JD Edwards EnterpriseOne

Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with Autonomous Database Learning Path

1.0

JD Edwards EnterpriseOne
Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with Autonomous Database Learning Path

1.0

Part Number: G36617-01

Contents

Preface

Welcome to the JD Edwards EnterpriseOne documentation.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at *http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc* .

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit *http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info* or visit *http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs* if you are hearing impaired.

Related Information

For additional information about JD Edwards EnterpriseOne applications, features, content, and training, visit the JD Edwards EnterpriseOne pages on the JD Edwards Resource Library located at:

http://learnjde.com

Conventions

The following text conventions are used in this document:

Convention	Meaning
Bold	Boldface type indicates graphical user interface elements associated with an action or terms defined in text or the glossary.
Italics	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`Monospace`	Monospace type indicates commands within a paragraph, URLs, code examples, text that appears on a screen, or text that you enter.
> Oracle by Example	Indicates a link to an Oracle by Example (OBE). OBEs provide hands-on, step- by-step instructions, including screen captures that guide you through a process using your own environment. Access to OBEs requires a valid Oracle account.

1 Introduction

Overview

This learning path describes how to deploy JD Edwards EnterpriseOne Release 9.2 using One-Click Provisioning on Oracle Cloud Infrastructure with Linux using an Oracle Autonomous Database. It is important to note that this document is primarily concerned with the basic requirements for an installation of JD Edwards EnterpriseOne on Oracle Cloud Infrastructure. It is not intended as a blueprint for operations in environments such as Development, Test, or Production.

Important Consideration: For best performance, Oracle strongly recommends creating VMs for all JD Edwards components within the same availability domain (AD). This includes the autonomous database itself. The recommended order of creation is to first create the Oracle Autonomous Database. Then based on the AD in which the system automatically creates the database, you should specify that AD for each of the JD Edwards EnterpriseOne server components that you subsequently create. These components include the Enterprise Servers (logic and batch), HTML Web Servers, AIS Web Servers, and Deployment Servers. Further, it is also strongly recommended to keep Production and Non-Production environments in that same AD.

Upon completion of this learning path, you will have a working deployment of JD Edwards EnterpriseOne on Linux in Oracle Cloud Infrastructure with an Oracle Autonomous Database. This process includes the post-installation tasks and the administration of your deployment.

Oracle Support

You can accelerate your provisioning to Oracle Cloud Infrastructure and streamline your process with Oracle Support by entering a planning Service Request (SR) before you start the provisioning process on Oracle Cloud Infrastructure (OCI). Details on how to open the Service Request are on “My Oracle Support, E1: OCI: How to Open A Service Request (SR) For An Oracle Cloud Infrastructure Planning Session” (*Doc ID 2348382.1*).

Oracle Cloud Infrastructure User Interface

The user interface for the Oracle Cloud Infrastructure is constantly evolving. As a result, the screens depicted in this tutorial may not exactly coincide with the current release. This tutorial is routinely updated to include any functional changes to the JD Edwards EnterpriseOne implementation process on the Oracle Cloud Infrastructure. When such updates are made, any differences between the screens in the tutorial and the user interface are reconciled.

Before You Begin

This section provides information and resource requirements critical to understand prior to using One-Click Provisioning on Oracle Cloud Infrastructure.

Fundamentals

You can use the Oracle Cloud Infrastructure (also called "IaaS" - Infrastructure as a Service) to deploy JD Edwards EnterpriseOne using the JD Edwards EnterpriseOne One-Click Provisioning Server. This server features a web-based JD Edwards Provisioning Console interface that enables the provisioning of a fully functional suite of interconnected servers within Oracle Cloud Infrastructure. The required core servers are the Database Server (either as an Oracle Cloud Infrastructure Compute instance, as an Oracle Cloud Infrastructure database service, or as an Oracle Autonomous

Database), the Enterprise Server, the HTML Web Server, and the Application Interface Service (AIS) Server. The One-Click Provisioning Server includes the JD Edwards Server Manager Console (SMC). One-Click Provisioning will also deploy the required Deployment Server into a Microsoft Windows environment.

All servers running in the Oracle Cloud Infrastructure are virtual machines (VMs) that are functionally equivalent to their non-VM physical on-premises machine counterparts.

For all servers you should use the highest version available that is specified in the Oracle Certifications for JD Edwards EnterpriseOne for One-Click Provisioning.

The following servers can be deployed by One-Click Provisioning:

- One or more Database Servers either in Compute or as an Oracle Cloud Infrastructure database service
- One or more Enterprise Servers
- One or more Standard JAS Servers (as standalone instances within WebLogic, not clustered)
- One or more Dedicated HTML Servers (as standalone instances within WebLogic, not clustered)
- One or more AIS Servers (as standalone instances within WebLogic, not clustered)

> **Note:** One-Click Provisioning for Oracle Cloud Infrastructure only supports the deployment of each JD Edwards EnterpriseOne server in a separate VM. You cannot combine all the servers into a single VM instance or server, nor can you combine multiple servers into a single VM instance.

Prior to deploying the web components using One-Click, users must separately install their own licensed version of WebLogic Server. This learning path includes the prerequisites for WebLogic setup for use with JD Edwards EnterpriseOne.

One-Click provisioning can only deploy this machine into a Microsoft Windows Server:

- A single Deployment Server

> **Note:** For Linux systems, to maintain optimal performance, the Microsoft Windows environment must reside in the same Availability Domain (AD) in an Oracle Cloud Infrastructure Region as the companion Linux-based servers. You must have a Deployment Server to apply ESUs (updates), perform Package Builds and commit to future upgrades.

> **Note:** For Cloud implementations, it is imperative that you reuse the SSH key created for your Oracle Compute Cloud Service account for all deployments of JD Edwards EnterpriseOne. Otherwise, any subsequent SSH key creation may cause any previously provisioned instances to be orphaned.

Standard JAS and Dedicated HTML for AIS Servers

Beginning with One-Click Provisioning based on JD Edwards Tools Release 9.2.5, the JD Edwards EnterpriseOne One-Click architecture for HTML and AIS Servers introduces the concept of Standard JAS and Dedicated HTML for AIS Servers.

Using Quick Start Mode, users are required to select a single HTML Server that is dedicated to a specific AIS Server.

Using Advanced Mode, in addition to configuring additional Dedicated HTML and AIS server pairs, users can add a Standard JAS Server that is not dedicated to an AIS Server. A Standard JAS Server functions as a traditional HTML Server handling all non-AIS requests. To scale by pathcode, multiple pairs of HTML and AIS servers can be configured.

Likewise, multiple Standard JAS Servers can be configured for use with or without load balancers. Using one or more Standard JAS Servers is recommended for Production environments.

The following diagram depicts the architecture and following that is a description of the characteristics of each Standard JAS, Dedicated HTML, and AIS server.

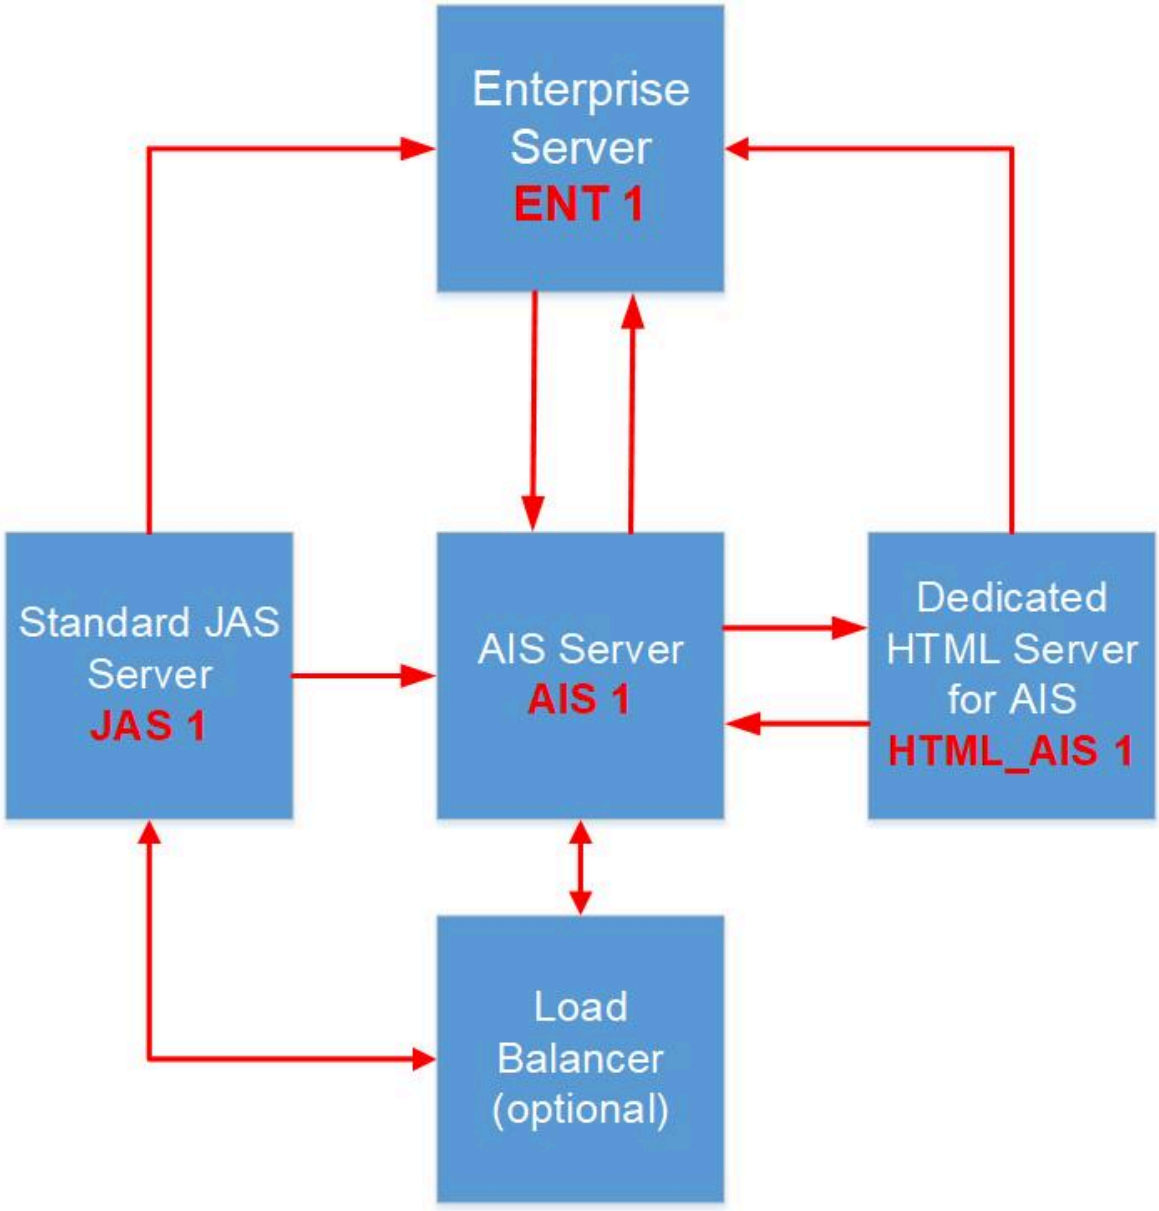

- **Standard JAS Server**

  Note: The Standard JAS Server can only be deployed using the Advanced Mode.

  Referencing the architecture diagram, the Standard JAS Server is JAS 1 and has these characteristics:

  - Can be load balanced (using LBaaS or other)
  - Configured for one-way communication to a specific Enterprise Server (ENT 1)
  - Configured for one-way communication to an AIS Server (AIS 1)

    - Not paired with an AIS Server to perform AIS runtime
    - Not configured to communicate with the Dedicated HTML Server for AIS (HTML_AIS 1)

- **AIS Server**

  Referencing the the pod architecture diagram, the AIS Server is AIS 1 and has these characteristics:

    - Can be load balanced (using LBaaS or other)
    - Paired with a Dedicated HTML Server for AIS (HTML_AIS 1)
    - Configured for two-way communication with both an Enterprise Server (ENT 1) and, for AIS runtime, to the Dedicated HTML Server for AIS (HTML_AIS 1)

- **Dedicated HTML Server for AIS**

  > **Note:** The Quick Start mode can only deploy a single Dedicated HTML Server for AIS. You must use the Advanced Mode to deploy a Standard HTML Server.

  Referencing the the pod architecture diagram, the Dedicated HTML Server for AIS is HTML_AIS 1 and has these characteristics:

    - Cannot be load balanced (using LBaaS or other)
    - Paired to a specific AIS Server (AIS 1)
    - Configured for one-way communication with an Enterprise Server (ENT 1)
    - Configured for two-way communication with a specific AIS Server (AIS 1) for AIS runtime.

Standalone Deployment Server

You can use Advanced Mode to install a Standalone Deployment Server, which is usually defined as part of the required workflow during Quick Start mode. Using this optional flow you can provision Deployment Server by itself without having to orchestrate any other server. However, you must ensure that you know the exact host name of the core JD Edwards EnterpriseOne servers that are, or will be, in the same environment as the Standalone Deployment Server. These core servers include the Database Server, Enterprise Server, HTML Server, and AIS Server. It is not necessary that of any of these core JD Edwards EnterpriseOne servers exists at the time of orchestrating the Standalone Deployment Server. At any time prior or subsequent to the deployment of a Standalone Deployment Server, you can orchestrate the core JD Edwards EnterpriseOne servers. After you have deployed a Standalone Deployment Server, you must ensure that you perform the procedures in the OBE of this Learning Path entitled: ***Performing Post Installation for the Standalone Deployment Server.***

Server Manager Console Centralized Configuration

As part of One-Click Provisioning, Centralized Configuration in the Server Manager Console is enabled automatically. One-Click adds the server group specific INI settings when provisioning individual servers. Post provisioning, the administrator is responsible for applying the group changes on the Centralized Configuration enabled environments.

Process Flow

The following is a process flow diagram for One-Click Provisioning on Oracle Cloud Infrastructure.

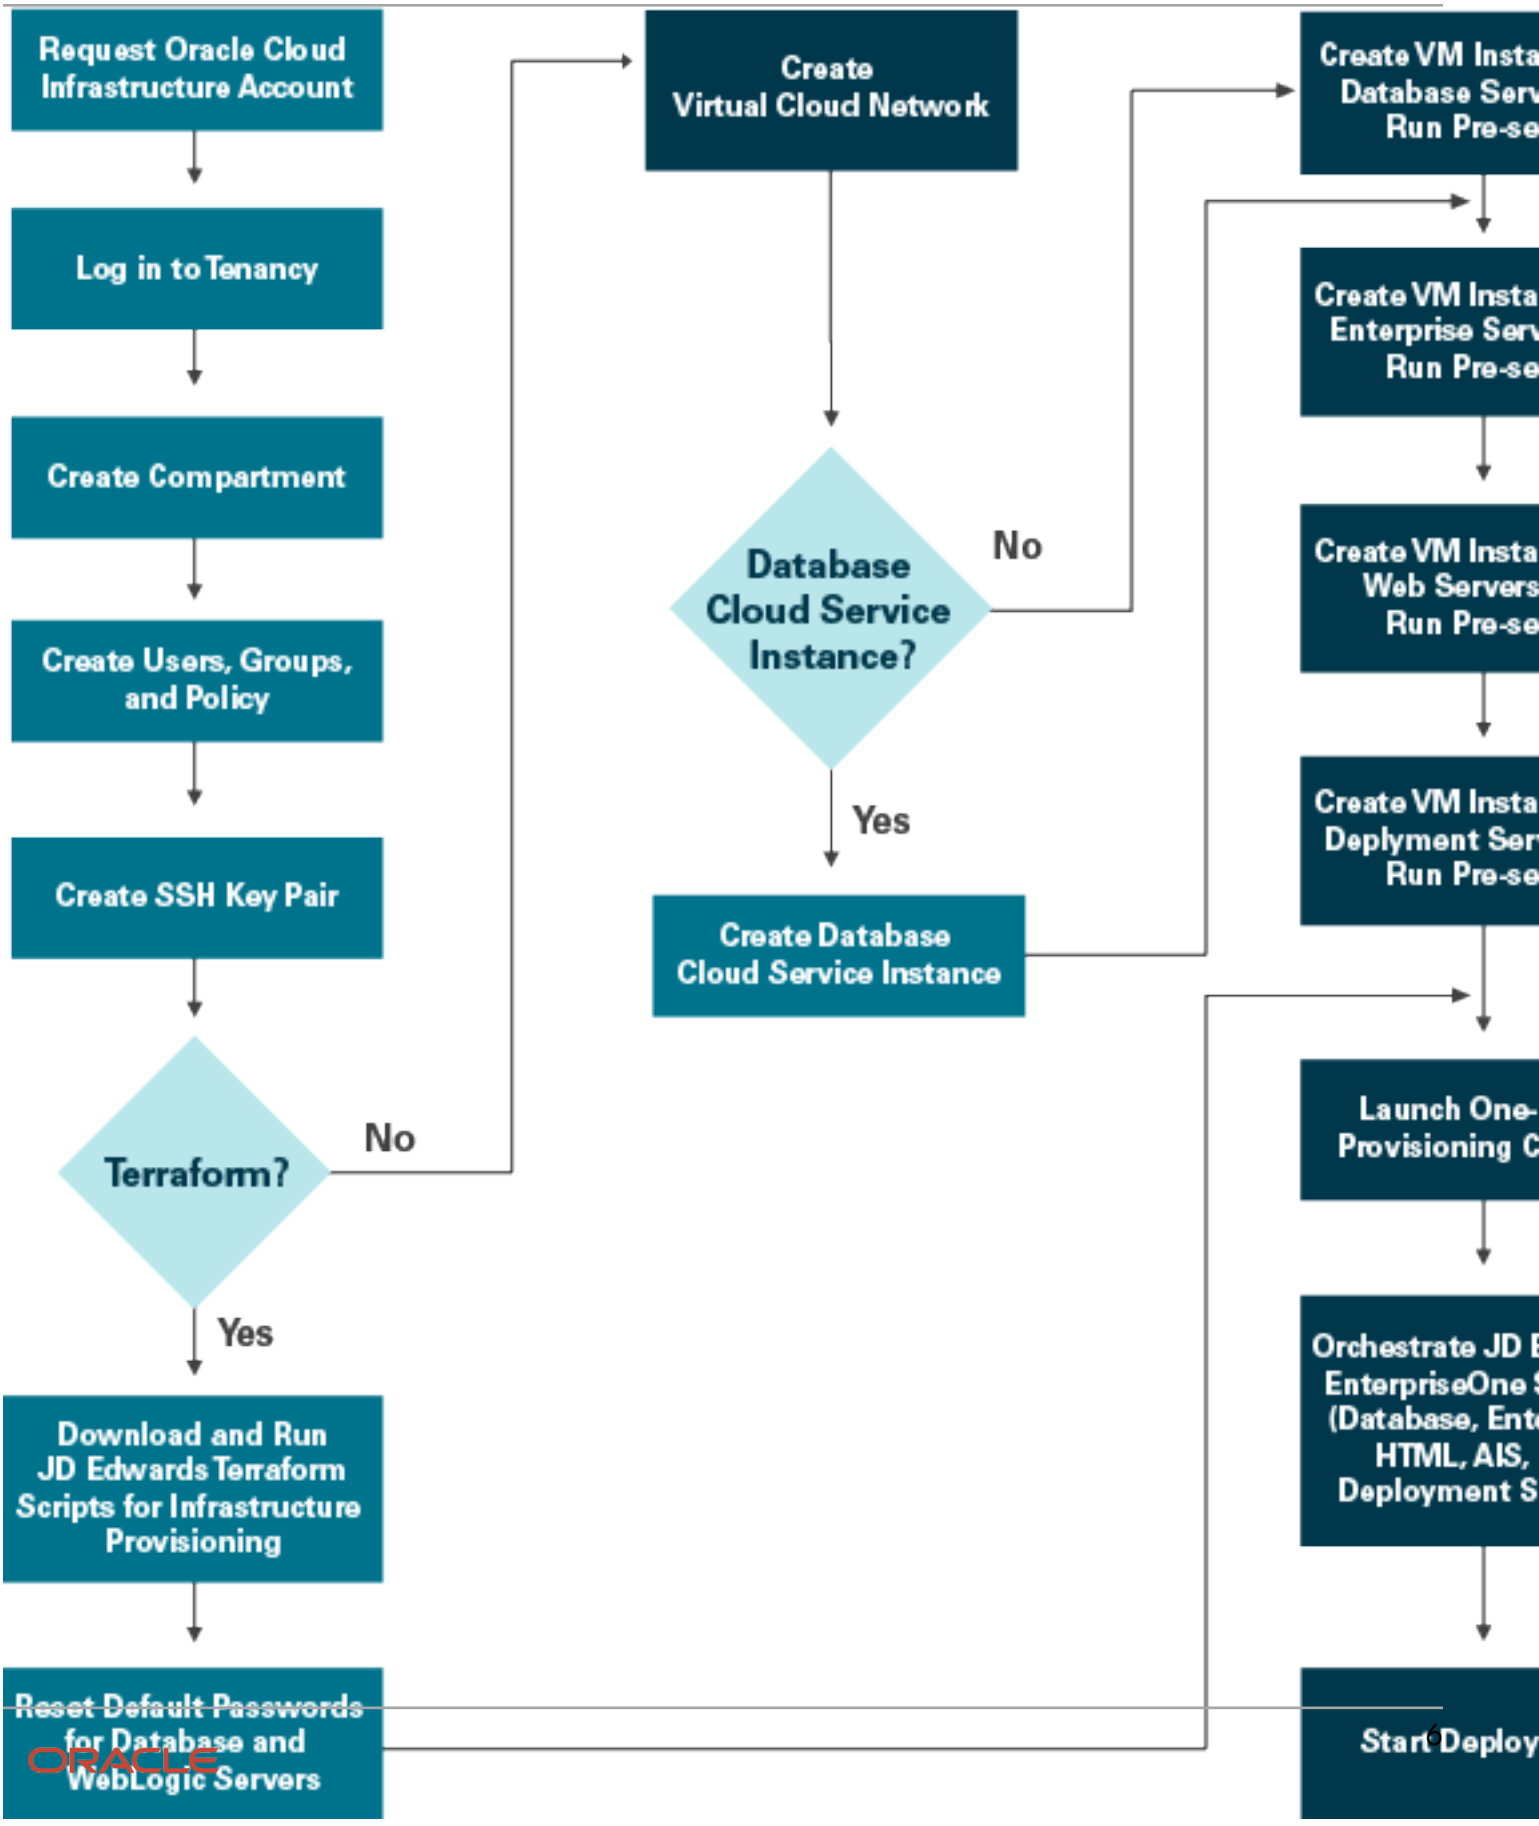
Request Oracle Cloud Infrastructure Account
Log in to Tenancy
Create Compartment
Create Users, Groups, and Policy
Create SSH Key Pair
Terraform?
No
Yes
Download and Run JD Edwards Terraform Scripts for Infrastructure Provisioning
Reset Default Passwords for Database and WebLogic Servers
Create Virtual Cloud Network
Database Cloud Service Instance?
No
Yes
Create Database Cloud Service Instance
Create VM Insta Database Serv Run Pre-se
Create VM Insta Enterprise Serv Run Pre-se
Create VM Insta Web Servers Run Pre-se
Create VM Insta Deplyment Ser Run Pre-se
Launch One- Provisioning C
Orchestrate JD E EnterpriseOne S (Database, Ente HTML, AIS, Deployment S
Start Deploy

Supported Software Versions

The following table lists the supported software versions for Oracle and Microsoft components running on Oracle Cloud Infrastructure:

Supported Software Versions	
Operating System	
- Provisioning Server	Oracle Enterprise Linux 9.6
- Oracle Database Server in Compute	Oracle Enterprise Linux 9.6
- Oracle Database Server in DB Systems	Oracle Enterprise Linux 8.10 (Oracle 26ai: current, subject to change)
- Oracle Autonomous Database	Autonomous Transaction Processing on Dedicated Exadata Infrastructure - Version 26ai
- Enterprise Server	Oracle Enterprise Linux 9.6
- Oracle WebLogic Server	Oracle Enterprise Linux 9.6
- Deployment Server	Microsoft Windows Server 2022
Oracle Database	Oracle 19C (Compute Database) Oracle 26ai (Database System, Oracle Autonomous Database)
Oracle WebLogic Server	14.1.1.0
Oracle WebLogic Server Patches for 14.1.1.0	p28186730, p38412913 Tip: You can verify your patch level running this command from the **<OH>/OPatch** location on your WebLogic Server: **./opatch lspatches**
Java Development Kit (JDK) (Required for Oracle WebLogic Server)	Version 1.8.0 up to Update 471
Load Balancer (Optional)	OCI Load Balancer
JD Edwards EnterpriseOne One-Click System	Tools Release 9.2.26.1 64-bit ESU up to JN21495

	UDO up to UDO_9.2_10719 + UN26_UDO_Bundle Planner JN21409 Data Pack is DP0289201

Prerequisite Knowledge and Skills

You must have a fundamental understanding of the Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at the sites:

- *Oracle Cloud Infrastructure*

Prerequisite Tools and Resources

You must obtain the PuTTY tool (*http://www.putty.org*) for generating SSH key pairs on the client machine that you will use to connect to any Linux server deployed by One-Click Provisioning.

2 Upgrading from Prior Releases of EnterpriseOne

Upgrade Considerations

If you are planning an Applications Upgrade to a One-Click installation of JD Edwards EnterpriseOne, you should be aware of important technical considerations.

These considerations are described in the Applications Upgrade Guide which describes an upgrade to the base Applications Release 9.2.

Refer to this section of the upgrade guide:

Technical Considerations for Applications Upgrade to a 64-bit Version of 9.2

3 Planning Your Deployment

Minimum Resource Requirements

The table below specifies the minimum resource requirements to install and run JD Edwards on Oracle Cloud Infrastructure using an Oracle Autonomous Database with an instance of Autonomous Transaction Processing on Dedicated Exadata Infrastructure (ATP-D). Your environment may require additional resources based on transaction volumes, number of users, availability requirement, integrations, and business requirements.

Note: For implementation of JD Edwards EnterpriseOne on Oracle Cloud Infrastructure, users can choose any Shape Series that provides:

- x86-compatible processors (such as Intel and AMD). Support for RISC (reduced instruction set) processors (such as ARM) is specifically excluded.
- Minimum of 2 OCPUs
- Minimum of 30 GB memory per OCPU

JD Edwards EnterpriseOne Server Type	Minimum Recommended			Notes
	oCPU	Memory	Storage Value (GB)	
Required JD Edwards Components				
One-Click Provisioning Server	2	30 GB	150 GB	This resource type includes Server Manager. It also includes an Oracle database client, which is required to connect to an ATP-D instance for provisioning the JD Edwards EnterpriseOne database.
Deployment Server	2	30 GB	210* GB For all the four path codes	Using One-Click Provisioning, customers must install all the four path codes. There is no automated way to add additional path codes post-deployment. * The storage volume provided here is in addition to the space required by the Windows operating system, which can be up to 45 GB. Therefore, the minimum recommended storage volume is 256 GB.
Autonomous Database	Not Configurable	Not Configurable	1024 GB	The minimum storage that can be allocated to an Oracle Autonomous Database is 1024 GB.

JD Edwards EnterpriseOne Server Type	Minimum Recommended			Notes
	oCPU	Memory	Storage Value (GB)	
				JD Edwards EnterpriseOne requires 50 GB for a single path code and shared data, and a further 20 GB is required for each additional path code. The path code requirements provided here are for demo data only and should be adjusted for required business data space.
Enterprise Server	2	30 GB	75 GB	
WebLogic Server • HTML Server for Web Client • HTML Server for AIS Server • AIS Server	-	-	30 GB	For the operating system and WebLogic Server
	2	30 GB	20 GB	For each web instance
	2	30 GB	20 GB	For each web instance
	2	30 GB	20 GB	For each web instance
Optional JD Edwards Components				
Business Services Server (BSSV)	2	30 GB	50 GB	For each web instance
Transaction Server for Real Time Events (RTEs)	2	30 GB	50 GB	For each web instance
Application Development Framework Server (ADF)	2	30 GB	50 GB	For each web instance
One View Reporting (OVR) Server and BI Publisher (BIP) Server	2	30 GB	50 GB is required for a single path code	Each additional path code requires 10 GB
Development Client	2	30 GB	100 GB	For each Development Client installation

Note: Optional components are not deployed by One-Click. However, the web components can be manually added through Server Manager and the Development Client can be added in a new Microsoft Windows instance using the traditional on-premises methodology.

Understanding Port Restrictions

This section provides an overview of the restricted ports that cannot be defined or used while creating any web component or server, or both. You should be aware of the restricted ports that cannot be defined or used while creating any web component or server, or both.

The specific port restrictions for any One-Click Provisioning deployment of JD Edwards EnterpriseOne are grouped as follows:

- One-Click Provisioning Console for JD Edwards
- All Internet Browsers
- Google Chrome and Mozilla Firefox Browsers

One-Click Provisioning Console for JD Edwards

- Any port below 1024 is restricted.

All Internet Browsers

The following are restricted ports enforced by the rules of any internet browser:

- 2049
- 4045
- 6000

Google Chrome and Mozilla Firefox Browsers

In addition to the above mentioned restricted ports for any internet browser, the Google Chrome and Mozilla Firefox browsers block specific ports which they deem as unsafe to use on HTTP/HTTPS protocol. These restricted ports are:

- 3659, // apple-sasl / PasswordServer
- 6665, // Alternate IRC [Apple addition]
- 6666, // Alternate IRC [Apple addition]
- 6667, // Standard IRC [Apple addition]
- 6668, // Alternate IRC [Apple addition]
- 6669, // Alternate IRC [Apple addition]

Note: It may be possible to configure Chrome and/or Firefox to change these restrictions.

Creating a Pre-installation Worksheet

Machine	**IP Address**
Provisioning Server and Server Manager (same machine)	https://__.__.__.__:3000 where **https** is the only supported browser protocol, and where **3000** is always the port, which you must include as part of the address
Machine	**Hostname**
Database Server in Compute or a DB System	
Autonomous Database (ADB)	ADB Admin User Password: _____ ADB WALLET: _____ Auth Token of User: _____ Tenancy: _____ Username: _____ Region: _____ Bucket: _____
HTML Server	
AIS Server	
Deployment Server	Hostname: _____ Location Name: _____

Note:
Do not use a fully qualified domain name for host names; you should only specify the first node of the domain name.

For the Database, Enterprise, and Deployment Servers, machine names are limited to a maximum of 15 **lower-case** alphanumeric characters as limited by the JD Edwards EnterpriseOne database table and application design.

For all servers, you cannot use special characters in the name, such as an underscore "_". The user interface will enforce this restriction.

User Account	Password
oracle	These users do not have a login password. Access is through SSH Keys only, logging in using the username **opc**. Enter SSH Key names here.
jde920	
em_manager_user	
sys, system	
weblogic	
JDE	
jde_admin	
Initial Windows Machine Administrator (temporary)	
Site Key Passphrase	
Server Manager Console	
SQL DB User (SA) **Note:** This is only applicable to Oracle Cloud Infrastructure with Microsoft Windows and On-Premises Microsoft Windows with SQL Server database.	

> **Note:** The password for ***any*** Windows user on ***any*** Windows machine must not contain the $ or ! characters; using these characters violates the Oracle password policy and will result in denied access.

Generating Secure Shell (SSH) Key Pairs on Your Local System

This section outlines the requirement to generate SSH key pairs on your local system, which can be UNIX or Windows.

> **Note:** All references to UNIX also apply to Linux.

If you already have an SSH key pair, you may use that to connect to your environment whether it is running On-Premises or in Oracle Cloud Infrastructure.

> **Tip:** The best practice is to create at least two SSH keys, a primary and a backup, because if for any reason an SSH Key is no longer valid, access to the instance or machine would be lost with no means to recover the access. A user cannot access an instance or machine without using an SSH Key.

For instances in Oracle Cloud Infrastructure, it is recommended you use the Oracle Cloud Shell to interface with a Oracle Cloud Infrastructure instance. Oracle Cloud Shell is browser-based, does not require installation or configuration of anything on your laptop, and works independently of your network setup. The below referenced Oracle documentation for generating SSH keys provides instructions for these options:

- Oracle Cloud Shell
- Apple MacOS
- Microsoft Windows 10
- Prior Microsoft Windows Versions
- SSH Keys for Linux

Refer to this Oracle tutorial that describes the process to generate an SSH key pair for Oracle Compute Cloud Service instances. Fundamentally, you can use these same procedures for On-Premises environments.

Generate SSH keys

4 Performing Setup Tasks in Oracle Cloud Infrastructure

Logging into Oracle Cloud Infrastructure

Supported Browsers

Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

To sign in to Oracle Cloud at *https://cloud.oracle.com*, you need:

- User name and password
- Your cloud account name

When your tenancy is provisioned, Oracle sends an email to the default administrator at your company with the sign-in credentials and URL. This administrator can then create a user account for each person who needs access to Oracle Cloud Infrastructure. Check your email or contact your administrator for your credentials and account name.

Signing In for the First Time

Links for signing in are also provided in your welcome email.

1. Open a supported browser and go to *https://cloud.oracle.com*.
2. Click **Sign In**.
3. Enter your **Cloud Account Name** and click **Next**.
4. Enter your user name and temporary password from your welcome email. You will be prompted to change your temporary password.

After you sign in, the Console Home page is displayed.

About the Console URL

Alternatively, you can sign in directly to Oracle Cloud Infrastructure using the Console URL. When you sign up to use Oracle Cloud Infrastructure, you receive a customized URL for your organization. For example:

https://console.us-ashburn-1.oraclecloud.com/?tenant=CompanyABC

If you instead use the base URL (https://console.us-ashburn-1.oraclecloud.com), you are prompted to specify your tenant (or cloud account name) on the sign-in page, along with your user name and password.

Creating a Compartment

This section shows you how to create a Compartment in Oracle Cloud Infrastructure. You must define a Compartment as part of core functionality in Oracle Cloud Infrastructure.

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

To create a Compartment for JD Edwards EnterpriseOne on Oracle Cloud Infrastructure, follow the steps in the *Creating a Compartment* section of the Oracle Cloud Infrastructure Documentation.

Creating a Group

This section shows you how to create a Group in Oracle Cloud Infrastructure.

For additional information on using Groups in Oracle Cloud Infrastructure, refer to the section entitled: **Add a New Group in the Oracle Cloud Infrastructure Console** in *Adding Groups and Users for Tenancies Federated with Oracle Identity Cloud Service*.

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

To create a group using the Oracle Cloud Infrastructure Console, follow the steps in the *Creating a Group* section of the Oracle Cloud Infrastructure Documentation.

Creating a User

This section shows you how to create users in Oracle Cloud Infrastructure. For additional information on creating users in Oracle Cloud Infrastructure, refer to the topic **Create a User** in *Adding Users*.

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

> **Note:** You must perform this task as an Oracle Cloud Infrastructure user who has sufficient permissions to create and update resources within Oracle Cloud Infrastructure. For more information, refer to *Adding Users* in *Oracle Cloud Infrastructure Documentation*.

> **Note:** You must create a Domain before creating a User. Refer to *Creating an Identity Domain* in Oracle Cloud Infrastructure Documentation.

To create a User, see *Creating a User* in Oracle Cloud Infrastructure Documentation.

Generating an Auth Token for a User

This section shows you how to generate an Auth token for a user in Oracle Cloud Infrastructure. This token is required for setting up a connection to an Oracle Autonomous Database.

For additional information on creating users in Oracle Cloud Infrastructure, refer to the topic "Create a User" in *Adding Users.*

Prerequisites

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

Generating an Auth Token for a User

An Auth Token is required for setting up a connection to an Oracle Autonomous Database.

To generate an Auth Token for a user, refer to this topic in *Oracle Cloud Infrastructure Documentation:*

Getting an Auth Token

Adding Users to Groups

This section tutorial shows you how to add users to groups in Oracle Cloud Infrastructure.

For additional information on using Groups in the Oracle Cloud Infrastructure, refer to the section entitled: Managing *"Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console"* in *Oracle Cloud Infrastructure Documentation*. For additional information on creating users in Oracle Cloud Infrastructure, refer to *Adding Users*.

Prerequisite

- You must have already created a user by following the process described in the preceding module "*Creating a User*" in this Learning Path.
- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

To add users to the respective groups, see *Adding a User to a Group* in Oracle Cloud Infrastructure Documentation.

Creating a Policy

This section describes the minimum setup IAM policies required to use JD Edwards EnterpriseOne Infrastructure Provisioning for Reference Architecture. The user who is running the Infrastructure Provisioning must have these policy settings for the group to which they belong. The tenancy administrator for Oracle Cloud Infrastructure is responsible for creating and assigning these requisite policies. If you are unsure of your policy settings you should check with the tenancy administrator.

> **Note:** As described above, this procedure may only be necessary in certain regions or for certain classes of subscribers or users in Oracle Cloud Infrastructure.

For additional information on using policies in Oracle Cloud Infrastructure, refer to the documentation for Oracle Cloud Infrastructure in the *Managing Policies*.

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

To create a policy using the Oracle Cloud Infrastructure Console, refer to *Creating a Policy* in Oracle Cloud Infrastructure Documentation.

In the Policy Builder section, click the **Customize/Advanced** button.

In the Policy Builder dialog, enter the following statements:

```
Allow group <group_name> to read announcements in tenancy

Allow group <group_name> to manage virtual-network-family in tenancy

Allow group <group_name> to manage load-balancers in tenancy

Allow group <group_name> to manage instance-family in tenancy

Allow group <group_name> to read app-catalog-listing in tenancy

Allow group <group_name> to manage volume-family in tenancy

Allow group <group_name> to manage file-family in tenancy

Allow group <group_name> to manage database-family in tenancy

Allow group <group_name> to manage autonomous-database-family in tenancy

Allow group <group_name> to manage keys in tenancy

Allow group <group_name> to manage dns in tenancy

Allow group <group_name> to manage object-family in tenancy

Allow group <group_name> to manage compartments in tenancy

Allow group <group_name> to manage tag-namespaces in tenancy

Allow group <group_name> to manage vaults in tenancy

Allow group <group_name> to manage secret-family in tenancy
```

where ***<group_name>*** is the name of group that you specified as described in the section of this Learning Path entitled "Creating a Group".

Creating a Virtual Cloud Network

This section shows you how to create a Virtual Cloud Network (VCN) in Oracle Cloud Infrastructure using the Start VCN Wizard.

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

Before you can launch an instance, you must have a Virtual Cloud Network (VCN) in Oracle Cloud Infrastructure. For more information refer to this topic: *Creating the VCN and Subnets to Use with Oracle Functions, if they don't exist already*.

Oracle JD Edwards recommends using the **Start VCN Wizard** to create a complete set of networking resources using the concept of regional networking, which includes route tables with private and public subnets across all Availability Domains (ADs) in your region.

On Virtual Cloud Networks in <your_compartment>, click the **Start VCN Wizard** button.

On Create a VCN with Internet Connectivity - Configuration, complete these fields in the **Basic Information** and **Configure VCN and Subnets** sections:

Basic Information

- VCN NAME
- COMPARTMENT

Configure VCN and Subnets

- VCN CIDR BLOCK (see below Note)
- PUBLIC SUBNET CIDR BLOCK (see below Note)
- PRIVATE SUBNET CIDR BLOCK (see below Note)

> **Note:** For CIDR block values, you can either use the example values or obtain the values from your network engineer.

DNS RESOLUTION

In this section, ensure the following check box is selected:

- USE DNS HOSTNAMES IN THIS VCN

To create a VCN, see *Creating a VCN* in the Orcale Cloud Infrastructure Documentation.

Creating Rules for a VCN Security List

This section shows you how to create rules for virtual cloud network (VCN) security list in Oracle Cloud Infrastructure for use with JD Edwards EnterpriseOne One-Click Provisioning.

Prerequisite

- You must have created a Virtual Cloud Network as described in the preceding section of this Learning Path entitled: " ***Creating a Virtual Cloud Network*** ".
- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

You must create rules for the Public subnet that was automatically created by the workflow wizard when you created the VCN.

> **Note:** If you have created a VCN using any other method, you must destroy that VCN and follow the documented procedure in this Learning Path.

Follow the below steps to create rules for a VCN security list:

1. Navigate to Networking > Virtual Cloud Networks
2. In the List Scope section, verify the correct COMPARTMENT is selected.
3. Click the link for the VCN you created and click on Subnet
4. On the list of subnets, click on the link to open the Public-subnet-<vcn_name> definition.
5. In Subnet Details, click on Security section and click on the link Default Security List for <vcn_name>.
6. On the Default Security List for <vcn_name> page click on security rules, where default rules are displayed, for the default rule for Port 22, it is recommended to edit the rule to specify a Source CIDR range of 123.123.123.123/32.
7. On the Add Ingress Rules dialog, add rules for ports that should be open to the Private subnet as shown in the following table.

Stateless Check box (always unchecked)	Source CIDR	IP Protocol (always TCP)	Source Port Range (Always All)	Destination Port Range	Comment
	Public subnet (see Note 1)			3000	One-Click Provisioning Console
				3389	Remote Desktop Protocol (RDP)
				8998	Server Manager Console https (SSL)
				8999	Server Manager Console http (non-SSL)
				User-specified SSL port for each HTML instance using Provisioning Console	HTML Server
				User-specified SSL port for each AIS instance using Provisioning Console	AIS Server
	Private subnet only (see Note 2)			1521	Oracle Database Server
				1433	SQL Server Database Server
				User-specified port at installation	WebLogic Admin Console - https (SSL)
				User-specified port at installation	WebLogic Admin Console - http (non-SSL)
				5150	Required for One-Click Provisioning to validate ports over the network.
				5985	Deployment Server Used for the winrm port for communication between Provisioning Server and Deployment Server
				6017-6026	Enterprise Server These ports depend on release level of EnterpriseOne and user specification.
				14501-14520	JMX Port (required for Server Manager Agent communication)

Note: Source CIDR. In this form, the IP address value represented by **Source CIDR** (where CIDR means Classless Inter-Domain Routing) is a function of the VCN. This is the source IP address from where connection is allowed on a particular port. The syntax x.x.x.x/x provides an IP address range. For example, 10.0.0.0/16 means 10.0.0.1 to 10.0.255.255 (where /16 is the bit length of the subnet mask), while 0.0.0.0/0 means all IP addresses.

Note: It is recommended to not open any port for all IP addresses. Instead you should only open ports to specific Public IP addresses by either setting your VCN or by using the OCI function allowlist (formerly whitelist) to specify a CIDR (range of IP addresses). For example, to open an IP port for a range of addresses such as 123.123.123.123, add 123.123.123.123/32 as source CIDR for the port.

Note: Private Network. These ports should not use CIDR values that are open to the Internet. They should be specified on a Private Network for machine-to-machine communications. For example, 10.0.0.0/16. Ensure that your Private subnet is configured so that the ports list in the above table are open.

For details, see *Creating a Security List* in Oracle Cloud Infrstructure Documentation.

5 Configuring the Database

Creating an Autonomous Transaction Processing Instance on Dedicated Exadata Infrastructure in Oracle Cloud Infrastructure

Oracle Cloud Infrastructure supports these Autonomous Database systems:

- Autonomous Transaction Processing
- Autonomous Data Warehouse

> **Note:** For use with JD Edwards EnterpriseOne, only the *Autonomous Transaction Processing* system is supported.

The Autonomous Transaction Processing system runs on:

- Dedicated Exadata Infrastructure
- Shared Exadata Infrastructure

> **Note:** For use with JD Edwards EnterpriseOne, only the *Dedicated Exadata Infrastructure* system is supported.

> **Note:** As already detailed in the Overview of this Learning Path, the recommended order of creation is to first create the Oracle Autonomous Database. Then based on the AD in which the system automatically creates the database, you should specify that AD for each of the JD Edwards EnterpriseOne server components that you subsequently create. These components include the Enterprise Servers (logic and batch), HTML Web Servers, AIS Web Servers, and Deployment Servers. Further, it is also strongly recommended to keep Production and Non-Production environments in that same AD.

Prerequisites

- Refer to the "Prerequisites" section in this Oracle Cloud Infrastructure documentation: *Creating Autonomous Transaction Processing on Dedicated Infrastructure.*
- You must have already created a compartment in which the Autonomous Database will reside. Refer to the module "Performing Setup Tasks in Oracle Cloud Infrastructure" of this Learning Path.

Creating an Autonomous Database Instance for Autonomous Transaction Processing Dedicated

Use one of the following two procedures to create an Autonomous Transaction Processing instance on Dedicated Exadata Infrastructure in an Autonomous Database running in Oracle Cloud Infrastructure.

You can access the Create Autonomous Database page by either of these means:

- Using Quick Actions

- Using the Main Navigation Menu

Using Quick Actions

1. To access the Create Autonomous Database page using Quick Actions, in the Quick Actions section click the **Create an ATP database** icon.

> **Note:** The selection of **Quick Actions** on the main entry into Oracle Cloud Infrastructure is variable; therefore this action may or may not be available in subsequent versions of this user entry.

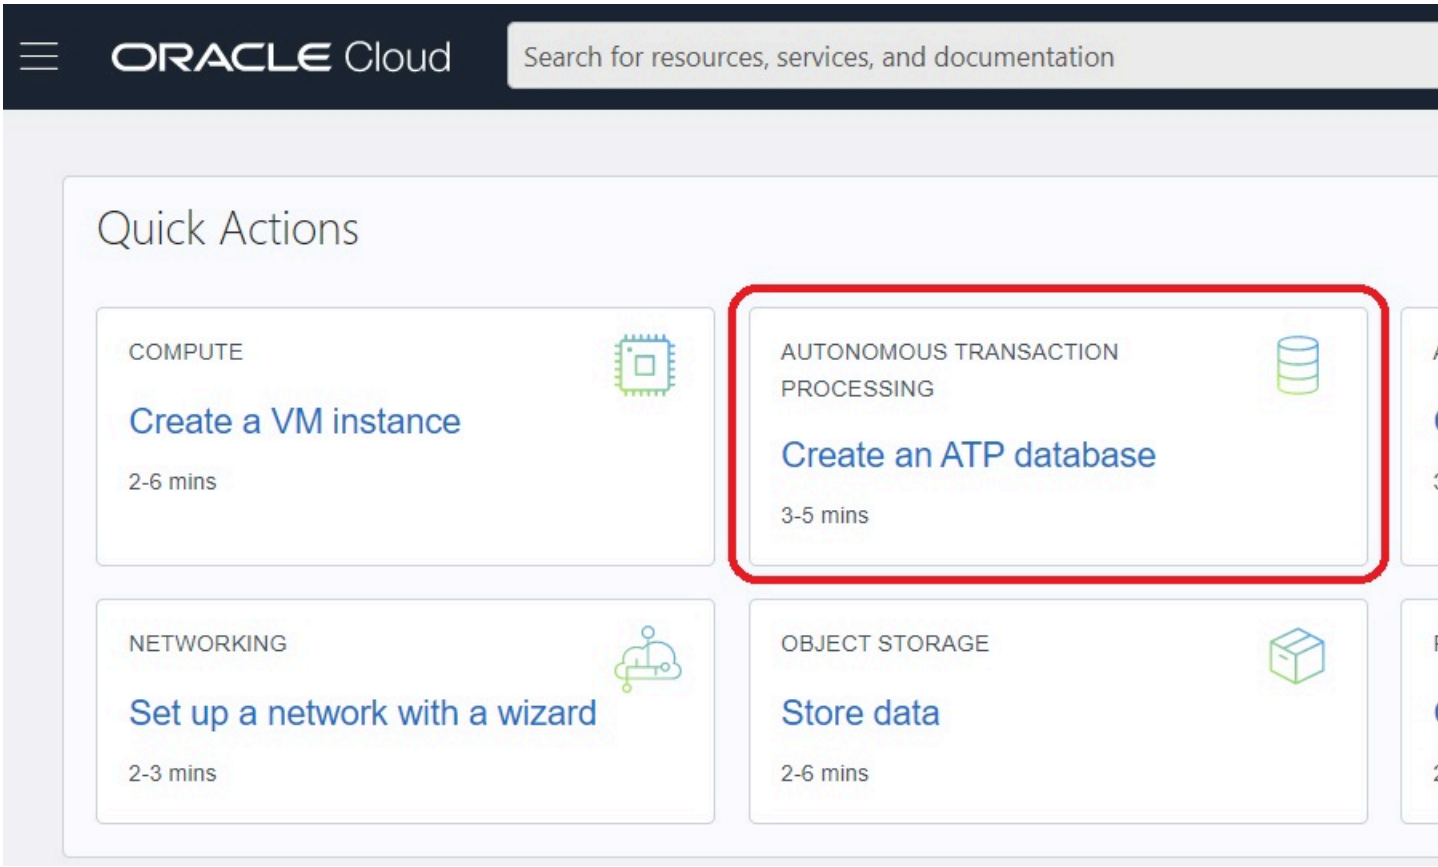

The **Create Autonomous Database** screen is displayed.

> **Note:** By default, the Autonomous Database creation workflow launches in your current compartment and you can use the network and subnet resources in that compartment. Otherwise, you can use the menu in the **Compartment** section to choose the compartment into which you want to create your Autonomous Transaction Processing instance. In this example, the user has launched the Quick Action process from a compartment that has been named as ATP.

2. Click the **Transaction Processing** tile.

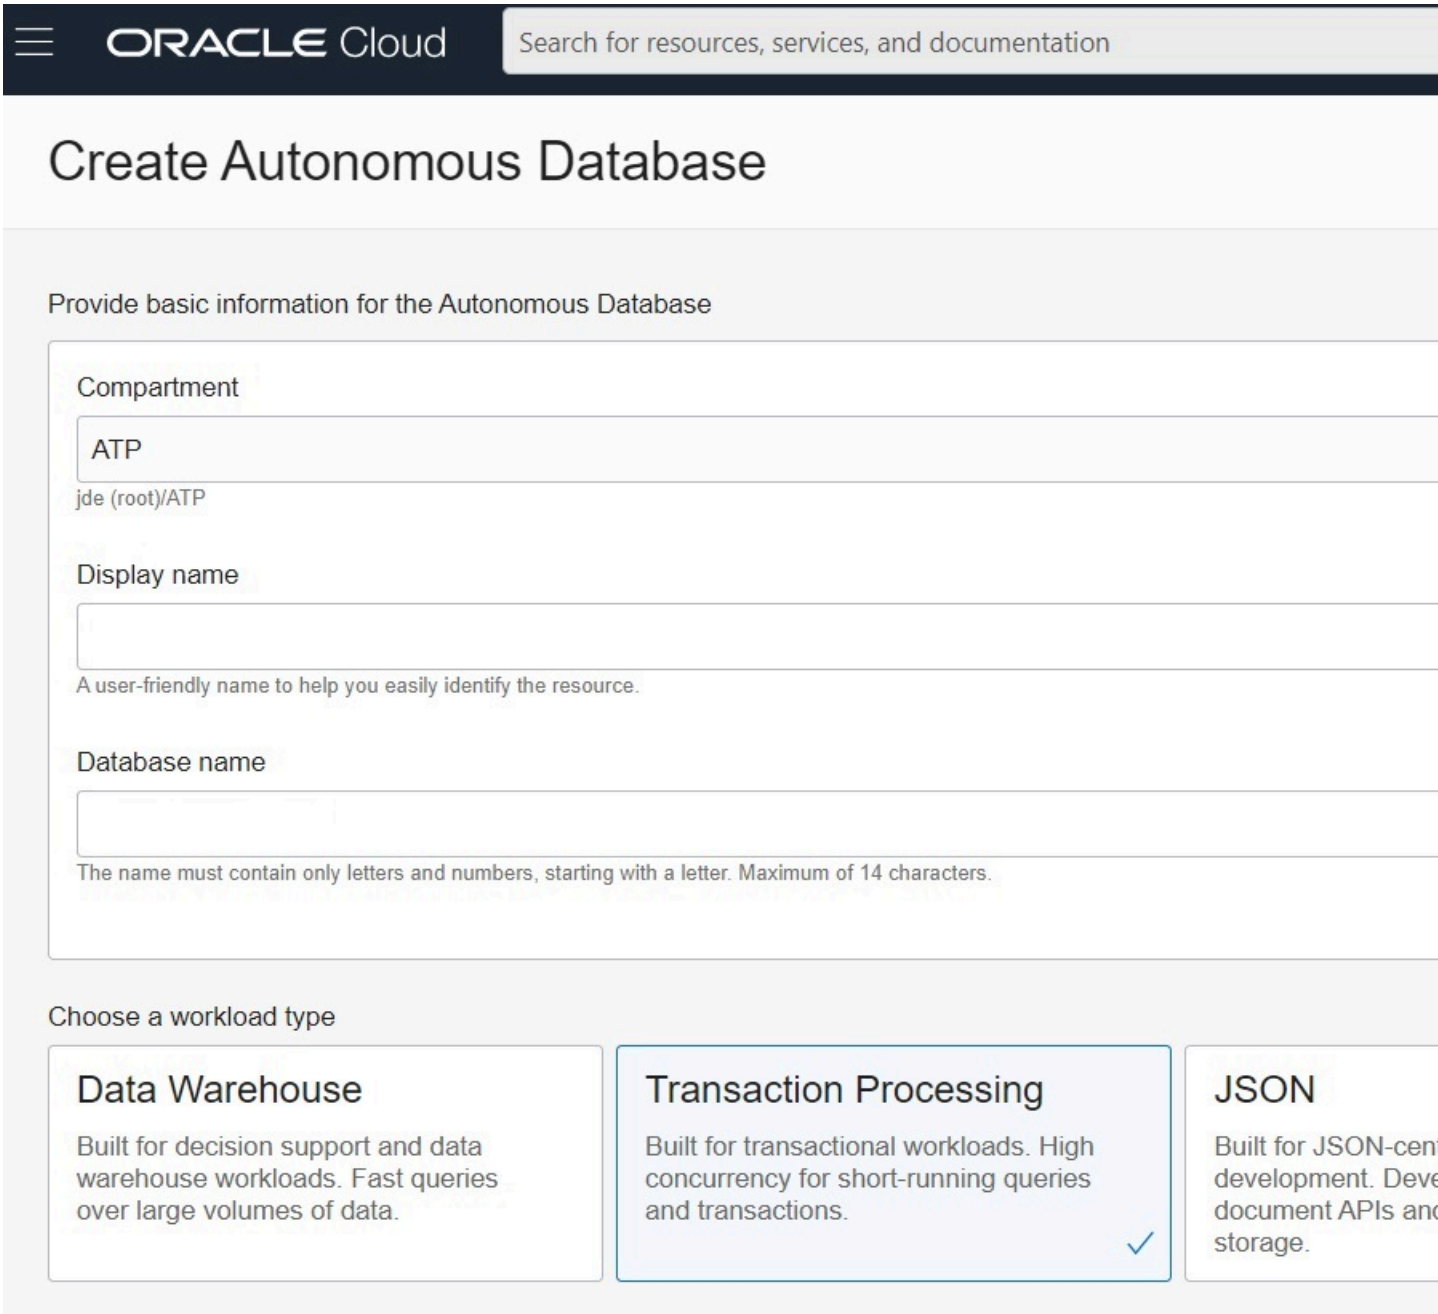

3. To continue with the creation of this instance, complete the procedure described in this Oracle Cloud Infrastructure document: Creating Autonomous Transaction Processing on Dedicated Infrastructure.

Using the Main Navigation Menu

1. On the Oracle Cloud Infrastructure Console Home page, click the navigation menu in the upper-left corner and select **Autonomous Transaction Processing**.

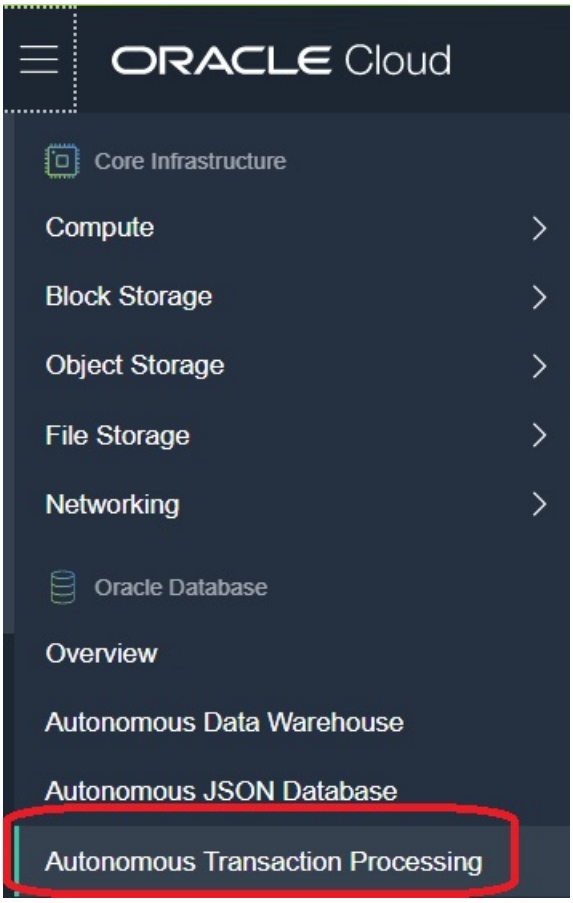

2. In the left pane, select **Autonomous Database**, then in the right pane, select **Create Autonomous Database**.

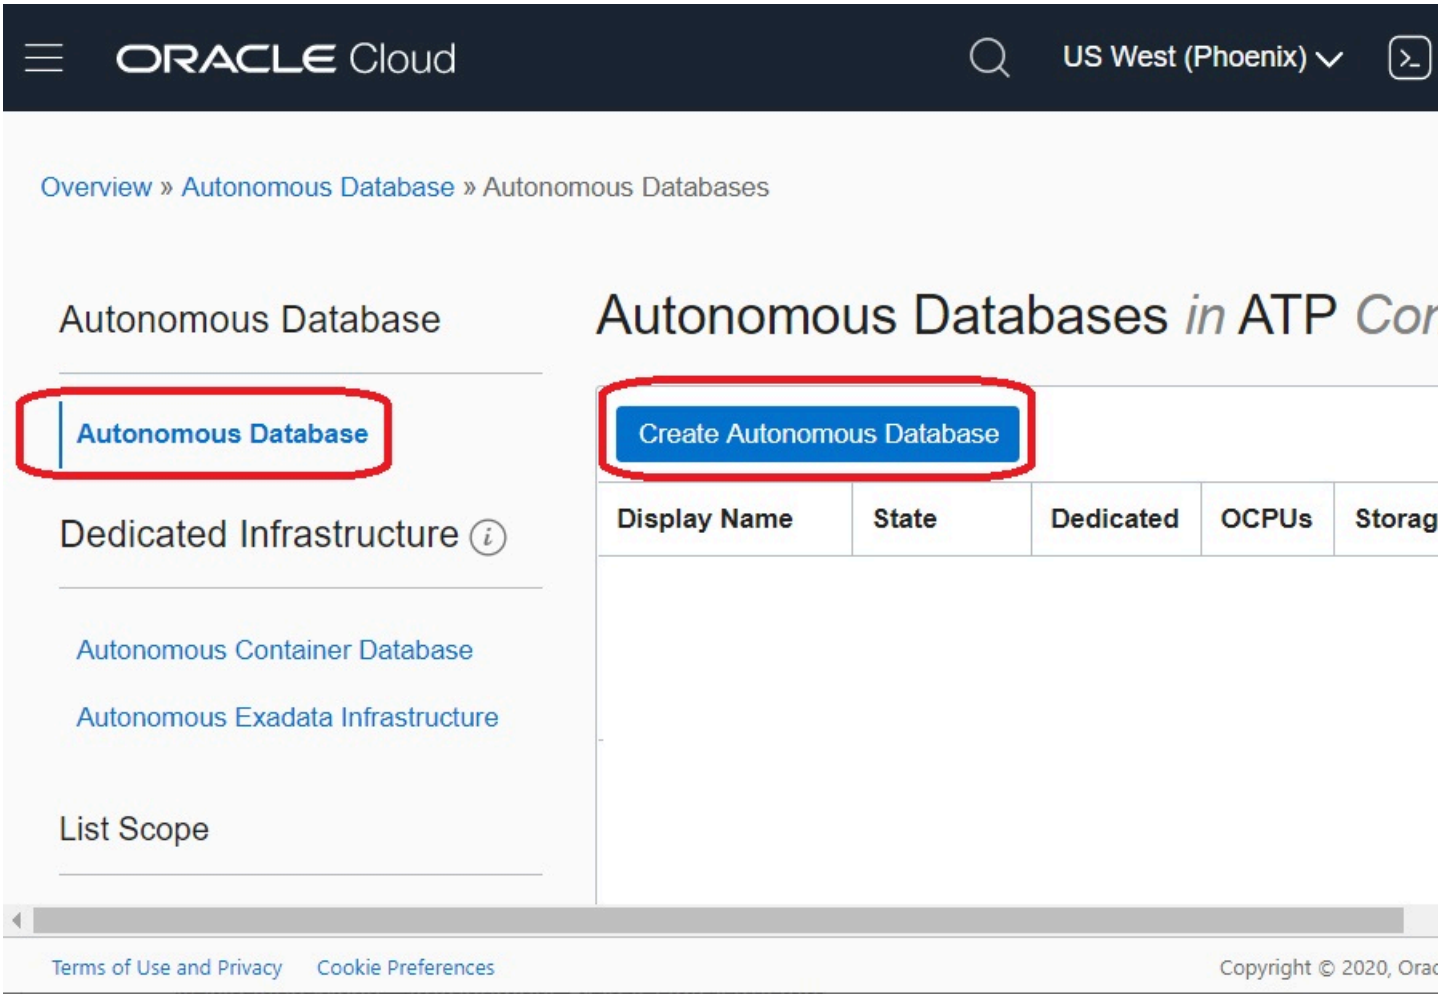

> **Note:** By default, the Autonomous Database creation workflow launches in your current compartment and you can use the network and subnet resources in that compartment. Otherwise, you can use the menu in the Compartment section to choose the compartment into which you want to create your Autonomous Transaction Processing instance. In this example, the user has launched the Quick Action process from a compartment that has been named as ATP.

