JD Edwards EnterpriseOne

Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with DB Systems using Reference Architecture Learning Path

1.0

JD Edwards EnterpriseOne
Deploying JD Edwards EnterpriseOne on Oracle Cloud Infrastructure on Linux with DB Systems using Reference Architecture Learning Path

1.0

Part Number: G36659-01

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 the automated processes for deploying JD Edwards EnterpriseOne Release 9.2 on Oracle Cloud Infrastructure on Linux with Database Cloud Service (DB Systems) using JD Edwards Infrastructure Provisioning and JD Edwards One-Click Provisioning. It is important to note that this document is primarily concerned with the basic requirements for an installation of JD Edwards EnterpriseOne into the Oracle Cloud Infrastructure.

Upon completion of this learning path, you will have a working deployment of a JD Edwards EnterpriseOne on Linux in Oracle Cloud Infrastructure with Database Cloud Service (DB Systems). This process includes 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 for functional changes to the JD Edwards EnterpriseOne implementation for the Oracle Cloud Infrastructure, at which time any differences in the user interface will be reconciled.

Before You Begin

Provides information and resource requirements critical to understand prior to using Infrastructure Provisioning on Oracle Cloud Infrastructure.

Key Concepts of Reference Architecture

This section describes these key concepts of reference architecture:

1. NAT Gateway
2. Bastion Hosts
3. Private Load Balancer (LBaaS)
4. Production Environments
5. Non-Production Environments
6. Disaster Recovery Environments
7. POD Architecture

NAT Gateway

Key characteristics of a NAT gateway include:

- The gateway gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections.
- The gateway is a networking technique commonly used to give an entire private network access to the internet without assigning each host a public IPv4 address. The hosts can initiate connections to the internet and receive responses, but not cannot receive inbound connections initiated from the internet.

When a host in the private network initiates an internet-bound connection, the NAT device's public IP address becomes the source IP address for the outbound traffic. The response traffic from the internet therefore uses that public IP address as the destination IP address. The NAT device then routes the response to the host in the private network that initiated the connection.

For additional information on NAT gateways, refer to Oracle Cloud Infrastructure Documentation:

https://docs.cloud.oracle.com/iaas/Content/Network/Tasks/NATgateway.htm

Bastion Hosts

Bastion hosts are an important part of the network security layer for both cloud and data center deployments. Combined with firewall policies, bastion hosts can protect your environment by blocking access to external management interfaces.

Because most of the infrastructure denies remote access, you need a method to log in to the servers located in the private subnets. You can establish a point-to-network VPN, but this method increases the complexity of and the management necessary for the setup. A secure and convenient method that you can use, is to connect to the bastion hosts by using the SSH protocol.

Private Load Balancer (LBaaS)

To isolate your load balancer (LBaaS) from the internet and simplify your security posture, you can create a private load balancer. The load balancing service assigns it to a private IP address that serves as the entry point for incoming traffic.

When you create a private load balancer, the service requires only one subnet to host both the primary and standby load balancers. The load balancer can be regional or AD-specific, depending on the scope of the host subnet. The load balancer is accessible only from within the VCN that contains the host subnet, or from any other network based on your security list rules.

The assigned floating private IP address is local to the host subnet. The primary and standby load balancers each require an additional private IP address from the host subnet.

If there is an availability domain outage, a private load balancer created in a regional subnet within a multi-AD region provides failover capability. A private load balancer created in an AD-specific subnet, or in a regional subnet within a single-AD region, has no failover capability in response to an availability domain outage.

Production Environment

Production environments include a set of resources that provides high availability and scalability in terms of resources and services. The production environment can be deployed in a single AD only.

Resources include:

- NAT Service

- Bastion Servers
- One-Click Provisioning Server
- RAC database servers (DBS only)
- Enterprise Servers
- Web Servers (HTML and AIS)
- LBaaS for Web Servers
- Microsoft Windows-based Deployment Server

Non-Production Environments

Non-production environments include a set of resources that provides JD Edwards EnterpriseOne servers that are not deployed in high availability abilities. The non-production environment can be deployed in a single AD only.

Resources include:

- One-Click Provisioning Server
- RAC Database Servers (DBS only)
- Enterprise Servers
- Web Servers (HTML and AIS)
- Microsoft Windows-based Deployment Server

Disaster Recovery Environments

Disaster Recovery (DR) environments include the same set of resources as Production environments with high availability and scalability in terms of resources and services. The Production environment can be deployed in a single AD only.

Resources in the DR environment include:

- NAT Service
- Bastion Servers
- RAC Database Servers (DBS created after enabling Oracle Data Guard for the PD environment)
- Enterprise Servers
- Web Servers (HTML and AIS)
- LBaaS for Web Servers (HTML and AIS)

POD Architecture

The JD Edwards EnterpriseOne One-Click architecture for HTML and AIS Servers introduces the concept of pods, which matches the deployment concept of JD Edwards EnterpriseOne Infrastructure Provisioning for Reference Architecture. You can use pods to scale JD Edwards EnterpriseOne, or to separate by pathcodes such as Development and Production. The POD architecture is recommended for high performance environments such as Production.

For purposes of this discussion regarding the definition of HTML and AIS Servers, a pod consists of a Logic Server, a Batch Server, a Standard HTML Server, a Dedicated HTML Server and an AIS Server. Pods are generally associated by a common pathcode. To scale by pathcode, multiple pods can be configured.

Fundamentals

You can use the Oracle Cloud Infrastructure (also called "IaaS" - Infrastructure as a Service) to deploy JD Edwards EnterpriseOne. This process features a web-based JD Edwards Infrastructure Provisioning Console interface that enables the provisioning of a fully functional suite of interconnected Linux-based servers and a Microsoft Windows machine based on JD Edwards Reference Architecture that is optimized for deployments in Oracle Cloud Infrastructure. Linux-based servers are required to deploy core JD Edwards EnterpriseOne servers are the Enterprise Server(s), the HTML Server(s), the Application Interface Service (AIS) Server, and the One-Click Provisioning Server that also includes the JD Edwards Server Manager Console (SMC). The database is provided by Oracle Cloud Infrastructure DB Systems (DBS) also called DBaaS (Database as a Service). A Microsoft Windows-based machine is required for Deployment server. On successful completion of Infrastructure provisioning, JD Edwards One-Click Provisioning will be used to deploy JD Edwards EnterpriseOne components.

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

Deploying JD Edwards EnterpriseOne Using Reference Architecture

JD Edwards provides you a reference architecture optimized for Oracle Cloud Infrastructure, configured with out-of-the-box security and high availability. Deploying JD Edwards EnterpriseOne using reference architecture is completely automated for Linux with DB systems.

Deploying JD Edwards EnterpriseOne Using Reference Architecture

Process Flow for One-Click Provisioning

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

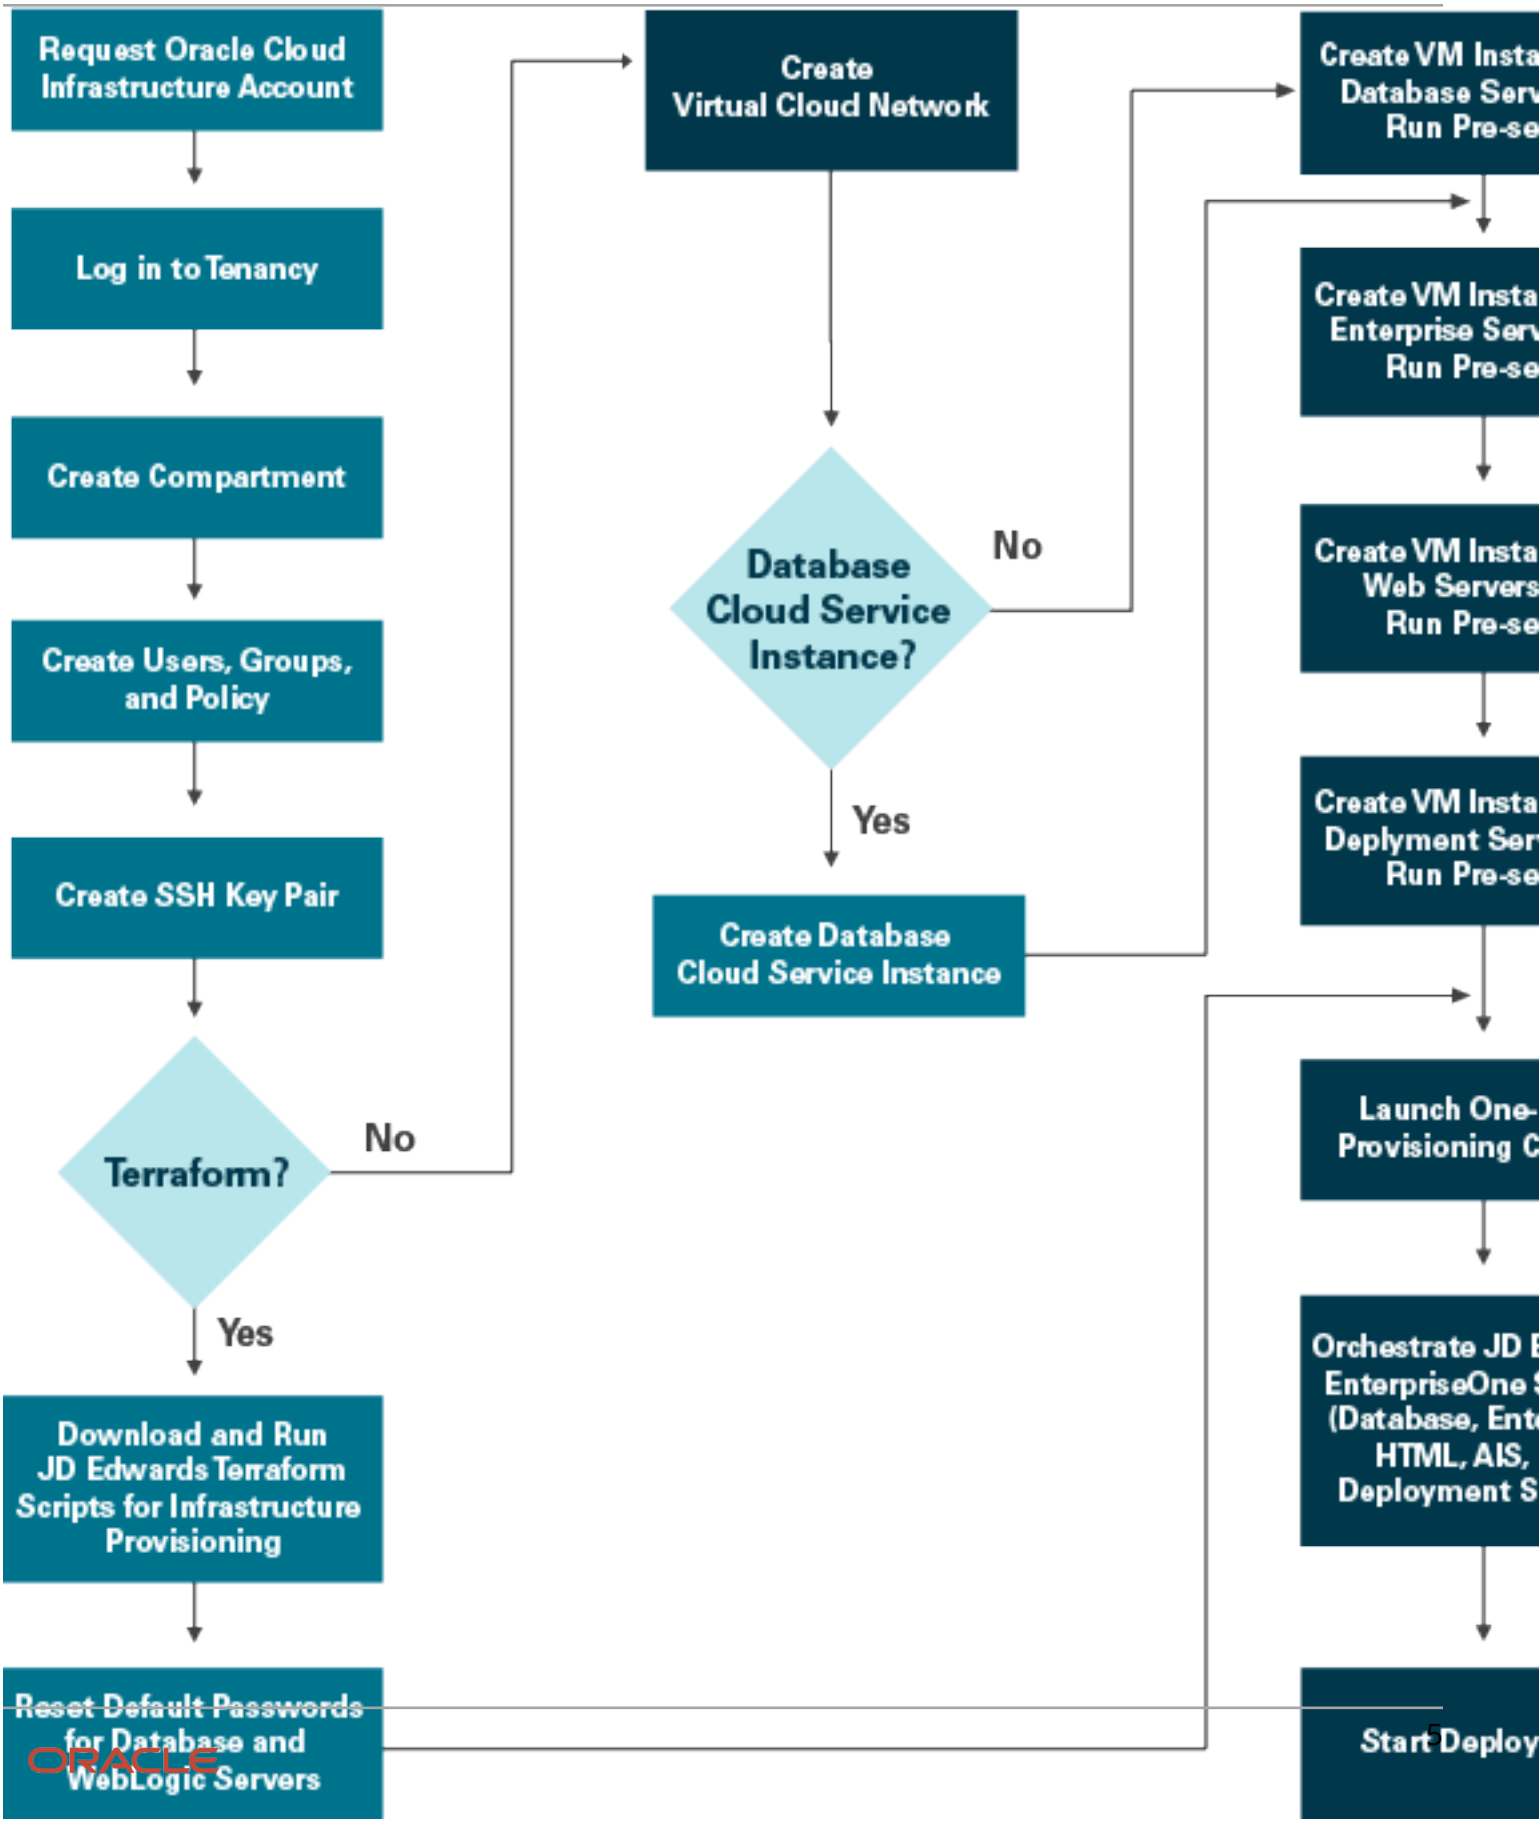
Request Oracle Cloud Infrastructure Account
Log in to Tenancy
Create Compartment
Create Users, Groups, and Policy
Create SSH Key Pair
Terraform?
No
Yes
Download and Run JD Edwards Terraform Scripts for Infrastructure Provisioning
Reset Default Passwords for Database and WebLogic Servers
Create Virtual 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

Minimum Resource Requirements

The table below specifies the minimum resource requirements to install and run JD Edwards On-Premises using an Oracle database (Linux and Microsoft Windows) or a Microsoft SQL Server database (Microsoft Windows). 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

Note: Within the **Specialty and previous generation** shape series, only these shapes are supported:

- VM.Standard2.2
- VM.Standard2.4
- VM.Standard2.8
- VM.Standard2.16
- VM.Standard2.24

JD Edwards EnterpriseOne Server Type	Minimum Recommended			Notes
	CPU	Memory (GB)	Storage Volume (GB)	
Required JD Edwards Components				
One-Click Provisioning Server	2	30	100	Includes Server Manager
Deployment Server	2	30	210 GB* for all four (4) pathcodes	Using One-Click, customers must install all four (4) path codes. There is no automated way to add additional path codes post deployment. * Storage volume space is in addition to that required by the Windows OS itself, which can be up to 45 GB.

				Therefore, the minimum recommended storage volume size is 256 GB.
Database Server	2	30	50 GB for a single pathcode and shared data	20 GB is required for each additional pathcode. These numbers are for demo data only and should be adjusted for expected required business data space.
Enterprise Server	2	30	75	
WebLogic Server • HTML Server for Web Client • HTML Server for AIS Server • AIS Server	-	-	30	For the OS and WebLogic Server
	2	30	20	Per Web Instance
	2	30	20	Per Web Instance
	2	30	20	Per Web Instance
Optional JD Edwards Components				
Business Services Server (BSSV)	2	30	50	Per Web Instance
Transaction Server for Real Time Events (RTE)	2	30	50	Per Web Instance
Application Development Framework Server (ADF)	2	30	50	Per Web Instance
One View Reporting (OVR) Server / BI Publisher Server (BIP)	2	30	50 GB is required for a single pathcode	10 GB is required for each additional pathcode
Development Client	2	30	100	Per 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-premise methodology.

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.

You must also obtain a VNC viewer such as TigerVNC (*https://tigervnc.org*). This product, or functional equivalent, is required in order to access the Infrastructure Provisioning Server that is part of the automated deployment of One-Click using reference architecture (OpenTofu).

2 Planning Your Deployment

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.

Generating Instance Key Pairs in openssh Format

This section shows you how to generate instance SSH key pairs in openssh format on your local system, which can be UNIX or Windows. Key pairs in this format are required if you need to connect directly to any Linux instance.

Multiple methods can be used to securely connect to Oracle Cloud Infrastructure for JD Edwards EnterpriseOne infrastructure provisioning. This connection is made using a key pair that consists of a public key and a private key. The following are the key pairs that are used for connecting to Oracle Cloud Infrastructure:

- **Instance Key Pair.** Secure Shell (SSH), which provides an encrypted login method, is a more secure replacement for Telnet for logging on to Oracle Cloud Infrastructure. Before you can use Infrastructure Provisioning <infrastructure provisioning> to create instances, you must generate instance key pairs and upload the SSH public key to Oracle Cloud Infrastructure. This SSH public key is used for authentication for any Oracle Cloud Infrastructure instance except the Bastion Server.
- **Bastion Host Key Pair**. This key pair is similar to an instance key pair and is strongly recommended to be used as a best practice for ensuring the highest level of security. That is, you should the same procedure to create a different set of **Bastion Host Key Pairs** for public production access to the Bastion host, while keeping access to other host instances securely separated.
- **Infrastructure Provisioning User Key Pair**. Instead of SSH key format, this key pair must be generated in Privacy Enhanced Mail (PEM) container format. The public key of this key pair must be added to the account for the infrastructure provisioning user. When this public key is uploaded, the system automatically creates a fingerprint that is displayed in the console of Oracle Cloud Infrastructure and is required for input into the Infrastructure Provisioning Console. To enable provisioning in Oracle Cloud Infrastructure, the private key of this key pair is required as an input in the Infrastructure Provisioning Console.
- **CA Certificates.** You must generate CA certificates to support Load Balancing as a Service (LBaaS), which is a core functionality that is deployed and configured by JD Edwards EnterpriseOne Infrastructure Provisioning. These certificates are used to configure LBaaS with SSL. The procedure is described in the section "Generating CA Certificates for Load Balancing as a Service (LBaaS)" of this Learning Path.

> **Tip:** The best practice is to create at least two SSH key pairs for each purpose, because if for any reason a single SSH key is no longer valid, access to the server would be lost permanently with no means to recover. You cannot access the server without using an SSH Key. Additional keys can be added manually after the instance is started.

> **Note:** Use caution if prompted to overwrite a previously generated SSH key. If you overwrite a key previously used to connect to a prior Oracle Cloud Infrastructure instance, you may permanently lose access to (that is, the ability to log in to) any prior Oracle Cloud Infrastructure instance that used that key.

The following diagram illustrates the security key architecture used by JD Edwards Infrastructure Provisioning.

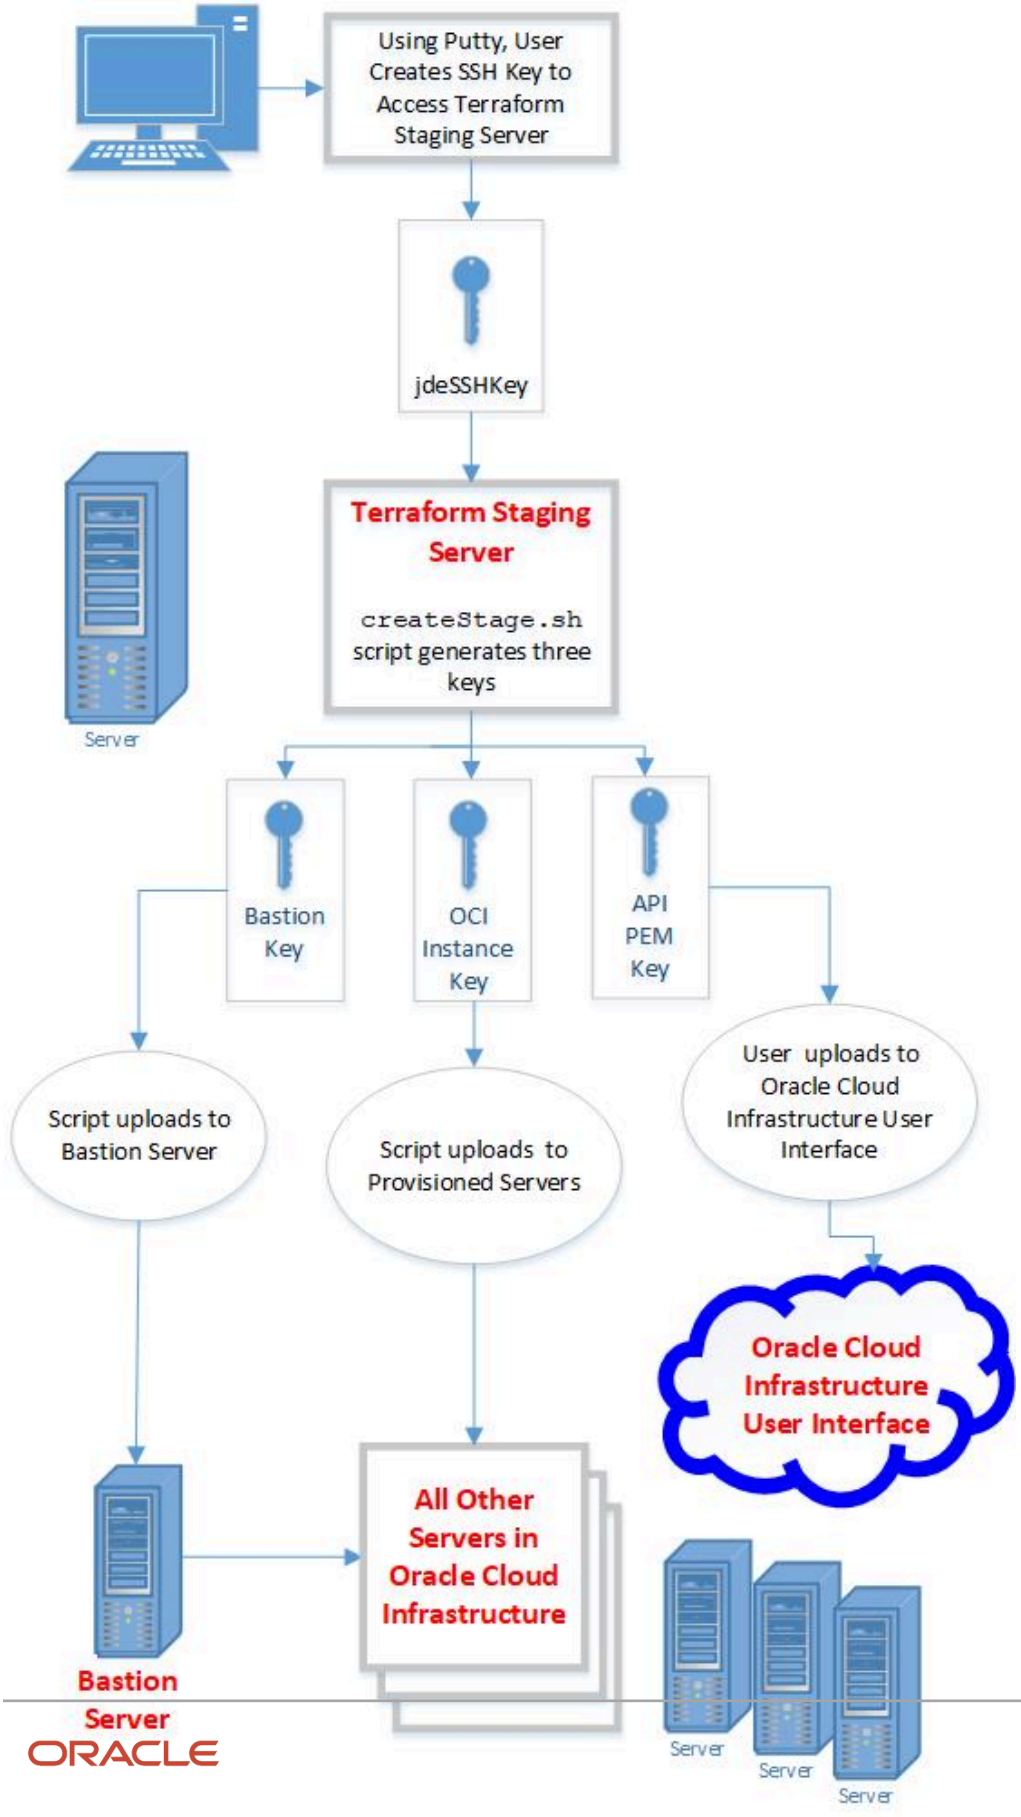
Using Putty, User Creates SSH Key to Access Terraform Staging Server
jdeSSHKey
Terraform Staging Server
createStage.sh script generates three keys
Server
Bastion Key
OCI Instance Key
API PEM Key
Script uploads to Bastion Server
Script uploads to Provisioned Servers
User uploads to Oracle Cloud Infrastructure User Interface
Oracle Cloud Infrastructure User Interface
All Other Servers in Oracle Cloud Infrastructure
Bastion Server
Server
Server
Server

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

Secure Shell (SSH), which provides an encrypted login method, is a more secure replacement for Telnet for logging on to Oracle Cloud Infrastructure. Before you create your Oracle Cloud Infrastructure instance, you must generate SSH key pairs and upload the SSH public keys to Oracle Cloud Infrastructure. These SSH public keys will be used for authentication when you log in to the instance. You must also create pairs of private keys, one pair for use by the One-Click Provisioning Server to create instances for JD Edwards EnterpriseOne servers and another pair to enable access to the instances. Below is a summary of the required SSH keys and their formats:

- **Private Key in .ppk Microsoft Windows Format**

  Required to connect from a Microsoft Windows machine to an Oracle Cloud Service instance including the Provisioning Server itself and also to connect to any provisioned server such as DBCS, Enterprise Server, and JCS servers (such as HTML and AIS servers).

  See Step 3 in the following procedure.
- **Public Key in .pub Format**See Step 4 in the following procedure.

Note: Do *not* set a passphrase for any SSH key.

Use this procedure to generate a Secure SHell (SSH) key on your local system and save the key on a file for uploading to Oracle Cloud Infrastructure.

1. Locate and run **puttygen.exe** in the PuTTY folder of your local Microsoft Windows computer.

2. Generate the key using the following steps:
   a. On the PuTTY Key Generator window, accept the default key type, SSH-2 RSA, and in the Number of bits in a generated key: field, ensure that the value is set to 2048.
   b. Click the Generate button.

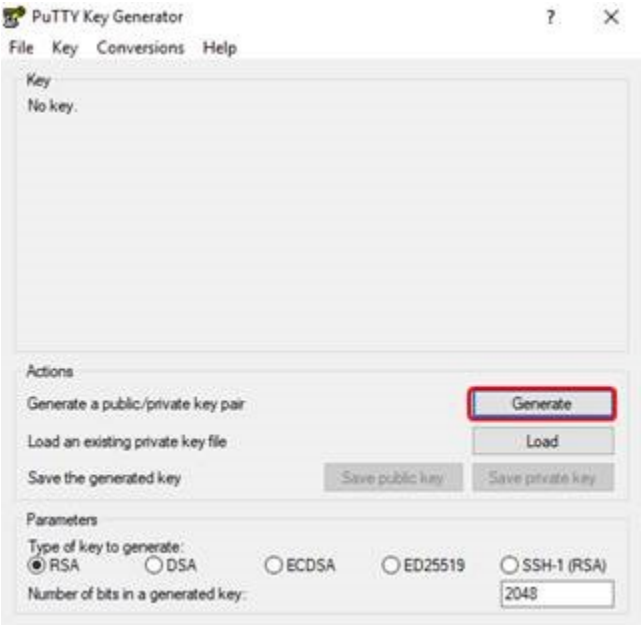

   c. After you click the Generate button, move your mouse around the blank area to generate randomness for the SSH key pair you generate.

3. Use this step to create a private key in .ppk Microsoft Windows format.

   a. In the PuTTY Key Generator dialog box, click the Save private key button to save your private key to the system.

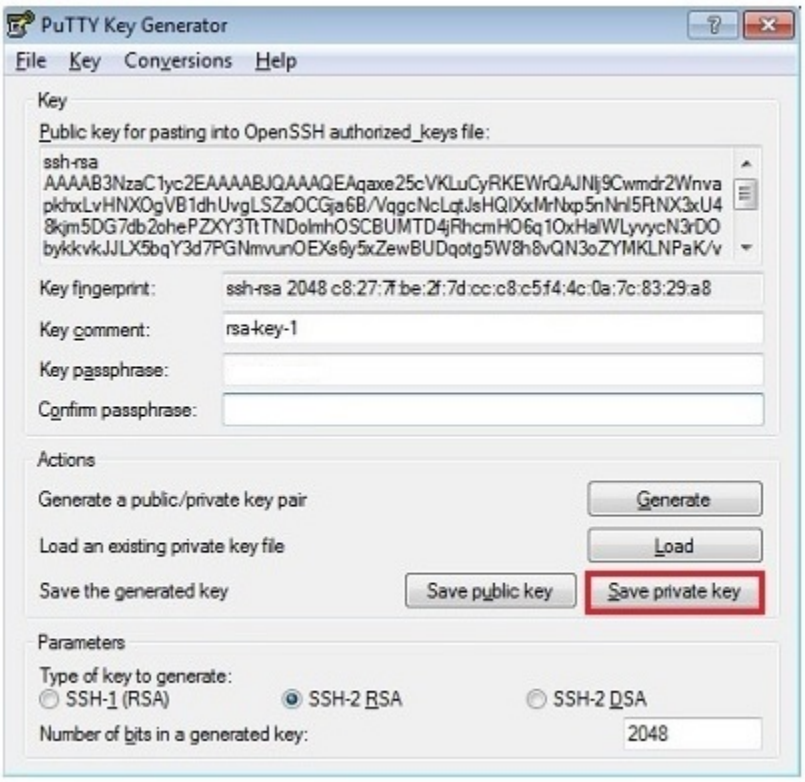

   b. On the PuTTYgen Warning dialog box, click the Yes button to confirm that you want to create the private key without a passphrase.

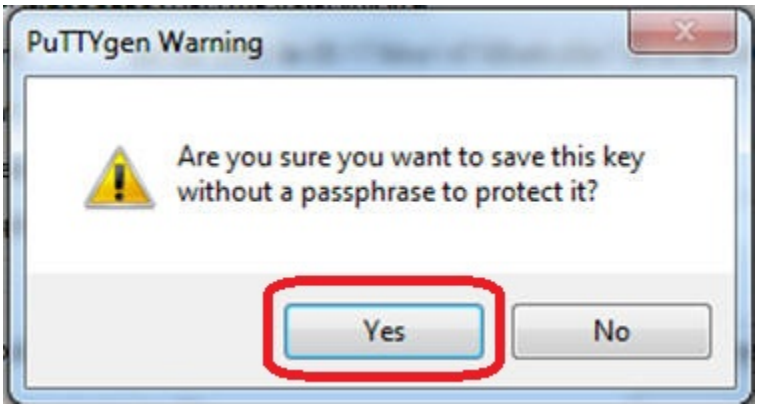

> **Note:** When you save this key to the local file system, you should give it a significant name such as like `jdeSSHKey.ppk`. It is also important that you ensure that this key has a `.ppk` extension.

