430_202439.log	critical 0	medium 0	low 208	errors 208



Example report from test migration with Fusion Apps (FA)

medium rows 2, messages 2							
Seq	Error code	Error message	Source	Confidence	Action	Context / SQL snippet	
1	ORA-31685	ORA-31685: Object type SYSTEM_GRANT:"AZ_MIG"."EXPORT FULL DATABASE" failed due to insufficient privileges. Failing sql is: GRANT EXPORT FULL DATABASE TO "AZ_MIG" W-1 Completed 11 SYSTEM_GRANT objects in 0 seconds W-1 Processing object type SCHEMA_EXPORT/ROLE_GRANT	ai	.8	investigate	ORA-31685: Object type SYSTEM_GRANT:"AZ_MIG"."EXPORT FULL DATABASE" failed due to insufficient privileges. Failing sql is: GRANT EXPORT FULL DATABASE TO "AZ_MIG" W-1 Completed 11 SYSTEM_GRANT objects in 0 seconds W-1 Processing object type SCHEMA_EXPORT/ROLE_GRANT	
2	ORA-39083	ORA-39083: Object type ROLE_GRANT failed to create with error:	ai	.8	investigate	ORA-39083: Object type ROLE_GRANT failed to create with error:	



Example report from test migration with Fusion Apps (FA)

low rows 323, messages 323

Seq	Error code	Error message	Source	Confidence	Action	Context / SQL snippet
1	ORA-31684	ORA-31684: Object type FUNCTION:"FUSION_GRC"."GRC_ACCESS_TYPE_B=" already exists	vector_cache	.95	ignore	ORA-31684: Object type FUNCTION:"FU SION_GRC"."G RC_ACCESS_TY PE_B=" already exists
2	ORA-31684	ORA-31684: Object type FUNCTION:"FUSION_GRC"."GRC_ACCESS_TYPE_TL=" already exists	vector_cache	.95	ignore	ORA-31684: Object type FUNCTION:"FU SION_GRC"."G RC_ACCESS_TY PE_TL=" already exists



What would help you?



Quick Demo

- Microsoft VS Code
- Github Copilot Chat
- SQL Developer Web Extensions for VS Code
- Any LLM

EXTENSIONS

Search Extensions in Marketplace

INSTALLED

**Cline**
Autonomous coding agent right in your IDE, capable of crea...
Cline 504ms

**GitHub Copilot Chat**
AI chat features powered by Copilot
GitHub 404ms

**Oracle Code Assist**
Oracle Code Assist
Oracle

**Oracle SQL Developer Extension for VSCode**
End-to-end development tools for your SQL and PL/SQL ap...
Oracle Corporation 91ms

**GitHub Copilot**
Your AI pair programmer
GitHub

RECOMMENDED

**SQLTools**
Connecting users to many of the most commonly used data...
Matheus Teixeira 6M 3.5
Install

**Vim**
Vim emulation for Visual Studio Code
vscodvim 8.3M 4
Install

MCP SERVERS

**MCP Servers**
Browse and install [Model Context Protocol \(MCP\)](#)

Untitled-1 X

1 `Generate code (M), or select a language (M). Start typing to dismiss or don't show this again.`

CHAT



Build with Agent

AI responses may be inaccurate.

[Generate Agent Instructions](#) to onboard AI onto your codebase.

Untitled-1

Describe what to build next

Agent GPT-5 mini

Ln 1, Col 1 Spaces: 4 UTF-8 LF () PL/SQL (Dedicated) ADB26ai

Key Learnings



- 1 Precise prompting is important
- 2 Never use a read/write user
- 3 Keep your brain turned on – always!!

Break

We will start again at 15:15

Cross-Platform Migrations


Pushing the limits



The Beast

Introduction

Who is who?





ANDREAS GROETZ

Oracle DBA Tech Lead

Entain Services Austria GmbH

Entain is one of the world's largest sports betting and gaming groups. Leveraging the power of the Entain Platform, they bring moments of excitement into their customers lives through more than 30 iconic brands such as bwin, Coral, Ladbrokes and many more.

