



Frequently Asked Questions for Oracle Autonomous Database with Oracle JD Edwards EnterpriseOne



List of Frequently Asked Questions and Answers for Oracle
Autonomous Database with Oracle JD Edwards EnterpriseOne

October, 2021 | Version 1.01
Copyright © 2021, Oracle and/or its affiliates
Confidential - Public

PURPOSE STATEMENT

This document provides answers to frequently asked questions for Oracle Autonomous Database with Oracle JD Edwards EnterpriseOne. It is intended solely to help you assess the business benefits of upgrading to Oracle JD Edwards Applications 9.2, Tools 9.2.6 and Oracle Autonomous Database and to plan your I.T. projects.

DISCLAIMER

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

TABLE OF CONTENTS

Purpose Statement	1
Disclaimer	1
Introduction	3
General Queries on Oracle Autonomous Database	3
What is the Oracle Autonomous Database?	3
What are the benefits of using the Oracle Autonomous Database with JD Edwards?	3
What is Oracle Autonomous Database on Dedicated Exadata Infrastructure?	4
What is Oracle Autonomous Database on Shared Exadata Infrastructure?	4
What are the components of the Oracle Autonomous Database on Dedicated Exadata Infrastructure?	4
I already have an existing license of Oracle Database Enterprise Edition, is it possible to use it to move to Oracle Autonomous Database?	4
Which are the Oracle Cloud Infrastructure regions where Oracle Autonomous Database is available?	5
Can I leverage Autonomous Database capabilities such as automatic indexing and automatic scaling?	5
Configuring JD Edwards with Oracle Autonomous Database	5
What is the JD Edwards Applications and Tools release that is supported for Autonomous Database?	5
Can my existing JD Edwards installation work with Oracle Autonomous Database?	5
I am an existing on-premise customer of Oracle JD Edwards, can I leverage Oracle Autonomous Database?	5
I am currently using Oracle Database Enterprise Edition with Oracle JD Edwards, what are the options of loading data into Oracle Autonomous Database?	6
Is Oracle JD Edwards EnterpriseOne 32-bit supported with Oracle Autonomous Database?	6
Which source database versions can be migrated to Oracle Autonomous Database for JD Edwards?	6
I would like to continue using my non-autonomous database instances for my non-production environments while using Oracle Autonomous Database for my production environments. Can I have this kind of setup?	6
Can I use my existing Oracle JD Edwards Server Manager to manage Oracle Autonomous Database?	7
Maintenance Activities with Oracle Autonomous Database on dedicated exadata infrastructure	7
What type of maintenance activities are currently supported for JD Edwards with Oracle Autonomous Database on Dedicated Exadata Infrastructure?	7
What is application continuity service drain time?	7
Can Oracle JD Edwards customers modify drain time on Oracle Autonomous Database on Dedicated Exadata Infrastructure?	7
After I configure my Oracle JD Edwards instance with Oracle Autonomous Database on Dedicated Exadata Infrastructure, can I expect that my unplanned maintenance and outages will be taken care of by the self-repairing and failover capabilities of the Autonomous Database?	7
How does the automatic patching of Oracle Autonomous Database take place?	8
Can I defer the maintenance patching so that it does not impact critical periods for my business?	8
I do not want to directly apply the patches on my production instance of Oracle Autonomous Database on Dedicated Exadata Infrastructure. How do I ensure that the patches are validated before applying them on the production instance?	8
Maintenance Activities with Oracle Autonomous Database on shared exadata infrastructure	8
What type of maintenance activities are currently supported for JD Edwards with Oracle Autonomous Database on Shared Exadata Infrastructure?	8
What is application continuity service drain time?	8
Can Oracle JD Edwards customers modify drain time on Oracle Autonomous Database on Shared Exadata Infrastructure?	9
After I configure my Oracle JD Edwards instance with Oracle Autonomous Database, can I expect that my unplanned maintenance and outages will be taken care of by the self-repairing and failover capabilities of the Autonomous Database?	9
How does the automatic patching of Oracle Autonomous Database take place?	9
Can I defer the maintenance patching so that it does not impact critical periods for my business?	9
I do not want to directly apply the patches on my production instance of Oracle Autonomous Database on Shared Exadata Infrastructure. How do I ensure that the patches are validated before they are applied automatically on the production instance?	9



INTRODUCTION

This document intends to provide answers to frequently asked questions on implementing Oracle Autonomous Database with Oracle JD Edwards Enterprise One. The document covers general queries related to Oracle Autonomous Database, queries specific to configuring Oracle JD Edwards with Oracle Autonomous Database, and queries about maintenance activities in Oracle Autonomous Database.