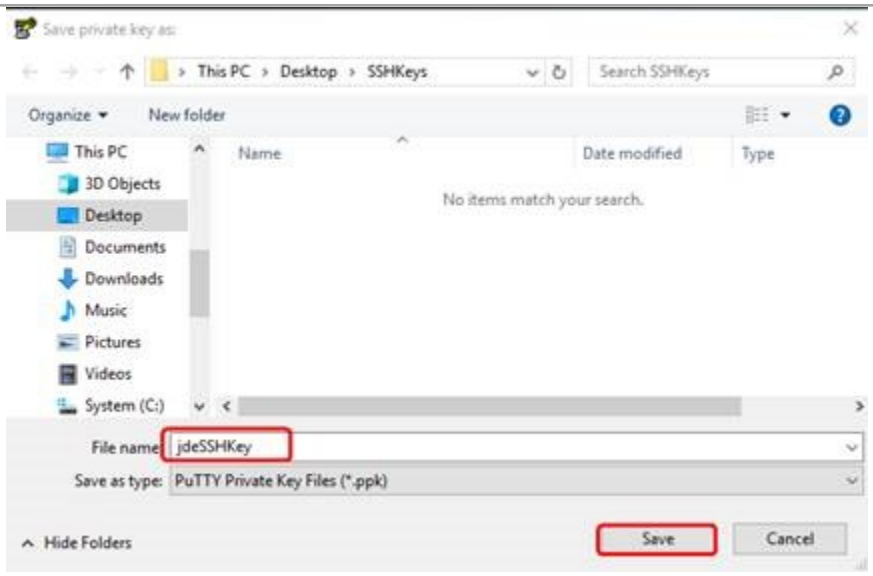

Note: Keep a record of the file name and location which you will need when you upload this key from a Microsoft Windows machine that is accessing the JD Edwards Reference Architecture Infrastructure Provisioning Server.

4. Use this step to create a public key in .pub format.

   a. On the PuTTY Key Generator dialog box, select all the characters in the Public key for pasting into OpenSSH authorized_keys file field.

      Note: Be sure you select all the characters, not just the ones you can see in the narrow window. If a scroll bar appears next to the characters, scroll through the entire window to select all the characters.

   b. On the selected text, right-click to see the context menu and select Copy.

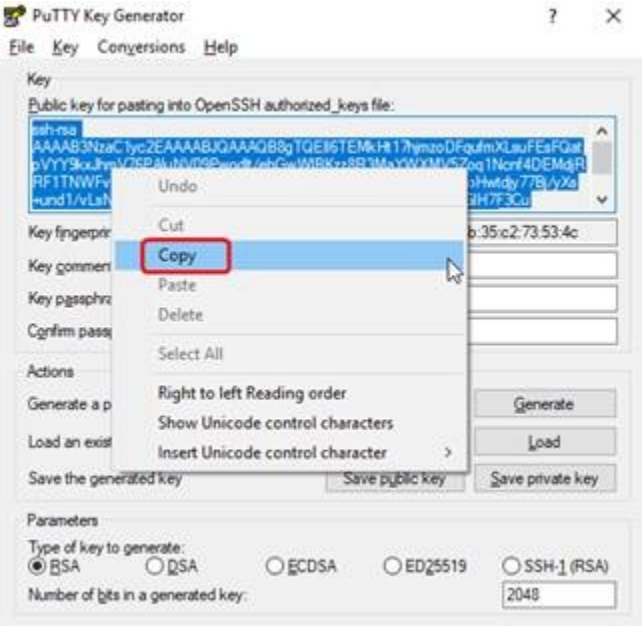

   c. Open a plain text editor (such as vi on UNIX or Notepad on Microsoft Windows) and paste the characters you just copied. Ensure that you paste the text at the first character in the text editor, and do not insert any line breaks.

   d. Save the plain text file as a file name with a .pub extension and keep a record of the file name.

      Tip: You should give this key a significant name such as jdeSSHKey.pub.

Note: You should not use the Putty Key Generator function to save the public key as a file, which is invoked by the Save Public Key button. Using this function introduces extraneous characters and strings to the key that are not compatible with Oracle Cloud Infrastructure.

Note: As a backup, in case the primary public key is lost or damaged, you should also generate a secondary key using this same procedure. Again you may use any file extension, but .pubbak is a useful convention to indicate that this file is a backup of the primary public key.Keep a record of the primary and secondary public key file names and their location. You will need to upload these public keys when you create an instance.

Note: You cannot access an instance in Oracle Cloud Infrastructure without a valid public key which you have successfully uploaded to Oracle Cloud Infrastructure during instance creation. Therefore you must upload the primary public key whenever you are creating an instance. Optionally, you can upload the backup secondary public key in the same step.

Generating Self-Signed Certificates for Load Balancing as a Service (LBaaS)

This section shows you how to generate the public, private, and device self-signed certificates that are required as input into the JD Edwards EnterpriseOne Infrastructure Provisioning Console for Oracle Cloud Infrastructure in support of Load Balancing as a Service (LBaaS), which is a core functionality for infrastructure provisioning.

Prerequisite

- Access to a machine running Oracle Enterprise Linux. For example, this can be the Linux instance that you created for the OpenTofu staging server in Oracle Cloud Infrastructure.

Generating Self-Signed Certificates

On any machine running Oracle Enterprise Linux, use this procedure to generate self-signed root and device certificates. You will be prompted for these certificates by the Infrastructure Provisioning Console, which runs on the OpenTofu staging server. Therefore, the OpenTofu staging server is a logical choice of a server on which to perform this procedure.

1. Generate a **rootCA private key** using this command:

   ```
   openssl genrsa -out rootCA.key 2048 -days 365
   ```

2. Generate **rootCA public key** using this command:

   ```
   openssl req -x509 -new -nodes -key rootCA.key -sha256 -days 365 -out rootCA.pem -subj "/C=<country>/
   ST=<state>/L=<location>/O=<company>/OU=<organization unit>/CN=localhost"
   ```

   where **C=<country>** is the country where you will be submitting the CA Certificates Service Request (CSR), and

   where **ST=<state>** is the state within the country you specified in `c=`, and

   where **L=<location>** is the location or city within the state, and

   where **O=<organization>** is the organization for which the CSR will be issued, and

   where **OU = <organization unit>** is the originzational unit of the organization

   For example, if the CSR is being requested for Oracle Corporation, for the JD Edwards EnterpriseOne division, in Denver, Colorado, United States, the command would be:

   ```
   openssl req -x509 -new -nodes -key rootCA.key -sha256 -days 365 -out rootCA.pem -subj "/C=US/ST=CO/
   L=Denver/O=Oracle Corporation/OU=EnterpriseOne/CN=localhost"
   ```

3. Create a **configuration file named** `device-csr.conf` that includes these sections and settings:

```
[req]

distinguished_name = req_distinguished_name

req_extensions = v3_req

prompt = no

[req_distinguished_name]

C = <country>

ST = <state>

L = <location>

O = <organization>

OU = <organization unit>

CN = localhost

[v3_req]

keyUsage = keyEncipherment, dataEncipherment

extendedKeyUsage = serverAuth

subjectAltName = @alt_names

[alt_names]

DNS.1 = localhost

DNS.2 = <logic Enterprise Server>

DNS.3 = <batch Enterprise Server>

DNS.4 = <HTML web server>
```

where **C=<country>** is the country where you will be submitting the CA Certificates Service Request (CSR), and

where **ST=<state>** is the state within the country you specified in `c=`, and

where **L=<location>** is the location or city within the state, and

where **O=<organization>** is the organization for which the CSR will be issued, and

where **OU = <organization unit>** is the unit of the organization, and

where **<logic Enterprise Server>** is the name of your JD Edwards Enterprise Server running application logic, which must be limited to 15 alphanumeric characters, and

where **<batch Enterprise Server>** is the name of your JD Edwards Enterprise Server running UBE batch processes, which must be limited to 15 alphanumeric characters, and

where **<HMTL Web Server>** is the name of your JD Edwards Web Server running HTML application logic, which must be limited to 15 alphanumeric characters

For example, a properly configured `device-csr.conf` file might look like this:

```
[req]

distinguished_name = req_distinguished_name

req_extensions = v3_req

prompt = no

[req_distinguished_name]

C = US

ST = CO

L = Denver

O = Oracle Corporation

OU = EnterpriseOne

CN = localhost

[v3_req]

keyUsage = keyEncipherment, dataEncipherment

extendedKeyUsage = serverAuth

subjectAltName = @alt_names

[alt_names]

DNS.1 = localhost

DNS.2 = logiclb

DNS.3 = batchlb

DNS.4 = weblb
```

4. Generate a **device key** using this command:

```
openssl genrsa -out device.key 2048 -days 365
```

5. Generate a **device CSR** using this command:

```
openssl req -new -key device.key -out device.csr -subj "/C=<country>/ST=<state>/L=<location>/
O=<organization>/OU=<organization_unit>/CN=localhost"
```

For example:

```
openssl req -new -key device.key -out device.csr -subj "/C=US/ST=CO/L=Denver/O=Oracle Corporation/
OU=EnterpriseOne/CN=localhost"
```

6. Generate a **device certificate** using this command, where you have created and configured the input file `device-csr.conf` file as described in the previous step in this procedure:

```
openssl x509 -req -in device.csr -CA rootCA.pem -CAkey rootCA.key -CAcreateserial -out
```

3 Performing Setup Tasks in Oracle Cloud Infrastructure

Logging In To Oracle Cloud

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 for Reference Architecture

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

On the Oracle Cloud Infrastructure Welcome page, you can click the *About Compartments* link for additional information.

Prerequisite

- 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 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 and an administrator account in the platform. For more information, refer to this site:
  - *Getting Started with Oracle Cloud*
- To access the Oracle Cloud Infrastructure Console, you must use a *supported browser*.

Creating a Compartment

To create a compartment for JD Edwards EnterpriseOne on Oracle Cloud Infrastructure:

1. As the user for which you will create a public key to access the OpenTofu staging server, log in to Oracle Cloud Infrastructure using the following URL, which leads to the tenancy for which you want to provision infrastructure using the JD Edwards EnterpriseOne Infrastructure Provisioning Console:

   ```
   https://console.<your_region-ad>.oraclecloud.com/#/a/
   ```

   For this tutorial, we will use this URL:

   ```
   https://console.us-ashburn-1.oraclecloud.com/#/a/
   ```

2. On the Oracle Cloud Infrastructure Console Home page, click the navigation menu in the upper-left corner.

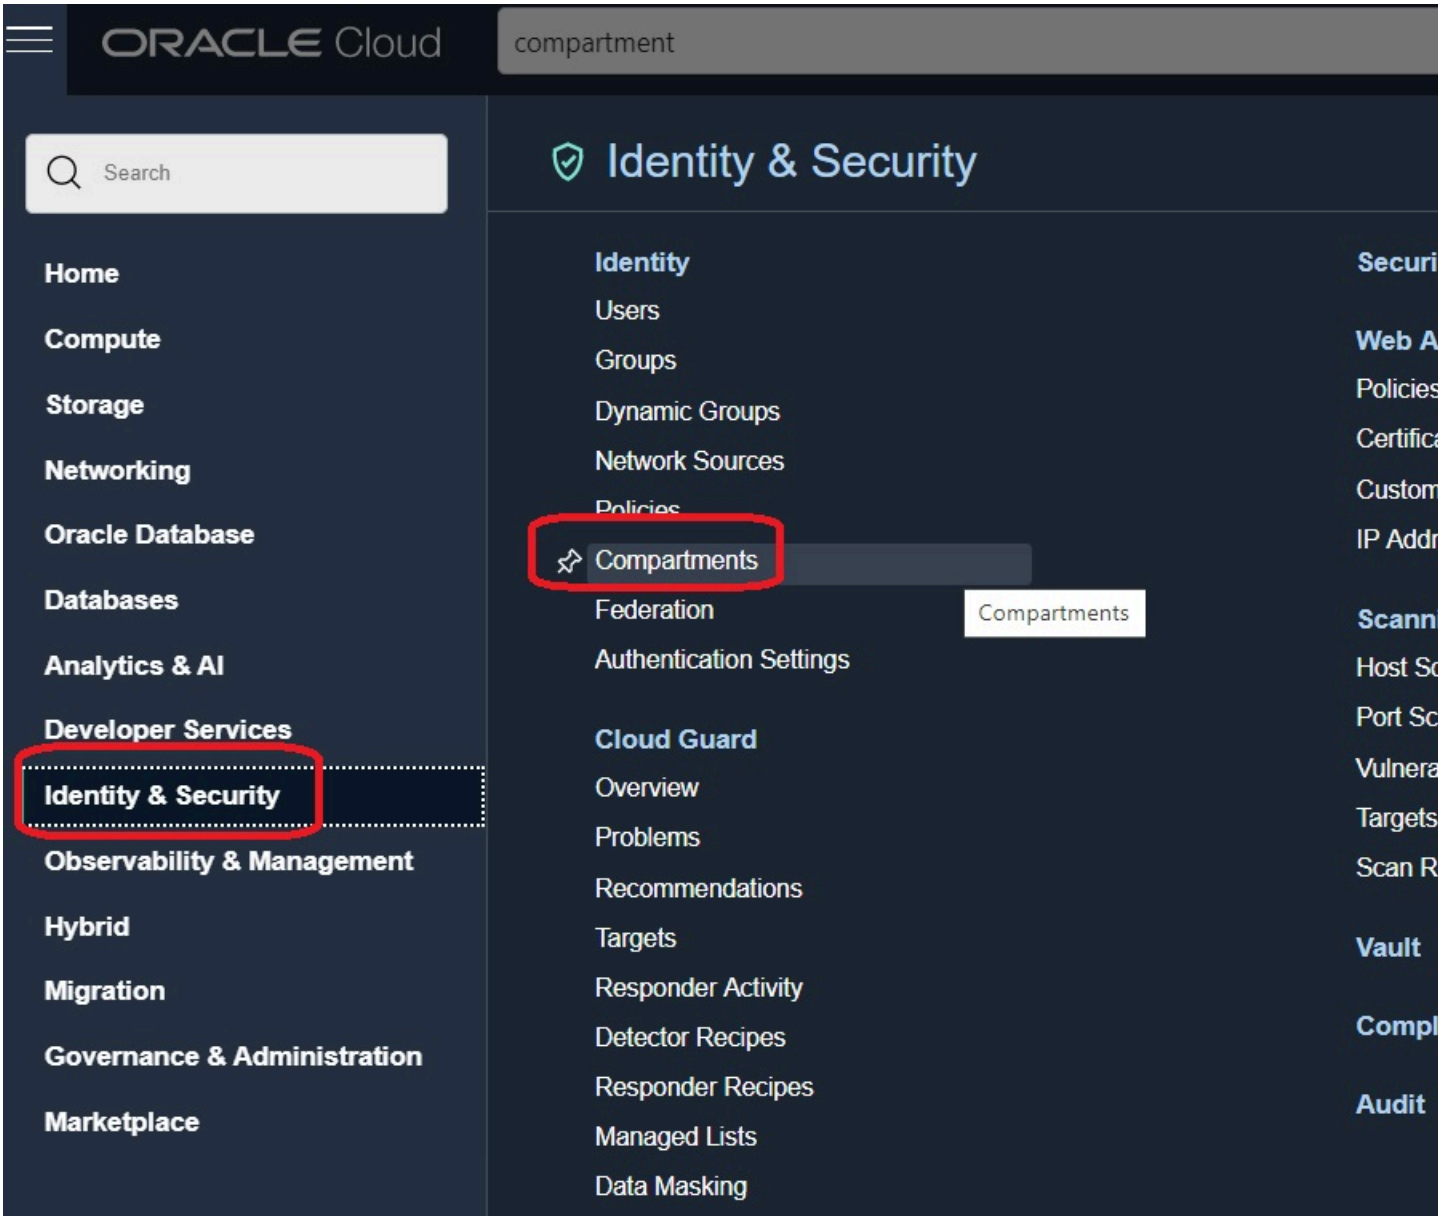

3. Click the **Identity & Security** section in the left pane, and select the **Compartments** service.

4. In the **Compartments** section, click the **Create Compartment** button.

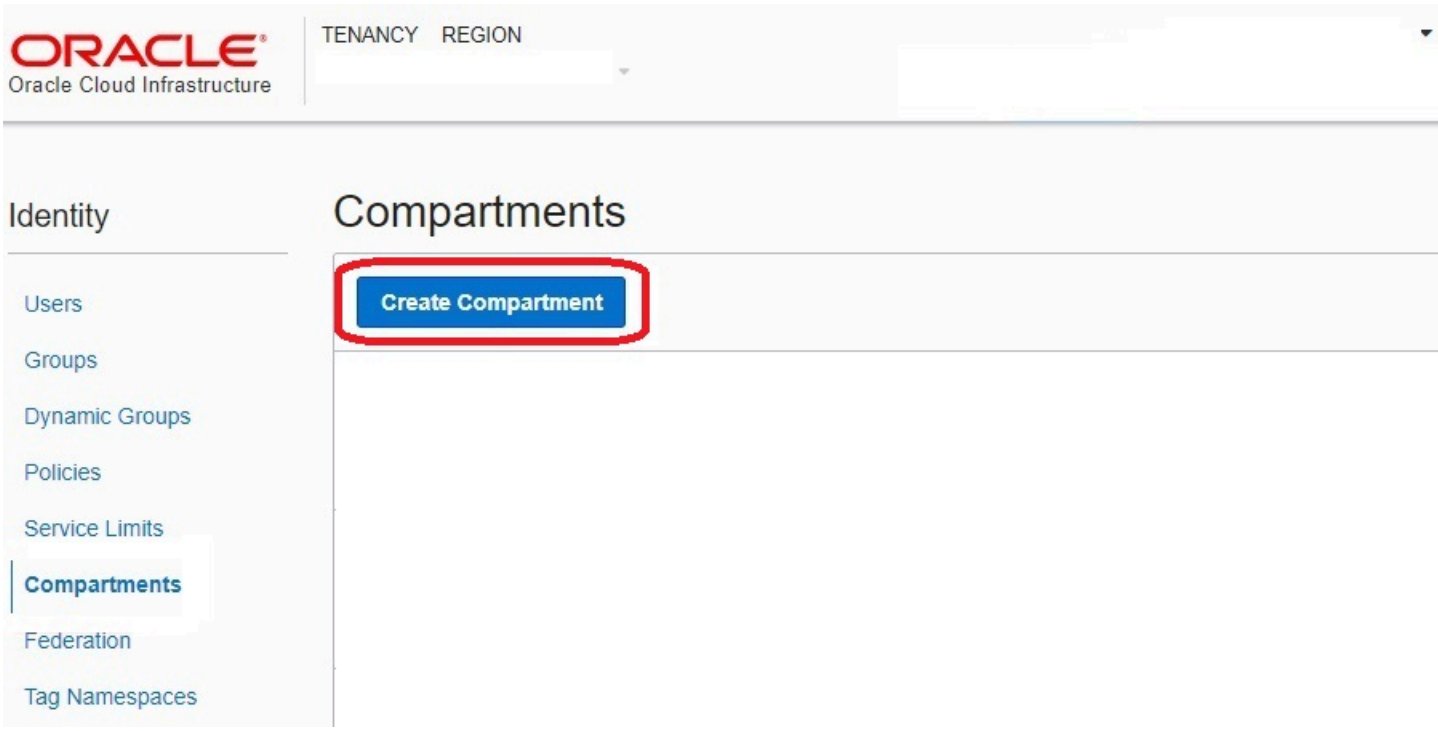

5. On the Create Compartment dialog box, complete these fields:
   - **Name**

     Enter a name for the compartment.
   - **Description**

     Enter a description for the compartment.
   - **Tags**

     Optionally you can enter tag information in these fields. For more information, click the link **Learn more about tagging** in the dialog box.

6. Click the **Create Compartment** button.

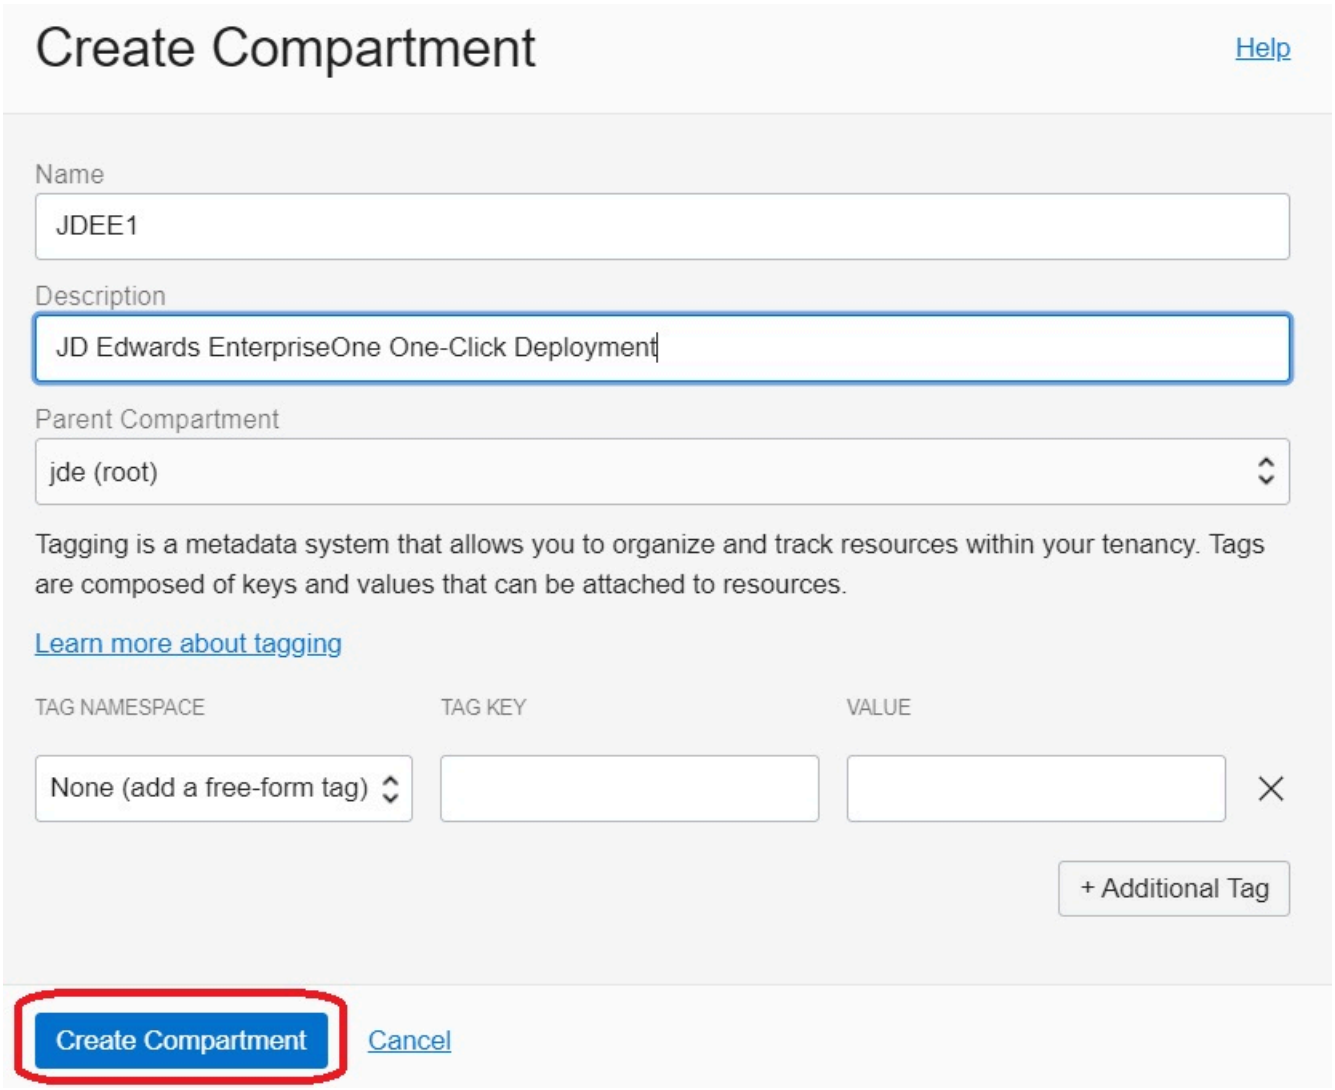

Creating a Virtual Cloud Network

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 for Infrastructure Provisioning

This section shows you how to create rules for a virtual cloud network (VCN) security list in Oracle Cloud Infrastructure to use the VCN for JD Edwards EnterpriseOne infrastructure provisioning.

Prerequisite

- You must have created a virtual cloud network by following the steps described in the preceding section "Creating a Virtual Cloud Network" of 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 and an administrator account in the platform. For more information, refer to this site: *Getting Started with Oracle Cloud*
- To access the Oracle Cloud Infrastructure Console, you must use a supported browser. See *Browser Issues*.