3. Click the **Create Autonomous Exadata Infrastructure** button.
4. To continue with the creation of this instance, complete the procedure described in this Oracle Cloud Infrastructure document: Creating Autonomous Transaction Processing on Dedicated Infrastructure.

Downloading a Database Wallet for Autonomous Transaction Processing on Dedicated Infrastructure

This section describes how to download the database wallet, which is required to connect to the Autonomous Database. Because there is no direct machine access for Autonomous Transaction Processing on Dedicated Infrastructure, you must connect to the underlying database by using:

- Database Wallet (which contains `tnsnames.ora and sqlnet.ora` related files)
- Connection Strings
- Web SQL Developer

For additional details, see Connecting to Autonomous Transaction Processing

Creating a Database Wallet for Autonomous Transaction Processing on Dedicated Infrastructure

This wallet is required as input to the Infrastructure Provisioning console during the orchestration to facilitate JD Edwards EnterpriseOne Infrastructure Provisioning.

1. On the Oracle Cloud Infrastructure home page, click the navigation menu in the upper-left corner and select **Autonomous Transaction Processing**.

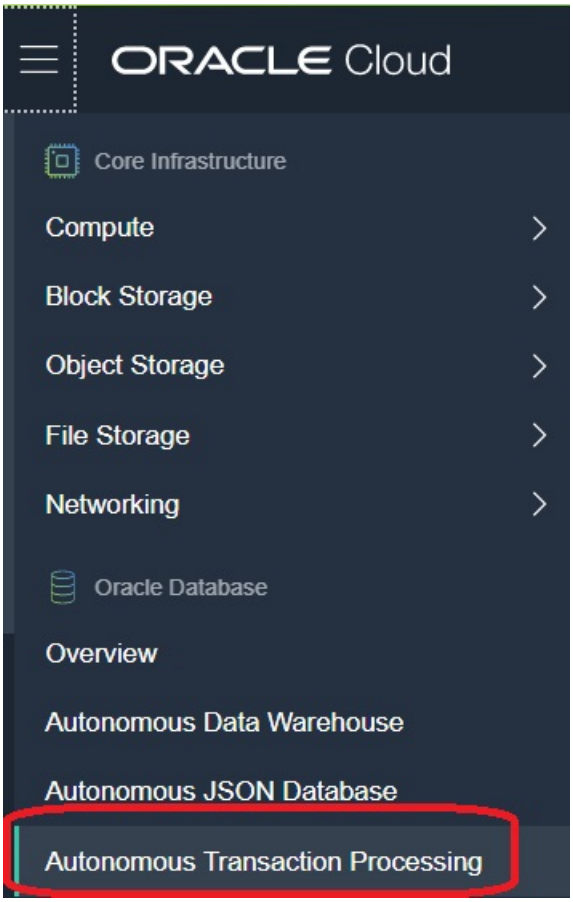

2. Click the instance of Autonomous Transaction Processing (ATP) that you previously created in the section of this learning path entitled "Creating an Autonomous Transaction Processing Instance on Dedicated Exadata Infrastructure in Oracle Cloud Infrastructure".

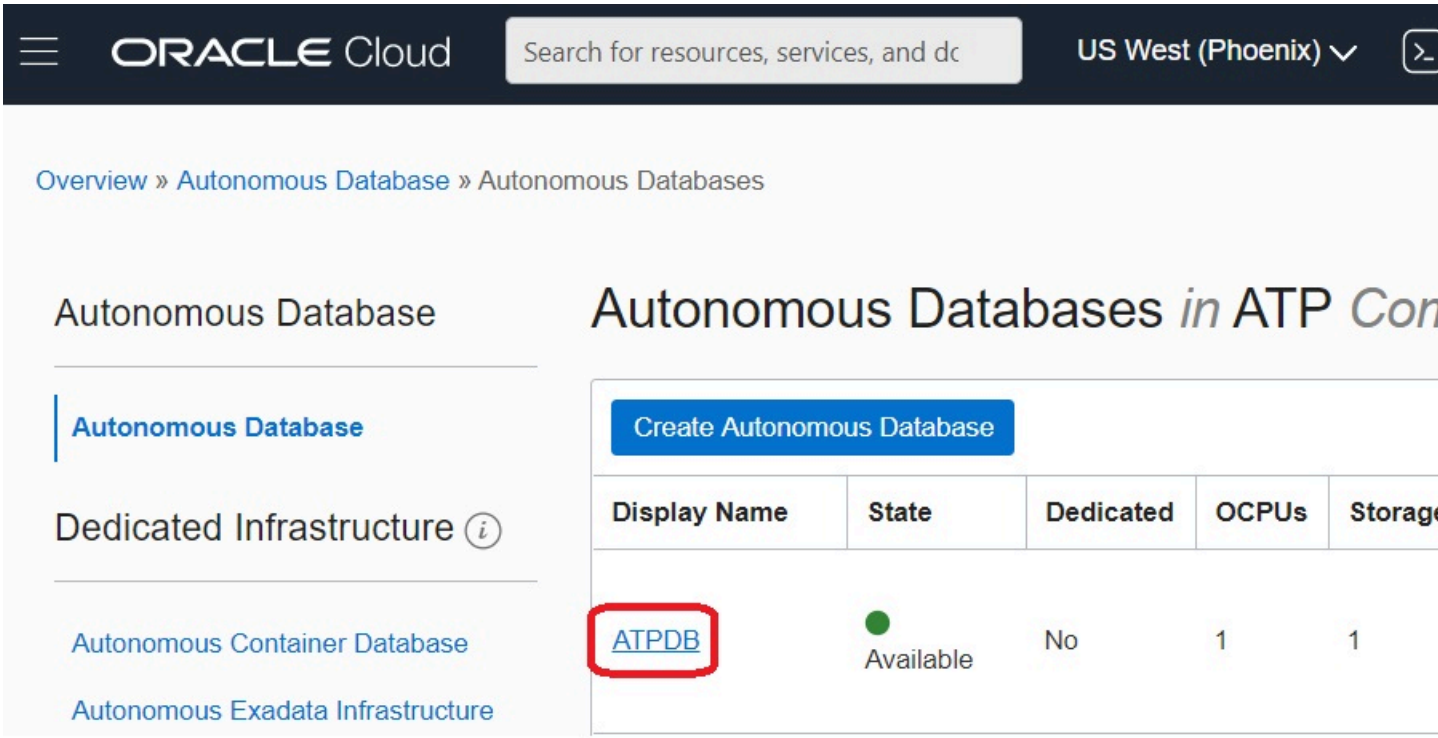

3. On the Autonomous Database Details section, click the **DB Connection** button.

> **Note:** The selected ATP instance must be in an **Available** state before you can download the database wallet. If the state is **Unavailable**, the **DB Connection** button is disabled.

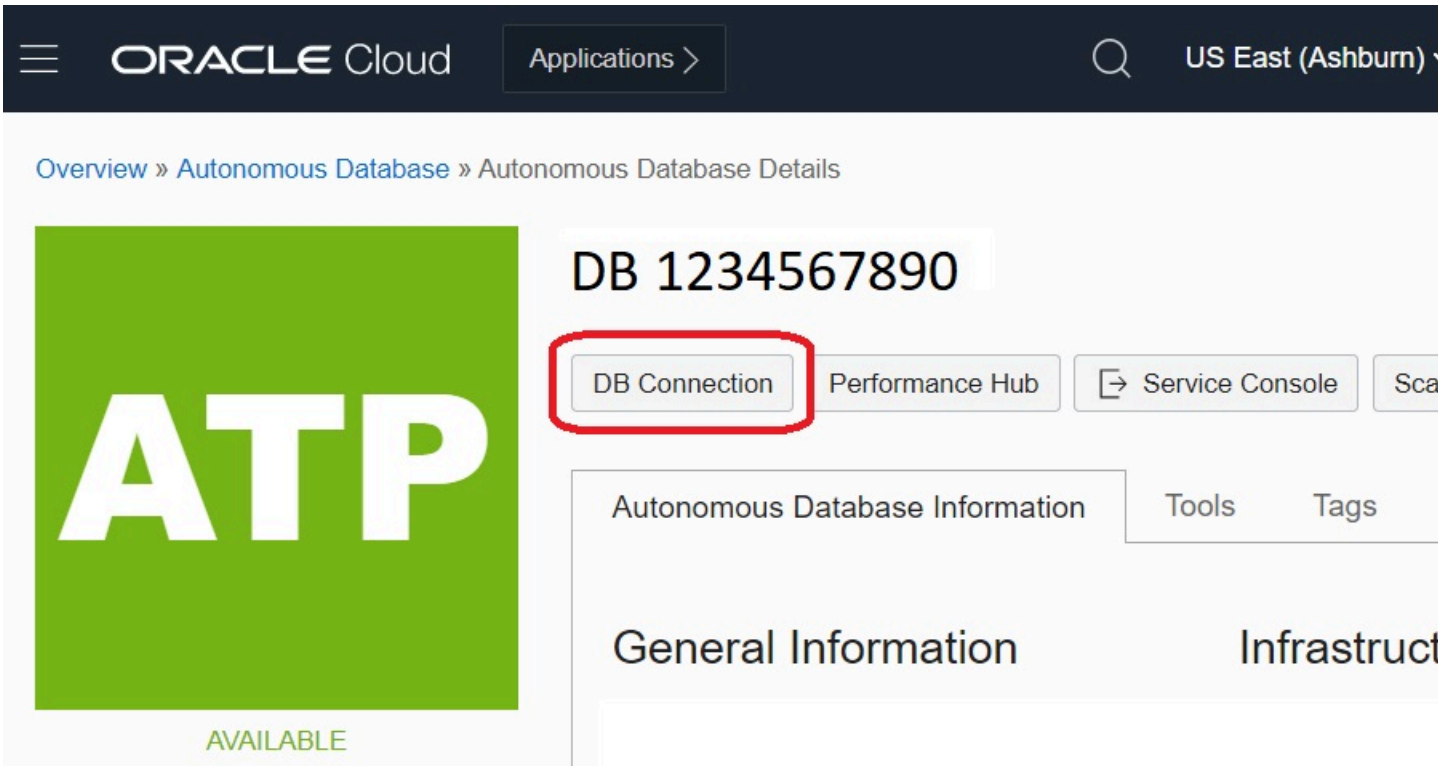

4. On the Database Connection dialog box, click the **Download Wallet** button.

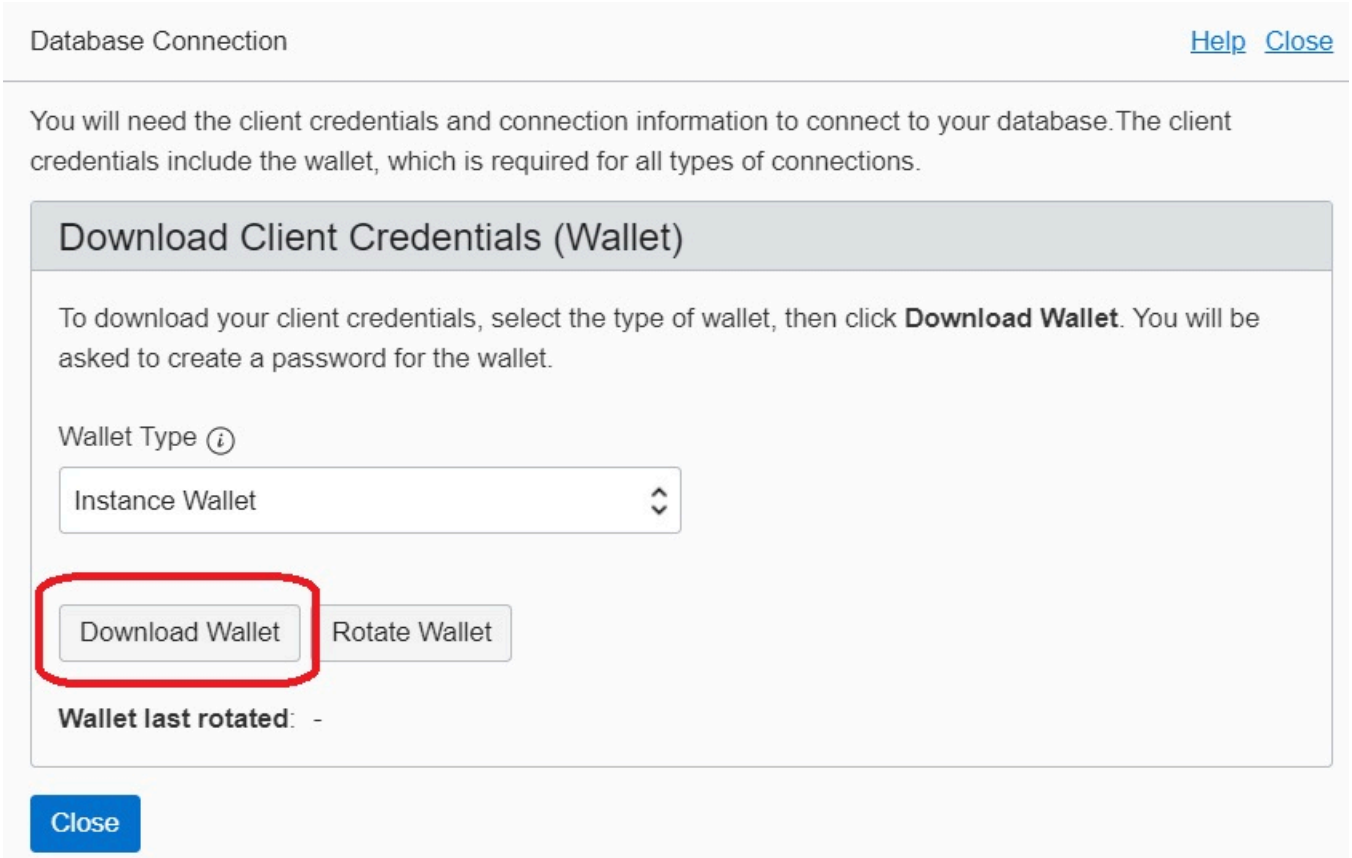

5. On the Download Wallet dialog box, provide a password to encrypt the keys inside the wallet. The password must have at least 8 characters and must include at least 1 letter and either 1 numeric character or 1 special character.

> **Note:** This password should be the same as the ADMIN user of ATP.

See *About Downloading Client Credentials* for information about the files included in the download. Take note of or copy the TNS names or connection strings you need for your connection.

See *About Connecting to Autonomous Databases* for information about making connections.

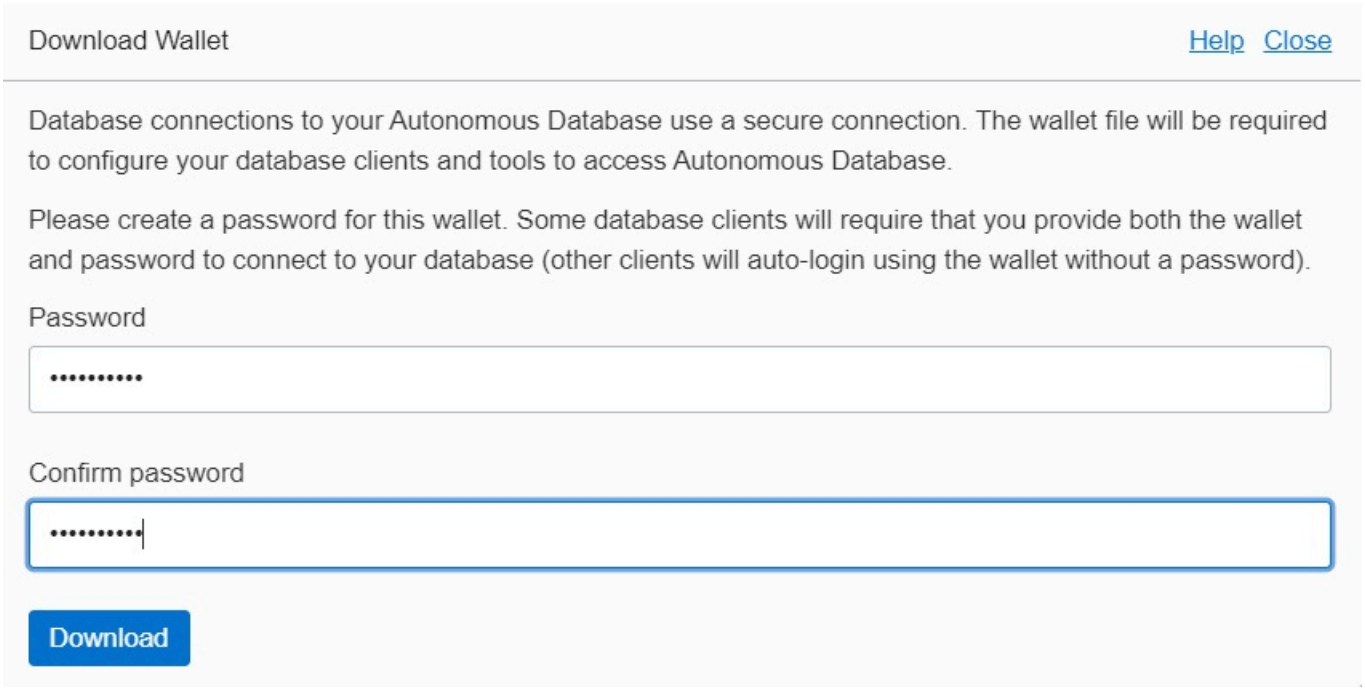

6. Click the **Download** button.
7. Save the downloaded compressed file containing client credentials to a secure location.

Creating an Object Storage Bucket for an Autonomous Database

This section describes how to create an Object Storage Bucket for use by an Autonomous Database.

In the Oracle Cloud Infrastructure Object Storage service, a bucket is a container for storing objects in a compartment within an Object Storage namespace. You must setup this function of the service for use by Oracle Autonomous Database using the instructions provided for the topic at this link:

Object Storage Buckets

6 Creating Linux Instances in Oracle Cloud Infrastructure

Creating the Linux Instance for the One-Click Provisioning Server

This section shows you how to create a Linux instance for the One-Click Provisioning Server for use by JD Edwards EnterpriseOne on Oracle Cloud Infrastructure. This instance includes the requisite Linux system.

Background

Oracle Cloud Infrastructure Compute lets you provision and manage compute hosts, known as instances.

> **Note:** While Oracle Cloud Infrastructure offers both Bare Metal and Virtual Machine instances, the current version of One-Click Provisioning for Oracle Cloud Infrastructure only supports Virtual Machine instances.

For additional information refer to *Creating an Instance* in the *Oracle Cloud Infrastructure Documentation*.

Creating a Linux Instance for the One-Click Provisioning Server

Use this procedure to create an instance for JD Edwards EnterpriseOne One-Click Provisioning Server.

1. On the **Oracle Cloud Infrastructure Console Home** page, click the navigation menu in the upper-left corner.
2. From the navigation menu, select **Marketplace > Applications**.
3. In the **Marketplace** page, under **All Applications**, using the Search to locate the listing for JD Edwards EnterpriseOne One-Click Provisioning Server Image.

   > **Tip:** To more easily locate this image, in the browser page, you can search for all or or part of the complete listing title, which is "JD Edwards EnterpriseOne One-Click Provisioning Server".

4. Click on the **Marketplace listing** for JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server Image.
5. On the Marketplace listing for JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server, in the **COMPARTMENT** field, use the drop down list to select the compartment that you previously created in the OBE of this Learning Path entitled ***Creating a Compartment*** .
6. Also, on the **Launch Instance page**, you must review the **Terms of Use** and click the check box indicating you accept the terms.
7. Click the **Launch Instance** button to continue.

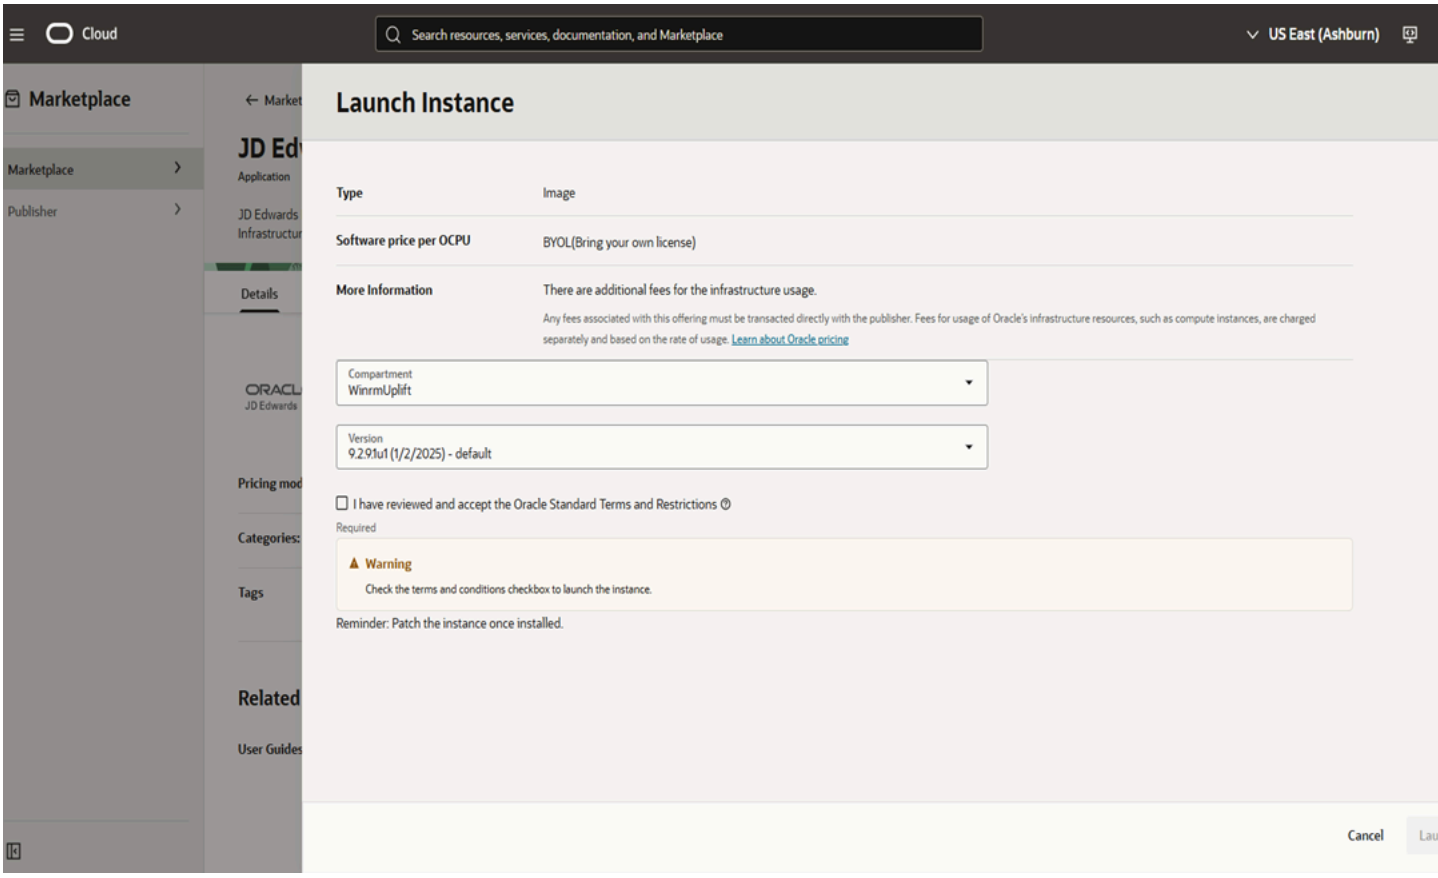

8. On the **Create Compute Instance** page, complete these fields:

   - ***Name***

     > **Note:** The system automatically populates this field. It is recommended that you change this name to a value that you might more easily recognize. You can change the name later. The name does not need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the instance.

     The name you enter here is the display name of the instance. This will be the host name of the JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server Image server.

     For example, the One-Click Provisioning Server might be called **jdeprov**.

     > **Note: Special Naming Restrictions**: Ensure that the host name of the JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server Image instance contains only alphanumeric values. You cannot use special characters such as the dollar sign ($), exclamation point (!), underscore (_), vertical pipe (|), at sign (@), and so on, in the host name.Machine names are limited to a maximum of 15 lower-case alphanumeric characters. This is the character count which is supported by the JD Edwards EnterpriseOne database table and application design. Do not use a fully qualified domain name for host names. You should only specify the first node of the domain name. If the existing hostname does not conform to these requirements, the runtime of JD Edwards EnterpriseOne will fail. If a host name does not conform to these requirements, you should permanently change the host name for your system.

   - ***Create in Compartment***

     Enter the name of the compartment in which you want to create the instance. By default, the currently selected Compartment is displayed.

   - ***Configure placement and hardware***

     Click on an availability domain (AD) to select it. You can only specify a single AD.

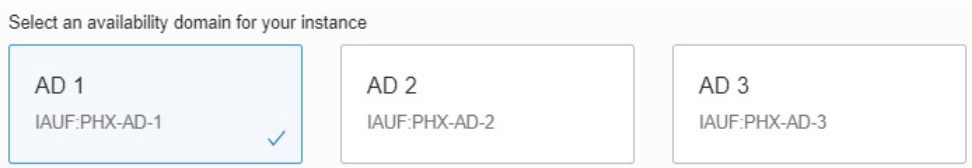

- ***Image***

  Because you already selected the image from Marketplace, the following image is automatically shown here. You should accept this value.

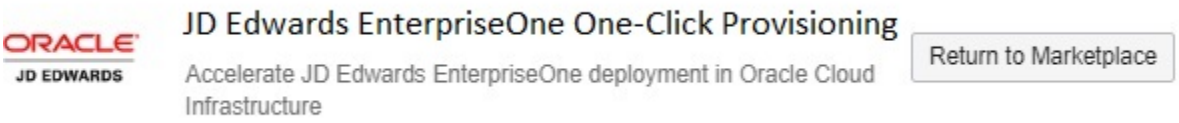

- ***Shape***

  Click the **Change Shape** button.

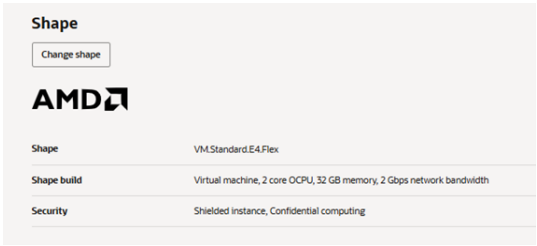

  On Browse All Shapes, make the following selections:

  - *Instance type*

    Virtual Machine
  - *Shape series*

    Supported shapes are those which meet the OCPU and memory requirements listed in the OBE of this Learning Path entitled: *Minimum Resource Requirements*

  Click the **Select Shape** button to save your selections.

- ***Configure networking***

  Verify, and if necessary, set your network settings for the instance of JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server Image. These settings include:

  - Network
  - Subnet
  - Subnet in *<Compartment>*

    > **Note:** Be sure to select this radio button: Assign a public IPV4 Address.

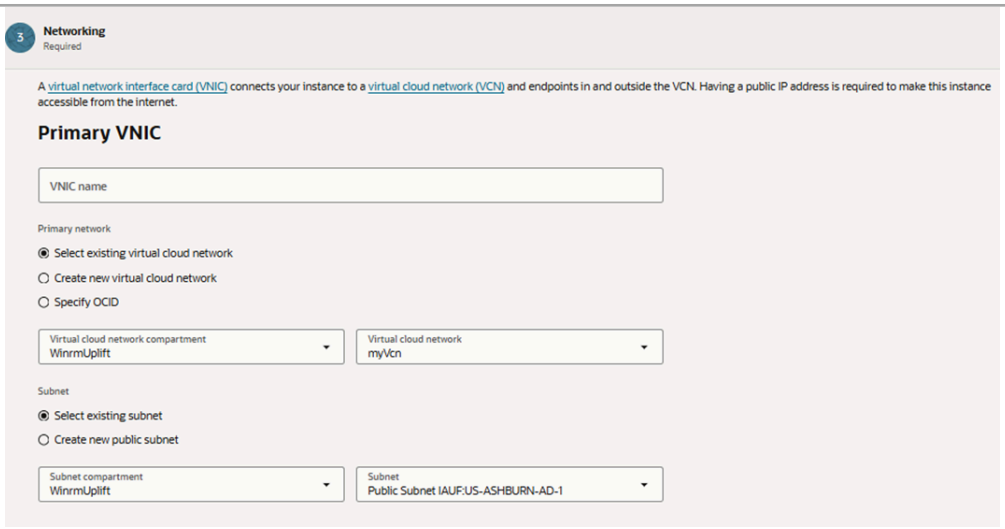

- ***Advanced Options***

  Expand the Advanced Options section and enter the Hostname. This value also appears as the Instance Name under the Basic Information section.

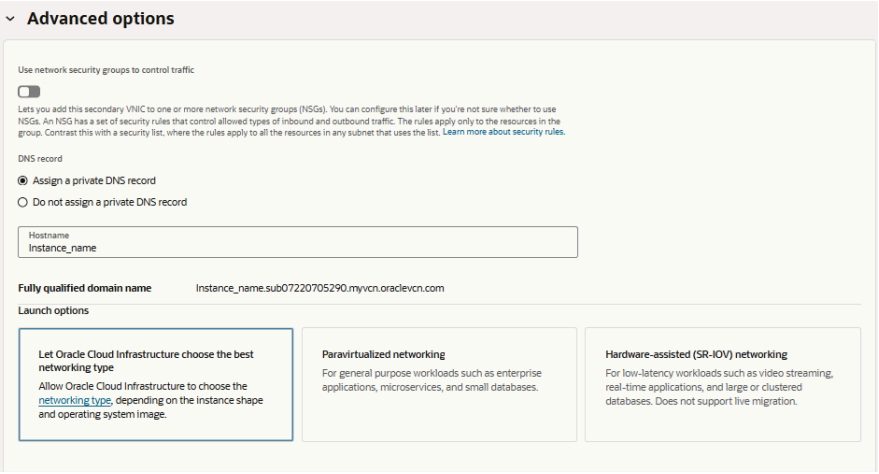

- ***Add SSH Keys***

  You must add the SSH keys that you previously created for use with JD Edwards EnterpriseOne JD One-Click Provisioning Server Image.

  You can either click the **Choose SSH key file** option to select a file with your public (.pub) key from your computer, or choose the **Paste SSH keys** option to paste the key.

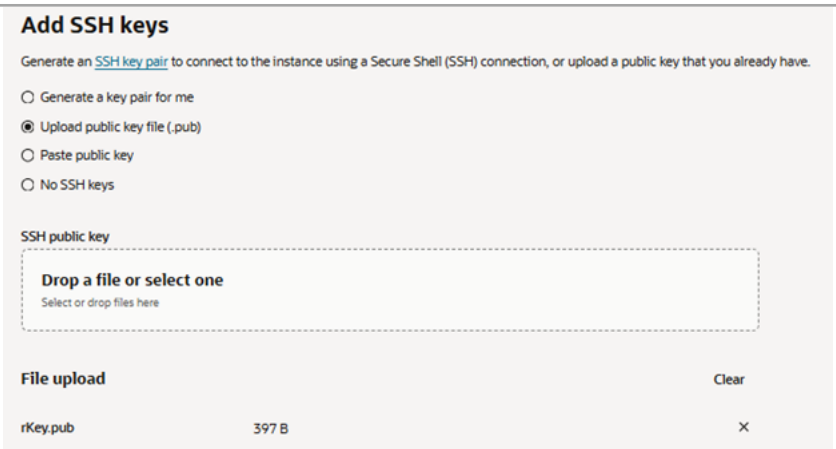

- ***Configure boot volume***

  One-Click with a Compute or DB Systems (DBS) Database

  You must specify a value of 100 GB. This size is required to support the install of the Provisioning Server with Pristine data (50 GB) with the remainder for use by a customer's normal data set. Further, you must enable this extended boot volume size as described below in the task Enabling the Custom Boot Volume Size.

  One-Click with an Autonomous Database (ADB)

  You must specify a value of 150 GB. This size is required to support the install of the Provisioning Server with Pristine data (50 GB) with the remainder for use by a customer's large data set. Further, you must enable this extended boot volume size as described below in the task Enabling the Custom Boot Volume Size.

  Optionally you can choose to enable encryption using the checkbox for **Use In-Transit Encryption**.
- ***Show Advanced Options***

  You do not need to set any Advanced Options.

9. Click the **Create** button to create the Oracle Cloud Infrastructure instance for JD Edwards EnterpriseOne JD Edwards EnterpriseOne One-Click Provisioning Server Image.

> **Note:** After the instance is created, the system assigns a **Public IP address**. Make a note of this address because you will need it to connect to the instance using the SSH keys in order to complete the first-time configuration of the One-Click Provisioning Server Image for JD Edwards EnterpriseOne as described in the following section.

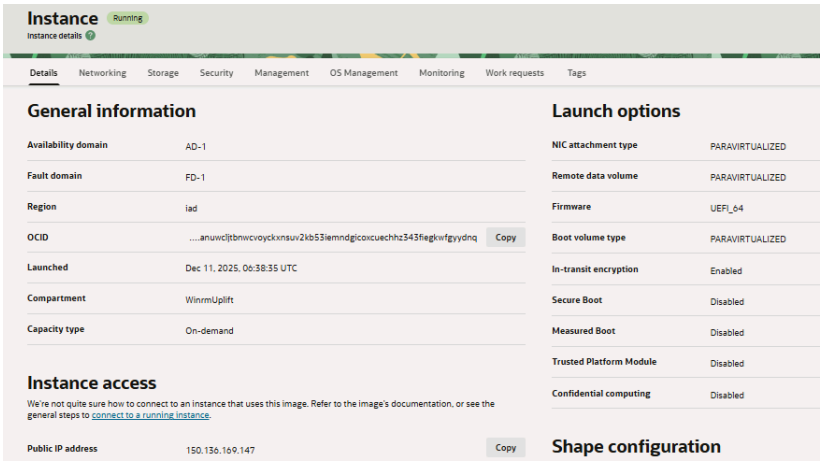

Enabling the Custom Boot Volume Size

If you followed the recommendation in this Learning Path and specified a Custom Boot Volume Size when you created the Linux instance for the Provisioning Server, you must run the following command (as root) to complete the resize functionality.

```
sudo /usr/libexec/oci-growfs -y
```

> **Note:** It is good practice to ensure that the latest utilities are installed on your operating system image.

> **Note:** The utility oci-growfs expands the root filesystem of the instance to its configured size. This command must be run as root. For additional information, refer to this Oracle documentation: *https://docs.oracle.com/en-us/iaas/Content/Compute/References/oci-growfs.htm*

Setting Up the Provisioning Server

> **Note:** Before executing any scripts on the Provisioning Server, you must run a mandatory patch as described in the MOS document: E1: OCI: Additional One-Click server deployment is failing on existing TLS-enabled SMC (Doc ID 3010098.1).

Use this procedure to set up the One-Click Provisioning Server.

1. Connect to Provisioning Server Instance with public IP address and private key using Putty.

2. Run the following command:

```
sudo ./prepareProvisionPrime.sh
```

> **Note:** You can safely ignore access or permission errors in the console output during script execution. Such errors do not affect the successful execution of the script.

> **Note:** The Provisioning Server will reboot after running prepareProvisionPrime.sh script.

3. Upon completion of the reboot, run the following command:

```
sudo ./setupProvisionPrime.sh
```

> **Note:** Do not reboot the Provisioning Server after running above setup command until you change password for Server Manager Console and WebLogic Server.

Creating Linux Instances as VMs in Oracle Cloud Infrastructure

This section shows you how to create Linux instances as virtual machines (VMs) in the Compute Service of Oracle Cloud Infrastructure.

Background

Oracle Cloud Infrastructure Compute enables you to provision and manage compute hosts, known as instances.

> **Note:** While Oracle Cloud Infrastructure offers both Bare Metal and Virtual Machine instances, the current version of One-Click Provisioning for Oracle Cloud Infrastructure only supports Virtual Machine instances.

For additional information refer to *Creating an Instance* in the *Oracle Cloud Infrastructure Documentation*.

In preparation for JD Edwards EnterpriseOne One-Click Provisioning, you must create Linux VM instances in Oracle Cloud Infrastructure Compute for each of these core JD Edwards EnterpriseOne servers:

- **Oracle Database Server**

  > **Note:** A Linux VM instance is required for this server if you are installing your own Oracle database (Bring Your Own License - BYOL) in Oracle Cloud Infrastructure Compute. Otherwise, you can use Oracle Database Services.

- **Enterprise Server**
- **WebLogic Server**

> **Note:** The image in the Marketplace for Oracle Cloud Infrastructure for the JD Edwards EnterpriseOne One-Click Provisioning Server includes the requisite Linux image.

Note: For best performance, Oracle strongly recommends creating VMs for all JD Edwards components within the same Availability Domain (AD).

Prerequisite

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Get to Know the Console*.
- You should have a fundamental understanding of Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription to Oracle Cloud Infrastructure and an Administrator account in the platform. For more information, refer to *Getting Started with Oracle Cloud*.
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

Creating Linux Instances as VMs

With these exceptions, the procedure for creating a Linux instance as a VM is specified in the section: *Creating the Linux Instance for the One-Click Provisioning Server:*

- Supported shapes are those which meet the OCPU and memory requirements listed in the section of this Learning Path entitled: *Minimum Resource Requirements*.
- Instead of having the requisite Operating System version automatically selected when you are installing from a Marketplace image for the Provisioning Server, you must manually select a specific image that is certified with JD Edwards EnterpriseOne. Refer to the *Supported Software Versions* section of this Learning Path for the supported release of Oracle Linux.

1. On Create Compute Instance, in the **Image** section, click the **Change Image** button.

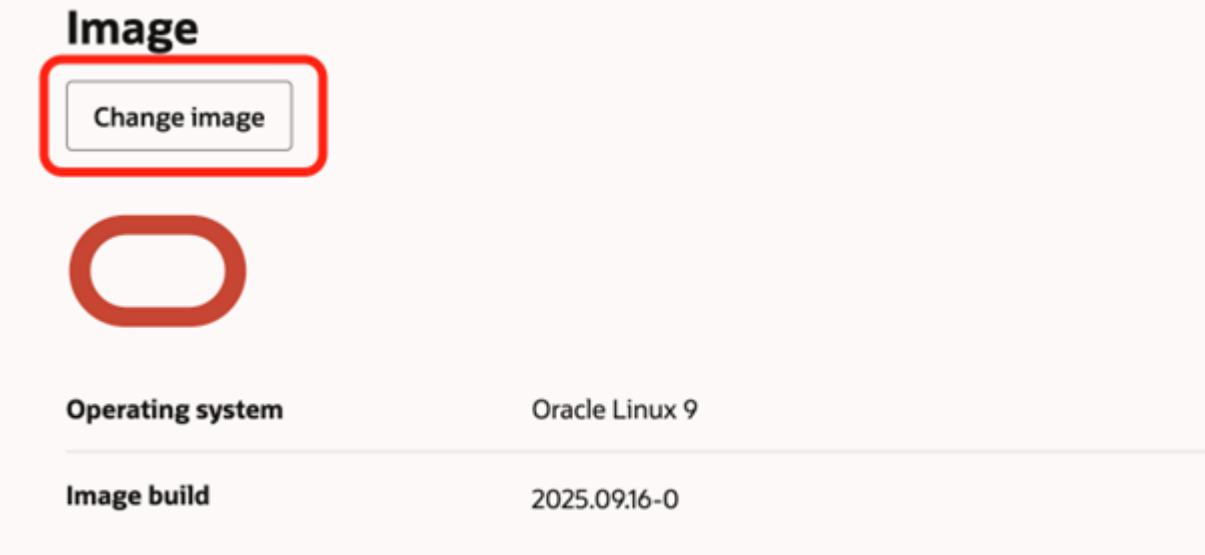

2. On Browse All Images:
   - In the **Image name** column, click the **Oracle 9 Linux** checkbox.
   - In the **Image build** column, use the pull-down to select the latest build.

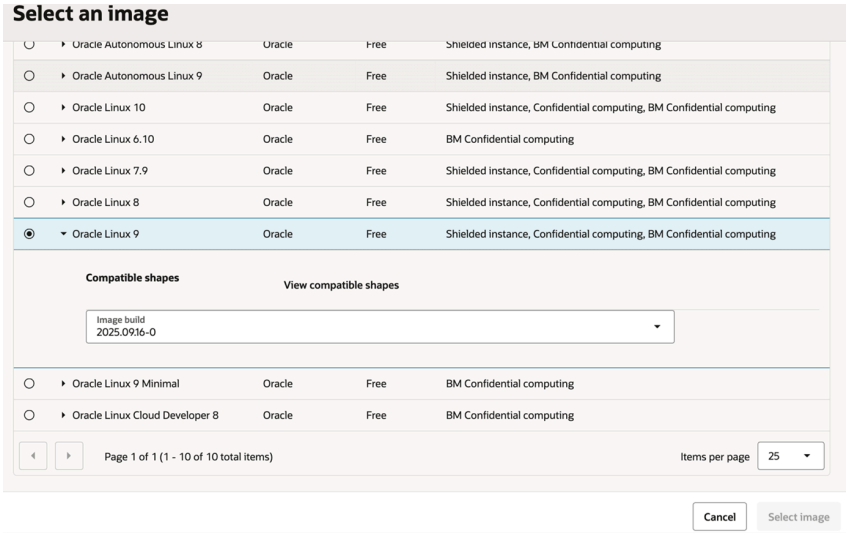

7 Creating Volume Storage

Using Volume Storage

This section provides an overview to using volume storage in Oracle Cloud Infrastructure for use by JD Edwards EnterpriseOne with One-Click Provisioning.

To understand your volume storage options in Oracle Cloud Infrastructure, refer to this link: *Overview of Block Volume*.

Using Storage: Block Volume and Boot Volume

Oracle Cloud Infrastructure supports two types of volumes:

- **Block Volumes**

  This storage is used to expand the storage capacity of Compute instances, to provide durable and persistent data storage that can be migrated across Compute instances, and to host large databases in Compute. To use this volume type you must create, attach, and connect the volume.

- **Boot Volumes**

  Instances that are created on the Oracle Cloud Infrastructure are allocated a fixed default amount of storage, called the Boot Volume. This default storage amount must be expanded to meet the requirements of JD Edwards Enterprise servers. These storage amount requirements are specified in the section of this Learning Path entitled *Minimum Resource Requirements* .

Selecting between boot volume or block volume storage is the responsibility of your system architect. For JD Edwards EnterpriseOne, at a minimum, you must create sufficient storage for these machines:

- Provisioning Server
- Database Server
- Enterprise Server
- WebLogic Server
- Deployment Server
- Development Client

Block Volume Storage

If you choose to use block volume storage, you must complete the following steps:

- *Creating a Block Volume*
- *Attaching a Volume*
- *Connecting a Volume*

Note: This section contains instructions for both Linux and Microsoft Windows environments.

Boot Volume Storage

If you choose to use boot volume storage, refer to the steps in the following document:

- *Boot Volumes*

8 Configuring the Linux Servers

Performing Common Setup for Linux Servers

This section describes the setup tasks that must be performed on each Linux-based server that will be deployed by JD Edwards EnterpriseOne One-Click Provisioning, excluding the One-Click Provisioning Server, which does not require manual setup. The remaining Linux servers used by JD Edwards EnterpriseOne may include a Compute Oracle Database Server. If you are using DB Systems instead of a Compute Oracle Database Server, you must perform the setup described in the Learning Path Deploying *JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with DBaaS (DB Systems)*.

Prerequisite

You must have created a Linux VM for each server that will be deployed by the JD Edwards EnterpriseOne One-Click Provisioning Server for the Oracle Cloud Infrastructure. At a minimum, this includes the Compute Database Server (except if you are using DBaaS - DB Systems), Enterprise Server, and WebLogic Server. Refer to the section "*Creating Linux Instances as VMs on the Oracle Cloud Infrastructure*" in this Learning Path.

General

Use the following general steps to confirm the completion of the common setup that is required for each VM instance that you created in Compute to support JD Edwards EnterpriseOne on the Oracle Cloud Infrastructure.

1. Ensure that the `/u01` mount point is created and mounted. This should have been done when you created Block Volume storage. Use this command to verify that the volume is mounted on `/u01`:

   ```
   df -h
   ```

   The returned results should be similar to this:

   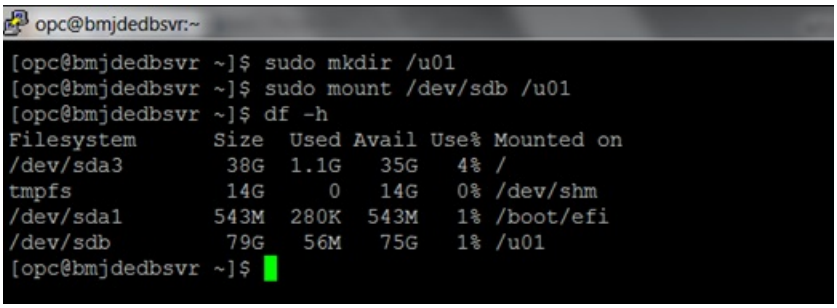

2. Ensure that the proper permissions are set on the /u01 mount using this command:

   ```
   sudo chmod 770 /u01
   ```

3. Ensure that the host name for the server contains only alphanumeric values. Special characters are not allowed, with the exception of a hyphen "_", which is allowed. Special characters are not allowed except a hyphen (-). Refer to **Special Naming Restrictions** in the section "*Creating Linux Instances as VMs in Oracle Cloud Infrastructure*" in this Learning Path.
4. Edit the `sshd_config` file using this command:

```
sudo vi /etc/ssh/sshd_config
```

5. Verify the following settings in the `/etc/ssh/sshd_config` file:
   a. Contains this setting:

      ```
      ClientAliveInterval 3600
      ```

      Ensure that the setting exists, is set properly to 3600, and is not commented out.
   b. Has the SSH connection over IPv6 **disabled** by having this setting:

      ```
      AddressFamily inet
      ```

      Ensure that the setting exists, is set properly to inet (not any), and is not commented out.
6. Edit the sudoers file using this command:

   ```
   sudo vi /etc/sudoers
   ```
7. Verify that the `/etc/sudoers` file either does not contain this setting or this setting is disabled in the file:

   ```
   Defaults requiretty
   ```
8. You must disable IPv6 using the following command:

   ```
   sudo vi /etc/sysctl.conf
   ```

   Ensure that the `/etc/sysctl.conf` file has these settings to disable IPv6 protocol:

   ```
   net.ipv6.conf.default.disable_ipv6=1

   net.ipv6.conf.all.disable_ipv6=1
   ```
9. Optionally, you can choose to edit the bash.profile to set the user default language for output to prevent unreadable characters from being inserted into messages and logs.

   ```
   sudo vi /root/.bash_profile
   ```

   Add this setting:

   ```
   LANG="C"
   ```

Enable Inbound Ports in the Firewall for Compute Instances

You must enable inbound ports in the firewall service for the Compute instance of each Linux server to enable the functionality of the Provisioning Server to provision each Linux server.

1. For each JD Edwards EnterpriseOne server, repeat this command as necessary to specify open ports in the firewall:

   ```
   sudo firewall-cmd --zone=public --add-port=<PORT>/tcp --permanent
   ```

   where *<PORT>* is the port number that must be open. These port numbers are listed in the following table.

Linux Server Firewall Port List	
Component	**Inbound Ports to Open**

Oracle Database	22 5150 *<DB_PORT>* 14502-14510
Enterprise Server	22 5150 6017-6022 14502-14510
Web Server	22 5150 *<WLS_ADMIN_PORT>* **See Note: 1** *<SSL_ACCESS_PORT>* **See Note 2** *<SSL_ACCESS_PORT-1>* **See Note 3** 14502-14520

> **Note: 1** This is the Admin Port on which the Admin Server is running. This value is set by the user while creating a WebLogic domain.

> **Note: 2** This SSL port must be opened to enable the Server Manager Console to complete the tasks.

> **Note: 3** This port number is equal to the value of the SSL port minus one. This port must be opened to enable the Server Manager to complete the provisioning tasks. For example, if you have specified the port value as 8081 in the One-Click Provisioning Console, you must open 8080 port (8081-1). Also, you must open a port for each of the WebLogic Servers on Oracle Cloud Infrastructure. For example, if you have the following environments and server combinations, you should open eight ports: four ports for SSL for the port numbers specified by using the Provisioning Console, and four ports for non-SSL. The value of the ports for non-SSL should be the value of the SSL port minus one.
>
> - DV HTML
> - PD HTML
> - DV AIS
> - PD AIS

2. After all the ports are opened in the firewall for each server, use these commands to reload the firewall:

```
sudo systemctl stop firewalld

sudo systemctl start firewalld
```

Disable SELINUX and Reboot Machines

After you have configured the firewall, use this procedure to disable SELINUX and reboot the machine for the changes to take effect.

1. Use this command to check the status of Security Enhanced Linux (SELINUX):

   ```
   sudo getenforce
   ```

   If the returned status is either **Enforcing** or **Permissive**, you must temporarily disable the extra security restriction provided by SELINUX by modifying this file:

   ```
   /etc/selinux/config
   ```

   Edit the `/etc/selinux/config` file using this command:

   ```
   sudo vi /etc/selinux/config
   ```

   Ensure that `SELINUX=` is set to this value:

   ```
   SELINUX=disabled
   ```

2. When you change any security settings, for example if you disable SELINUX as described in the preceding step, you must reboot the machine using this command:

```
sudo shutdown -r now
```

3. If you reboot to enable the security settings, after the reboot is complete, run sudo getenforce to confirm that the status of SELINUX is **Disabled**.

> **Note:** As a recommended best practice, after you have successfully completed JD Edwards EnterpriseOne One-Click Provisioning, you should reactivate SELINUX by changing the setting to either of the following: **SELINUX=Enforcing SELINUX=Permissive**

Set Up Users and Groups

By default, the Oracle Cloud Infrastructure has a predefined **opc** user and **opc** group. Perform these tasks to set up additional users and groups on each Linux server:

1. Create the **dba** and **oracle** groups using these commands:

```
sudo groupadd -g nnnn dba

sudo groupadd -g nnnn oracle
```

   where nnnn is the ID value with which the group will be created. You must use a free value that ideally should be above 1000. You can determine used group IDs with the id command.

2. Create the **oracle** user and add it to the to the **oracle** group using this command:

```
sudo useradd -d /home/oracle -m -s /bin/bash oracle
```

   > **Note:** This step is not required for DB Systems because this database is already set up by default.

3. Add the **opc** user to the **oracle** group and the **oracle** user to the **dba** group and the **oracle** group using these commands:

```
sudo usermod -a -G oracle opc

sudo usermod -a -G oracle oracle

sudo usermod -a -G dba oracle
```

4. Set `/u01` as owned by the **oracle** group using this command:

```
sudo chgrp oracle /u01
```

Install Requisite YUM Packages

You must install a specific set of packages from the YUM repository onto each Linux-based server (excluding the One-Click Provisioning Server) in the JD Edwards EnterpriseOne environment. The additional servers include:

- ***Compute Oracle Database Server***
- ***Enterprise Server***
- ***WebLogic Server***

> **Note:** Ensure that you have enabled your YUM repository. For Linux instances that are provisioned in Oracle Cloud Infrastructure, this is generally set to enabled by default. This functionality is required to install YUM packages that are prerequisites to set up each JD Edwards EnterpriseOne server. If you had not pre-enabled your YUM repository before you ran One-Click to set up the JD Edwards EnterpriseOne servers, then you must enable it and manually install and/or update all required YUM packages as listed below.

The following lists are the required packages sorted alphabetically. You can create a script to install the packages one at a time, or all at once. In either case, be sure to confirm that when each command executes, the command returns either of these results:

```
Complete
```

or

```
Nothing to do
```

> **Note:** If you choose to run all the YUM packages at once, using a single command, you may encounter errors. For example, a package may not install as expected due to dependency or other technical issues with the package or library. Further, this behavior is unpredictable and can vary according to system conditions. The resolution of such errors is the responsibility of your system administrator and is beyond the scope of this document. If you need help for resolving YUM issues, your Linux Admin should open a service request with the Oracle Linux support team. Additionally, various Linux OS images may be delivered with some of these packages already installed by default. You can use the complete list shown below for each server to cross-check which packages may already be installed and those that you will need to manually install.

> **Note:** It is recommended that you run YUM install commands as the root user. For example: `sudo su  yum install -y bind-utils`

Compute Oracle Database Server

Ensure the VM instance for the Compute Oracle Database Server contains these YUM packages:

> **Note:** For users who have set up servers using a previous version of Linux, the package names below with strikethrough are not required for OL9.6, while the names in **bold** are new requirements.

```
bind-utils

bc.x86_64

binutils.x86_64

~~compat-libcap1.x86_64~~

~~compat-libstdc++-33.x86_64~~

compat-openssl10.x86_64

gcc.x86_64

elfutils-libelf.x86_64
```

```
fontconfig.x86_64

gcc-c++.x86_64

glibc.i686

glibc.x86_64

glibc-devel.x86_64

ksh.x86_64

libaio.x86_64

libaio-devel.x86_64

libasan.x86_64

liblsan.x86_64

libgcc.x86_64

libnsl.x86_64

libstdc++.x86_64

libstdc++-devel.x86_64

libX11.x86_64

libXau.x86_64

libxcb.x86_64

libXext.x86_64

libXi.x86_64

libXtst.x86_64

libvirt-libs.x86_64

libXrender.x86_64

libxcrypt-compat.x86_64

libibverbs.x86_64

librdmacm.x86_64

make.x86_64

policycoreutils.x86_64

policycoreutils-python-utils.x86_64

sysstat.x86_64

smartmontools.x86_64

unzip.x86_64
```

~~zip.x86_64~~

Enterprise Server

> **Note:** Usually, the order in which you run YUM updates is not a factor. However, there is a known dependency on the installation order for certain packages, therefore you should install packages in the order shown below. If other dependencies arise, they should be resolved by your Linux administrator.

Ensure that the VM instance for the Enterprise Server contains these YUM packages:

> **Note:** For users who have set up servers using a previous version of Linux, the package names below with strikethrough are not required for OL9.6, while names in **bold** are new requirements.

binutils.x86_64

~~**compat-libcap1.x86_64**~~

~~**compat-libstdc++-33.i686**~~

~~**compat-libstdc++-33.x86_64**~~

compat-openssl11.x86_64

elfutils-libelf-devel.x86_64

file

gcc.x86_64

gcc-c++.x86_64

~~**glibc.i686**~~

glibc.x86_64

~~**glibc-devel.i686**~~

glibc-devel.x86_64

ksh.x86_64

~~**libaio.i686**~~

libaio.x86_64

~~**libaio-devel.i686**~~

libaio-devel.x86_64

libasan.x86_64

liblsan.x86_64

libnsl2.x86_64

libnsl2-devel.x86_64

~~libgcc.i686~~

libgcc.x86_64

~~libnsl.x86_64~~

~~libstdc++.i686~~

libstdc++.x86_64

libstdc++-devel.x86_64

~~libX11.i686~~

libX11.x86_64

~~libXau.i686~~

libXau.x86_64

~~libxcb.i686~~

libxcb.x86_64

~~libXext.i686~~

libXext.x86_64

~~libXi.i686~~

libXi.x86_64

~~libXtst.i686~~

libXtst.x86_64

make.x86_64

~~ncompress~~

oracle-database-preinstall-19c.x86_64

sysstat.x86_64

unixODBC.x86_64

~~unixODBC-devel.x86_64~~

unzip.x86_64

~~zlib.i686~~

zlib.x86_64

zip.x86_64

~~zlib-devel.i686~~

zlib-devel.x86_64

nss-softokn-freebl.x86_64

~~nss-softokn-freebl.i686~~

WebLogic Server

Ensure that the VM instance for the WebLogic Server contains these YUM packages:

> **Note:** For users who have set up servers using a previous version of WebLogic Server, the package names below with strikethrough are not required for the currently supported version, while two (2) names in red are new requirements.

bind-utils

~~glibc.i686~~

glibc.x86_64

glibc-devel.x86_64

ksh.x86_64

~~libnsl.x86_64~~

net-tools

unzip.x86_64

zip.x86_64

zlib-devel.x86_64

Edit the resolv.conf File to Specify the DNS Domain Name

You must ensure that the search setting in the `/etc/resolv.conf` file specifies the DNS domain name of the Availability Domain to which all JD Edwards EnterpriseOne servers belong. While logged in as the root user (the owner of the resolv.conf file), you must add a line with the following syntax to the `resolv.conf` file:

```
$ vi /etc/resolv.conf

search <DNS_Domain_Name> <subnet1> <subnet2> <subnet3>.<DNS_Domain_Name>
```

Generally each Availability Domain contains at least three subnets. You must include an exact search string for each of those subnets.

For example, your subnets can look like this:

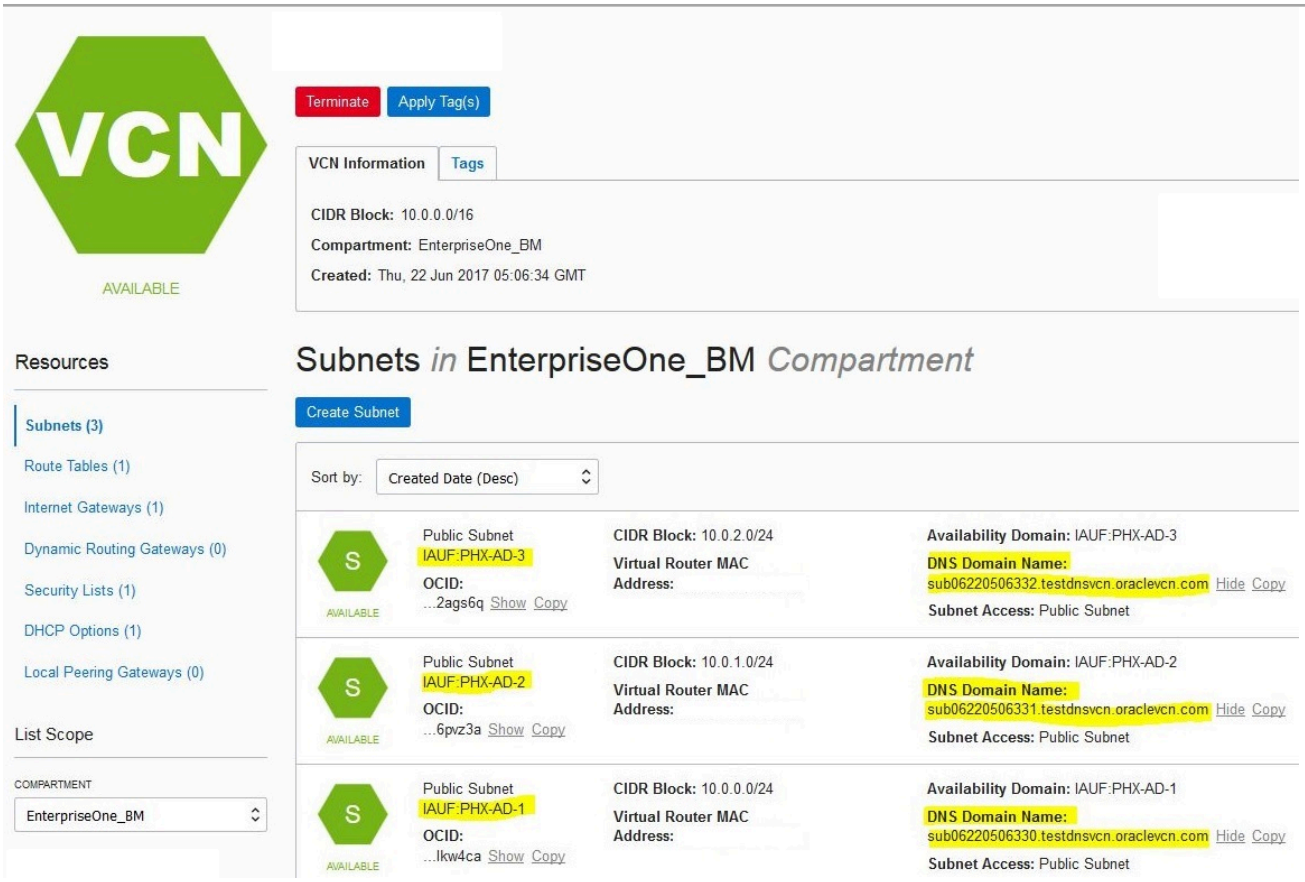

Using the above example, assuming that your JD Edwards EnterpriseOne Servers are running in the Phoenix Availability Domain, you would edit your `/etc/resolve.conf` file to add a line with a search setting similar to the setting provided below to include all the three domains:

```
search testdnsvcn.oraclevcn.com sub06220506330.testdnsvcn.oraclevcn.com sub06220506331.testdnsvcn.oraclevcn.com sub06220506332.testdnsvcn.oraclevcn.com
```

Change the Maximum Transmission Unit (MTU) Setting

The recommended MTU setting is 1500. Use this procedure to check, and if necessary change, the current MTU setting.