GENERAL QUERIES ON ORACLE AUTONOMOUS DATABASE

What is the Oracle Autonomous Database?

Oracle Autonomous Database (ADB) is the world's first autonomous data management system in the cloud to deliver automated patching, upgrades, and tuning—including all routine database maintenance tasks while the system is running—without human intervention. This new autonomous database cloud is self-driving, self-securing, and self-repairing, and thereby eliminates manual database management and human errors. Oracle Autonomous Database is fully elastic: you simply specify the number of OCPUs and the storage capacity for the database. At any time, you may scale—increase or decrease either the OCPUs or the storage capacity. For more details on the capabilities of Oracle Autonomous Database, refer to [Oracle Autonomous Database Features](#) and [Oracle Autonomous Database Strategy](#).

What are the benefits of using the Oracle Autonomous Database with JD Edwards?

Oracle Autonomous Database provides a level of performance and reliability that manually managed databases cannot deliver. Compared to a manually managed database, the Autonomous Database costs less to run and maintain, provides higher availability, reduces human error, and is more secure—all contributing to increased performance. Oracle JD Edwards customers can leverage the capabilities of the Oracle Autonomous Database and reap significant benefits from it.

- **Eliminate Human Labor:** The Oracle Autonomous Database eliminates human labor for database administration activities like provisioning, securing, monitoring, tuning and troubleshooting databases. This results in reduced costs and saved human labor, which JD Edwards customers can leverage for accelerating innovation in their core business areas.
- **Increased system availability:** The Autonomous Database automatically recovers itself from any physical failure either at the data center or server level. It updates the database, operating system, virtual machine, clusterware, hypervisor, or firmware in a rolling fashion without any application downtime. With zero application downtime, Autonomous Database provides JD Edwards customers with increased system availability and business continuity, eliminating any impact on business operations.
- **Reduced Risk:** The Autonomous Database protects itself from vulnerabilities by applying database, operating system and security patches every quarter. The patches are applied on a rolling fashion across the nodes of the database cluster ensuring zero application downtime. This ensures that JD Edwards customers are secure from any kind of vulnerabilities.
- **Leveraging future Oracle innovation:** The Autonomous Database with its advanced capabilities is a perfect stepping-stone for JD Edwards customers to leverage and benefit from future Oracle innovation. This would

catalyze the long-term process of redefining and reimagining IT management at an organizational level. By adopting the Oracle Autonomous Database, JD Edwards customers set a solid foundation for leveraging any future value-driven innovation from Oracle that provides them significant business benefits.

For more details on the business benefits, Refer: [Oracle JD Edwards EnterpriseOne with Oracle Autonomous Database](#)

What is Oracle Autonomous Database on Dedicated Exadata Infrastructure?

Oracle JD Edwards supports the Oracle Autonomous Database on Dedicated Exadata Infrastructure starting from tools release 9.2.4.3. This support enables you to provision Autonomous Databases into your own dedicated Exadata cloud infrastructure, instead of into a shared infrastructure with other tenants. With this service, you can exclusively use the Exadata hardware. The Oracle Autonomous Database Dedicated service uses the Autonomous Exadata Infrastructure that offers multitenant database architecture, enabling you to create and manage multiple Autonomous Databases within a single database system. Oracle Autonomous Database dedicated is also available on Exadata Cloud@Customer for customers who want to deploy Autonomous Database in their data center. As of September 2021, Autonomous Exadata Infrastructure (X8 or X7) is available in three options - quarter rack, half rack and full rack. For the latest available configurations, refer to [ADB Hardware Shapes](#).