Creating Rules for a VCN Security List

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

1. Navigate to Networking, Virtual Cloud Networks.

2. In the List Scope section, verify that the correct compartment is selected.

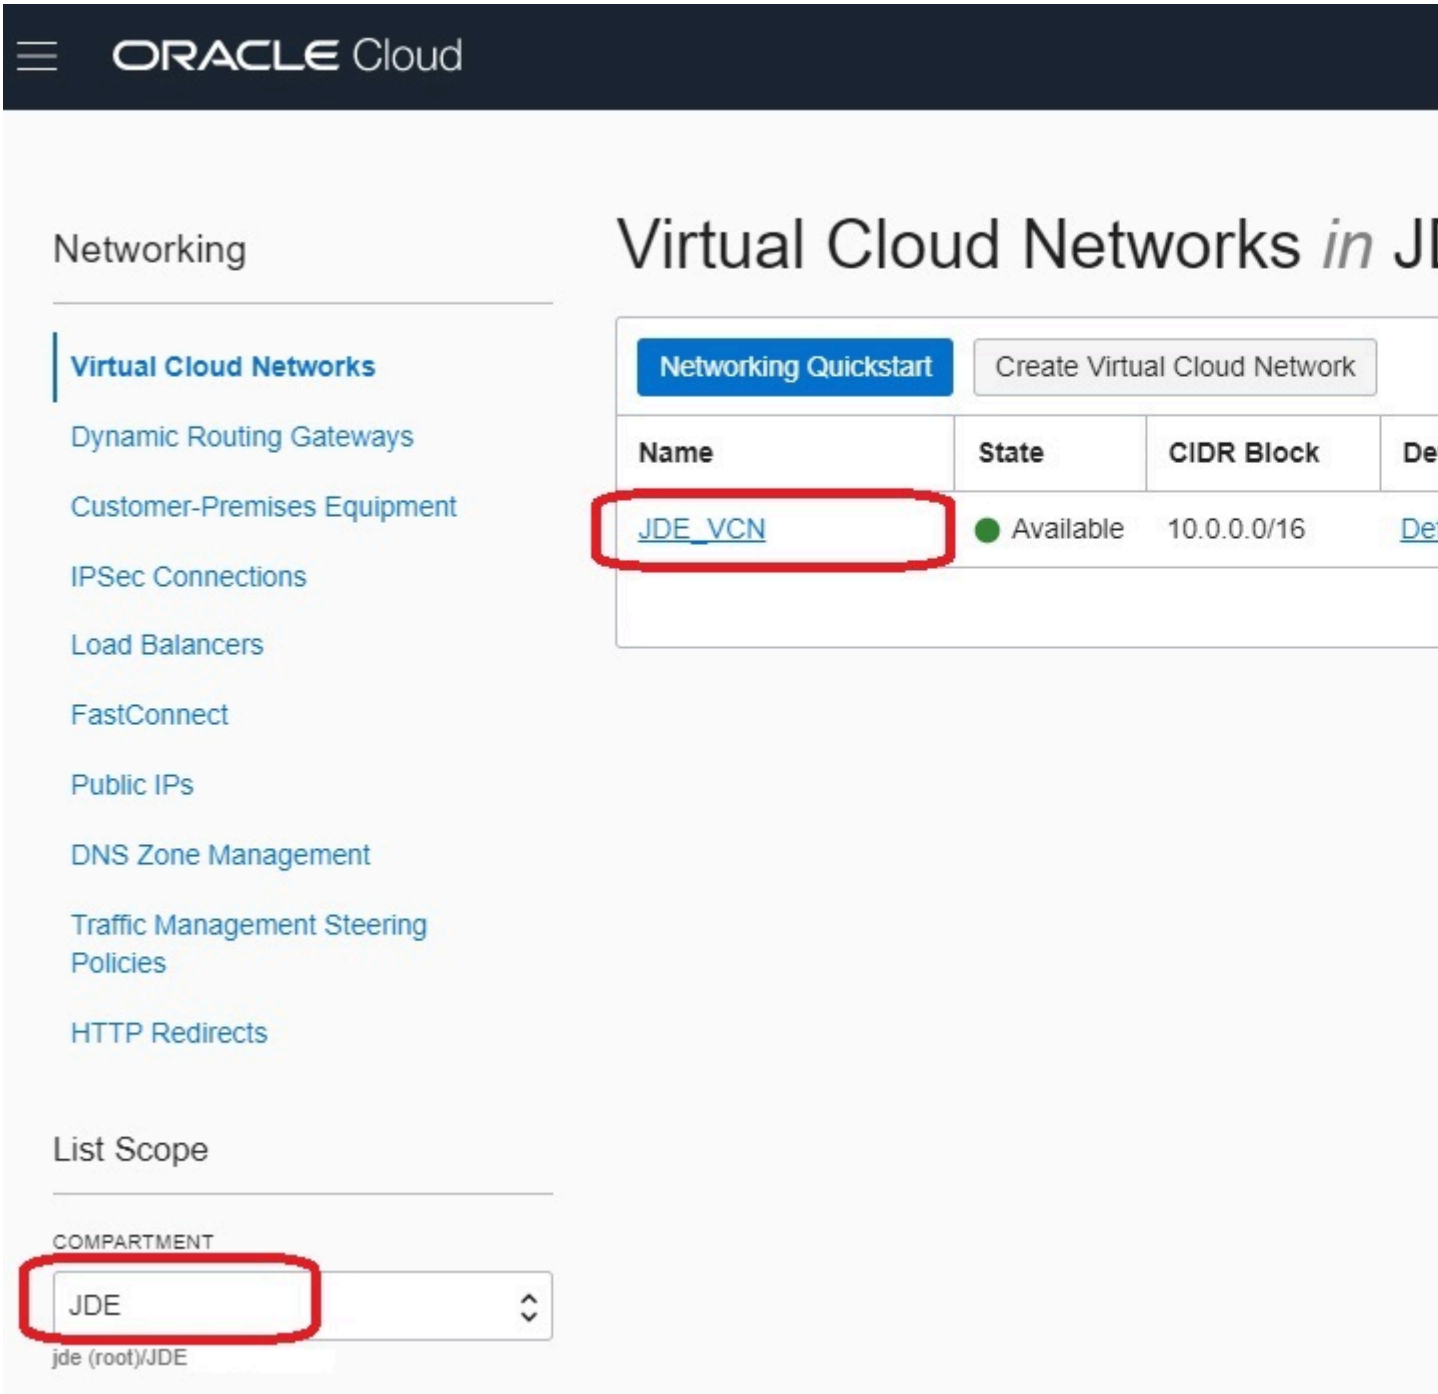

3. Click the link for the VCN you created.

4. In the Subnets in JDE_ Compartment, section, click the Public-subnet-<vcn_name> link.

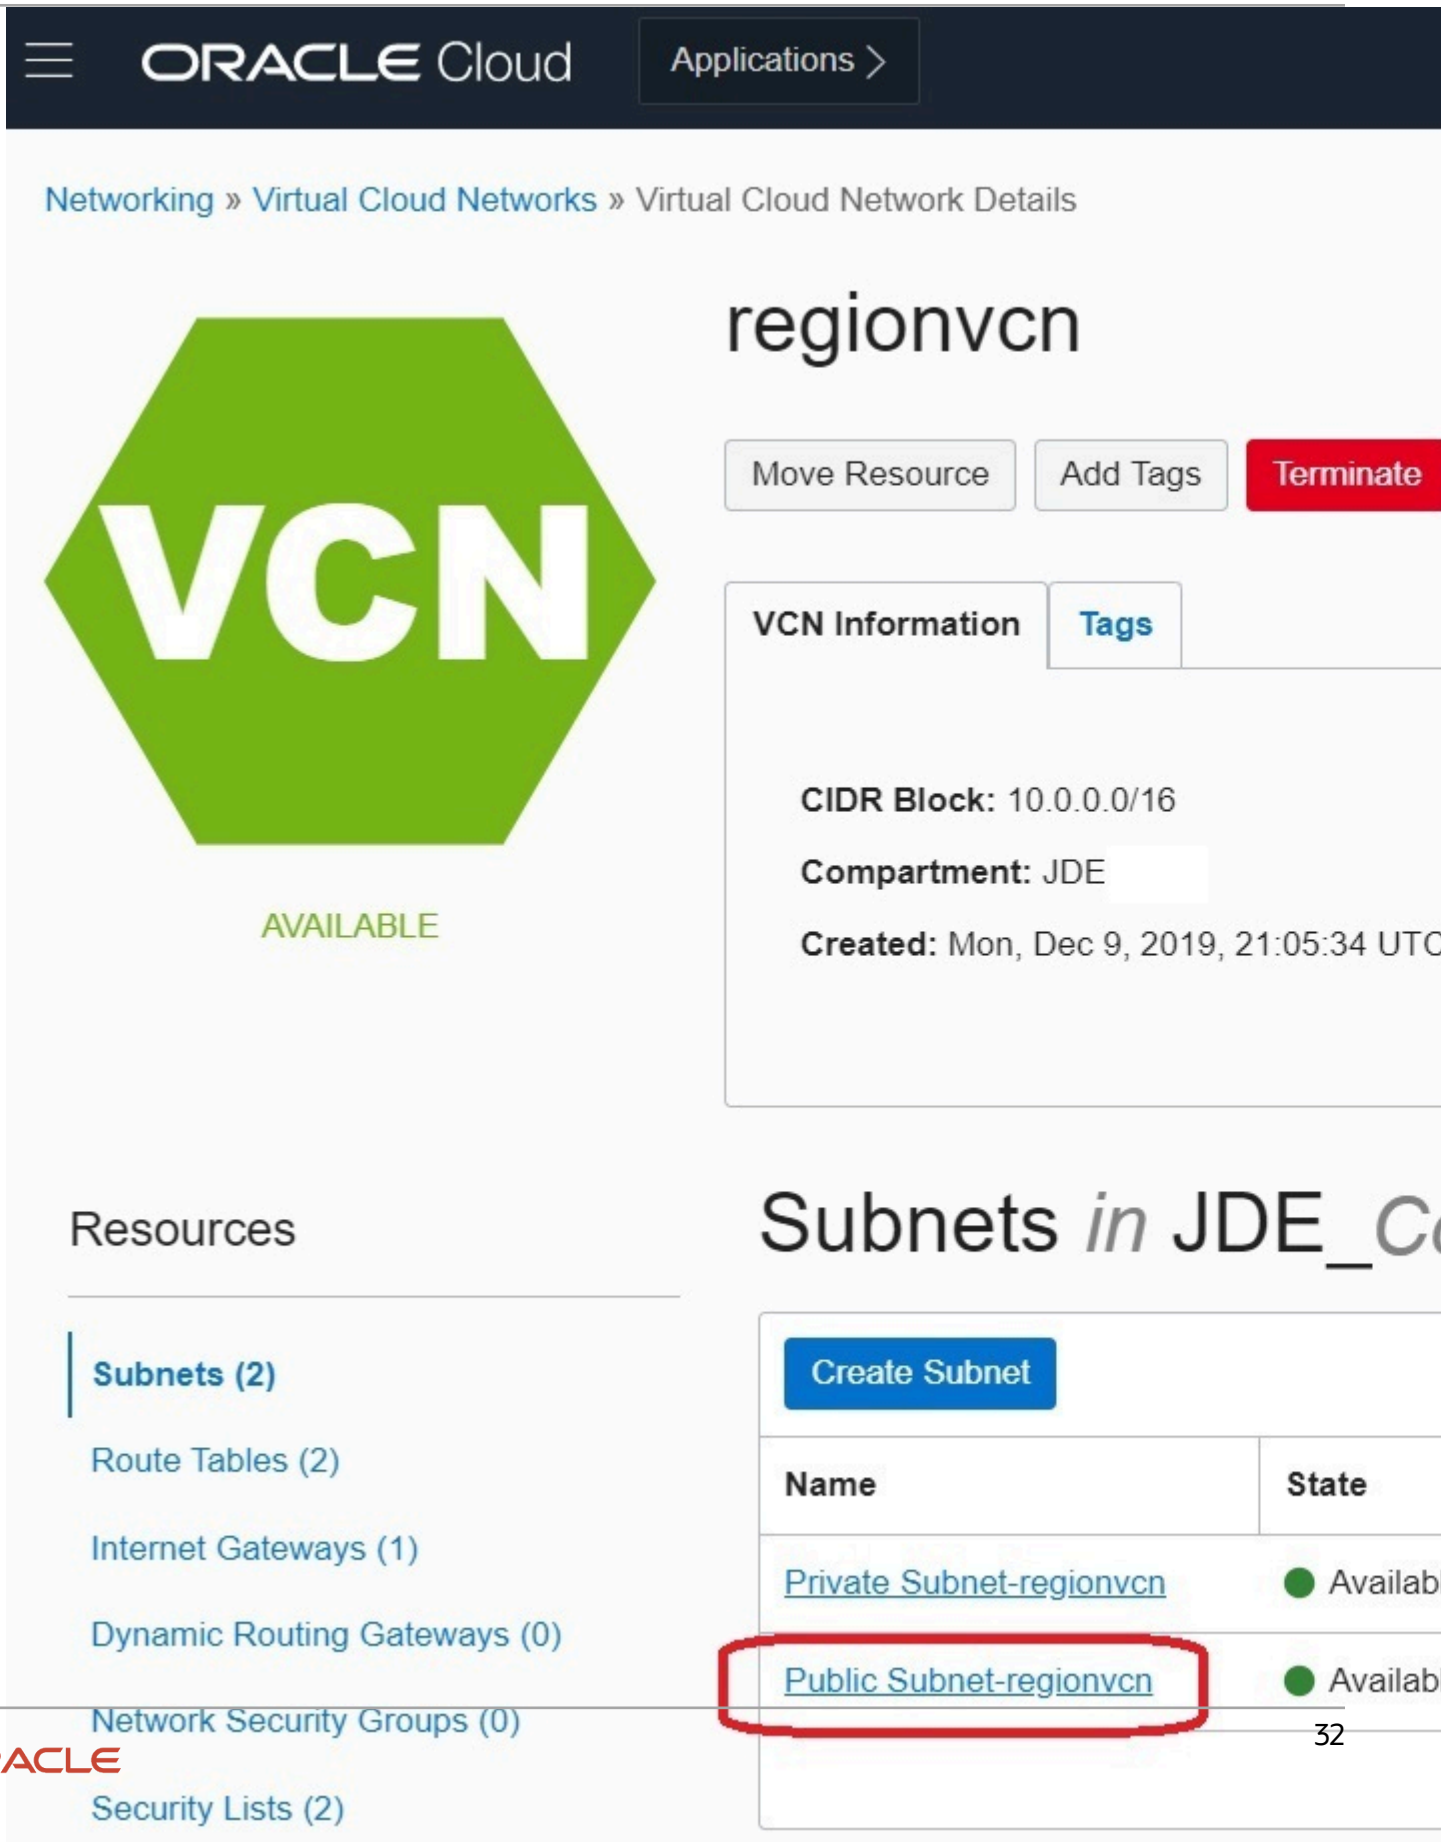
ORACLE Cloud
Applications
Networking » Virtual Cloud Networks » Virtual Cloud Network Details
regionvcn
VCN
AVAILABLE
Move Resource
Add Tags
Terminate
VCN Information
Tags
CIDR Block: 10.0.0.0/16
Compartment: JDE
Created: Mon, Dec 9, 2019, 21:05:34 UTC
Resources
Subnets (2)
Route Tables (2)
Internet Gateways (1)
Dynamic Routing Gateways (0)
Network Security Groups (0)
Security Lists (2)
Subnets in JDE_C
Create Subnet
Name
State
Private Subnet-regionvcn
Public Subnet-regionvcn

5. In the Security Lists section, click the link Default Security List for <vcn_name>.

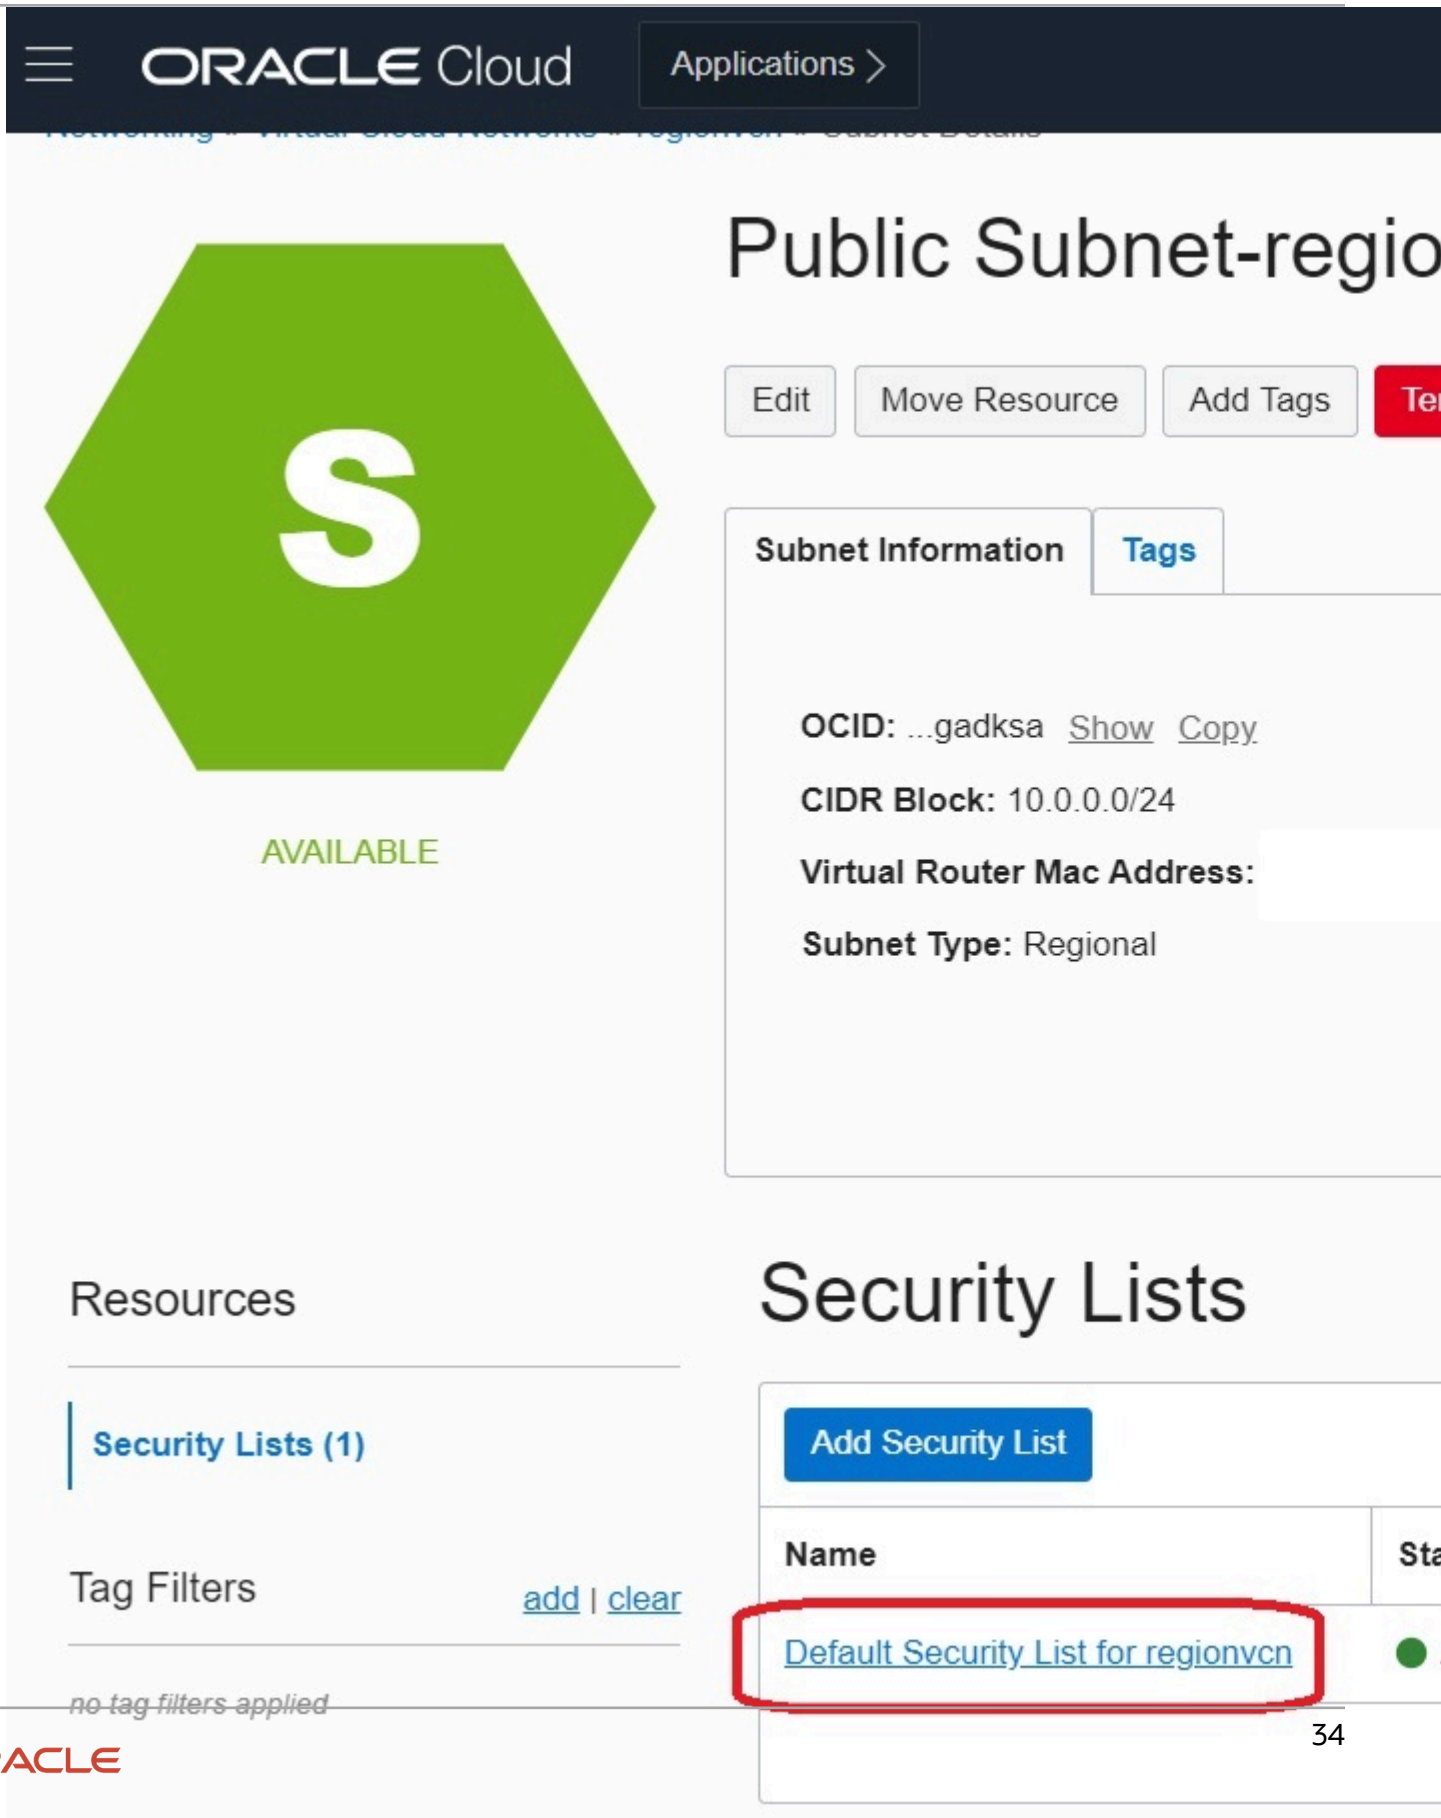
ORACLE Cloud
Applications
Public Subnet-regio
Edit
Move Resource
Add Tags
Subnet Information
Tags
OCID: ...gadksa Show Copy
CIDR Block: 10.0.0.0/24
Virtual Router Mac Address:
Subnet Type: Regional
AVAILABLE
Resources
Security Lists (1)
Tag Filters
add | clear
no tag filters applied
Security Lists
Add Security List
Name
Default Security List for regionvcn

6. To create additional rules that are required for JD Edwards EnterpriseOne One-Click Provisioning, click the Add Ingress Rule button.

   Because this VCN is only for the infrastructure staging server, you only need to open port 5901 with a source CIDR of 0.0.0.0/0. This is the listen port of the VNC Server.
7. Use these values to define the new rule for the listen port of the VNC server:

   Source CIDR: 0.0.0.0/0

   Destination Port Range: 5901
8. Click the Save Security List Rules button to complete the setup for Ingress Rules.
9. You can accept the default Egress Stateful rule that allows Egress to all destinations, all protocols, and all traffic for all ports.

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.

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".

4 Creating a Linux VM

Obtaining the Image and Creating an Instance for the OpenTofu Staging Server

This section shows you how to obtain the image for the requisite Linux system and create a Linux instance for the One-Click Provisioning Server for use by JD Edwards EnterpriseOne on 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 supports only virtual machine instances.

For additional information refer to *Creating an Instance* in the OCI documentation.

Prerequisites

- For instructions on using the different types of credentials while working with Oracle Cloud Infrastructure, refer to *Security Credentials* in the Oracle Cloud Infrastructure documentation.
- You must complete the steps described in the section *Required Keys and OCIDs* in the Oracle Cloud Infrastructure documentation.
- 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 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 this site:
  - *Getting Started with Oracle Cloud*
- To access the Oracle Cloud Infrastructure Console, you must use a *supported browser*.

Obtaining the Image for a OpenTofu Staging Server

Use this procedure to obtain the image and create a Linux instance for the OpenTofu staging server.

1. On **Oracle Cloud Infrastructure Console**, click the navigation menu in the upper-left corner.

2. From the navigation menu, select **Marketplace**.

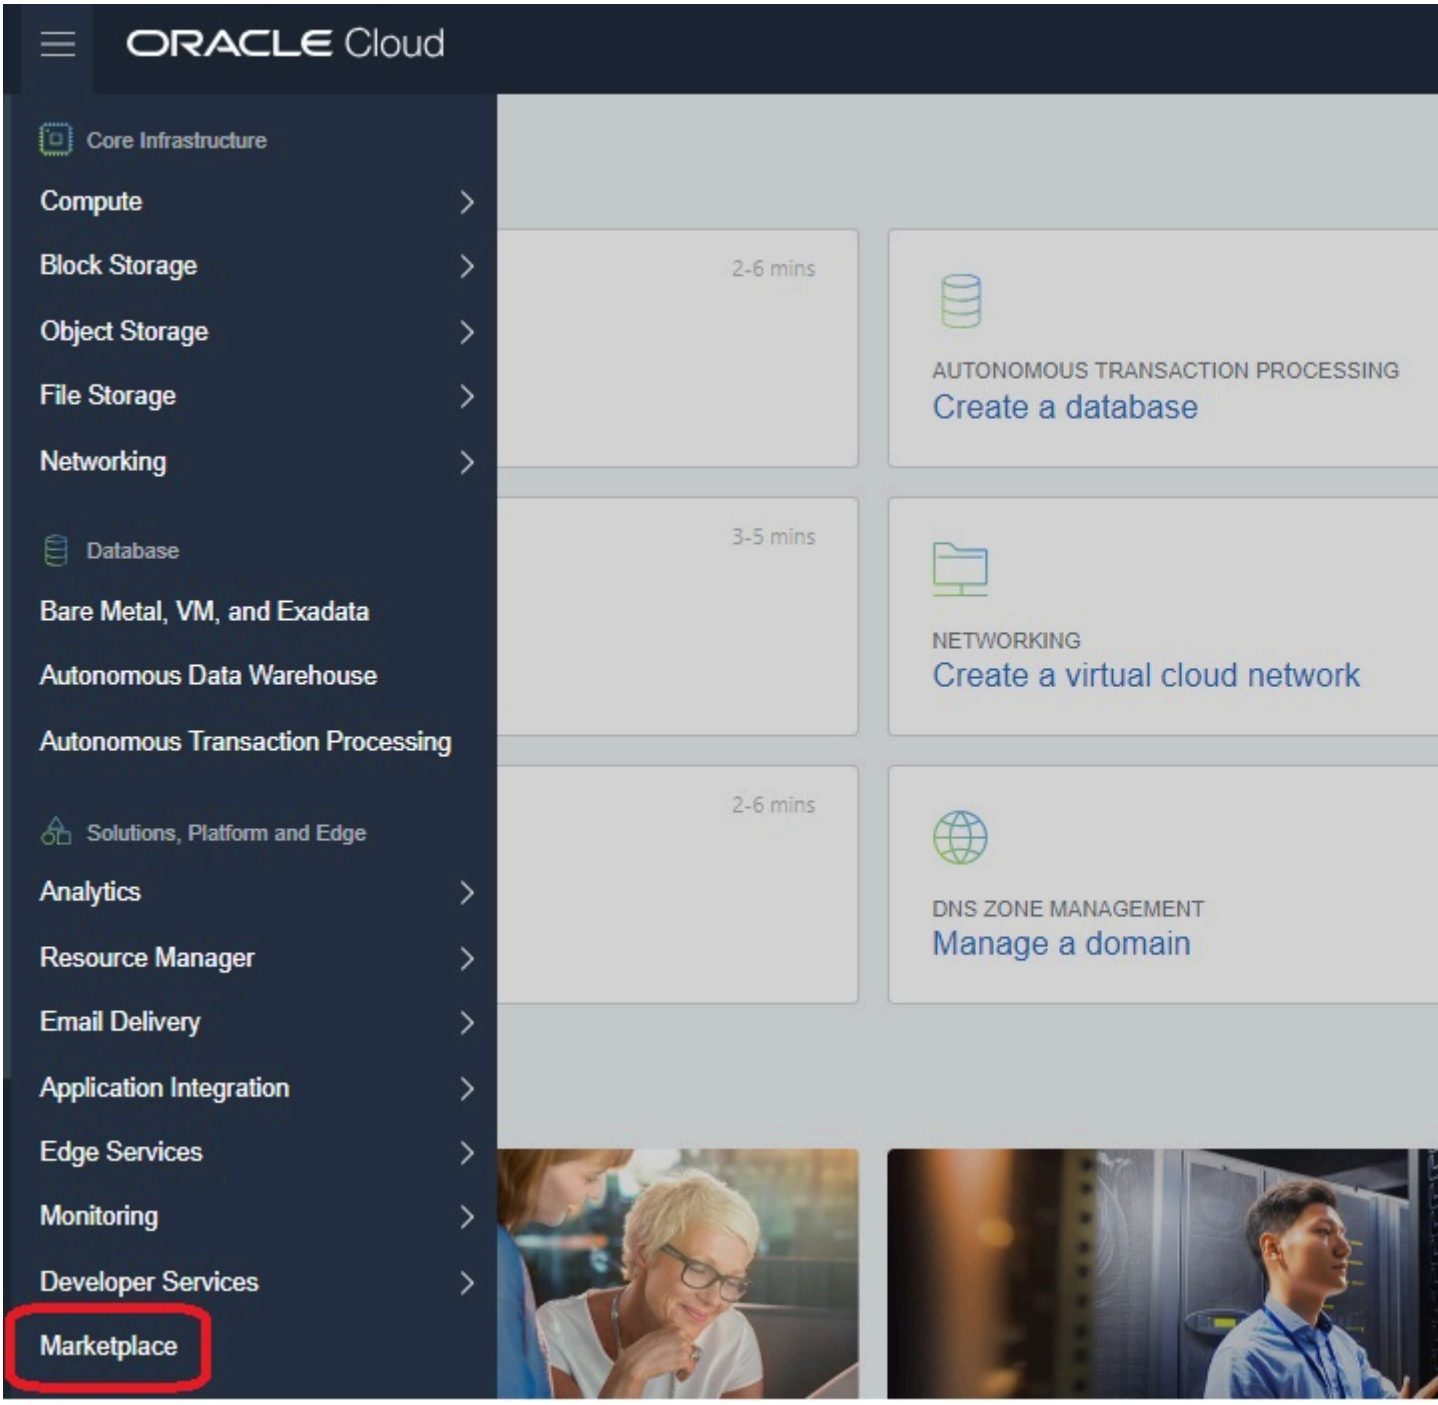

3. In the **Marketplace** page, under **All Applications**, locate the listing for the JD Edwards EnterpriseOne JD Edwards EnterpriseOne Reference Architecture image.

> **Tip:** To more easily locate this image, in the browser page, search for the entire or part of the listing title, which is "JD Edwards EnterpriseOne Reference Architecture (OpenTofu)".

4. Click the Marketplace listing for the JD Edwards EnterpriseOne Reference Architecture image.
5. On the Marketplace listing for JD Edwards EnterpriseOne Reference Architecture, in the **Compartment** field, use the drop-down list to select the compartment that you previously created in the OBE "Creating a Compartment" of this Learning Path.
6. On the Launch Instance page, review the Terms of Use and click the check box to accept the terms.

7. Click the **Launch Instance** button to continue.

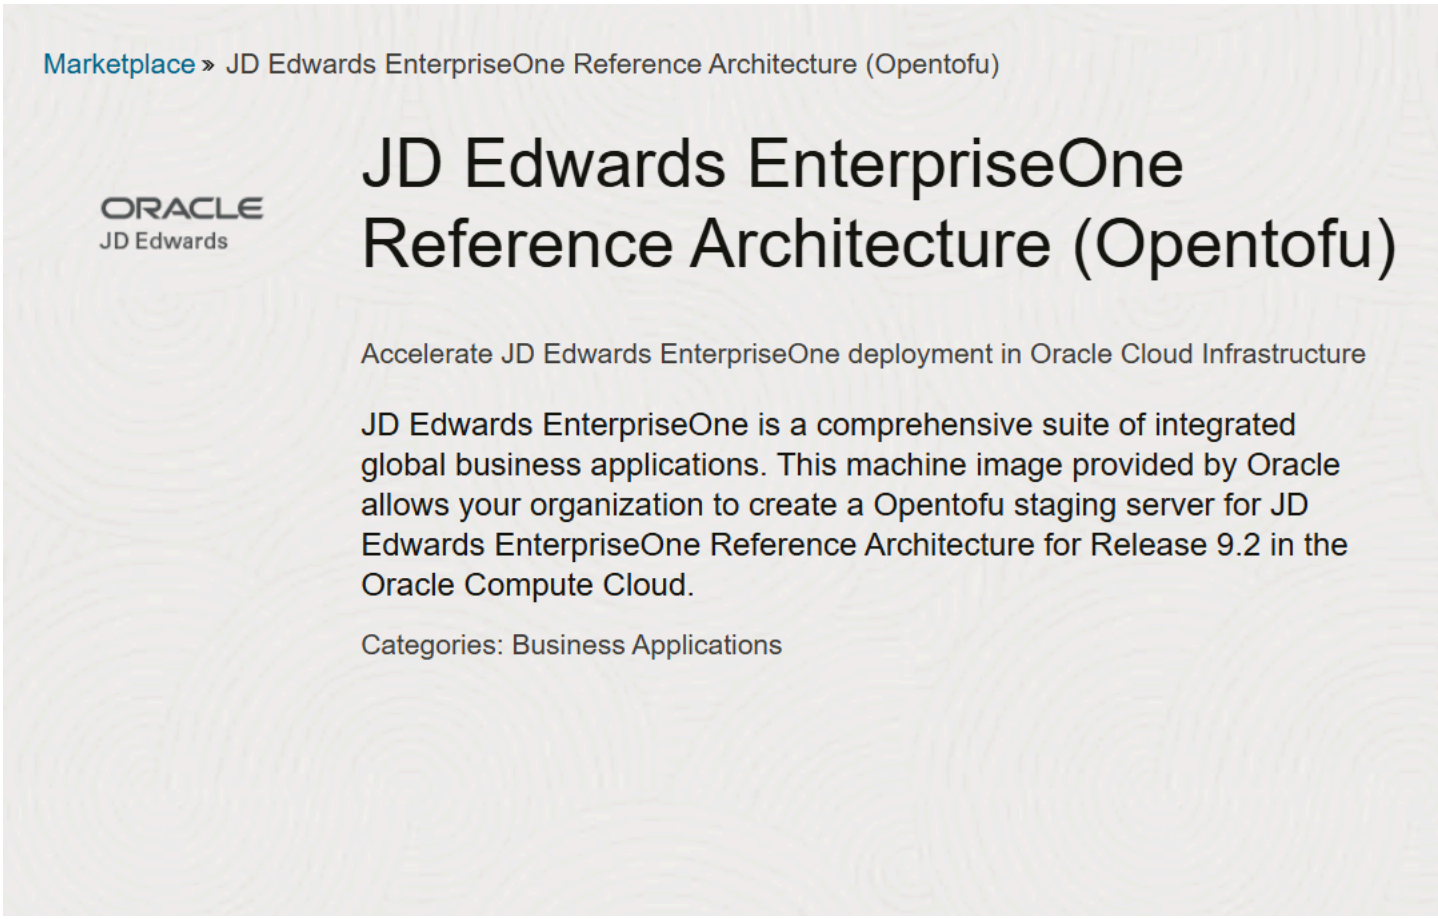

8. On the **Create Compute Instance** page, complete these fields:
   - **Name**

     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 again 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 Reference Architecture image server. For example, the staging server for reference architecture might be called **jderefarch**.

     Note: Ensure that the host name of the JD Edwards EnterpriseOne Reference Architecture image instance contains only alphabetic characters and numbers. 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 must be alphanumeric with only lowercase letters and should not exceed 15 characters. This is the character count supported by the JD Edwards EnterpriseOne database table and application design. Do not use a fully qualified domain name for a host name. You should only specify the first node of the domain name. If the existing host name 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 an availability domain (AD) to select it. You can only specify a single AD.

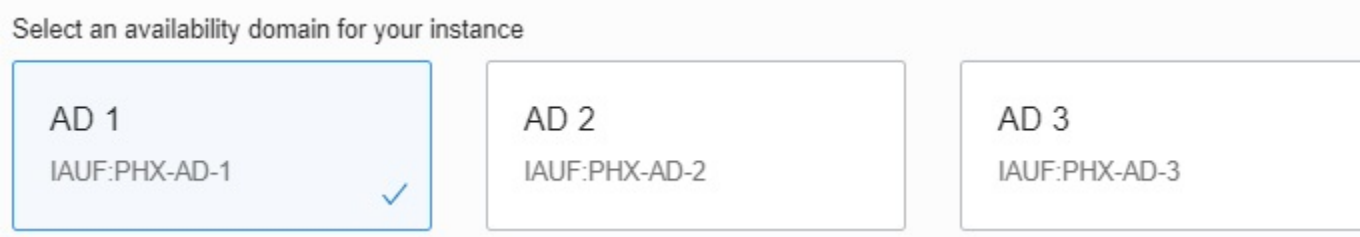

- **Image**

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

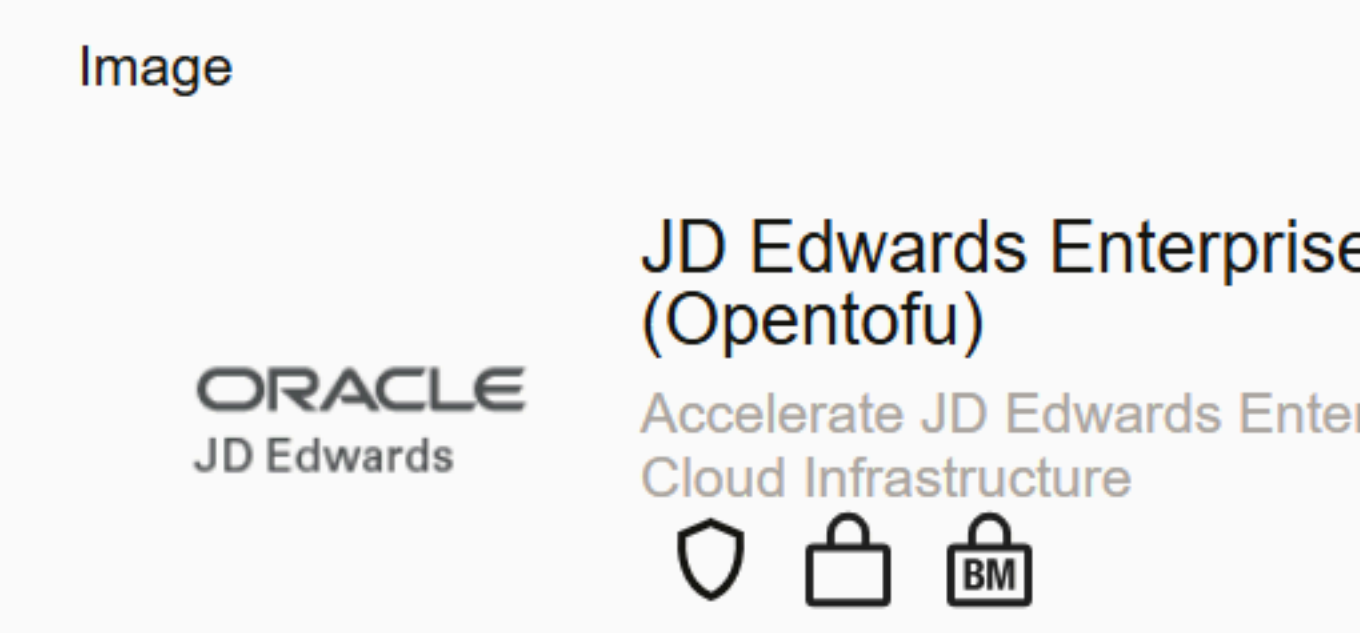

- **Shape**
  - Click the **Change Shape** button.