1. Run this command to check current value of MTU:

   ```
   Ifconfig
   ```

2. If the MTU is not set to 1500, then run below commands to set MTU value:

   ```
   sudo sed -i -e "\$a ifconfig ens3 mtu 1500 up" /etc/rc.d/rc.local

   sudo sed -i -e "\$a ifconfig ens3 mtu 1500 up" /etc/rc.local

   sudo chmod +x /etc/rc.d/rc.local
   ```

3. Reboot the machine to enable this change.

Setting Up the Provisioning Server for an Oracle Autonomous Database

This section describes the prerequisite tasks and the script that must be run to set up the Provisioning Server for an Oracle Autonomous Database.

Prerequisites

- You must have deployed the Provisioning Server as described in the section "*Creating the Linux Instance for the One-Click Provisioning Server*" of this Learning Path.
- You must have allocated storage volume to the Linux VM for the Provisioning Server.
- You must have performed the tasks described in the section "*Performing Common Setup for All Linux Servers*" of this Learning Path.

General

Use the following general steps to set up the Provisioning Server for an Oracle Autonomous Database.

1. Ensure that you have performed all the tasks described in the tutorial "Performing Common Setup for All Linux Servers" of this Learning Path, including the steps to create groups and users as well as assigning proper permissions to `/u01`.
2. Ensure that you have created Block Volume storage with **150** GB for use by the Provisioning Server. This space is required:
    - To install an Oracle database client for connection to the Oracle Autonomous Database.
    - To extract the database package so that the data can be transferred to the Oracle Autonomous Database.

    You should have already created and attached the requisite Block Storage if you followed the instructions in the module of this Learning Path entitled: ***Creating Storage for Linux Instances*** .

    Use this command to view your available storage space on `/u01`:

    ```
    df -h
    ```

    You should have allocated 150 GB on `/u01`.

    > **Note:** If the boot volume size is not 150 GB, you can either redo the instructions to create the instance with a boot volume size of 150 GB, or optionally you can resize the boot volume by following the instructions in "*Resizing a Volume*" in *Oracle Cloud Infrastructure Documentation*.

3. Run this script to install the necessary prerequisites to the Oracle database client and its required package to enable the Provisioning Server to connect to ADB:

    ```
    sudo ./setupPrAutoDB.sh
    ```

Setting Up the Enterprise Server

This section shows you how to set up the Enterprise Server for JD Edwards EnterpriseOne on the Oracle Cloud Infrastructure.

Prerequisite

You must have created a Linux VM for the Enterprise Server as described in the module of this Learning Path entitled: *Creating Linux Instances as VMs on the Oracle Cloud Infrastructure*.

The Linux VM for the Enterprise Server must be allocated with storage volume space, the amount of which is specified in the module of this Learning Path Entitled: **Before You Begin**.

You must have performed the functions described in the module of this Learning Path entitled: **All Servers Common Setup**.

Setting Up the Enterprise Server

1. Ensure you have performed all the tasks described in the module of this Learning Path entitled: *Common Setup for Linux Servers* including the steps to create groups and users as well as assigning proper permissions to `/u01`.
2. Use this command to ensure that the `/u01` mount point has proper permissions to allow the One-Click Provisioning process to subsequently create requisite folders:

   ```
   sudo chmod 770 /u01
   ```
3. Create the **oneworld** group using this command:

   ```
   sudo groupadd -g nnnn oneworld
   ```

   where *nnnn* is the ID value with which the group will be created. You must use a free value which ideally should be above 1000. You can determine used group IDs with the id command.
4. Create the **jde920** user and make the KornShell the default shell (using the -s switch with /bin/ksh), using this command:

   ```
   sudo useradd -d /home/jde920 -m -s /bin/ksh jde920
   ```

   > **Note:** This specific user, with access to ksh, is required in order to start and stop the Enterprise Server.

5. Add the **jde920** user to the **oracle** and **oneworld** group using these commands:

   ```
   sudo usermod -a -G oracle jde920

   sudo usermod -a -G oneworld jde920
   ```
6. Ensure that there is sufficient available storage for the Enterprise Server, YUM packages, and the archive (compressed and extracted). For Production environments, this space should be allocated as a Block Volume. You should have already created and attached the requisite Block Storage if you followed the instructions in the

modules of this Learning Path entitled: **Create Additional Block Volumes for Linux Instances** and **Attach the Block Volumes for Linux Instances**.

Use this command to view your available storage space on `/u01`:

```
df -h
```

where you should have allocated sufficient size when you created the VM instance for the Enterprise Server in the amount specified in the *Before You Begin* section of this Learning Path.

7. Edit the `/etc/sysctl.conf` file:

```
sudo vi /etc/sysctl.conf
```

8. Ensure that this setting exists:

```
kernel.sem = 1250 32000 100 1024
```

9. Update `.bash_profile` of the **root, opc**, and **oracle** users with below export values.

> **Note:** This setting is required on the Oracle database target machine, which is the Enterprise Server, in order to enable the installation of the Oracle database client.

```
sudo vi /root/.bash_profile
export CV_ASSUME_DISTID=OL7

sudo vi /home/opc/.bash_profile
export CV_ASSUME_DISTID=OL7

sudo vi /home/oracle/.bash_profile
export CV_ASSUME_DISTID=OL7
```

Setting Up the WebLogic Server

This section shows you how to set up the WebLogic Server for JD Edwards EnterpriseOne on the Oracle Cloud Infrastructure.

Prerequisite

- You must have created a Linux VM for the WebLogic Server.
- The Linux VM for the WebLogic Server must be allocated with storage volume space, the amount of which is specified in the module of this Learning Path Entitled: **Before You Begin**.
- You must have performed the functions described in the module of this Learning Path entitled: **All Servers Common Setup**.
- If you are installing WebLogic Server in a Compute instance, complete sections 1 through 4.
- If you have installed WebLogic Server using an image provided by Oracle Cloud Infrastructure, complete sections 1 through 5.

General

The following general prerequisites are required for the Oracle WebLogic Server:

1. Ensure you have performed all the tasks described in module of this Learning Path entitled: *Performing Common Setup for All Linux Servers* including the steps to create groups and users as well as assigning proper permissions to `/u01`.
2. Use this command to ensure that the `/u01` mount point has proper permissions to allow the One-Click Provisioning process to subsequently create requisite folders:

   ```
   sudo chmod 770 /u01
   ```
3. Ensure the user with which Oracle WebLogic Server was installed is part of the **oracle** group. Use this command to verify:

   ```
   groups username
   ```

Install Oracle WebLogic Server and a JDK

You must manually pre-install Oracle WebLogic Server (WebLogic Server) and a JDK for use by WebLogic Server.

> **Note:** For the JDK installation, you must ensure that it is installed under the `/u01` mount point. Further you must ensure that the JDK is installed by the same user that has been or will be used to install WebLogic Server. For example, if you use the **oracle** user to install the JDK, you must use the **oracle** user to install WebLogic Server, or vice versa depending on the order you install these products.

Customers can then use JD Edwards One-Click Provisioning to deploy the JD Edwards EnterpriseOne web components into the pre-installed WebLogic Server domain. Refer to the Oracle Certifications for JD Edwards EnterpriseOne One-Click Provisioning on OCI for details on supported versions of WebLogic Server.

The complete list of prerequisites are below:

- Compute instance must have a supported version of WebLogic Server pre-installed.The password must be 8 to 30 characters long, must contain at least one number, and optionally any number of either the Pound Sign (#) or Underscore (_) special characters.

  > **Note:** The password for any WebLogic Server user *cannot* contain these special characters: Dollar Sign ($) Exclamation Mark (!) Ampersand (&) Using any of the above special characters violates the Oracle Cloud password policy and will result in denied access.

- Compute instance must also have a supported JDK version installed.
- WebLogic Server must be defined with at least one domain.
- WebLogic Server must be in a running state and able to connect to the Admin Console.
- WebLogic Server must have a running Node Manager process for the existing domain.
- WebLogic Server must be installed as a user that is part of the **oracle** group.

Post Installation of WebLogic Server

After you have installed WebLogic Server, prior to using the Provisioning Console for JD Edwards EnterpriseOne, you must manually configure WebLogic Node Manager to disable SSL. Otherwise the provisioning will fail.

1. Locate the `nodemanager.property` file which is typically found in this location:

`<ORACLE_HOME>/user_projects/domains/<DOMAINNAME>/nodemanager/nodemanager.properties`

2. Edit the `nodemanager.property` file so set this property to disable SSL:

   SecureListener=false
3. Save the `nodemanager.property` file.
4. You must set the listen address must be set exactly same as value of **ListenAddress** property of nodemanager.properties file. You must set the listen port 5556, which is non-secure (instead of 5557, which is secure) in both Node Manager properties file and in node manager using the WebLogic Server user interface. For instructions refer to this procedure: *Changing the Oracle HTTP Server Listen Ports*.
5. Restart the WebLogic Node Manager for the setting to take effect.

Ensure WebLogic Server is Patch Current

Refer to the Oracle Certifications for One-Click for WebLogic Server to verify that the latest available patches are applied. Such patches can affect the ability of One-Click Provisioning to deploy web components into WebLogic Server.

Configuring a WebLogic Server Instance Installed from the Marketplace of Oracle Compute Infrastructure

This section is only required if you have installed WebLogic Server from the Marketplace of Oracle Cloud Infrastructure.

The following describes feature support for an installation of WebLogic Server from the Marketplace for Oracle Cloud Infrastructure:

- **Standard or Enterprise Edition**

  Both of these editions are supported.

- **Number of Nodes**

  Only a single node is supported.

- **Load Balancer**

  This feature is not supported by JD Edwards One-Click Provisioning.

After the WebLogic Server is installed from the Marketplace of Oracle Cloud Infrastructure, you must perform further configuration as described below.

1. Connect to your instance of WebLogic Server.
2. Access the WebLogic Admin Console.
3. Navigate to Environments, Machines.
4. Select the tab for Configuration, Node Manager.
5. Set Type as PLAIN.
6. Save the selection.
7. Log in to WebLogic Server instance and edit the `/etc/hosts` file to add these two lines:

   ```
   <Private ip> localhost localhost.localdomain localhost4 localhost4.localdomain4

   <Private ip> <FQDN> <short host name>
   ```
8. Run this command to enable root account:

```
sudo usermod -s /bin/bash root
```

9 Preparing the Microsoft Windows Deployment Server

Creating a Windows VM in Oracle Cloud Infrastructure

This section describes how to create a Microsoft Windows instance in Oracle Cloud Infrastructure.

Oracle Cloud Infrastructure Compute lets you provision and manage compute hosts, known as instances.

> **Note:** While Oracle Cloud Infrastructure offers both Bare Metal and Virtual Machine instances, the current version of One-Click Provisioning for Oracle Cloud Infrastructure only supports Virtual Machine instances.

For additional information refer to *Creating an Instance*.

In preparation for JD Edwards EnterpriseOne One-Click Provisioning, each installation **must** create a Microsoft Windows VM instance in Oracle Cloud Infrastructure Compute for the Deployment Server, which is a JD Edwards EnterpriseOne core server.

Creating a Windows VM Instance in Oracle Cloud Infrastructure

With the following exceptions, the procedure for creating a Microsoft Windows instance as a VM is specified in the OBE in this Learning Path entitled "Creating the Linux Instance for the One-Click Provisioning Server".

Exception 1: Change Image for Microsoft Windows Server

Instead of having the requisite Operating System version automatically selected when you are installing from a Marketplace image for the Provisioning Server, you must manually select a specific image that is certified with JD

Edwards EnterpriseOne. Refer to the "Supported Software Versions" section of this Learning Path for the supported release of Microsoft Windows Server.

1. On Create Compute Instance, in the **Image** section, click the **Change Image** button.

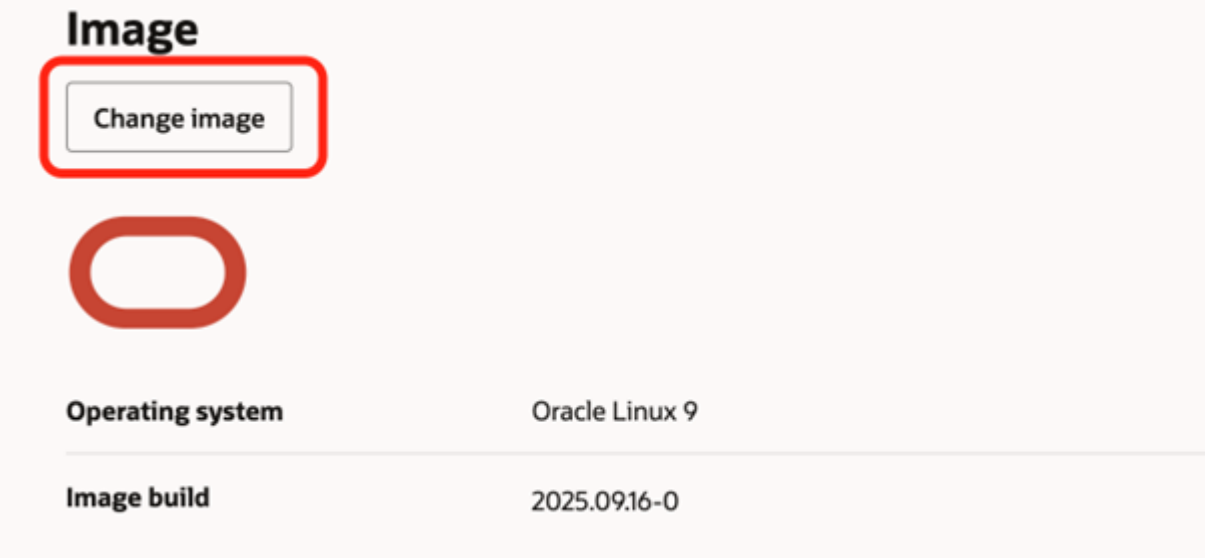

2. On Browse All Images:
   - In the **Image name** column, click the **Windows** checkbox.
   - In the **OS version** column, use the pull-down to select the supported version as specified in the "Supported Software Versions" section of this Learning Path. For example, **Server 2019 Standard**.
   - In the **Image build** column, use the pull-down to select the latest build. For example, **2021.03.17-0**.
3. Click the **Select Image** button.

Exception 2: Login Credentials

Upon creating this Microsoft Windows instance, both a user name and an initial password will be generated for you. They will be available on the details screen for the newly launched Instance. You must create a new password upon logging into the instance for the first time.

Exception 3: Sizing the VM

You must size the VM for the Deployment Server per the Minimum Resource Requirements, which currently states a value of 256 GB.

When you are creating the Windows VM instance, in the "Boot volume" section:

1. Select the option “Specify a custom boot volume size”.
2. In the “Boot volume size (GB)” field, enter the value `256` .

3. If it is selected by default, remove the selection for option “Use in-transit encryption”. Neither encryption checkbox should be selected.

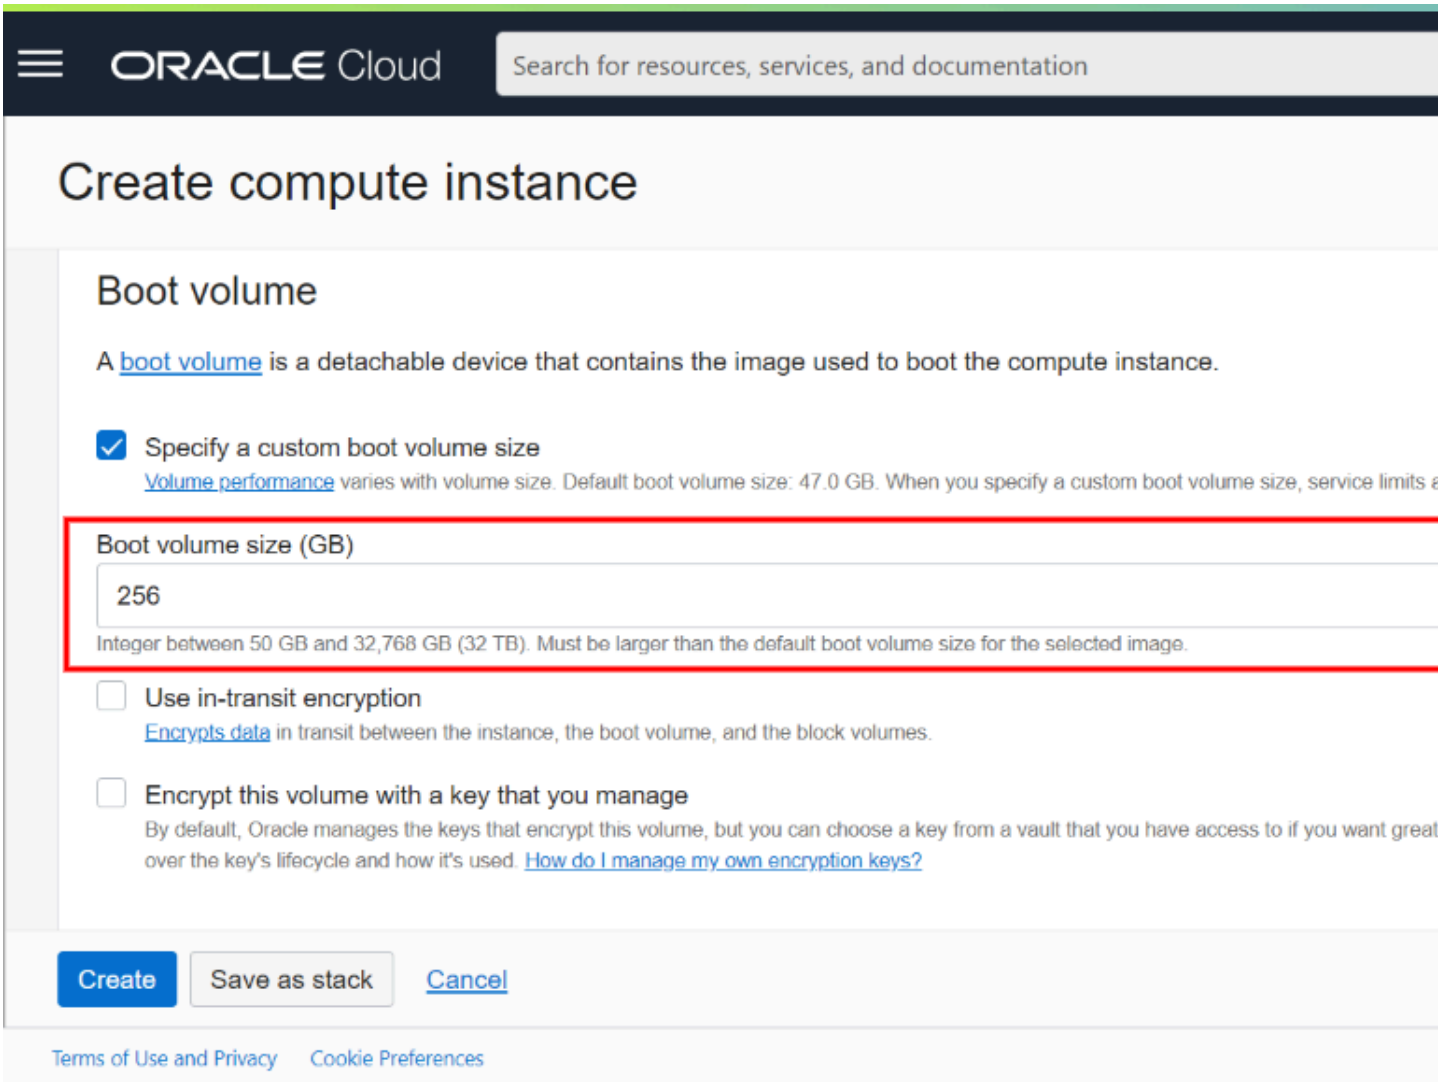

10 Configuring the Microsoft Windows Deployment Server

Logging in to the Windows VM

This section shows you how to log in to the Microsoft Windows virtual machine (VM) in Oracle Cloud Infrastructure.

Prerequisite

You must have previously created a Microsoft Windows VM in Oracle Cloud Infrastructure as described in the previous module of this Learning Path entitled: *Creating a Windows VM in Oracle Cloud Infrastructure*.

Log in to the Windows VM

1. Connect to the Microsoft Windows VM in Oracle Cloud Infrastructure with a Remote Desktop Protocol (RDP) connection using this command in a Command Window:

   ```
   mstsc /f
   ```

2. On Remote Desktop Connection, complete these fields:

   - ***Computer***

     Enter the Public IP Address which can be derived from the Instance Details for the Microsoft Windows Instance for the Deployment Server in the **Public IP Address** field.

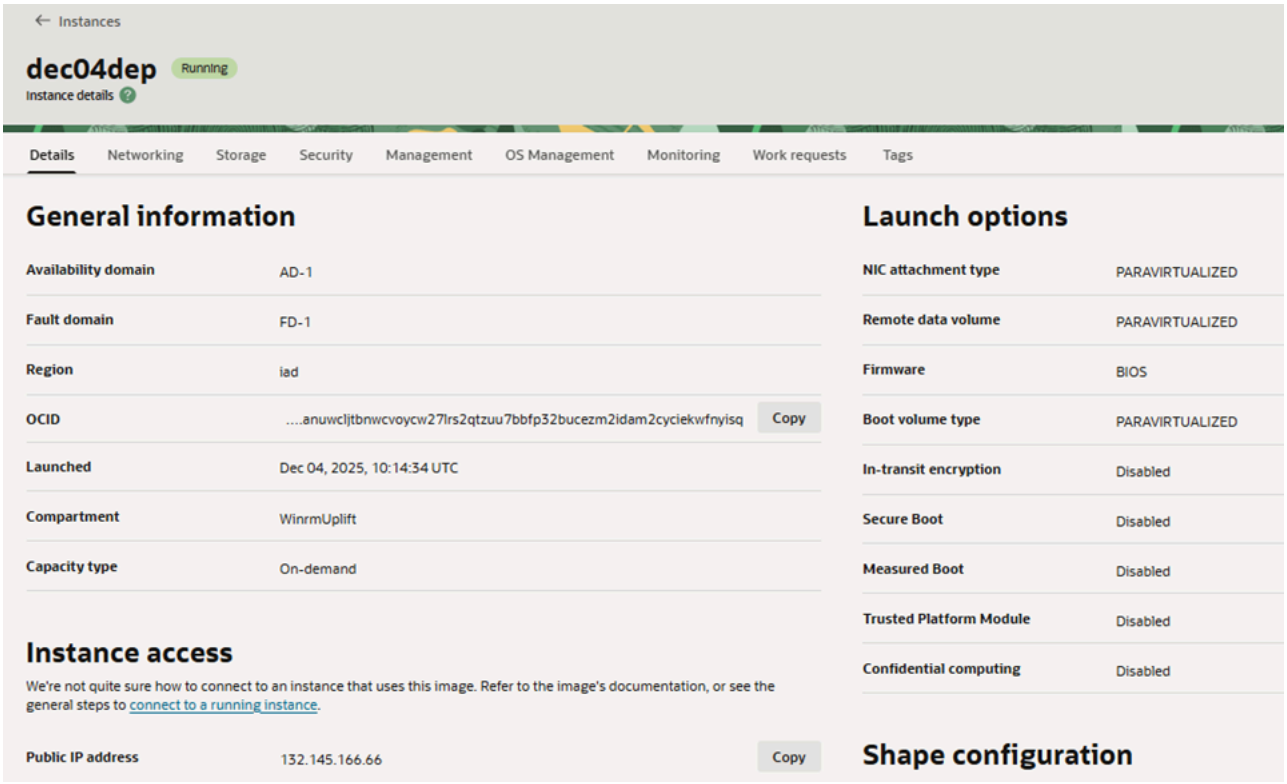

   - ***User name***

     Enter the value **opc**.

3. Click the **Connect** button.

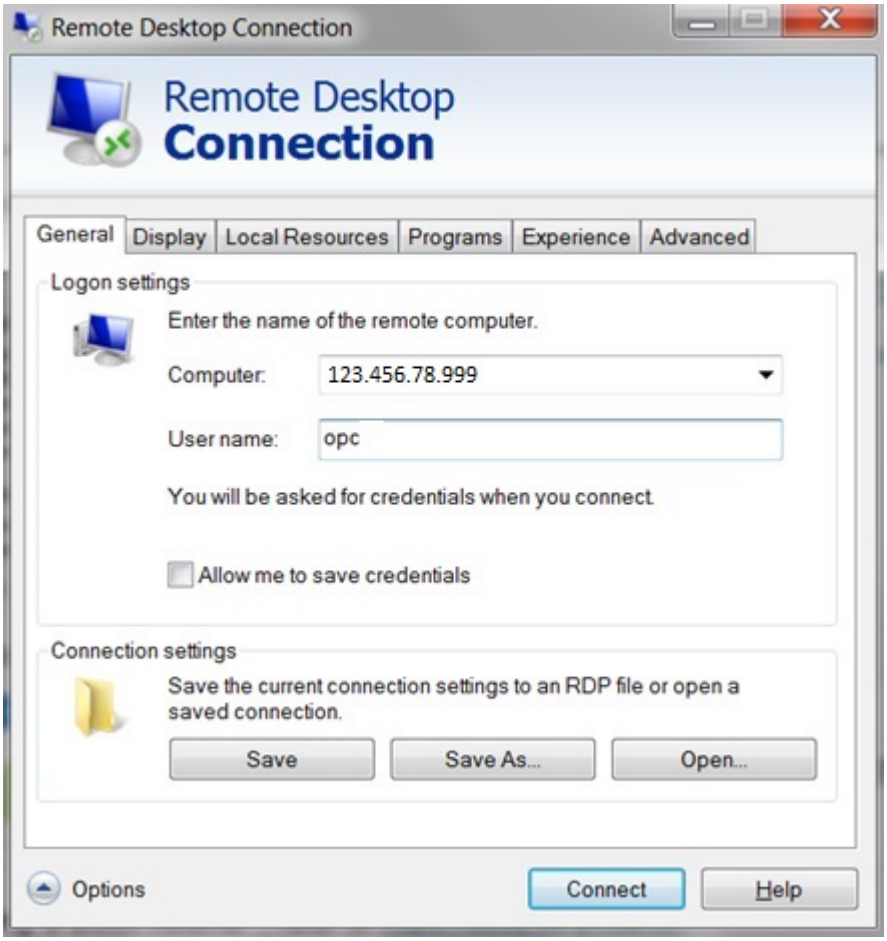

4. You are prompted to enter a password. This is the password that the system assigned upon creation of the Windows VM. On the **Instance Details** screen for the Windows VM that you created, locate the **Initial password** field.

5. In the **Initial password** field, click to expand more options and click the **Copy** option to copy the password.

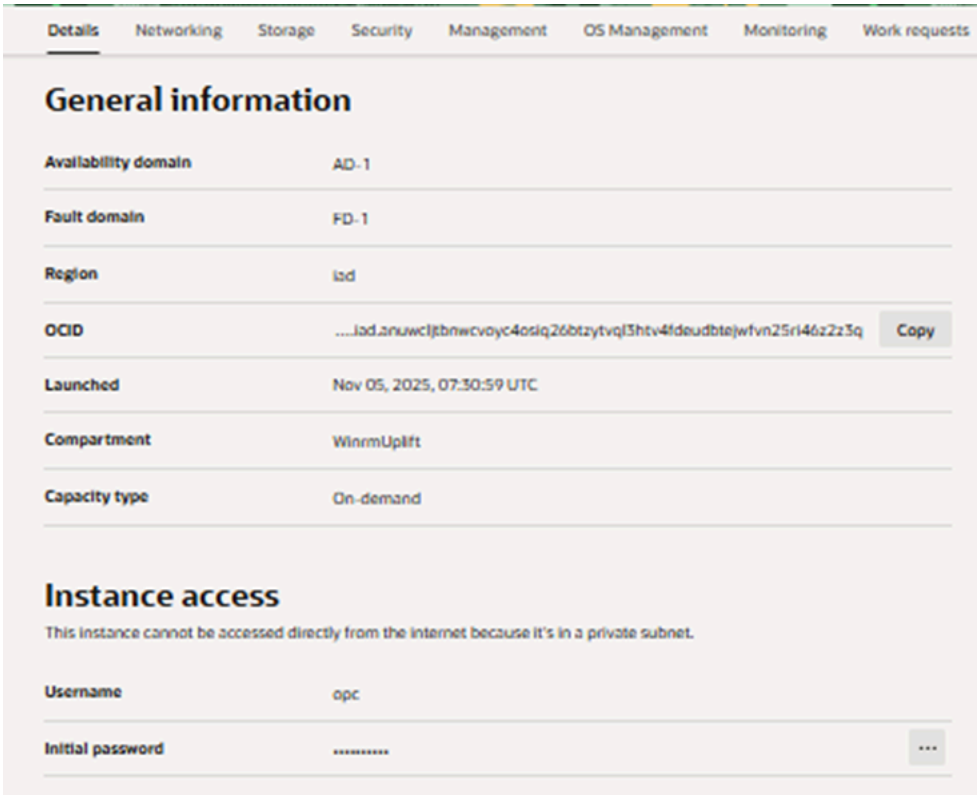

6. On Remote Desktop Connection, enter `opc` as the user and paste the copied password.

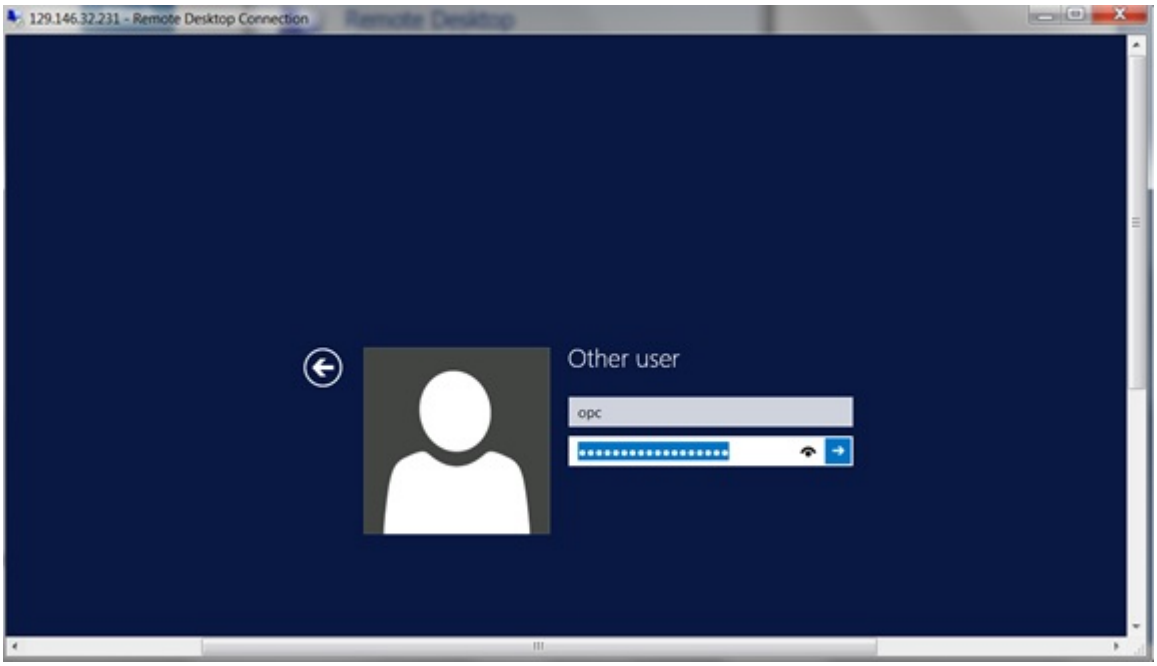

7. After you enter the credentials the first time you log in, you are prompted to change the password:

The user's password must be changed before signing in.

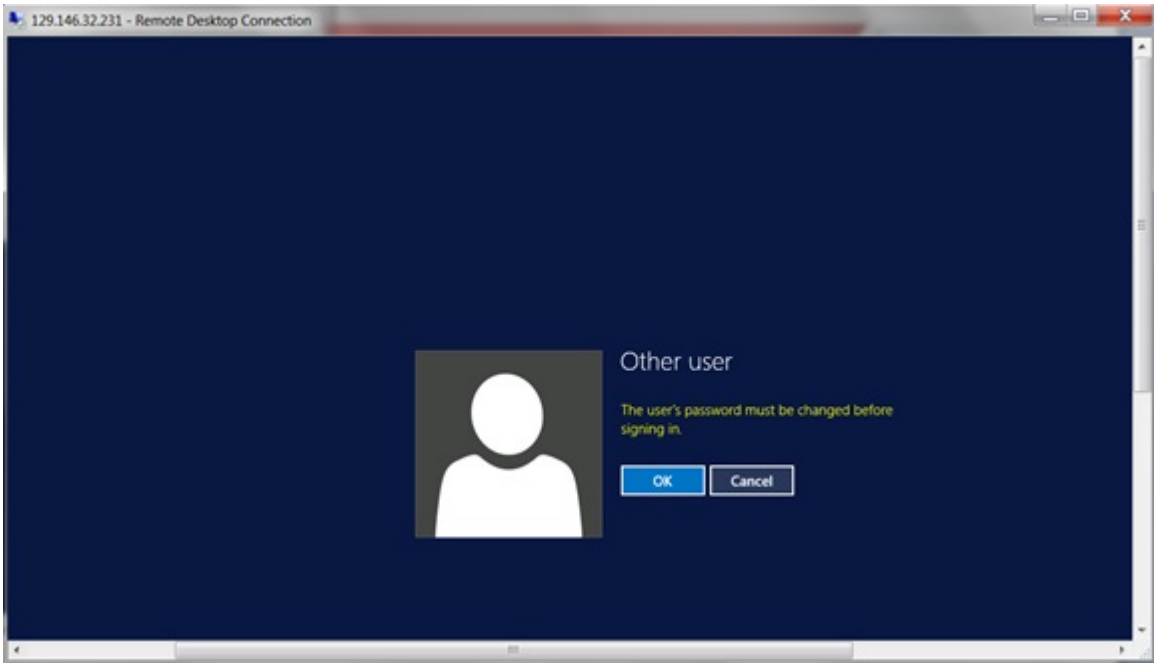

8. Change the password for the **opc** user following the Microsoft Windows policy for creating a user password.

> **Note:** The Password must contain at least 12 characters and must be comprised of at least one lower case letter, one upper case letter, one numeric digit, and one special character (for example; special characters could be _ @ ~ ! # % * +. Be sure and make a note of this password because you will need to use this same password later in the One-Click Provisioning Console.

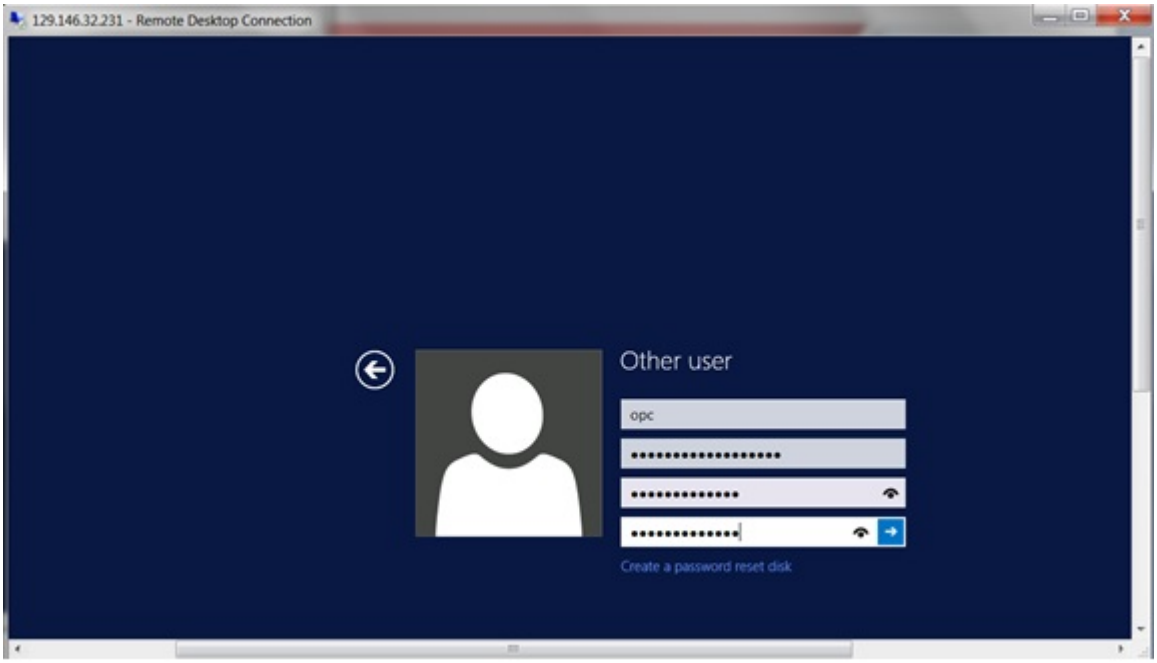

Running Commands to Change Microsoft Windows Settings

This section shows you how to run commands to change Microsoft Windows settings.

Prerequisites

- You must have access to a physical or virtual Windows machine dedicated for use as a JD Edwards EnterpriseOne Deployment Server.
- You must be able to connect to the Microsoft Windows machine using a Remote Desktop Protocol (RDP) session from a Microsoft Windows client machine.

Enable Remote Command Execution Through Microsoft PowerShell

From Microsoft PowerShell, use this command to enable remote command execution:

```
winrm quickconfig -q
```

Enable Inbound Ports in the Firewall

If you have the Microsoft Windows firewall enabled for any profile (public, private, default), which is recommended, in order for One-Click Provisioning to deploy to the Deployment Server and for JD Edwards EnterpriseOne runtime to function properly, you will need to explicitly open Inbound and Outbound ports.

Use this procedure to open the Inbound ports for **each** Microsoft Windows instance.

1. Go to Administrative Tools > Windows Firewall with Advanced Security.
2. In the left pane, highlight **Inbound Rules**.
3. In the right pane, Actions, Inbound Rules, click **New Rule** ...
4. On Rule Type, select the **Port** radio button as the type of rule to create.
5. Click the **Next** button.
6. On Protocol and Ports:
    - You can accept the default value of **TCP** for the protocol to which this rule applies.
    - Choose the radio button for **Specific local ports** and, for ***each*** Windows Server, enter each of the ports shown in the following table, separated by a comma.

Windows Server Firewall Port List	
Component	**Inbound Ports to Open**
Deployment Server	445 3389 5150 5985 6017-6022 14502-14510

The following example shows the ports you should specify for the Deployment Server.

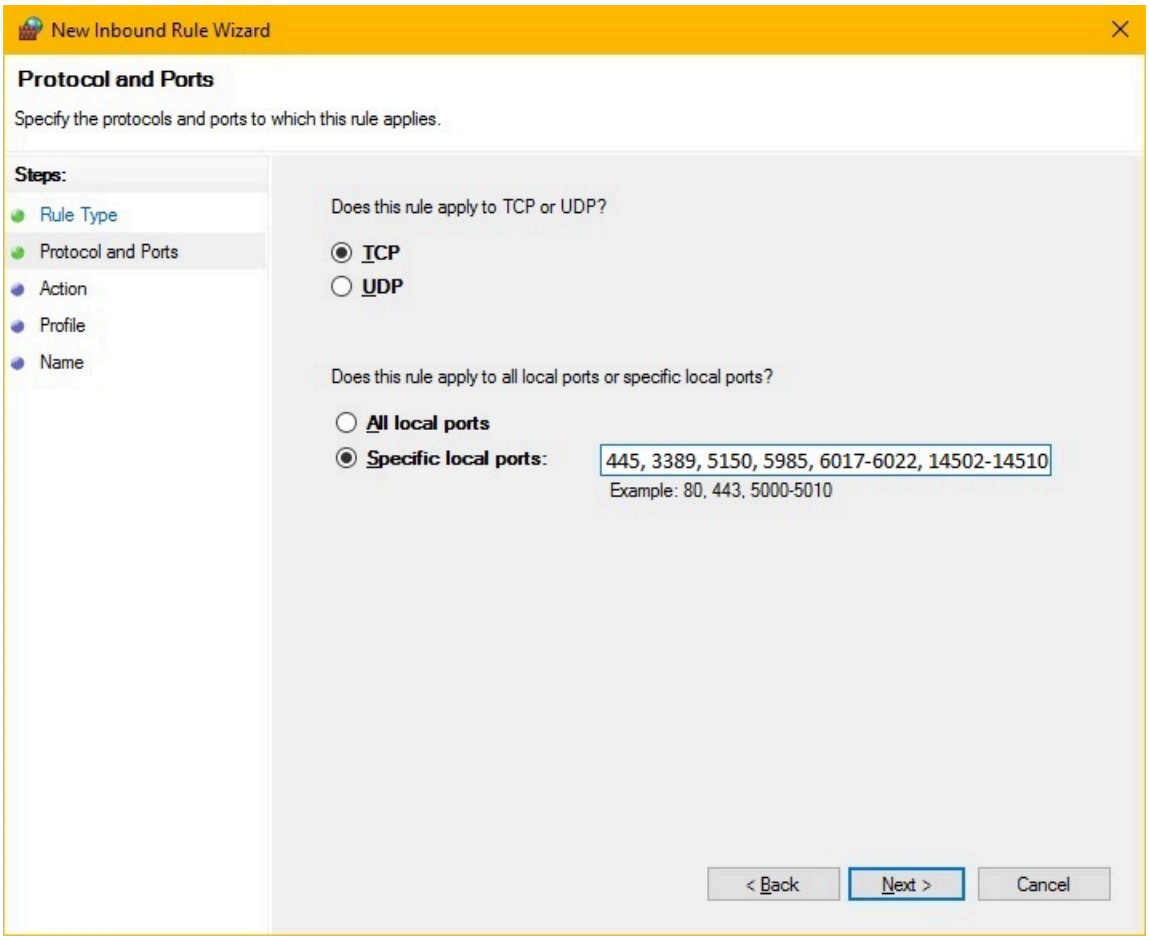

7. Click the **Next** button.
8. On Action, you can accept the default value which is **Allow the connection**.
9. Click the **Next** button.
10. On Profile, select all firewall profile options **(Domain, Private, Public)**.
11. Specify a name for the rule. For example, **JDESMC_RDP**.
12. Click the **Finish** button to save the rule and exit the wizard.

Enable Outbound Ports in the Firewall

If you have the Microsoft Windows firewall enabled for any profile (public, private, default), which is recommended, in order for One-Click Provisioning to deploy to the Deployment Server and for JD Edwards EnterpriseOne run-time to function properly, you will need to explicitly open Inbound and Outbound ports.

Use this procedure to open the Outbound ports in your Microsoft Windows instance.

1. Go to Administrative Tools > Windows Firewall with Advanced Security.
2. In the left pane, highlight **Outbound Rules**.

3. In the right pane, Actions, click **New Rule** ...
4. On Rule Type, select the **Port** option as the type of rule to create.
5. Click the **Next** button.
6. On Protocol and Ports:
   - You can accept the default value of **TCP** for the protocol to which this rule applies.
   - Choose the option for **All remote ports**.

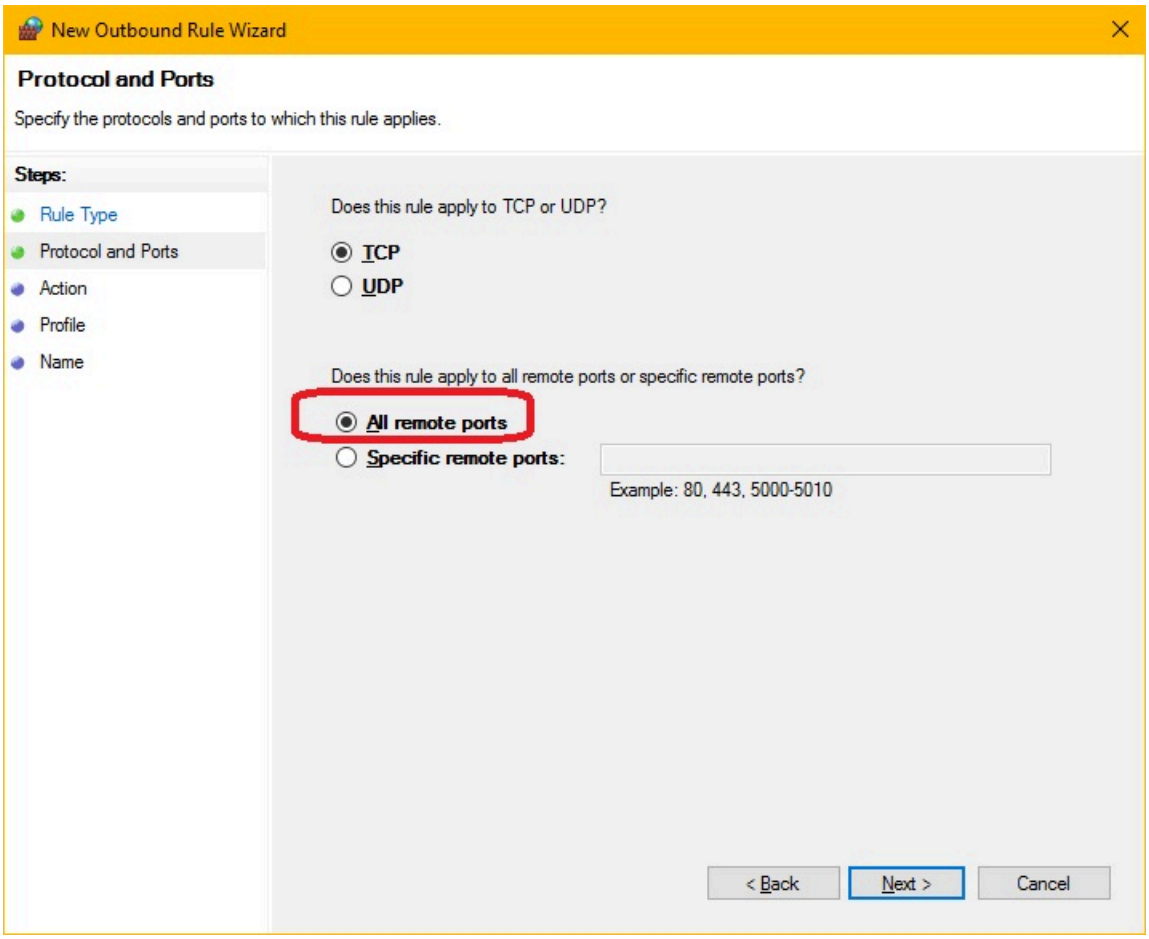

7. On Action, you can accept the default value which is **Allow the connection**.
8. Click the **Next** button.
9. On Profile, select all firewall profile options **(Domain, Private, Public)**.
10. Specify a name for the rule.
11. Click the **Finish** button to save the rule and exit the wizard.

Configure Settings for Ethernet Connections

You must configure settings for Ethernet connection to specify the domain name of the Domain Name System (DNS) for the Availability Domain to which all JD Edwards EnterpriseOne servers belong.

For example, if your subnets look like that shown in the example below, you must configure your Network Settings using this procedure:

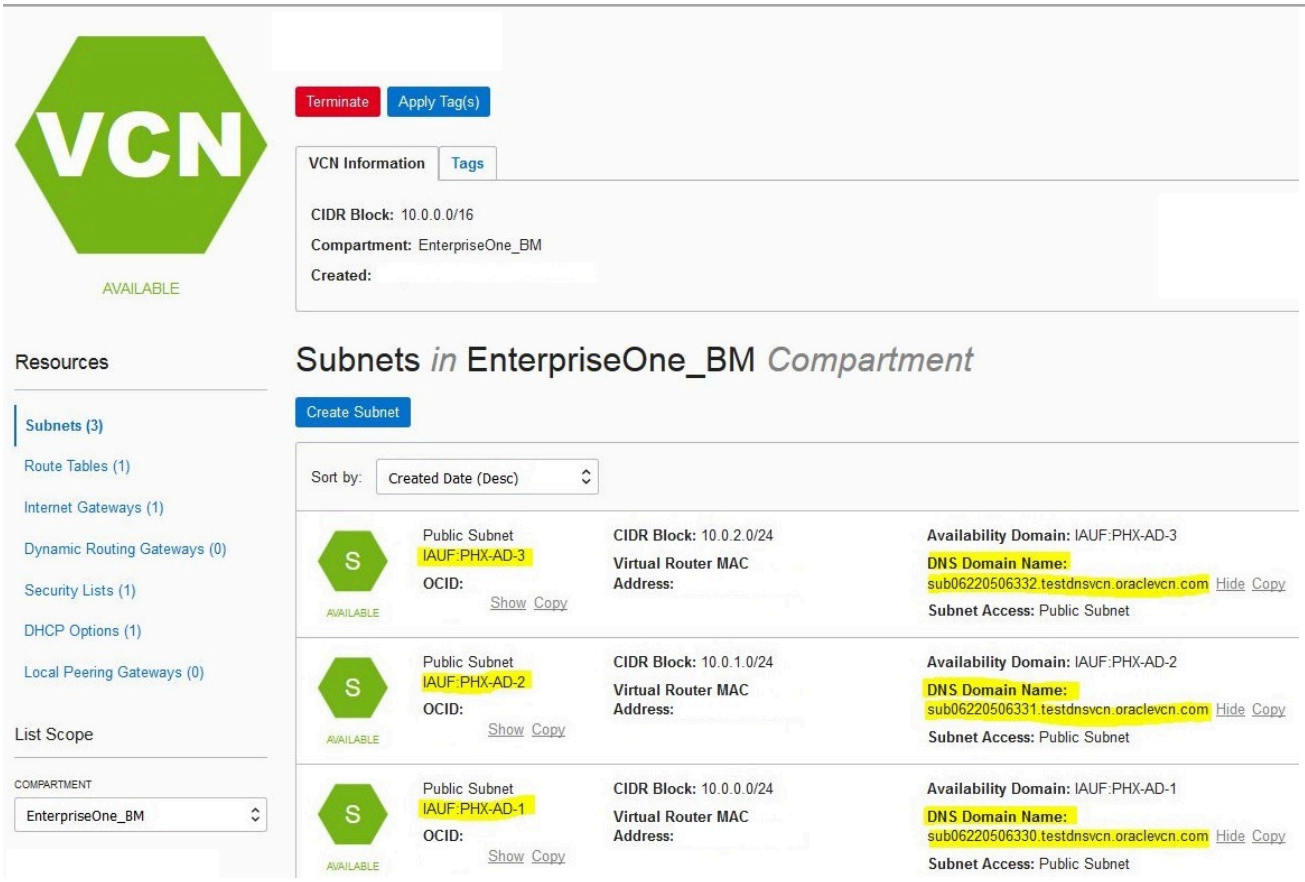

1. Open **Network and Sharing Center**.

2. On Ethernet settings, in Network > Connections, click **Ethernet** to open **Ethernet Status**.

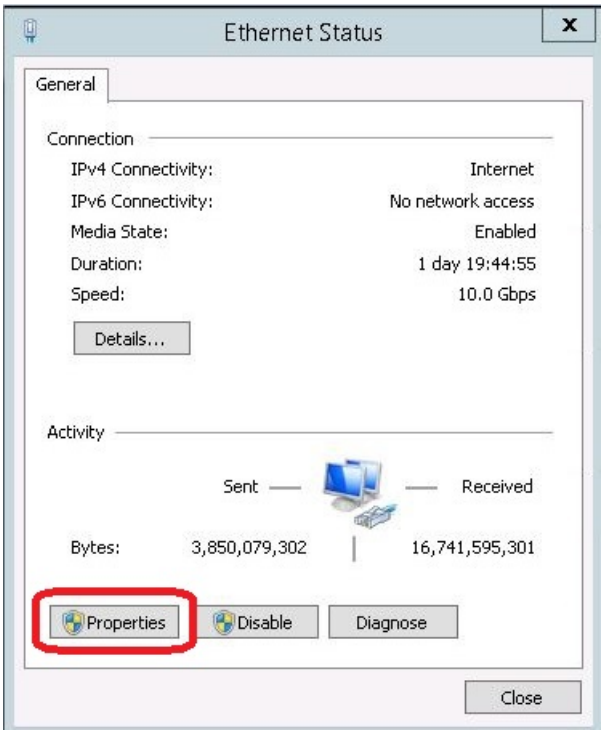

3. On Ethernet Status, click the **Properties** button.

4. On Ethernet Properties, highlight this line: **Internet Protocol Version 4 (TCP/IPv4)** and click the **Properties** button.

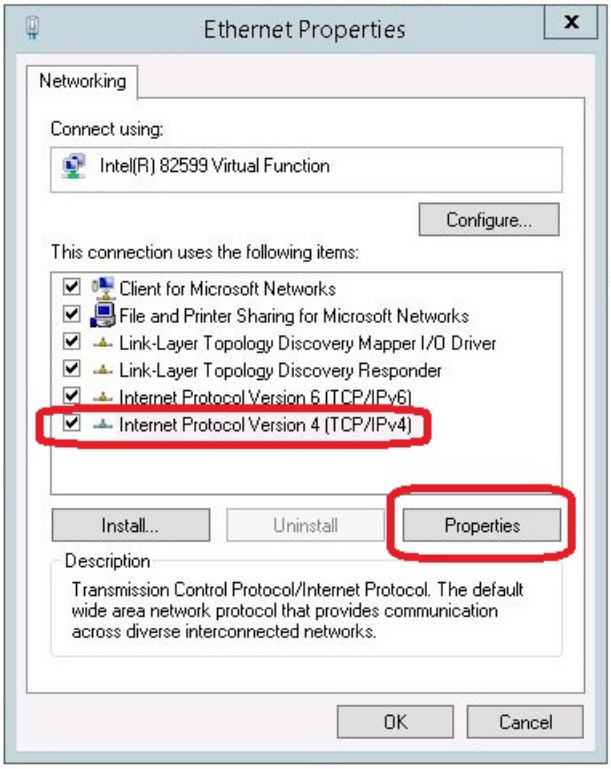

5. On Internet Protocol Version 4 (TCP/IPv4) Properties, click the **Advanced** button.

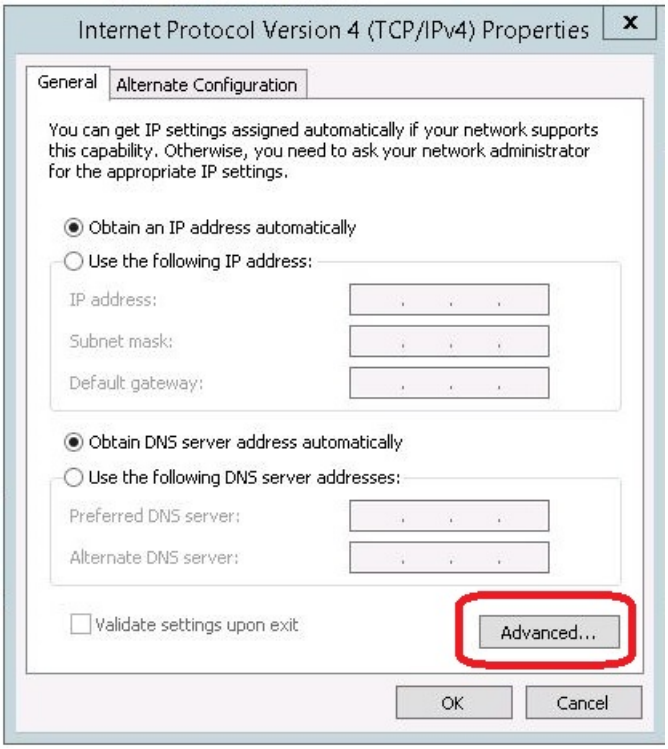

6. On Advanced TCP/IP Settings, click to enable this option: **Append these DNS suffixes (in order):** and click the **Add** button.
7. On the TCP/IP Domain Suffix dialog, enter the value of the DNS Domain Name for your Availability Domain. For example, assuming your subnets were as shown in the preceding screen showing the subnets for each Availability Domain and all servers are created in **sub0622506330.testdnsvcn.oraclevcn.com** subnet, you would enter this value for the suffix:

   ```
   sub0622506330.testdnsvcn.oraclevcn.com
   ```
8. Click the **Add** button to add the DNS suffix.

9. Add this same subnet (in this example, **sub0622506330.testdnsvcn.oraclevcn.com**) in the field labelled **DNS suffix for this connection**.

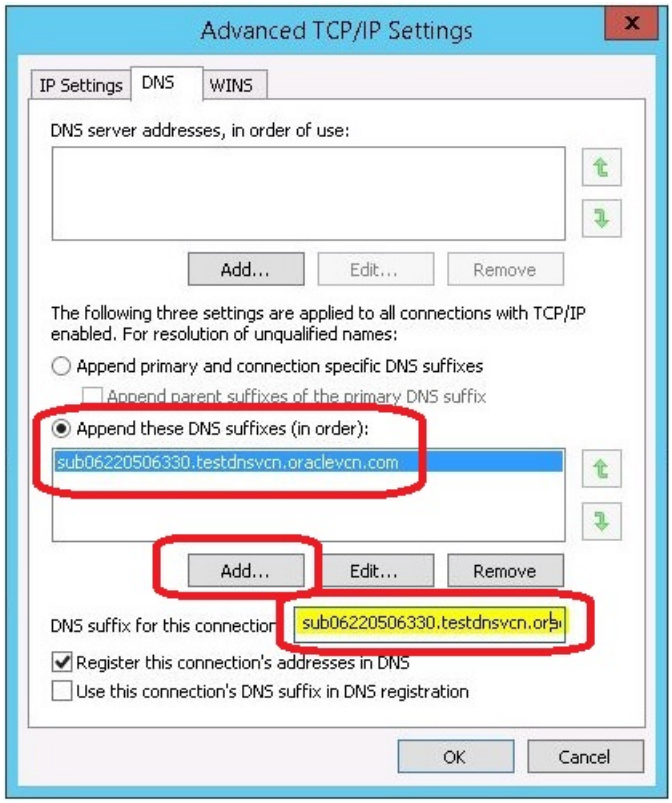

10. Click the **OK** button to accept the values and exit the **Network and Sharing Center**.

Change Security Option

If you chose a user other than **opc** when for the One-Click Provisioning deployment of your Deployment Server, use this procedure to change the Microsoft Windows security option so that user will be recognized by JD Edwards EnterpriseOne.