What is Oracle Autonomous Database on Shared Exadata Infrastructure?

Oracle JD Edwards supports the Oracle Autonomous Database on Shared Exadata Infrastructure starting from tools release 9.2.6. This enables you to provision Autonomous Databases on a shared infrastructure with other tenants. With this option, you provision and manage only the Autonomous Database, while Oracle deploys and manages the Exadata infrastructure. Customers get up to a maximum of 128 CPU cores with Oracle Autonomous Database. To create a Shared Autonomous Database, you create an Oracle Autonomous Database, specifying "Transaction Processing" as its workload type and "Shared Infrastructure" as its deployment type.

What are the components of the Oracle Autonomous Database on Dedicated Exadata Infrastructure?

The Oracle Autonomous Database on Dedicated Exadata Infrastructure has the following components:

- **Autonomous Exadata Infrastructure**—a resource that allocates an available Oracle Exadata database machine to you. Its primary purpose is to act as a bridge between the hardware and software components of your dedicated infrastructure. You must create at least one Autonomous Exadata Infrastructure resource before you can create any of the other kinds of dedicated infrastructure resources.
- **Autonomous Container Database**—a resource that provides a container for your autonomous databases. You can create multiple Autonomous Container Database resources in a single Autonomous Exadata Infrastructure resource, but you must create at least one before you can create any autonomous databases. Multiple autonomous databases can be created within a single container database.
- **Autonomous Database**—a resource that is a user database. This resource can be considered as a pluggable database. To create a dedicated Autonomous Database, you create an autonomous database specifying "Transaction Processing" as its workload type and "Dedicated Infrastructure" as its deployment type. You can create multiple autonomous databases in a single Autonomous Container Database resource.
- **Compartment**—a collection of related resources. Compartments are fundamental components of Oracle Cloud Infrastructure for organizing and isolating your cloud resources.
- **Autonomous Backup**—a resource that is a backup of an Autonomous Database. Oracle creates this resource for you automatically.

I already have an existing license of Oracle Database Enterprise Edition, is it possible to use it to move to Oracle Autonomous Database?

You can provision a license-included version of ADB, or you can provision a Bring Your Own License (BYOL) instance. If you wish to use the BYOL model, then you must apply the current database licenses to your ADB service. If you run Oracle Database Enterprise Edition with either the Multitenant option or the Multitenant with Real Application Clusters option, then you can activate Oracle Autonomous Database using the options listed at [Oracle Autonomous Database pricing page](#). Contact your Oracle technology representative for additional information and your specific licensing.

Which are the Oracle Cloud Infrastructure regions where Oracle Autonomous Database is available?

ADB is available in worldwide Oracle Cloud Infrastructure data regions. Examples include Ashburn, Phoenix, Toronto, Montreal, London, Frankfurt, Zurich, Amsterdam, Jeddah, Sao Paulo, Tokyo, Osaka, Seoul, Sydney, and Mumbai. For the latest list of the Oracle Cloud Infrastructure data regions, the number of which is rapidly increasing, refer to [Data Regions for Platform and Infrastructure Services](#).

Can I leverage Autonomous Database capabilities such as automatic indexing and automatic scaling?

No. Currently JD Edwards does not support automatic indexing and automatic scaling with Oracle Autonomous Database.

CONFIGURING JD EDWARDS WITH ORACLE AUTONOMOUS DATABASE

What is the JD Edwards Applications and Tools release that is supported for Autonomous Database?

Oracle JD Edwards supports Oracle Autonomous Database on Dedicated Exadata Infrastructure starting from tools release 9.2.4.3 and Oracle Autonomous Database on Shared Exadata Infrastructure starting from tools release 9.2.6. Oracle recommends customers should always adopt the latest tools release and the latest certified versions of Oracle Database, Oracle Linux and Oracle Weblogic server to leverage Oracle Autonomous Database. For information about latest certified platforms, please refer [My Oracle Support Certifications](#).

Can my existing JD Edwards installation work with Oracle Autonomous Database?