Entain operates on over 140 licenses across 40+ territories and employs over 29,000 talented workforce. Entain is listed on the London Stock Exchange and is a constituent of the FTSE 100 Index.

Entain



Ladbrokes



BET**MGM**

sportingbet

CORAL 
GET CLOSER TO THE ACTION

bwin 

EUROBET 


SuperSport

party ker

Foxy BINGO

Challenges

What is special, what makes it so complex?



Migration Challenges



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges

180TB
size



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges

15TB
redo/day



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash

Migration Challenges



SPARC SuperCluster



ZDLRA



Exadata X9M Extreme Flash




5 Physical Standby DBs

Local, and in different region, 2500km away

Constraints

Limiting factors, and other things to know





Up to 15TB redo/day is beyond what Oracle GoldenGate will be able to synch


The system is highly active 24 x 7 x 365



Photo by Dave Hoefler on Unsplash

Very large database, very constrained downtime

- 180+ TB database size
- 5-6 TB growth/month
- Every minute of downtime costs \$\$\$



Migrating tablespaces upfront or separately **definitely** not an option

- Way too many cross-dependencies
- Tablespaces aren't isolated



Every complex Oracle data type you can imagine is used

- XML binary types
- Nested partitioned tables
- Evolved object types



Tight downtime window

- Dry run: 2 hours outage approved
 - Tablespace read-only
 - Full Transportable Export
- Live migration: 13 hours approved



The available **Oracle V4 PERL** migration scripts would have been worked technically, but ...

- No section-size backup support
- No standby backup support
- No selective PDB migration support


Migration






Typically, we use Full Transportable Export/ Import for large cross-endian migrations

Concept



Scripts for Incremental Backup Automation

Backup / restore / recover



- No multisection backups
- No encrypted tablespaces
- Inefficient parallelism
- Incomplete multitenant support




M5 is the next-generation cross-platform transportable tablespace procedure

- New RMAN functionality combined with Full Transportable Export/Import
- MOS [KB144840](#)

M5 Migration Script

The new migrations scripts superseding the V4 PERL scripts



```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP
    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;
}
```

```
# source database
RUN
{
  ALLOCATE CHANNEL d1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL d2 DEVICE TYPE DISK FORMAT '...';
  BACKUP
    FILESPERSET 1
    SECTION SIZE 64G
    TAG UP19_L0_240206101548
    TABLESPACE <list-of-tablespace>;
}
```

target database

RUN

```
{
  ALLOCATE CHANNEL DISK1 DEVICE TYPE DISK FORMAT '...';
  ALLOCATE CHANNEL DISK2 DEVICE TYPE DISK FORMAT '...';
  RESTORE ALL FOREIGN DATAFILES TO NEW FROM BACKUPSET
    '<backup-set-1>',
    '<backup-set-2>',
    ...
    '<backup-set-n>'
};
```


Benefits

M5 procedure supports:

- Encrypted tablespaces
- Multisection backups
- Migrating multiple databases into the same CDB simultaneously
- Compressed backup sets
- Better parallelism

Requirements

- Source and target database must
 - be 19.18.0 or higher
 - use Data Pump Bundle Patch



Always use the latest version of M5 script


- Download from [KB144840](#)



Use Block Change Tracking for faster incremental backups

- Check the License Guide for details

Migration Plan



M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster



Local standby




ZDLRA




Exadata
X9M Extreme Flash

M5 Workflow



- Download M5 script from [KB144840](#)
- Configure shared NFS
- Edit `dbmig_ts_list.txt`
- Edit `dbmig_driver.properties`
- Create new, empty target database

M5 Workflow



SPARC
SuperCluster




ZDLRA




Exadata
X9M Extreme Flash

M5 Workflow



- Start initial level 0 backup
 - Use driver script `dbmig_driver_m5.sh` `L0`
- Driver script creates a restore script
 - Restore using `restore_L0_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow