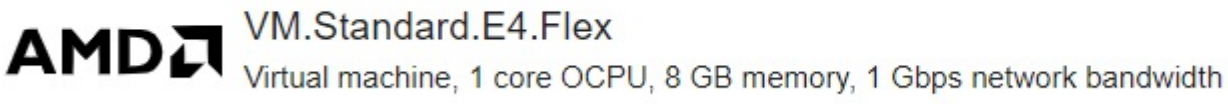

  On Browse All Shapes, make the following selections:
  - Instance type

Virtual Machine

- Shape series

  Specialty and previous generation
- Shape Name

  VM.Standard2.2

  This is the recommended shape for the OpenTofu Staging Server, which minimally includes 2 OCPUs and 30 GB memory.

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

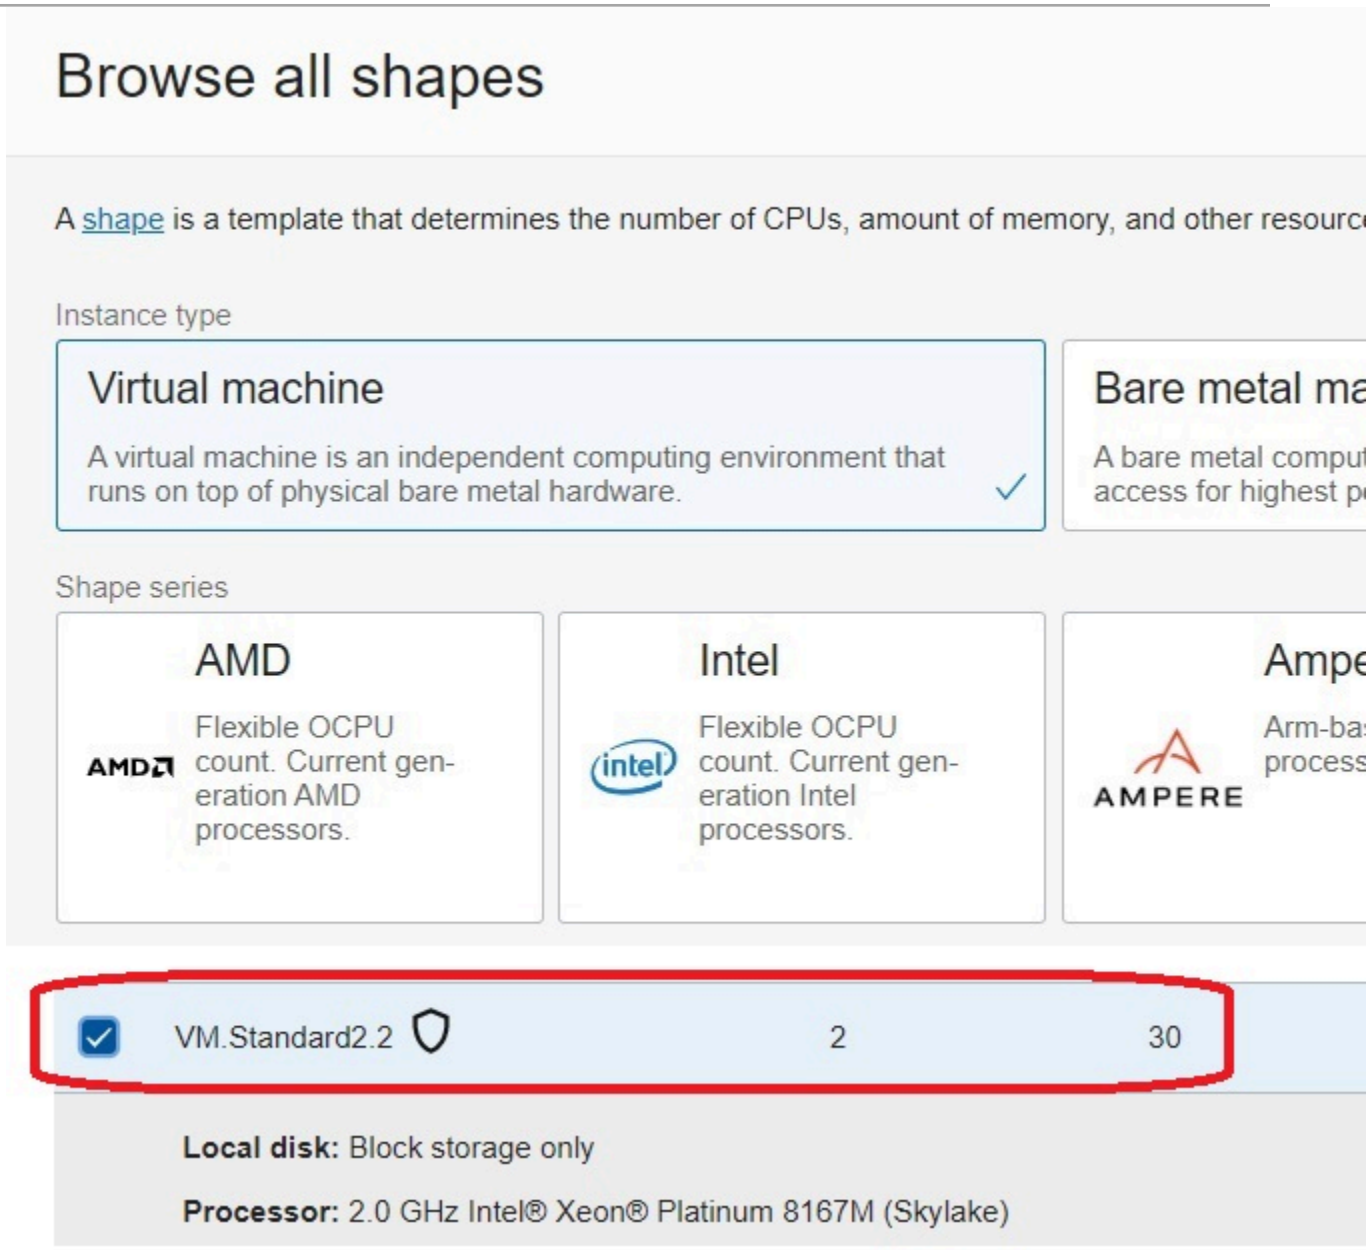

- **Configure networking**

  Verify, and if necessary, set your network settings for the instance of the JD Edwards EnterpriseOne Reference Architecture image. These settings include:

  - Virtual cloud network compartment
  - Virtual cloud network
  - Subnet compartment
  - Subnet

Note: You must select the option **Assign a public IP Address**.

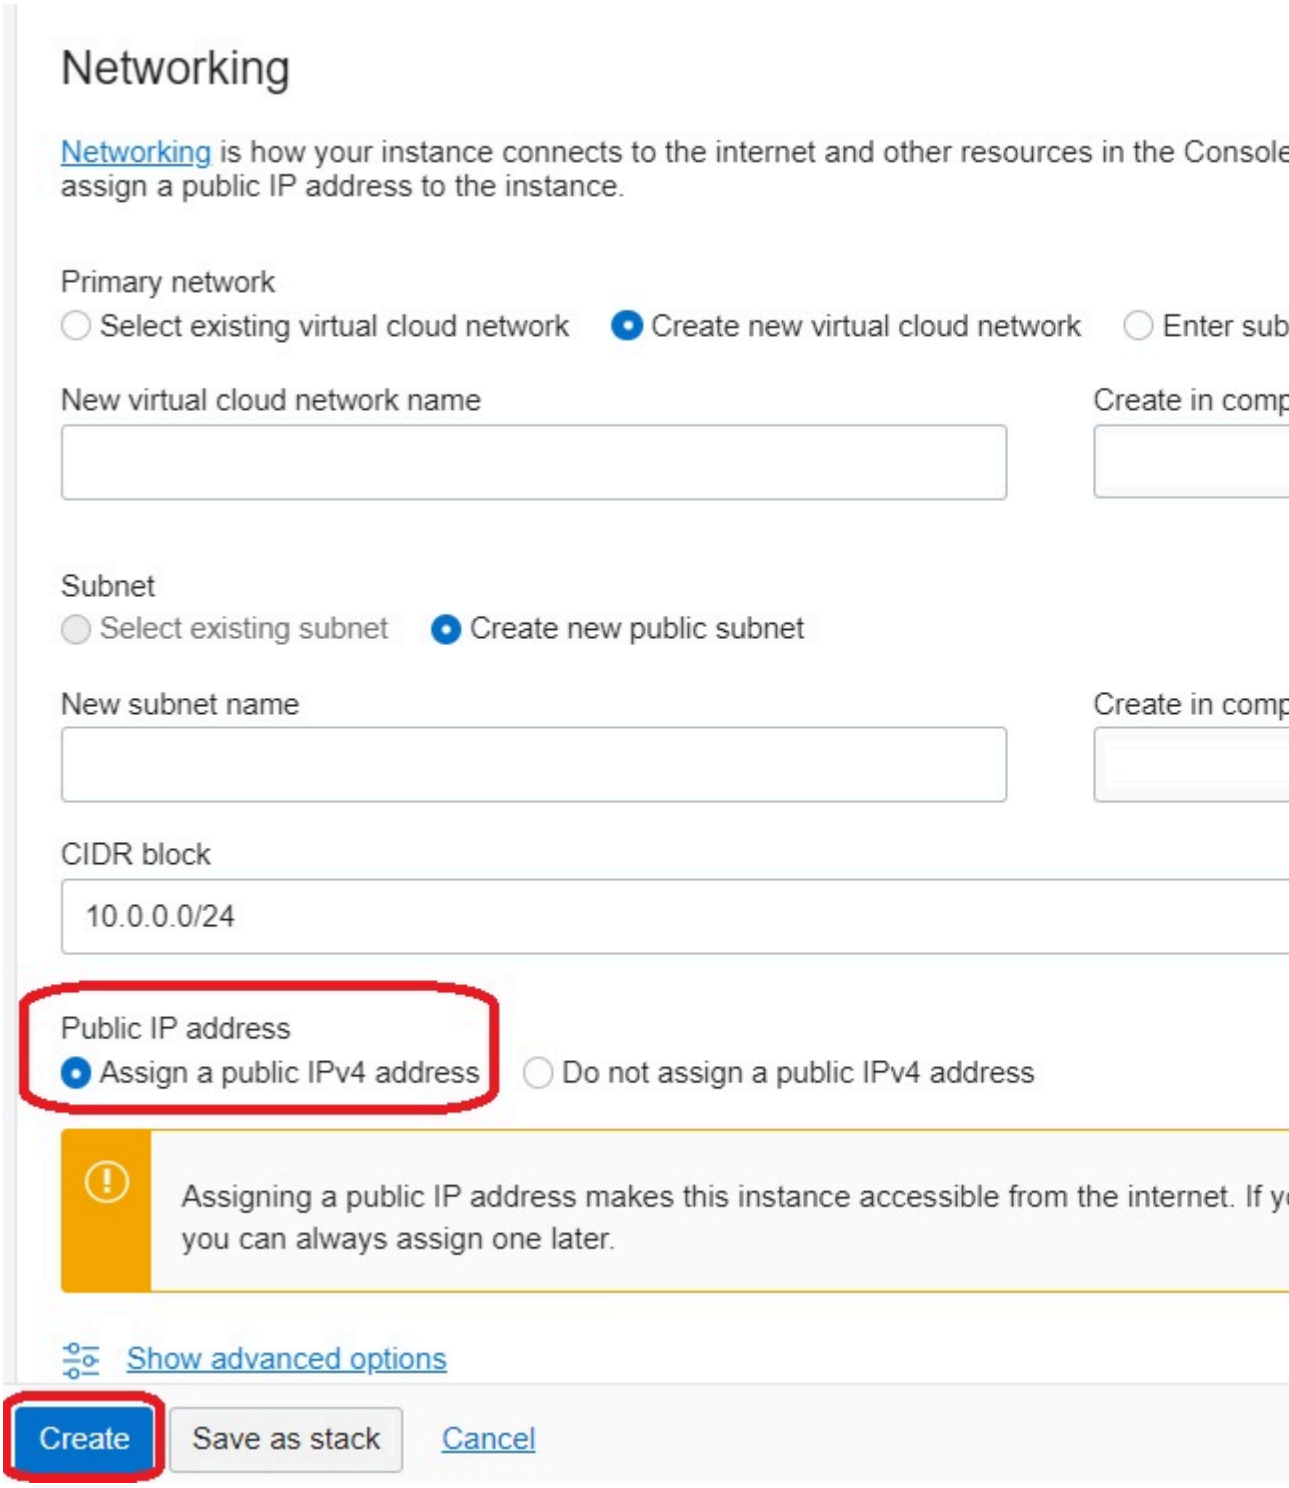

- **Add SSH Keys**

You must add the SSH keys that you previously created for use with the JD Edwards EnterpriseOne Reference Architecture image.

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

> **Note:** If you followed the recommendation in the previous section "Generating Instance Key Pairs in openssh Format", this file is named jdeSSHKey.pub.

- **Configure Boot volume**
  - Click the **Specify a custom boot volume size** check box.
  - Enter a numeric value in the size field. For the JD Edwards EnterpriseOne Reference Architecture image, the minimum recommended value is **100**.
  - Optionally you can choose to enable encryption using the checkbox for **Use In-Transit Encryption**.
  - Do not select any other check boxes.

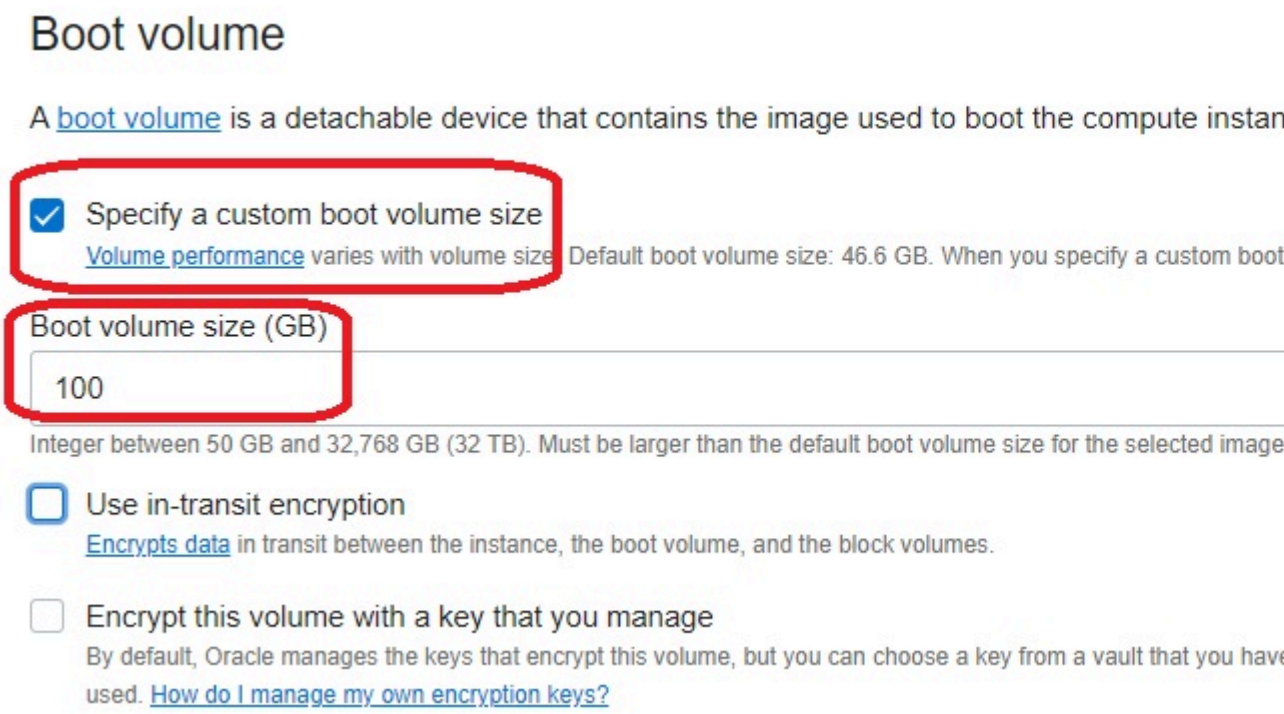

- **Show Advanced Options**

  You do not need to set any Advanced Options.

9. Click the **Create** button to create the Oracle Cloud Infrastructure instance for the JD Edwards EnterpriseOne Reference Architecture 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 to complete the first-time configuration of the Reference Architecture image for JD Edwards EnterpriseOne as described in the following section.

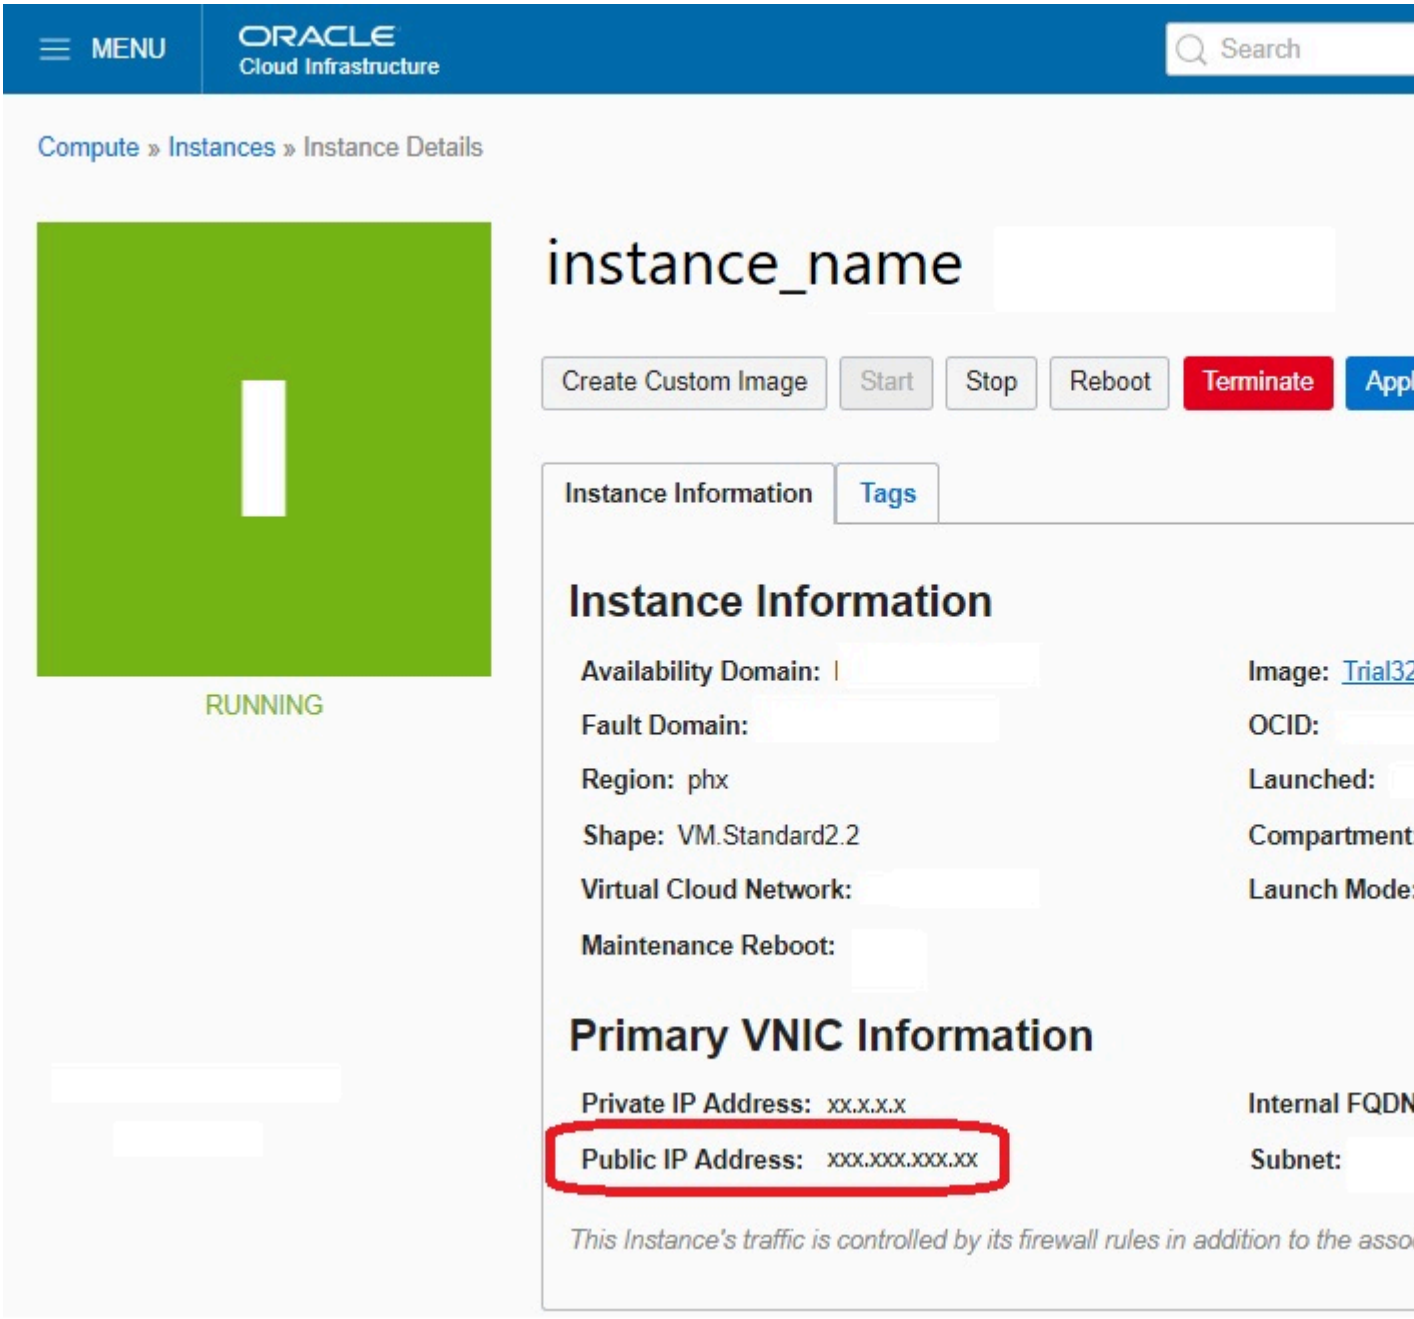

5 Setting Up the OpenTofu Staging Server

Logging in to the Linux Instance for the OpenTofu Staging Server

This section describes how to log in to the Linux instance that you created for the OpenTofu staging server for JD Edwards EnterpriseOne One-Click Provisioning on Oracle Cloud Infrastructure.

You must log in to the OpenTofu staging server to perform the subsequent setup at the Linux operating system level.

Prerequisite

- You must have previously created a Linux instance for the OpenTofu staging server by following the steps described in the section "Obtaining the Image and Creating an Instance for the OpenTofu Staging Server" of this Learning Path.
- You must have the open source Putty terminal emulator program installed on the client machine that you will use to connect to the OpenTofu staging server.

Logging in from a Microsoft Windows System

Use this procedure to log in to the Linux instance that you created for the OpenTofu staging server.

1. On your local machine, launch the PuTTY terminal emulator application.

2. PuTTY, in the Sessions category, complete these fields:
   - Host Name (or IP Address)

     Use the public IP address that was assigned by Oracle Cloud Infrastructure when you created the Compute instance for the OpenTofu staging server.
   - Port

     Use Port 22. <Use port 22.>
   - Connection Type

     Ensure that the **SSH** option is selected.

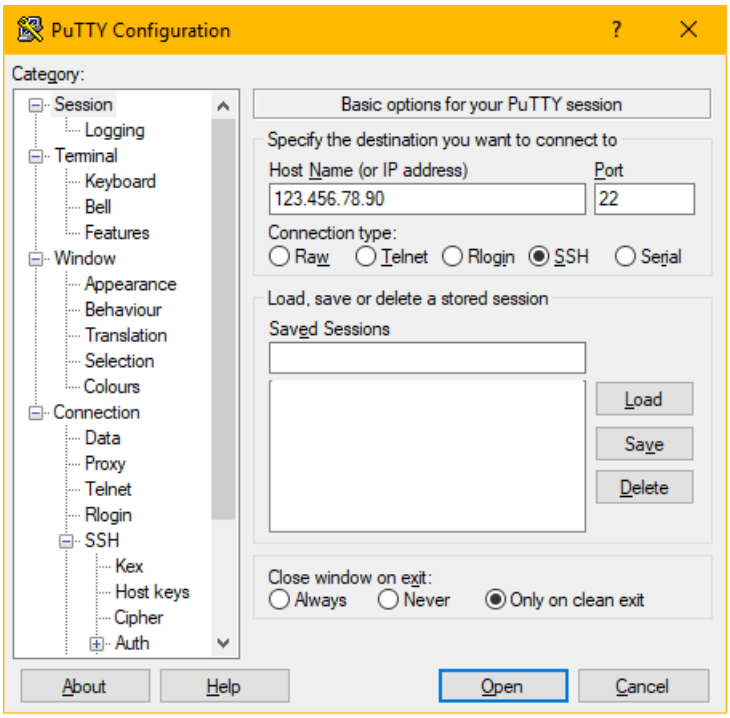

3. In PuTTy, navigate to **Connection**, **SSH**, **Auth category**, and browse to the location of your local private key for authentication.

> **Note:** If you followed the recommendation in the previous section "Generating Instance Key Pairs in openssh Format", this file is named `jdeSSHKey.ppk`.

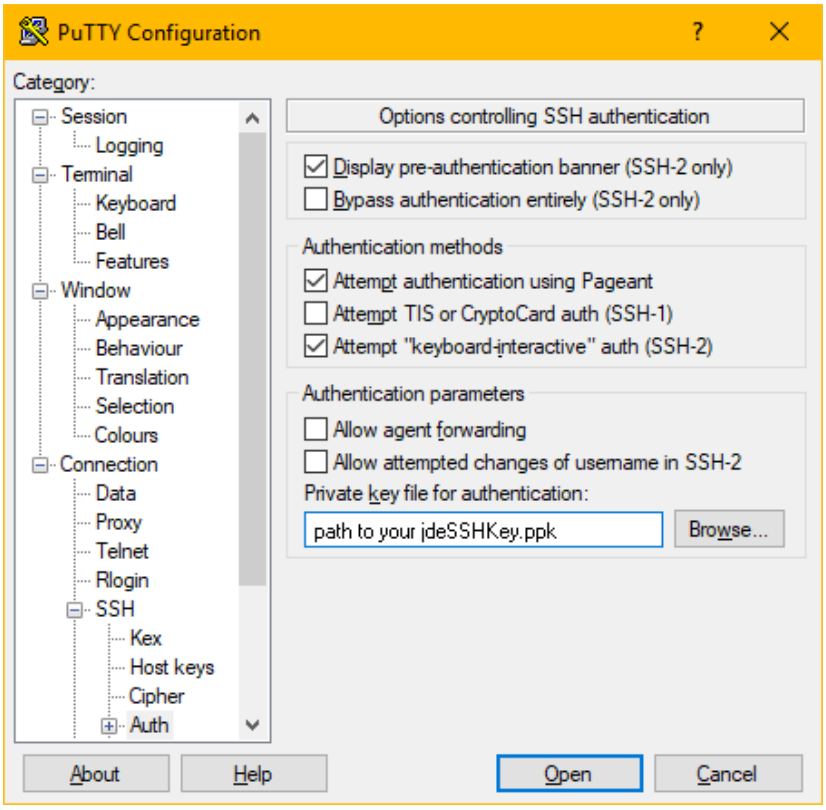

4. In PuTTy, navigate to Connection, Data, and in the Login details section enter opc as the value for Auto-login username.

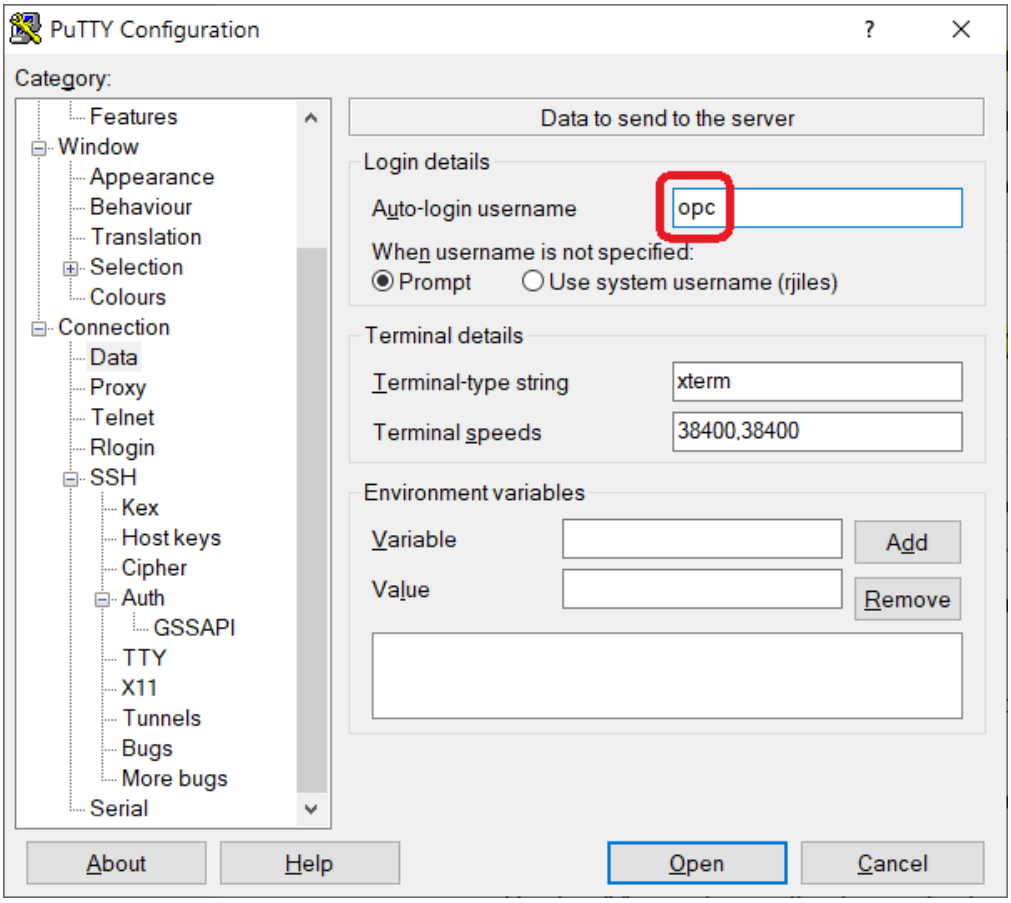

5. Click the Open button to establish a connection between your local machine and the Linux instance in Oracle Cloud Infrastructure for the OpenTofu staging server.

Logging in from a UNIX or UNIX-Like Host

Use this procedure to log in to the Linux instance that you created for the OpenTofu staging server.

1. The format of this command can vary depending upon your UNIX server. Enter one of the following commands:

```
ssh -l opc <public_ip_address_of_the_instance>
```

-or-

```
ssh -i /path_to_private_key opc@<public_ip_address_of_the_instance>
```

> **Note:** To determine the public IP address of your instance, go to Instance Information for the instance that you created for the OpenTofu staging server. This is illustrated in an example in the previous section in this Learning Path.

2. If the message Are you sure you want to continue connecting (yes/no)? is displayed, select **yes**.
3. When prompted, enter the passphrase you had provided for your SSH key pair.

4. On first login, you will be prompted for the public IP address. Enter the IP address used in Step 2 to connect to this server. This step is only done once on the first login.