Yes, if you are an existing Oracle JD Edwards customer on Oracle Cloud Infrastructure, you can provision an Oracle Autonomous Database instance either on Shared or Dedicated Exadata Infrastructure. If you are considering Oracle Autonomous Database on Shared Exadata Infrastructure, then you have to migrate data manually from your existing database to the Autonomous Database, and perform manual configurations to integrate the autonomous instance with the Oracle JD Edwards installation. If you are considering Oracle Autonomous Database on Dedicated Exadata Infrastructure, then you can leverage the One-Click provisioning automation to integrate the autonomous instance with the Oracle JD Edwards installation and then migrate data manually from your existing database to the Autonomous Database.

If your Oracle JD Edwards instances currently do not run on Oracle Cloud Infrastructure, you will need to lift and shift your instances. If you want to migrate to Oracle Autonomous Database on Dedicated Exadata infrastructure, you can use Oracle JD Edwards One-Click Provisioning to provision the Oracle JD Edwards environment with an Oracle Autonomous Database instance with automated configurations to integrate the autonomous instance with the Oracle JD Edwards installation. If you are an on-premise customer but do not want to move to public cloud, you can leverage the benefits of a subscription-based pay-per-use Oracle Autonomous Database by migrating to Oracle Autonomous Database on Exadata Cloud@Customer. For details, refer to: [Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with Autonomous Database](#)

If you want to migrate to Oracle Autonomous Database on Shared Exadata infrastructure, you can use Oracle JD Edwards One-Click Provisioning to provision the Oracle JD Edwards environment with a Compute Database or a database service and then provision an Oracle Autonomous Database instance, migrate data manually from your existing database to the Autonomous Database, and perform manual configurations to integrate the autonomous instance with the Oracle JD Edwards installation. For details, refer to: [Manual deployment of Oracle JD Edwards EnterpriseOne with Oracle Autonomous Database on Shared Exadata Infrastructure](#)

I am an existing on-premise customer of Oracle JD Edwards, can I leverage Oracle Autonomous Database?

Yes. However, Oracle Autonomous Database is a cloud-only service and is available on Oracle Cloud Infrastructure (Shared and Dedicated) and Exadata Cloud@Customer (Dedicated only).

If you are an on-premise JD Edwards customer and considering a move to Oracle Autonomous Database on Shared Exadata infrastructure, then you can use Oracle JD Edwards One-Click Provisioning to provision the Oracle JD Edwards environment with a Compute Database or a database service using the steps described in [Upgrading JD Edwards EnterpriseOne on a One-Click Provisioned Target Environment on Linux](#). You then need to manually provision an Oracle Autonomous Database instance, migrate data manually from your existing database to the Autonomous Database, and perform manual configurations to integrate the autonomous instance with the JD Edwards installation.

If you are an on-premise JD Edwards customer and considering a move to Oracle Autonomous Database on Dedicated Exadata infrastructure, then you can use Oracle JD Edwards One-Click Provisioning to provision the Oracle JD Edwards environment with an Oracle Autonomous Database instance with automated configurations to integrate the autonomous instance with the Oracle JD Edwards installation using the steps described in [Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with Autonomous Database](#) and migrate data manually from your existing database to the Autonomous Database using the steps described in [Manual Deployment of Oracle JD Edwards EnterpriseOne with Oracle Autonomous Database](#)

If you are currently on-premise with a supported Oracle Database version but not looking to move to public cloud, you can leverage the benefits of a subscription-based pay-per-use Oracle Autonomous Database by migrating to Oracle Autonomous Database Dedicated on Exadata Cloud@Customer.

I am currently using Oracle Database Enterprise Edition with Oracle JD Edwards, what are the options of loading data into Oracle Autonomous Database?

You can migrate your data into Oracle Autonomous Database in multiple ways:

- You can use Data Pump Import to load data from dump files. Data Pump Import is integrated with Oracle Cloud Infrastructure Object Storage and Oracle Cloud Infrastructure Object Storage Classic. You can upload your dump files to one of these object storage locations and import data from those dump files.
- You can also use client-side tools such as SQL*Loader, or your in-house scripts to load data from files residing on your client machines.
- Another option is to use Oracle tools and cloud services like GoldenGate, Oracle Data Integrator (ODI), Goldengate Cloud Service, and Data Integration Platform Cloud.
- Any third-party ETL tool that can connect to ADB using the security credentials wallet can also be used to load data.