SPARC
SuperCluster




ZDLRA





Exadata
X9M Extreme Flash

M5 Workflow



- Start level 1 incremental backup
 - Use driver script `dbmig_driver_m5.sh` [L1](#)
- Driver script creates a restore script
 - Restore using `restore_L1_<source_sid>_<timestamp>.cmd`
- Check logs
- [Repeat as often as desired](#)

M5 Workflow




SPARC
SuperCluster




Exadata
X9M Extreme Flash

M5 Workflow




- Maintenance window begins
- Read-only sessions can still use the database

M5 Workflow




M5 Workflow



- Start final level 1 incremental backup
 - Use driver script `dbmig_driver_m5.sh` [L1F](#)
 - Sets tablespaces read-only
 - Performs level 1 incremental backup
 - Starts Data Pump full transportable export
- Optionally, shuts down source database

M5 Workflow



- Driver script created a restore script
 - Restore using `restore_L1F_<source_sid>_<timestamp>.cmd`
- Check logs

M5 Workflow

Configure

Level 0

Level 1

Outage

Final
Backup

Final
restore

Export
Import



SPARC
SuperCluster




ZDLRA




Exadata
X9M Extreme Flash

M5 Workflow



- Copy Data Pump dump file to *DATA_PUMP_DIR*
- Use import driver script in test mode
 - Start `impdp.sh <dump_file> <restore_log> test`
- Check generated parameter file
 - Use `impdp.sh <dump_file> <restore_log> run`
- Check Data Pump log file

Wanna try it out?



Oracle LiveLabs – Run the lab just inside your browser!




The Project Plan



Timelines and the Run Book



Overall Project Timeline



Key to Success: Runbook

Complex projects absolutely require a detailed runbook

ID	Task	Status	Responsible Primary Person	Responsible Secondary Person	Predecessor	Start Time (CEST)	Duration (hh:mm)	End Time (CEST)	Start Time (IST)	End Time (IST)	Actual Start Time (CEST)	Actual Duration	Actual End Time (CEST)	Comments - Blocker
----	------	--------	----------------------------------	------------------------------------	-------------	----------------------	---------------------	--------------------	---------------------	-------------------	--------------------------------	--------------------	------------------------------	-----------------------

- This run book covered over 200 individual tasks

Timeline Live Migration

Task	22:00	23:00	0:00	1:00	2:00	3:00	4:00	5:00	6:00	7:00	8:00




How many people were involved?



Migration Issues

Some of them...



Where it started ...



PoC first FTEX **export**:

01-OCT-21 05:32:36.275: Job "SYSTEM"."SYS_EXPORT_FULL_01" successfully completed at
Fri Oct 1 05:32:36 2021 elapsed 0 04:25:22

PoC first FTEX **import**:

05-OCT-21 01:48:59.534: Job "SYSTEM"."SYS_IMPORT_FULL_01" completed with 103000
error(s) at Tue Oct 5 01:48:59 2021 elapsed 3 18:34:09

The way forward ...

80 SRs opened and solved in various areas

Migration
ZDLRA
Standby
Infrastructure

18 one-off patches

5 merges

Daily calls with Oracle, countless evening / night / weekend hours

Many areas required special attention....

Optimizer Statistics

Scheduler jobs

Resource Manager

Cross Schema objects

AQ

Evolved Types/partitioned nested tables

Binary XML

Standby DBs

...