1. Log in to the Deployment Server as the user you configured, which is other than the **opc** user.
2. Open the Microsoft Windows **Local Security Policy** program from Start > Run, or from a Command Prompt:

   ```
   secpol.msc
   ```
3. On Local Security Policy, navigate to Local Policies > Security Options.
4. Locate and edit this setting:

   Network Security: LAN Manager authentication level

5. Use the pulldown to select this setting:

   Send NTLM V2 response only

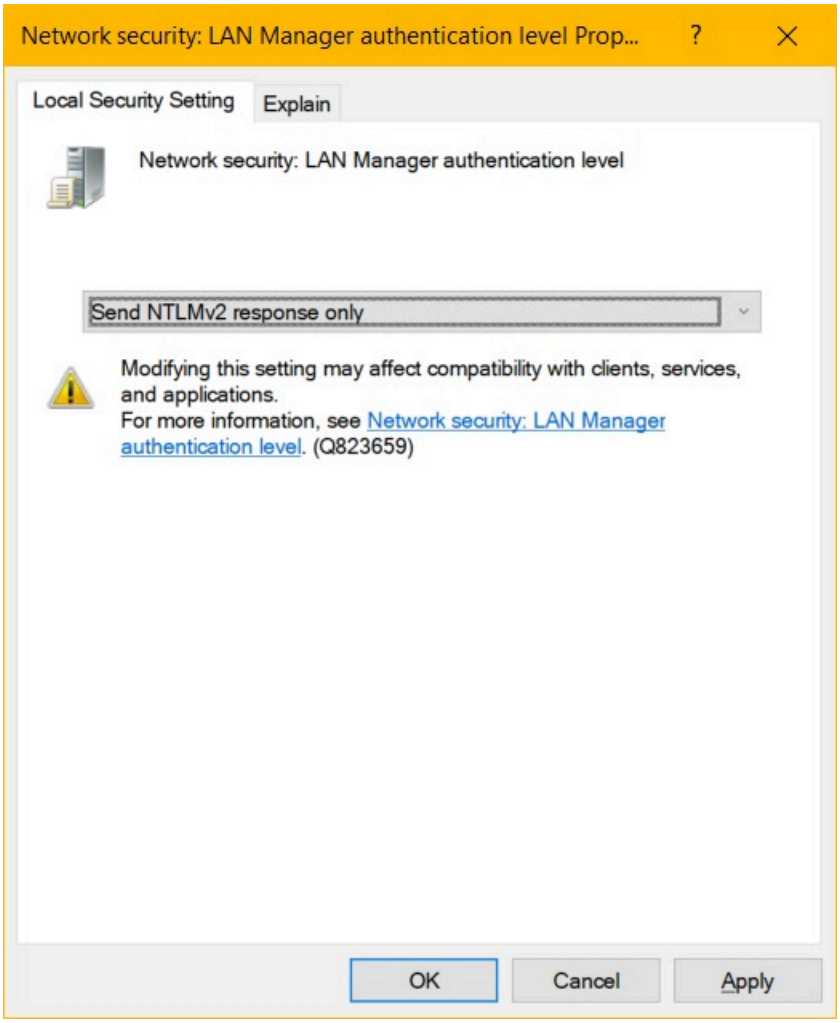

6. Click the **Apply** button to save the setting.

Change the Maximum Transmission Unit (MTU) Setting

The recommended MTU setting is 1500. Use this procedure to check, and if necessary, change the current MTU setting.

1. Open Windows PowerShell as Administrator and run the following command (as a single contiguous line) to check the current value of the MTU setting:

```
Get-NetIPInterface | where {($_.InterfaceAlias -eq "Ethernet") -and ($_.AddressFamily -eq "IPv4") -and
($_.NlMtu -Gt 0)}
```

2. If the returned value is not set to 1500, run this command (as a single contiguous line) to set the MTU value to 1500:

```
Get-NetIPInterface | where {($_.InterfaceAlias -eq "Ethernet") -and ($_.AddressFamily -eq "IPv4") -and
($_.NlMtu -Gt 0)} | Set-NetIPInterface -NlMtuBytes 1500
```

11 Using the One-Click Provisioning Server

Configuring CA Certificates for One-Click UI (Optional)

This section shows you how to configure CA Certificates for the One-Click User Interface. This task is optional.

JD Edwards EnterpriseOne One-Click Provisioning can be configured with a supplied CA certificate. However, after initial installation you must obtain your own CA certificates. These must be certificates that are verified by a verified CA authority such as Entrust and Symantec Corporation.

Prerequisite

You must obtain CA certificates and their chain of deliverables.

Configuring CA Certificates for One-Click UI (Optional)

Use this procedure to configure CA certificates for the One-Click Provisioning user interface. This step is not required for initial installation but is required for subsequent usage of JD Edwards EnterpriseOne in Oracle Cloud Infrastructure.

1. Combine the the contents of `rootca.pem` and `intermediateca.pem` certificates to a file named `ca-cert.pem` file as per order from the CA authority.
2. Rename `.net.key` to `key.pem`.
3. Rename `<variable>.crt` to `cert.pem`.
4. Copy `ca-cert.pem`, `key.pem` and `cert.pem` to this directory:

   ```
   /E1CloudConsole/keys
   ```
5. Restart the Provisioning Server:
   - **Linux-based Provisioning Server**

     Restart the E1CloudConsole.service using following command:

     ```
     sudo systemctl stop E1CloudConsole.service sudo systemctl start E1CloudConsole.service
     ```
   - **Microsoft Windows-based Provisioning Server**

     Use these steps to restart the user interface:

     - Navigate to this directory:

       ```
       \JDE\bin
       ```
     - Open Powershell with Administrator rights.
     - Run this script:

       ```
       .\E1CloudConsoleWin.ps1
       ```

Accessing the JD Edwards Provisioning Console

This section shows you how to access the JD Edwards Provisioning Console.

As described in the the subsequent OBEs in this Learning Path, you will access and use the Provisioning Console to set up a completely functional EnterpriseOne environment using the JD Edwards One-Click Provisioning Console by performing these three steps:

- Configure – Provide the Server Manager details.
- Orchestrate – Create the deployment plan.
- Deploy – Initiate the scripts for the automated provisioning of the EnterpriseOne system.

Prerequisites

- You should have a fundamental understanding of the Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription and an Administrator account to Oracle Cloud Infrastructure as described at this site: *Getting Started with Oracle Cloud Applications*
- To access the Provisioning Console, you must use a supported browser. See **Supported Browsers** in *Troubleshooting Signing In to the Console*.

Accessing the JD Edwards Provisioning Console

The recommended browsers for accessing the One-Click Provisioning Console are:

- Google Chrome
- Mozilla Firefox

  Open the browser and connect to the URL using this syntax:

  ```
  https://xxx.xxx.xxx.xxx:3000
  ```

  Where:

  - **https** is the only supported browser protocol.
  - **XXX.XXX.XXX.XXX** is the public IP address for the One-Click Provisioning Server Instance. This IP address is shown in the Oracle Cloud Infrastructure user interface on the Instance Details screen for the instance that you created for the One-Click Provisioning Server, as shown in the below example:

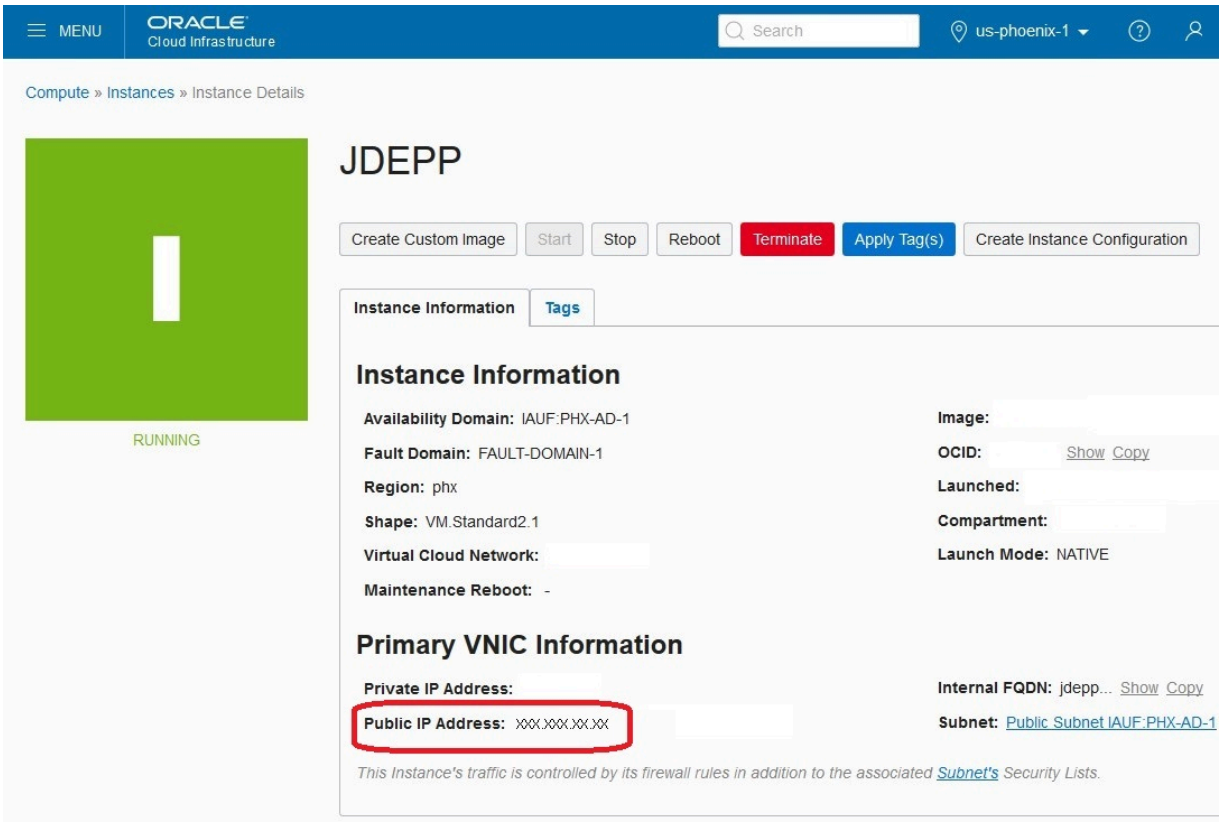

  - **3000** is the port on which the One-Click Provisioning Server is running, and which you must include as part of the address.

Note: Should you exit the Provisioning Console or experience a timeout, the next time you enter the Provisioning Console it displays the point where you left off. When you click the **Configure** icon, you will be prompted for the same credentials that you entered during your first access to the One-Click Provisioning Console.

Configuring the Server Manager Account

This tutorial shows how to configure the Server Manager account in JD Edwards One-Click Provisioning Console.

Configuring the Server Manager Account

Configuring the Server Manager Account

This section shows how to configure the Server Manager account in JD Edwards One-Click Provisioning Console.

If you are a new user, you are required to change the administrator passwords for WebLogic Server and Server Manager Console. Remember the Server Manager password you entered in the Change Password window for future logins.

Prerequisite

Downloaded the One-Click archive files from the Oracle Software Delivery Cloud (also called OSDC or E-Delivery). For more information see the section entitled: ***Setting Up the Provisioning Server*** .

Configure the Server Manager Account

Use this procedure to provide the account information.

1. On Welcome to the JD Edwards Provisioning Console, click the **Configure** icon.

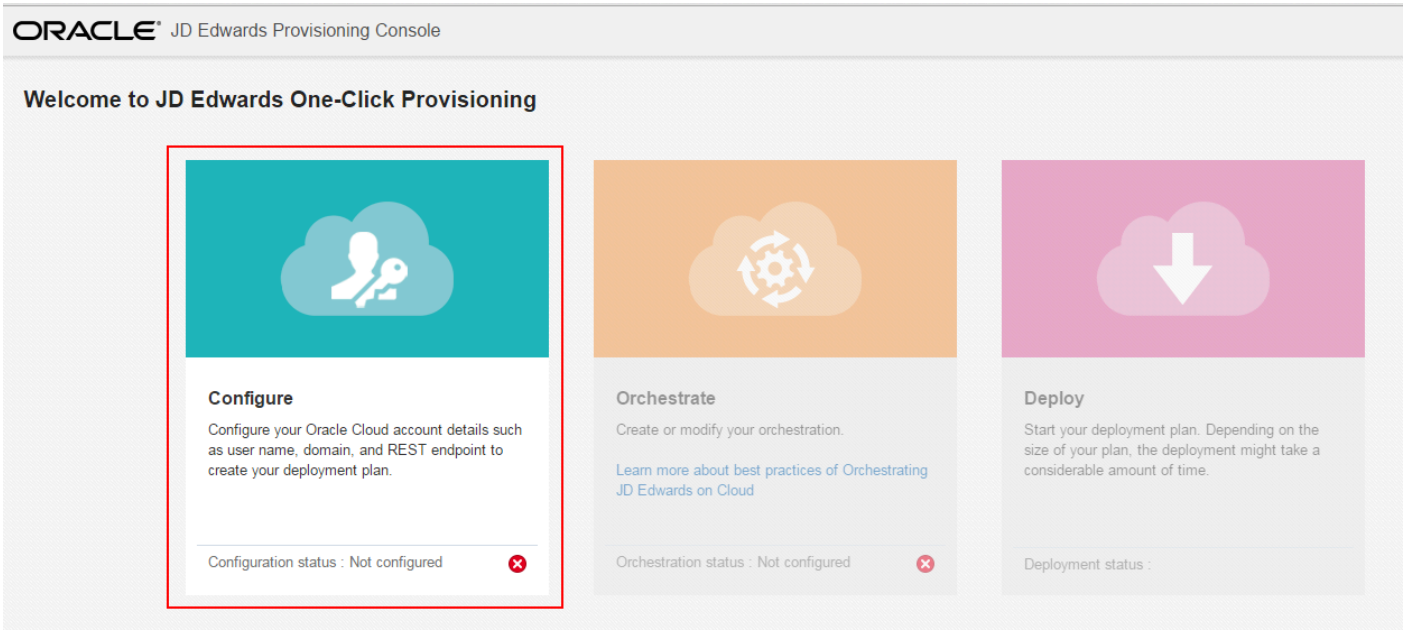

2. If you are a new user, as prompted by the JD Edwards Provisioning Console, you are also required to change the administrator passwords for WebLogic Server and Server Manager Console.

   On Change Password, enter the passwords for the Server Manager Administrator. The password must only have numbers, alphabets, and special characters (@,!,$,_,#), and is between 8 and 30 characters long.

   Also, enter the password for the WebLogic Server. The password must start with a letter, is between 8 and 30 characters long, contains at least one number, and optionally, any number of special characters (#,_). For example, Ach1z0#d

   > **Note:** The password for any WebLogic Server user cannot contain the $ or ! character. Using either of these characters violates the Oracle password policy and will result in denied access.

   > **Tip:** Valid values for the passwords are displayed in the tooltip when you click the field.

   Click the **OK** button.

   > **Note:** The system displays a message indicating that is changing default passwords; this may take a few minutes to complete before the next screen is displayed.

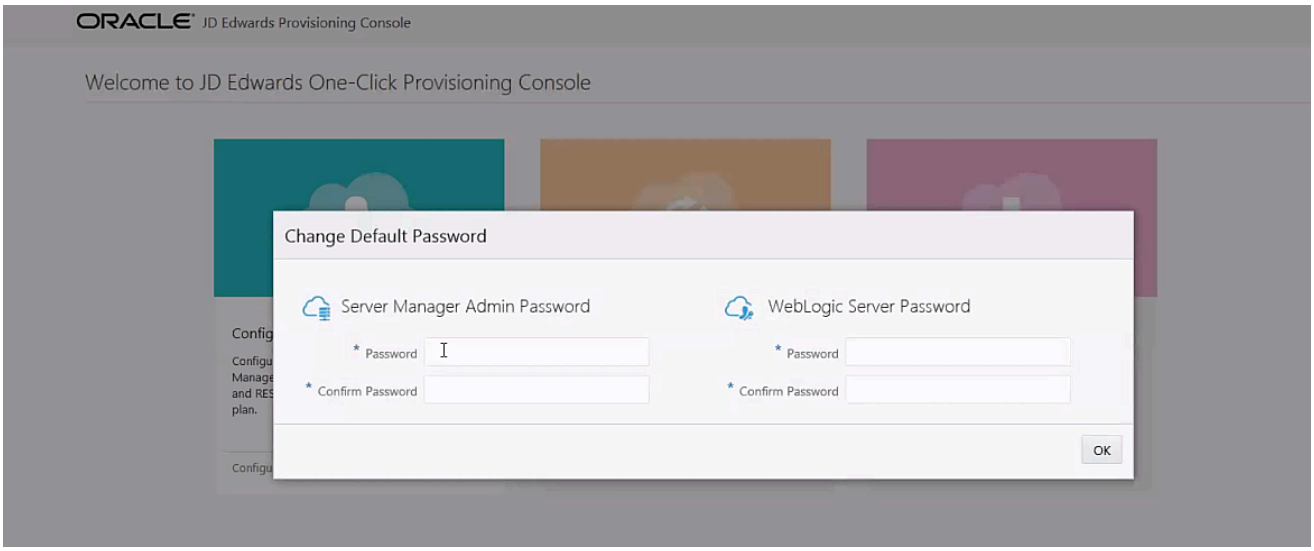

3. On the Success window, click the **OK** button.

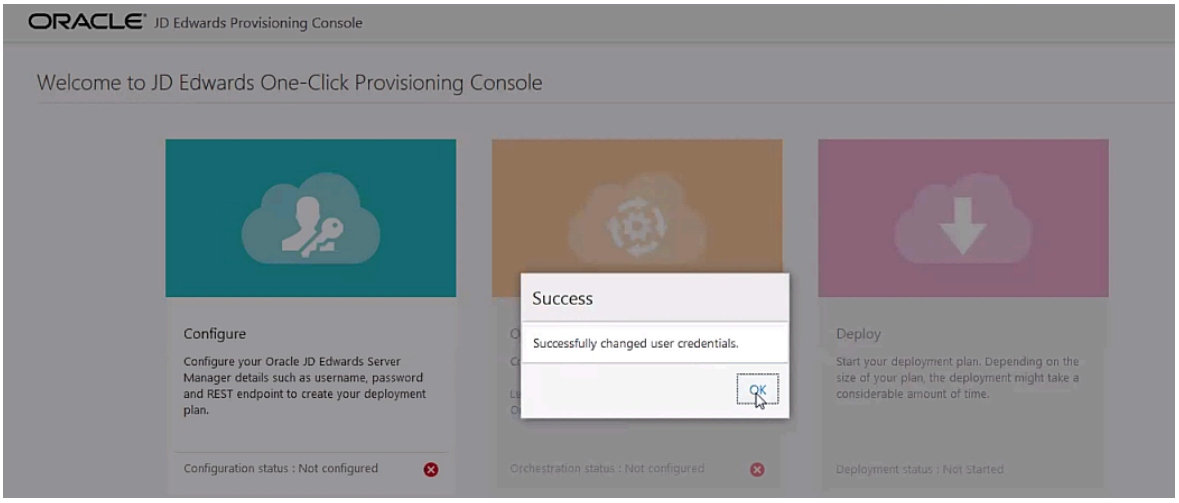

4. On Oracle JD Edwards Server Manager Details, reenter the Admin password for the Server Manager Console. The SMC Endpoint is pre-populated automatically.

> **Note:** It is recommended that you record this value on the Pre-Install Worksheet, which you created as described in the companion document to this tutorial in the section entitled: Create the **Pre-Install Worksheet**.

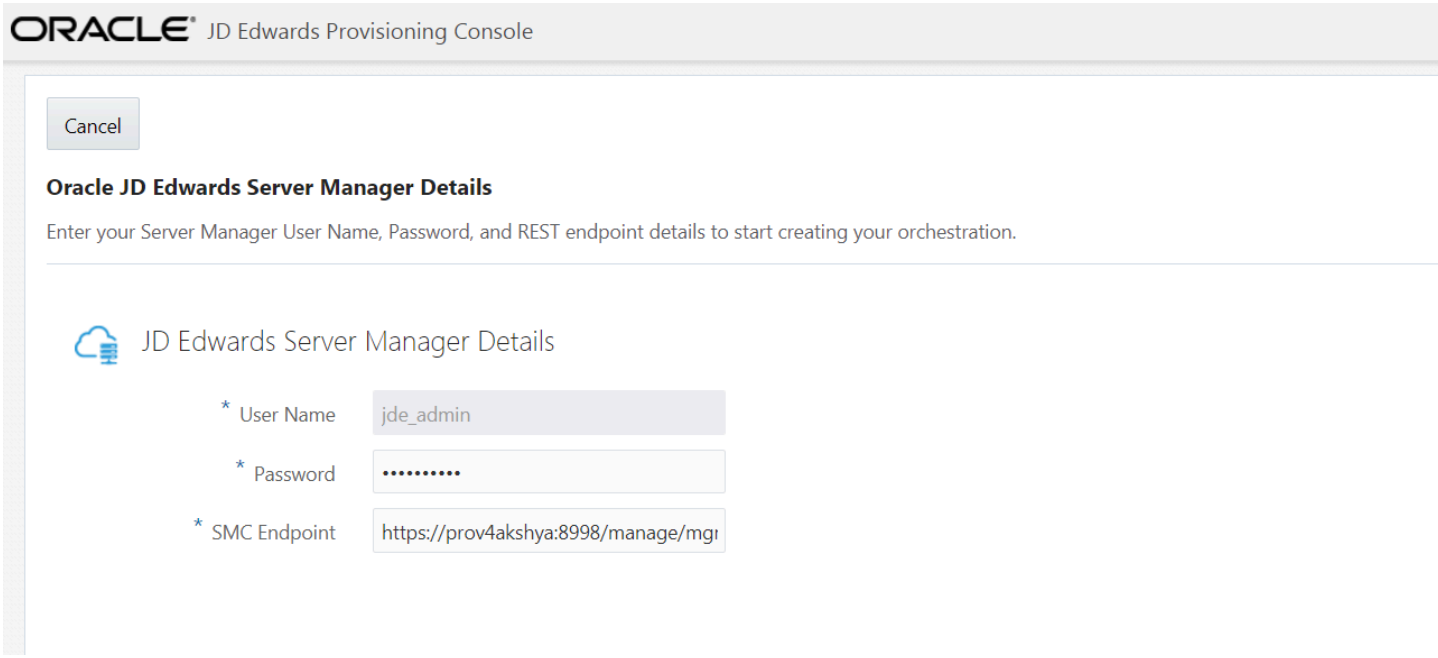

5. Click the **Save Configuration** button.

   The system will take some time to authenticate the Server Manager Endpoint and Credentials. When the authentication is verified, click the **OK** button on the **Congratulations** box.

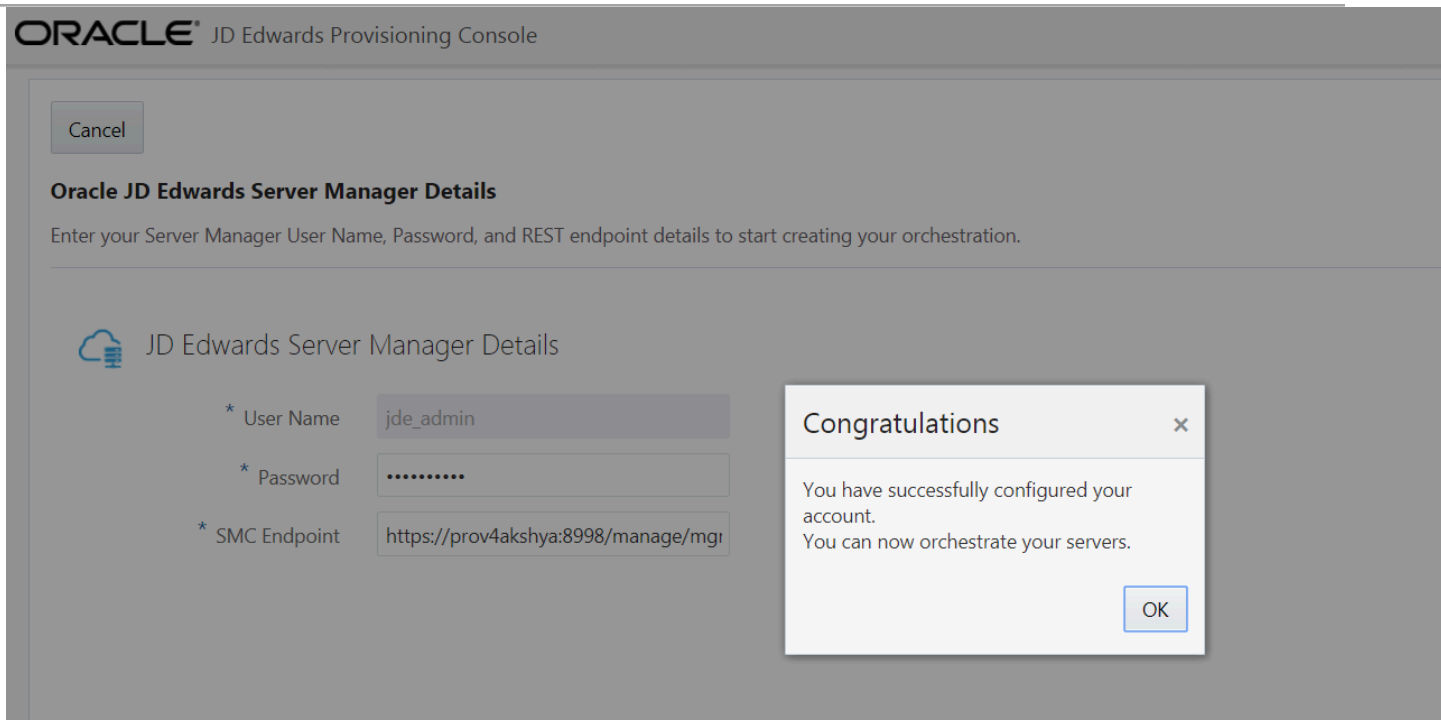
ORACLE JD Edwards Provisioning Console
Cancel
Oracle JD Edwards Server Manager Details
Enter your Server Manager User Name, Password, and REST endpoint details to start creating your orchestration.
JD Edwards Server Manager Details
* User Name
jde_admin
* Password
* SMC Endpoint
https://prov4akshya:8998/manage/mgr
Congratulations
You have successfully configured your account.
You can now orchestrate your servers.
OK

12 Creating a Deployment Plan

Orchestrating Using Quick Start Mode

This section shows you how to orchestrate a a basic deployment plan using the Quick Start mode. For performance reasons, creating an orchestration by using only the Quick Start mode is not recommended for production environments.

Use the Quick Start mode in the JD Edwards One-Click Provisioning Console to create a plan to deploy all the core components of JD Edwards EnterpriseOne on Oracle Cloud Infrastructure.

Refer to the "*Fundamentals*" section of this Learning Path for a description of the various servers that can be deployed by One-Click Provisioning.

Prerequisite

You must have configured the administrator passwords for the WebLogic Server and the Server Manager Console in the Configure section of the JD Edwards One-Click Provisioning Console.

Orchestrating a Quick Start Deployment Plan

To orchestrate a Quick Start deployment plan:

1. On the Welcome to JD Edwards One-Click Provisioning Console page, click the Orchestrate icon.

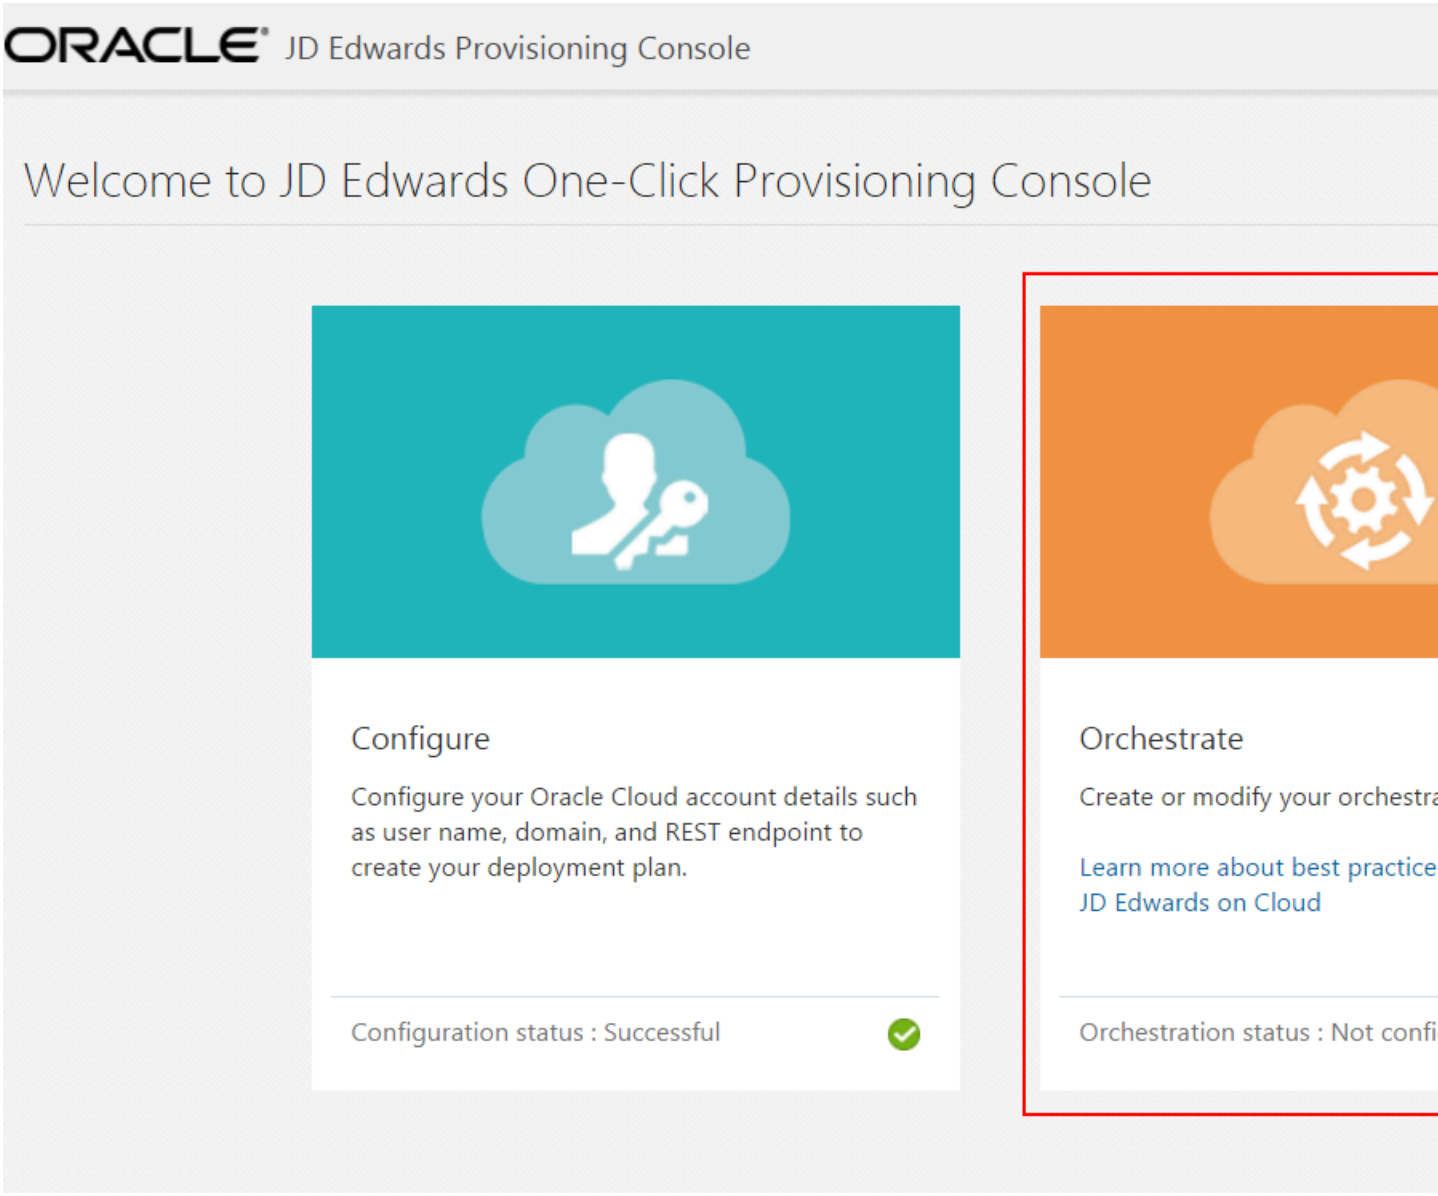

2. If this is the first pass through a JD Edwards One-Click Provisioning orchestration, the following Global Settings screen will appear first. In the SSH Private Key section, click the View/Edit Private Key button.

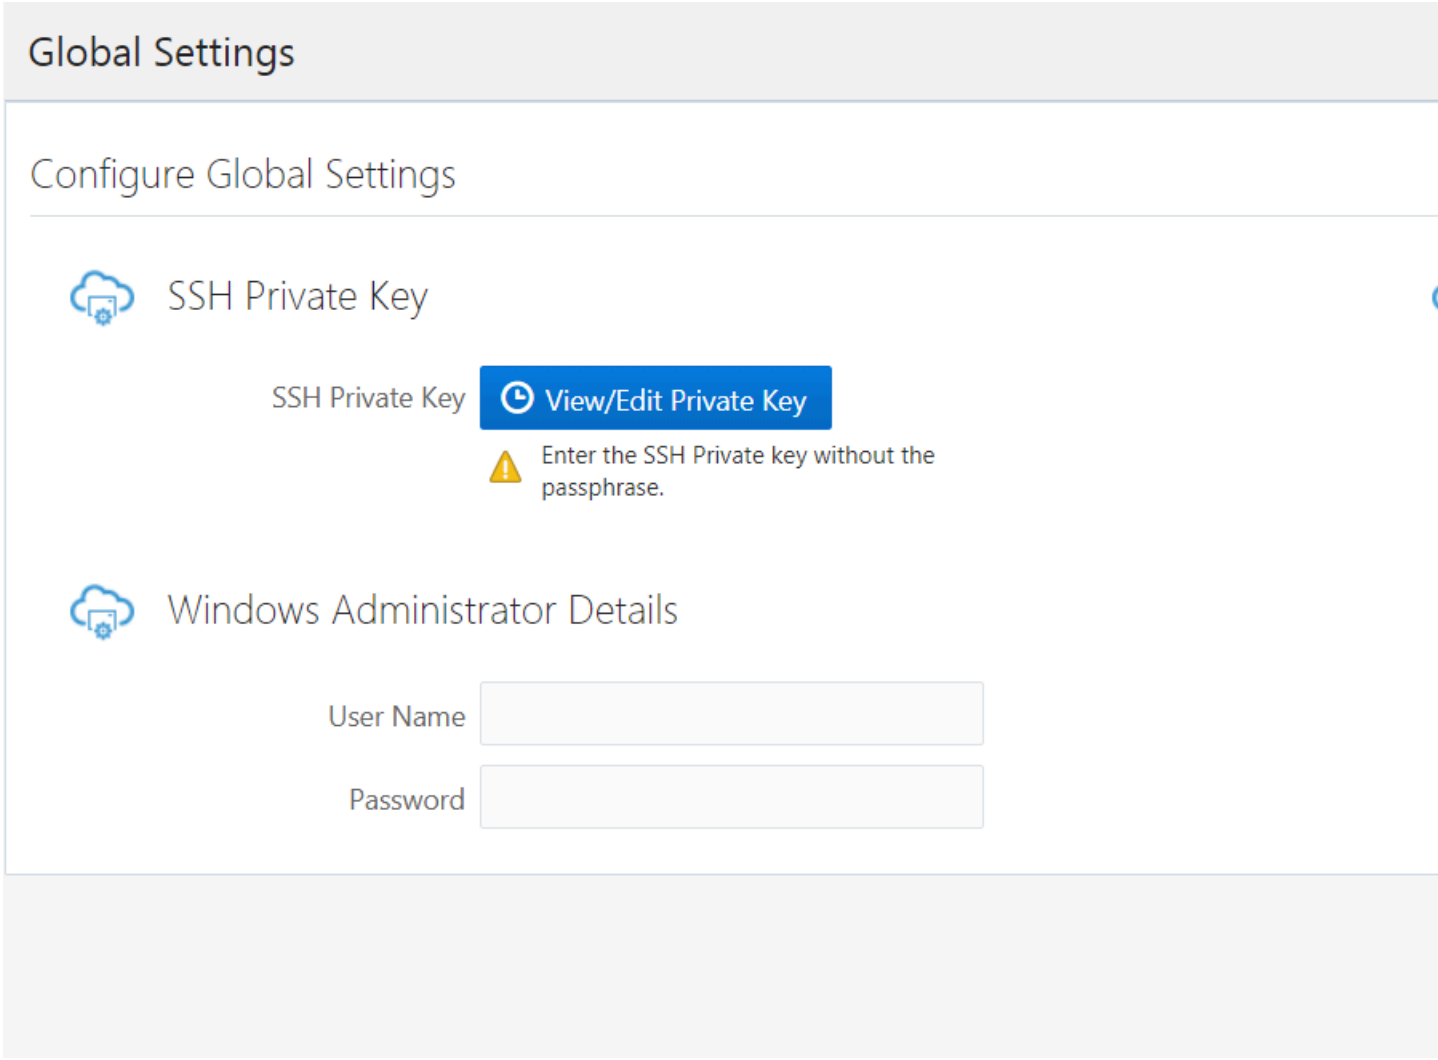

3. To access all the instances that are provisioned using this tool, on the Private Key input for VM access window, you must either specify the values for the SSH private key text, or browse and select the file that contains the SSH private key contents.

   To select a file, select the option SSH Private Key File and then click the Choose File button. For more information regarding SSH Keys, refer to the section "Generate Secure SHell (SSH) Key Pairs on Your Local System" of this Learning Path.

   > **Note:** The One-Click Provisioning Console validates the private keys and you cannot save the Global Settings if this field is empty.

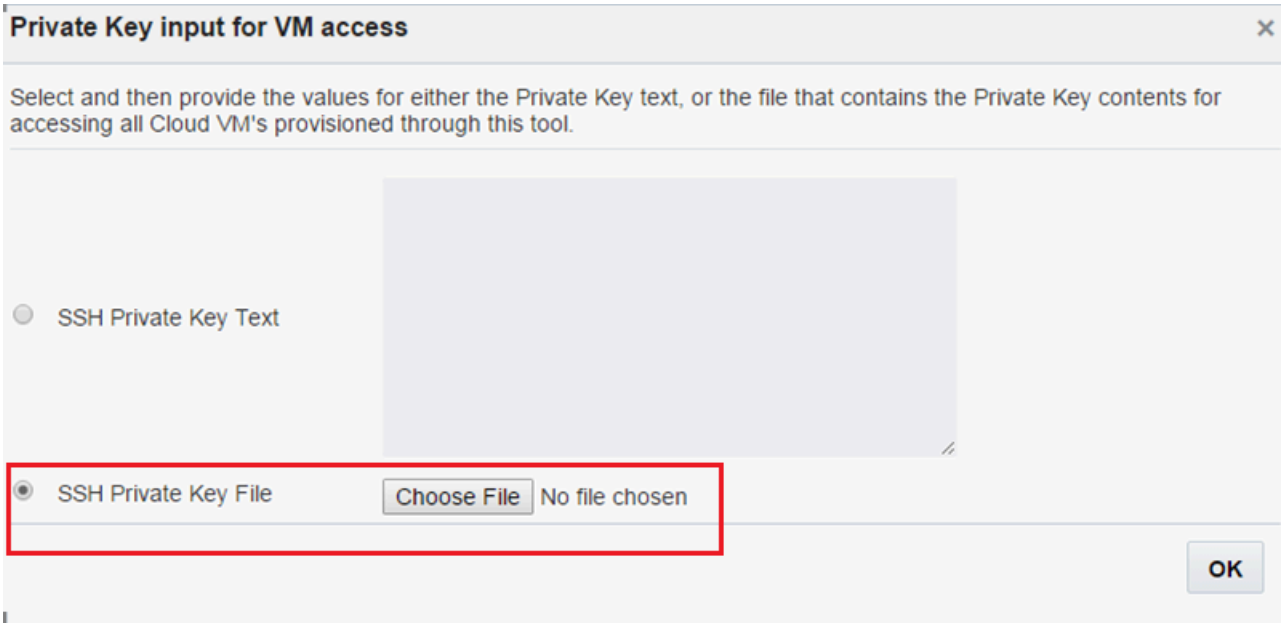

4. Browse and select the appropriate file, and then click OK.

> **Note:** If you click the View/Edit Private Key button again, you can see the Public Key for VM access window with the SSH Public Key value in the SSH Public Key text field. To change the private key, click the Change button and provide the new value.

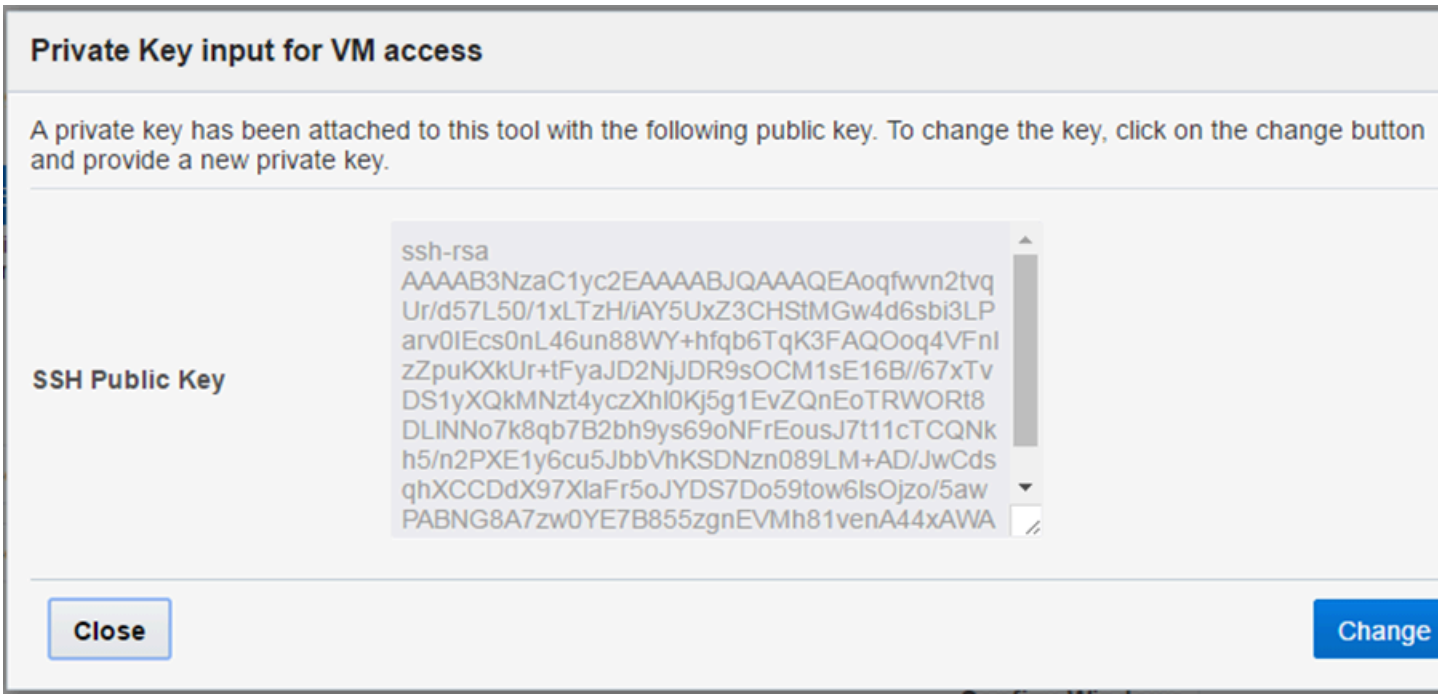

5. Click the Close button.
6. In the Windows Administrator Details section, enter the Windows user name and password. Ensure that the user name is opc and that this user has the administrative privileges.

> **Note:** You must enter the same password for this Windows Server that you previously specified in the section of this OBE entitled: Logging in to the Windows VM.

7. In the Set EnterpriseOne Passwords section, enter and then confirm these passwords
   - **JDE User Password**
     Create the password for JD Edwards EnterpriseOne. In support of the long password functionality, the password must be between 12 and 30 characters. It can contain letters and numbers, and can only include this special character: underscore (_).
   - **Site Key Passphrase**
     Enter the passphrase for generating the site key. The passphrase must start with a letter, end with either a letter or a number, must be between 8 and 40 characters, and contain at least 2 uppercase letters, 2 lowercase letters, 2 numbers, and 2 underscore (_) symbols.

   Note: The conditions to set the passwords appear in a tooltip when you click each field.

   Note: It is highly recommended that you keep a record of these critical passwords. If you have not already done so, you must record these values on the Preinstallation Worksheet.

8. Click the Save button to exit the Global Settings screen.
9. On the Choose your Orchestration Template window, click the Quick Start icon to create your Quick Start JD Edwards deployment plan.

   Note: Refer to the applicable sections of this guide for instructions on using the Advanced, Export, and Import orchestration functions.

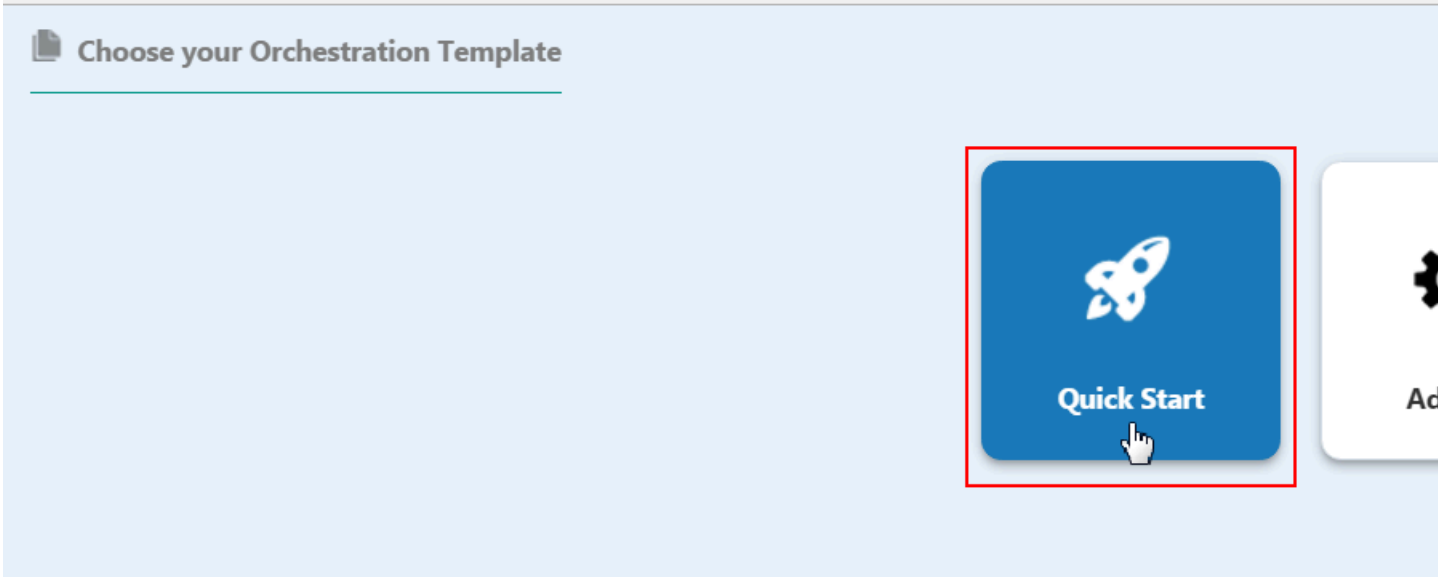

10. On the Database Server Instance window, complete these fields to create and configure the Database Server instance.
    Server Configuration
    - Database Server Type
      The Database Server Type is displayed as Oracle Database.

- ATP-D

  Enable this selector button for ATP-D.
- Platform

  This field is disabled and it is automatically populated as Linux.
- Instance Name

  Create an instance name for your database instance.
- Host Name

  Enter the host name.

Database Configuration

- DB Admin Password

  Enter the password of the database administrator.
- DB Wallet

  Click the Browse button to locate and select the DB Wallet that you created by following the steps in the section "Downloading a Database Wallet for Autonomous Transaction Processing on Dedicated Infrastructure" of this Learning Path.

JD Edwards Database Configuration

- JDE DB Install Directory

  Enter the installation path.

  Path Rules. All directories in the specified path must preexist, except the last directory. Therefore you must manually create the directory structure except for the last directory, which the Provisioning Server

deployment process creates. For example, if you specify /u01/ORCL/INSTALL, the /u01/ORCL directory must preexist and the Provisioning Server deployment creates the /INSTALL directory.

- Schemas

  Click the Schemas field and select the schemas from the auto-suggest text. The schemas available are: Shared, Development, Prototype, Production, and Pristine with Demo Data.

  > **Note:** It is mandatory to add the Shared schema.

  > **Note:** At this point, you should ensure that you specify all the schemas you plan to use. The schemas you choose to install on the Database Server can only be deployed once, which is specified at this point in the Provisioning Console. You can use the Provisioning Console to programmatically add additional schemas after deploying the orchestration.

- Demo Data

  Click the Demo Data field and select the demo data available from the auto-suggest text. Demo data will be available depending on the schema selected. For example, if you select the schema as Development, the Development demo data will be available.

OCI Object Storage Details

- Tenancy

  Enter the tenancy where you have access to the Object Storage Service in Oracle Cloud Infrastructure.
- User Name

  Enter the user name with which you can access the Object Storage Service in Oracle Cloud Infrastructure.
- Auth Token

  Enter the Auth Token for the given user. This token is used to upload the JD Edwards database dump files into the Object Storage Service in Oracle Cloud Infrastructure.

  For additional details, see the section "*Getting an Auth Token for a User*."
- Region

  OCI Region
- Bucket

  Enter the bucket name that you have previously created for use with Oracle Cloud Infrastructure Object Storage Service.

  For additional details, see the section "*Creating an Object Storage Bucket for an Autonomous Database*" in this Learning Path.

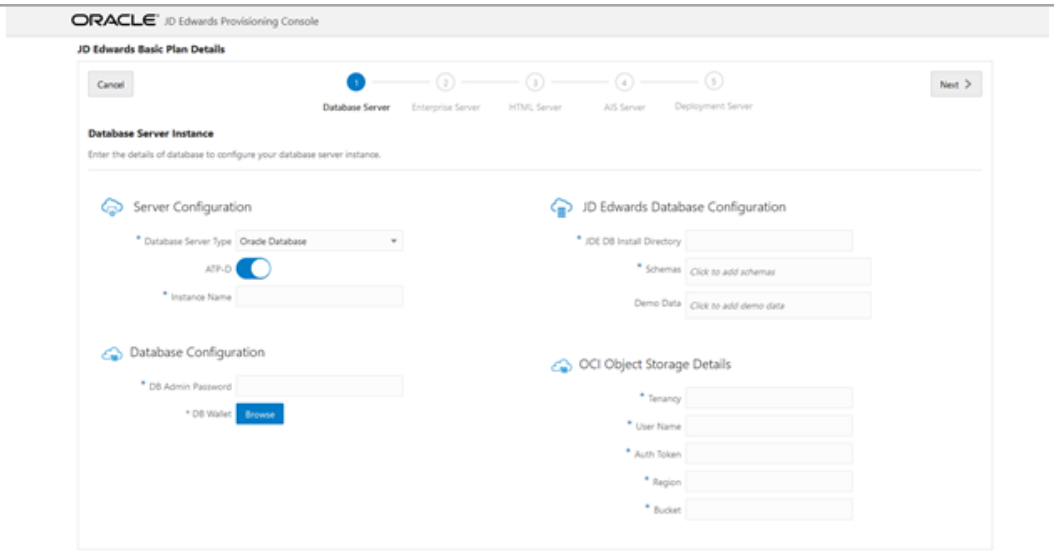

11. Click the Next button. Because you will be connecting to an existing Oracle Cloud Infrastructure Service database instance, the system validates all the inputs that are provided. If the validation is successful, the JD Edwards Enterprise Server page is displayed.

12. In the Enterprise Server Instance section, complete these fields to create and configure the Enterprise Server instance.

Server Configuration

- Platform

  This field is disabled and it is automatically populated as Linux.
- Instance Name

  Create an instance name for the Enterprise Server. The conditions to set the instance name is displayed in the tooltip when you click the field.
- Host Name

  Enter the host name.

Enterprise Server Preferences

- Server Type

  Select one or both of the available Server Types for this Enterprise Server.

  Single Enterprise Server: If you are deploying only a single Enterprise Server, select both logic and batch servers.

  Multiple Enterprise Servers: If you are deploying multiple Enterprise Servers, at least one server must be specified as a logic server for each pathcode. The other servers can be specified as batch servers.
- Pathcodes

  Click the Available Pathcodes field and select the pathcodes from the auto-suggest text. The four available pathcodes are: Development, Prototype, Pristine, and Production.

  > **Note:** It is good practice to select pathcodes here that correlate to the schemas you selected for the Database Server. The Provisioning Console programmatically enforces this correlation. If you select pathcodes on the Enterprise Server that are a superset of the database schemas you selected, the Enterprise Server will not be able to access the data required to function correctly. The pathcodes you choose to install on the Enterprise Server can only be deployed once, which is specified at this point in the Provisioning Console. You can use the Provisioning Console to programmatically add additional schemas after the deploying the orchestration.

- Oracle JDBC Driver Details

  This driver is required for connectivity between the Enterprise Server and the Oracle database server.

  Click the Browse button to select each of the required components for the Oracle JDBC driver. For example:

  - odbc8.jar
  - ons.jar
  - ucp.jar

  > **Note:** Refer to Oracle Certifications for the version of the supported driver and associated components.

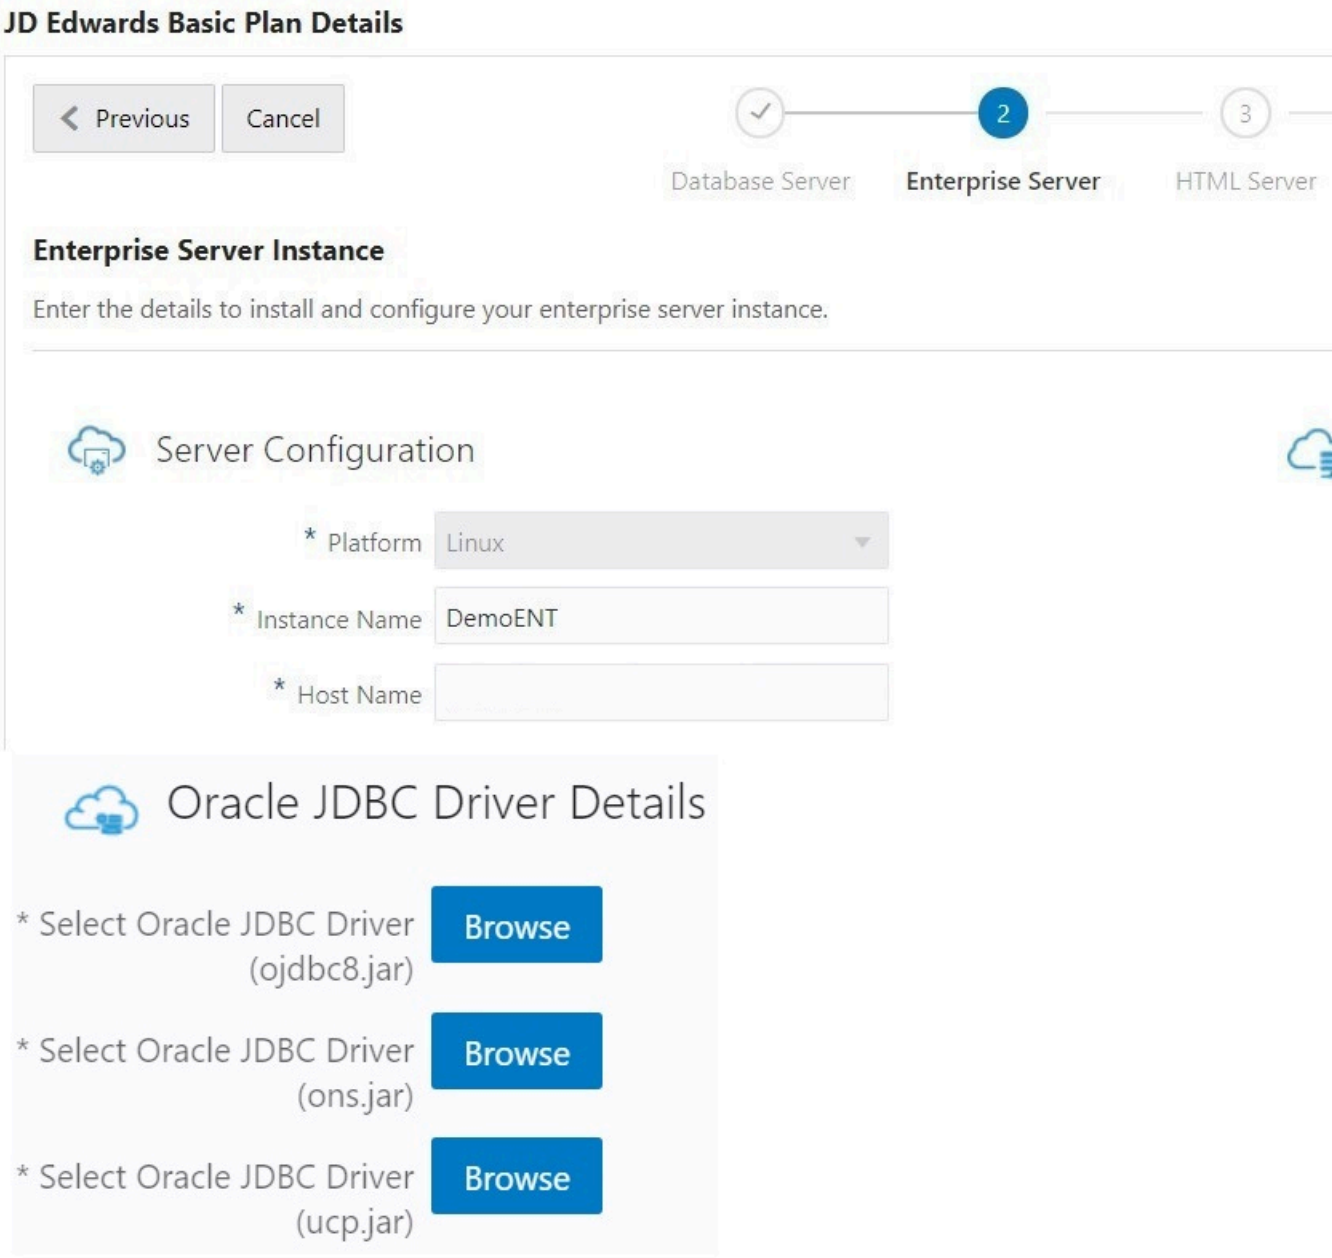

- Click the Next button. Because you will be connecting to an existing Oracle Cloud Infrastructure Service EnterpriseOne instance, the system validates all the inputs that are provided. If the validation is successful, the JD Edwards HTML Server page is displayed.
- In the HTML Server Instance section, complete these fields to create and configure the HTML Server instance.

  Server Configuration

  - Platform

This field is disabled and it is automatically populated as Linux.

- Instance Name

  Create the instance name of the HTML Server instance.

- Host Name

  Enter the host name.

- Port

  For this server, enter a unique (available) port that will use an SSL connection. This port must have a numerical value between 1024 and 65535. This port number is used by HTTPS to create a container and deploy the web component. For the port number you enter here, ensure that the

port for one less is also available. That is, if you specify Port 8081, also ensure that Port 8080 is available.

> **Note:** For each SSL port that you open in the firewall, you must also open a companion port for non-SSL access required for Server Manager. The companion port must have a numeric value that is one less than the value specified for the SSL port. For example, if you specify a port value of 8081 for SSL, in the firewall you must also open a port one less than that value, that is, in this case you must open Port 8080. For more information, refer to the subsection "Enable Inbound Ports in the Firewall" in the OBE "Performing Common Setup for All Microsoft Windows Servers" of this Learning Path.

Web Server Preferences

Pathcode

Select the required pathcode from the drop-down menu.

> **Note:** Using the Quick Start mode, you can only specify a dedicated HTML Server for AIS. If you want to create a standard HTML Server, which is strongly recommended for production environments, you must use the Advanced Deployment mode. For a description of each HTML server type, refer to the"Fundamentals" section of this Learning Path.

> **Note:** Each dedicated HTML and AIS Server pair can support only one pathcode. If you want additional HTML instances to support additional pathcodes, you must configure additional HTML Server pairs using the Advanced mode of the Provisioning Console. For more information, refer to the section "Orchestrating Using Advanced Mode" of this Learning Path.

WebLogic Details

- User Name

  Enter the user name.
- Password

  Enter the WebLogic Server password.
- Admin Port

  Enter the port number to access the WebLogic Administration Console.
- Install Path

  Enter the installation path of the WebLogic instance.
- JDK Install Path

  Enter the JDK installation path.

  > **Note:** Each dedicated HTML and AIS Server pair can support only one pathcode. If you want additional HTML instances to support additional pathcodes, you should configure additional pathcoodes using the Advanced mode of the Provisioning Console. Refer to the section "Orchestrating Using Advanced Mode" of this Learning Path.