For more details on these methods, refer to [Loading Data into Autonomous Database](#).

Is Oracle JD Edwards EnterpriseOne 32-bit supported with Oracle Autonomous Database?

No. Only Oracle JD Edwards 64-bit is supported with Oracle Autonomous Database.

Which source database versions can be migrated to Oracle Autonomous Database for JD Edwards?

Currently, all Oracle database versions that are supported by Oracle JD Edwards Tools Release 9.2.6 are supported for migration to Oracle Autonomous Database. If your existing Oracle Database version is not 19c, then you need to upgrade your Oracle Database Client on the Enterprise Server to 19c after migrating to Oracle Autonomous Database. You should use the Autonomous Database regional wallet in your Enterprise Server and the Database Server to connect to your Autonomous Database instance.

I would like to continue using my non-autonomous database instances for my non-production environments while using Oracle Autonomous Database for my production environments. Can I have this kind of setup?

While this kind of setup is possible, Oracle JD Edwards recommends that you move all your non-production environments as well to Oracle Autonomous Database.

Can I use my existing Oracle JD Edwards Server Manager to manage Oracle Autonomous Database?

No. Server Manager does not support any management activities for Oracle Autonomous Database. All operations must be done through the Oracle Cloud Infrastructure console only.

MAINTENANCE ACTIVITIES WITH ORACLE AUTONOMOUS DATABASE ON DEDICATED EXADATA INFRASTRUCTURE

What type of maintenance activities are currently supported for JD Edwards with Oracle Autonomous Database on Dedicated Exadata Infrastructure?

For Oracle Autonomous Database on Dedicated Exadata Infrastructure, JD Edwards currently supports planned maintenance activities such as database patching and operating system updates using extended drain time (these activities are explained in detail in the following sections of this FAQ). Unplanned maintenance is currently not supported. Customers will be able to provide a predefined drain time in the database to drain active connections to other nodes during planned maintenance activities. For more details, refer to: [Planned Database Maintenance of Oracle Autonomous Database with JD Edwards EnterpriseOne](#)

What is application continuity service drain time?

Application continuity service drain time is the time taken for active connections on one database node to drain to another active database node. During drain time, new database connections are not permitted to the node that is draining the connections and all new connection requests are redirected to the other active nodes that are not draining connections. When all the active connections are drained, the node will be shut down for planned maintenance.

Application continuity service drain time ensures that there are no in-flight transactions or operations on a particular database node during maintenance and hence prevents transactions from being abruptly aborted resulting in data integrity issues. It also prevents data corruption by avoiding partial completion of transactions. It eliminates the additional time and effort required to resubmit the transactions and revert changes made by the aborted transaction. Therefore, you can carry out planned maintenance activities on Oracle Autonomous Database without disrupting the ongoing transactions or operations.

Can Oracle JD Edwards customers modify drain time on Oracle Autonomous Database on Dedicated Exadata Infrastructure?

Yes. Oracle JD Edwards customers can configure the drain time to last up to 8 hours enabling connections related to long-running transactions to complete and ensure data integrity before the node is brought down for maintenance. Drain time enables Oracle JD Edwards customers to carry out planned maintenance activities on their Oracle Autonomous Database on Dedicated Exadata Infrastructure without disrupting their ongoing transactions or operations. Oracle Autonomous Database is by default deployed on a Real Application Cluster(RAC) and hence Oracle JD Edwards recommends that you configure the application continuity service drain time irrespective of the number of OCPUs of your Oracle Autonomous Database. For information on how to determine the ideal drain time and configure it in Oracle Autonomous Database, refer to [Planned Database Maintenance of Oracle Autonomous Database with JD Edwards EnterpriseOne](#)

After I configure my Oracle JD Edwards instance with Oracle Autonomous Database on Dedicated Exadata Infrastructure, can I expect that my unplanned maintenance and outages will be taken care of by the self-repairing and failover capabilities of the Autonomous Database?