Photo by Justin Chm on Unsplash

Issue 1 | Long Running Metadata Import

Fix applied to remedy export errors

- BUG 34201281 - MERGE ON DATABASE RU 19.12.0.0.0 OF 33963454 34052641

Result:

- Now Full Transportable import alone took over 6 days (!!)

```
08-JUN-22 16:21:17.887: W-1 Processing object type DATABASE_EXPORT/.../PROCACT_INSTANCE
14-JUN-22 18:56:58.813: W-1      Completed 108 PROCACT_INSTANCE objects in 527737 seconds
...
14-JUN-22 19:15:50.016: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed with 316 error(s)
                        at Tue Jun 14 19:15:49 2022 elapsed 6 06:45:38
```

Issue 1 | Long Running Meta Import

Long running action identified via tracing:

```
UPDATE "POSTMAN"."T_MAIL_LOG"  
      SET "C_CVAR"=SYS_REMAP_XMLTYPE("C_CVAR")
```

- 300+ million rows

Issue in internal package DBMS_CSX_INT

- Fast merge of XMLTYPE is not happening as expected
 - *Reason:* Incorrect internal check query
- Tokens between source and target are not identical
 - *Reason:* Different Endianness

Issue 1 | Long Running Meta Import

Solution:

- Use workaround from [KB152248](#) in UPGRADE mode

```
25-JUL-22 12:28:40.813: Job "SYSTEM"."SYS_IMPORT_TRANSPORTABLE_01" completed  
with 317 error(s) at Mon Jul 25 12:28:40 2022 elapsed 0 04:33:03
```

Issue 2 | Metadata API and Nested Tables

Full transportable import errors out for a **nested partitioned table**

```
PLS-00172: string literal too long
```

```
ORA-39151: Table "DBA_XY"."X_GAMES" exists.
```

```
All dependent metadata and data will be skipped due to table_exists_action
```


Root cause was a string overflow in the Metadata API

- Data Pump creates the index, and then alters it – here the overflow happened
- Side effect was a misleading error message

Issue 3 | Data Pump Bundle Patch

Many of the TTS and Metadata fixes got included into Data Pump Bundle

- [\(KB107134\) Data Pump Recommended Proactive Patches For 19.10 and Above](#)



Issue 4 | Evolved Object Types

Evolved TYPES can lead to Data Pump errors during transportable import:

```
ORA-39083: Object type TABLE:"APPUSER"."CARS" failed to create with error:  
ORA-39218: type check on object type "APPUSER"."CAR_TYPE" failed  
ORA-39216: object type "APPUSER"."CAR_TYPE" hashcode or version number mismatch
```

Further Information:

- [Blog Post: Understand why Data Pump errors with evolved types](#)



Using evolved types in table definitions

--Create a new type. The type is now version 1

--Use the type in a table

```
CREATE TYPE CAR_INFO_TYPE IS OBJECT (model VARCHAR2(40));  
CREATE TABLE CARS (id number, car_info car_info_type);  
INSERT INTO CARS VALUES (1, car_info_type('Volvo V90'));
```

The type is now evolving

--Make a change to the type. The type is now version 2

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE horsepower NUMBER CASCADE NOT INCLUDING TABLE DATA;  
INSERT INTO CARS VALUES (2, car_info_type('BMW 530i', 250));
```

Existing data is not updated

--Make another change to the type. The type is now version 3

```
ALTER TYPE CAR_INFO_TYPE ADD ATTRIBUTE color VARCHAR2(20) CASCADE NOT INCLUDING TABLE DATA;  
INSERT INTO CARS VALUES (3, car_info_type('Hyundai Sonata', 160, 'Black'));
```

Evolved Types