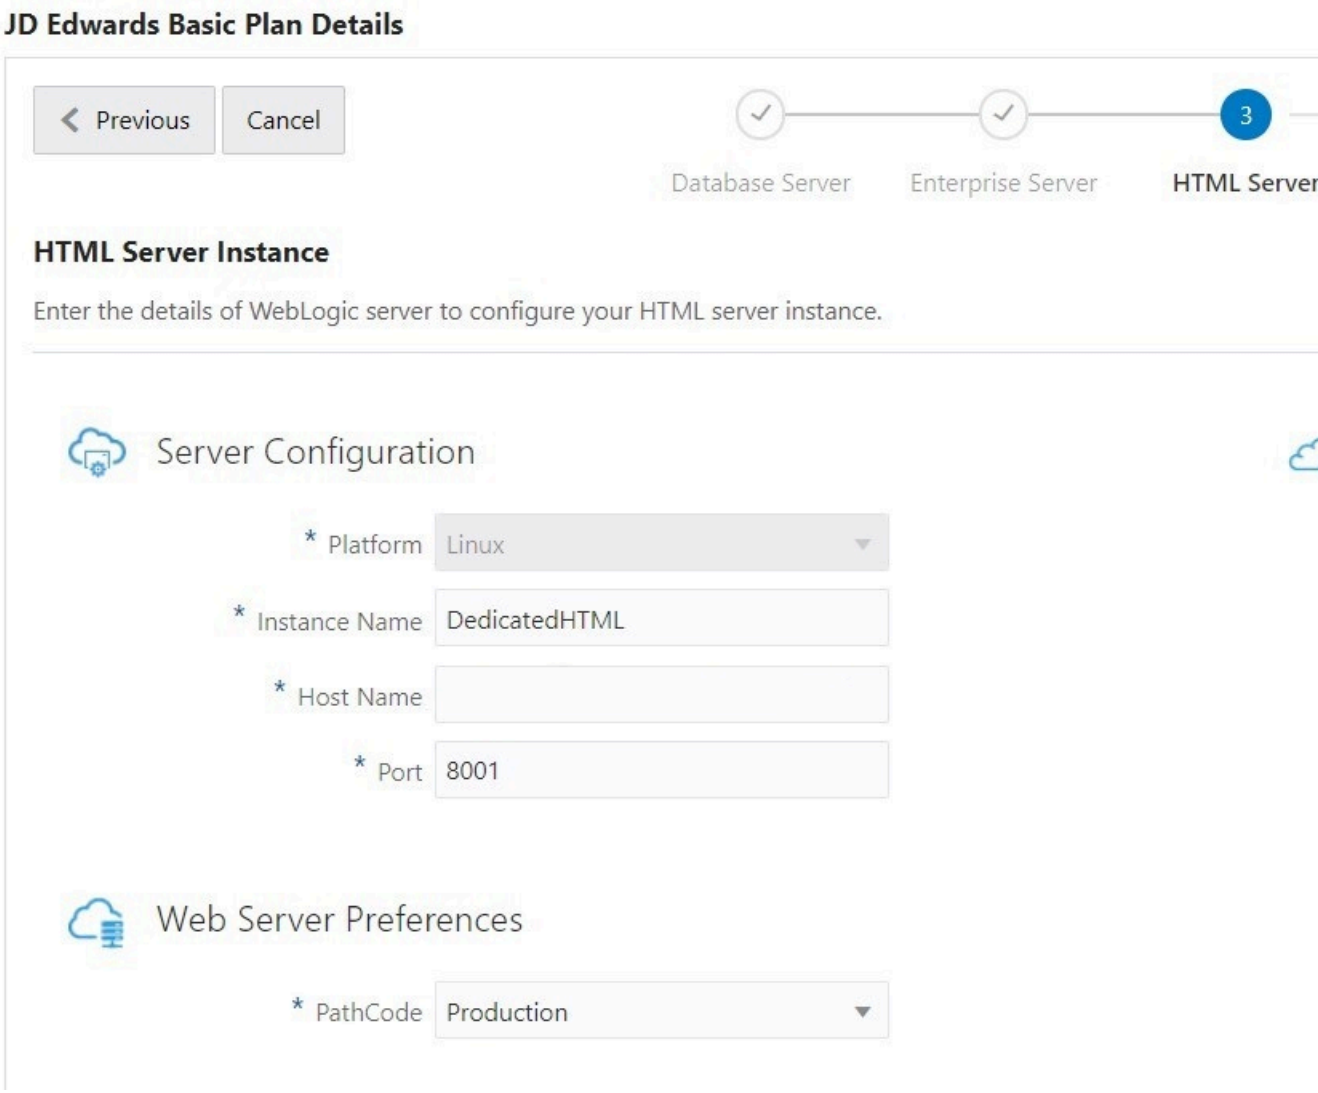

- Click the Next button. The system validates the inputs. If the validation is successful, the AIS Server Instance page is displayed.
- In the AIS Server Instance section, complete these fields to configure your AIS Server instance.

  Same as HTML Server

  This option is selected by default. You cannot deselect it because this AIS Server must be paired with the dedicated HTML Server that was configured in the preceding step.

  Server Configuration

  - Platform

    This field is disabled and it is automatically populated as Linux.
  - Instance Name

Create an instance name for the WebLogic Server.

- Host Name

  Enter the host name.

- Port

  For this server, specify a unique (available) port that will use a SSL connection. This port must have a numerical value between 1024 and 65535. This port number is used by HTTPS to create a container and deploy the web component. For the port number you enter here, ensure that the

port for one less is also available. That is, if you specify Port 8081, also ensure that Port 8080 is available.

> **Note:** For each SSL port that you open in the firewall, you must also open a companion port for non-SSL access required for Server Manager. The companion port must have a numeric value that is one less than the value specified for the SSL port. For example, if you specify a port value of 8081 for SSL, in the firewall you must also open a port one less than that value; in this case you must open Port 8080. Refer to the subsection "Enable Inbound Ports in the Firewall" in the OBE "Performing Common Setup for All Microsoft Windows Servers" for this Learning Path.

Web Server Details

- Type

  This field is disabled and is automatically populated as AIS Server.
- HTML Server Instance

  Use the drop-down menu to choose the dedicated HTML Server for AIS for this pod.

WebLogic Details

- User Name

  Enter the user name.
- Password

  Enter the WebLogic password.
- Admin Port

  Enter the admin port number.
- Install Path

  Enter the installation path of the WebLogic instance.
- JDK Install Path

  Enter the JDK installation path.

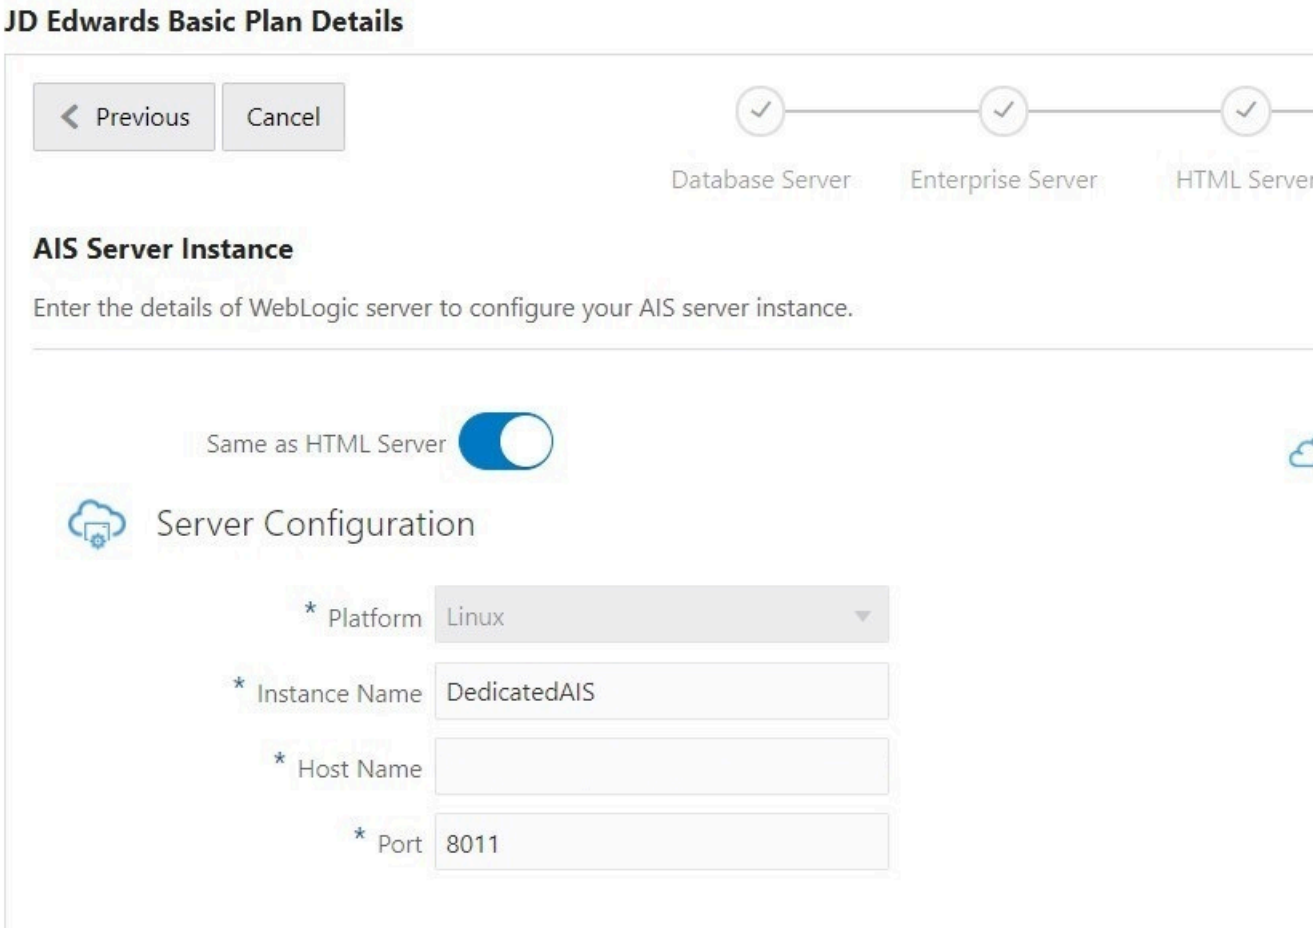

- Click the Next button. The system validates the inputs. If the validation is successful, the JD Edwards Deployment Server page is displayed.
- On the JD Edwards Deployment Server page, complete these fields to create and configure your Deployment Server instance.

Server Configuration

- Instance Name

  Create an instance name for the Deployment Server instance. The conditions to set the instance name is displayed in the tooltip when you click the field.
- Host Name

  Enter the host name.
- Windows User

  Enter the name of the windows user.
- Windows Password

  Enter the password of the Windows user.

- Location

  Enter the location.

  This value is the base location for your JD Edwards EnterpriseOne machines. For example, typical values might be a city name (such as Denver or Austin), a geographical region name (such as US or India), or a general location name (such as Corporate).
- Installation Drive

  Enter the drive for the installation.
- Pathcodes

  This field is automatically populated.

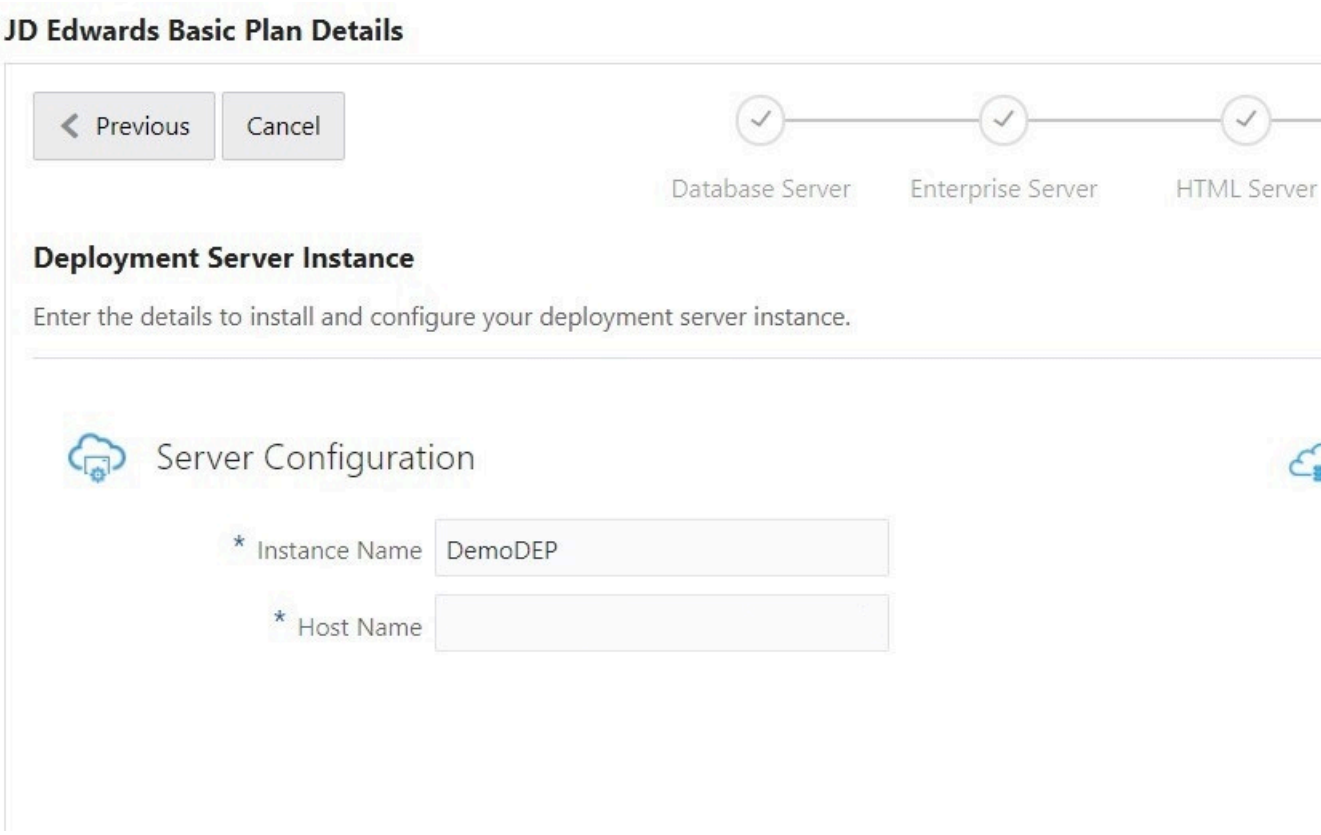

- Click the Finish button.

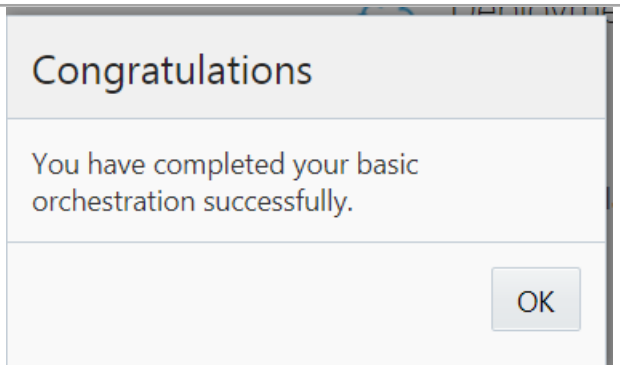

- Click the OK in the success message window.
- The system displays the summary window. In the summary window, you can see that one instance is created for each of the servers. Click the Back button to deploy your services by following the steps described in the section "Deploy an Orchestration" of this Learning Path.

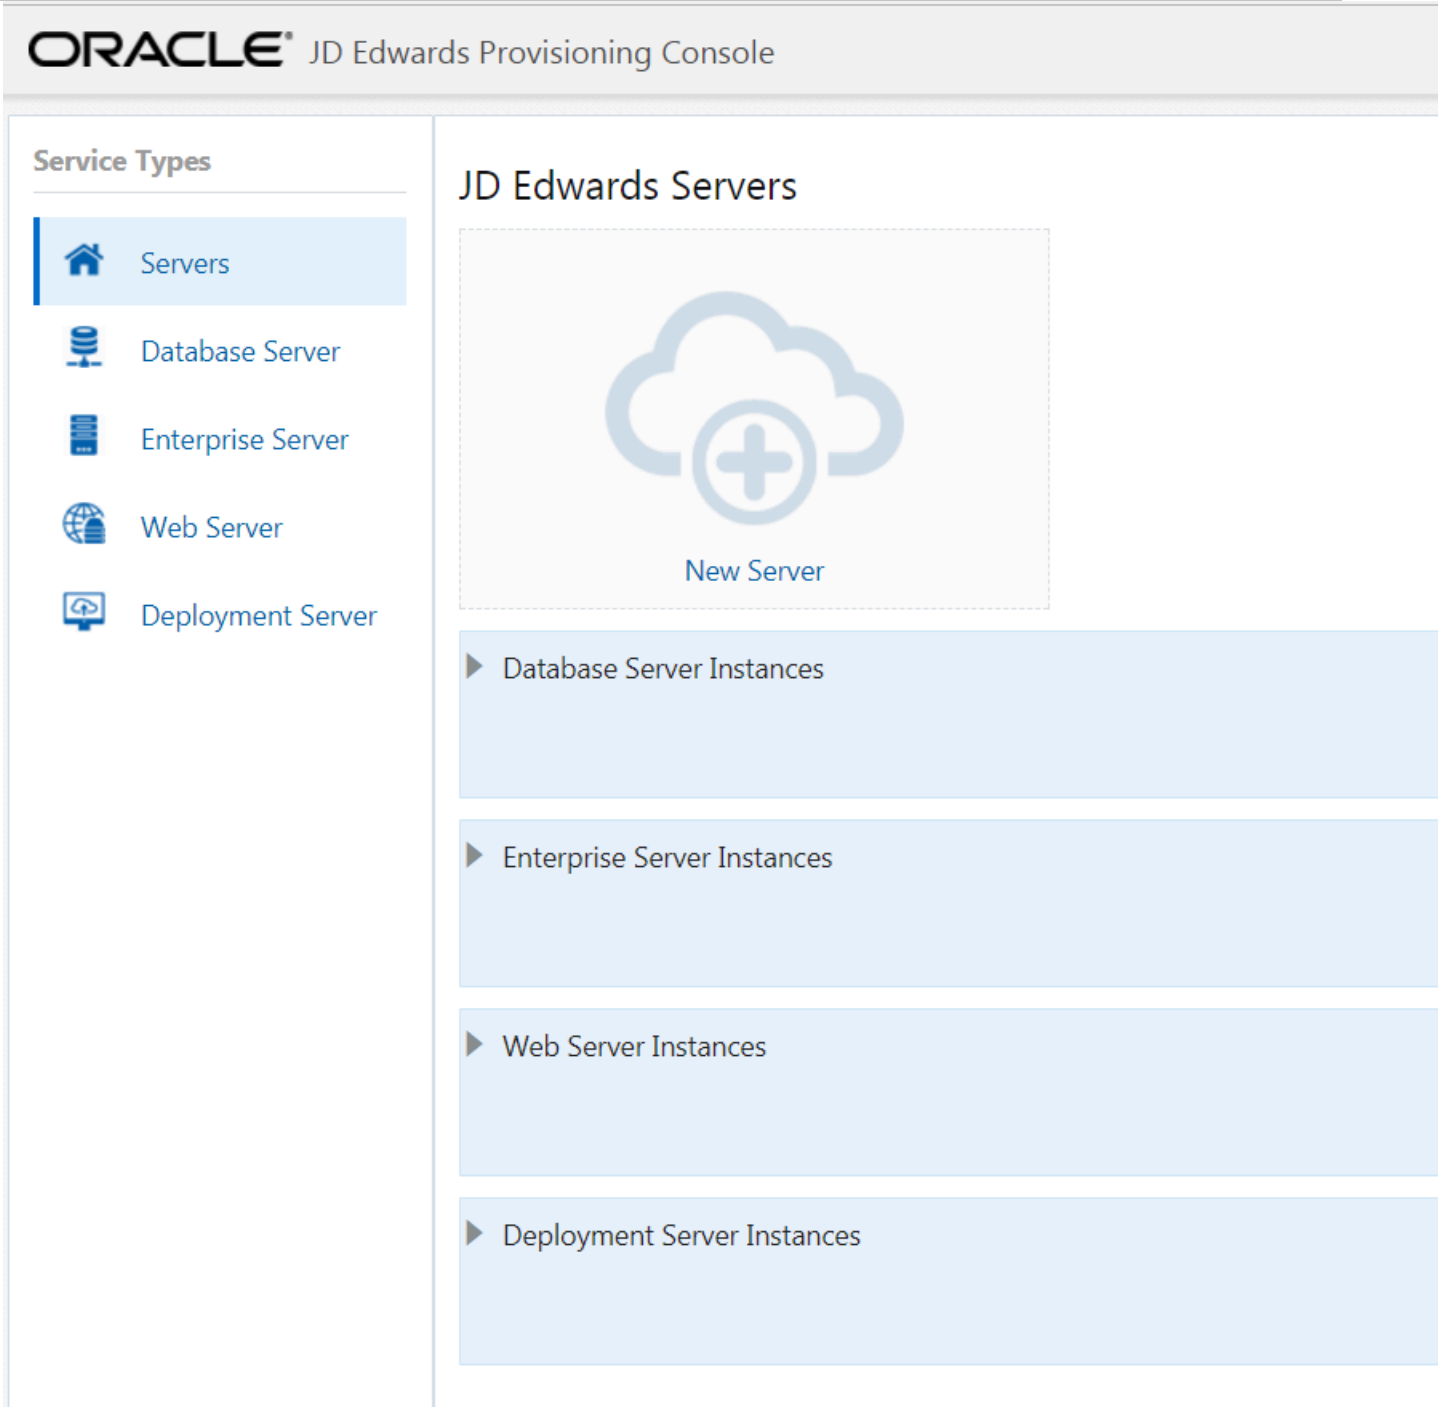
ORACLE JD Edwards Provisioning Console
Service Types
Servers
Database Server
Enterprise Server
Web Server
Deployment Server
JD Edwards Servers
New Server
Database Server Instances
Enterprise Server Instances
Web Server Instances
Deployment Server Instances

> **Note:** For this Quick Start deployment plan, you can click Options from the menu bar to change your global settings or to reset your settings (that is, to delete your configuration details, global settings, and orchestration data). Also optionally, you can further customize a completed Quick Start deployment plan using the Advanced mode of the JD Edwards Provisioning Console. Refer to the section "Orchestrate Using Advanced Mode" of this Learning Path.

Orchestrating Using Advanced Mode

This section section shows how to orchestrate a deployment plan using the advanced mode on Linux using the JD Edwards One-Click Provisioning Console.

You can orchestrate a Deployment Plan using Advanced mode, which allows you to:

- Modify and add servers to an existing Orchestration that was created using the Quick Start wizard.
- Create a new Orchestration with no guidance from a wizard. Unlike using the Quick Start wizard, the Advanced mode will ***not*** guide you through the required sequence of machines to satisfy dependencies of the configuration. For example, the allowable pathcodes for an Enterprise Server are dependent on the schemas previously selected for the Database Server.

Prerequisite

- You must have configured the administrator passwords for WebLogic Server and Server Manager Console in the Configure section of the JD Edwards One-Click Provisioning Console.
- Before orchestrating an Advanced Deployment Plan, the recommended practice is to first orchestrate a basic Deployment Plan using the Quick Start mode. The Quick Start procedure is documented in the section of this document entitled: Orchestrate Using Quick Start Mode. After this Quick Start orchestration is created, you can use the Advanced mode to modify, delete, or add on additional server instances.
- Alternately, experienced users can use Advanced Mode to create an orchestration without assistance from a structured wizard, which assumes they are aware of the required sequence of creation and the inter dependencies.

Orchestrating an Advanced Deployment Plan

> **Note:** The following procedure explains how to modify and add servers to an existing orchestration that was created using the Quick Start mode.

The required sequence of adding instances to an orchestration and the rules related to the machines and pathcodes is as follows:

1. **Database Server**

   You can create five database instances if you select one schema per instance.
2. **Enterprise Server**

   One to many Enterprise Servers can be created. At least one Enterprise Server must be created with selected pathcodes available from the available schemas that were selected for the Database Server. An Enterprise Server must be configured before you can add HTML Servers. If multiple Enterprise Servers are deployed, at

least one must be configured as a Logic Server per pathcode. Additional Enterprise Servers can be added to an Orchestration and they can be deployed after your initial plan is deployed.

3. **Web Servers**

   HTML Server. Web Servers include the JD Edwards EnterpriseOne HTML Server. There are two types of HTML Servers: **Standard JAS** and **Dedicated HTML for AIS**.

   One to many HTML Server instances can be created and each will be associated with one specific pathcode that is available on the Enterprise Server. Additional HTML Servers can be added to an Orchestration and they can be deployed after your initial plan is deployed.

   Refer to the *Fundamentals* section of this Learning Path for a description of each type of HTML Server.

   AIS Server. Another type of Web Server for JD Edwards EnterpriseOne is the AIS Server, which must be installed and configured along with a Dedicated HTML Server for AIS.

   Zero to many AIS Server instances can be created and each will be associated with a specific HTML Server instance. Additional AIS Server instances can be added on to your plan and deployed after your initial plan is deployed.

   > **Note:** If you do not specify at least one AIS Server, the full functionality of certain JD Edwards EnterpriseOne applications will not be available. If you used the Quick Start mode to configure a basic environment, the workflow required the inclusion of this server.

   > **Note:** You cannot specify a single HTML Server instance to also support an AIS Server; that is, you cannot combine two servers in the same instance. You must create a separate instance for each AIS Server.

4. **Deployment Server**

   A single Deployment Server can be created and all pathcodes can be selected regardless of pathcodes selected for your runtime servers.

   > **Note:** For any orchestration created or modified using the Advanced Mode Deployment Plan, you can click the **Options** function from the menu bar to change your Global Settings or to Reset your settings (that is, to delete your configuration details, global settings, and orchestration data)

   .

To use the JD Edwards Provisioning Console to orchestrate an Advanced Deployment Plan:

1. On the JD Edwards Provisioning Console, click the **Orchestrate** icon.

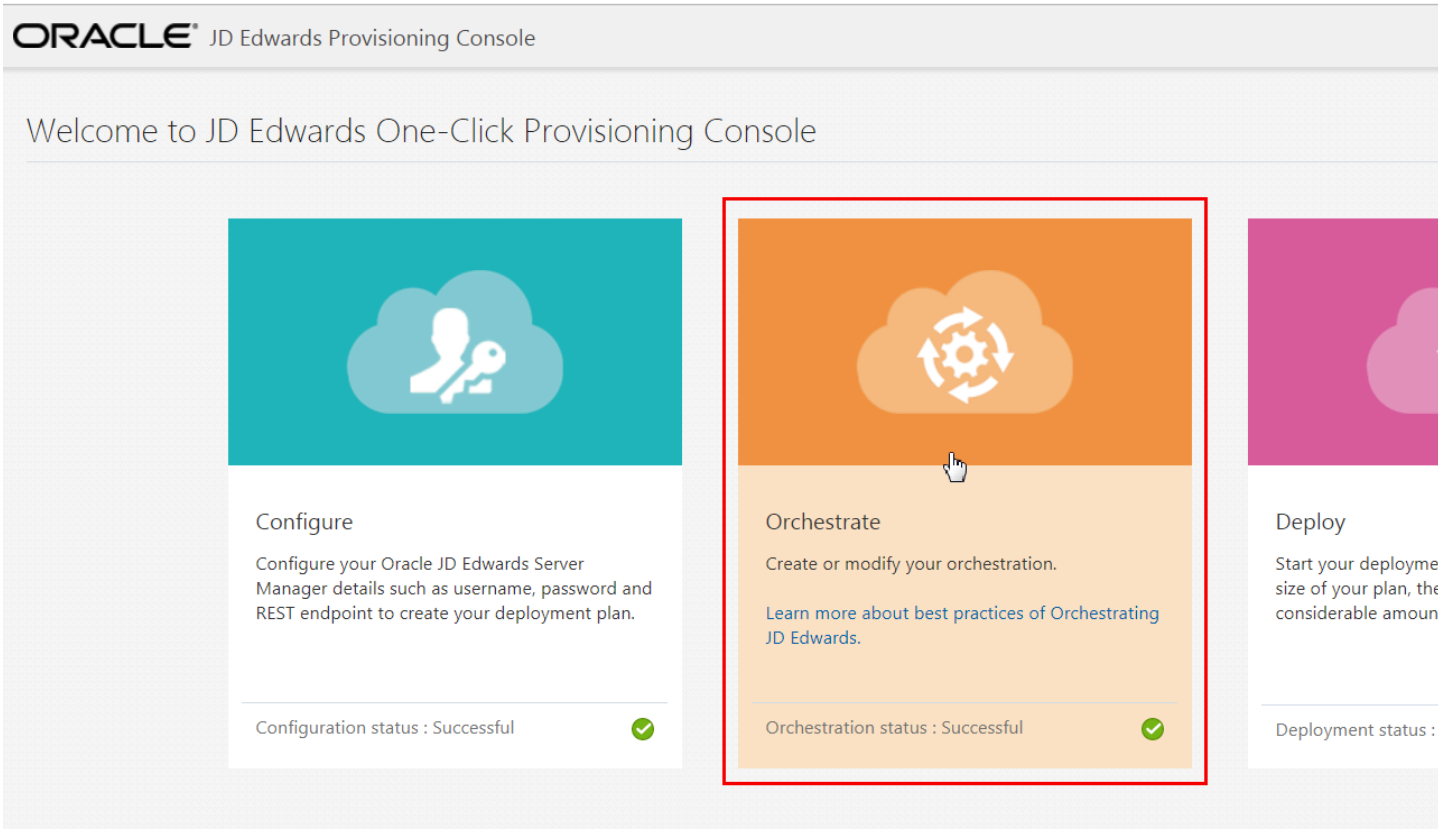

2. On JD Edwards Servers, click on the instance of an existing server, or click the **New Server** icon to add a new JD Edwards service. Alternately you can select the desired server from **Service Types** on the left tab, and then click **New Server** to add a New Server.

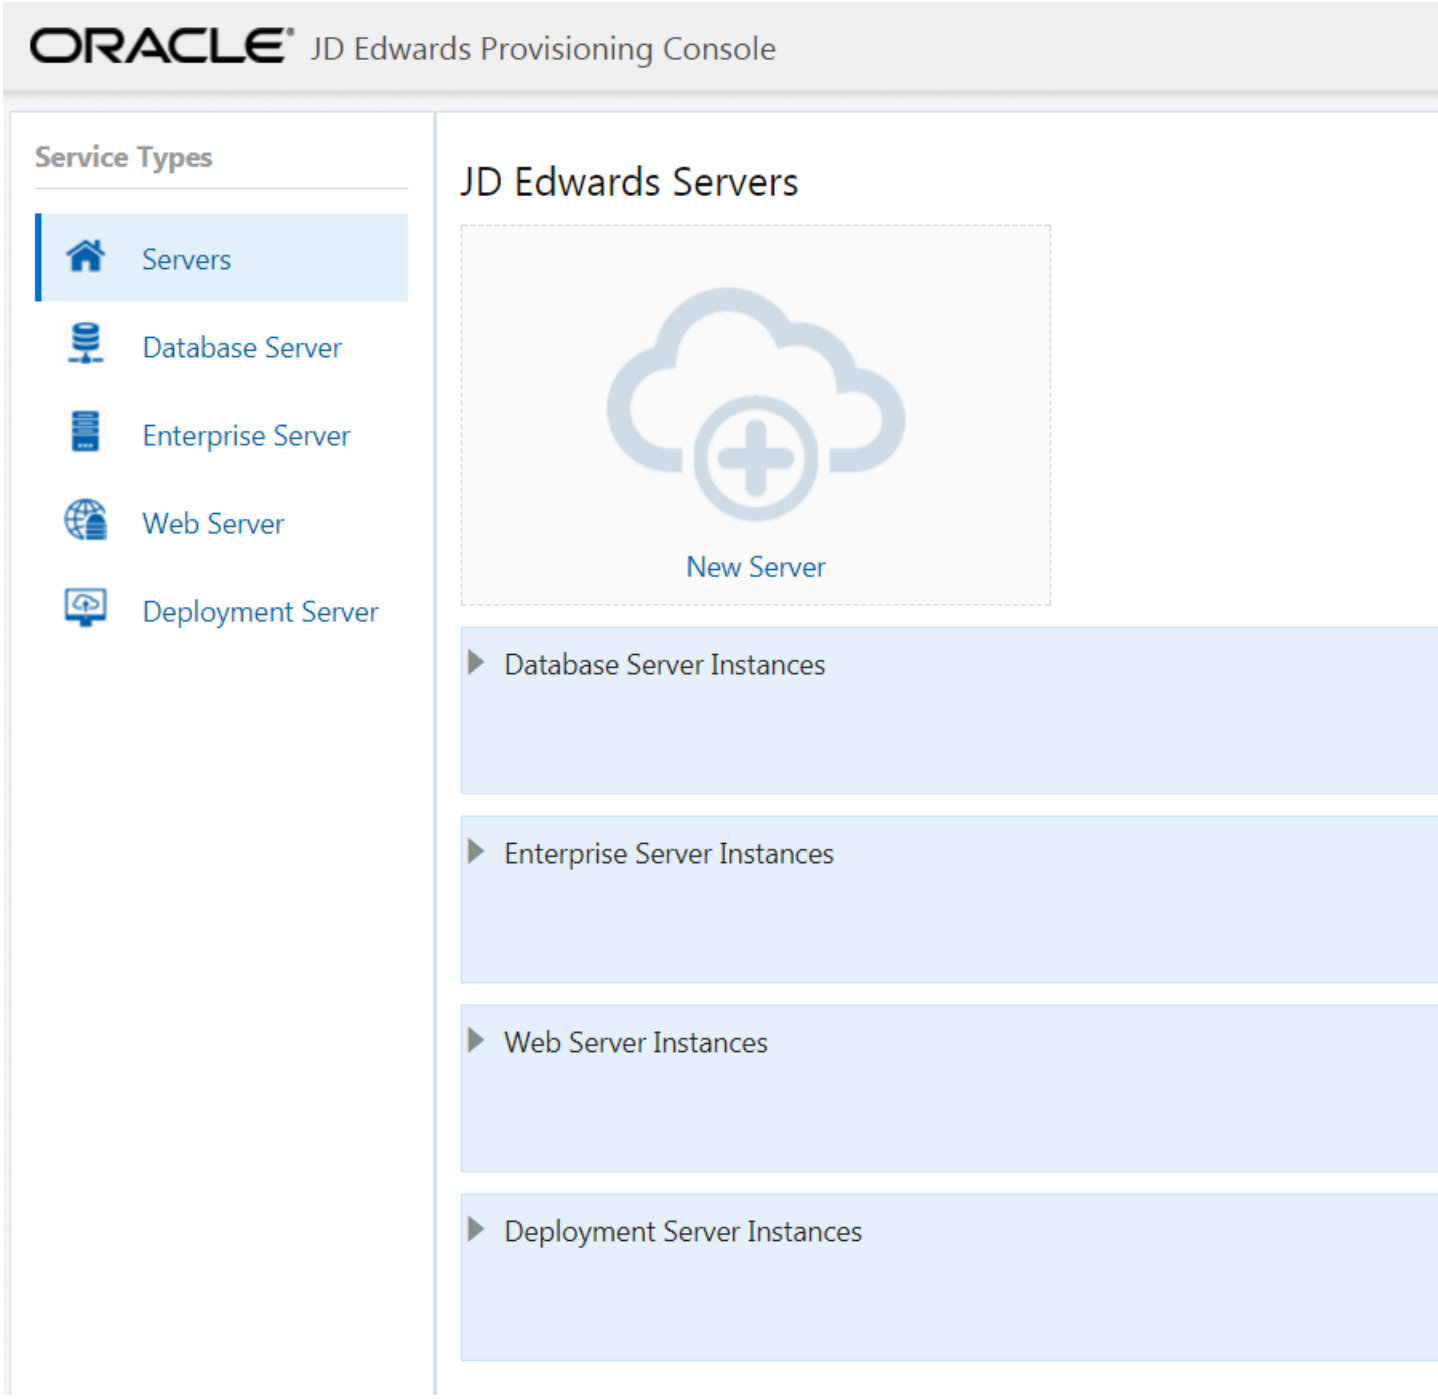

Note: The remainder of this procedure includes steps to either add (where allowed) or modify (existing) these instances:

- Database Server
- Enterprise Server
- Web Server (for example, HTML Server and AIS Server)
- Deployment Server*

* The Provisioning Console will not allow you to add more than one of these server types because only one server of this type is supported per deployment.

Database Server

You can create five database instances if you select one schema per instance. For example, if you select Development and Shared schemas for an instance, you can create three more database instances. For one deployment instance, you can have a maximum number of five schemas distributed across one or many database instances.

You can provision the Database Server instance with the available schemas as required. The following schemas are available for the database instance:

- Development
- Shared (required)
- Prototype
- Production
- Pristine

To add a Database Server:

1. Click the **New Server** icon and select Database Server from the Select a Server Component window.

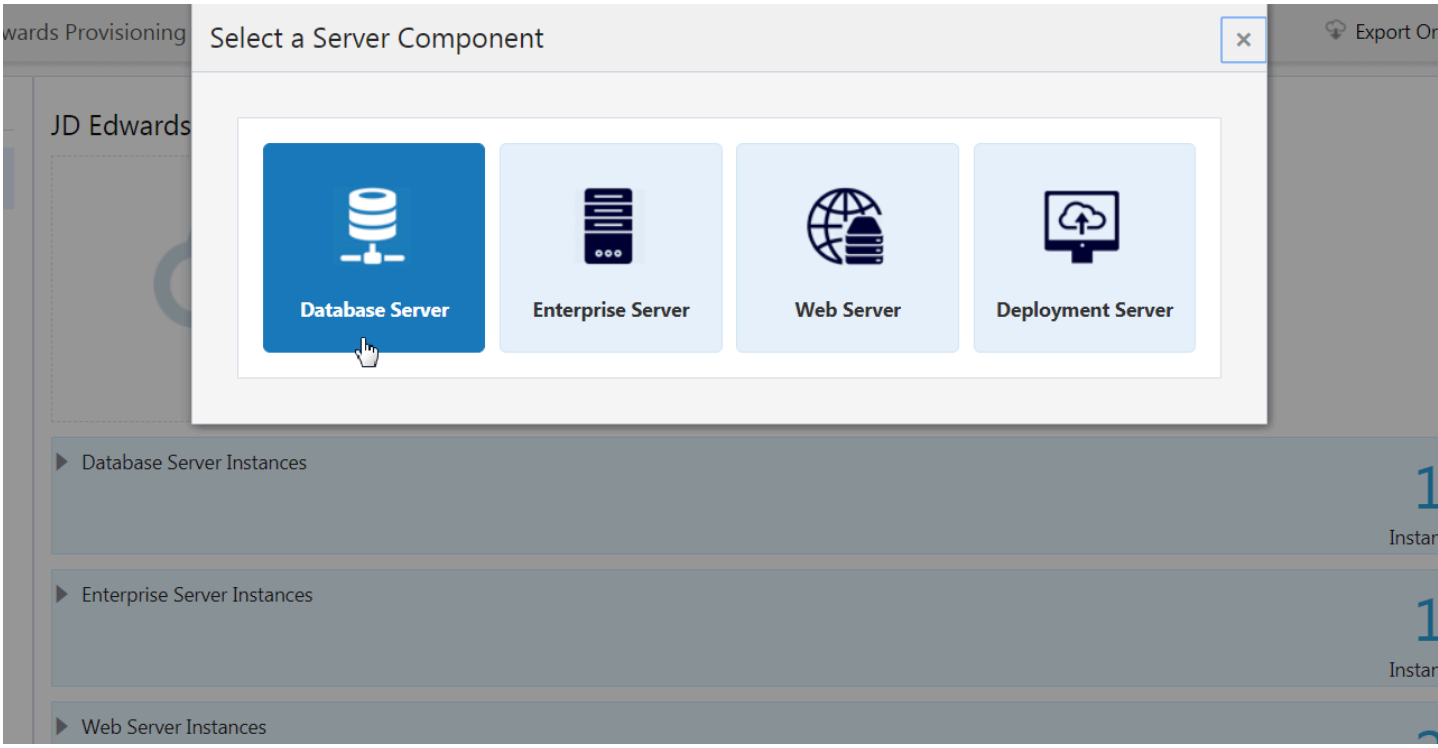

2. On Database Server Instance page, complete the following fields to create and configure the Database Server instance.

   The below sections and supported screens are separated for Standard Oracle Database and Oracle Autonomous Database.

   Standard Oracle Database

   Server Configuration

   - *Database Server Type*

     The Database Server Type is populated by default as Oracle Database.
   - *ATP-D*

     You should only enable the ATP-D option if you are using an Oracle Autonomous Database Dedicated; this database is only supported in Oracle Cloud Infrastructure. This functionality is described in the

Learning Path "Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with Autonomous Database."

A description of the fields specific to Autonomous Database are presented below the first figure below.

- *Platform*

  This field is disabled and it is automatically populated as Linux.
- *Instance Name*

  Create an instance name for your database instance.
- *Host Name*

  Enter the host name.

Database Configuration

- *DB Install Path*

  Enter the DB installation path.
- *DB Admin Password*

  Enter the password of the database administrator.
- *Net Service Name*

  Enter the net service name.

JD Edwards Database Configuration

- *Use ASM feature*

  Enable this option if you are using RAC DB as your Database Server.

  If you disable ASM in your Orchestration, you must enter valid values for the install, table, and index directories for your Oracle database. For example:

  - `/u01/DataDB`
  - `/u01/ORATABLE`
  - `/u01/ORAINDEX`

  If you enable ASM in your Orchestration, you must enter valid values for your DISK group. By default the values for are assumed to be DATA. Otherwise, you can enter any other name that you have created. An example screen is shown below as Database Server Instance - ASM Enabled.
- *JDE DB Install Directory*

  Enter the installation path.

  Path Rules. All directories in the specified path must preexist, ***except*** the last directory in the path. Therefore you must manually create the directory structure except for the last directory, which the Provisioning Server deployment process creates. For example, if you specify /u01/ORCL/INSTALL,

the /u01/ORCL directory must preexist and the Provisioning Server deployment creates the /INSTALL directory.

- *JDE DB Table Directory*

  Enter the path to install the table data.

  Path Rules: All directories in the specified path must preexist, ***except*** the last directory in the path. Therefore you must manually create the directory structure except for the last directory, which the Provisioning Server deployment process creates. For example, if you specify /u02/ORCL/TABLE, the /u02/ORCL directory must preexist and the Provisioning Server deployment creates the /TABLE directory.

- *JDE DB Index Directory*

  Enter the path to install the indexes.

  Path Rules: All directories in the specified path must preexist, ***except*** the last directory in the path. Therefore you must manually create the directory structure except for the last directory, which the

Provisioning Server deployment process creates. For example, if you specify /u03/ORCL/INDEX, the /u03/ORCL directory must preexist and the Provisioning Server deployment creates the /INDEX directory.

- *Schemas*

  Click the Schemas field and select the schemas you want from the auto-suggest text. The schemas available are: Shared, Development, Prototype, Production, and Pristine with Demo data.

  > **Note:** Note: It is mandatory to add the Shared schema.

  > **Note:** At this point, you should ensure that you specify all the schemas you might plan to use. The schemas you choose to install on the Database Server can only be deployed once, which is specified at this point in the Provisioning Console. You can use the Provisioning Console to programmatically add additional schemas after deploying the orchestration.

- *Demo Data*

  Click the **Demo Data** field and select the demo data from the auto-suggest text. Demo data is available depending on the schema selected. For example, if you select the schema as Development, the Development demo data will be available.

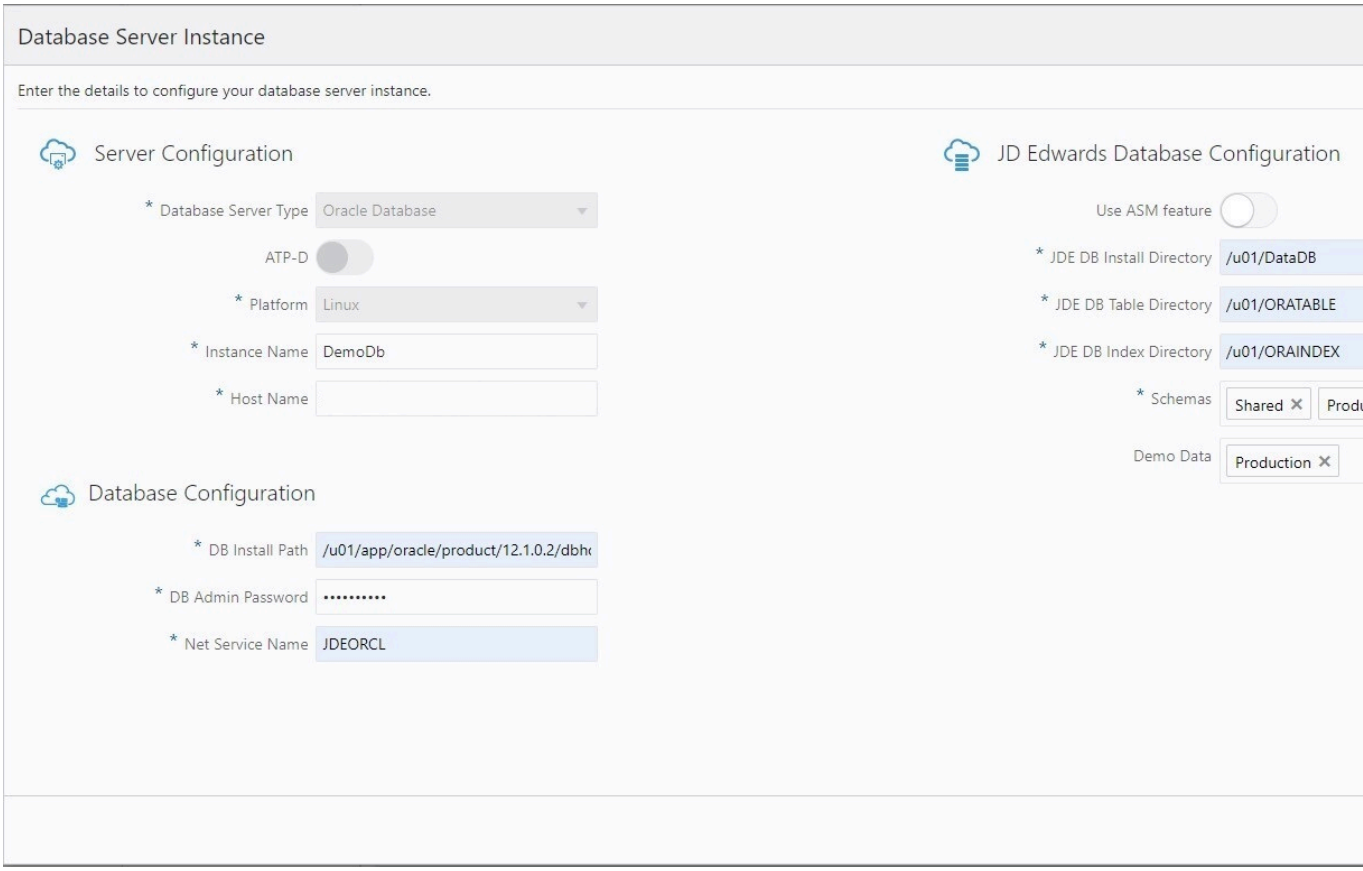

Oracle Autonomous Database

Server Configuration

- *Database Server Type*

  The Database Server Type is displayed as Oracle Database.
- *ATP-D*

  Enable this selector button for Oracle Autonomous Database ATP-D.
- *Platform*

  This field is disabled and it is automatically populated as Linux.
- *Instance Name*

  Create an instance name for your database instance.
- *Host Name*

  Enter the host name.

Database Configuration

- *DB Admin Password*

  Enter the password of the database administrator.
- *DB Wallet*

  Click the **Browse** button to locate and select the DB Wallet that you created by following the steps in the section "Downloading a Database Wallet for Autonomous Transaction Processing on Dedicated Infrastructure" of this Learning Path.

JD Edwards Database Configuration

- *JDE DB Install Directory*

  Enter the installation path.

  Path Rules. All directories in the specified path must preexist, except the last directory. Therefore you must manually create the directory structure except for the last directory, which the Provisioning Server

deployment process creates. For example, if you specify /u01/ORCL/INSTALL, the /u01/ORCL directory must preexist and the Provisioning Server deployment creates the /INSTALL directory.

- *Schemas*

  Click the Schemas field and select the schemas from the auto-suggest text. The schemas available are: Shared, Development, Prototype, Production, and Pristine with Demo Data.

  > **Note:** Note: It is mandatory to add the Shared schema.

  > **Note:** point, you should ensure that you specify all the schemas you plan to use. The schemas you choose to install on the Database Server can only be deployed once, which is specified at this point in the Provisioning Console. You can use the Provisioning Console to programmatically add additional schemas after deploying the orchestration.

- *Demo Data*

  Click the **Demo Data** field and select the demo data available from the auto-suggest text. Demo data will be available depending on the schema selected. For example, if you select the schema as Development, the Development demo data will be available.

OCI Object Storage Details

- *Tenancy*

  Enter the tenancy where you have access to the Object Storage Service in Oracle Cloud Infrastructure.
- *User Name*

  Enter the user name with which you can access the Object Storage Service in Oracle Cloud Infrastructure.
- *Auth Token*

  Enter the Auth Token for the given user. This token is used to upload the JD Edwards database dump files into the Object Storage Service in Oracle Cloud Infrastructure.

  For additional details, see the section *Getting an Auth Token*
- *Region*

  OCI Region
- *Bucket*

  Enter the bucket name that you have previously created for use with Oracle Cloud Infrastructure Object Storage Service.

  For additional details, see the section "Managing Buckets" in this Learning Path.

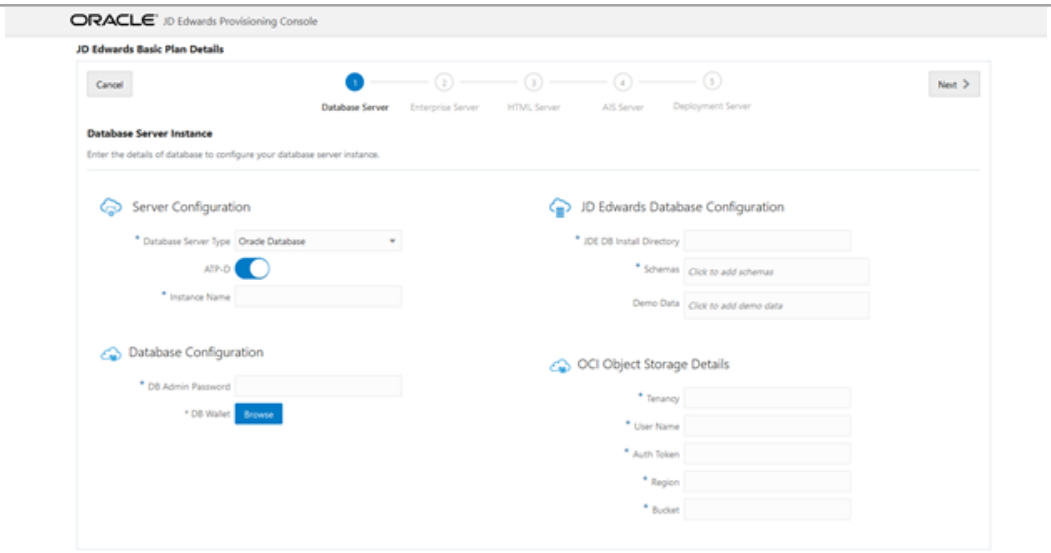

3. You should now be able to view multiple instances of the Database Server.

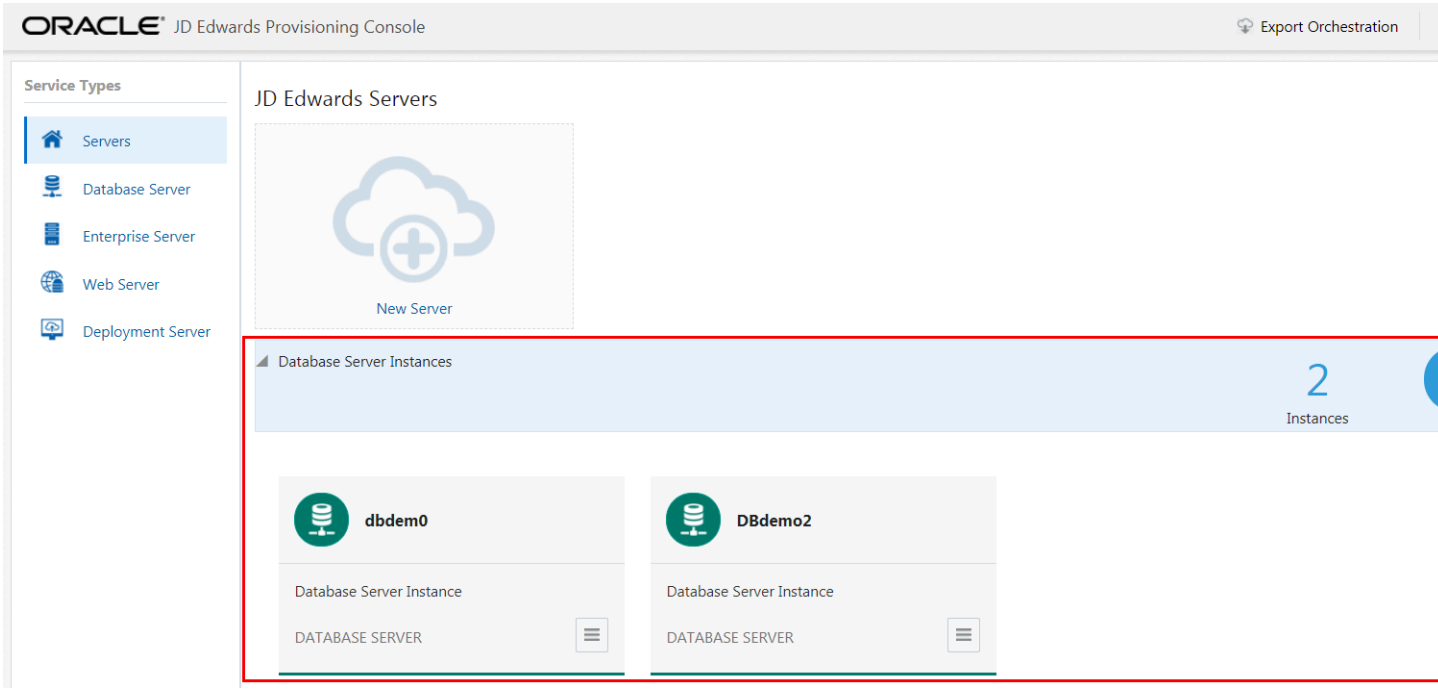

4. If a Database Server exists, click **View** from the Application Options tab to view the existing configuration for the Database Server. To modify the instance configuration use the Application Options tab and choose the **Edit** option.

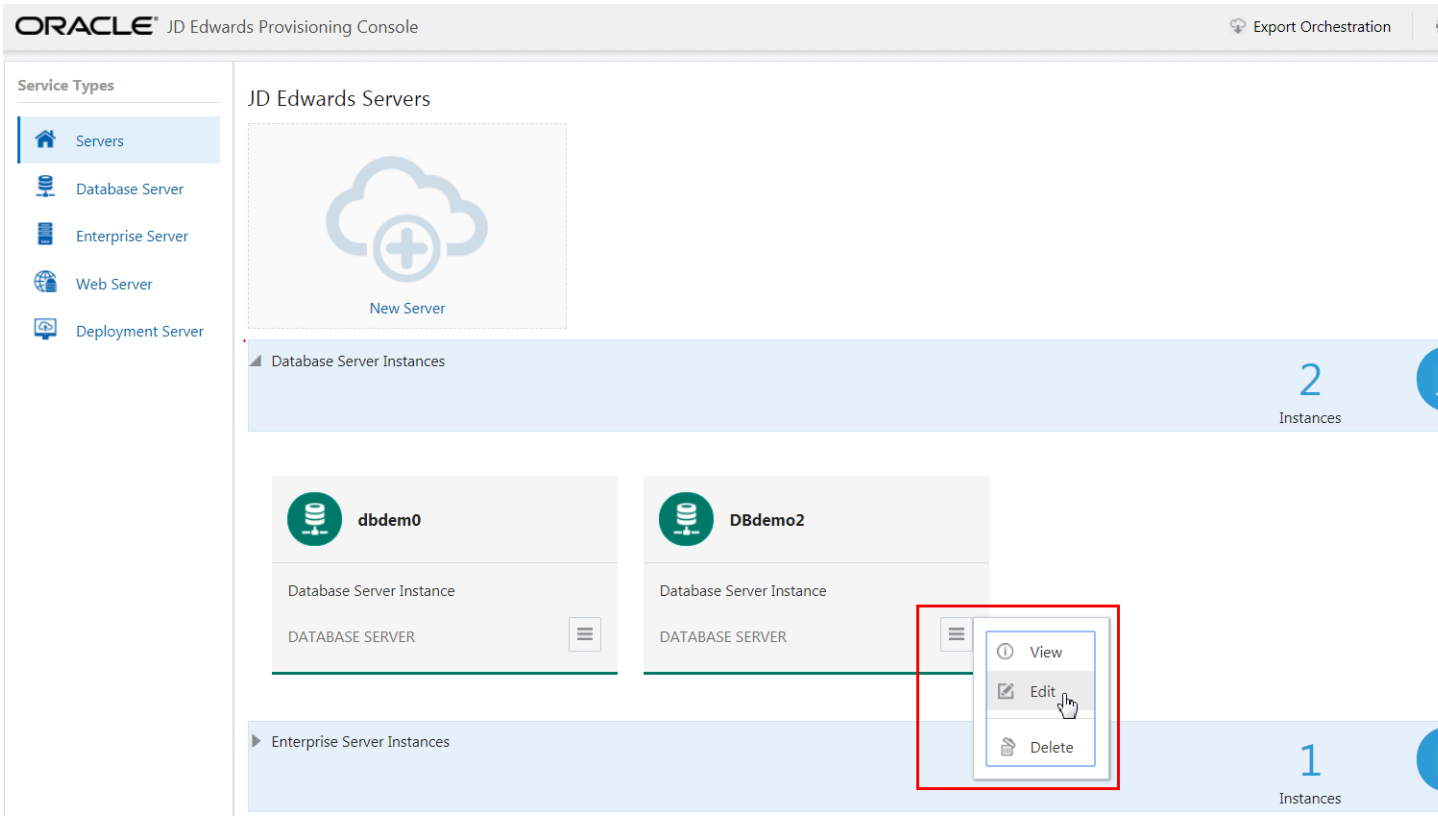

5. To delete the instance, use the Application Options tab to select **Delete.**

Enterprise Server

You can define any number of Enterprise Server instances. If you only want one Enterprise Server, you should define it to run both Logic and Batch. If you want to define multiple Enterprise Servers, at least one must be a Logic Server per pathcode.

To add an Enterprise Server:

1. Click the **New Server** icon and select **Enterprise Server** from the Select a Server Component window.

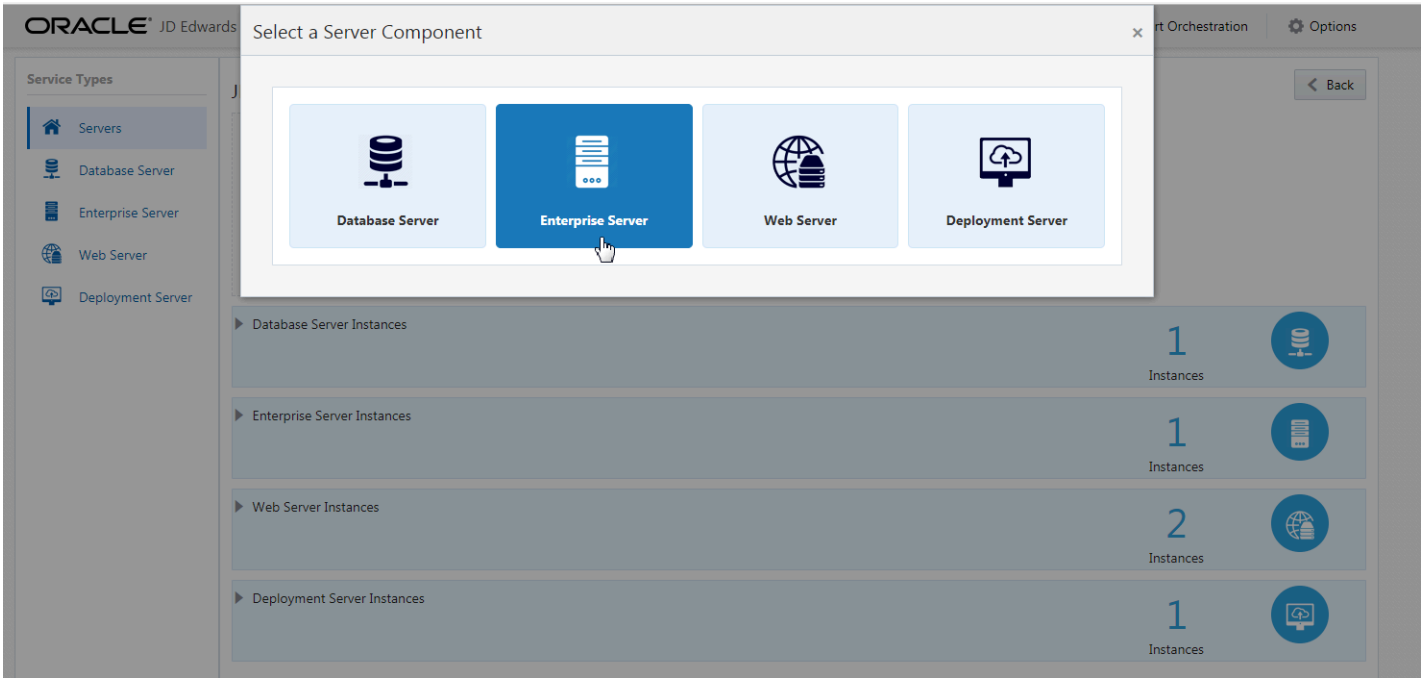

2. On Enterprise Server Instance page, complete these fields to create and configure the Enterprise Server instance.

Server Configuration

- *Platform*

  This field is disabled and it is automatically populated as Linux.
- *Instance Name*

  Create an instance name for the Enterprise Server. The conditions to set the instance name is displayed in the tooltip when you click the field.
- *Host Name*

  Enter the host name.