> **Note:** If you are not prompted for the public IP address on the first login, log out and log back in after a few minutes.

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 OpenTofu Staging Server, you must run the following commands to complete the resize functionality.

> **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.

> **Note:** Be sure to issue the commands in this order:

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

sudo yum install python3-oci-sdk python3-oci-cli

sudo yum update oci-utils

sudo systemctl start ocid.service
```

For additional information, refer to this Oracle documentation:

https://docs.oracle.com/en-us/iaas/Content/Compute/References/oci-growfs.htm

Performing First-Time Configuration of the OpenTofu Staging Server

This section shows you how to perform the first-time configuration of JD Edwards EnterpriseOne OpenTofu Staging Server that is deployed as an image into Oracle Cloud Infrastructure. Before you can use the OpenTofu staging server, you need to perform the first-time configuration to specify certain port and password values.

Prerequisite

You must have logged in to the staging server by following the steps described in the section "Logging in to the Linux Instance for the OpenTofu Staging Server" of this Learning Path.

Performing First-Time Configuration of the OpenTofu Staging Server

When you log in to the OpenTofu Staging Server instance for the first time, or at any time thereafter without previously entering all the values into the first-time configuration script, the system automatically invokes a mandatory first-time configuration script. You must enter and confirm password values as described in this procedure.

> **Note:** If you only want to view contents on the OpenTofu Staging Server, there is no requirement to enter password information.

1. You must manually launch the first-time configuration script using this command:

   ```
   sudo ./createStage.sh
   ```

   

2. As the first-time configuration script of the OpenTofu staging server runs, the system displays the following prompt. At this prompt, you must set and confirm a password for the VNC server that this first-time configuration script installs and configures on the OpenTofu staging server.

   The password values you enter are masked. As with all password values, you should keep a record of the values that you enter because these values are encrypted and cannot subsequently be retrieved or decrypted.

   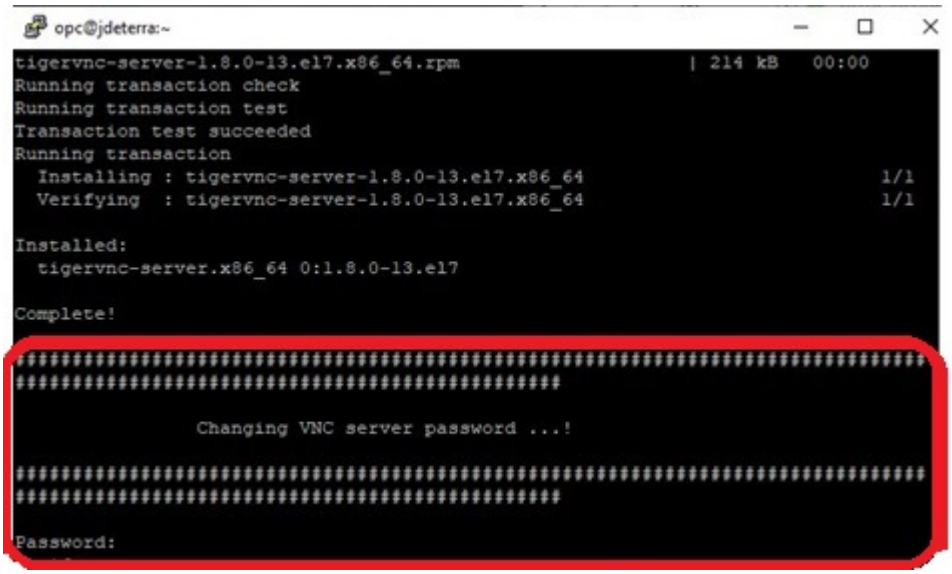

3. Next the first-time configuration scripts prompts you to enter the password for the opc user. This password will be used by the system to authenticate the opc user that is created by the script.

   The password values you enter are masked. As with all password values, you should keep a record because the values that you enter because these values are encrypted and cannot subsequently be retrieved or decrypted.

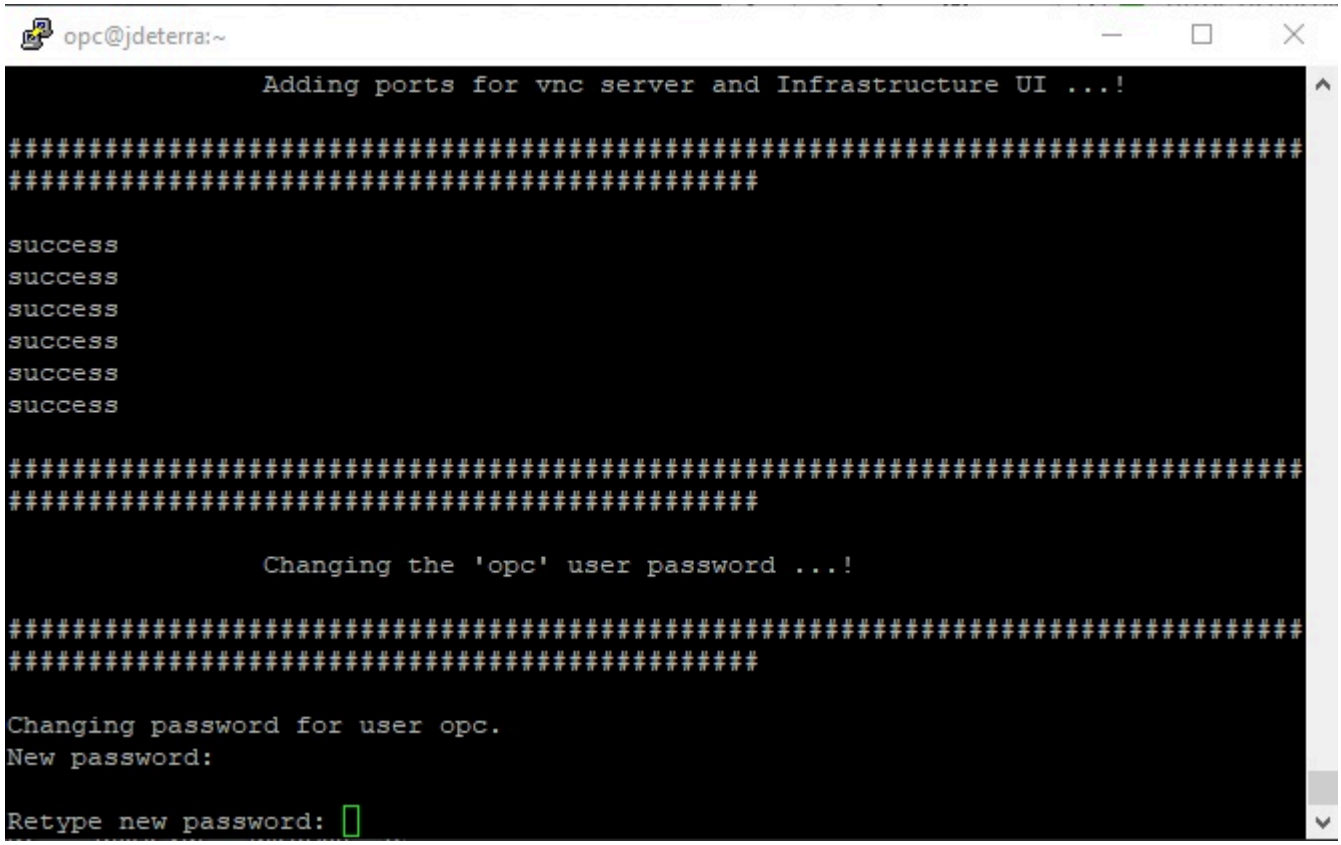

After the script completes running, the system reboots the instance for the OpenTofu staging server and starts the Infrastructure Provisioning Console.

To access the console, you must log back in to the OpenTofu staging server using a VNC viewer as detailed in the next section: Accessing the OpenTofu Staging Server Through the Virtual Network Computing (VNC) Viewer.

Accessing the OpenTofu Staging Server Through the Virtual Network Computer (VNC) Viewer

This section shows you how to access the OpenTofu staging server using a virtual network computer (VNC) viewer. This method of connection is required to perform subsequent tasks in this Learning Path.

Prerequisite

You must have run the Linux commands to install and set up the VNC server by following the steps described in the previous section "Running Linux Commands to Set Up the OpenTofu Staging Server" of this Learning Path.

Accessing the OpenTofu Staging Server Through the Virtual Network Computing (VNC) Viewer

To use the JD Edwards EnterpriseOne Infrastructure Provisioning Console, you must be able to access the OpenTofu staging server using the VNC server that you previously installed. In this Learning Path, the VNC server that was installed is the open source application TigerVNC. In this case, you must use TigerVNC Viewer to connect to that server on the OpenTofu staging server.

1. Start the TigerVNC Viewer to connect to the VNC server. For example, for TigerVNC, the following dialog box is displayed:

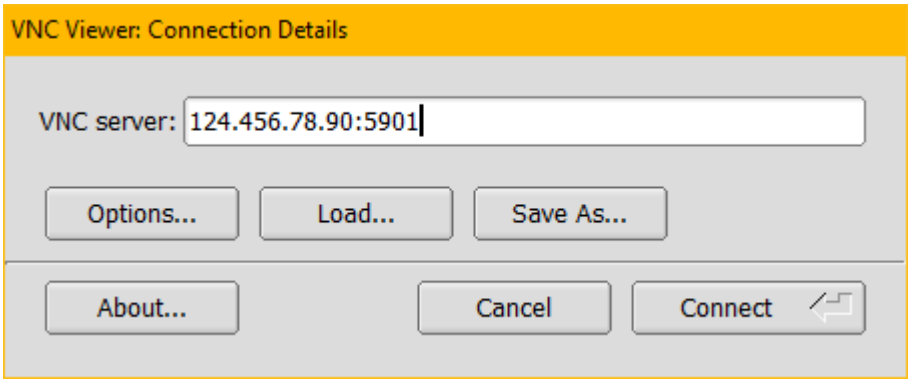

Note: If the system displays the following pop-up message, click the **Cancel** button to proceed.

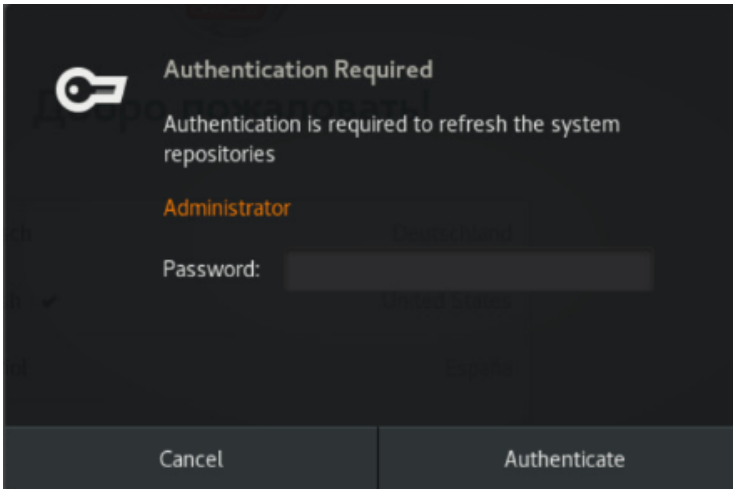

2. Enter the IP address and port of your OpenTofu staging server.
3. Click the **Connect** button.

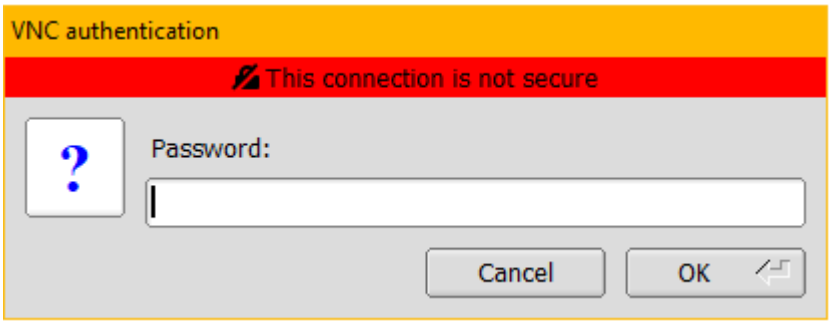

4. On the VNC authentication dialog box, enter the VNC server password that you created in the section "Performing First-Time Configuration of the OpenTofu Staging Server".