```
SELECT * FROM CARS;
```



CARS	
1	car_info_type v1: Volvo V90
2	car_info_type v2: BMW 530i, 250
3	car_info_type v3: Hyundai Sonata, 160, Black



DICTIONARY	
car_info_type v1	model
car_info_type v2	model, horsepower
car_info_type v3	model, horsepower, color



Data Pump recreates types during Full Transportable Export/Import

Evolved Types

- To avoid data corruption,
Data Pump must recreate the exact same type evolution in target database
- Due to implementation restrictions,
it is not always possible to recreate the exact same type evolution
- In such situations, to avoid corruption,
Data Pump reports ORA-39218 or ORA-39216 on import

Evolved Types | Possible Solutions

- 1 Conventional Data Pump export
- 2 Manually recreate `type` in target database with matching evolution
- 3 Recreate `type` without evolution before export

[Blog post](#) with details

Issue 5 | Advanced Queueing

Source database

Queue table

`<queue_table_name>`

`AQ$_<queue_table_name>_E`
`AQ$_<queue_table_name>_I`
`AQ$_<queue_table_name>_T`
`AQ$_<queue_table_name>_F`
`AQ$_<queue_table_name>_C`
`AQ$_<queue_table_name>_D`
`AQ$_<queue_table_name>_G`
`AQ$_<queue_table_name>_H`
`AQ$_<queue_table_name>_L`
`AQ$_<queue_table_name>_P`
`AQ$_<queue_table_name>_S`
`AQ$_<queue_table_name>_V`



Queue
infrastructure

Target database

`<queue_table_name>`

`AQ$_<queue_table_name>_E`
`AQ$_<queue_table_name>_I`
`AQ$_<queue_table_name>_T`
`AQ$_<queue_table_name>_F`

Issue 5 | Advanced Queueing

Queue tables and underlying objects may change during import

- COMPATIBLE **during creation** of queue tables matters
- COMPATIBLE **during import** matters as well
- [\(KB115340\) Understanding How AQ Objects Are Exported And Imported](#)
- [Blog post: Changing data types in queue tables during import](#)

Options:

- Recreate the queue tables with "old" COMPATIBLE setting
- Benefit from new COMPATIBLE setting and test the application



Take into account when comparing source and target databases' object count

- [\(KB115340\) Understanding How Advanced Queueing \(AQ\) Objects Are Exported And Imported](#)



Data Pump does not start queues

- Manually start queues after migration
- Use `DBMS_AQADM.START_QUEUE`

Issue 6 | Default Tablespaces

Due to a security fix export wants to write into the user's default tablespace

- [Bug 27692190](#)
- But default, tablespaces are read-only while full transportable export runs

Workaround:

- Change default tablespace for all users to SYSTEM
- Full Transportable Export
- Full Transportable Import
- Revert default tablespaces back to original in target (and source)

Issue 7 | Exporting Statistics

Exporting statistics is slow using
`DBMS_STATS.EXPORT_SCHEMA_STATS`




10046 trace reveals long runtime on:
`EXP_STAT$`
`EXP_OBJ$`


Expression Statistics for Auto-Indexing

Default: `_column_tracking_level=53`

Issue 7 | Exporting Statistics



Issue 7 | Exporting Statistics



Issue 8 | Auditing

Tablespaces containing auditing tables **can't be set read-only**

Data Pump always unloads the audit records into the dump file

- Huge audit trail will lead to a huge dump file and longer outage

Options:

- Export audit records, and eventually import them afterwards
- Archive audit records, purge the audit trail

Key Learnings



- 1 Practice, practice, practice
- 2 Create a detailed runbook
- 3 Work together

AI changes everything



Harvard Business Review

***“AI won’t replace humans -
but humans with AI will
replace humans without AI”***

Karim Lakhani





AI Database



Using AI

Filter



Year



Genre



Nominations

Awards

Question

Message

Motivational movies about athletics

35

Search Type

☒ Text Search

☐ Vector Search

Search

Results

1 - 10 of 10

	Title	Awards	Year	Nominations	Genre
	Goodbye Columbus	NONE	1969	Academy Award for Best Writing, Adapted Screenplay	Romance, Comedy, Family
	The Formula	NONE	2002	NONE	Unknown
	Batman Returns	NONE	1992	Academy Award for Best Makeup and Hairstyling, Academy Award for Best Visual Effects, MTV Movie Award for Best Villain	Film-Noir, Action, Family, Fantasy