Enterprise Server Preferences

- *Server Type*

  Select one or both of the available server types for this Enterprise Server.

  Single Enterprise Server. If you are deploying only a single Enterprise Server, select both Logic and Batch as the server types.

  Multiple Enterprise Servers. If you are deploying multiple Enterprise Servers, at least one server must be specified as a Logic server for each pathcode. The other servers can be specified as Batch servers.
- *Pathcodes*

  Click the Available Pathcodes field and select the pathcodes required from the auto-suggest text. The four available pathcodes are: Development, Prototype, Pristine, and Production.

  > **Note:** Important: It is good practice to select pathcodes here that correlate to the schemas you selected for the Database Server. The Provisioning Console programmatically enforce this correlation. If you select pathcodes on the Enterprise Server that are a superset of the database schemas you selected, the Enterprise Server will not be able to access the data required to function correctly. In the Provisioning Console, the pathcodes that you specify at this point for installation on the Enterprise Server can be deployed only once. You can use the Provisioning Console to programmatically add additional schemas after deploying the orchestration.

- **Oracle JDBC Driver Details**

  This driver is required for connectivity between the Enterprise Server and the Oracle database server.

  Click the Browse button to select each of the required components for the Oracle JDBC driver. For example:

  - odbc8.jar
  - ons.jar
  - ucp.jar

Note: Refer to Oracle Certifications for the version of the supported driver and associated components.

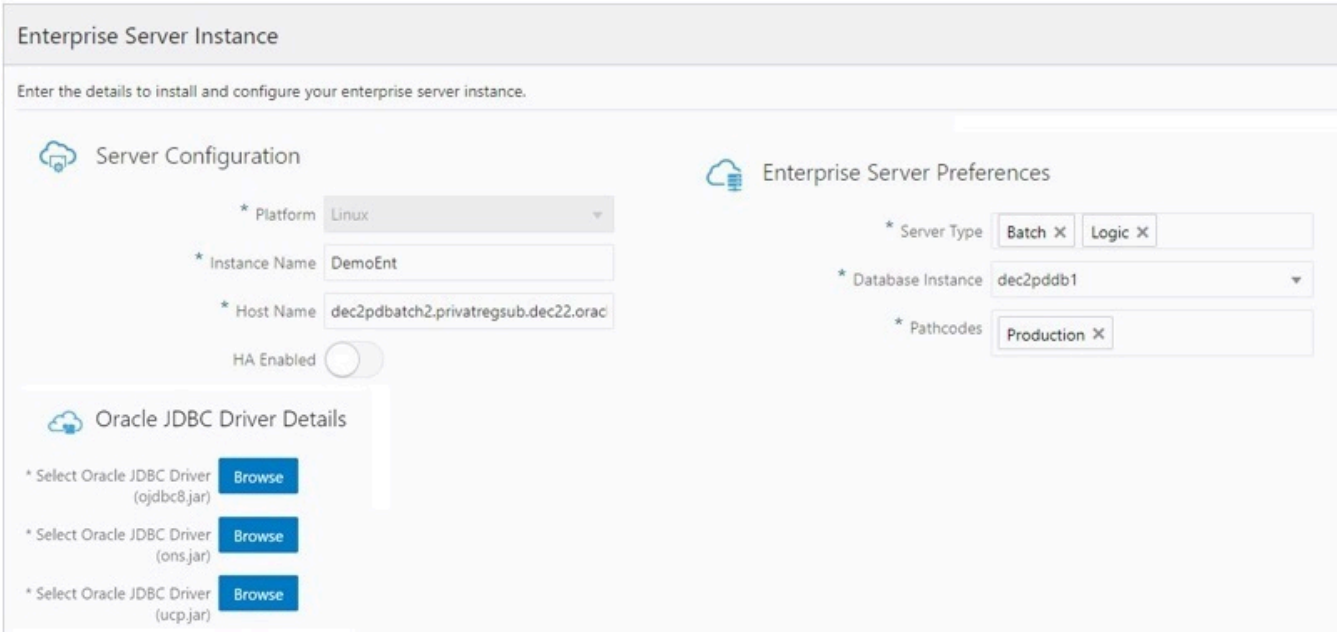

3. You should now be able to view multiple instances of the Enterprise Server.

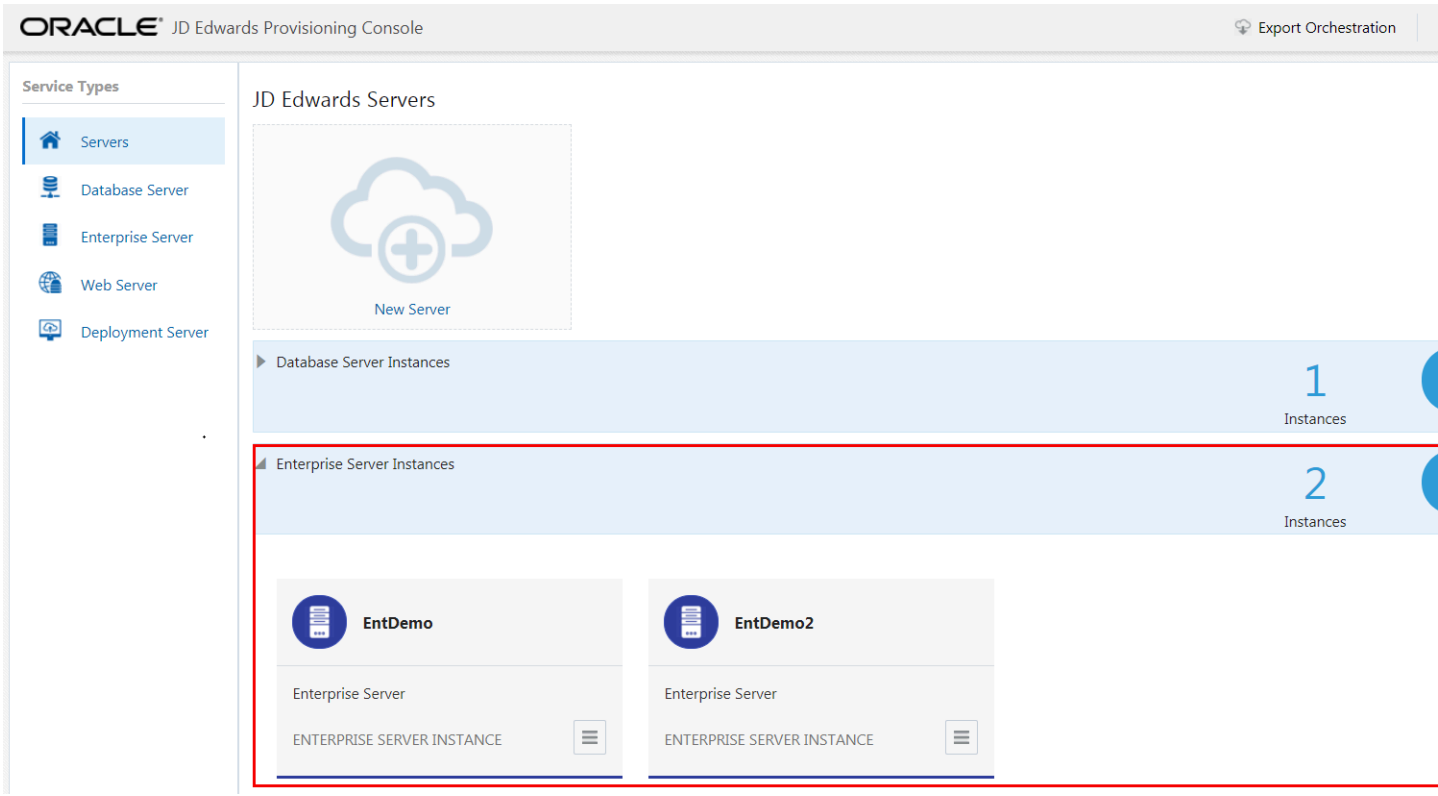

4. If an Enterprise Server exists, click **View** from the Application Options tab to view the existing configuration for the Enterprise Server. To modify the instance configuration use the Application Option tab and choose the **Edit** option. Use the Clone option to **clone** the Enterprise Server instance.

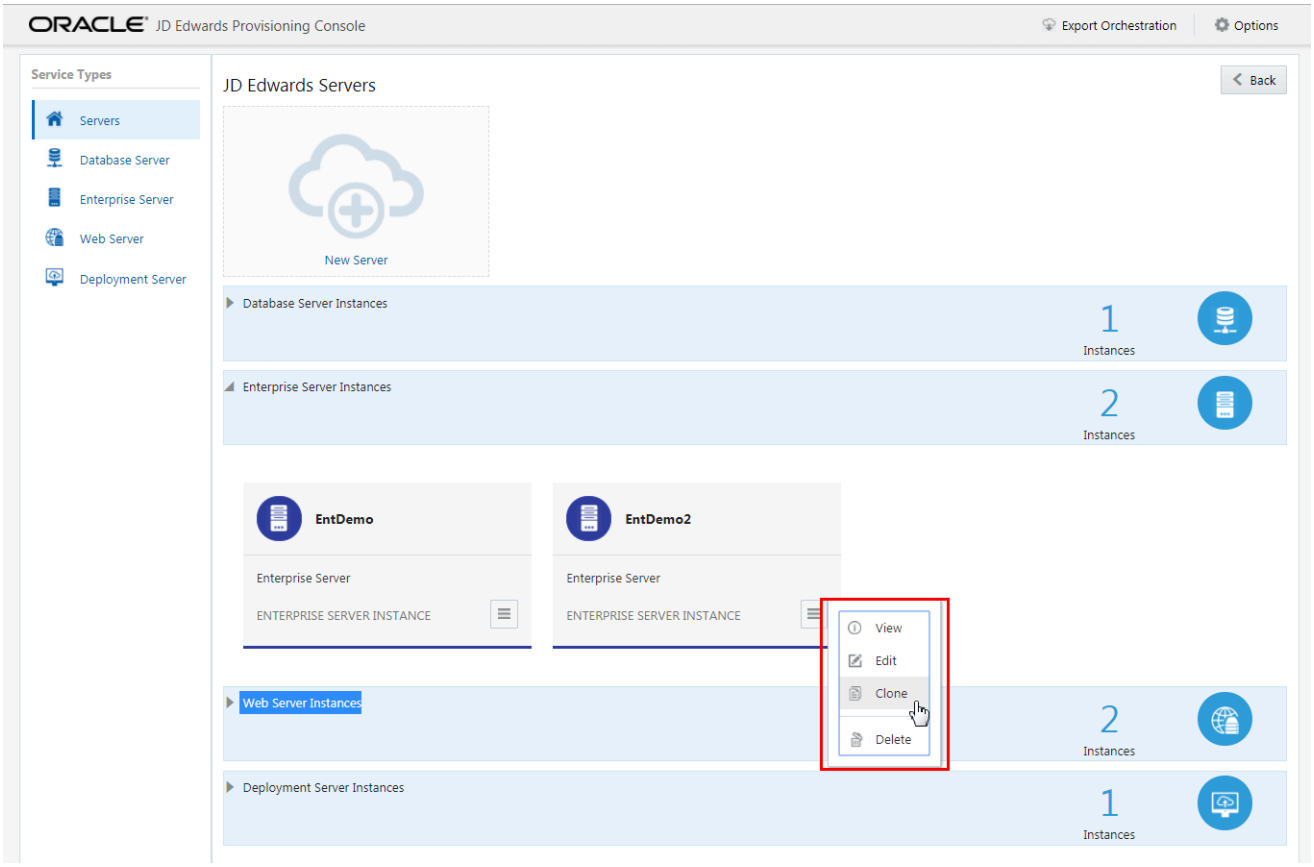

5. To delete the instance, use the Application option tab to select **Delete**.

Web Server

The Instance Type for Web Servers can be any of the available servers types that are selectable from the drop-down list; however, you must have at least one configured Dedicated HTML Server saved prior to creating an associated AIS instance.

You can define any number of Web Server instances.

To add or modify a Web Server instance:

1. Select the Service Type for Web Server.
   - To create a new Web Server, click the **New Server** icon and then the **Web Server** icon and complete the required fields.
   - To modify an existing Web Server, click **View** from the Application Actions tab to view the existing configuration for the Web Server. To modify the instance configuration use the action tab and choose the **Edit** option.

2. On the Web Server Instance, you can configure each of these Web Server types:
   - HTML Server (Dedicated HTML Server for AIS)

     At least one of these servers must be specified in a pair with an AIS Server.
   - Application Interface Services (AIS) Server

     At least one of these servers must be specified in a pair with a Dedicated HTML Server for AIS.
   - Standard JAS Server

     This is a traditional JAS Server, which is optional and can be none to many.

     > **Note:** For Oracle Cloud Infrastructure only, you can select the HA Enabled option if required, and when prompted, should enter the Virtual Host Name, and then click OK.

> **Note:** For a description of each HTML server type (Dedicated for AIS or Standard), refer to the Fundamentals section of this Learning Path.

> **Note:** Because multiple Web Servers instances can run on the same WebLogic Server running in Oracle Cloud Infrastructure, you MUST specify different ports for each instance.

3. On the HTML Server Instance page, complete these fields to create and configure the HTML Server instance.

   Server Configuration

   - *Platform*

     This field is disabled and it is automatically populated as Linux.
   - *Instance Name*

     Create the instance name of the HTML Server instance.
   - *Host Name*

     Enter the host name.
   - *Port*

     Enter a unique (available) port number for this server that will use an SSL connection. This port number must be between 1024 and 65535. This port number is used by HTTPS to create a container and deploy

the web component. Ensure the availability of a port that is one less than the port number that you enter here. That is, if you specify port 8081, you must also ensure that port 8080 is available.

> **Note:** For each SSL port that you open in the firewall, you must also open a companion port for non-SSL access required for Server Manager. The numeric value for the companion port must be one less than the value specified for the SSL port. For example, if you specify a port value of 8081 for SSL, in the firewall you must also open a port one less than that value; in this case you must open port 8080. Refer to the section "Enable Inbound Ports in the Firewall for Compute Instances" in the OBE "Performing Common Setup for All Linux Servers" of this Learning Path.

Web Server Preferences

- *Pathcode*

  Select the required pathcode from the drop-down menu.

  > **Note:** Using the Quick Start mode, you can specify only a dedicated HTML Server for AIS. If you want to create a standard HTML Server, which is strongly recommended for Production environments, you must use the Advanced Deployment mode. For a description of each HTML Server type, refer to the section "Fundamentals" of this Learning Path.

  > **Note:** Each dedicated HTML Server and AIS Server pair can support only one pathcode. If you want additional HTML instances to support additional pathcodes, you must configure additional HTML Server pairs using the Advanced deployment mode of the Provisioning Console. For more information, refer to the OBE "Orchestrating Using Advanced Mode" of this Learning Path.

WebLogic Details

- *User Name*

  Enter the user name.
- *Password*

  Enter the WebLogic Server password.
- *Admin Port*

  Enter the port number to access the WebLogic Administration Console.
- *Install Path*

  Enter the installation path of the WebLogic instance.
- *JDK Install Path*

  Enter the JDK installation path.

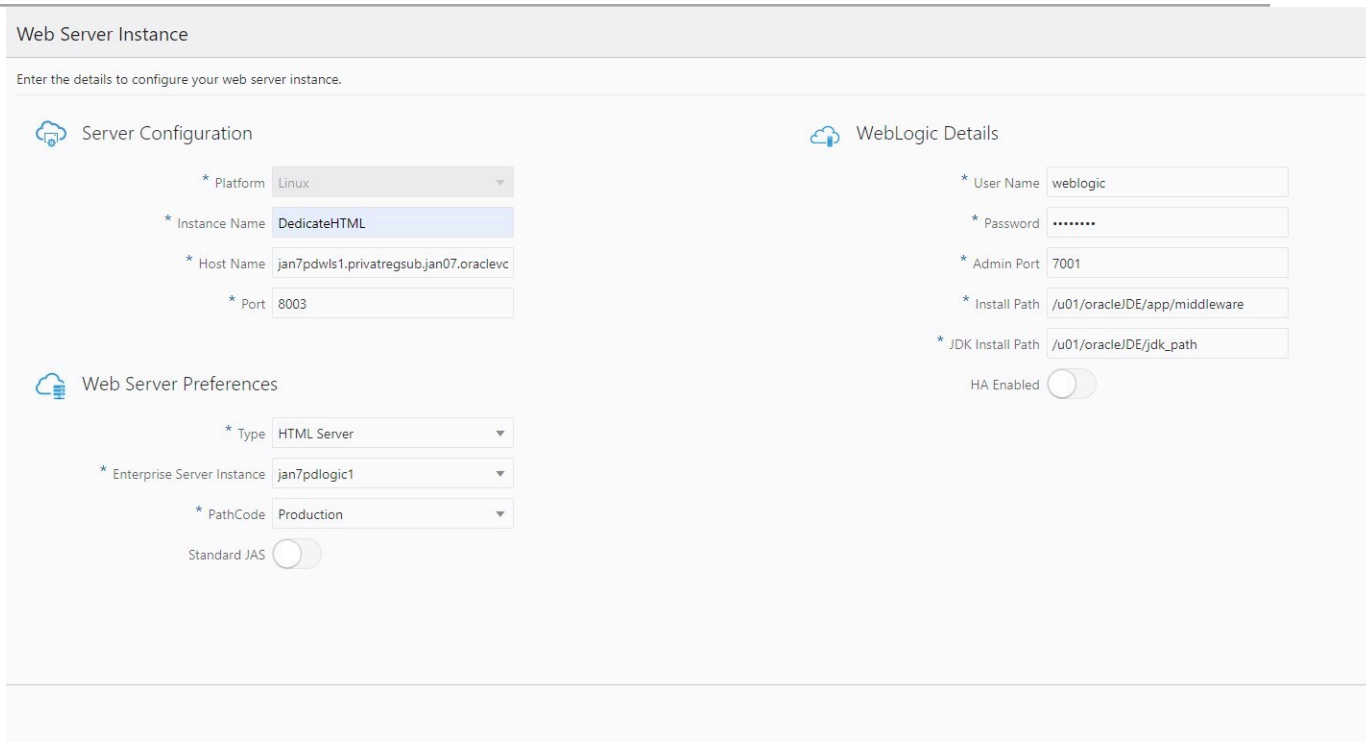
Web Server Instance
Enter the details to configure your web server instance.
Server Configuration
* Platform Linux
* Instance Name DedicateHTML
* Host Name jan7pdwls1.privatregsub.jan07.oraclevc
* Port 8003
WebLogic Details
* User Name weblogic
* Password ••••••••
* Admin Port 7001
* Install Path /u01/oracleJDE/app/middleware
* JDK Install Path /u01/oracleJDE/jdk_path
HA Enabled
Web Server Preferences
* Type HTML Server
* Enterprise Server Instance jan7pdlogic1
* PathCode Production
Standard JAS

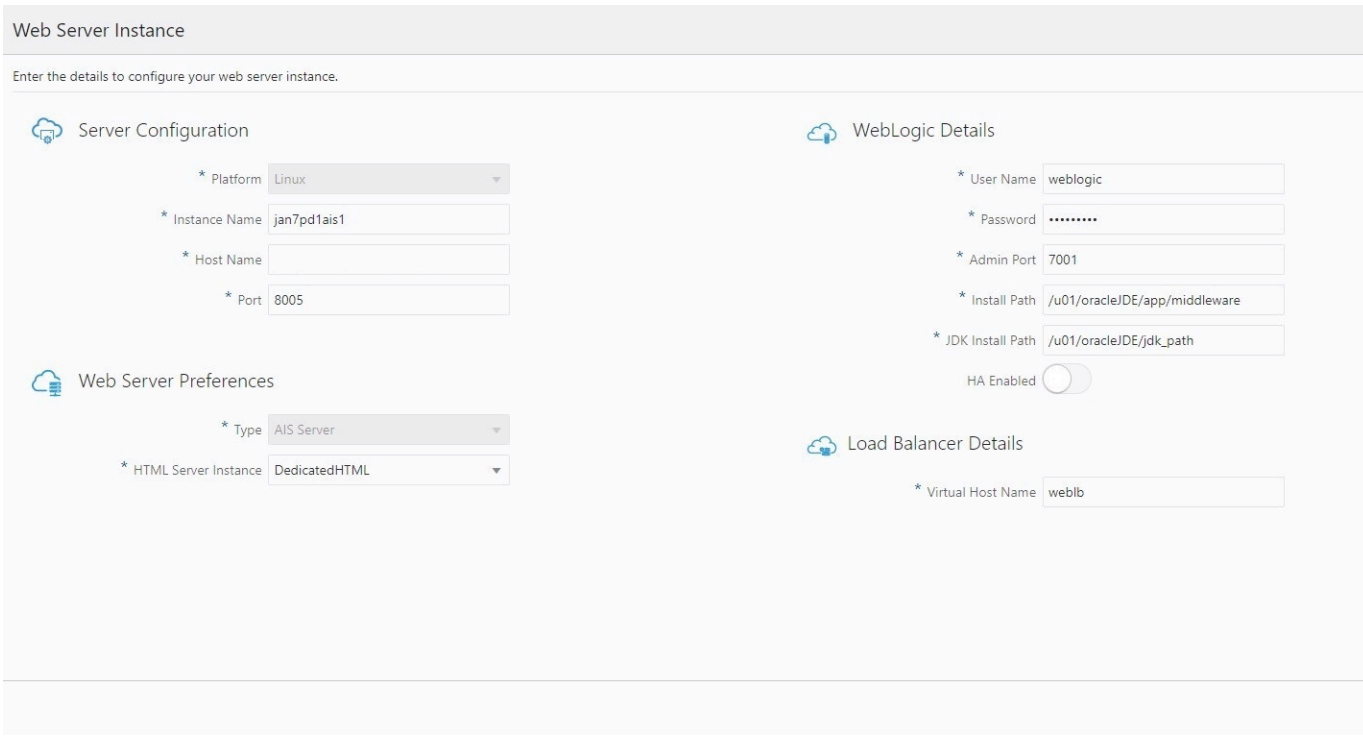
Web Server Instance
Enter the details to configure your web server instance.
Server Configuration
* Platform Linux
* Instance Name jan7pd1ais1
* Host Name
* Port 8005
WebLogic Details
* User Name weblogic
* Password ••••••••••
* Admin Port 7001
* Install Path /u01/oracleJDE/app/middleware
* JDK Install Path /u01/oracleJDE/jdk_path
HA Enabled
Web Server Preferences
* Type AIS Server
* HTML Server Instance DedicatedHTML
Load Balancer Details
* Virtual Host Name weblb

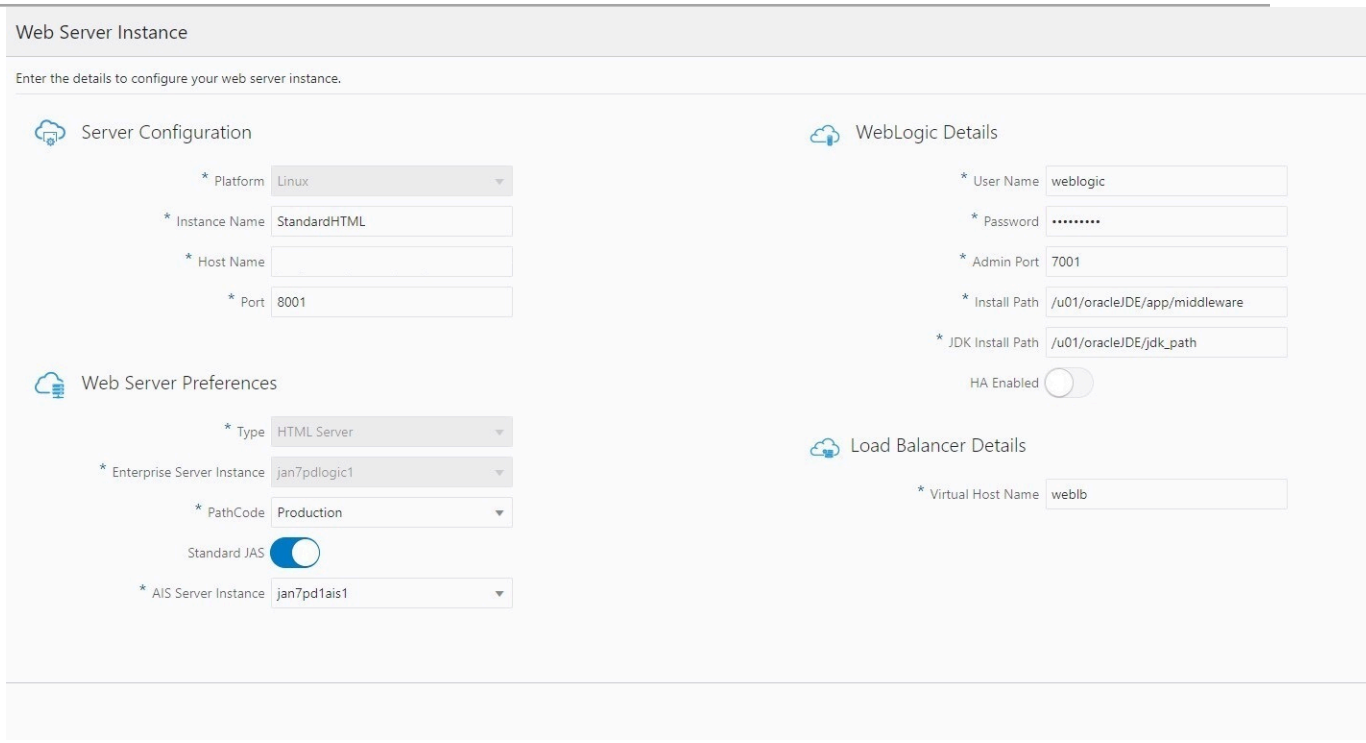

4. Click the **OK** button.

5. Verify the Web Server instances you modified or added is displayed in the JD Edwards Servers window.

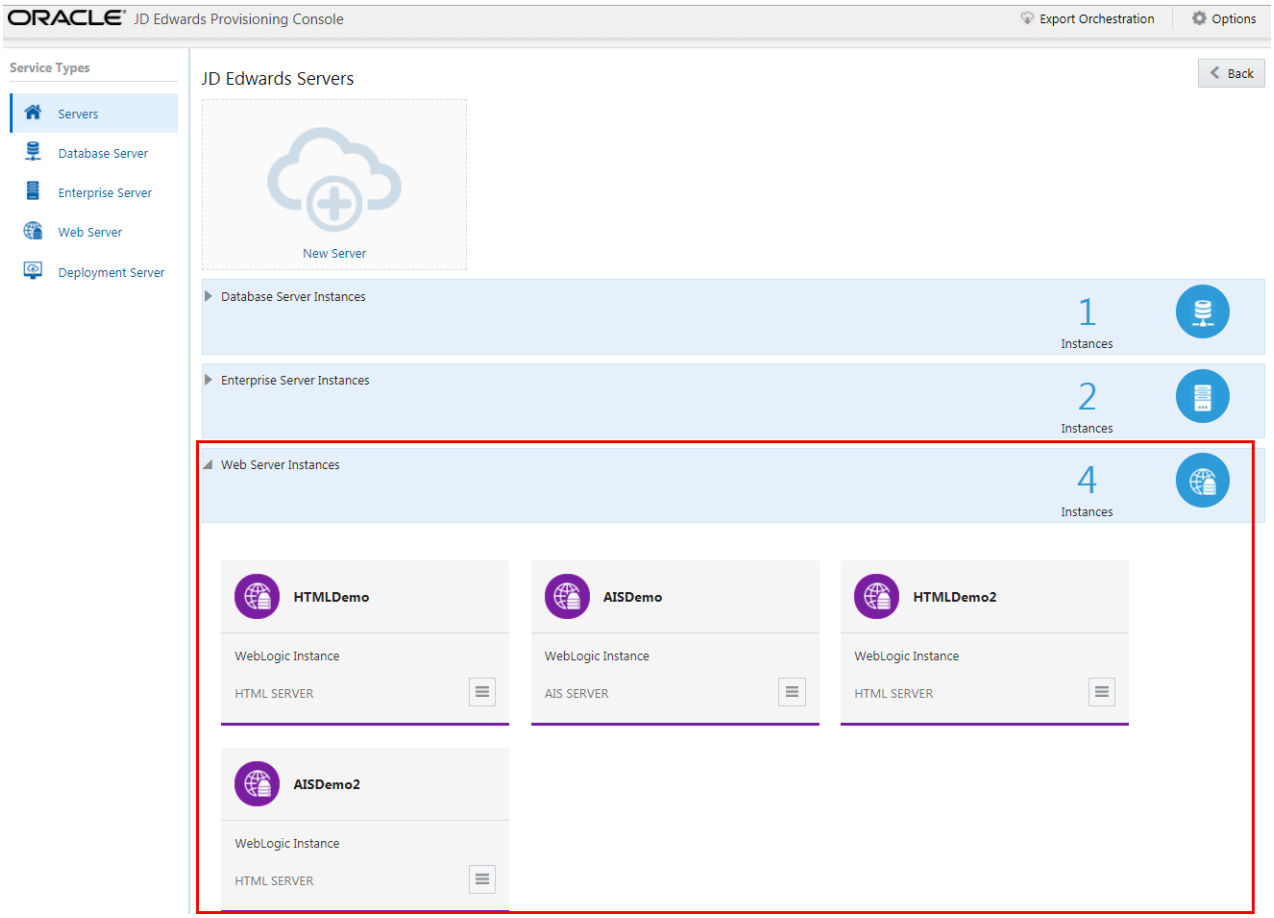

6. To delete any web instance, use the Application Options tab to select **Delete.**

Deployment Server

You can use only one Deployment Server per deployment. If you attempt to add more than one Deployment Server, the Provisioning Console displays an error.

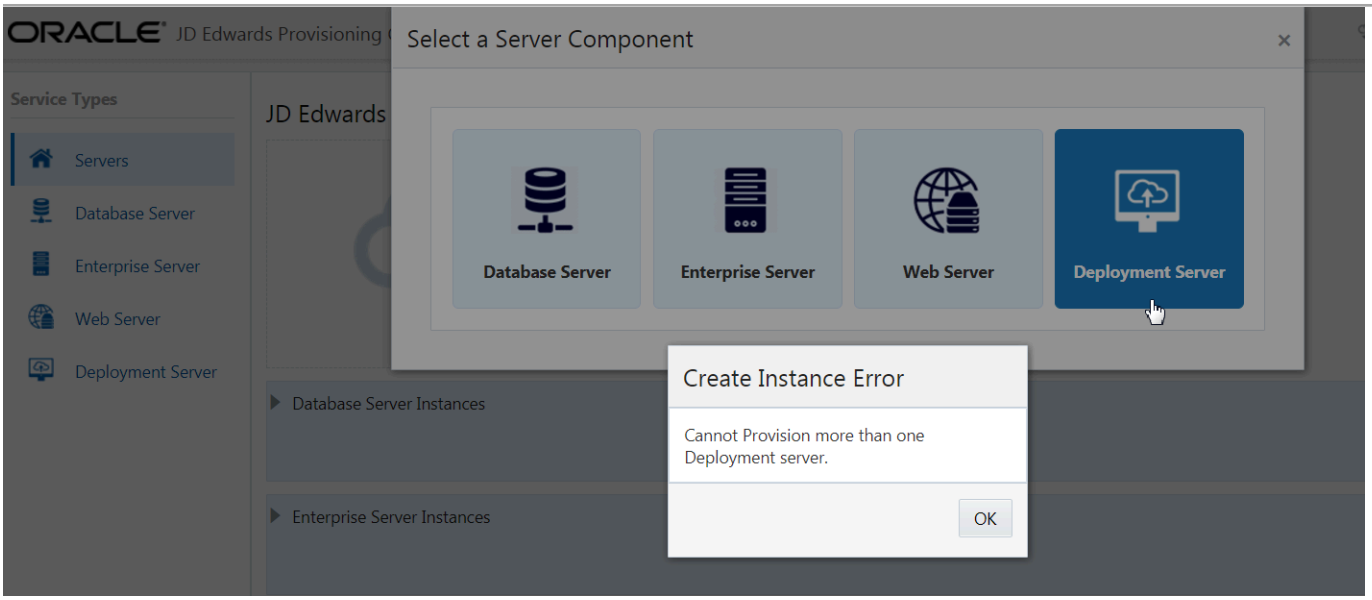
ORACLE JD Edwards Provisioning
Service Types
Servers
Database Server
Enterprise Server
Web Server
Deployment Server
JD Edwards
Database Server Instances
Enterprise Server Instances
Select a Server Component
Database Server
Enterprise Server
Web Server
Deployment Server
Create Instance Error
Cannot Provision more than one Deployment server.
OK

To modify the Deployment Server instance:

1. On JD Edwards Servers, click the existing Deployment Server instance, click the **Applications Options** icon, and then click **Edit**.

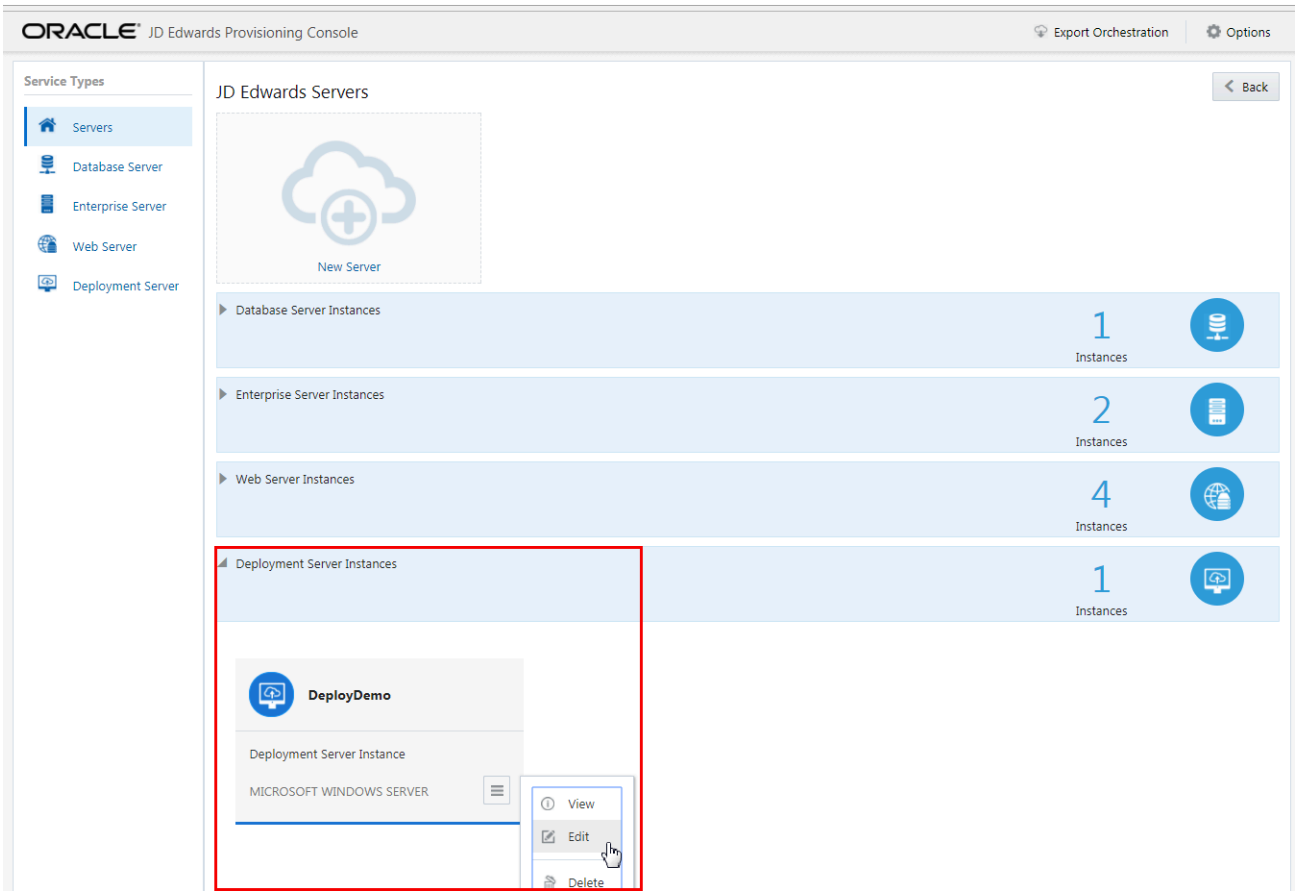

2. On JD Edwards Deployment Server page, complete these fields to create and configure your Deployment Server instance.

Server Configuration

- Instance Name

  Create a name for the Deployment Server instance. The conditions to set the instance name is displayed in the tooltip when you click the field.
- Host Name

  Enter the host name.
- Windows User

  Enter the name of the Windows user.
- Windows Password

  Enter the password of the Windows user.

Deployment Server Preferences

- *Location*

  Enter the location.

  This value is the base location for your JD Edwards EnterpriseOne machines. For example, typical values might be a city name (such as Denver or Austin), a geographical region name (such as US or India), or a general location name (such as Corporate).
- *Installation Drive*

  Enter the drive for the installation.
- *Pathcodes*

  This field is automatically populated.

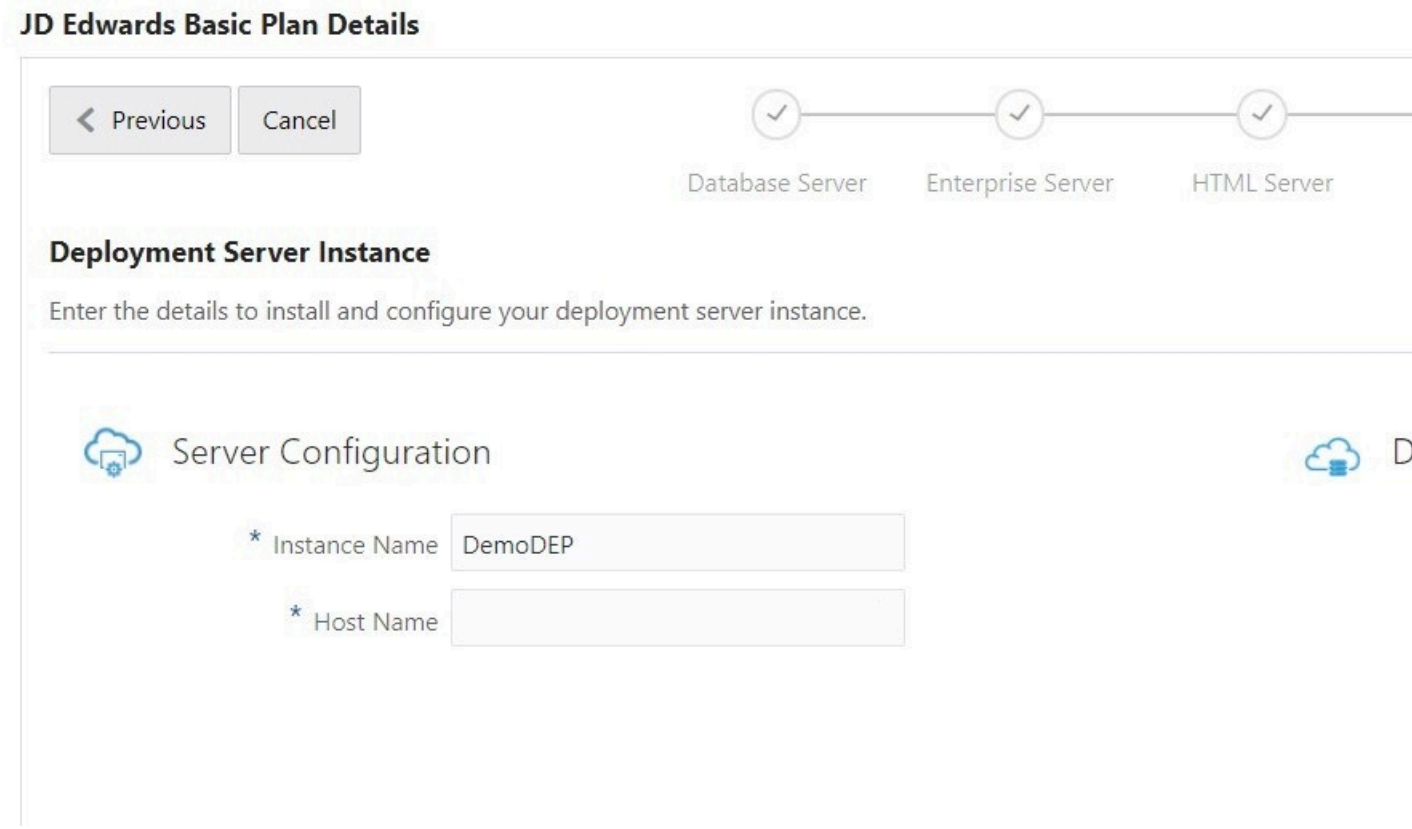

3. Verify that the Deployment Server instance you modified is displayed in the JD Edwards Servers window.
4. You can choose the **Delete** option in the action tab to delete the Deployment Server instance. After you delete the existing Deployment Server instance, you can click the **New Server** icon, and then select Deployment Server from the Select a Component page to add a new Deployment Server instance.

Note: To deploy an orchestration, refer to the section of this tutorial entitled: ***Deploying an Orchestration.***

13 Deploying JD Edwards EnterpriseOne

Deploying an Orchestration

This tutorial shows how to deploy an orchestration on Oracle Cloud Infrastructure on Linux using the JD Edwards One-Click Provisioning Console.

Deploying an Orchestration

Deploying an Orchestration

This section shows you how to deploy an Orchestration.

You can create a Quick Start or an Advanced Deployment Plan in the Orchestrate section of the JD Edwards One-Click Provisioning Console. When you start your deployment, the system initiates the scripts for the automated provisioning of the EnterpriseOne system.

Prerequisite

A completed Quick Start or an Advanced Deployment Plan created using the JD Edwards One-Click Provisioning Console.

Deploying an Orchestration

This procedure describes how to deploy an orchestration, whether it is a Quick Start or an Advanced mode Deployment Plan.

1. After you create a Deployment Plan, from the JD Edwards Provisioning Console, click the **Deploy** icon.

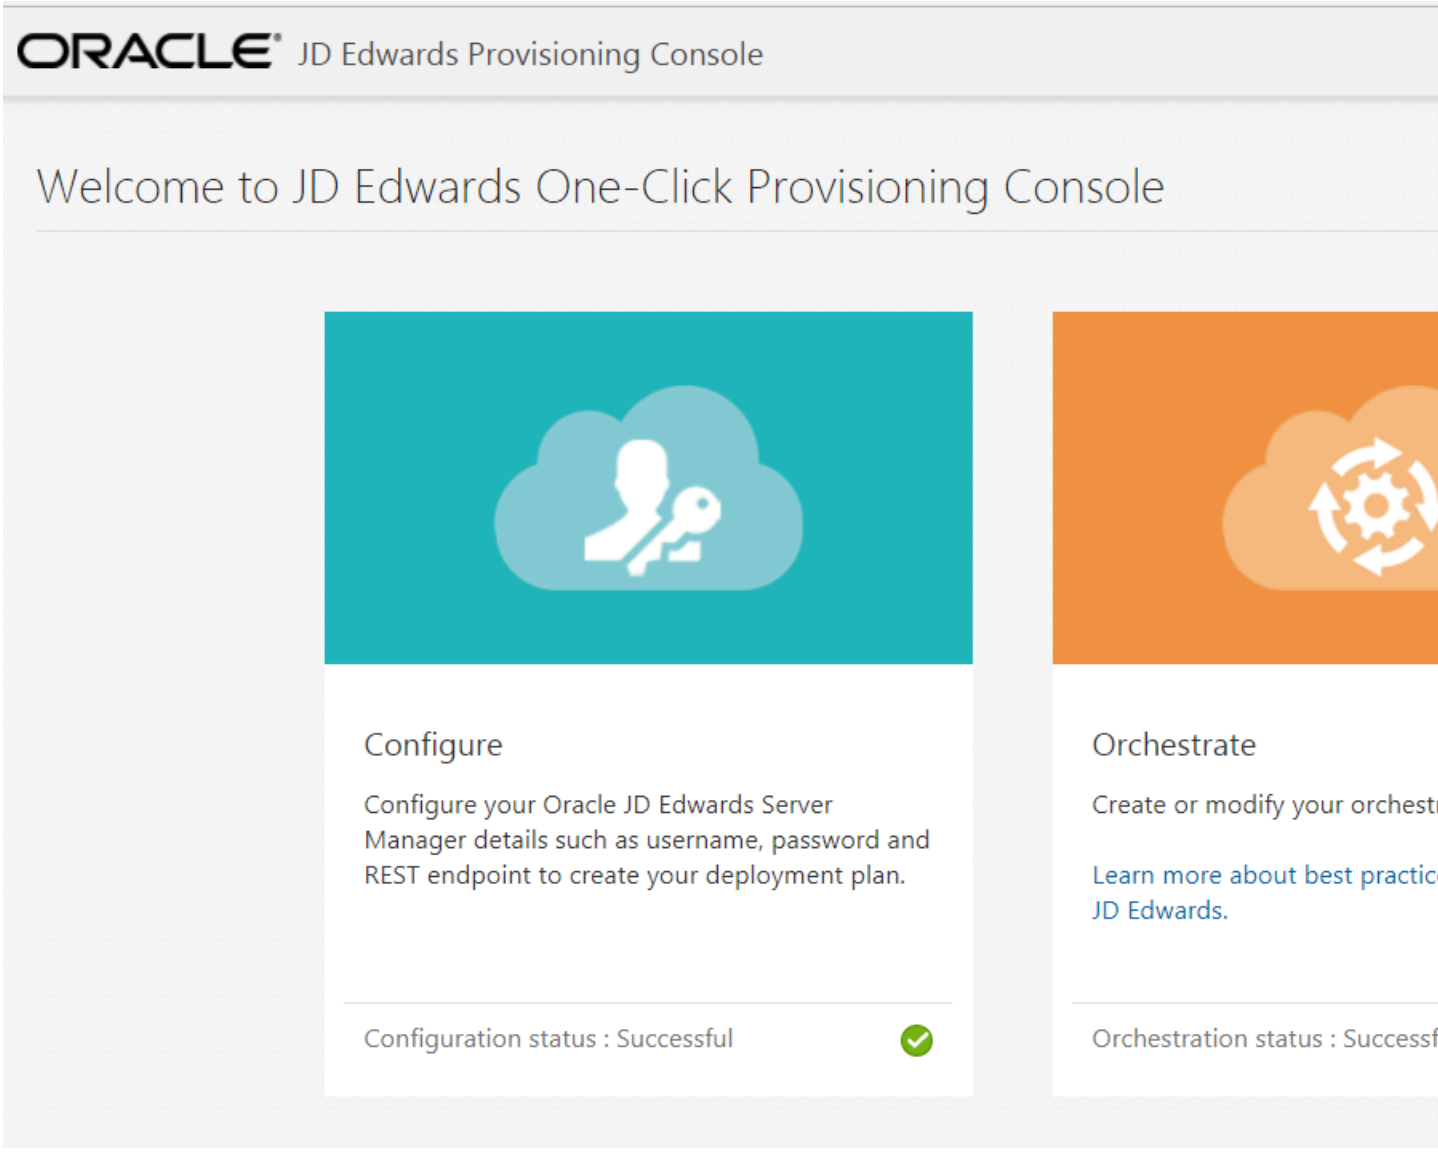

2. To view the Account details and Global Settings Summary, in the **Deployment Details** tab click the "**>**" icon for each server in the Deployment Details to see the details of the servers you provisioned.

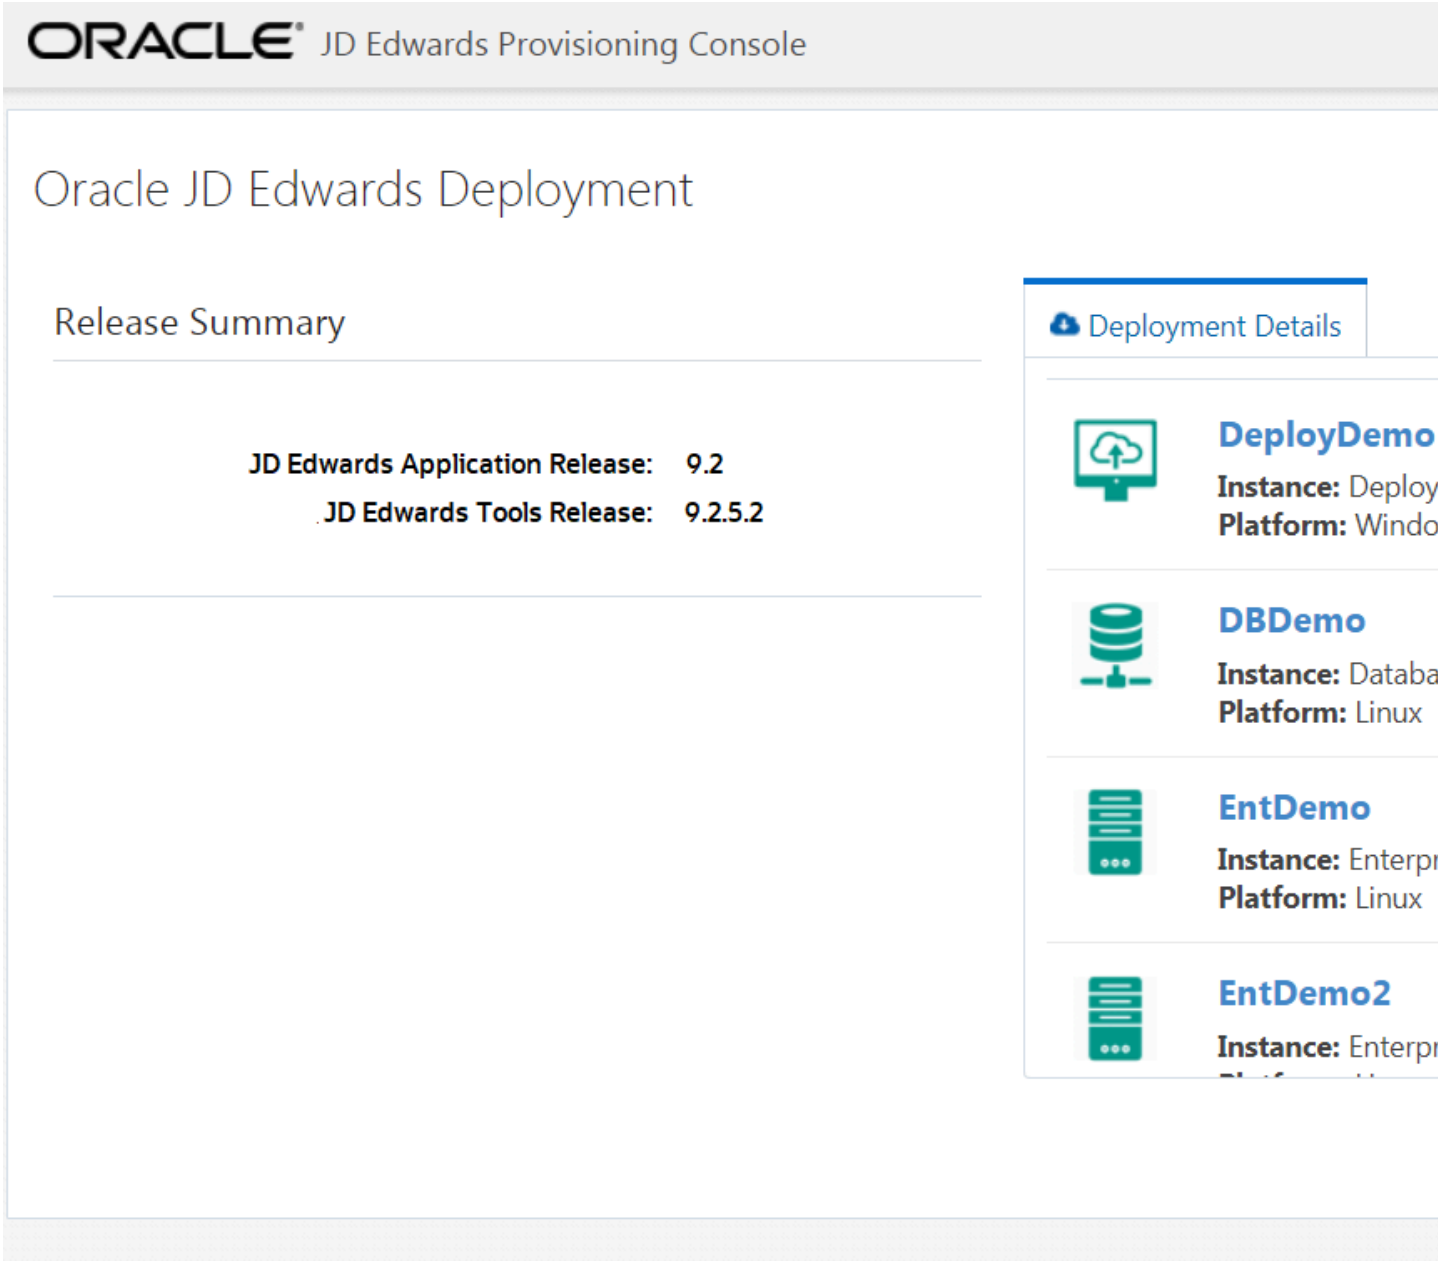

3. Click the **Back** button if required to make additional changes to the server instances.
4. To start the deployment, click the **Start Deployment** button.

5. On Deployment Status, you can see all the instances and task details and their progress. You can also view the log or summary of each of the tasks. The time taken to deploy the servers depends on your customization.

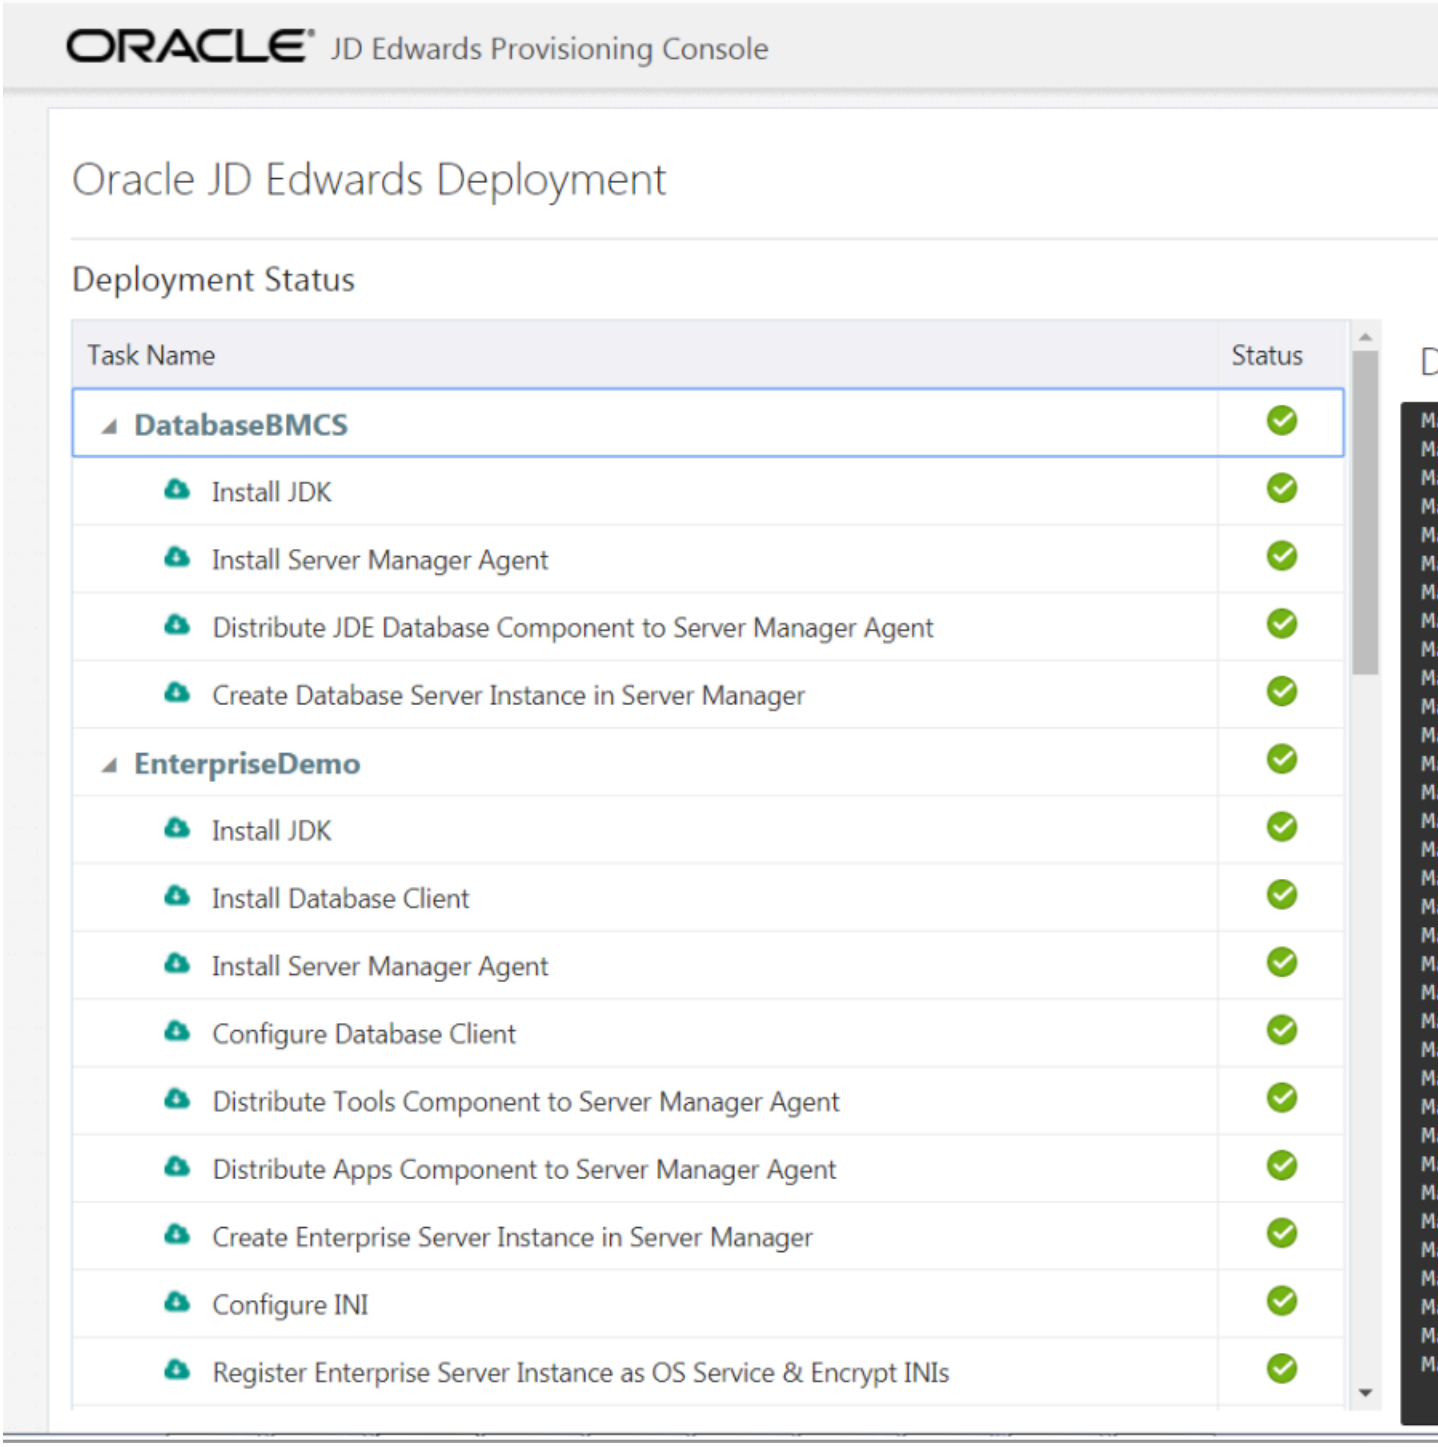

6. Click the **Back** button if you want to go back to the Deployment Summary window.

Adding Additional Pathcodes Post Deployment

You can add additional pathcodes to the Database Server instance after the deployment is successful.