No. Currently Oracle JD Edwards is not supported for handling unplanned outages with Oracle Autonomous Database on Dedicated Exadata Infrastructure and hence self-repairing and failover capabilities cannot be leveraged. In some cases, the failover could happen but you might have to restart the JD Edwards servers and resubmit the transactions. For more details on the current behavior of JD Edwards on database node failover, refer to [Oracle JD Edwards EnterpriseOne and Oracle Database Real Application Cluster \(RAC\) Processing \(Doc ID 2605829.1\)](#).

How does the automatic patching of Oracle Autonomous Database take place?

For Oracle Autonomous Database on Dedicated Exadata Infrastructure, the patching is carried out at the Exadata infrastructure layer and the container database layer. Using the Dedicated Exadata Infrastructure, you can choose a 4-hour maintenance window during which the maintenance activities can be carried out. The database nodes will be patched one after another and automatically brought back to normal state once the process is completed. The patching can start and stop at any time within the 4-hour window and the entire 4 hours might or might not be consumed for the patching activity based on the patch being applied. The order of patching is random and is automatically determined. Some of the nodes might be automatically grouped to enable parallel patching to reduce the time taken for patching.

You have the flexibility to define the time, day, week and month of the year when the patching needs to take place through the Oracle Cloud Infrastructure console. Although there is no downtime for patching, it is recommended that you provide this window during off-business hours or low-operations hours. You can also view the maintenance history and the details of the upcoming maintenance windows through the Oracle Cloud Infrastructure console.

Can I defer the maintenance patching so that it does not impact critical periods for my business?

Yes. You can defer the maintenance patching up to a period of two quarters beyond which the patches will be automatically applied onto your Oracle Autonomous Database on Dedicated Exadata Infrastructure instance.

I do not want to directly apply the patches on my production instance of Oracle Autonomous Database on Dedicated Exadata Infrastructure. How do I ensure that the patches are validated before applying them on the production instance?

To validate the patches, it is suggested that you apply the patches on an Oracle JD Edwards test environment that is configured with Oracle Autonomous Database and is separated from the production environment. The database for the test instance can be on a separate container database from the one on which the production instance of Oracle Autonomous Database is residing. Having the instances on separate container databases enables you to have independent patching schedules on the test and production instances of the database, thereby enabling you to validate patches by setting the patching schedule for the test environment prior to the patching schedule for the production. For more details validating patches, refer to [Manual Deployment of Oracle JD Edwards EnterpriseOne with Oracle Autonomous Database](#)

MAINTENANCE ACTIVITIES WITH ORACLE AUTONOMOUS DATABASE ON SHARED EXADATA INFRASTRUCTURE

What type of maintenance activities are currently supported for JD Edwards with Oracle Autonomous Database on Shared Exadata Infrastructure?

For Oracle Autonomous Database on Shared Exadata Infrastructure, JD Edwards currently supports planned maintenance activities such as database patching and operating system updates using automated kernel reconnection (these activities are explained in detail in the following sections of this FAQ). Unplanned maintenance is currently not supported. Customers will be able to automatically connect their JD Edwards kernels to active database nodes during planned maintenance activities. For more details, refer to: [Planned Database Maintenance of Oracle Autonomous Database on Shared Exadata Infrastructure with JD Edwards EnterpriseOne](#)

What is application continuity service drain time?

Application continuity service drain time is the time taken for active connections on one database node to drain to another active database node. During drain time, new database connections are not permitted to the node that is draining the connections and all new connection requests are redirected to the other active nodes that are not draining connections. When all the active connections are drained, the node will be shut down for planned maintenance.

Application continuity service drain time ensures that there are no in-flight transactions or operations on a particular database node during maintenance and hence prevents transactions from being abruptly aborted resulting in data integrity issues. It also prevents data corruption by avoiding partial completion of transactions. It eliminates the additional time and effort required to resubmit the transactions and revert changes made by the aborted transaction. Therefore, you can carry out planned maintenance activities on Oracle Autonomous Database without disrupting the ongoing transactions or operations.