> **Note:** After you are connected to the console, the first time you connect to the VNC, you will see GNOME prompts that you must answer before using the servers. These prompts are described in steps 5 through 10 below.

5. In the Welcome section, select the preferred language and click the Next button. In the following example, English is selected.

6. In the Typing section, select the keyboard layout that you want to use as the input method and click the Next button.

   In the following example, English (US) is selected.

7. In the Privacy section, enable the Location Services option and click the Next button.

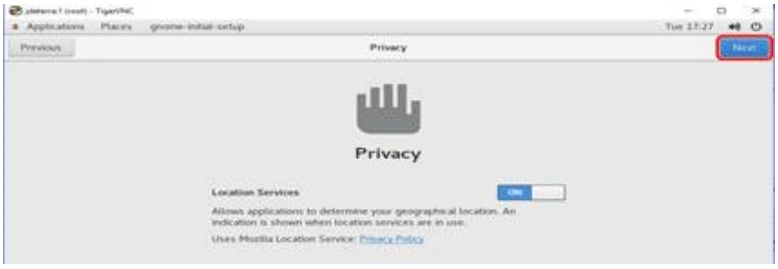

8. In the Online Accounts section, because for this use you do not need to import Online Accounts, click the Skip button.

9. In the Ready to Go section, click the Start Using Oracle Linux Server button to complete the GNOME initial setup.

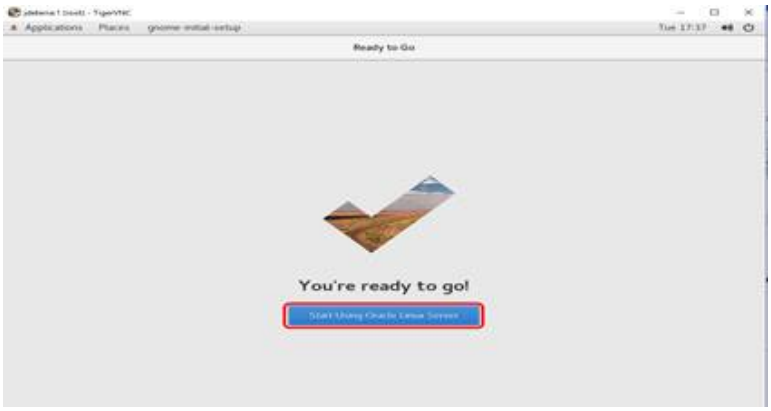

10. In the section Getting Started—GNOME Help, click the X icon to close the window.

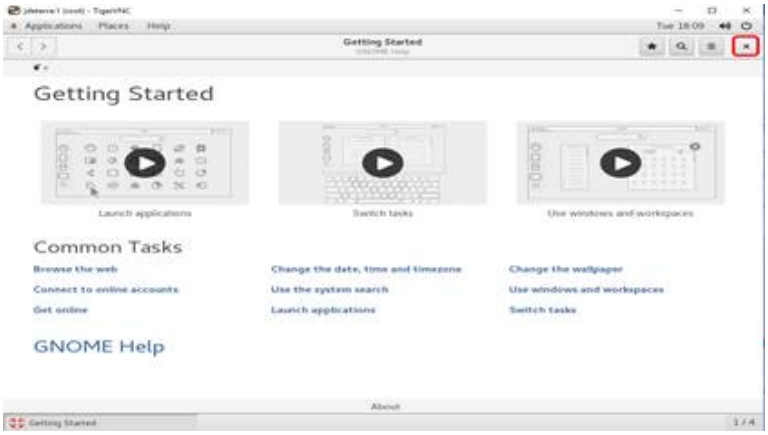

At this point, you are connected to the VNC server on the OpenTofu staging machine.

Note: The above steps 5-10 apply to the first-time connection to the VNC. For subsequent connections, you will be prompted to enter the password for the opc user before you can proceed.

Gathering Essential Information from Oracle Cloud Infrastructure for Infrastructure Provisioning

This section shows you how to gather essential Oracle Cloud Infrastructure information for the infrastructure provisioning user that will be running the JD Edwards EnterpriseOne Infrastructure Provisioning Console.

Worksheet Table - Details to Gather

This worksheet table lists the details you need to gather. For reference you can enter the details in the value column.

Detail Item	Detail Value
User OCID	
User SSH Key (PEM)	
User API Key Fingerprint	
Tenancy OCID	
Region Name	
Compartment OCID	

Prerequisite

To perform the steps in this tutorial, you must have:

- Created a Linux instance for the OpenTofu staging console by following the steps described in the previous section "Creating a Linux Instance for the OpenTofu Staging Server" in this Learning Path.
- Refer to Oracle Cloud Infrastructure documentation for additional details on: *Managing User Credentials*.

Locating the Generated Public Key for the Infrastructure Provisioning User

From the OpenTofu staging server, use this procedure to locate the generated public key (also called a PEM key) for the user that will be running the JD Edwards EnterpriseOne Infrastructure Provisioning Console. This user can be the current user or any user that you want. If the user does not exist, create a user by following the standard procedure described in Oracle Cloud Infrastructure documentation for Create a User in *Adding Users*. Also refer to Oracle Cloud Infrastructure documentation for additional details on: *Managing User Credentials*.

1. Using PuTTY, connect to the OpenTofu staging server and run this command to generate the content of the key:

   ```
   cat /home/opc/keys/ociApiKeyPublic.pem
   ```

2. Make a record of the output of the above command for use in a task in the next section "Adding the Public Key as an API Key for the Infrastructure Provisioning User".

Note: From Tools Release 9.2.4.1, the Infrastructure Provisioning Console does not prompt for any key.

Adding the Public Key as an API Key for the Infrastructure Provisioning User

Use this procedure to add the public key as an API key to Oracle Cloud Infrastructure for the user that will be running the Infrastructure Provisioning Console.

1. As the user for which you generated a public key in the preceding step, log in to Oracle Cloud Infrastructure using the following URL for the tenancy in which you want to provisioning infrastructure:

   ```
   https://console..oraclecloud.com/#/a/" format="html
   ```

   For example:

   ```
   https://console.us-ashburn-1.oraclecloud.com/#/a/
   ```

2. On the Oracle Cloud Infrastructure console, click the navigation menu in the upper-left corner.

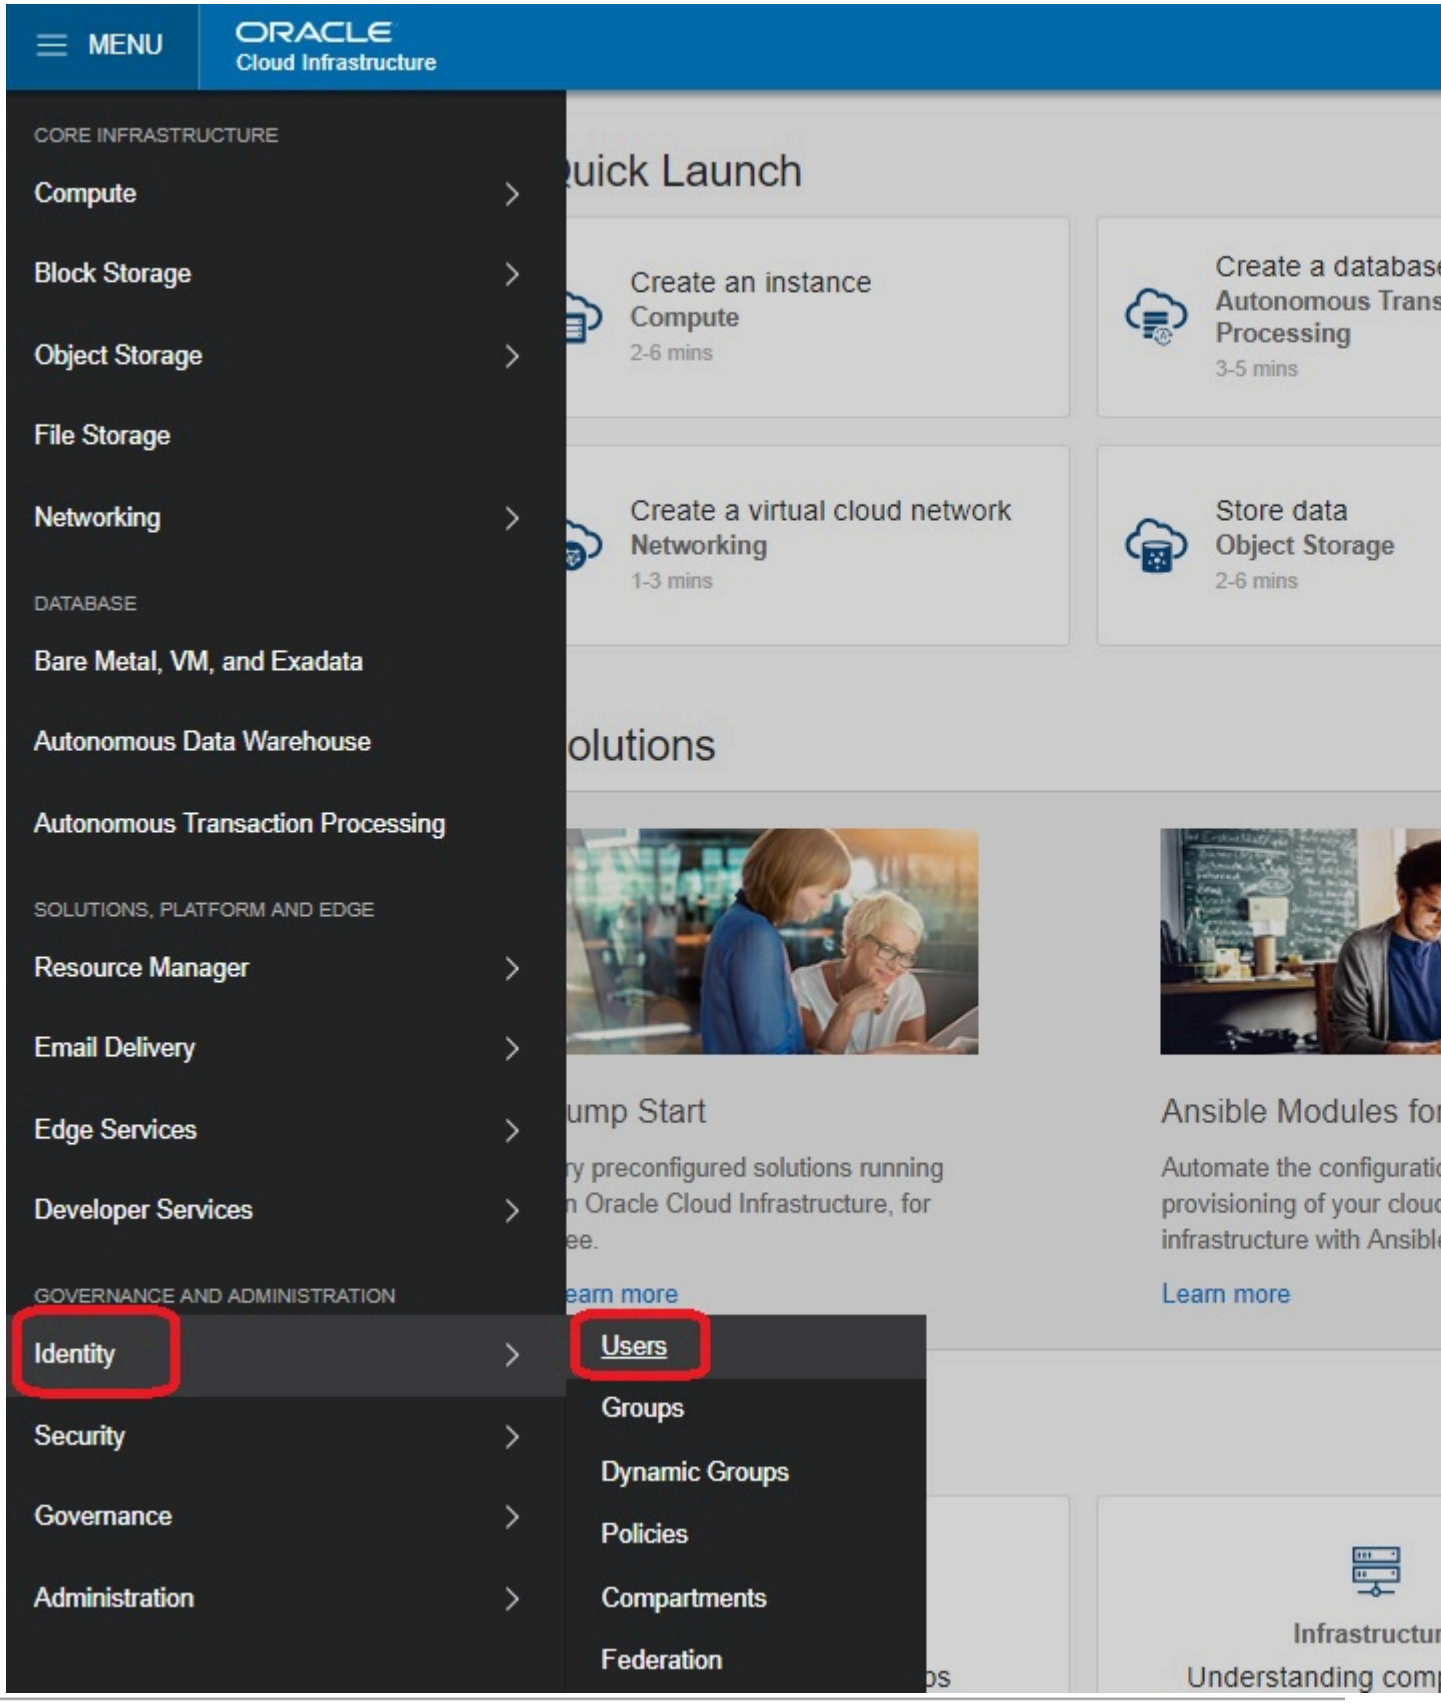

3. From the navigation menu, go to the Identity section and select Users.
4. Click the link for the user which you have designated as the infrastructure provisioning user.

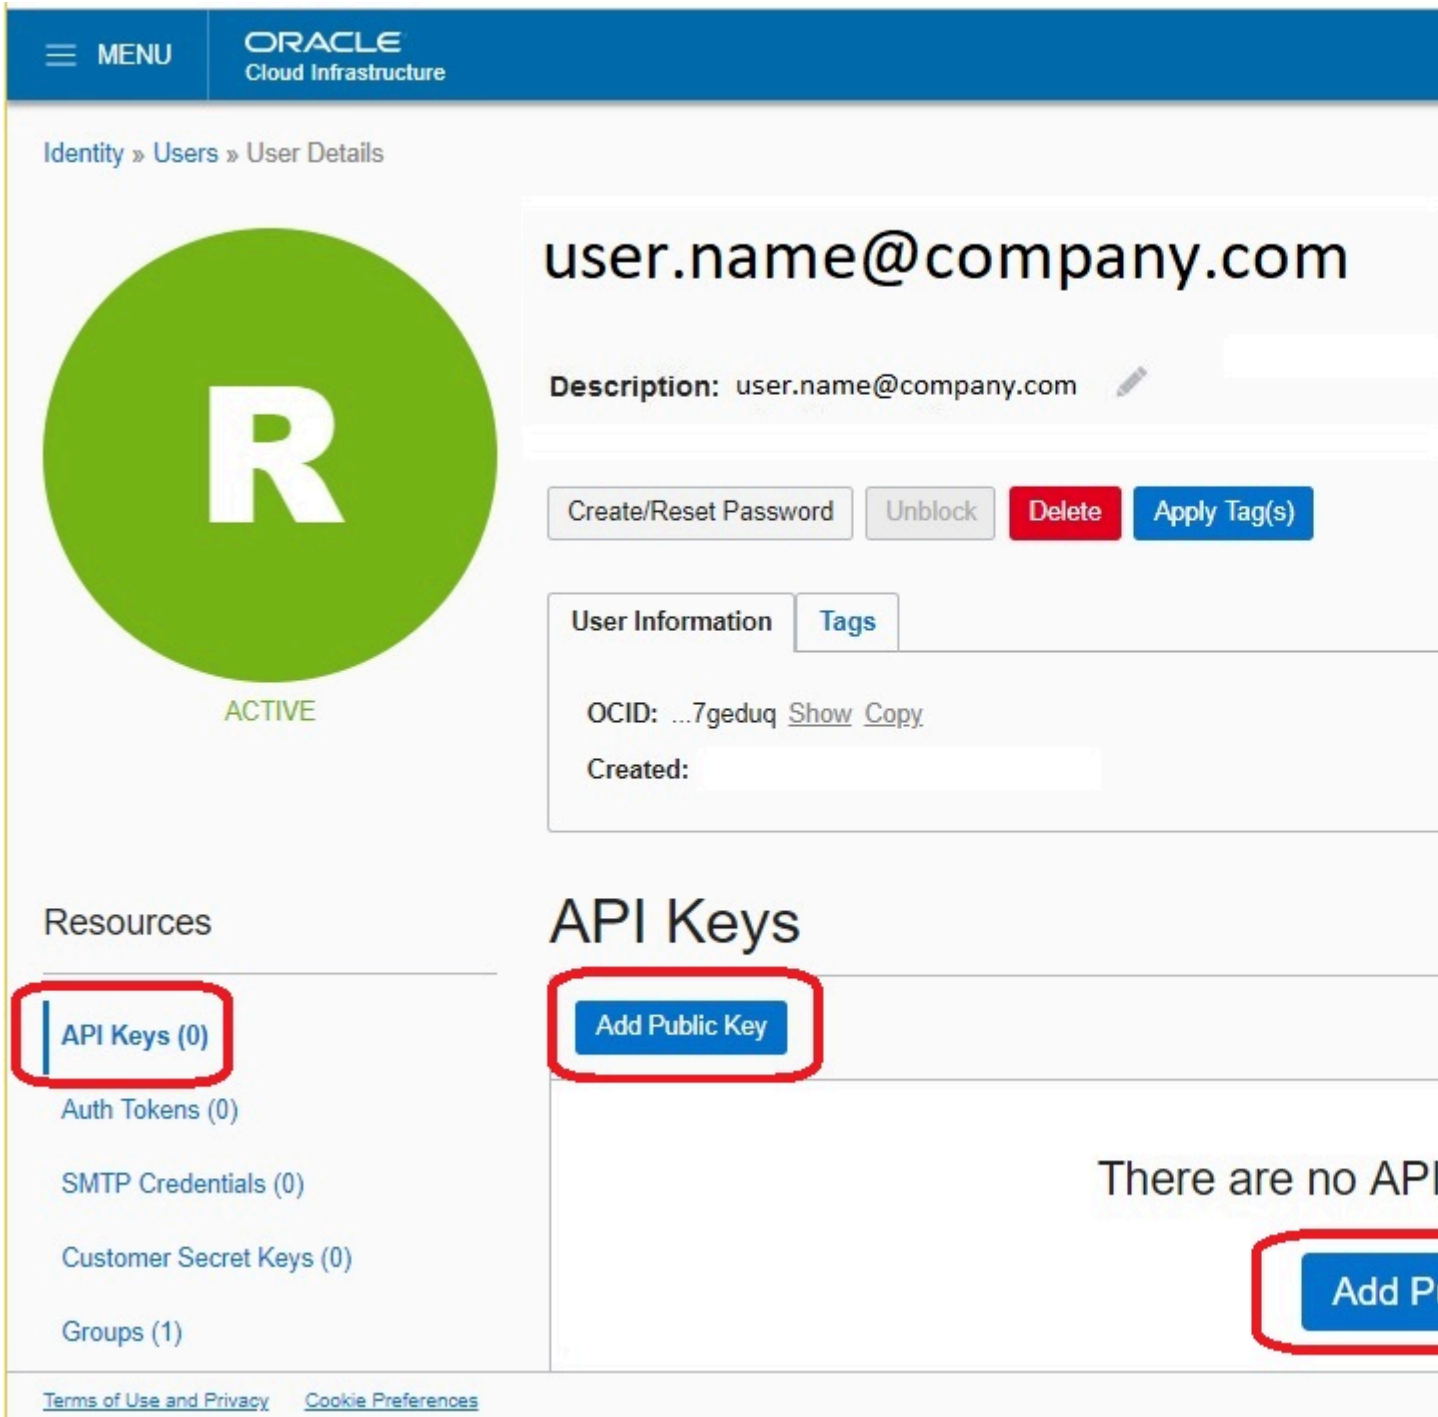

5. In the API Keys section, click the **Add Public Key** button.

6. In the PUBLIC KEY section, copy and paste the public key that you received in the previous section "Adding the Public Key as an API Key for the Infrastructure Provisioning User".

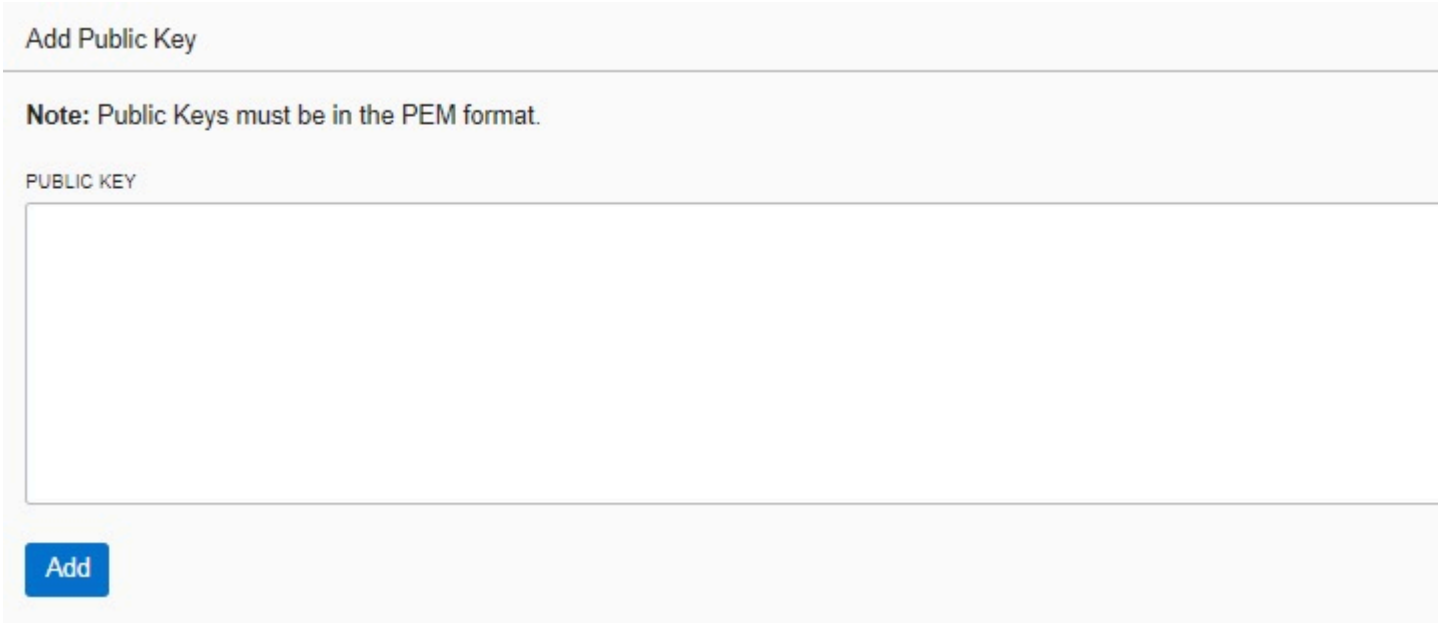

7. Click the **Add** button.

Note: Ensure that you copy the complete contents of the key including the BEGIN PUBLIC KEY and END PUBLIC KEY text.

1. Logged in as the infrastructure provisioning user, on the Oracle Cloud Infrastructure console, click the navigation menu in the upper-left corner.

Locat

After you
fingerpri

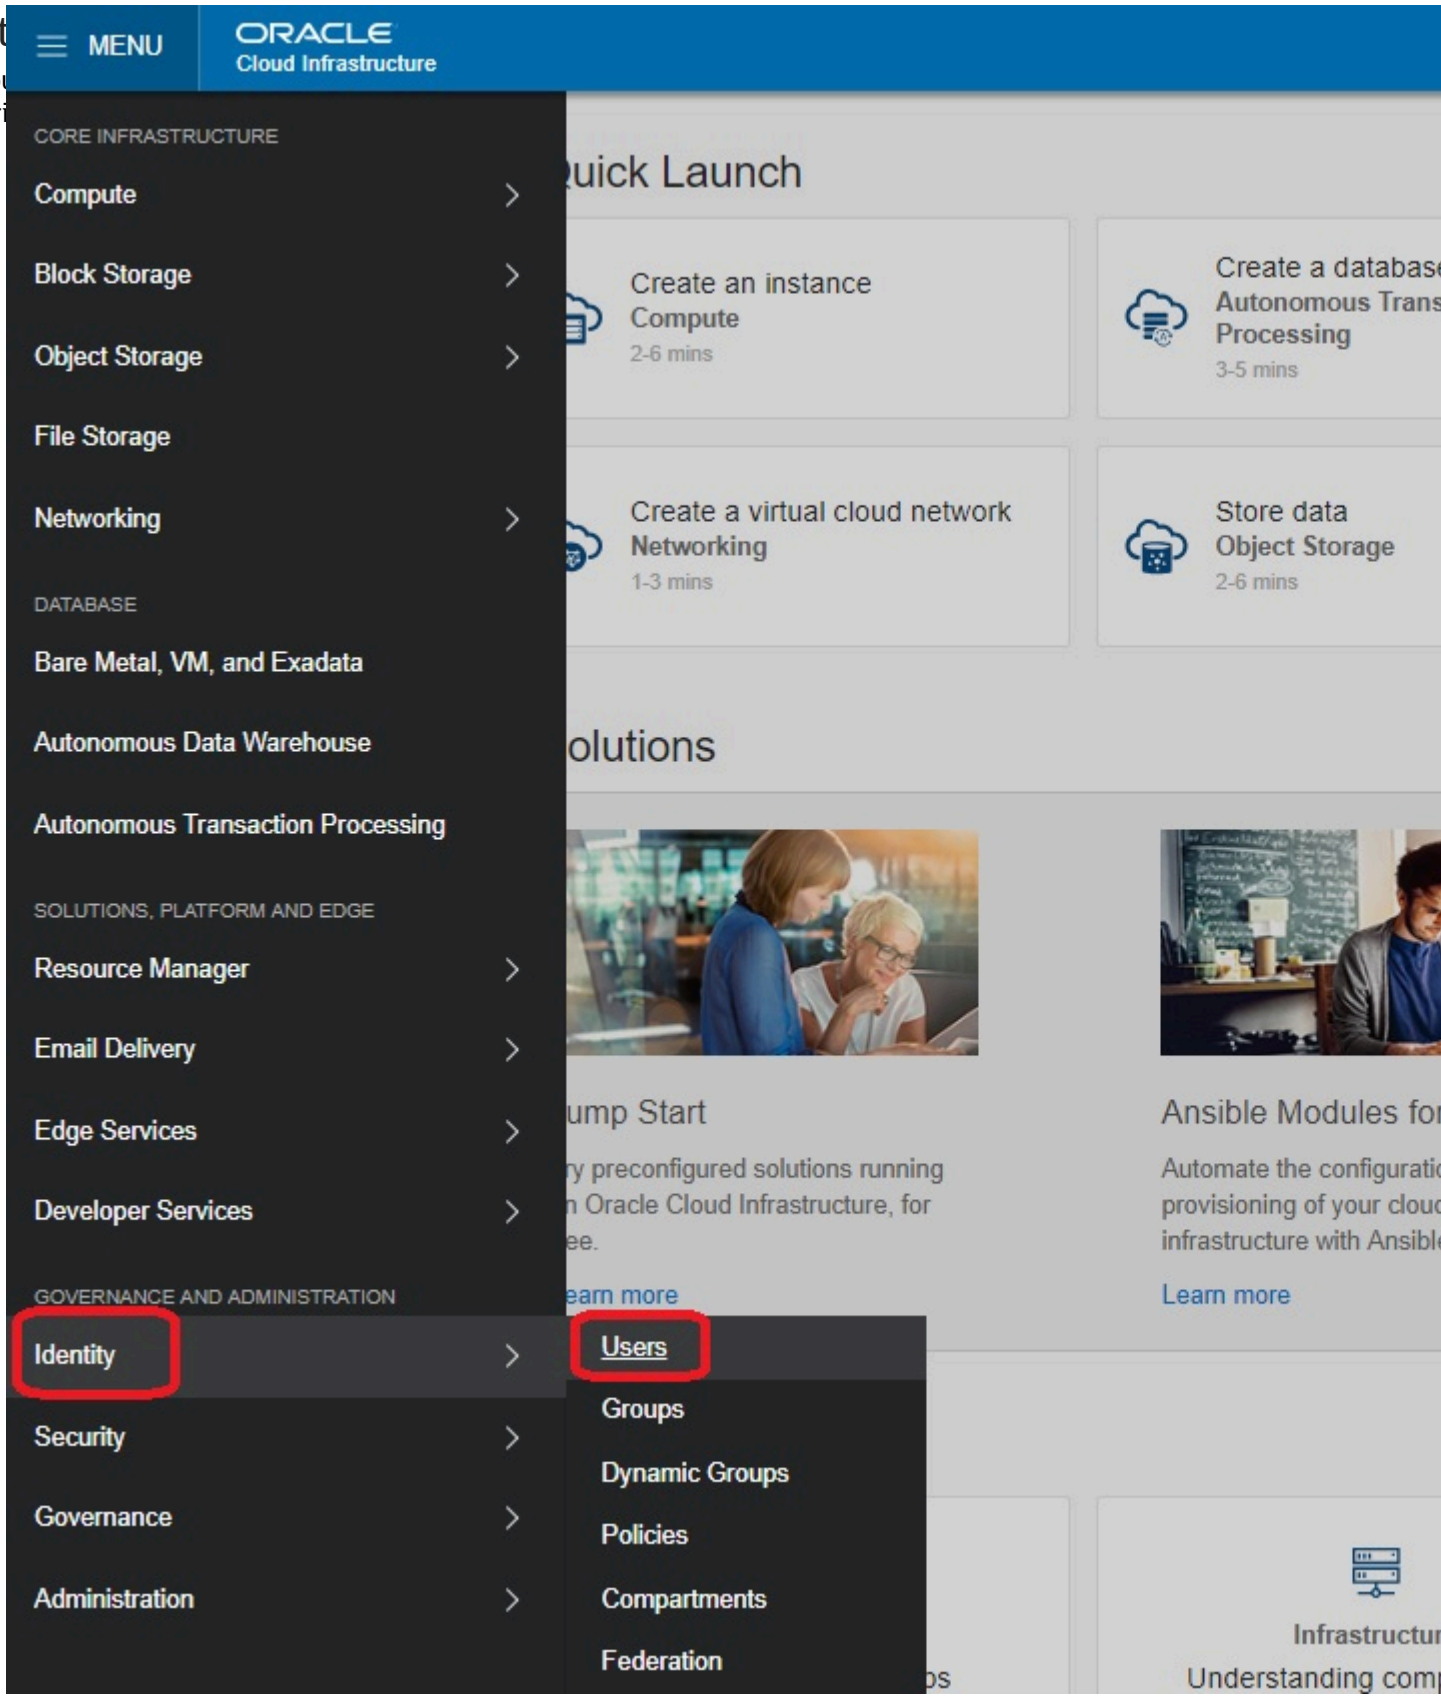

2. From the navigation menu, go to the Identity section, and select Users.