To add additional pathcodes:

1. On Welcome to the JD Edwards Provisioning Console, click the **Orchestrate** icon.
2. On JD Edwards Servers, click on **Database Server Instances**.
3. Select the Database Server instance, and then use the action tab to choose **Edit Schema** option.

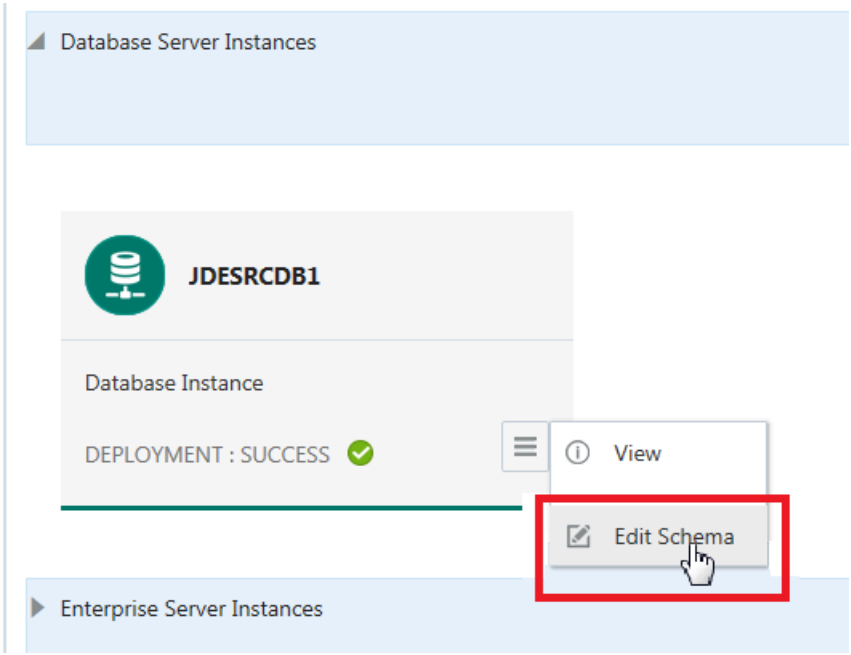

4. On Existing Database Instance Details, select the available pathcodes. An error message window appears if you attempt to delete the previously installed schema or the demo data.
5. Click the **OK** button.

6. Now you can see the Deployment status as **MODIFIED**.

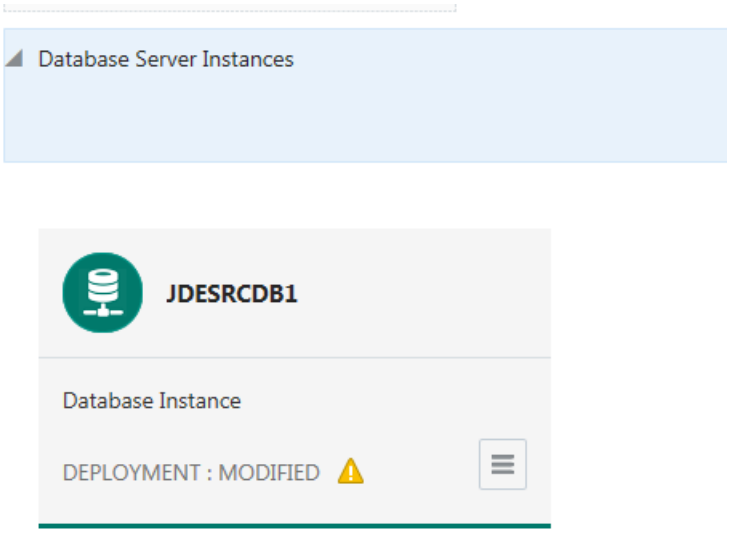

7. Click the **Back** button, and then click **Deploy** icon to start your modified deployment.

Orchestrating an Advanced Deployment Plan by Adding Web Servers and Enterprise Servers Post Deployment

This section shows you how to orchestrate an advanced deployment plan by adding Web Servers and Enterprise Servers post deployment.

You can add additional Web Servers and Enterprise Servers after you complete your Deployment. You cannot add any other server type after the deployment.

Prerequisite

Successfully completed the Deployment of a Quick Start or an Advanced Deployment Plan using the JD Edwards One-Click Provisioning Console. These tasks are described in preceding sections of this Learning Path.

Orchestrating an Advanced Deployment Plan by Adding Web Servers and Enterprise Servers Post Deployment

To add additional servers post deployment:

1. On JD Edwards One-Click Provisioning Console, click the **Orchestrate** icon .
2. On JD Edwards Servers, click the **New Server** icon, and in the Select a Component window, select **Web Server**.
3. On Web Server Instance Details, enter appropriate values for your installation.
4. Click the **OK** button. Similarly, you can add new Enterprise Servers.
5. On JD Edwards Servers, click the **Back** button.
6. On **JD Edwards One-Click Provisioning Console**, click the **Deploy** icon.

7. On **JD Edwards Deployment** window, review the deployment summary. You can verify the newly added servers in the **Deployment Details** tab.

> **Note:** To deploy an orchestration, refer to the section of this guide entitled: ***Deploy an Orchestration.***

14 Saving and Reusing Deployment Plans

Exporting an Orchestration

This section shows you how to export an orchestration.

You can use the export and import functions of the JD Edwards Provisioning Console to save an existing orchestration (export) and to reuse (import) the saved orchestration.

Prerequisite

A completed Quick Start Deployment Plan or an Advanced Deployment Plan created using the JD Edwards One-Click Provisioning Console.

Exporting an Orchestration

To export an orchestration:

1. On the JD Edwards Provisioning Console window, click the **Orchestrate** icon.

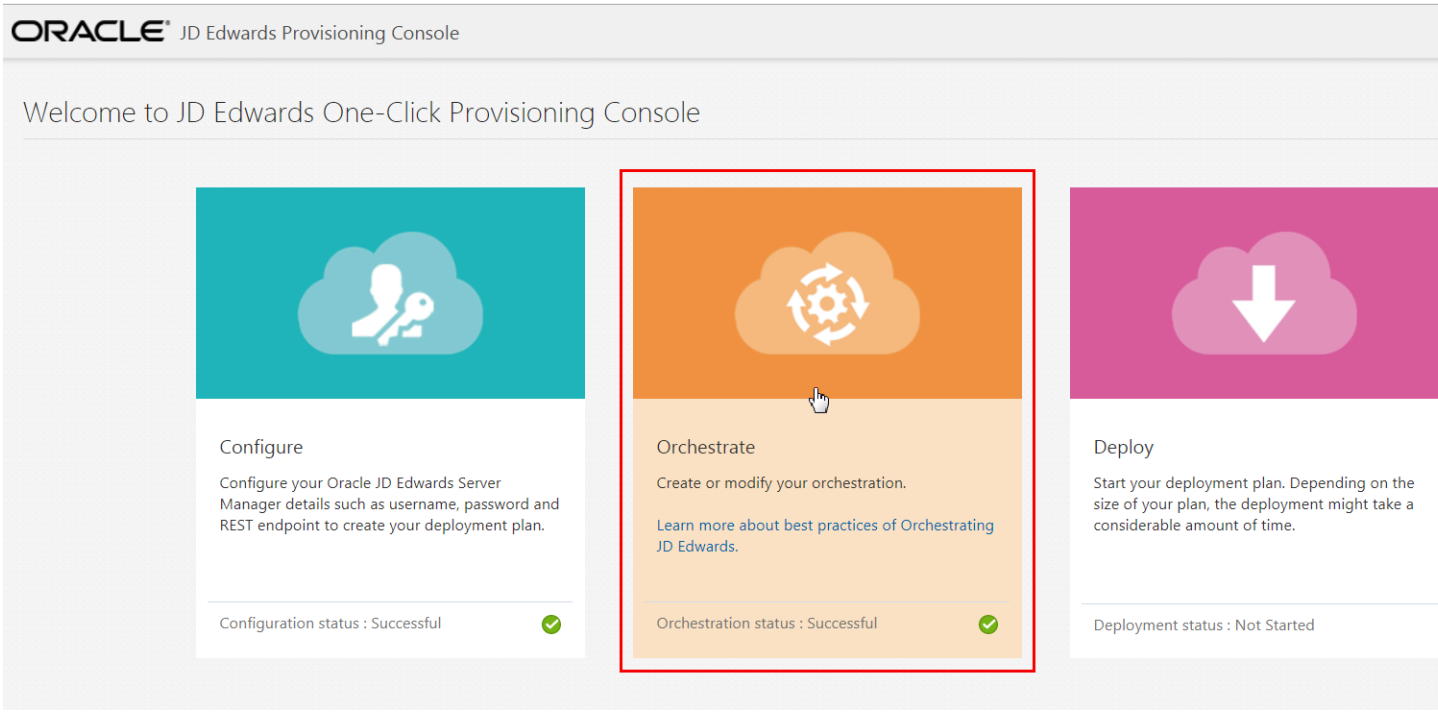

2. To export an existing orchestration, on an existing orchestration, from the title bar select **Export Orchestration**.

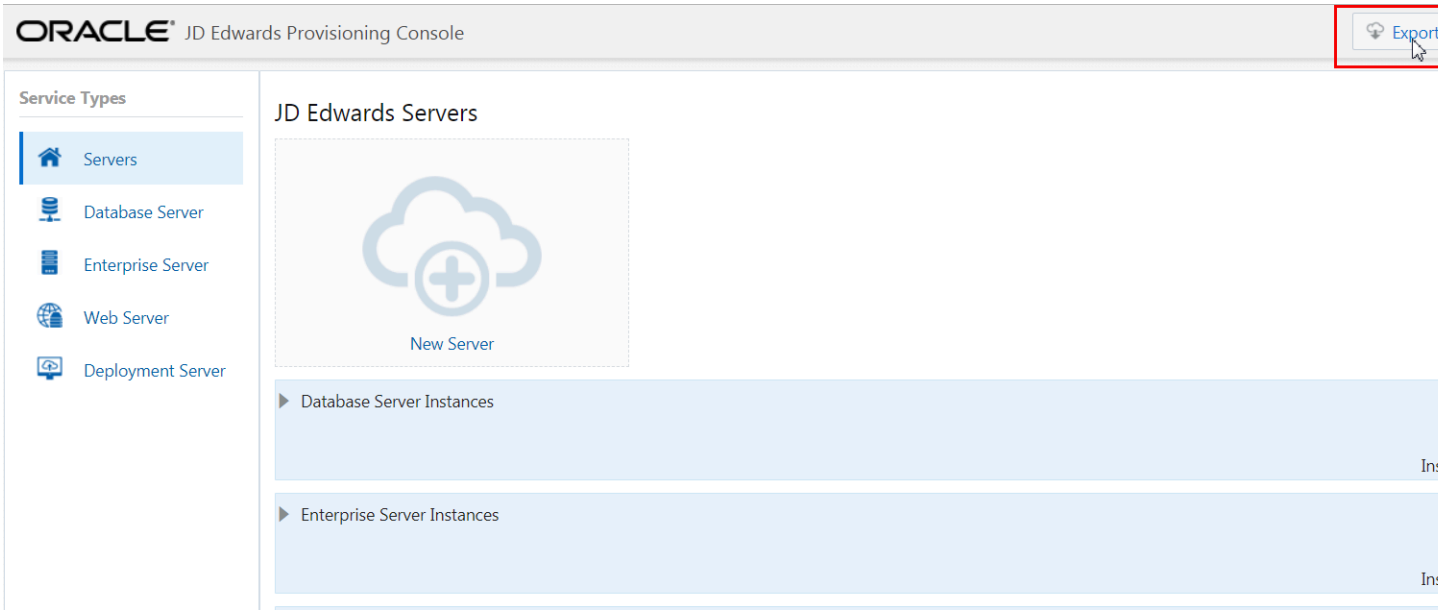

3. On the **Save As** dialog box, choose a path and file name for the exported orchestration.

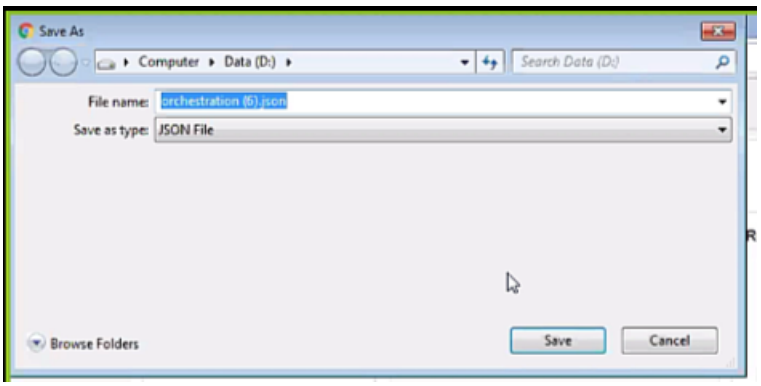

4. Click the **Save** button.

Note: If you want to view the raw contents of the .json file, you can open the file in an ASCII editor as shown in the following screenshot.

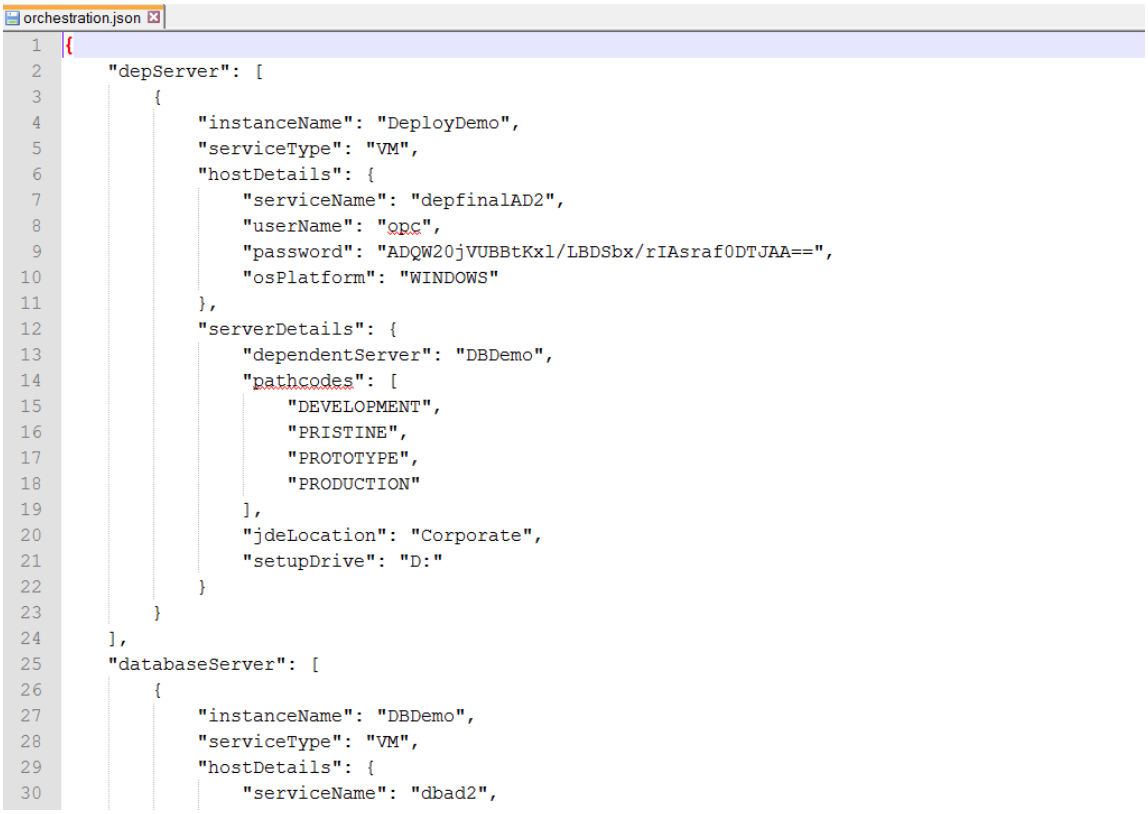
orchestration.json

```
{
    "depServer": [
        {
            "instanceName": "DeployDemo",
            "serviceType": "VM",
            "hostDetails": {
                "serviceName": "depfinalAD2",
                "userName": "opc",
                "password": "ADQW20jVUBBtKxl/LBDSbx/rIAsraf0DTJAA==",
                "osPlatform": "WINDOWS"
            },
            "serverDetails": {
                "dependentServer": "DBDemo",
                "pathcodes": [
                    "DEVELOPMENT",
                    "PRISTINE",
                    "PROTOTYPE",
                    "PRODUCTION"
                ],
                "jdeLocation": "Corporate",
                "setupDrive": "D:"
            }
        }
    ],
    "databaseServer": [
        {
            "instanceName": "DBDemo",
            "serviceType": "VM",
            "hostDetails": {
                "serviceName": "dbad2",
```

Importing an Orchestration

This section shows you how to import an orchestration.

You can use the export and import functions of the JD Edwards Provisioning Console to save an existing orchestration (export) and to reuse (import) the saved orchestration.

Prerequisite

You must have configured the administrator passwords for WebLogic Server and Server Manager Console in the Configure section of the JD Edwards One-Click Provisioning Console.

Importing an Orchestration

To import an existing exported orchestration:

1. On the JD Edwards Provisioning Console, click the **Orchestrate** icon.

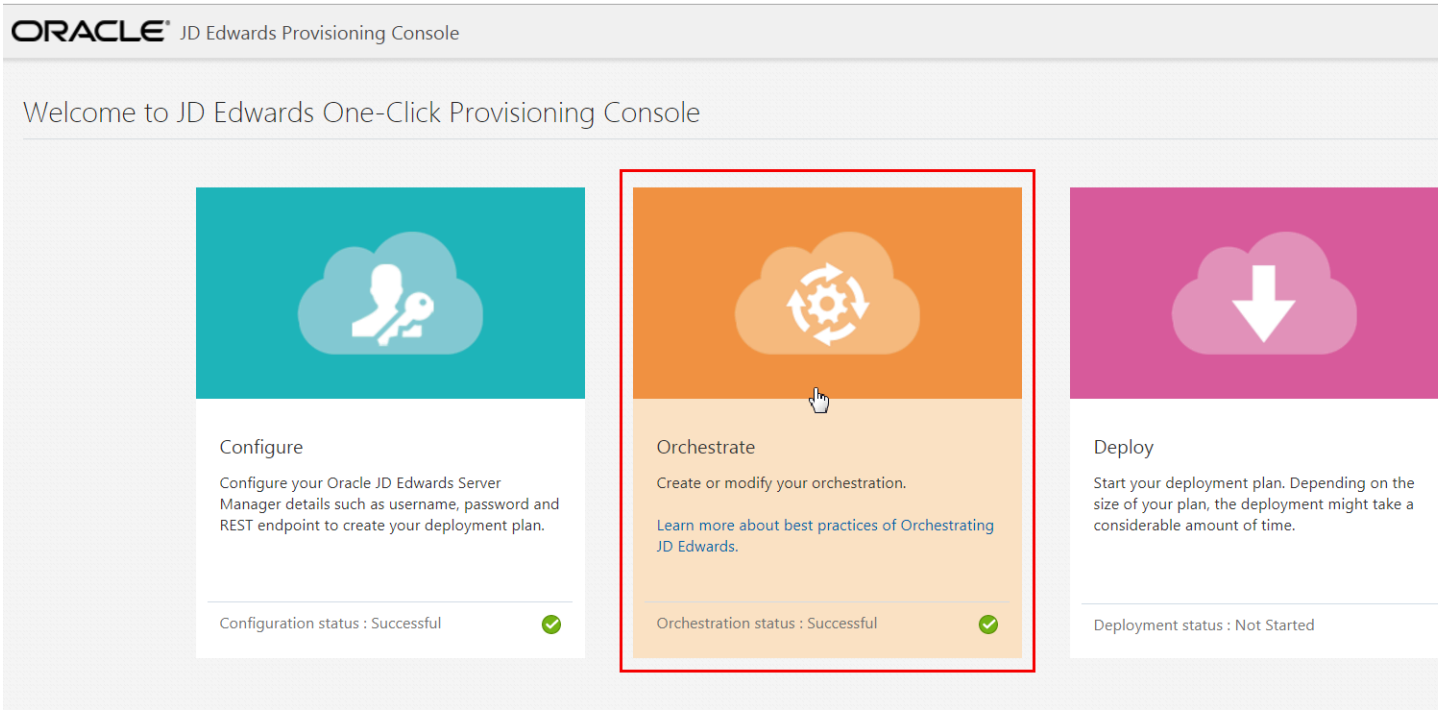

2. Click the **Import** icon.

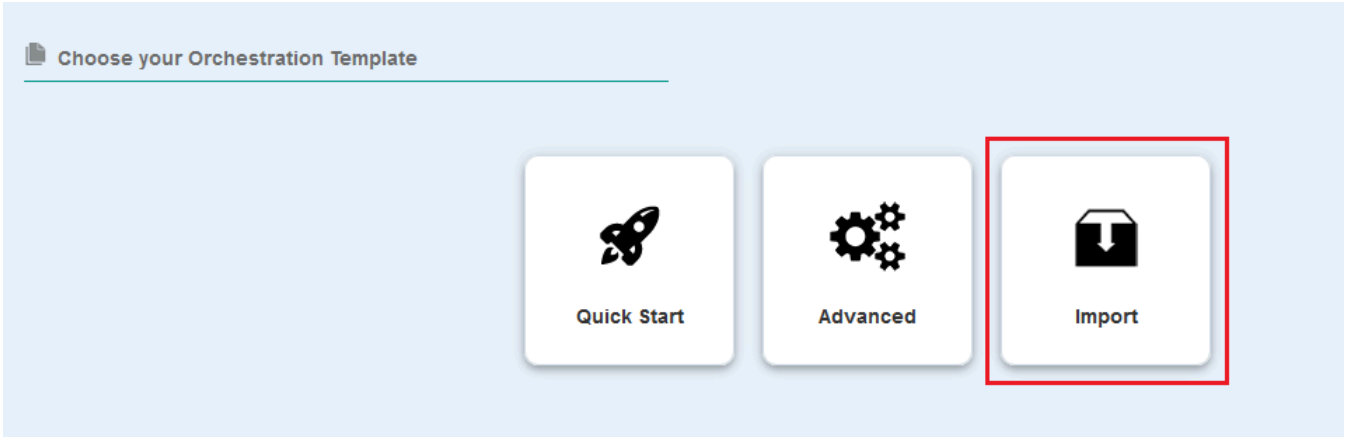

3. On the Import One-Click Provisioning Console Orchestration window, click **Choose File**, and then browse and select the exported `.json` file.

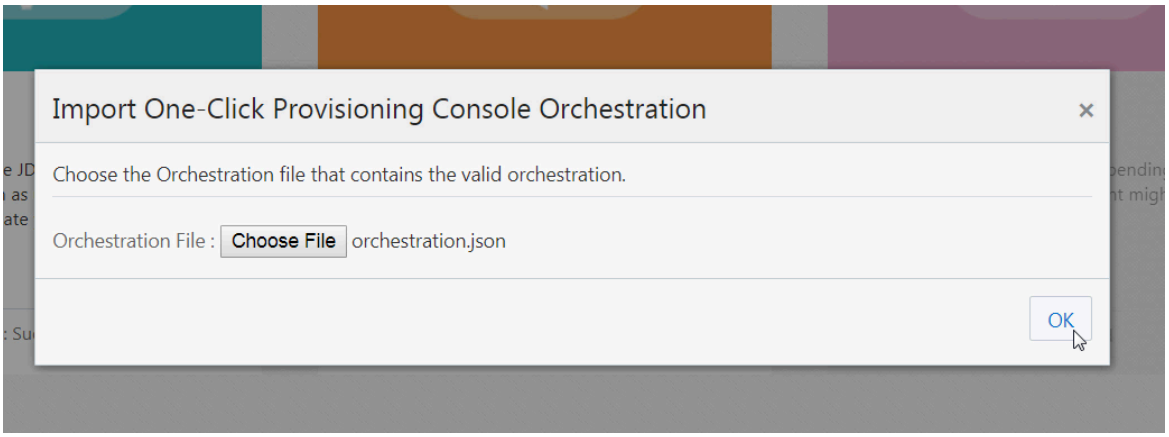

Note: If an error exists in the orchestration file you imported, the system displays the details of that error. Also, the system displays an error as illustrated in the following screenshot, if password authentication is required. You must edit the instance, enter the required password, and then save the instance.

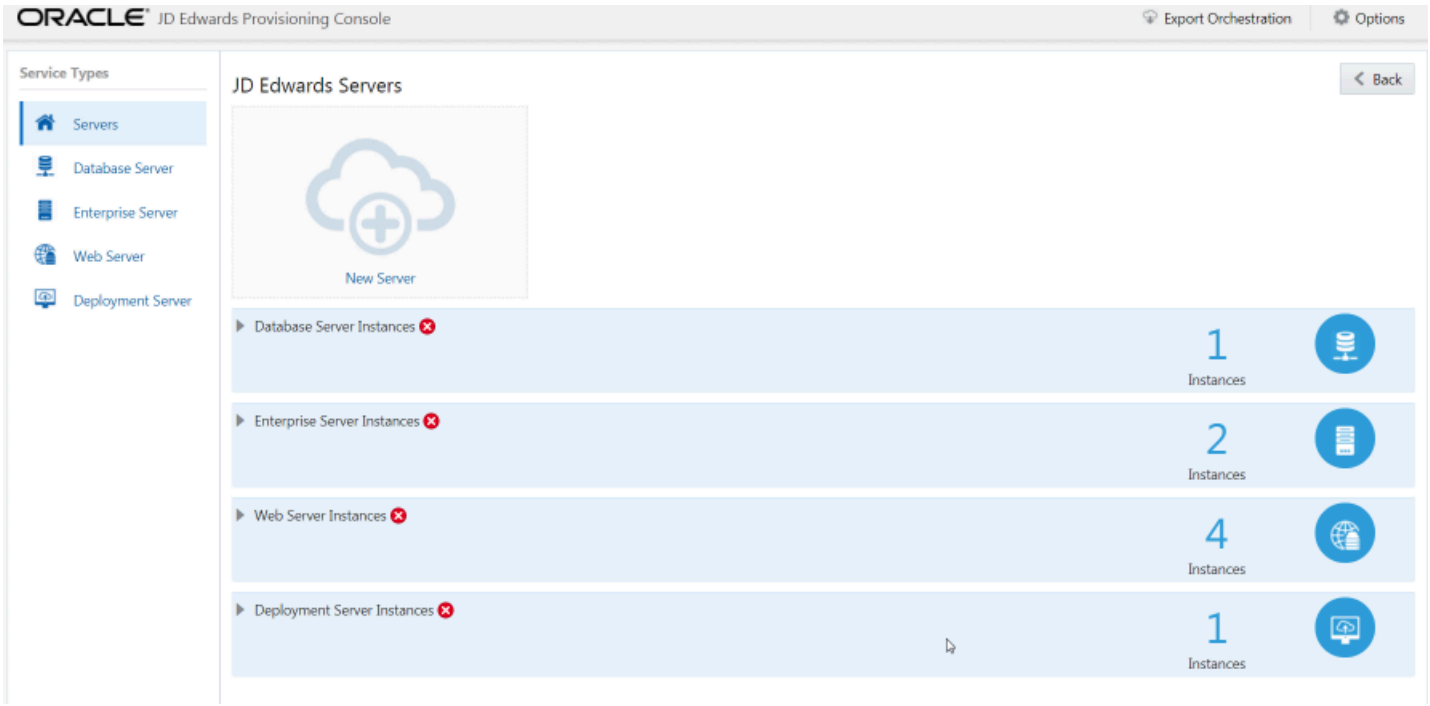

15 Configuring JD Edwards Components Post Deployment

Performing Post Installation for the Deployment Server

This section shows you how to perform post installation for the Deployment Server.

After you successfully provision all the Servers using the One-Click Provisioning Console, you must build packages on your Deployment Server.

Prerequisite

The complete Visual Studio product (which includes the runtime, the compiler, and associated tools) must be purchased and licensed from Microsoft.

- JD Edwards EnterpriseOne Applications Release 9.2 requires runtime libraries and the full product for Visual Studio.

General

The Deployment Server that is deployed by the Provisioning Server includes all the required third-party products including a JDK, E1Local Oracle database, and the EnterpriseOne database client. If you will be performing package builds, refer to the subsection entitled: Package Build Considerations.

Accessing the Deployment Server

You can access your Deployment Server using Microsoft Windows Remote Desktop Protocol (RDP).

- For One-Click Provisioning, you will need the Public IP address of the Deployment Server and the password.
- For Infrastructure Provisioning, you will need to connect as described in the section entitled "Connecting to a Windows Host in a Private Network Through the Bastion Host".

For information regarding the Public IP address, refer to the section of this Learning Path entitled: ***Accessing the JD Edwards EnterpriseOne Servers Using Their Public IP Addresses.***

The password for the Deployment Server was assigned when you input values in the **Deployment Server Instance** screen in the preceding sections of this Learning Path that are titled: ***Orchestrate a Quick Start Deployment Plan*** . If you followed the recommendation, this password should be recorded on the ***Pre-Install Worksheet*** .

Package Build Considerations

In order to build packages on your Deployment Server, you will need to:

- Install Microsoft Windows Visual Studio and Windows SDK
- Update Visual Studio Version in the jde.ini File
- Refresh CNC Data in JDEPLAN
- Build a New Client Package

- Build a New Server Package

Install Microsoft Windows Visual Studio and Windows SDK

The One-Click Provisioning Server delivers a Deployment Server to Windows 2022 Standard with most of the software already installed with the exception of Microsoft Windows Visual Studio and Microsoft Windows Software Development Kit (SDK). Refer to the JD Edwards Deployment Server Certification page for current updates on supported software versions and software prerequisites. The Certification page can be accessed through the Oracle Support Portal:

https://www.oracle.com/support/index.html

The Visual Studio runtime libraries (which are partial products with no development tools) for each supported release of Visual Studio are freely available from the Microsoft Download Center. The complete Visual Studio product (which includes the runtime and the compiler and associated tools) must be purchased and licensed from Microsoft.

As of the general availability of JD Edwards EnterpriseOne One-Click Provisioning for Tools Release 9.2, the following versions and associated links and navigation are valid.

- **Visual Studio 2022 Full Product**

  > **Note:** You need Visual Studio 2022, which is a licensed product from Microsoft, if you plan on building packages on the Deployment Server. *https://visualstudio.microsoft.com/downloads/*

- **Microsoft Software Development Kit (SDK) for Windows 11**

  > **Note:** Although the SDK download is labelled as Microsoft Windows 11, the same download is applicable to both Windows Client 11 and Windows Server 2022. This SDK is specifically required for any Microsoft Windows-based machine that is building JD Edwards EnterpriseOne packages.Windows Software Development Kit Version 10.0.22621.0 *https://msdn.microsoft.com/en-us/windows/downloads/sdk-archive*

Update Visual Studio Version in the jde.ini File

Ensure that you set the correct version of Visual Studio in the `[JDE_CG]` section of the `jde.ini` file on the Deployment Server. For details, refer to this document on Oracle Technology Network (OTN) for instructions:

JD Edwards EnterpriseOne Development Client Installation Guide for Oracle WebLogic Server (WLS) and WebSphere Application Server (WAS) Express.

- Understanding the Development Client Installation
- Installing Compiler, Linker, and Software Development Kit (SDK)
- Updating the jde.ini File

Refresh CNC Data in JDEPLAN

The ESU process has changed to include automatic generation of OCM mappings and tables for tables that are not in Business Data. In order for this process to work correctly, you must update the CNC information in JDEPLAN before applying any Tools-related ESUs. Also if you have multiple Enterprise Servers, Database Servers, or both, you must perform this procedure in order to synchronize data between the System and Planner pathcodes before you proceed further.

1. Sign into JDEPLAN.

2. Run R9840C, and copy system/planner information

   from: System – 920

   to: Planner – 920

   Note: R9840C will copy any missing CNC from System to Planner.

3. Run R98403A, choosing version XJDE0004, for syncing the data from System – 920 to Planner – 920 database.
4. Change the Data Selection to F98611.
5. Change the Processing Options as follows:
   - Source Data Source = System – 920
   - Target Data Source = Planner – 920
   - Proof Mode = 1
   - Replace Duplicates = Y
   - Clear Table = N

Using these processing options, running the R98403A will add any missing data source definitions in JDEPLAN with the correct definitions from System – 920.

Build a New Client Package

You must build a new Client package in order to obtain source files for the path code and to build future update packages against.

Note: Update packages against the delivered FA packages (PS920FA, PY920FA, PD920FA, and DV920FA) is not supported.

Due to space considerations on the `D:\` drive, you will need to move the `D:\Software` directory and all of its contents to the `C:\` drive or another computer. For example:

1. Add a new directory under the `C:\` drive labeled:

   `C:\SoftwareBKUP`
2. Copy the `D:\Software` directory (and all its contents) to the `C:\SoftwareBKUP` directory.
3. Delete the `D:\`Software directory.

Complete the following steps to build a full package by following the standard procedure to build a full package with the following important distinctions.

Note: The creation of update packages against the delivered FA packages (PS920FA, PY920FA, PD920FA, and DV920FA) is not supported. In order to build update packages in the future, you must build and deploy a new full package.

1. On the last screen of the Build Definition, on Package Assembly – [Package Build Revisions], select the Compression Options tab.
2. Ensure the **Compress Foundation** check box is selected. This is required to recompress the system folder on the Deployment Server to create a new systemcomp directory that includes a JRE. This JRE is required for use by the Development Client.

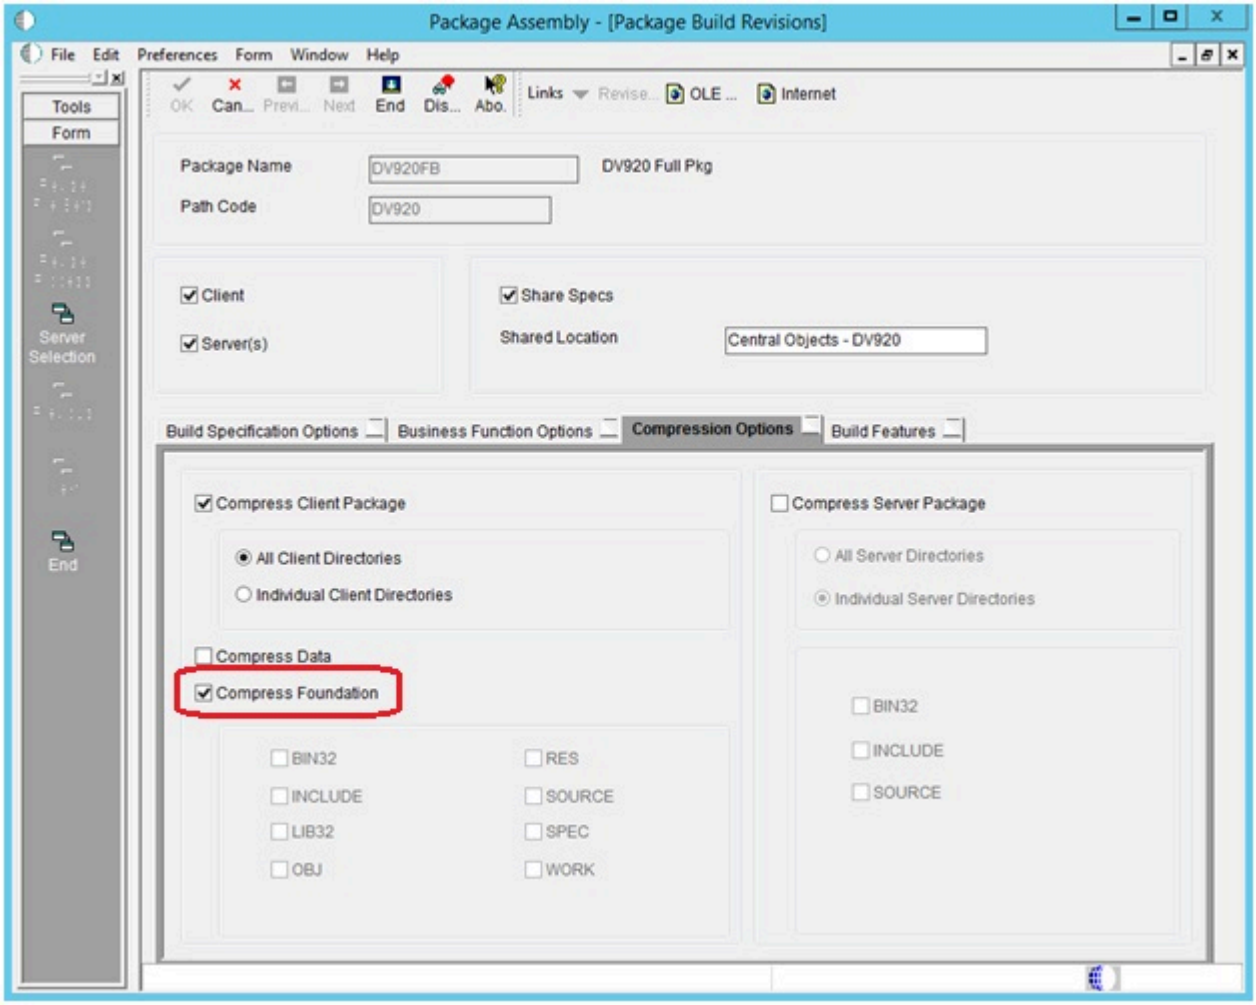

Note: If you encounter NER failures during your full package build (server side only), to resolve these failures refer to the My Oracle Support Doc Id 1950295.1 at this link: *https://support.oracle.com/epmos/faces/DocumentDisplay?id=1950295.1*

Build a New Server Package

As a best practice it is strongly recommended that you build a new server package. For multiple Enterprise Server environments, the procedures in this section are mandatory.

On Package Build, Package Build Revisions, in addition to the normal compression options, you must select these additional options during the package build process for any full package:

- **Compress Server Package**

- **All Directories**

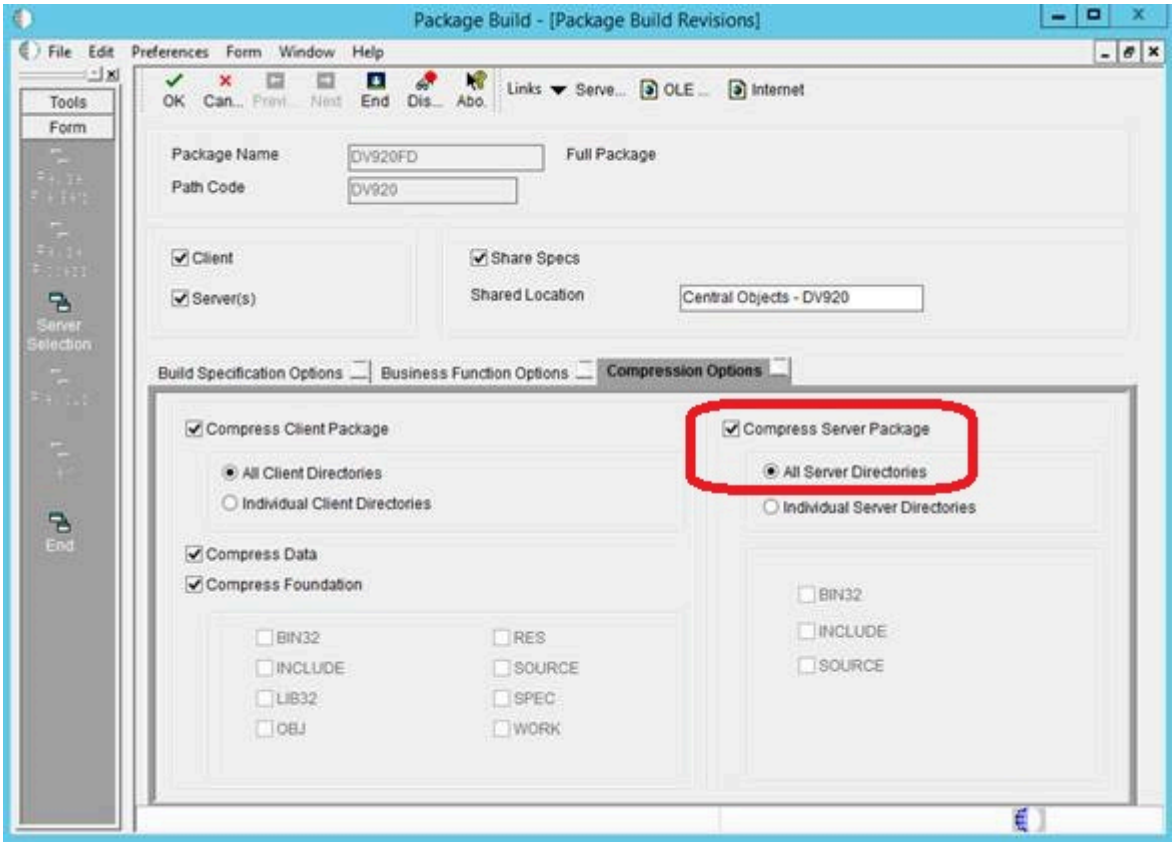

Performing Post Installation for the Standalone Deployment Server with Oracle Database

This section shows you how to perform the post installation tasks for the standalone Deployment Server with Oracle database.

> **Note:** This section is *only* applicable if you have deployed a standalone Deployment Server with Oracle database using One-Click Provisioning.

You can use One-Click Provisioning to install a standalone Deployment Server. A standalone Deployment Server is deployed when you first use One-Click Provisioning to provision a Deployment Server only, and then use One-Click Provisioning to provision other servers such as the Database Server, Enterprise Server, and HTML Server. This section describes how to complete the postinstallation steps for such a standalone Deployment Server.

If you have used One-Click Provisioning to deploy a standalone Deployment Server (that is, you did not use One-Click to initially deploy a complete system), you must perform the following procedure on your Deployment Server to enable access by JD Edwards EnterpriseOne.

Enable DEP920 Login for an Oracle Database

> **Note:** This section is only applicable if your deployment of JD Edwards EnterpriseOne includes an Oracle Database Server.

After you have used One-Click Provisioning to deploy the complete suite of JD Edwards EnterpriseOne servers, you need to enable EnterpriseOne access to the Oracle database from the Deployment Server.

1. Access the Deployment Server as described in the preceding section "Accessing the Deployment Server".
2. Edit the jde.ini file, which is typically located in this directory:

   ```
   c:\Windows
   ```
3. Edit the jde.ini file to provide a valid value for the **SecurityServer**= setting. This is the machine name for your Enterprise Server. For example:

   ```
   [SECURITY]

   SecurityServer=oraes
   ```
4. Save and close the `jde.ini` file.
5. Navigate to the location of the Oracle Client that is installed on your Deployment Server. For example:

   ```
   C:\JDE\oracle1212\product\client_1\network\admin
   ```
6. In the above directory, edit the `tnsnames.ora` file to add the listener details. You can copy these details from the `tnsnames.ora` file on your Server Manager machine or your Enterprise Server machine. This `tnsnames.ora` file is located in the same location on each machine. From the Server Manager, use this navigation to find this `tnsnames.ora` file on either the Server Manager machine or the Enterprise Server:

   From the HTML Server, select DATABASE, and then select Tnsnames.

   For example, this is the listener section you should copy:

   ```
   JDEORCL =
    (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = oradb.subnet.vcn.com)(PORT = 1521))
    (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = jdeorcl.subnet.vcn.com)
    )
   )
   ```
7. Save and close the `tnsnames.ora` file.
8. Navigate to the location of the E1Local Network, which is typically:

   ```
   C:\JDE\DEP\Oracle\E1Local\NETWORK\ADMIN
   ```
9. In the above directory, edit the `tnsnames.ora` file to add the same listener details that you obtained and copied in a previous step in this procedure.
10. Save and close the `tnsnames.ora` file.
11. Validate the configuration of the tnsnames.ora file by running `ActiveConsole.exe` and logging in to the DEP920 environment.

Considerations for Development Client

This section explains the required considerations before you create the Development Clients.

Prerequisites

- If you want to install a JD Edwards EnterpriseOne Development Client that will work with JD Edwards EnterpriseOne One-Click Provisioning, you must manually provision an additional Microsoft Windows machine. After the Microsoft Windows machine is provisioned, the process to set up and install a Development Client is the same regardless of where the client is installed.
- You must also install all the required third-party software and build a full client package on the Deployment Server.
- Refer to the *JD Edwards EnterpriseOne Development Client Installation Guide* for detailed instructions.

Considerations for Development Client

Consider the following requirements before you create the Development Clients:

- Microsoft Windows machines must be on the same private network on which the JD Edwards EnterpriseOne Database Server and Enterprise Server exist, and the machines must be able to communicate with those servers.
- Ensure that a full package build is completed on the Deployment Server. This package build creates the Development Client package installation. The delivered FA package will technically install a Development Client, but such a Development Client is not suitable for most developer activities.
- Ensure that IPv6 is disabled on the Microsoft Windows machine on which you will install the JD Edwards EnterpriseOne Development Client.
  Run this command to disable IPv6:
  ```
  reg add hklm\system\currentcontrolset\services\tcpip6\parameters /v DisabledComponents /t REG_DWORD /d 0xFFFFFFFF
  ```
  Use **ipconfig** to verify that the status of IPv6 is set to disabled.

Obtaining and Installing CA Certificates in the Oracle WebLogic Servers and the Deployment Server

This section shows you how to obtain and install CA Certificates in the Oracle WebLogic Servers and the Deployment Server.

The deployment of JD Edwards EnterpriseOne One-Click Provisioning includes temporary Certificate Authority (CA) certificates. Because these certificates are set to expire at preset and non-extendable times, you must obtain and install your own CA certificates. These must be certificates that are verified by a verified CA authority such as Entrust and Symantec Corporation.

Prerequisite

Installed Java Keystore.

The following outlines the general procedure to create a Keystore and to generate a Certificate Signing Request (CSR).

1. In your local environment, obtain and install a Java Keystore. This is a repository for security certificates – either authorization certificates or public key certificates – plus corresponding private keys. These keys are used for SSL encryption by the Oracle WebLogic Server. A file with extension jks serves as keystore.
2. From the Keystore, generate a Certificate Signing Request (CSR).
3. Export the Certificate Signing Request (CSR).
4. Validate the CSR. For example, you could use the validation tools provided by Symantec such as "checker".
5. Submit the CSR to the Certificate Authority such as Entrust or Symantec Corporation.
6. Upon return receipt, import the validated certificates to the Keystore for each server. That is, each server that must communicate with another must have its own certificate ***plus*** that of the target server. In this case, the HTML Server must have its own certificate plus that of the AIS Server, and vice versa.
7. Logged in as the WebLogic Adminstrator, you must manually modify each of these instance of Oracle WebLogic Server within your One-Click deployment to use the new Keystore:
   - Server Manager Console
   - Each instance of a JD Edwards EnterpriseOne HTML Server (JAS)
   - Each instance of a JD Edwards EnterpriseOne AIS Server
8. You should also modify the parameters using Server Manager to use the https connection for communication between the HTML Server and the AIS Server.
   a. In the **HTML instance**, modify the following Web Runtime parameters to use https, a fully qualified domain name, and https port:

   b. In the **AIS instance**, modify the following HTML Server parameters to use https, a fully qualified domain name, and https port:

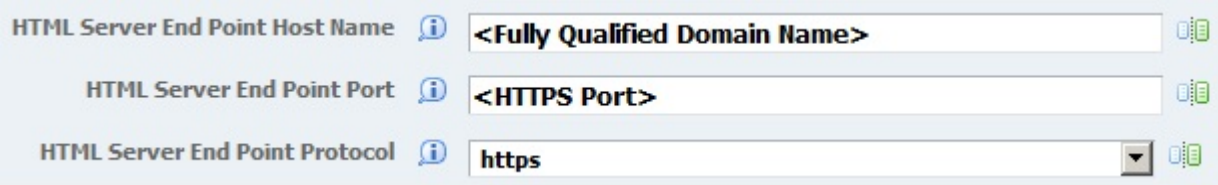

Tip: After you have SSL configured and tested, it is recommended that you disable all the non-SSL ports.

For additional details on working with CA certificates on your Oracle WebLogic Server, refer to this guide: *Administering Security for Oracle WebLogic Server.*.

Enabling Oracle Database Connection Pooling

Because this functionality applies equally to Oracle Cloud Infrastructure and On-Premises, beginning with Tools Release 9.2.6 this content resides in the ***Server Manager Guide*** , in the chapter entitled: "Enabling Oracle Database Connection Pooling (Release 9.2.6)".

16 Accessing the Provisioned Servers

Performing Post-Provisioning Tasks

This section shows you how to perform the post-provisioning tasks.

Prerequisite

Successfully completed the Deployment of a Quick Start or an Advanced Deployment Plan using the JD Edwards One-Click Provisioning Console.

Enabling Enterprise Server Predefined Ports

You must enable Predefined Ports in the `jde.ini` file on the Enterprise Server. Use Server Manager to ensure this setting exists:

```
[JDENET]

enablePredefinedPorts=1
```

For Release 9.2, the `serviceNameListen` value is 6017. The next port for use by net services on the Enterprise Server is defined by the value `serviceNameListen+1` until the value set by `maxNetProcesses=` is reached, where the default is value is 6 net processes. Therefore, the default value of 6 net process means you must open ports in the firewall on the Enterprise Server 6018-6023 for use by net processes. These ports cannot be in use by any other process running on the Enterprise Server. If they are already in use, then whatever program is using them must be reconfigured to use different ports.

Starting the Database in the Secondary Node in RAC

If you are using DB Systems for your database, and if you have enabled RAC within it, if you have followed the recommendation in this Learning Path you have previously stopped the database in the Secondary Node in RAC. At this point you must start the database in the Secondary Node for RAC. The recommended procedure is to use a line command as shown below.

```
srvctl start instance -d <DB_UNIQUE_NAME> -i <DB_NAME>
```

where `<DB_UNIQUE_NAME>` is the name you determined in the preceding OBE in this Learning Path entitled: xxxx and in the section within that OBE entited: ***Setting Up the Secondary Node.***

where `<DB_NAME>` is name you have given the database and which is appended with a 2. For example, if you gave the database name as ORCL, then the database name in the Secondary Node is named ORCL2.

> **Note:** At this point both nodes of the RAC in DB Systems are running. Therefore, when transactions are submitted to the database by any JD Edwards EnterpriseOne server, the transaction can go to either node (Primary or Secondary). If for any reason either node goes down, the transactions from JD Edwards EnterpriseOne will be served by the node that is running. However, if the node that went down was the node in which the database connection failed while it was running a transaction, then there is lag of operation while the system switches the transactions to the node that is running. Further, if the Database is down on both nodes, and if in start up of the nodes there is delay of more than 75 seconds, then you must manually restart both the HTML Server and the Enterprise Server after the database is properly started.

To view details of the Secondary Node in RAC:

1. On the Oracle Cloud Infrastructure Console Home page, click the **Navigation** Menu in the upper-left corner.

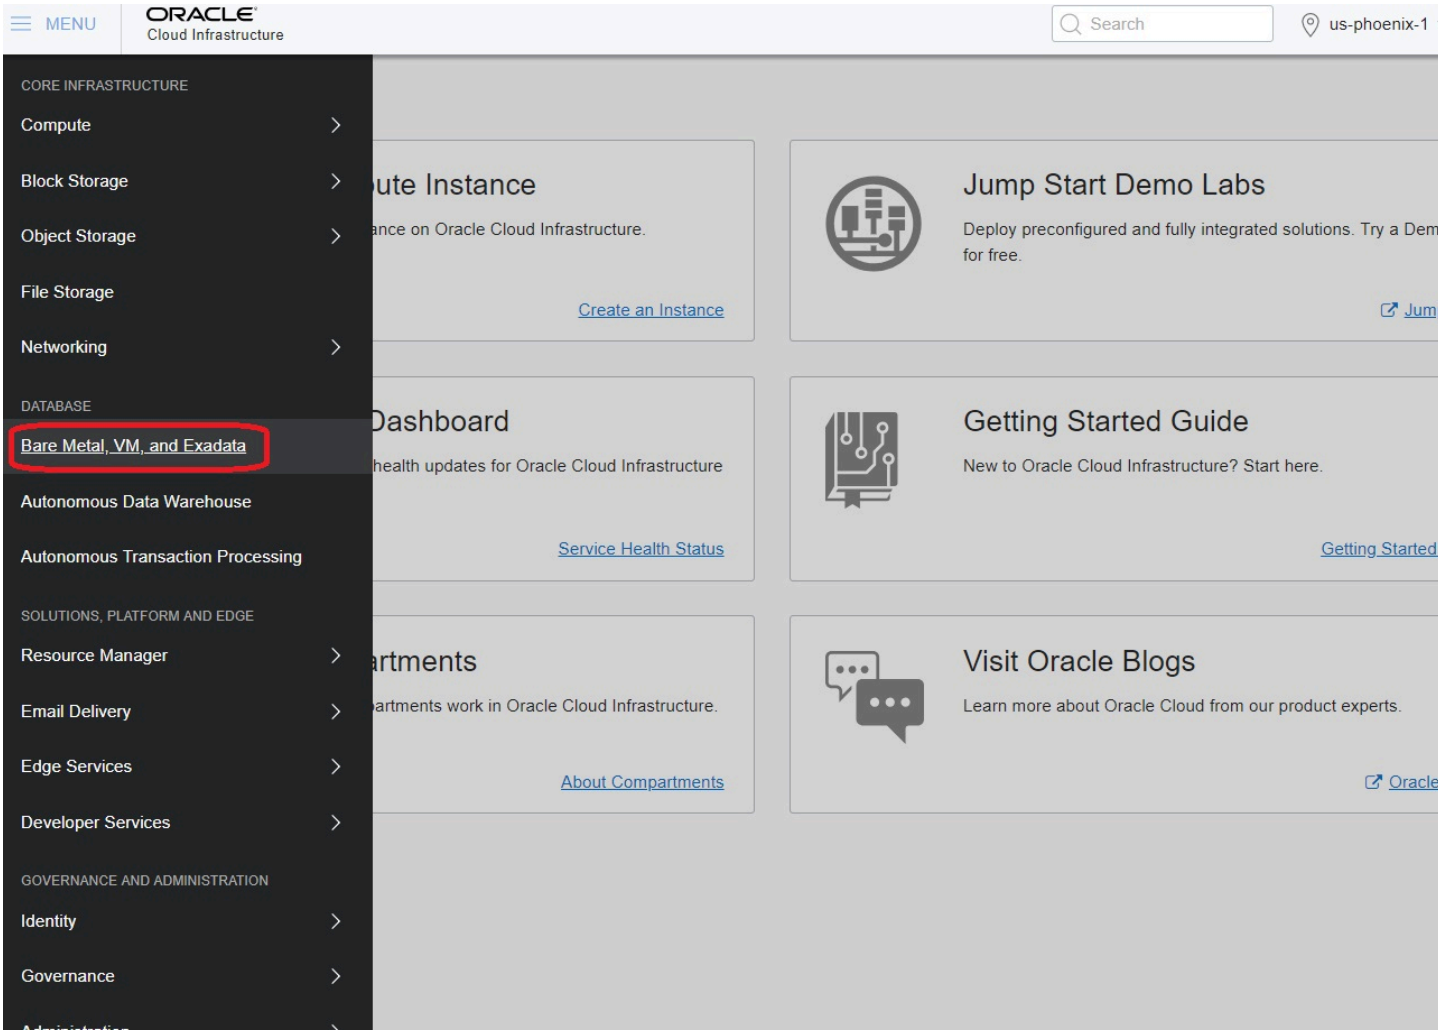

2. From the Navigation Menu, in the **Database** section, click to select **Bare Metal**, **VM**, and **Exadata**.
3. In the left panel ensure that DB Systems is selected, and that under **List Scope** that the applicable **COMPARTMENT** is selected which contains the DB System with RAC that you created.
4. Click the link for your DB System with RAC.

5. In the **Resources** section in the left pane, click **Nodes**.

Details for two Nodes are displayed as shown below:

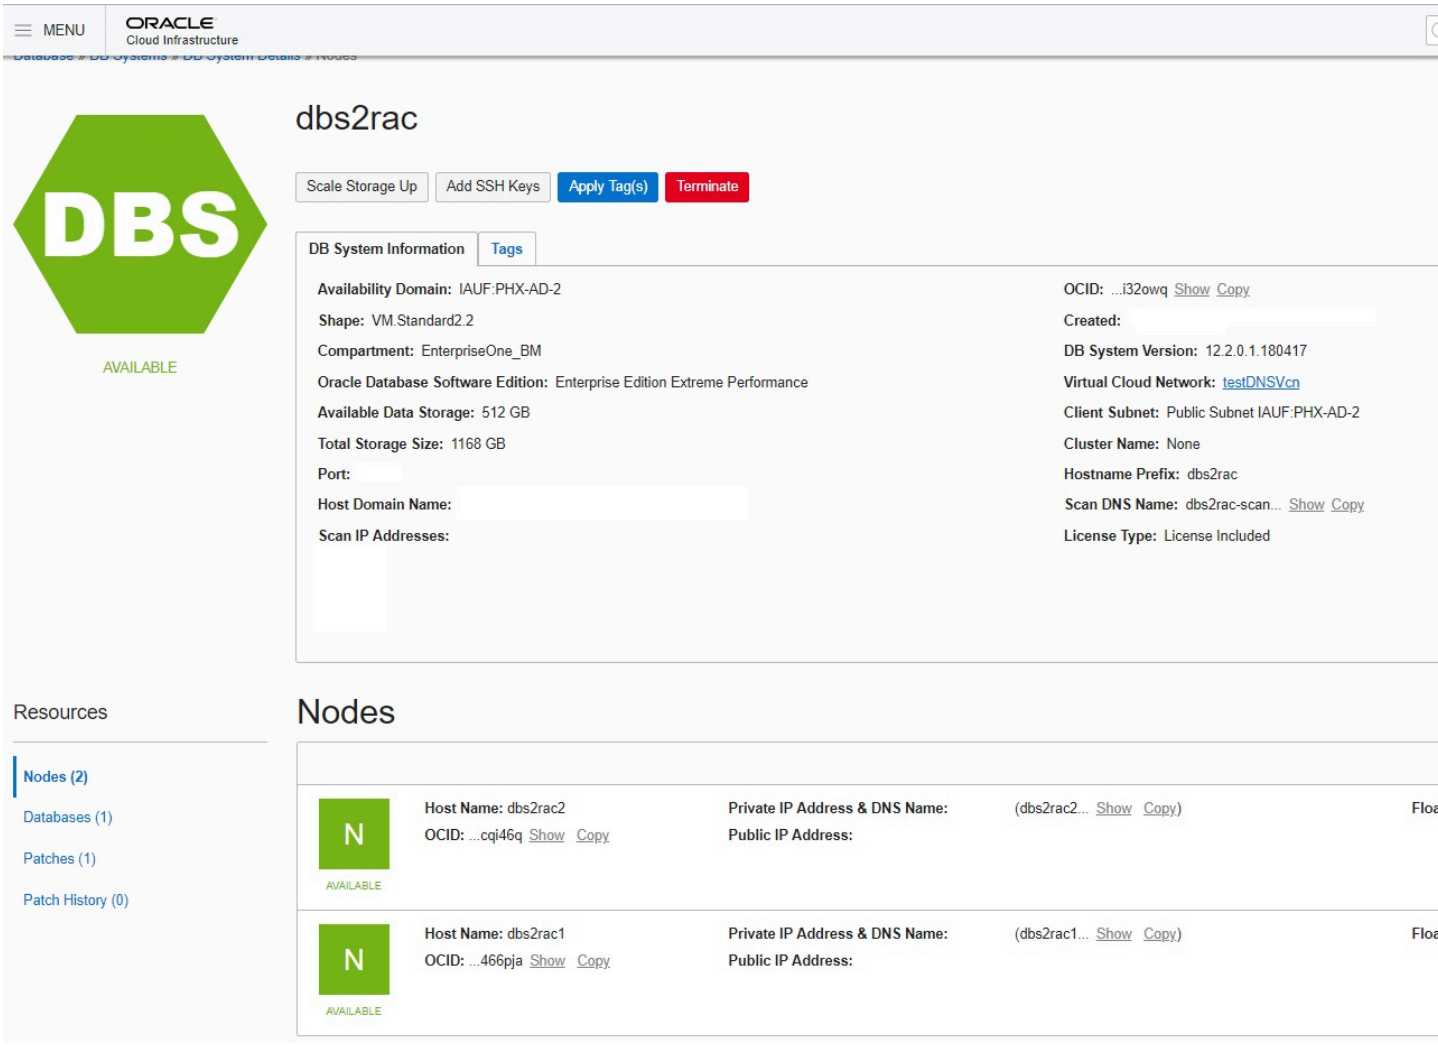

Setting up Additional Security Configurations

As part of meeting the CIS Benchmarks for secure Linux Machines, you can further secure your One-Click Provisioned Linux environments by performing the below steps:

These settings in the `sshd_config` file are located at `/etc/ssh`.