Can Oracle JD Edwards customers modify drain time on Oracle Autonomous Database on Shared Exadata Infrastructure?

No. Oracle Autonomous Database on Shared Exadata Infrastructure has a default drain time of 5 minutes and this cannot be modified by Oracle JD Edwards customers.

Oracle JD Edwards customers should leverage the automatic kernel reconnection capability of JD Edwards available from tools release 9.2.6 onwards to reconnect the EnterpriseOne kernels from nodes being brought down for maintenance to nodes that are active and receiving connections. Automatic kernel reconnection enables Oracle JD Edwards customers to avoid disruptions to their ongoing transactions or operations during planned maintenance activities on their Oracle Autonomous Database on Shared Exadata Infrastructure. For information on how to configure and leverage automatic kernel reconnection, refer to: [Automated Kernel Reconnection](#)

After I configure my Oracle JD Edwards instance with Oracle Autonomous Database, can I expect that my unplanned maintenance and outages will be taken care of by the self-repairing and failover capabilities of the Autonomous Database?

No. Currently Oracle JD Edwards is not supported for handling unplanned outages with Oracle Autonomous Database on Shared Exadata Infrastructure and hence self-repairing and failover capabilities cannot be leveraged. In some cases, the failover could happen but you might have to restart the JD Edwards servers and resubmit the transactions. For more details on the current behavior of JD Edwards on database node failover, refer to [Oracle JD Edwards EnterpriseOne and Oracle Database Real Application Cluster \(RAC\) Processing \(Doc ID 2605829.1\)](#).

How does the automatic patching of Oracle Autonomous Database take place?

For Oracle Autonomous Database on Shared Exadata Infrastructure, the patching is carried out by Oracle automatically every weekend. The maintenance window has a duration of 2 hours. The database nodes will be patched one after another and automatically brought back to normal state once the process is completed. The patching can start and stop at any time within the 2-hour window and the entire 2 hours might or might not be consumed for the patching activity based on the patch being applied. The order of patching is random and is automatically determined. Some of the nodes might be automatically grouped to enable parallel patching to reduce the time taken for patching.

You do not have any control or flexibility to define the time, day, week and month of the year when the patching needs to take place. Patching will be automatically done by Oracle on the scheduled time. Although there is no downtime for patching, it is recommended that you keep your operations low during the scheduled maintenance window. You can also view the maintenance history and the details of the upcoming maintenance windows through the Oracle Cloud Infrastructure console.

Can I defer the maintenance patching so that it does not impact critical periods for my business?

No. Maintenance patching on your Oracle Autonomous Database on Shared Exadata Infrastructure takes place automatically on the scheduled time every weekend. You do not have any choice to define your custom maintenance windows.

I do not want to directly apply the patches on my production instance of Oracle Autonomous Database on Shared Exadata Infrastructure. How do I ensure that the

patches are validated before they are applied automatically on the production instance?

To validate the patches, it is suggested that you have an Oracle JD Edwards test environment that is configured with Oracle Autonomous Database Shared separated from the production environment. The database for the test instance can be provisioned on an Oracle Autonomous Database Shared instance with an early patch level separated from the production instance of Oracle Autonomous Database Shared that is on regular patch level. Although the early patch level schedule also applies the patches automatically, having the instances on separate Oracle Autonomous Database Shared instances enables you validate patches on the test environment prior to the automatic patching on the production.

As of September 2021, early patching of Oracle Autonomous Database Shared is only available in Ashburn, Phoenix, Frankfurt, London regions of Oracle Cloud Infrastructure. If you want to validate your patches before being applied to production, then you are required to have both the test and production instances of your Oracle Autonomous Database Shared in the same region where early patching is available. You should be using a regional wallet instead of an instance specific wallet in this setup. For more details on validating patches, refer to [Manual deployment of Oracle JD Edwards EnterpriseOne with Oracle Autonomous Database on Shared Exadata Infrastructure](#)

CONNECT WITH US

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com).

Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2021, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Frequently Asked Questions for Oracle Autonomous Database with Oracle JD Edwards EnterpriseOne
October, 2021