3. Click the link for the user which you have designated as the infrastructure provisioning user.
4. Under Resources, in the API Keys section, locate the Fingerprint field.

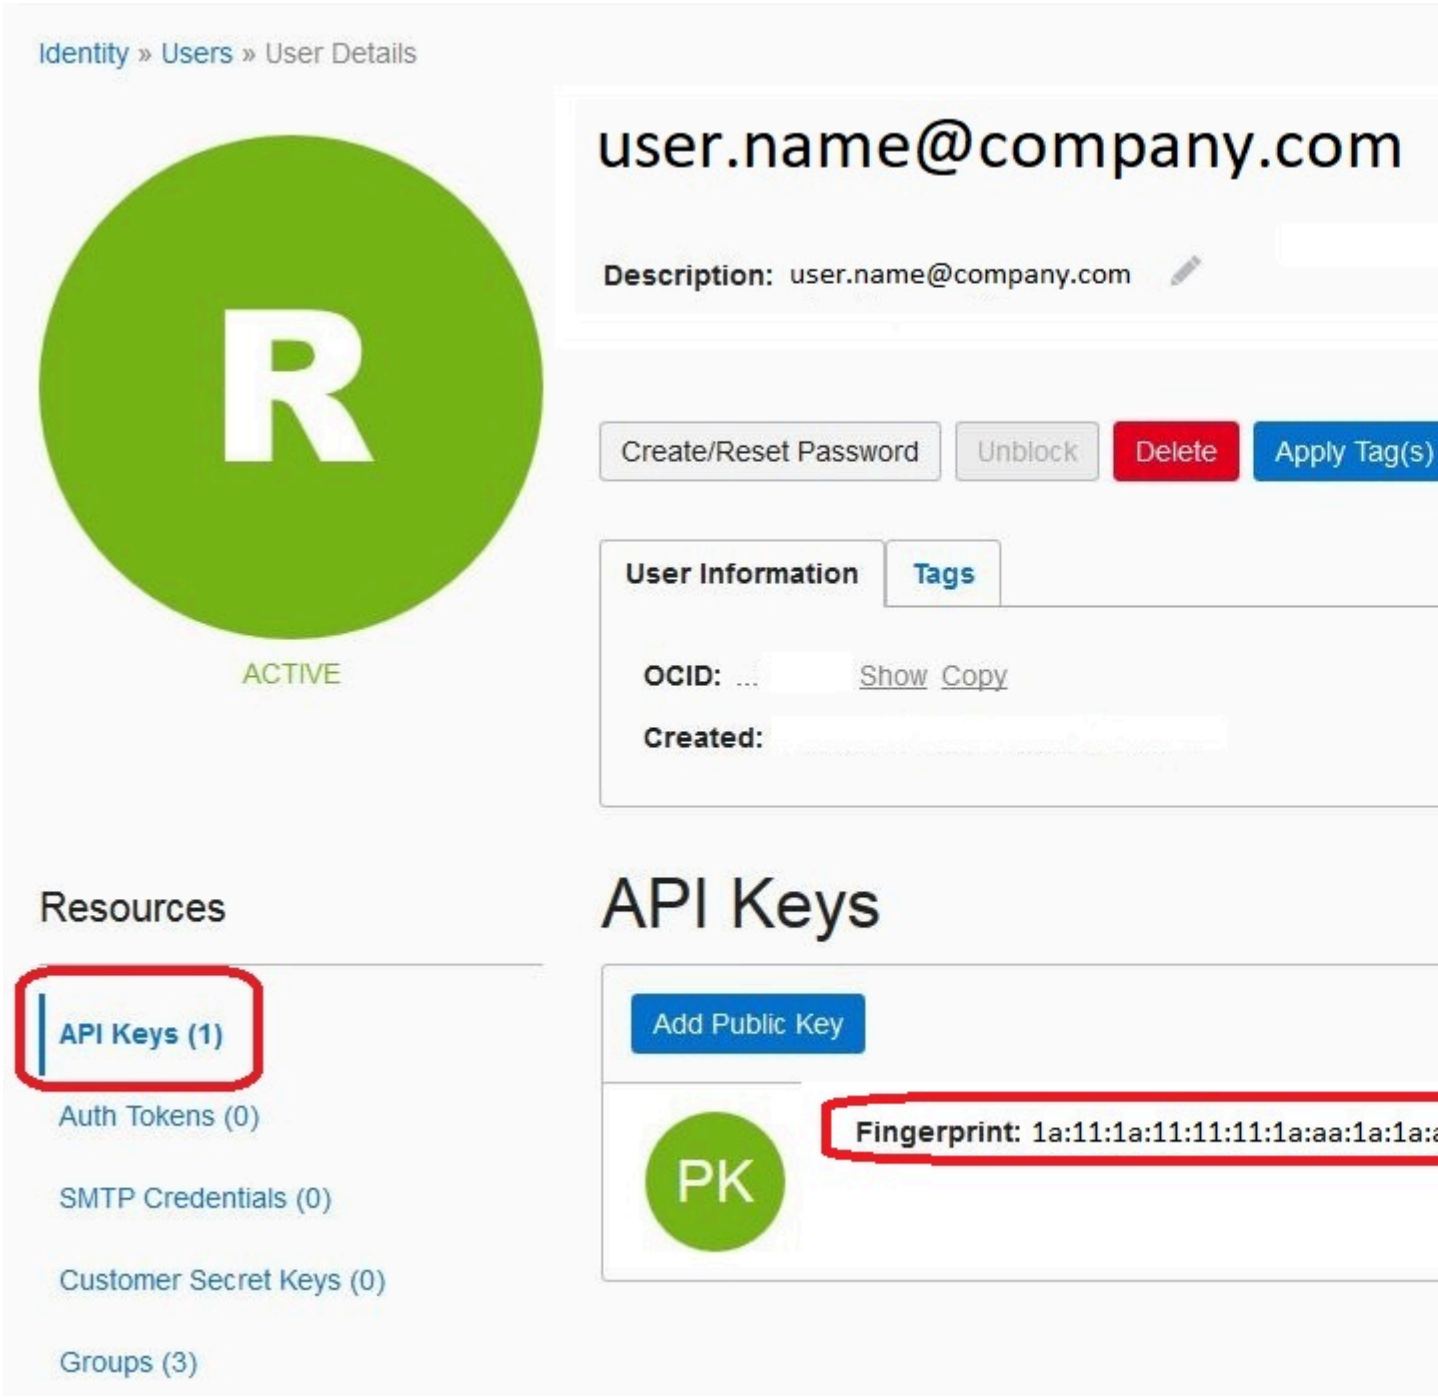

5. Make a record of the fingerprint for use as an input for a subsequent task in the section "Using the Infrastructure Provisioning Console" in this Learning Path.

1. Logged in as the infrastructure provisioning user, on the Oracle Cloud Infrastructure console, click the navigation menu in the upper-left corner.

Deter

Use this

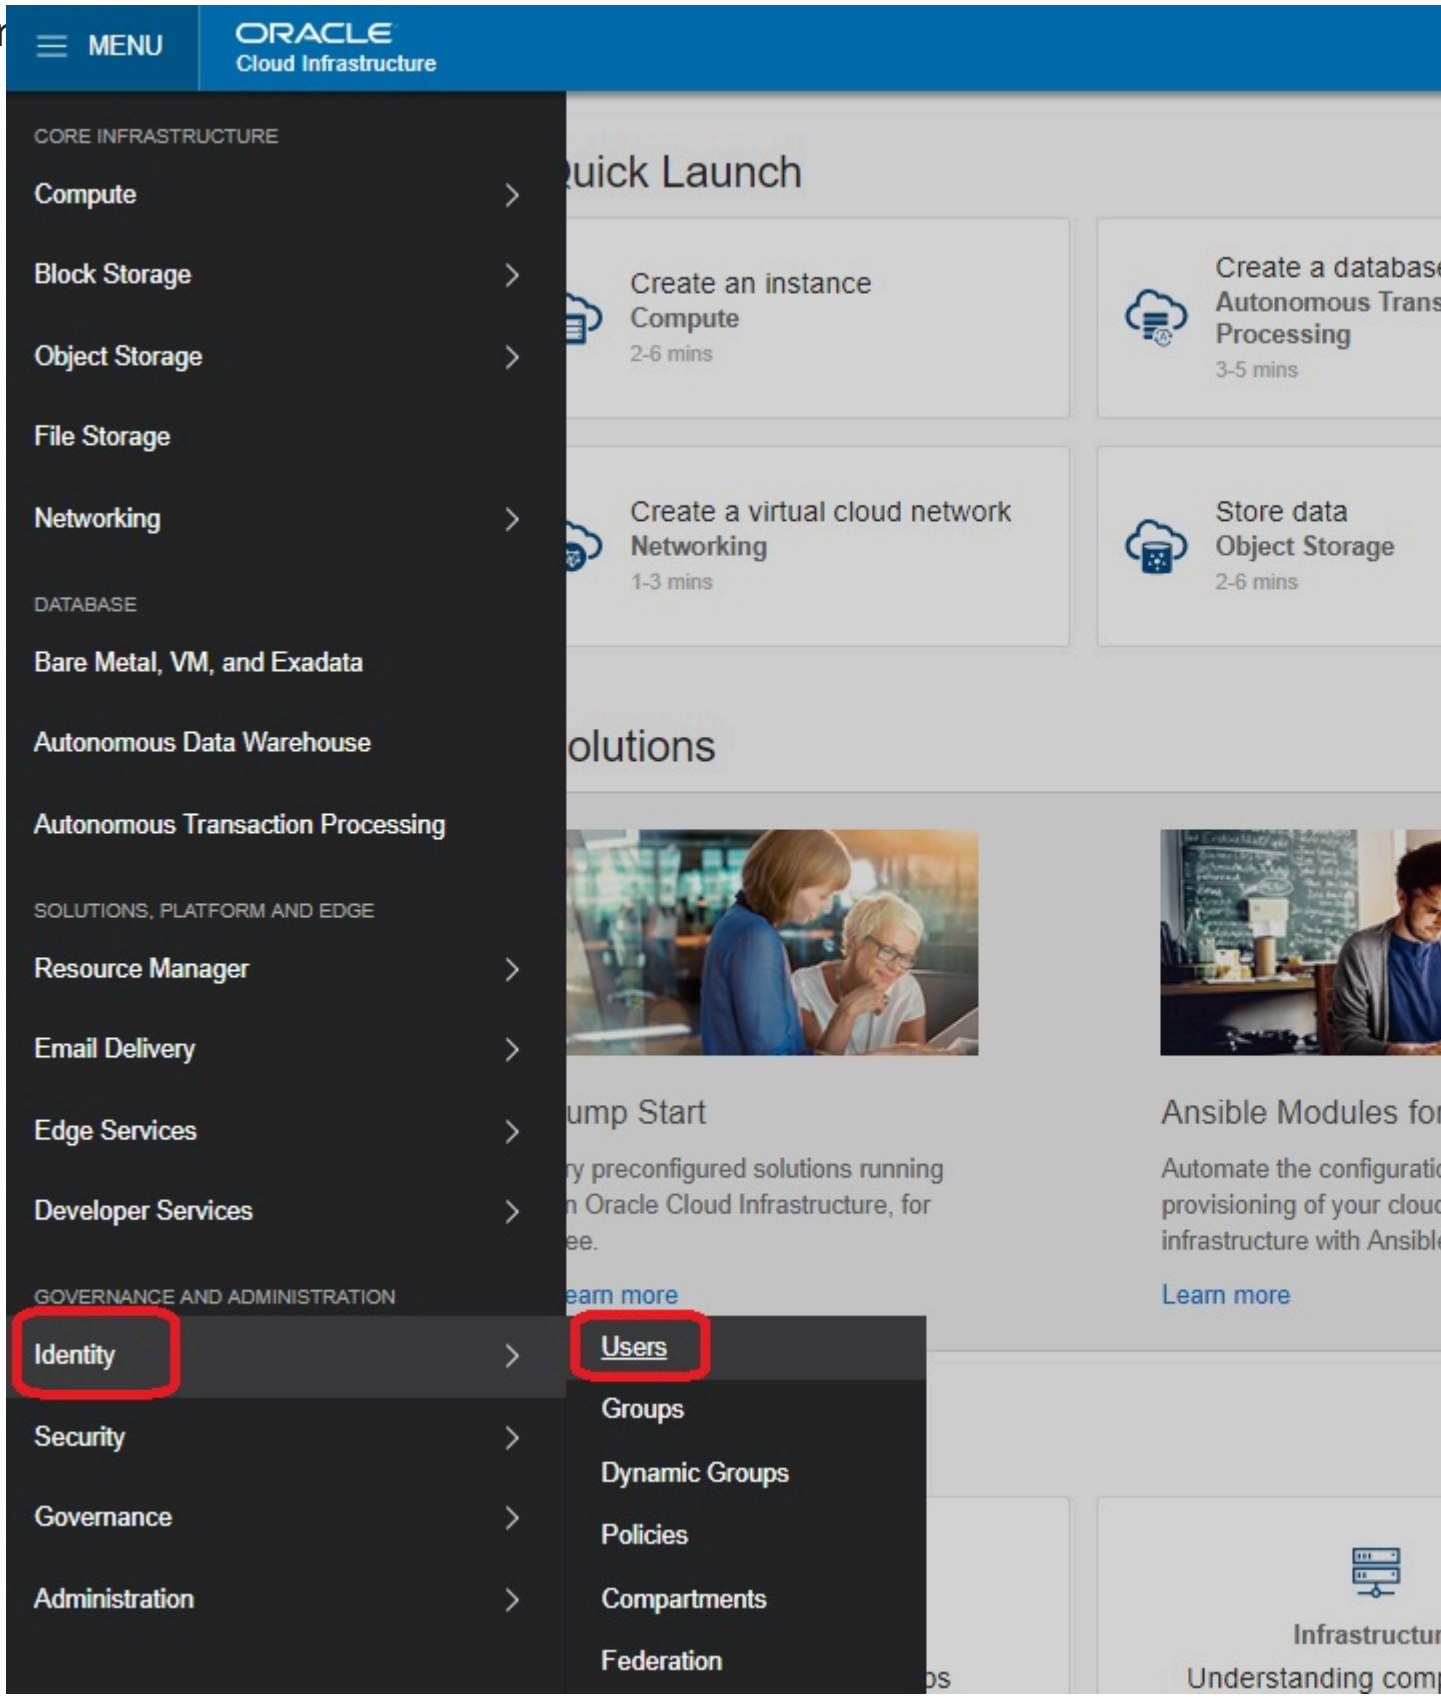

2. From the navigation menu, go to the Identity section and select Users.
3. Click the link for the user that you have designated as the infrastructure provisioning user.

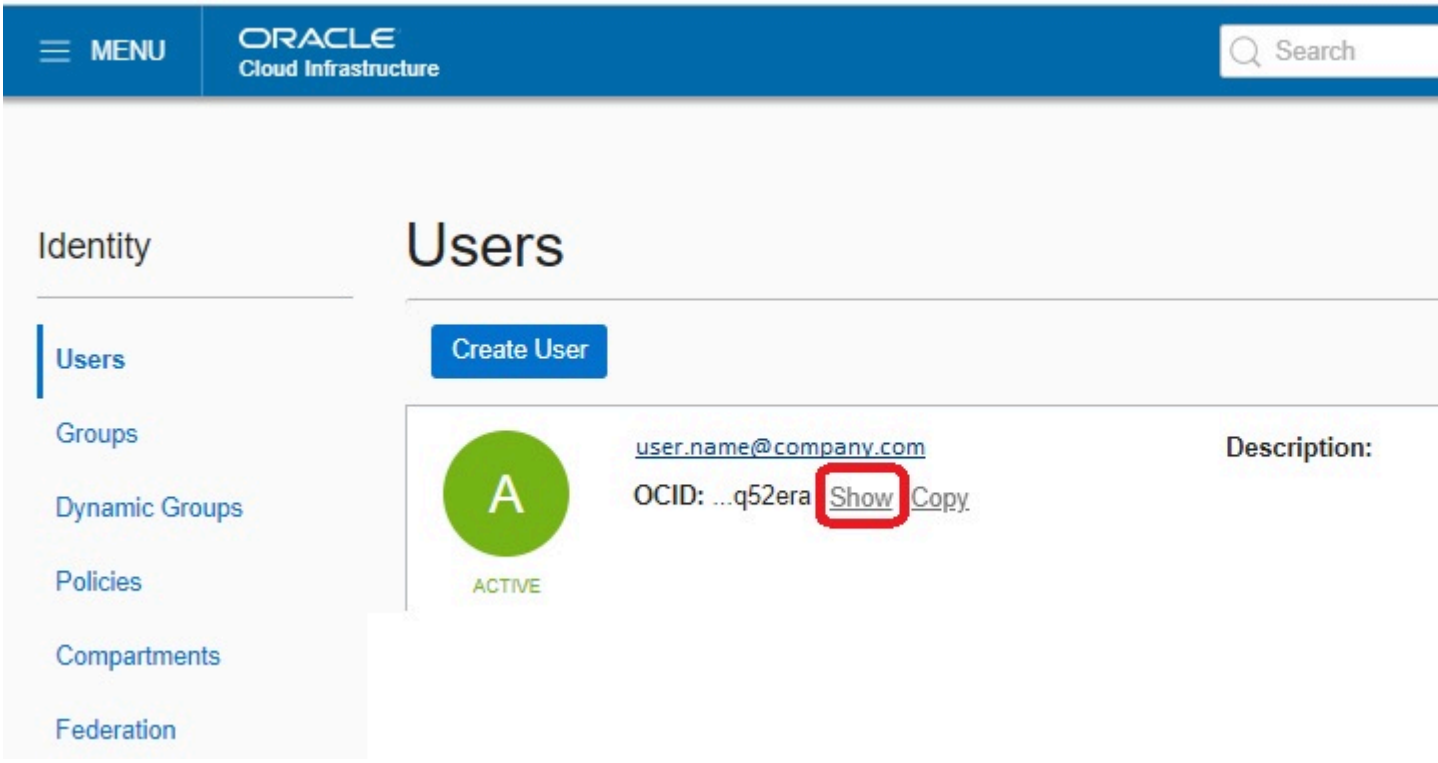

4. Click the Show link to see the OCID of the selected user.
5. Make a record of the OCID for use as an input for a subsequent task in the section "Using the Infrastructure Provisioning Console" in this Learning Path.

Locating the Tenancy OCID and Region for the Infrastructure Provisioning Compartment

Use this procedure to locate the tenancy OCID and region for the infrastructure provisioning compartment:

1. Logged in as the infrastructure provisioning user, on the Oracle Cloud Infrastructure console, click the navigation menu in the upper-left corner.

2. From the navigation menu, go to the Administration section, and select Tenancy Information.

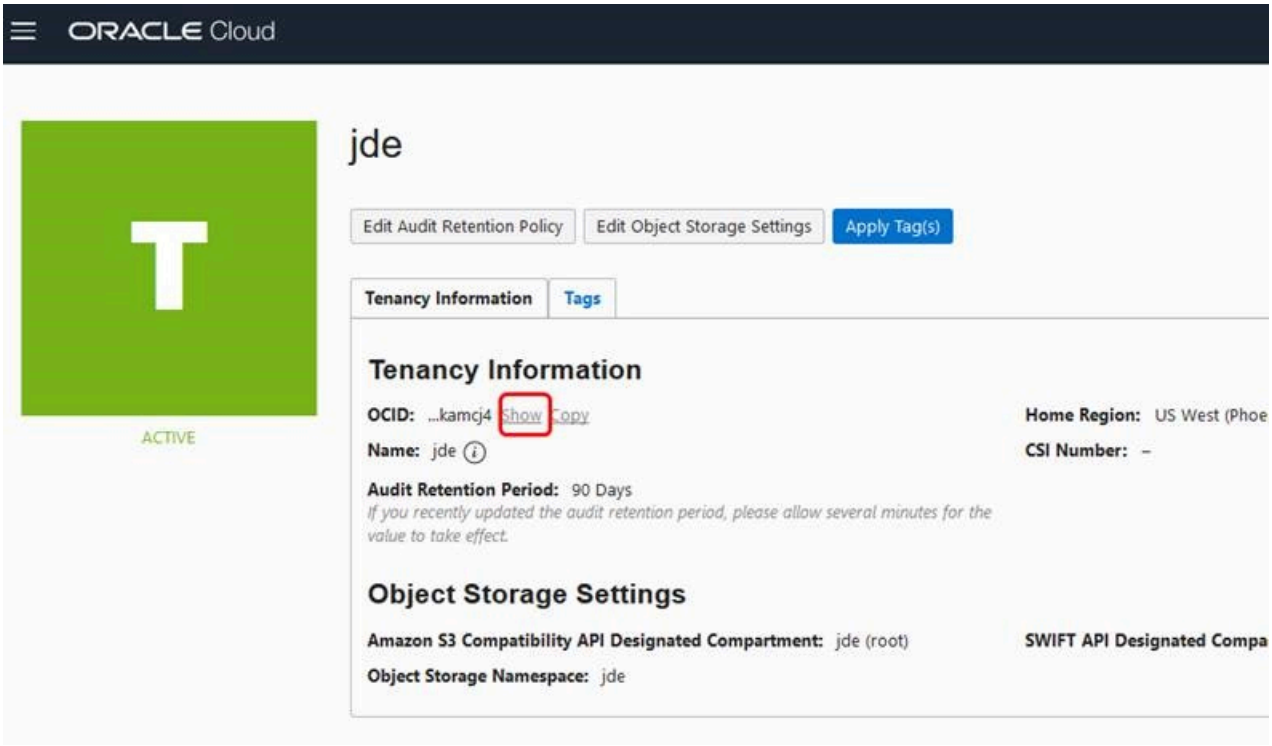

3. Click the Show link to see the OCID of the selected tenancy.
4. Make a record of the OCID for use as an input for a subsequent task in the section "Using the Infrastructure Provisioning Console" in this Learning Path.
5. In the Tenancy Information section, also locate the region that will host the instances to be provisioned. Make a note of the region for use as an input for a subsequent task in the section "Using the Infrastructure Provisioning Console" in this Learning Path.

Locating the Compartment OCID for Infrastructure Provisioning

Use this procedure to locate the compartment OCID for infrastructure provisioning:

1. Logged in as the infrastructure provisioning user, on the Oracle Cloud Infrastructure console, click the navigation menu in the upper-left corner.

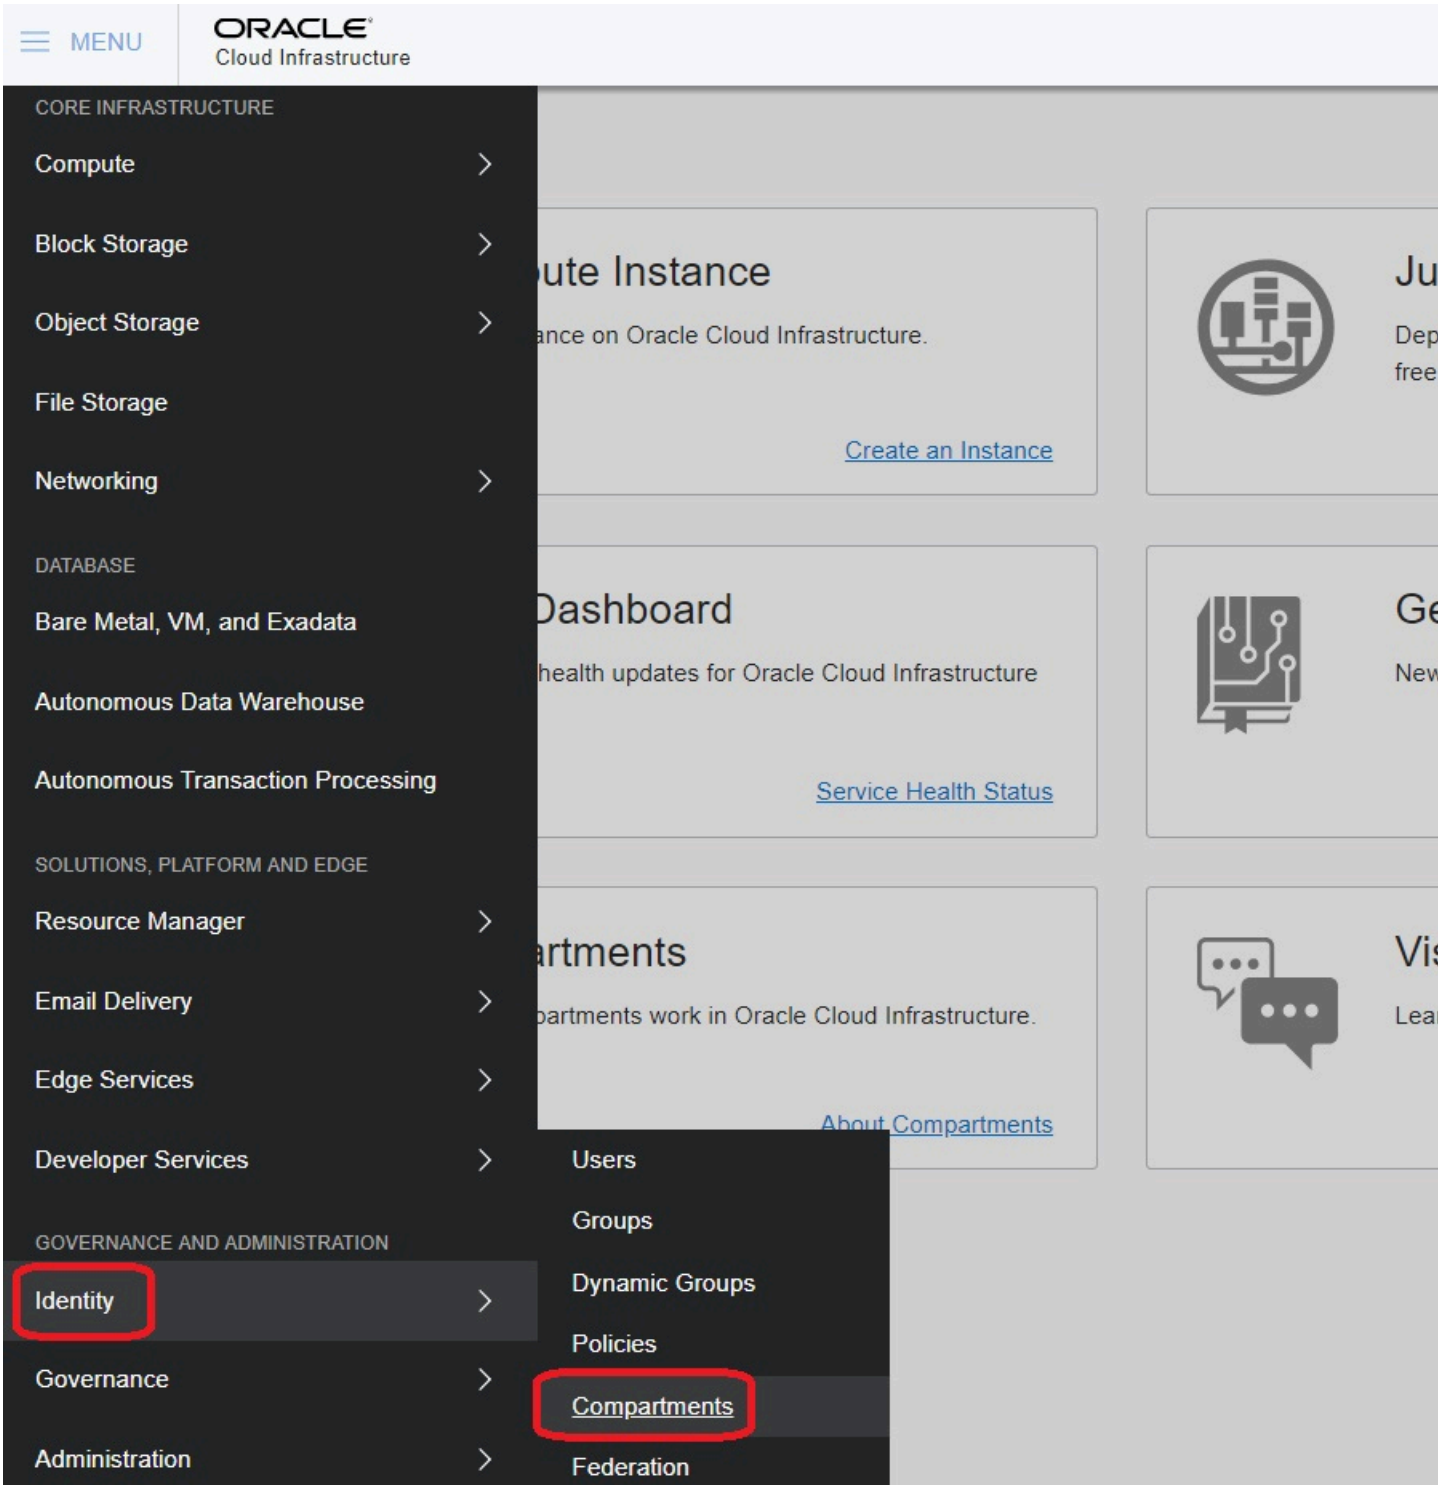

2. From the navigation menu, go to the Identity section, and select the Compartments service.

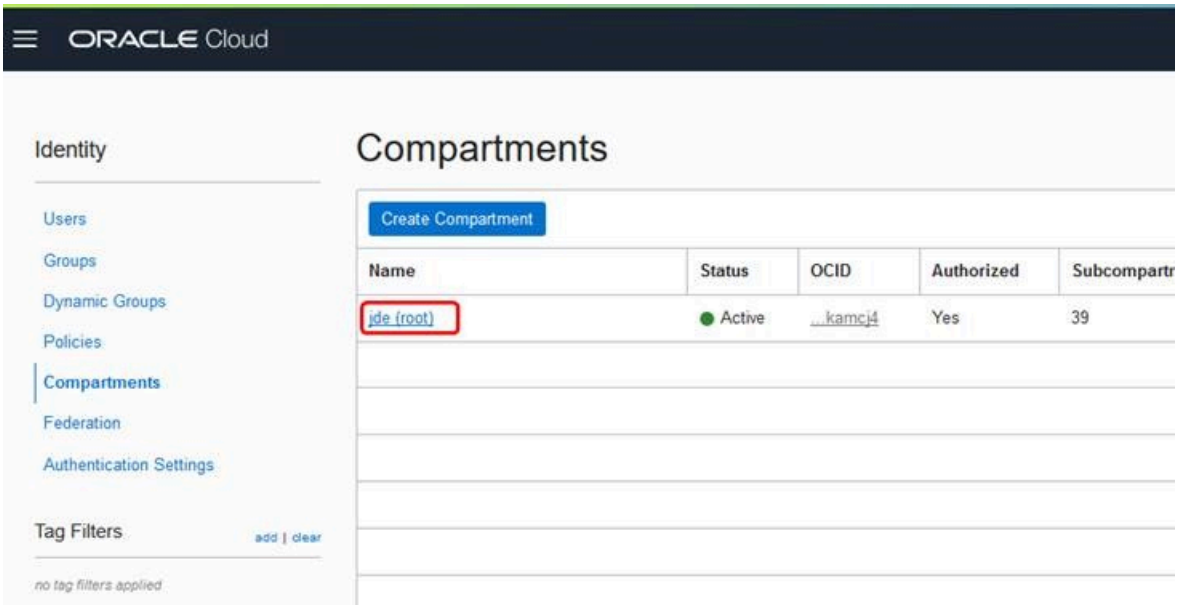

3. Click the link for the compartment that you have designated as the infrastructure provisioning compartment.

4. Click the Show link to see the OCID of the selected compartment.

5. Make a record of the OCID for use as an input for a subsequent task in the section "Using the Infrastructure Provisioning Console" in this Learning Path.

6 Using the Infrastructure Provisioning Console

Completing Configuration in the JD Edwards Infrastructure Provisioning Console

This section shows you how to use the JD Edwards Infrastructure Provisioning Console to configure details for these functional components for JD Edwards EnterpriseOne as delivered by OpenTofu for Linux instances in Oracle Cloud Infrastructure:

- Network
- Bastion Server
- Provisioning Server
- Shared Database
- Deployment Server
- Production Environment
- Non-Production Environment
- Disaster Recovery Environments

Prerequisites

- Details for the Oracle Cloud Infrastructure User, Tenancy, Region, and Compartment – see *"Gathering Essential OCI Information for Infrastructure Provisioning"* in this Learning Path.
- Logged into the JD Edwards Infrastructure Provisioning Console. For more information, see *"Logging into the JD Edwards Infrastructure Provisioning Console"* in this Learning Path.

Completing Configuration in the JD Edwards Infrastructure Provisioning Console

Use this procedure to complete configuration in the JD Edwards Infrastructure Provisioning Console on the Staging Server, which you must access through VNC viewer:

1. Access the JD Edwards Infrastructure Provisioning Console using a browser and the following URL: `https://localhost:3000`

   > **Note:** Because you are using localhost, the system prompts you to accept an insecure connection.

2. On Insecure Connection, Your connection is not secure, click the **Advanced** button and then click the **Add Exception** button.

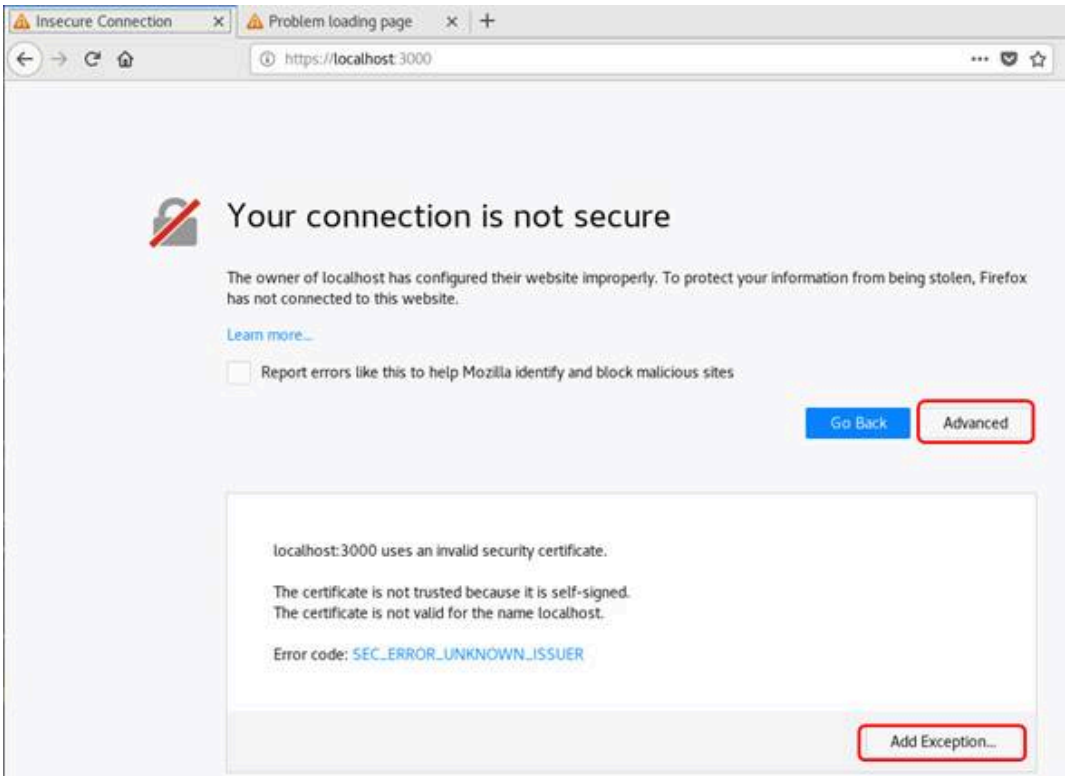

3. On Add Security Exception, click the **Confirm Security Exception** button.

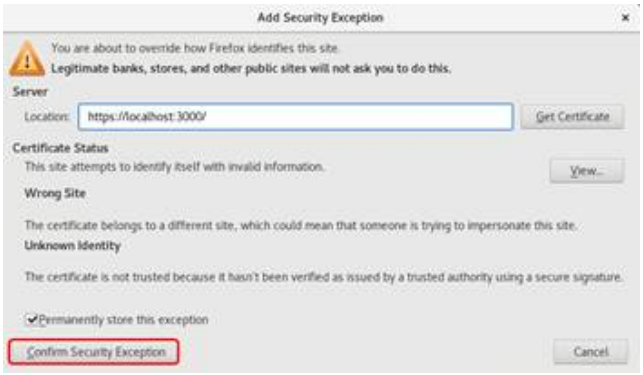

Note: After you add the Security Exception, the Infrastructure Provisioning Console is displayed.

4. On Network Information, enter valid values as illustrated in the following example:

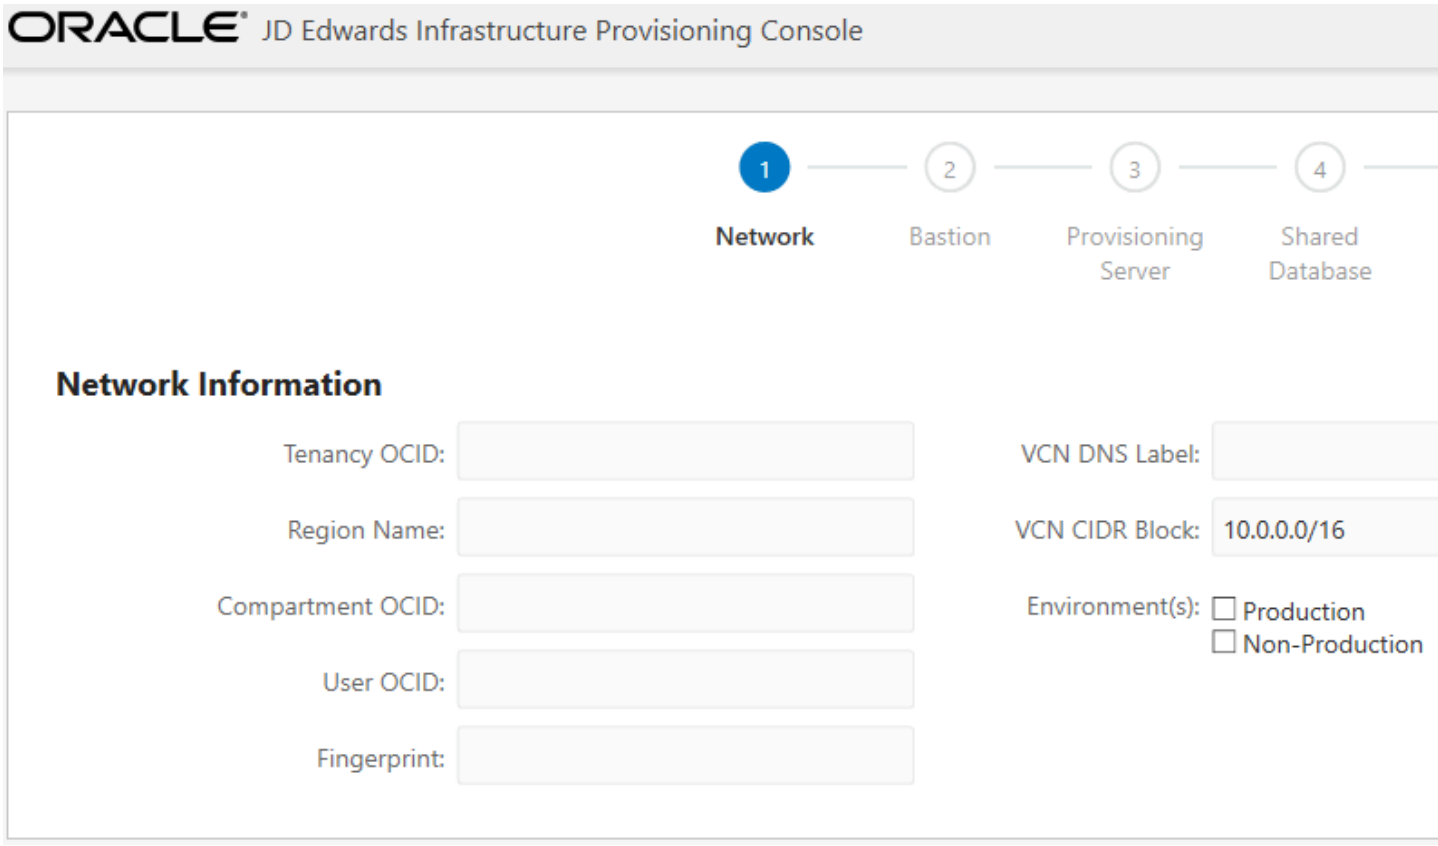

- ***Tenancy OCID***

  This is the Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). It is included as part of the resource's information in both the Console and the API. You can find your tenancy's OCID in the Console

at the top right user icon. The tenancy OCID looks something like this (notice the word "tenancy" in it):