- **Disable TCP Forwarding**

  Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

  ```
  AllowTcpForwarding=no
  ```

- **Disable Root Login**

  Edit the `/etc/ssh/sshd_config` file to set the parameter as follows:

  ```
  PermitRootLogin=no
  ```

  Ensure that there are no other entries that might override the above settings. For the settings to reflect, restart the SSH Service using the below command:

  ```
  sudo systemctl restart sshd
  ```

Closing Port 22 on Public IP for Enhanced Security

For enhanced security, it is highly recommended to close Port 22 (SSH) in both the Security List (VCN/Subnet level) and the Network Security Group (NSG) after provisioning is complete. This minimizes potential vulnerabilities and aligns with a stronger security posture.

Follow the below steps to close Port 22 in Security List and NSG:

- Security List:
  a. Navigate to the OCI Console > Networking > Virtual Cloud Networks (VCN)
  b. Select the VCN associated with your instance.
  c. Open the Security List attached to the subnet.
  d. Locate the ingress rule for Port 22 (TCP) and delete it.

- Network Security Group (NSG):
  a. Navigate to the OCI Console > Networking > Network Security Groups
  b. Select the NSG attached to your instance.
  c. Locate the ingress rule allowing Port 22 (TCP) and delete it.

Accessing the JD Edwards EnterpriseOne Servers Using Their Public IP Addresses

This section shows you how to access the JD Edwards EnterpriseOne Servers using their public IP addresses.

Prerequisites

- The user interface for the Oracle Cloud Infrastructure Console is constantly evolving. For the most up-to-date descriptions and navigation, refer to *Using the Console.*
- You should have a fundamental understanding of the Oracle Cloud Infrastructure. It is highly recommended that you review the extensive collateral information, including training, at this site: *Oracle Cloud Infrastructure*
- You must have a subscription and an Administrator account to Oracle Cloud Infrastructure as described at this site: *Getting Started with Oracle Cloud*
- Successfully completed the Deployment of a Quick Start or an Advanced Deployment Plan using the JD Edwards One-Click Provisioning Console.

After the successful deployment of all the Servers according to the Orchestration created using the JD Edwards One-Click Provisioning Console, you can access these deployed JD Edwards EnterpriseOne Servers and Development Client Using their public IP addresses:

- Server Manager Console
- HTML Server
- AIS Server
- Deployment Server

Server Manager Console

After the Server Manager Console is started, you can access it using its public IP address using this syntax:

http://<Public IP Address for Server Manager>:8999/manage

For example:

http://111.11.11.111:8999/manage

> **Note:** You should always use HTTP to access the Server Manager Console for normal operations. Although the Server Manager Console is also enabled for HTTPS/SSL protocol, such access is restricted for internal use as part of the REST/API functionality, and should not be used for normal operations.

To acquire the public IP address assigned to the instance:

1. Access Oracle Cloud Infrastructure and go to **Compute**, and select the **instance** menu.The system displays the list of instances.
2. Click the **host name** of your provisioning server in the displayed list. For example, click ProvisionPrime.

The system now displays the details of the instance where you can find the public IP address of the instance.

HTML Server

After your HTML Server is deployed, you can access it using its public IP address using this syntax:

https://<Public IP Address for HTML Server>:<port as defined in orchestration>/jde/E1Menu.maf

> **Note:** The file name E1Menu.maf is case-sensitive and must be specified exactly as shown here.

For example:

https://111.11.11.111:8001/jde/E1Menu.maf

To acquire the public IP address assigned to the HTML instance:

1. Access Oracle Cloud Infrastructure and go to **Compute**, and select the **instance** menu.The system displays the list of instances.
2. Click the **host name** of your provisioning server in the displayed list. For example, click WeblogicServer.

The system now displays the details of the instance where you can find the public IP address of the instance.

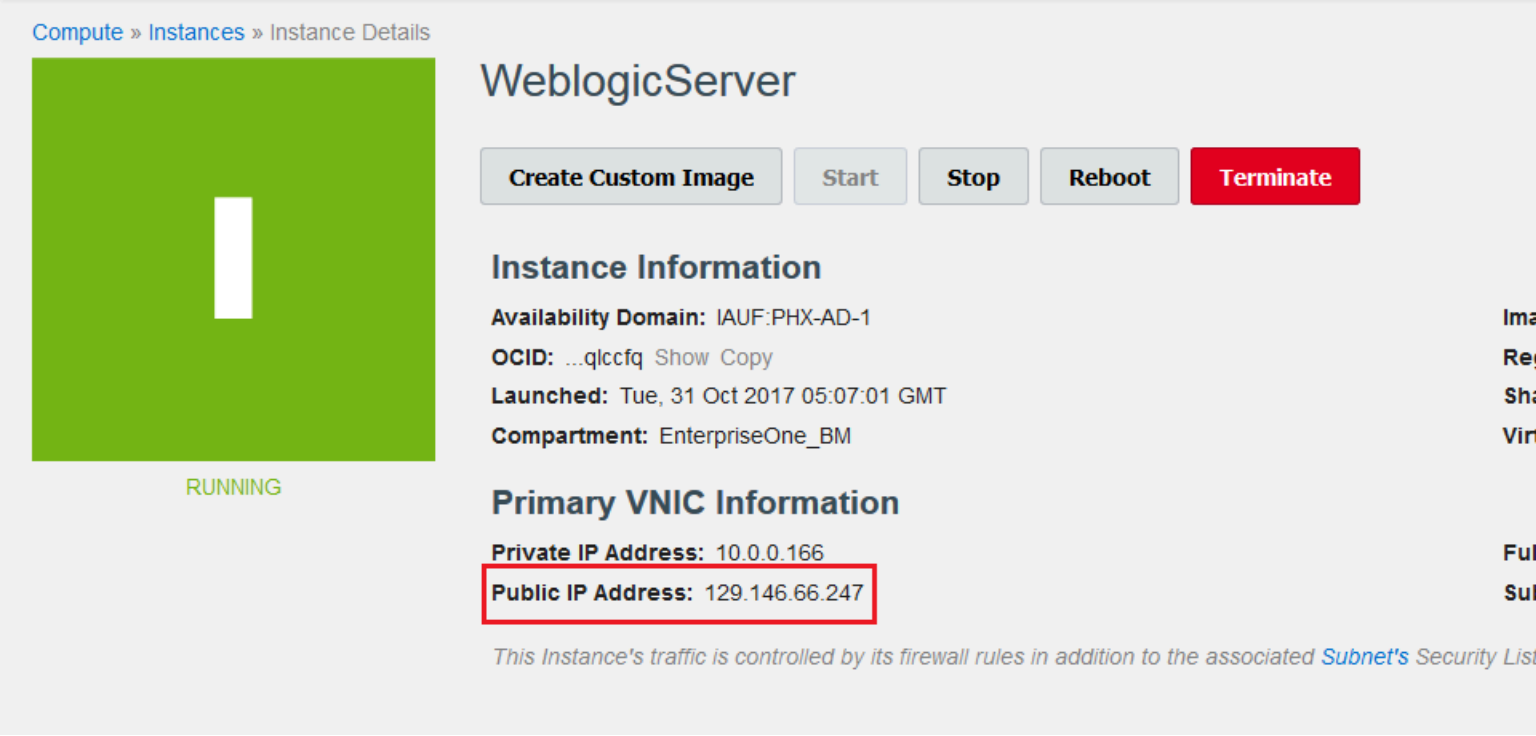

AIS Server

After your AIS Server is deployed, you can access it using its public IP address using this syntax:

https://<Public IP Address for AIS Server>:<port as defined in orchestration>/jderest/defaultconfig

For example:

https://111.11.11.111:8002/jderest/defaultconfig

To acquire the public IP address assigned to the AIS instance:

1. Access Oracle Cloud Infrastructure and go to **Compute**, and select the **instance** menu.The system displays the list of instances.
2. Click the **host name** of your provisioning server in the displayed list. For example, click WeblogicServer.

The system now displays the details of the instance where you can find the public IP address of the instance.

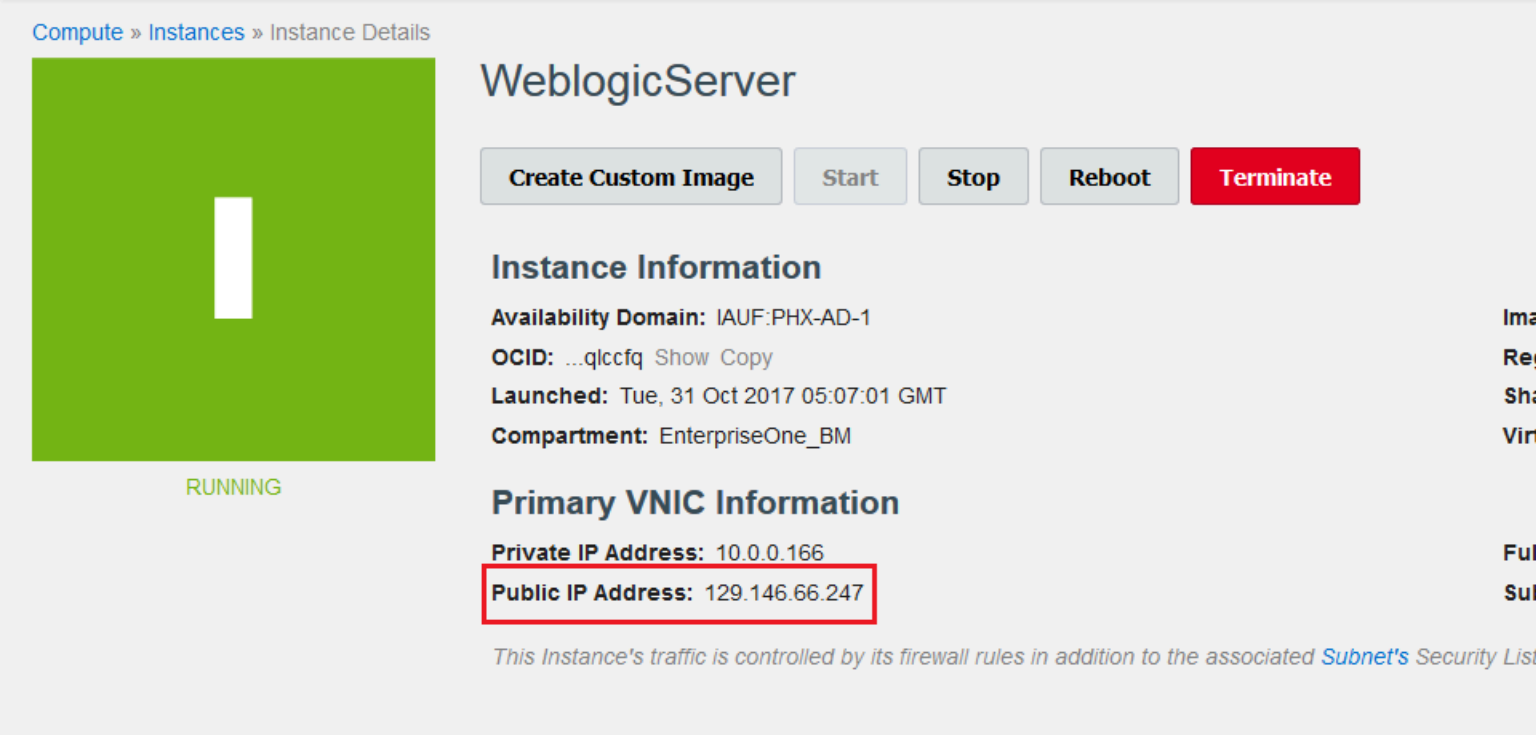

Deployment Server

You can access your Deployment Server using Microsoft Windows Remote Desktop Protocol (RDP). To do so, you will need the public IP address of this instance.

To acquire the public IP address assigned to the Deployment server instance:

1. Access Oracle Cloud Infrastructure and go to **Compute**, and select the **instance** menu.The system displays the list of instances.
2. Click the **host name** of your provisioning server in the displayed list. For example, click DeploymentServr.

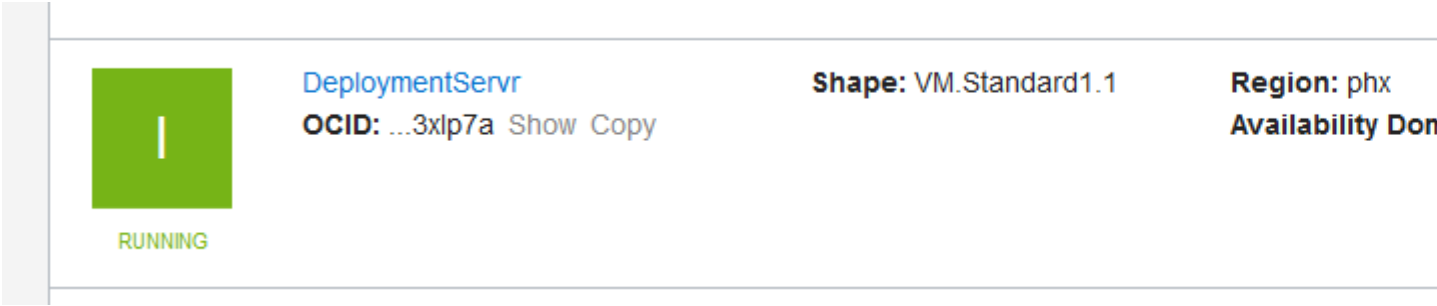

The system now displays the details of the instance where you can find the public IP address of the instance.

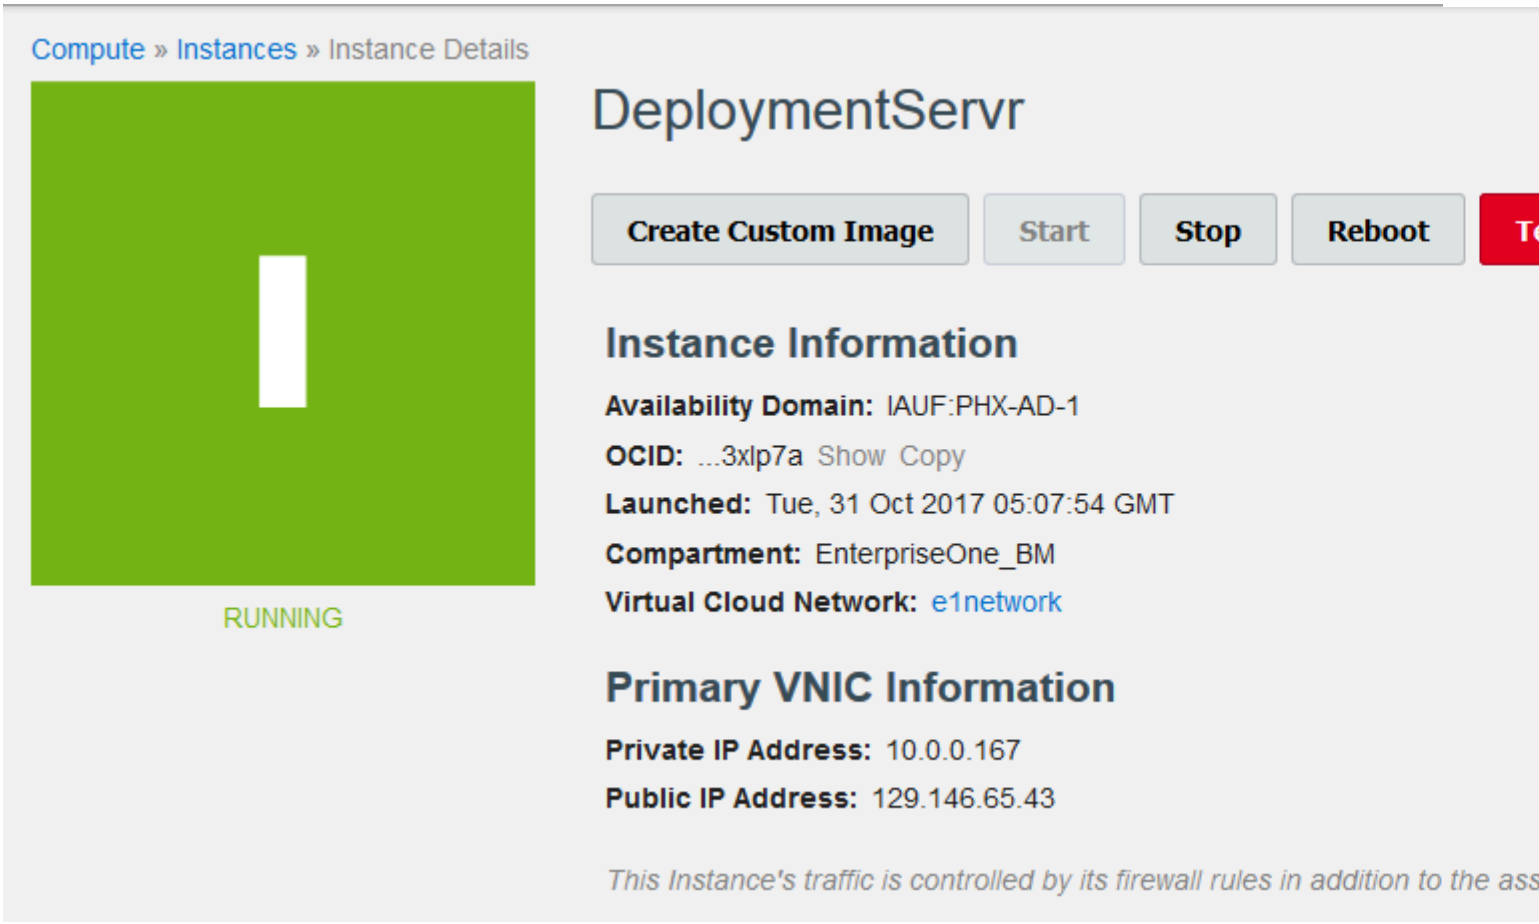
Compute » Instances » Instance Details
DeploymentServr
Create Custom Image
Start
Stop
Reboot
RUNNING
Instance Information
Availability Domain: IAUF:PHX-AD-1
OCID: ...3xlp7a Show Copy
Launched: Tue, 31 Oct 2017 05:07:54 GMT
Compartment: EnterpriseOne_BM
Virtual Cloud Network: e1network
Primary VNIC Information
Private IP Address: 10.0.0.167
Public IP Address: 129.146.65.43
This Instance's traffic is controlled by its firewall rules in addition to the ass

17 Troubleshooting Your One-click Deployment

Troubleshooting

This section shows you the list of troubleshooting topics and resolutions.

One-Click Provisioning Console Troubleshooting Tips

This section contains the following troubleshooting tips for the EnterpriseOne One-Click Provisioning Console:

- Accessing a Previously Saved Provisioning Console Configuration
- Console Administration
- Console Service Status
- Start the Console
- Stop the Console
- View the Console Logs

Accessing a Previously Saved Provisioning Console Configuration

Symptom:	You cannot access a previously created Quick Start Deployment Plan.
Cause:	After you created a Quick Start Deployment Plan (either complete or partial), you exited the browser used to connect to the One-Click Provisioning Console.
Resolution:	When you next access the Provisioning Console, you will be prompted to enter the same credentials that you entered for the first access to the console. Those credentials are the password for the Server Manager for JD Edwards EnterpriseOne. After the credentials are validated, you can access the previously saved configuration.

Console Administration

The One-Click Provisioning Console process is configured as a system service in the One-Click Provisioning Server image. The configuration file (E1CloudConsole.service) is located in this directory:

```
/etc/systemd/system/E1CloudConsole.service
```

Console Service Status

Run following command in the shell to display the One-Click Provisioning Console service status:

```
$ sudo systemctl status E1CloudConsole
```

Start the Console

Run the following command in the shell to start the One-Click Provisioning Console service.

```
$ sudo systemctl start E1CloudConsole
```

Run the following command to verify the service status within same machine:

```
$ curl -kI https://localhost:3000
```

The response from the Curl command should be similar to this:

```
[opc@multitest ~]$ curl -kI https://localhost:3000//</span>\\\\ HTTP/1.1 200 OK
X-DNS-Prefetch-Control: off
X-Content-Type-Options: nosniff
X-Frame-Options: SAMEORIGIN
Strict-Transport-Security: max-age=86400
X-Download-Options: noopen
X-XSS-Protection: 1; mode=block
Accept-Ranges: bytes
Cache-Control: public, max-age=0
Last-Modified:
ETag: W/"5862-156cad16600"
Content-Type: text/html; charset=UTF-8
Content-Length: 22626
Date:
Connection: keep-alive
```

Outside the machine, you can access the One-Click Provisioning Console application using the below URL:

```
https://<#PUBLIC_IP#>:3000/
```

Stop the Console

Run following command in the shell to stop the One-Click Provisioning Console service.

```
$ sudo systemctl stop E1CloudConsole
```

View the Console Logs

Run following command in the shell to view the logs for the One-Click Provisioning Console.

```
$ sudo journalctl -u E1CloudConsole
```

Press **Page-Down** to scroll through the logs.

Enter **Shift + G** to scroll to end of the log.

> **Note:** The One-Click Provisioning Console logs are completely regenerated each time the Console is run.

Enterprise Server Provisioning Fails - Communication Failure

Symptom:	Provisioning fails to deploy the Enterprise Server with an error returned indicating a communication failure.
Cause:	A timeout may have occurred during communications between the One-Click Provisioning processes and the Oracle Cloud Infrastructure Services.
Resolution:	Restart the deployment of One-Click Provisioning.

Enterprise Server Provisioning Fails - Health Check Fails

Symptom:	Porttest remotely executes, but fails.
Cause:	Possibly a timing issue if the kernels are not starting up fast enough to execute a porttest successfully.
Resolution:	If host name command in the Enterprise Server returns the FQDN then set the short hostname by using the following command: `sudo hostnamectl set-hostname <short host name>` Rerun the Deployment Server again. It should quickly skip to the Ent Health Check and pass the health check.

HTML Server Intermittently Unavailable

Symptom:	There is an intermittently occurring issue with the HTML server provisioned on the WebLogic Server going down and users not being able to access the EnterpriseOne HTML login page. Generally when this issue occurs, the below error is displayed in the WebLogic Server logs or in the HTML Server logs: **java.lang.OutOfMemoryError: PermGen space**

Resolution:	To bring up the EnterpriseOne HTML server when it goes down with the above error, try one of the following steps:
	
	1. Restart the EnterpriseOne HTML Server instance from the Server Manager Console and then check if the EnterpriseOne HTML login page is accessible.
	2. If the JD Edwards EnterpriseOne HTML Server login page is not accessible even after restart, then set the below JVM arguments in the HTML Server from the WebLogic Server console and then restart the HTML Server.
	a. Log in to the WebLogic Server console.
	b. Go to the Servers on the environment tab and then select the HTML Server.
	c. Go to Server Start tab and click the **Lock & Edit** button in the left upper corner.
	d. In the Arguments section, change the values for the following arguments from their existing value to 1024m:
	
	`-XX:PermSize=1024m`
	
	`-XX:MaxPermSize=1024m`
	e. In the Arguments section, go to the end and enter one line space and paste the below arguments:
	
	`-XX:+CMSClassUnloadingEnabled`
	
	`-XX:+CMSPermGenSweepingEnabled`
	
	`-XX:+UseConcMarkSweepGC`
	f. Click the **Save** button.
	g. Click the **Activate Changes** button in the left upper corner.
	h. Restart the EnterpriseOne HTML Server.
	3. If the JD Edwards EnterpriseOne HTML Server login page is still not accessible, then free up some resources (OCPU and memory) from the cloud domain and restart the EnterpriseOne HTML Server from the Server Manager console.

Server Manager Console Fails to View Log File and Download Log for Enterprise Server and HTML Server

Symptom:	Server Manager Console Fails to View Log File and Download Log for Enterprise Server and HTML Server.
Resolution:	You should always use the HTTP protocol to access the Server Manager Console for normal operations. Although the Server Manager Console is also enabled for HTTPS/SSL protocol, such access is restricted for internal use as part of the REST/API functionality, and should not be used for normal operations.

PDB on Compute Not Open

The Console checks to ensure the Pluggable Database (PDB) on the Oracle Cloud Infrastructure is open (set to READ WRITE) as you begin to provision it. If the PDB is not open, the Console scripts will issue commands to open it. However, if the PDB is still closed (not set to READ WRITE) when the deployment commences, the deployment will fail. Use the following procedure to determine if the PDB is open, and if not how to set it to open.

1. Connect to the database using this command:

   ```
   sqlplus '/as sysdba';
   ```
2. Verify the pluggable database **JDEORCL** is open using this command:

   ```
   select OPEN_MODE from v$pdbs where NAME = 'JDEORCL';
   ```
3. If the **OPEN_MODE** is set to **READ WRITE**, then the PDB is ready to be provisioned by the JD Edwards EnterpriseOne One-Click Provisioning Server.
4. If the **OPEN_MODE** is not **READ WRITE**, then the PDB is not open. Open the PDB using this command:

   ```
   alter pluggable database JDEORCL open;
   ```

Not Able to View and Download Server Manager Home Logs When Server Manager is SSL-Enabled

When Server Manager is SSL-enabled, the users may not be able to view and download the Server Manager Home logs. The root cause may be because the proper security certificates are not available. For instructions on how create and import security certificates, refer to these documents:

- *JD Edwards EnterpriseOne Tools Server Manager Guide*
- *JD Edwards EnterpriseOne Tools Security Administration Guide*

The JD Edwards EnterpriseOne HTML Server Login Fails

If the JD Edwards EnterpriseOne HTML Server login fails, run the port test on the JD Edwards EnterpriseOne Enterprise Server. If the port test fails, check the host entries in the /etc/hosts file. See **Edit the /etc/hosts File for Connectivity** section in this tutorial.

JAVA BSFN Fails in Web Server

If the host name command in the JAS Server returns the FQDN then set the short hostname by using the following command.

```
sudo hostnamectl set-hostname<short host name>
```

The JD Edwards Enterprise Server - Port Test Fails

If the port test fails on the JD Edwards EnterpriseOne Enterprise Server, check the host entries in the /etc/hosts file. See **Edit the /etc/hosts File for Connectivity** section in this tutorial.

Resolving Hostname Issues During Provisioning

If you encounter hostname issues during orchestration because the provisioning was run against machines with the long hostnames, FQDNs, or both, then use the following procedure to resolve the issue:

1. On the operating system level, change the hostnames of all servers in the on-premises architecture to use a short hostname without the FQDN.
2. On the database level, change all the hostnames in JDE tables to use short hostnames without using the FQDN.
3. Ping all the machines with short hostnames and verify that the ping test is successful.

See *Understanding Machine Names* for more information.

Health Checks Fail as Final Step in OCI Provisioning

If the health checks fail during the final step in OCI Provisioning (as shown in the example below), it is likely because the predefined ports for the Enterprise Server are either not enabled or not properly defined in the VCN.

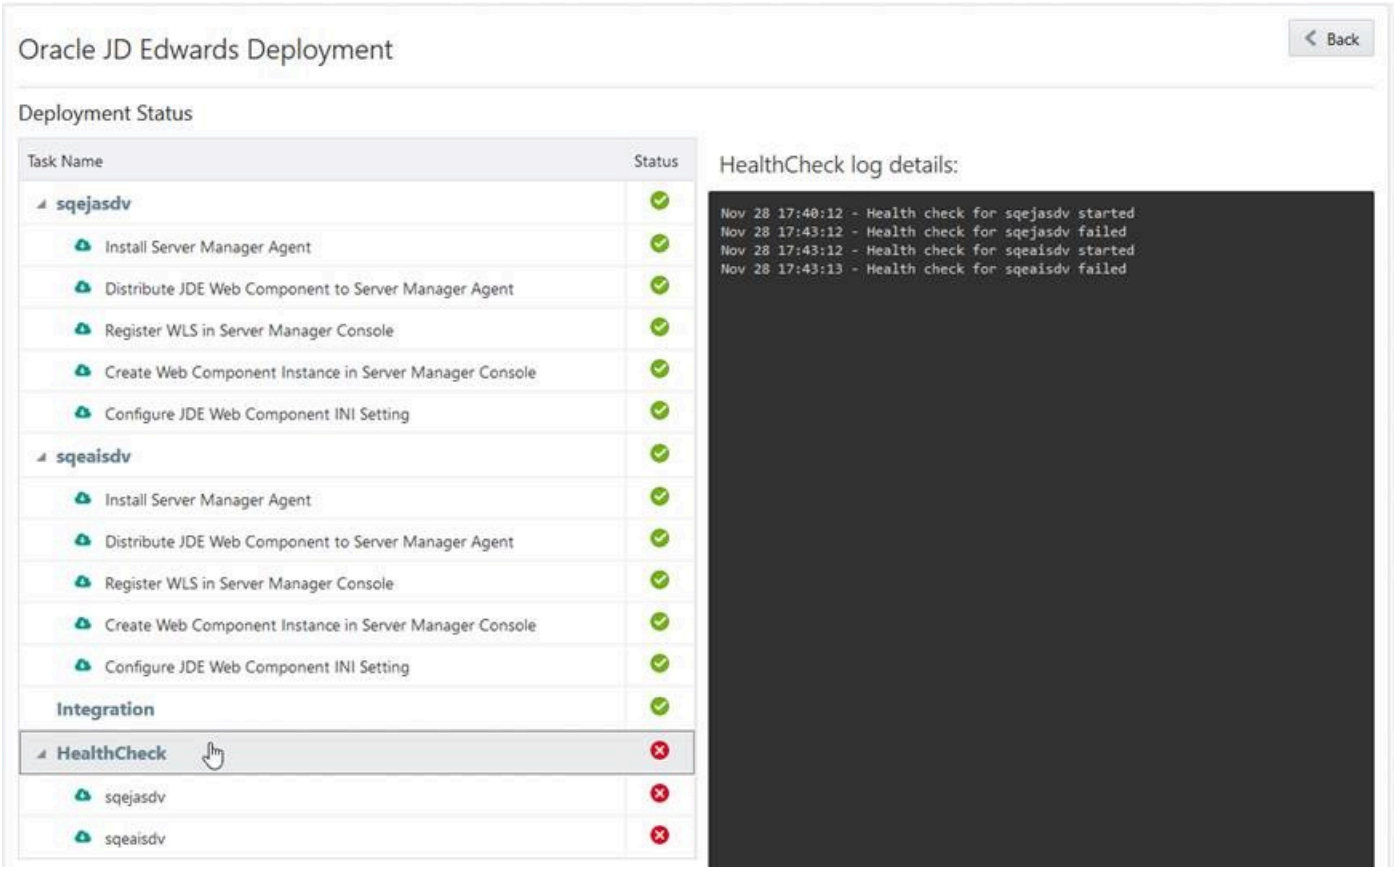

Security Server JdeSocket Or SocketException Errors When Logging In To EnterpriseOne

If a JD Edwards EnterpriseOne Server is deployed inside a firewall-protected network, you must enable predefined ports so that all NET server ports are allocated within the predefined port range.

Error ORA-00904: "NRUSER1": Invalid Identifier for Table F980051

The JAS log files displays the following error after you install EnterpriseOne using the JD Edwards One-Click Provisioning Console with Tools 9.2.2.4:

```
[SEVERE] - [JDBJ] SQLException occurred in the SQLPhysicalConnection.select():Auto: | Table or View Name
= F980051 - Data Source[0] = System - 920 java.sql.SQLSyntaxErrorException: ORA-00904: "NRUSER1": invalid
identifier

[SEVERE] - [BASE] com.jdedwards.database.base.JDBException: [SQL_EXCEPTION_OCCURRED] An SQL exception occurred:
ORA-00904: "NRUSER1": invalid identifier . java.sql.SQLSyntaxErrorException: ORA-00904: "NRUSER1": invalid
identifier

[SEVERE] - [BASE] com.jdedwards.database.base.JDBException: [SQL_EXCEPTION_OCCURRED] An SQL exception occurred:
ORA-00904: "NRUSER1": invalid identifier . com.jdedwards.database.base.JDBException: [SQL_EXCEPTION_OCCURRED] An
SQL exception occurred: ORA-00904: "NRUSER1": invalid identifier
```

Cause:

The delivered image for One-Click Provisioning has incorrect OCM mappings for the F980051 table. The OCM for this table is pointing to the System - 920 database data source where the structure of the table is not consistent with the table specs.

Solution:

1. Check both the System and the Server Map data source OCM mappings and change the OCM mappings as shown in the following list:
   Use the following correct OCM mappings for F98005* tables:
   - F980051 (Bell Status) mapped to Control Tables - <Environment>
   - F980052 (Notification Execution History) mapped to Control Tables - <Environment>
   - F980053 (Subscription Metadata) mapped to Control Tables - <Environment>
   - F980054 (Notification Queue) mapped to System - 920
   - F980055 (Notifications Constants) mapped to System - 920
   - F980056 (Notifications Offline Repository) mapped to System - 920
   - F980057 (Notifications Offline Repository Archive) mapped to System – 920
2. Restart the Enterprise service.

For more information, see Doc ID 2408472.1 on My Oracle Support.

Error in Server Manager Console Enterprise Server Process Detail Graph

Symptom:	Unable to display the results of the selected monitor. The monitor is no longer available for viewing.

Cause:	The monitors were registered for each of the Enterprise Server instances but the Server manager Console did not identify these monitors, and hence caused an issue related to the graphs.
Resolution:	Verify monitors.xml located in the path \SCFMC\targets\home\ , and check if the entries are recorded for all the required monitors. If the graphs are not visible for a particular monitor even if the monitor's entry is recorded in the monitors.xml file, then restart the Server Manager Console.

Orchestrating JD Edwards EnterpriseOne Servers Returns IPV6 Is Not Disabled

Use this procedure if One-Click Provisioning returns this message for any Linux server that you are provisioning:

- IPV6 Is Not Disabled
  a. Ensure that /etc/sysctl.conf file contains below list of entries:

     ```
     net.ipv6.conf.default.disable_ipv6 = 1

     net.ipv6.conf.all.disable_ipv6 = 1
     ```

  b. If the above settings do not exist, you must update the `/etc/sysctl.conf` file to include the above settings.

     > **Note:** A mandatory reboot is necessary in order for the settings to take affect. and reboot the machine.

  c. Verify the output of below command to ensure there is no entry for the string "inet6"

     ```
     sudo ifconfig | grep inet6
     ```

  d. If the string "inet6" continues to occur in the output of the above command output, verify your settings and ensure that the machine has been rebooted.

Regenerating and Reinstalling Self-Signed Certificates for One-Click Provisioning

This section shows you how to regenerate and reinstall Self-Signed Certificates for One-Click Provisioning.

The deployment of JD Edwards EnterpriseOne One-Click Provisioning includes temporary self-signed certificates. By design, these certificates are set to expire at one-year intervals from when they were generated.

> **Note:** After self-signed certificates are expired, Oracle strongly recommends that you generate and install your own CA certificates as described in this section. These must be certificates that are verified by a verified CA authority such as Entrust and Symantec Corporation.

Prerequisite

You must have an installation of Java Keystore.

Generating Self-Signed Certificates on Linux

Use this procedure to generate self-signed certificates on Linux.

1. Log into Provisioning Server.
2. Run the following commands (where each bulleted item is one contiguous line):
   - `sudo -i`
   - `mv /u01/jde920/.vm_unconfigured /u01/jde920/vm_unconfigured`
   - `/u01/jdk1.8.0/jre/bin/keytool -delete -alias cert -keystore "/u01/jdk1.8.0/jre/lib/security/cacerts" -storepass *****`
   - `/u01/CertGen/ConfigureCertKey_CC.sh`
   - `/u01/CertGen/ConfigureCertKey_SMC.sh <WebLogic Admin Password>`
   - `mv /u01/jde920/vm_unconfigured /u01/jde920/.vm_unconfigured`

> **Note:**
> - The `storepass` value ***** in above commands is the WebLogic Admin password.
> - The `generated cert.pem` file is located in: `/u01/E1CloudConsole/keys`
> - The log file path is: `/var/log`

Importing the Self-Signed Certificate into Target Machines on Linux

The self-signed certificate file that you generated in the previous steps must be imported into these target machines that were deployed by JD Edwards EnterpriseOne One-Click Provisioning:

- HTML Server
- Enterprise Server
- Oracle Database Server

HTML Server

Use this procedure to import the self-signed certificate into the HTML Server.

1. Run this command:

   ```
   sudo -i
   ```
2. Get the `cert.pem` file from Provisioning Server, which you generated in the previous procedure in this section entitled: "Generating Self-Signed Certificates on Linux".

3. Use this command is determine if a certificate with an alias of **smcert** is already imported to `jdk/jre/lib/security/cacerts`, where this command is a single contiguous line:
   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```
   If the results of the above command indicate that an alias of **smcert** certificate is already imported to the JDK/JRE lib location, use this command to remove it, where this command is a single contiguous line:
   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -delete -file cert.pem -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```
   > **Note:** The storepass value ***** in above commands is the default password for Java trustStore.

4. Use these commands to import the certificate that you generated in the previous procedure of this section entitled: "Generate Self-Signed Certificates" to `/u01/jde_home/SCFHA/jdk`, where each command is a single contiguous line:
   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -import -file cert.pem -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```
   > **Note:** The storepass value ***** in above commands is the default password for Java trustStore.

5. Use this command is determine if a certificate with an alias of **smcert** is already imported for the JDK path:
   ```
   /u01/oracleJDE/jdk_path/jre/bin/keytool -list -v -alias smcert -keystore "/u01/oracleJDE/jdk_path/jre/lib/security/cacerts" -storepass *****
   ```
   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:
   ```
   /u01/oracleJDE/jdk_path/jre/bin/keytool -delete -alias smcert -keystore "/u01/oracleJDE/jdk_path/jre/lib/security/cacerts" -storepass *****
   ```
   > **Note:** The storepass value ***** in above commands is the WebLogic Admin password.

6. Use these commands to import the certificate that you generated in the previous procedure of this document entitled: "Generate Self-Signed Certificates" to /u01/oracleJDE/jdk_path, where each command is a single contiguous line:
   ```
   /u01/oracleJDE/jdk_path/jre/bin/keytool -import -file cert.pem -alias smcert -keystore "/u01/oracleJDE/jdk_path/jre/lib/security/cacerts" -storepass *****
   /u01/oracleJDE/jdk_path/jre/bin/keytool -list -v -alias smcert -keystore "/u01/oracleJDE/jdk_path/jre/lib/security/cacerts" -storepass *****
   ```
   > **Note:** The storepass value ***** in above commands is the WebLogic Admin password.

Enterprise Server

Use this procedure to import the self-signed certificate into the HTML Server.

1. Run this command:

`sudo -i`

2. Get the `cert.pem` file from Provisioning Server, which you generated in the previous procedure in this section entitled: "Generating Self-Signed Certificates on Linux".
3. Use this command is determine if a certificate with an alias of **smcert** is already imported to `jdk/jre/lib/security/cacerts`, where this command is a single contiguous line:

   ```
   /u01/jde920/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde920/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   /u01/jde920/jde_home/SCFHA/jdk/jre/bin/keytool -delete -alias smcert -keystore "/u01/jde920/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the default password for Java trustStore.

4. Use these commands to import the certificate that you generated in the previous procedure of this section entitled: ***Generate Self-Signed Certificates*** to `/u01/jde920/jde_home/SCFHA/jdk`, where each command is a single contiguous line:

   ```
   /u01/jde920/jde_home/SCFHA/jdk/jre/bin/keytool -import -file cert.pem -alias smcert -keystore "/u01/jde920/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   ```
   /u01/jde920/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde920/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the default password for Java trustStore.

5. Use this command is determine if a certificate with an alias of **smcert** is already imported for the JDK path:

   ```
   /u01/jdk8_32/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jdk8_32/jre/lib/security/cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of smcert certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   /u01/jdk8_32/jre/bin/keytool -delete -alias smcert -keystore "/u01/jdk8_32/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the Site Key password.

6. Use these commands to import the certificate that you generated in the previous procedure of this section entitled: "Generate Self-Signed Certificates" to `/u01/jdk8_32`, where each command is a single contiguous line:

   ```
   /u01/jdk8_32/jre/bin/keytool -import -file cert.pem -alias smcert -keystore "/u01/jdk8_32/jre/lib/security/cacerts" -storepass *****
   ```

   ```
   /u01/jdk8_32/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jdk8_32/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the Site Key password.

Database Server

Use this procedure to import the self-signed certificate into the Oracle Database Server. Note that is only applicable if you are using Oracle Compute service for your Database Server; it not applicable if you are using the Oracle Database Service (DBS).

1. Log into the Oracle Compute Database Server.
2. Get the cert.pem file from Provisioning Server, which you generated in the previous procedure in this section entitled: "Generating Self-Signed Certificates on Linux".
3. Run this command:

   ```
   sudo -i
   ```
4. Use this command is determine if a certificate with an alias of **smcert** is already imported to jdk/jre/lib/security/cacerts, where this command is a single contiguous line:

   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -delete -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the default password for Java trustStore.
5. Use these commands to import the certificate that you generated in the previous procedure of this document entitled: "Generate Self-Signed Certificates" to /u01/jde_home/SCFHA/jdk, where each command is a single contiguous line:

   ```
   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -import -file cert.pem -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****

   /u01/jde_home/SCFHA/jdk/jre/bin/keytool -list -v -alias smcert -keystore "/u01/jde_home/SCFHA/jdk/jre/lib/security/cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the default password for Java trustStore.

Generating Self-Signed Certificates on Microsoft Windows

Use this procedure to generate self-signed certificates on Microsoft Windows.

1. Log into Provisioning Server.
2. Open Windows Powershell **As Administrator**.

3. Run the following commands, where each command is a single contiguous line:

```
ren <drive>\JDE\PP\jde920\.vm_unconfigured <drive>\JDE\PP\jde920\vm_unconfigured

<drive>\JDE\jdk1.8_64\jre\bin\keytool -delete -alias cert -keystore "<drive>\JDE\jdk1.8_64\jre\lib
\security\cacerts" -storepass *****
```

> **Note:** The storepass value ***** in above command is the WebLogic Admin password.

```
$env:OPENSSL_CONF += "C:\JDE\bin\openssl.cnf"

<drive>\JDE\PP\CertGen\ConfigureCertKey_CC.ps1

<drive>\JDE\PP\CertGen\ConfigureCertKey_SMC.ps1 <WebLogic Admin Password>

ren <drive>\JDE\PP\jde920\vm_unconfigured <drive>\JDE\PP\jde920\.vm_unconfigured
```

Importing the Self-Signed Certificate into Target Machines on Microsoft Windows

The self-signed certificate file that you generated in the previous steps must be imported into these target machines that were delivered by JD Edwards EnterpriseOne One-Click Provisioning:

- All Servers
- HTML Server
- Enterprise Server

All Servers

Use this procedure to import the self-signed certificate into all Microsoft Windows servers.

1. Log into the each Microsoft Windows server.
2. Get the `cert.pem` file that is located on the Provisioning Server. You generated this key using the previous procedure in this section entitled: "Generating Self-Signed Certificates on Microsoft Windows".

   On the Provisioning Server, this file is located at this location:

   ```
   <drive>\JDE\PP\E1CloudConsole\keys
   ```
3. On each machine, use this command is determine if a certificate with an alias of **smcert** is already imported to `<drive>\JDE\jde_home\SCFHA\jdk`, where this command is a single contiguous line:

   ```
   <drive>\JDE\jde_home\SCFHA\jdk\jre\bin\keytool -list -v -alias smcert -keystore

   "<drive>\JDE\jde_home\SCFHA\jdk\jre\lib\security\cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   <drive>\JDE\jde_home\SCFHA\jdk\jre\bin\keytool -delete -alias smcert -keystore "<drive>\JDE\jde_home
   \SCFHA\jdk\jre\lib\security\cacerts" -storepass *****
   ```

   > **Note:** The storepass value ***** in above commands is the default password for Java trustStore.

4. On each machine, use these commands to import the certificate that you generated in the previous procedure of this section entitled: "Generate Self-Signed Certificates" to `<drive>\JDE\jde_home\SCFHA\jdk`, where each command is a single contiguous line:

```
<drive>\JDE\jde_home\SCFHA\jdk\jre\bin\keytool -import -file cert.pem -alias smcert -keystore "

<drive>\JDE\jde_home\SCFHA\jdk\jre\lib\security\cacerts" -storepass *****

<drive>\JDE\jde_home\SCFHA\jdk\jre\bin\keytool -list -v -alias smcert -keystore

"<drive>\JDE\jde_home\SCFHA\jdk\jre\lib\security\cacerts" -storepass *****
```

> **Note:** The storepass value ***** in above commands is the default password for Java trustStore.

HTML Server

Use this procedure to import the self-signed certificate into the HTML Server.

1. Log into the HTML server.
2. Get the `cert.pem` file that is located on the Provisioning Server. You generated this key using the previous procedure in this section entitled: "Generating Self-Signed Certificates on Microsoft Windows".

   On the Provisioning Server, this file is located at this location:

   ```
   <drive>\JDE\PP\E1CloudConsole\keys
   ```
3. Use this command is determine if a certificate with an alias of **smcert** is already imported to `C:\Program Files\Java\jdk1.8.0_201`, where this command is a single contiguous line:

   ```
   <JDK_path>\jre\bin\keytool -list -v -alias smcert -keystore "

   <JDK_path>\jre\lib\security\cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   <JDK_path>\jre\bin\keytool -delete -alias smcert -keystore "

   <JDK_path>\jre\lib\security\cacerts" -storepass *****
   ```

   > **Note:** The storepass value ***** in above commands is the WebLogic Admin password.
4. Use these commands to import the certificate that you generated in the previous procedure of this document entitled: "Generate Self-Signed Certificates" to C:\Program Files\Java\jdk1.8.0_201 , where each command is a single contiguous line:

   ```
   <JDK_path>\jre\bin\keytool -import -file cert.pem -alias smcert -keystore "<JDK_path>\jre\lib\security
   \cacerts" -storepass *****

   <JDK_path>\jre\bin\keytool -list -v -alias smcert -keystore "<JDK_path>\jre\lib\security\cacerts" -
   storepass *****
   ```

   > **Note:** The storepass value ***** in above commands is the WebLogic Admin password.

Enterprise Server

Use this procedure to import the self-signed certificate into the HTML Server.

1. Log into the Enterprise Server.
2. Get the `cert.pem` file that is located on the Provisioning Server. You generated this key using the previous procedure in this section entitled: "Generating Self-Signed Certificates on Microsoft Windows".

   On the Provisioning Server, this file is located at this location:

   ```
   <drive>\JDE\PP\E1CloudConsole\keys
   ```
3. Use this command is determine if a certificate with an alias of smcert is already imported to <drive>\JDE \jdk8_32, where this command is a single contiguous line:

   ```
   <drive>\JDE\jdk8_32\jre\bin\keytool -list -v -alias smcert -keystore "

   <drive>\JDE\jdk8_32\jre\lib\security\cacerts" -storepass *****
   ```

   If the results of the above command indicate that an alias of **smcert** certificate is already imported, use this command to remove it, where this command is a single contiguous line:

   ```
   <drive>\JDE\jdk8_32\jre\bin\keytool -delete -alias smcert -keystore

   "<drive>\JDE\jdk8_32\jre\lib\security\cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the Site Key password.
4. Use these commands to import the certificate that you generated in the previous procedure of this document entitled: "Generate Self-Signed Certificates" to <drive>\JDE\jdk8_32, where each command is a single contiguous line:

   ```
   <drive>\JDE\jdk8_32\jre\bin\keytool -import -file cert.pem -alias smcert -keystore"

   <drive>\JDE\jdk8_32\jre\lib\security\cacerts" -storepass *****

   <drive>\JDE\jdk8_32\jre\bin\keytool -list -v -alias smcert -keystore "

   <drive>\JDE\jdk8_32\jre\lib\security\cacerts" -storepass *****
   ```

   Note: The storepass value ***** in above commands is the Site Key password.

18 Performing Basic Administration

Starting and Stopping Services

This section shows you how to start or stop services.

In general you should use Server Manager to start or stop services, or to check the status of services, for all JD Edwards EnterpriseOne servers.

Prerequisite

A deployment of JD Edwards EnterpriseOne.

Starting and Stopping Services

You can use these commands to manually start, stop, or check the status of the **jde-sm** service for the Server Manager Console Server:

```
sudo su - root

service jde-sm start

service jde-sm stop

service jde-sm status
```

Configuring the Public IP Address as an Internal Hostname on Client Machine

This section shows you how to configure the public IP address as an internal hostname on client machine.

Without this step, the JD Edwards EnterpriseOne application websites can only be accessed using the Public IP Address in a browser.

Prerequisite

A deployment of JD Edwards EnterpriseOne.

Configuring the Public IP Address

As a convenience, in addition to access using a Public IP address, you can also access JD Edwards EnterpriseOne application websites using the hostname. The steps given below assume that you have not configured a public web entry point for the JD Edwards EnterpriseOne application tier. In that case you need to manually add the public IP address and internal DNS hostname of the JD Edwards EnterpriseOne application tier in the hosts file of the client machine.

Alternatively, these steps can be replaced by contacting your site's network administrator and having the address and hostname setup in DNS.

Access from a UNIX Host

Modify the `/etc/hosts` file as shown below. For example:

- public IP: **123.456.789.999**
- public hostname: **demo.company.com**
- Using the above example values, the `/etc/hosts` file must contain this line:

  123.456.789.999 demo demo.company.com

Access from a Microsoft Windows System

Modify the `C:\Windows\System32\drivers\etc\hosts` file as shown below. For example:

- public IP: **123.456.789.999**

19 Considering Optional Administrative Tasks

Understanding JD Edwards EnterpriseOne Security

This section provides an overview of the JD Edwards EnterpriseOne security.

A minimal JD Edwards EnterpriseOne security definition has been shipped with your Database Server.

Prerequisite

A deployment of JD Edwards EnterpriseOne.

JD Edwards EnterpriseOne Security

Follow the instructions in the *JD Edwards EnterpriseOne Applications Release 9.2 Installation Guide for Oracle on UNIX* (in the chapter entitled: Performing Post Installation Tasks, in the section entitled: Working With Signon Security and Table Creation Security to change the passwords within EnterpriseOne for JDE and PS920 so they match any changes you make to the passwords for the Oracle Database users.

Additionally, for table creation security you should use the Datasource Master application using the Database Privilege row exit. For further details refer to the *JD Edwards EnterpriseOne Tools Security Administration Guide Release 9.2.*

JD Edwards One-Click Provisioning provides a preconfigured environment with sample data, user-defined content, roles, and security permissions. The JDE user id you used to sign on is associated with the SYSADMIN role, which has a very broad set of permissions. As such, the JDE user will have access to a large number of applications, EnterpriseOne pages, and other content. UDO View Security has been enabled for 9.2.1.0. It may be necessary to adjust security accordingly. To modify or set up the permissions for the JDE user or other users that you add to this environment, refer to *Provisioning User and Role Profiles* in the JD Edwards EnterpriseOne Tools Security Administration Guide Release 9.2.

Configuring AIS Server Depending on How You Manage Users

This section shows you how to configure the AIS Server depending on how you manage users.

REST services on the JD Edwards EnterpriseOne AIS Server can use HTTP Basic Authentication for access. Support for HTTP Basic Authentication is enabled out of the box and is required to run the EnterpriseOne Orchestrator Client, create custom Java calls from orchestrations, and use the AIS client Java API (versions 1.2.1.x and higher).

Prerequisite

A deployment of JD Edwards EnterpriseOne.

Configuring AIS Server

When the AIS Server is deployed on Oracle WebLogic Server, the Oracle WebLogic Server may require the following additional configuration depending on how you manage users:

- If you are maintaining a user registry in Oracle WebLogic Server that matches the user registry in EnterpriseOne, with identical sets of user names and passwords in each system, you do NOT need to modify your configuration.
- If you are NOT maintaining identical sets of users in Oracle WebLogic Server and EnterpriseOne, then you need to perform the following steps to modify your Oracle WebLogic Server configuration. This ensures that Oracle WebLogic Server will not intercept HTTP Basic Authentication credentials passed from the REST service.

1. In the WebLogic Server domain for your AIS Server, in the Config directory, find the `config.xml` file.
2. Add this configuration as the last line within the `<security-configuration>` element, just before the `</security-configuration>` tag:

   ```
   <enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-credentials>
   ```
3. Restart the AIS Server for the changes to take effect.

The following is an example of this configuration in the `<security-configuration>` element:

```
<node-manager-password-encrypted>{AES}tzAokzTHACTNNmkuutLPQEpP8bfk7Ble24vmoycooic=</node-manager-password-encrypted>

<enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-credentials>

</security-configuration>

<server>
```

Editing the Java.security File (Optional)

This section provides a description of an optional procedure that describes how to edit the **Java.security** file in a JDK to restore certain settings in server instances that were deployed by One-Click Provisioning.

Prerequisite

A deployment of JD Edwards EnterpriseOne.

Editing the Java.security File

The deployment of a JDK to servers using JD Edwards EnterpriseOne One-Click Provisioning removes the string **anon, NULL** from the property **jdk.tls.disabledAlgorithms** in this file:

```
\jre\lib\security\Java.security
```

Optionally, after the One-Click deployment is complete, you can add the string back in the JDK for each deployed server. For example, you can append the string **anon, NULL** in the Java.security file in the **jdk.tls.disabledAlgorithms** section as shown in the **bolded** string below:

```
jdk.tls.disabledAlgorithms=SSLv3, RC4, DES, MD5withRSA, DH keySize < 1024, \ EC keySize < 224, 3DES-EDE-CBC, anon, NULL
```

For additional details, refer to *Java.security File* in the Server Manager Guide.

Using the Linux Logical Volume Manager (LVM) to Manage Free Disk Space

This section shows you how to use the Linux Logical Volume Manager (LVM) to Manage Free Disk Space.

The following procedures demonstrate how to extend the /u01 disk file system. An LVM command summary is also provided.

Prerequisite

- An understanding of Linux system administration.
- Appropriate sudo permissions at the Linux file system level.
- Appropriate access and authority to the Oracle Cloud Infrastructure environment.

Using the Linux Logical Volume Manager (LVM) to Manage Free Disk Space

The following procedures demonstrate how to extend the /u01 disk file system by:

- Recreating the partition table to include the free space
- Resizing the physical volume
- Extending the physical volume to the existing logical volume

Recreating the Partition Table

To make the free disk usable, you must recreate the partition table. This is done using the fdisk command as illustrated in the following procedure.

To recreate the a partition table:

1. Log into the instance as the opc user and switch to the root user using this command:

   ```
   sudo -i
   ```

2. Review available disk space in /u01 prior to adding the new disk space by running this command:

   df -h /u01

   Example:

   ```
   df -h /u01
   Filesystem Size Used Avail Use% Mounted on /dev/mapper/vg_jde00-lv_jde00
   6.9G 3.0G 3.7G 45% /u01
   ```

3. Run these commands to show unallocated space:

   parted /dev/xvdb

   print free

quit

Example:

parted /dev/xvdb

GNU Parted 2.1

Using /dev/xvdb

Welcome to GNU Parted! Type 'help' to view a list of commands. (

parted) print free

Number Start End Size Type File system Flags

32.3kB 1049kB 1016kB Free Space

1 1049kB 538MB 537MB primary ext3 boot

2 538MB 20.7GB 20.2GB primary lvm

3 20.7GB 28.2GB 7509MB primary lvm

28.2GB 39.7GB 11.5GB Free Space

(parted) quit

4. To modify the partition table, run fdisk on the disk device. For example, in the example used in this guide is `xvdb. fdisk /dev/xvdb`
5. Next, at the prompts from the fdisk command, enter the following command letters to delete the current partition number 3, and recreate primary partition #3 with all the remaining free sectors:

```
d
3
n
p
3
<enter>
<enter>
```

Example:

fdisk /dev/xvdb

> **Note:** DOS-compatible mode is deprecated. It's strongly recommended to switch off the mode (command 'c') and change display units to sectors (command 'u').

Command (m for help): d

Partition number (1-4): 3

Command (m for help): n

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): 3

First cylinder (2520-4830, default 2520):

Using default value 2520

Last cylinder, +cylinders or +size{K,M,G} (2520-4830, default 4830):

Using default value 4830

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

> **Note:** Re-reading the partition table failed with error 16: Device or resource busy. The kernel still uses the old table. The new table will be used at the next reboot or after you run partprobe(8) or kpartx(8) Syncing disks.

6. This process extended the /dev/xvdb partition to include all of the free space. This new disk will be added to the /u01 logical volume.The system now needs to be rebooted before the partition table change will be recognized by the system. Use the reboot command to reboot now. Continue with the next section after rebooting.

Resizing the Physical Volume with the Partition Table Changes

After you resize the partition table as described in the previous procedure, the disk image is now ready to become a physical volume that can be added to a logical volume.

To resize the physical volume with the new partition space:

1. After the system finishes rebooting, log back into the instance as the opc user and switch to the roof user using this command: `sudo -i`

2. Execute the pvresizecommand as shown below:

```
pvresize /dev/xvdb3
```

Example:

pvresize /dev/xvdb3

Physical volume "/dev/xvdb3" changed<

1 physical volume(s) resized / 0 physical volume(s) not resized

See Also

For more information refer to the Linux man pages for the pvresize command.

Extending the Volume Group with the New Physical Disk

To extend the volume group with the new physical disk:

1. You must determine the name of the volume group to which the disk will be added by using the vgs and lvscan commands. In the below example, the name of the volume group with free space is vg_jde00, and it is located in this directory: /dev/mapper/vg_jde00-lv_jde00

   Example:

   vgs
   VG #PV #LV #SN Attr VSize VFree
   vg_apol 1 7 0 wz--n- 18.78g 32.00m
   vg_jde00 1 1 0 wz--n- 17.70g 10.71g
   The vgs command shows the 10.71g of free space is in volume group vg_jde00

   Example:

   lvscan
   ACTIVE '/dev/vg_jde00/lv_jde00' [6.99 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_root' [4.00 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_home' [1.00 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_tmp' [6.00 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_var' [1.50 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_var.log' [2.00 GiB] inherit
   ACTIVE '/dev/vg_apol/lv_var.log.audit' [256.00 MiB] inherit
   ACTIVE '/dev/vg_apol/lv_swap' [4.00 GiB] inherit
   The lvscan shows that volume group vg_jde00 has one logical volume named lv_jde00. This is where the free disk space will be put in this example. Take the volume group name (vg_jde00) and the logical volume name (lv_jde00) and confirm the location of the logical volume in the /dev/mapper directory.

   Example:

   ls /dev/mapper/vg_jde00-lv_jde00
   /dev/mapper/vg_jde00-lv_jde00
2. Using the volume group name found in the previous step, extend the volume group by using the lvextend command. The below example uses **/dev/mapper/vg_jde00-lv_jde00**.

   lvextend --resizefs -l +100%FREE /dev/mapper/vg_jde00-lv_jde00

   Example:

   lvextend --resizefs -l +100%FREE /dev/mapper/vg_jde00-lv_jde00
   Size of logical volume vg_jde00/lv_jde00 changed from 6.99 GiB (1790 extents) to 17.70 GiB (4531 extents).

Logical volume lv_jde00 successfully resized

resize2fs 1.43-WIP (20-Jun-2013)

Filesystem at /dev/mapper/vg_jde00-lv_jde00 is mounted on /u01; on-line resizing required

old_desc_blocks = 1, new_desc_blocks = 2

Performing an on-line resize of /dev/mapper/vg_jde00-lv_jde00 to 4639744 (4k) blocks.

The filesystem on /dev/mapper/vg_jde00-lv_jde00 is now 4639744 blocks long.

3. Confirm the space is now available for use by running this command:

```
df -h /u01
```

Example:

df -h /u01

Filesystem Size Used Avail Use% Mounted on

/dev/mapper/vg_jde00-lv_jde00

18G 3.0G 14G 18% /u01

See Also

For more information refer to the Linux man pages for the lvextend command.

Logical Volume Manager Command Summary

This section provides a summary of the Logical Volume Manager commands used in the previous steps of this tutorial.

[fdisk]

This is the disk utility command.

To resize the number 3 primary partition, use this syntax where subsequent values are entered at the command prompts:

```
fdisk /dev/xvdb
d
3
n
p
3
<enter>
<enter>
w
```

[pvresize]

Use this utility to resize a disk or partition for use by LVM. For example:

```
pvresize /dev/xvdb3
```

[vgscan]

Use this command to scan all disks for volume groups and rebuild caches (returns names).

[vgextend]

Use this utility to resize a disk or partition for use by LVM. For example:

```
vgextend VolGroup00 /dev/hdc1
```

[lvextend]

Use this command to extend the size of a logical volume. For example:

```
vlvextend --resizefs -l +100%FREE /dev/mapper/vg_jde00-lv_jde00
```

20 Upgrading your One-Click Provisioned Environment

Upgrade Learning Path

The following learning path will guide you through the upgrade process for your One-Click provisioned environment:

Upgrading JD Edwards EnterpriseOne on a One-Click Provisioned Target Environment on Linux