Why *Batman Returns*?

“Motivational Movies about Athletics”

NJQ15253, *Batman Returns*, 1992, https://upload.wikimedia.org/wikipedia/en/8/83/Batman_returns_poster2.jpg, 'Film-Noir, Action, Family, Fantasy', 'Academy Award for Best Makeup and Hairstyling, Academy Award for Best Visual Effects, MTV Movie Award for Best Villain', , *Batman Returns* is a 1992 American superhero film directed by Tim Burton and produced by Denise Di Novi and Burton, based on the DC Comics character Batman. The sequel to the 1989 film *Batman* , it is the second installment of Warner Bros. initial *Batman* film series, and stars Michael Keaton as Bruce Wayne / Batman, alongside Danny DeVito, Michelle Pfeiffer, Christopher Walken, Michael Gough, Pat Hingle and Michael Murphy. In *Batman Returns* , Batman faces the Penguin, who plots to kill all of Gotham City's firstborn sons, while dealing with Catwoman, who seeks vengeance against Max Shreck, a corrupt tycoon who allies with the Penguin to bring Gotham City under his control. Burton originally did not want to direct another *Batman* film. Warner Bros. developed a script with Sam Hamm which had the Penguin and Catwoman going after hidden treasure. Burton agreed to return after they granted him more creative control and replaced Hamm with Daniel Waters. Wesley Strick was later chosen to do an uncredited rewrite shortly before filming. This included normalizing dialogue, fleshing out the Penguin's

motivations and master plan, and removing scenes due to budget concerns. Strick continued working as the on-set writer through filming. Annette Bening was originally cast as Catwoman, but became pregnant and was replaced with Pfeiffer. *Batman Returns* was released on June 19, 1992. It grossed \$266.8 million worldwide on a total budget of \$80 million and received positive reviews. Critics praised its action sequences, performances, Danny Elfman's score, effects and villains, although its dark tone and high level of violence for a PG-13 film, was criticized. The film was nominated for two Academy Awards: Best Visual Effects and Best Makeup, as well as two BAFTA awards. A stand- alone sequel, *Batman Forever* , was released in 1995, with Val Kilmer replacing Keaton as Batman. An alternate comic book continuation, which ignores the events of the subsequent films, will be published by DC Comics starting in July 2021. Keaton is also set to reprise the role of Batman in the DC Extended Universe beginning with *The Flash* (2022).

Filter



Year



Genre



Nominations

Awards

Question

Message

Motivational movies about athletics

35

Search Type



Text Search




Vector Search

Search

Results

1 - 10 of 10

	Title	Awards	Year	Nominations	Genre
	The Do-Deca-Pentathlon	NONE	2012	NONE	Comedy
	Champions	NONE	2023	NONE	Comedy,Sport
	Across the Tracks	NONE	1991	NONE	Drama,Action
	Hustle	NONE	2022	NONE	Sport



Filter



Year

2023



Genre

Sport



Nominations

Awards

Question

Message

Motivational movies about athletics

35

Search Type



Text Search



Vector Search


Search

Results


1 - 3 of 3						
		Title	Awards	Year	Nominations	Genre
		Champions	NONE	2023	NONE	Comedy,Sport
		Sweetwater	NONE	2023	NONE	Biography,Sport
		80 for Brady	NONE	2023	NONE	Sport,Comedy

Key new AI data type: **AI Vector**

Stores semantic content
of a complex object



Oracle AI Database
can **store** AI vectors and
use them to **find**
similar objects



AI needs **business data** to answer your questions

Finding relevant business data
requires searching **private data**
with **traditional data search**
and **AI Vector Search**




AI Database **unifies** search across traditional data and AI vectors



Does my dental plan cover
braces for my 19-year-old?

Find semantically similar
benefit documents
that are relevant to the
employee



Bring AI search to
both operational and
analytical systems

with just **one extra**
line of SQL



Does my dental plan cover
braces for my 19-year-old?

```
SELECT document
FROM   benefits_documents
WHERE  doc_plan_id IN (SELECT plan_id from emp_plans...)
ORDER BY VECTOR_DISTANCE(:question_vec,
                           document_vec);
```

Unify AI Vector Search
with **all** types of data,
not just relational



Vector



XML

{JSON}



Spatial



Graph



Relational



Text

Easily pass results of search to LLM to answer the question



Does my dental plan cover braces for my 19-year-old?


AI Vector Search



Yes, your “Dental Plus” plan covers braces for your children until they are 21 years old.


**Retrieval
Augmented
Generation**

Vector Database



- New data type: **VECTOR**
- A vector is
*a numerical representation of data
that captures key features and relationships of the data*

Vector Database



```
CREATE TABLE movies (  
    id          NUMBER,  
    description CLOB,  
    photo       BLOB,  
    my_vector   VECTOR(768, FLOAT32));
```

It's just an array ...



↑
The dimension is
"how many numbers"

Vector Database




-
- Your data is already in the database
 - How do you perform a semantic search instead of a literal search?

Comparing numbers is trivial

`100 > 50 = true`

How to compare complex data types?



How to search for search complex data types?



Vector Database



-
- Search vector data using vector indexes
 - Calculate the vector distance



Type of movie



Genre







$$\begin{aligned} &\text{Distance (Euclidean Squared)} \\ &= (3-2)^2 + (1-6)^2 + (2-2)^2 + (8-3)^2 \end{aligned}$$

Vector Database




```
SELECT    ...  
FROM      movies  
ORDER BY  vector_distance(movie1, movie2, EUCLIDEAN_SQUARED);
```






How do you convert your data
into vectors?

Vector Database



- Embedding models transform your data into a vector



10,000 GPU ?

```
DBMS_DATA_MINING.import_onnx_model(  
    model_name => 'All-MiniLM-L6-v2',  
    model_data => 'All-MiniLM-L6-v2.onnx'  
    ...  
);
```




```
INSERT INTO movies
VALUES (:id,
        :description,
        :photo,
        VECTOR_EMBEDDING(All-MiniLM-L6-v2, :description));
```



Vector Database



Vector Database



Vector Database



Vector Database






Included in any edition
of Oracle AI Database 26ai

- No extra license required

Oracle AI Database 26ai Annotations

Annotations add **intended usage** to schema elements such as tables, views, columns
User-friendly display names, identify sensitive data, etc.



Annotations can be used by low code frameworks to **generate** user interfaces




AI Database



Using AI

Architecting **Agentic AI** into the database


- Select AI Agent
- Private Agent Factory
- SQLcl MCP Server



MCP Server

Model Context Protocol

Enable powerful and
simplistic AI



EXTENSIONS

Search Extensions in Marketplace

INSTALLED



Cline

Autonomous coding agent right in your IDE, capable of crea...

Cline

504ms



GitHub Copilot Chat

AI chat features powered by Copilot

GitHub

404ms



Oracle Code Assist

Oracle Code Assist

Oracle



Oracle SQL Developer Extension for VSCode

End-to-end development tools for your SQL and PL/SQL ap...

Oracle Corporation

91ms



GitHub Copilot

Your AI pair programmer

GitHub

RECOMMENDED



SQLTools

Connecting users to many of the most commonly used data...

Matheus Teixeira

6M

3.5

Install



Vim

Vim emulation for Visual Studio Code

vscodevim

8.3M

4

Install

MCP SERVERS



MCP Servers

Browse and install [Model Context Protocol \(MCP\)](#)

Untitled-1 X

1 `Generate code (M), or select a language (M). Start typing to dismiss or don't show this again.`

CHAT



Build with Agent

AI responses may be inaccurate.

[Generate Agent Instructions](#) to onboard AI onto your codebase.


 +  Untitled-1

Describe what to build next

Agent  GPT-5 mini 

Ln 1, Col 1 Spaces: 4 UTF-8 LF () PL/SQL (Dedicated) ADB26ai

YouTube Channel



<https://www.youtube.com/@upgradenow>

- 600+ videos
- New videos every week
- No marketing
- No buzzword
- All tech



SCAN ME

Thank You