```
ocid1.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdsq
```

- ***Region Name***

  Enter the Region Identifier of the Region Name where the resource is located. In the user interface this value is labeled as Region Identifier as shown below. For example, for the Region Name US West Phoenix, the Region Identifier that you would enter here would is us-phoenix-1.

- ***Compartment OCID***

  Enter the OCID of the Compartment. On OCI, access the Identity tab, and then click Compartments to view the Compartment OCID.

- ***User OCID***

  Enter the user OCID. This is the OCID in the User Information tab of the api.user window. See the section Oracle Cloud Identifiers (OCIDs) in this document.

- ***Fingerprint***

  Enter the Fingerprint of the Infrastructure Provisioning user. See "Gathering Essential OCI Information for Infrastructure Provisioning" in this Learning Path.

- ***VCN DNS Label***

  Enter the value of the **DNS Label** for your VCN.

  > **Note:** Length of the field cannot exceed 15 characters. The name must start with a letter and must contain only alphanumeric characters. The value should only specify the globally unique domain name; do not include the prefix "www" or "local".

  > **Note:** You cannot specify duplicate VCN names. You should not use the same name that was given when creating the VCN previously described in this Learning Path in the task entitled "Performing Setup Tasks in Oracle Cloud Infrastructure: Creating a Virtual Cloud Network".

- ***VCN CIDR Block:***

  Enter the VCN CIDR address. The VCN CIDR blocks indicates the network addresses that can be allocated to the resources.

  > **Note:** Ensure that the CIDR block represents the correct private IP Address range.

- ***Environment for this Plan:***

  Select the environment for this plan. The available options are: Production with Disaster Recovery and Non-Production.

Note: Only the following scenarios are supported: **Supported Scenarios:**

- Production and Non-Production together in the same run
- Production alone, not followed by any other run
- Production on one run, followed by a second Non-Production run
- Non-Production alone, not followed by any other run
- Non-Production run, followed by a second Production run
- Production with Disaster Recovery together in the same run

Note: Non-Supported Scenarios:

- Production in Single Run and Disaster Recovery in second Run

5. Click **Next**.

   Note: If a blank screen displays, it is because there might be a slight delay until the fields on this form are populated.

6. On Bastion Information, once the screen is populated, enter valid values as illustrated in the following example:
   - *Host Name Prefix*

     Enter the Bastion host name prefix.

     The Bastion instance is created using the prefix value entered in this field. The length of the prefix must not exceed 6 characters.

     For example, if the prefix value is entered as `pdjde`, the server names will be created as follows:

     Bastion Host: `pdjdebas` [first three character of domain][1-N]

     where **[xxx]** is the first three character of domain, and

     where **[N]** is the number of servers, which can be 1-N.

     Note: The value for this label cannot be the same as you have given on Bastion Details page.

   - *Shape*

     Select the Bastion instance shape. For use with Infrastructure Provisioning, the minimum recommended shape for the instance for the Bastion Server is any x86-compatible shape series that supports 1 OCPU and 30 GB of memory.
   - *Number of OCPUs*

     Specify the number of OCPUs. For this Bastion server the recommended value is 1.
   - *Amount of Memory (GB)*

     For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory allowed is either 1 GB or a value matching the number of OCPUs, whichever is greater.

- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.

- *Burst Size*

  For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.

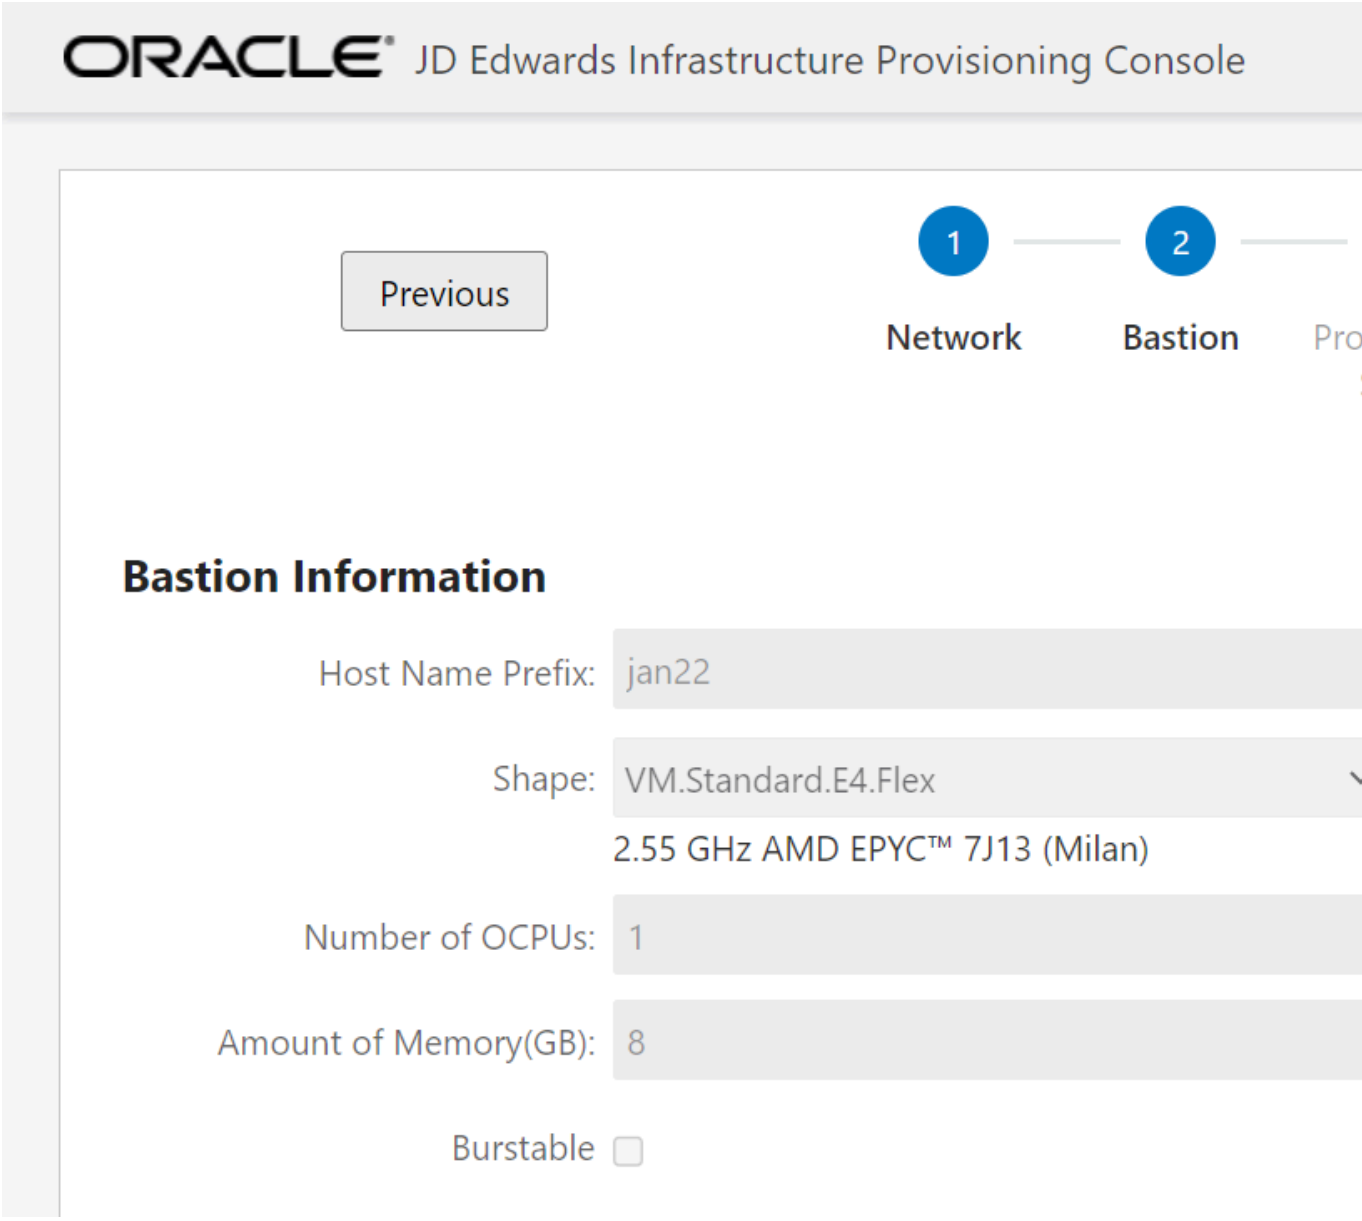

7. Click **Next**.

8. On Provisioning Server Information, enter valid values as illustrated in the following example:
   - *Shape*

     The minimum recommended shape is any x86-compatible shape series that supports 2 OCPUs and 32 GB of memory.
   - *Number of OCPUs*

     Specify the number of OCPUs. The minimum number of OCPUs recommended for the Provisioning Server is 2.
   - *Amount of Memory (GB)*

     For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for Provisioning Server is 32 GB.
   - *Burstable*

     For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

     If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
   - *Burst Size*

     For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.
   - *Block Volume (GB)*

     Select the block storage size. It is recommended to use at least 80 GB of storage size.

     > **Note:** Additional block storage size is attached to the /u01 mount point.

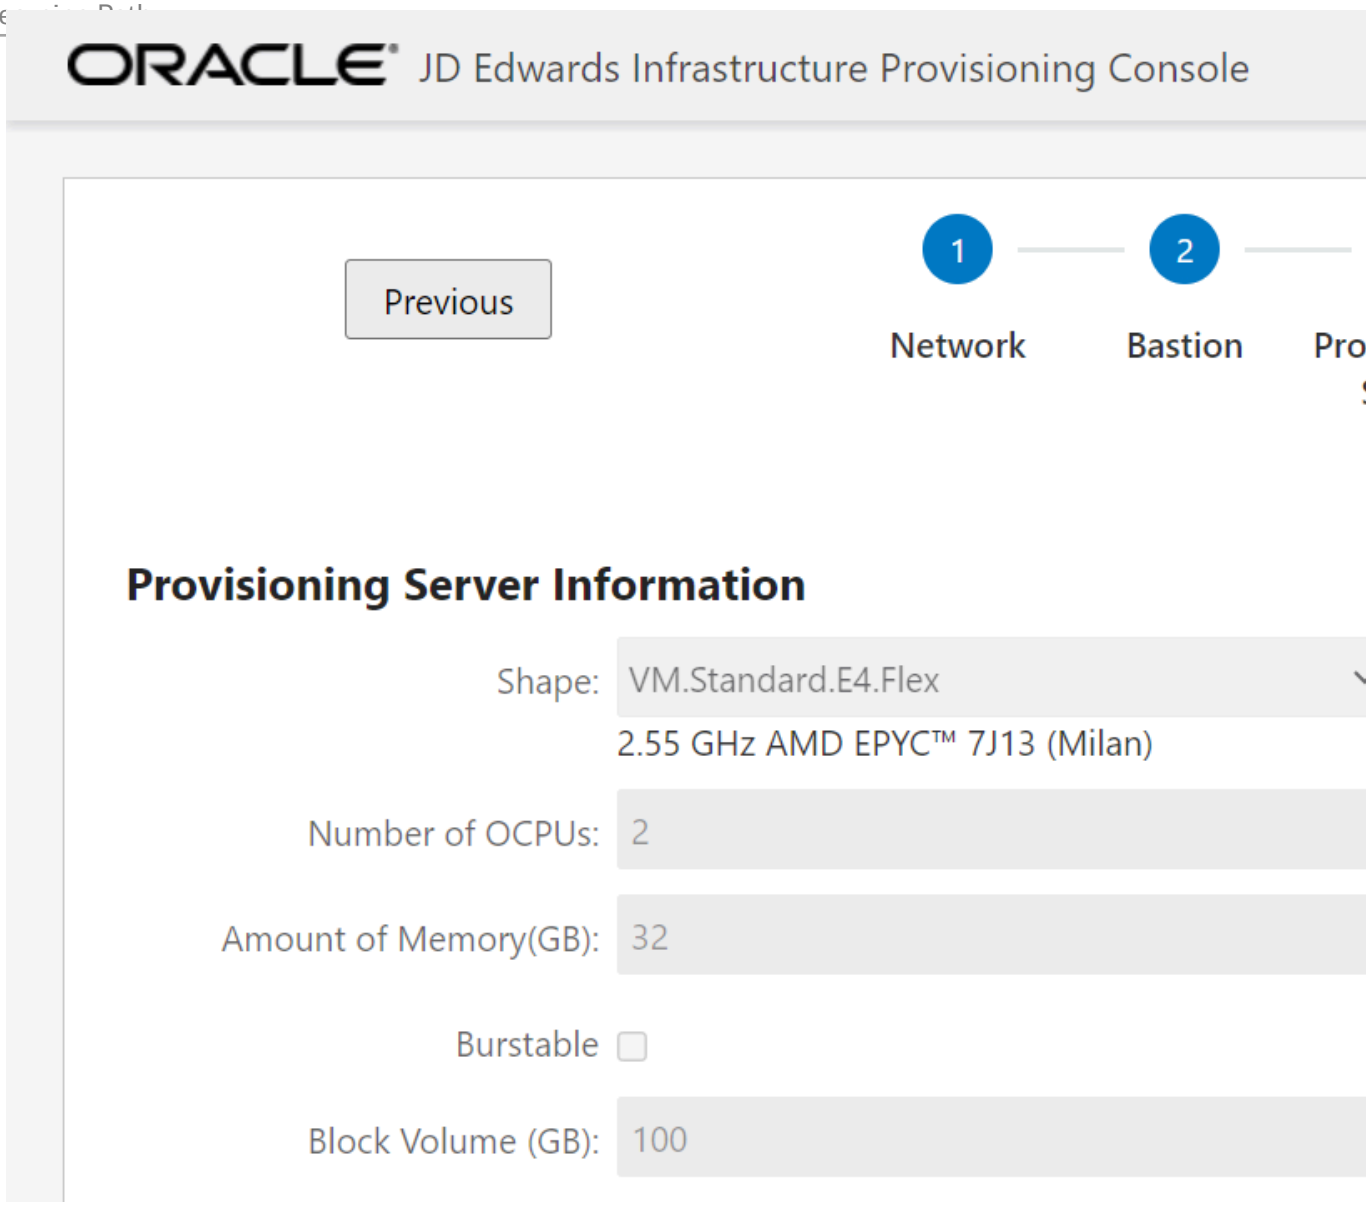
ORACLE JD Edwards Infrastructure Provisioning Console
Previous
1
2
Network
Bastion
Pro
Provisioning Server Information
Shape:
VM.Standard.E4.Flex
2.55 GHz AMD EPYC™ 7J13 (Milan)
Number of OCPUs:
2
Amount of Memory(GB):
32
Burstable
Block Volume (GB):
100

9. Click **Next**.

10. On Shared Database Information, select valid values as illustrated in the following example:
    - *Create Separate Database Server for Shared DB*

      Select this option to create an additional Database Server for the Shared DB. The number of additional Database Server will be same as the number of Database Servers entered for the selected pathcode.
    - *Create Shared Database along with Pathcode*

      If you select this option, an additional Database Server will not be created for the shared DB.

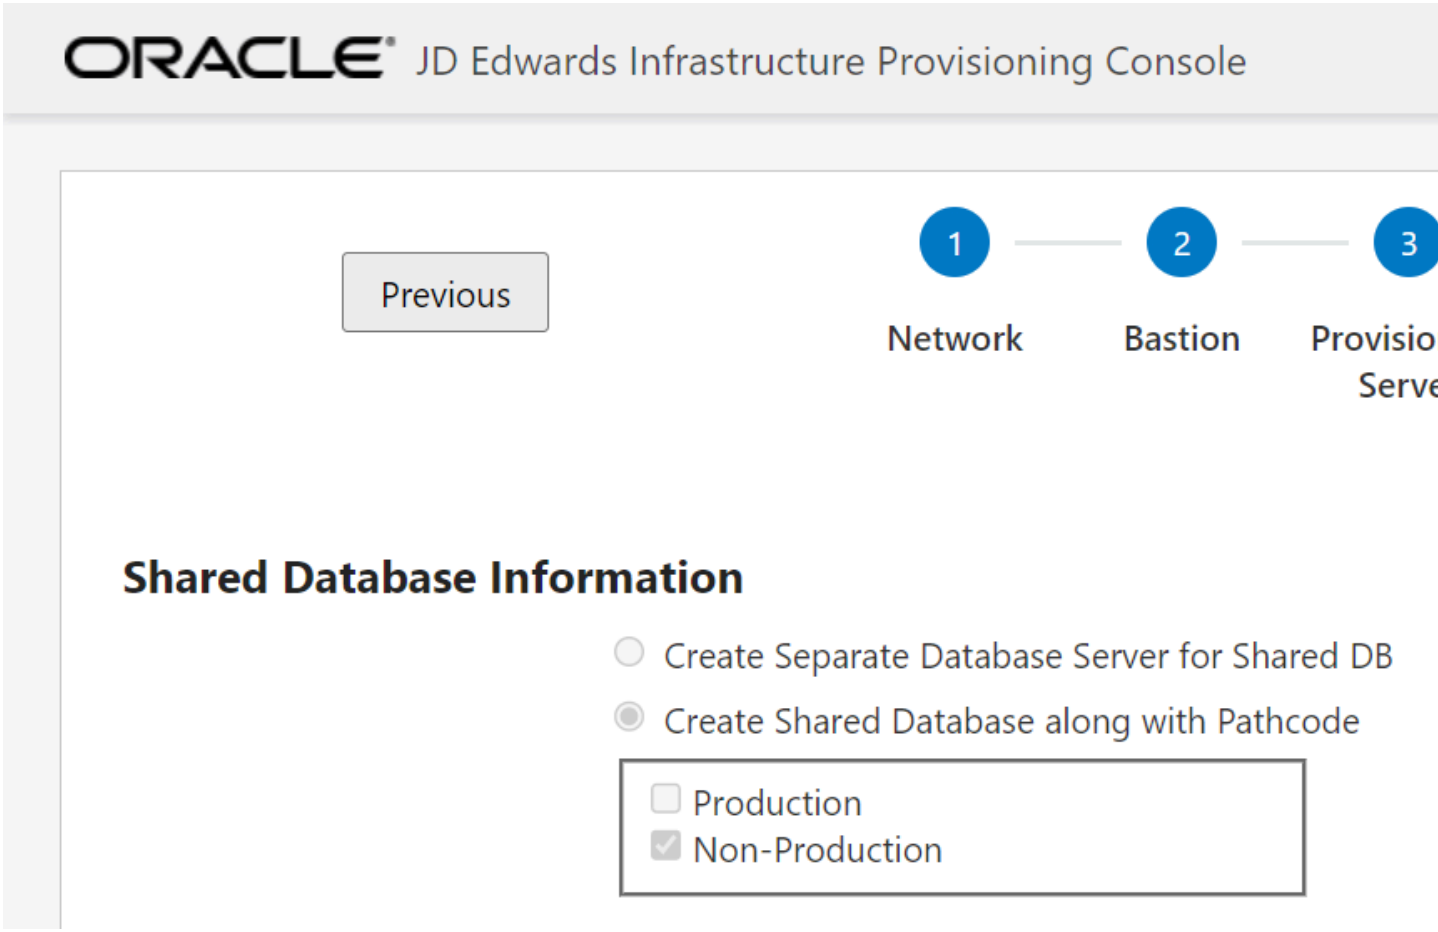

11. Click **Next**.

12. On Deployment Server Information, enter valid values as illustrated in the following example:
    - *Create Deployment Server*

      Choose from the options as required. The available options are: Yes and No.
    - *Deployment Server Instance Shape List*

      The recommended shape is any 86x-compatible shape series that supports 2 OCPUs and 32 GB of memory.
    - *Number of OCPUs*

      Specify the number of OCPUs. The minimum number of OCPUs recommended for Deployment Server is 2.
    - *Amount of Memory (GB)*

      For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the Deployment Server is 32 GB.
    - *Burstable*

      For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

      If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
    - *Burst Size*

      For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.
    - *Deployment Server Instance Password*

      Enter the Password for Deployment Server Instance.

      The password value must contain a minimum of twelve (12) characters which must contain any three (3) of these characteristics:

      i. Upper case letters
      ii. Lower case letters
      iii. Numeric digits
      iv. Special characters _ @ ~ ! # % * +) (} {] [. ?

      > **Note:** You should make a note of this password because this same password will be required in order to log into this Deployment Server instance.

    - *Size (DB)*

      Enter the size for the disk space for the Deployment Server. The minimum recommended value is 250 GB.

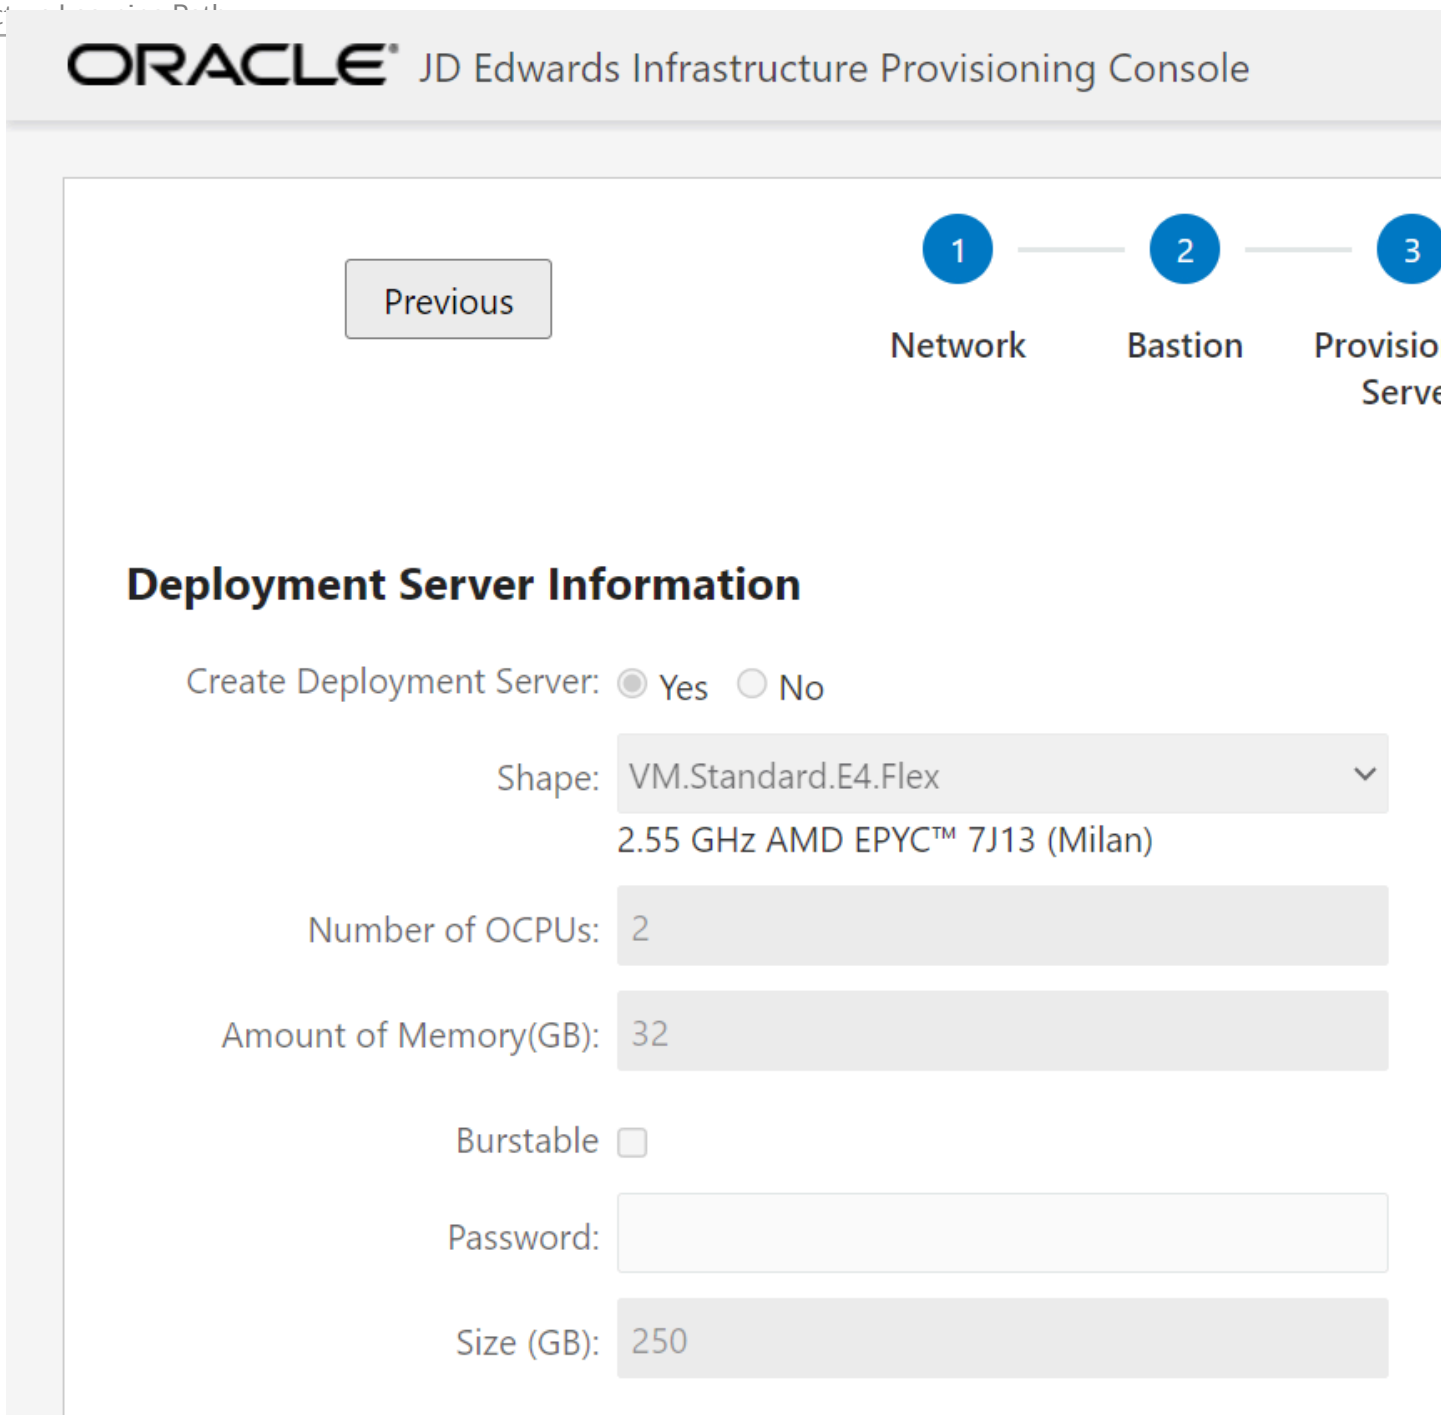

13. Click **Next**.

> **Note:** The following step is only applicable if you are provisioning a Production environment. Otherwise, you can skip this step.

14. This step is only applicable if you are configuring a Production environment (see **Note** above).

On Production Environment Information, enter valid values as illustrated in the following example:

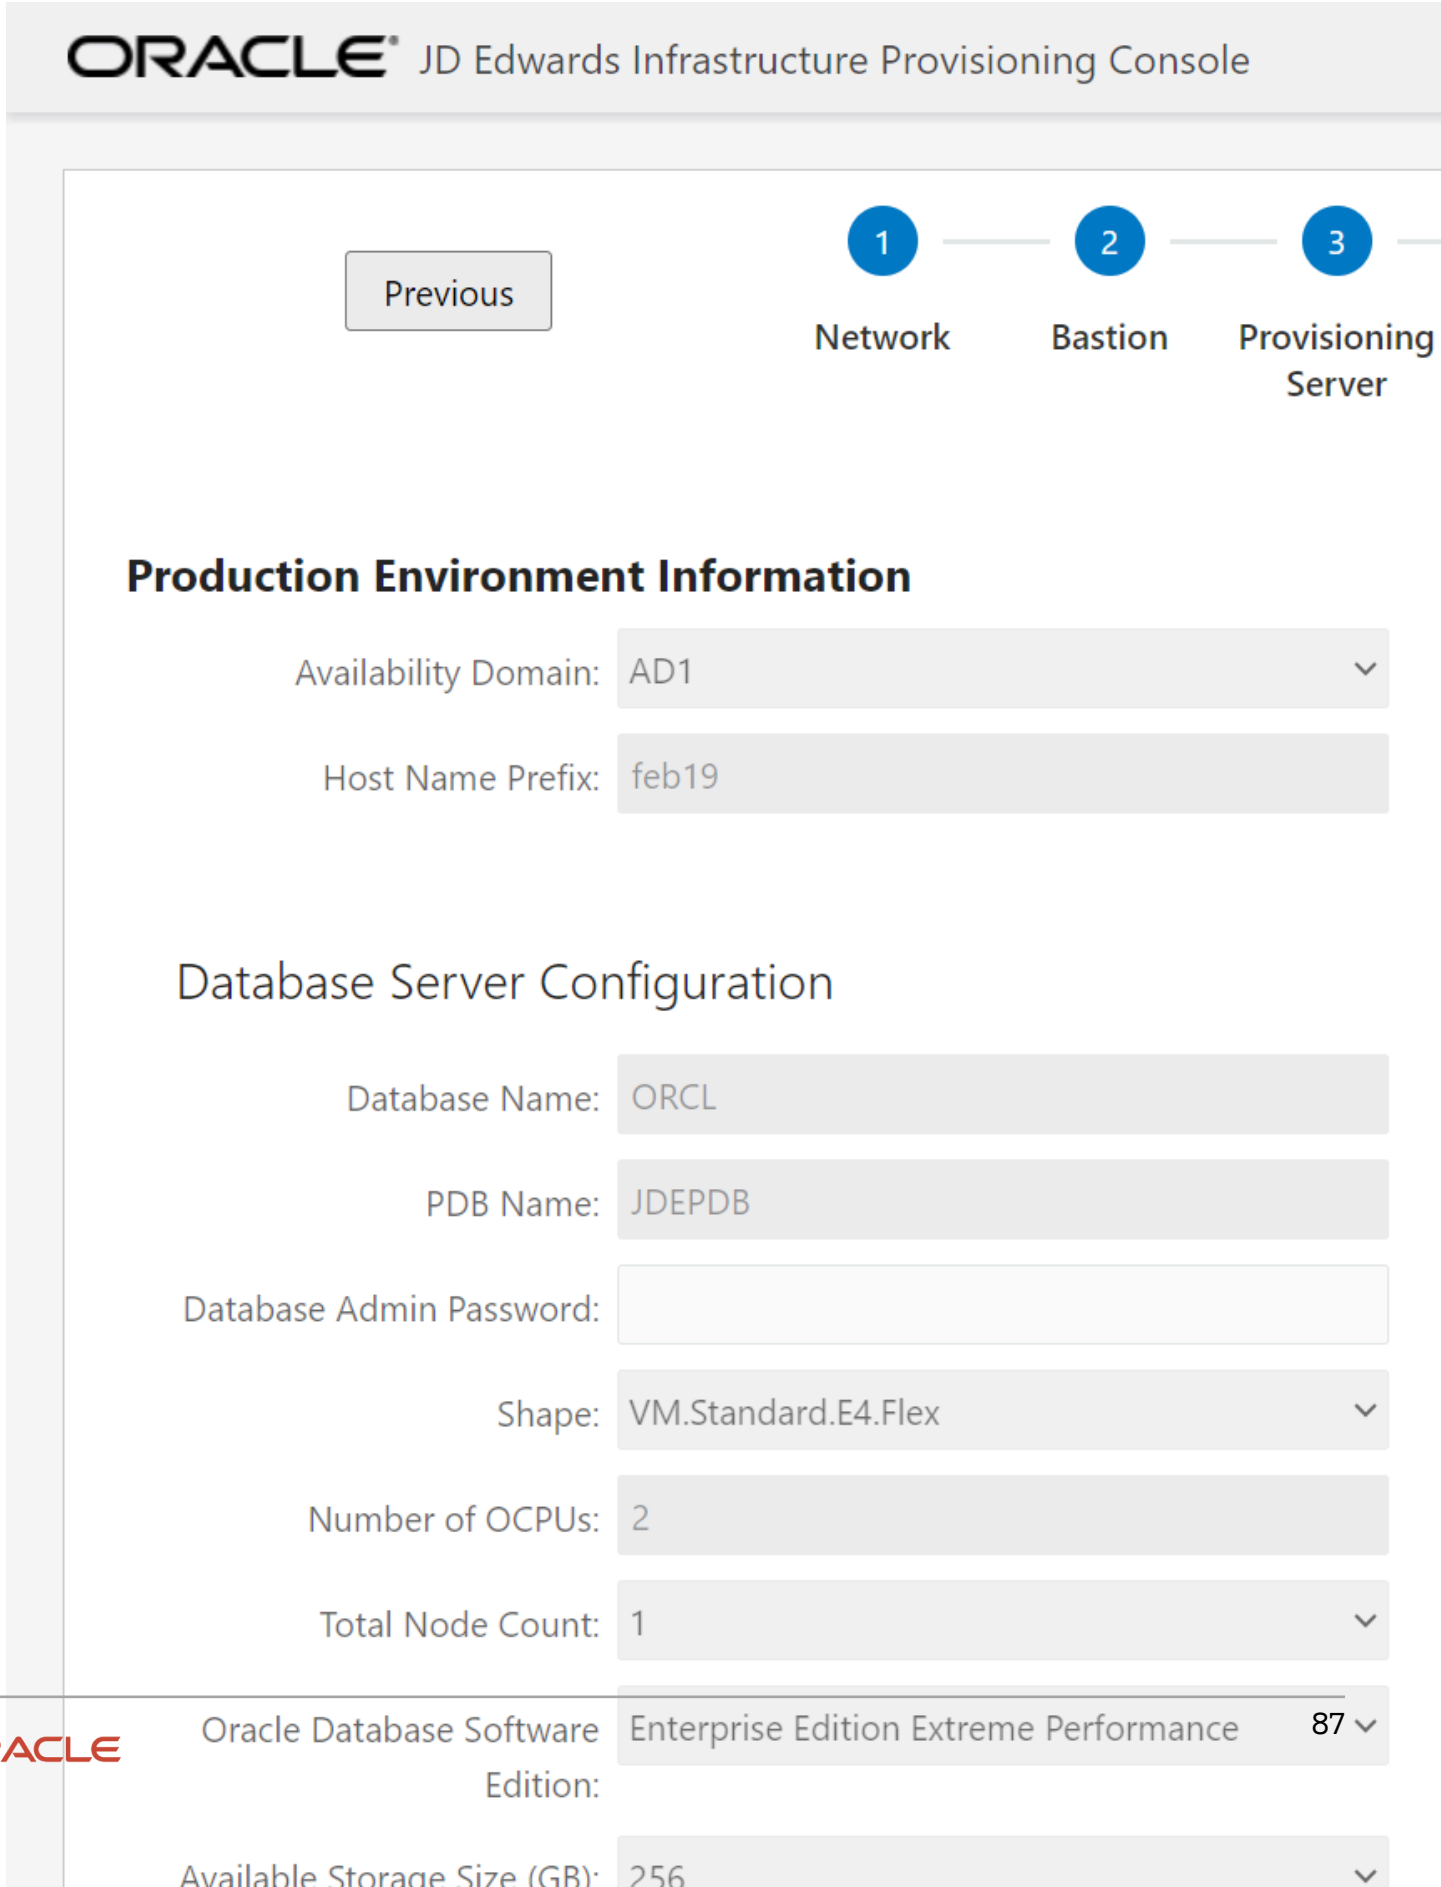

- *Availability Domain*

  Select the required availability Domain for the production servers. The available domains are: AD1, AD2, and AD3.

- *Host Name Prefix*

  Enter the host name prefix for the production servers.

  All the instances will be created with the prefix value entered in this field. The length of the prefix must not exceed 6 characters.

  For example, if the Prefix value is entered as pdjde,the instances are created with the following names:

  Database System: pdjdedb[1-N]

  Logic Server: pdjdelogic[1-N]

  Batch Server: pdjdebatch[1-N]

  Weblogic Server: pdjdewls[1-N]

  Deployment Server: pdjdedep[1-N]

Database Server Configuration

- *Database Name*

  The value in this field is pre-populated.

- *PDB Name*

  The value in this field is pre-populated.

- *Database Admin Password*

  Enter the password for the Database Administrator. The password must be 9 to 30 characters long, must contain at least two upper case letters, two lower case letters, two numbers, and two special characters.

  > **Note:** You must remember to enter this same password if you want to edit the information on this page in case of an error during the provisioning. You can enter a new password or use this same password if you want edit the information or to increase the server count after the provisioning is successful. The new password will apply only to the newly added servers.

- *Shape*

  Select a supported shape from the pull-down list.

  > **Note:** Do not select a non-supported shape, such as Exadata, or else you will experience downstream run-time errors.

- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPUs recommended for the Database Server is 2.

- *Total Node Count*

Select the required DB node count from the drop-down list. The available options are 1 and 2.

- *Oracle Database Software Edition*

  Select the Database software edition from the drop-down list.

  Your selection depends on the value you have entered for the field **Total Node Count**. If you have selected a value of **Total Node Count=1**, you can select any edition. However, for JD Edwards EnterpriseOne, the recommended value is **Enterprise Edition Extreme Performance**. If you selected

a value for **Total Node Count=2**, the value **Enterprise Edition Extreme Performance** is automatically selected and other values are disabled.

- *Available Storage Size (GB)*

  Select the DB block storage size from the drop-down list.
- *License Type*

  This option is pre-selected. The available options are License Included and Bring your Own License (BYOL).
- *Database Version*

  Select the Database version from the drop-down list.
- *Database Workload*

  The value in this field is pre-populated.

POD Configuration

- *POD Count*

  A single POD count includes one Logic Server instance, one Batch Server instance, and one WebLogic Server instance. The minimum value for POD count is 2.

Enterprise Server Configuration

Logic ES Information

- *Shape*

  Select the shape list for the instance.
- *Block Volume (GB)*

  Enter the logic block storage size.

  > **Note:** Additional storage is available at `/u01`.

- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPUs recommended for the Logic Enterprise Server is 2.
- *Amount of Memory (GB)*

  For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the Logic Enterprise Server is 32 GB.
- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
- *Burst Size*

  For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.
- *Load Balancer Virtual Host Name*

  Enter the load balancer virtual host name for the logic Enterprise Server.

> **Note:** The virtual host name must be same as the name entered while creating the LBaaS certificate. For example, logiclb.

Batch ES Information

- *Shape*

  Select the shape list for the instance.
- *Block Volume (GB)*

  Enter the batch storage size.

  > **Note:** Additional storage is available at `/u01`.
- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPU recommended for the Batch Enterprise Server is 2.
- *Amount of Memory (GB)*

  For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the Batch Enterprise Server is 32 GB.
- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
- *Burst Size*

  For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.
- *Load Balancer Virtual Host Name*

  Enter the load balancer virtual host name for the Batch Enterprise Server.

  > **Note:** The virtual host name must be same as the name entered while creating the LBaaS certificate. For example, batchlb.

WebLogic Server Configuration

- *WebLogic Installation*

  Use the radio buttons to select the installation type for WebLogic Server. The options are:

  - *WebLogic Server Image*

    This option deploys and configures sufficient volume storage for the supported version of WebLogic Server image from the Marketplace of Oracle Cloud Infrastructure.

    > **Note:** Based on the POD count, one WebLogic Server instance will be orchestrated to deploy one HTML server, one AIS server, and one HTML server for the AIS server. The `ochestration.json` will be generated with these 3 sets of servers for one POD.

  - *Standard Installation (manual)*

This option deploys and configures the supported version of WebLogic Server. Included in this deployment is the creation of the supported Linux Compute instance as well as the required minimum storage.

- *WebLogic Admin Password*

  Enter the password for the WebLogic Administrator.

  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.

  > **Note:** You must remember to enter this same password if you want to edit the information on this page in case of an error during the provisioning. You can enter a new password or use this same password if you want edit the information or to increase the server count after the provisioning is successful. The new password will apply only to the newly added servers.

- *WebLogic Server Version*

  The value in this field is pre-populated based on the WebLogic Server version that is supported by this release of Infrastructure Provisioning. This value is for informational purposes only; it cannot be changed.

- *WebLogic Server License Type*

  Use the pull-down menu to select your WebLogic Server license type.

- *Shape*

  Select the WebLogic Server instance shape list.

- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPU recommended for the WebLogic Server is 2.

- *Amount of Memory (GB)*

  For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the WebLogic Server is 32 GB.

- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.

- *Burst Size*

For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.

- *Block Volume (GB)*

  Enter the WebLogic Server block storage size.
- *Load Balancer Virtual Host Name for Web Servers*

  Enter the Load Balancer virtual host name for web servers (HTML and AIS).

  > **Note:** The virtual host name must be same as the name entered while creating the LBaaS certificate. For example, weblb.

- *HTTPS Listen Port Range for Web Servers*

  Select the HTTPS listen port range for the Web Server.

  This port range is used to orchestrate the HTML Server, AIS Server, and HTML for AIS Server in the generated `orchestration.json` file for the PD pathcode.

  > **Note:** Both the http and the https ports will be assigned for this range. It is recommended to enter minimum of 6 ports need to be added for HTML, AIS and HTML for AIS server respectively for the http and https ports. There is no requirement to add any other ports. For example, if you enter 8000 to 8005 , then all these servers will use one port for each http and https ports of HTML, AIS, and HTML for AIS servers.

- *LBaaS Listen Port for HTML*

  Enter the LBaaS https port number for the HTML Server.

  > **Note:** You can access the HTML Server through the Load Balancer using the IP address of the Load Balancer and this port number.

- *LBaaS Listen Port for AIS*

  Enter the LBaaS https port number for the AIS Server.

  > **Note:** You can access the AIS Server through the Load Balancer using the IP address of the Load Balancer and this port number.

- *LBaaS Listen Port for HTML for AIS*

  Enter the LBaaS listen port number for the HTML Server associated with the AIS Server.

  > **Note:** You can access the HTML Server associated with the AIS Server through the Load Balancer using the IP address of the Load Balancer and this port number.

LBaaS Configuration

- *Flex Shape Min Bandwidth*

  Use the numeric input field or the slider to specify the minimum bandwidth of the LBaaS Flex shape. In the range of allowable values, the lowest minimum value is 10 Mbps.
- *Flex Shape Max Bandwidth*

Use the numeric input field or the slider to specify the maximum bandwidth of the LBaaS Flex shape. In the range of allowable values, the highest maximum value is 8000 Mbps.

- *Certificate Name*

  Enter the load balancer certificate name. This is the name of the Load Balancer Certificate created on OCI console.

- *Public Certificate*

  Enter the content in the Load Balancer Public Certificate. To enter the content, you should use an ASCII editor to cut and paste the value from the generated certificate into this field.

  > **Note:** To enter the content in this field, you should use an ASCII editor to cut and paste the value from the generated certificate. For example, **device.crt** that is created in the section entitled: *Generating CA Certificates for Load Balancing as a Service (LBaaS)*.

- *CA Certificate*

  Enter the content in the Load Balancer CA certificate.

  > **Note:** To enter the content in this field, you should use an ASCII editor to cut and paste the value from the generated certificate.

  > **Note:** Be sure to include the BEGIN and END statements in the cut and paste content.

  For example, **rootCA.pem** that is created in the section entitled: *Generating CA Certificates for Load Balancing as a Service (LBaaS)*.

- *Certificate Private Key*

  Enter the content in the Load Balancer Certificate Private Key.

  > **Note:** To enter the content in this field, you should use an ASCII editor to cut and paste the value from the generated key.

  > **Note:** Be sure to include the BEGIN and END statements in the cut and paste content.

  For example, **device.key** that is created in the section entitled: *Generating CA Certificates for Load Balancing as a Service (LBaaS)*.

- *Certificate Passphrase*

  This is an optional field. Enter the passphrase for the LBaaS certificate.

15. Click **Next**.

If you have selected **WebLogic Server Image** option for the deployment of WLS, the following dialog for License Agreement is displayed:

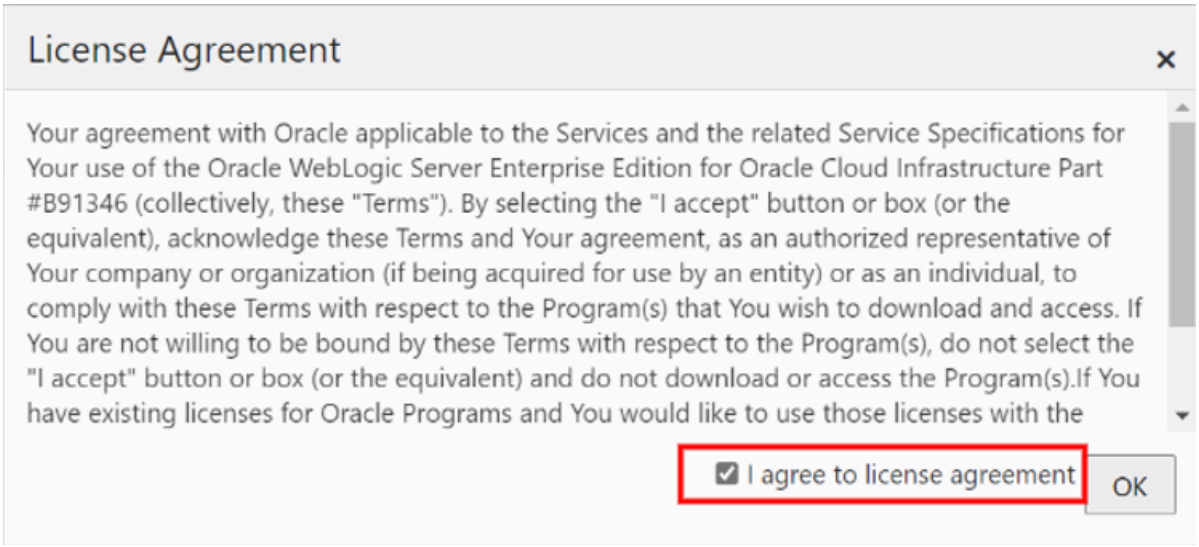

Note: You must review and agree to the terms of the License Agreement by clicking the "**I agree to license agreement**" checkbox before you can proceed with Infrastructure Provisioning. Click the **OK** button to proceed.

Note: The following step is only applicable if you are provisioning a Non-Production environment. Otherwise, you can skip this step.

16. This step is only applicable if you are configuring a Non-Production environment (see Note above).

On Non-Production Environment Information, enter valid values as illustrated in the following example:

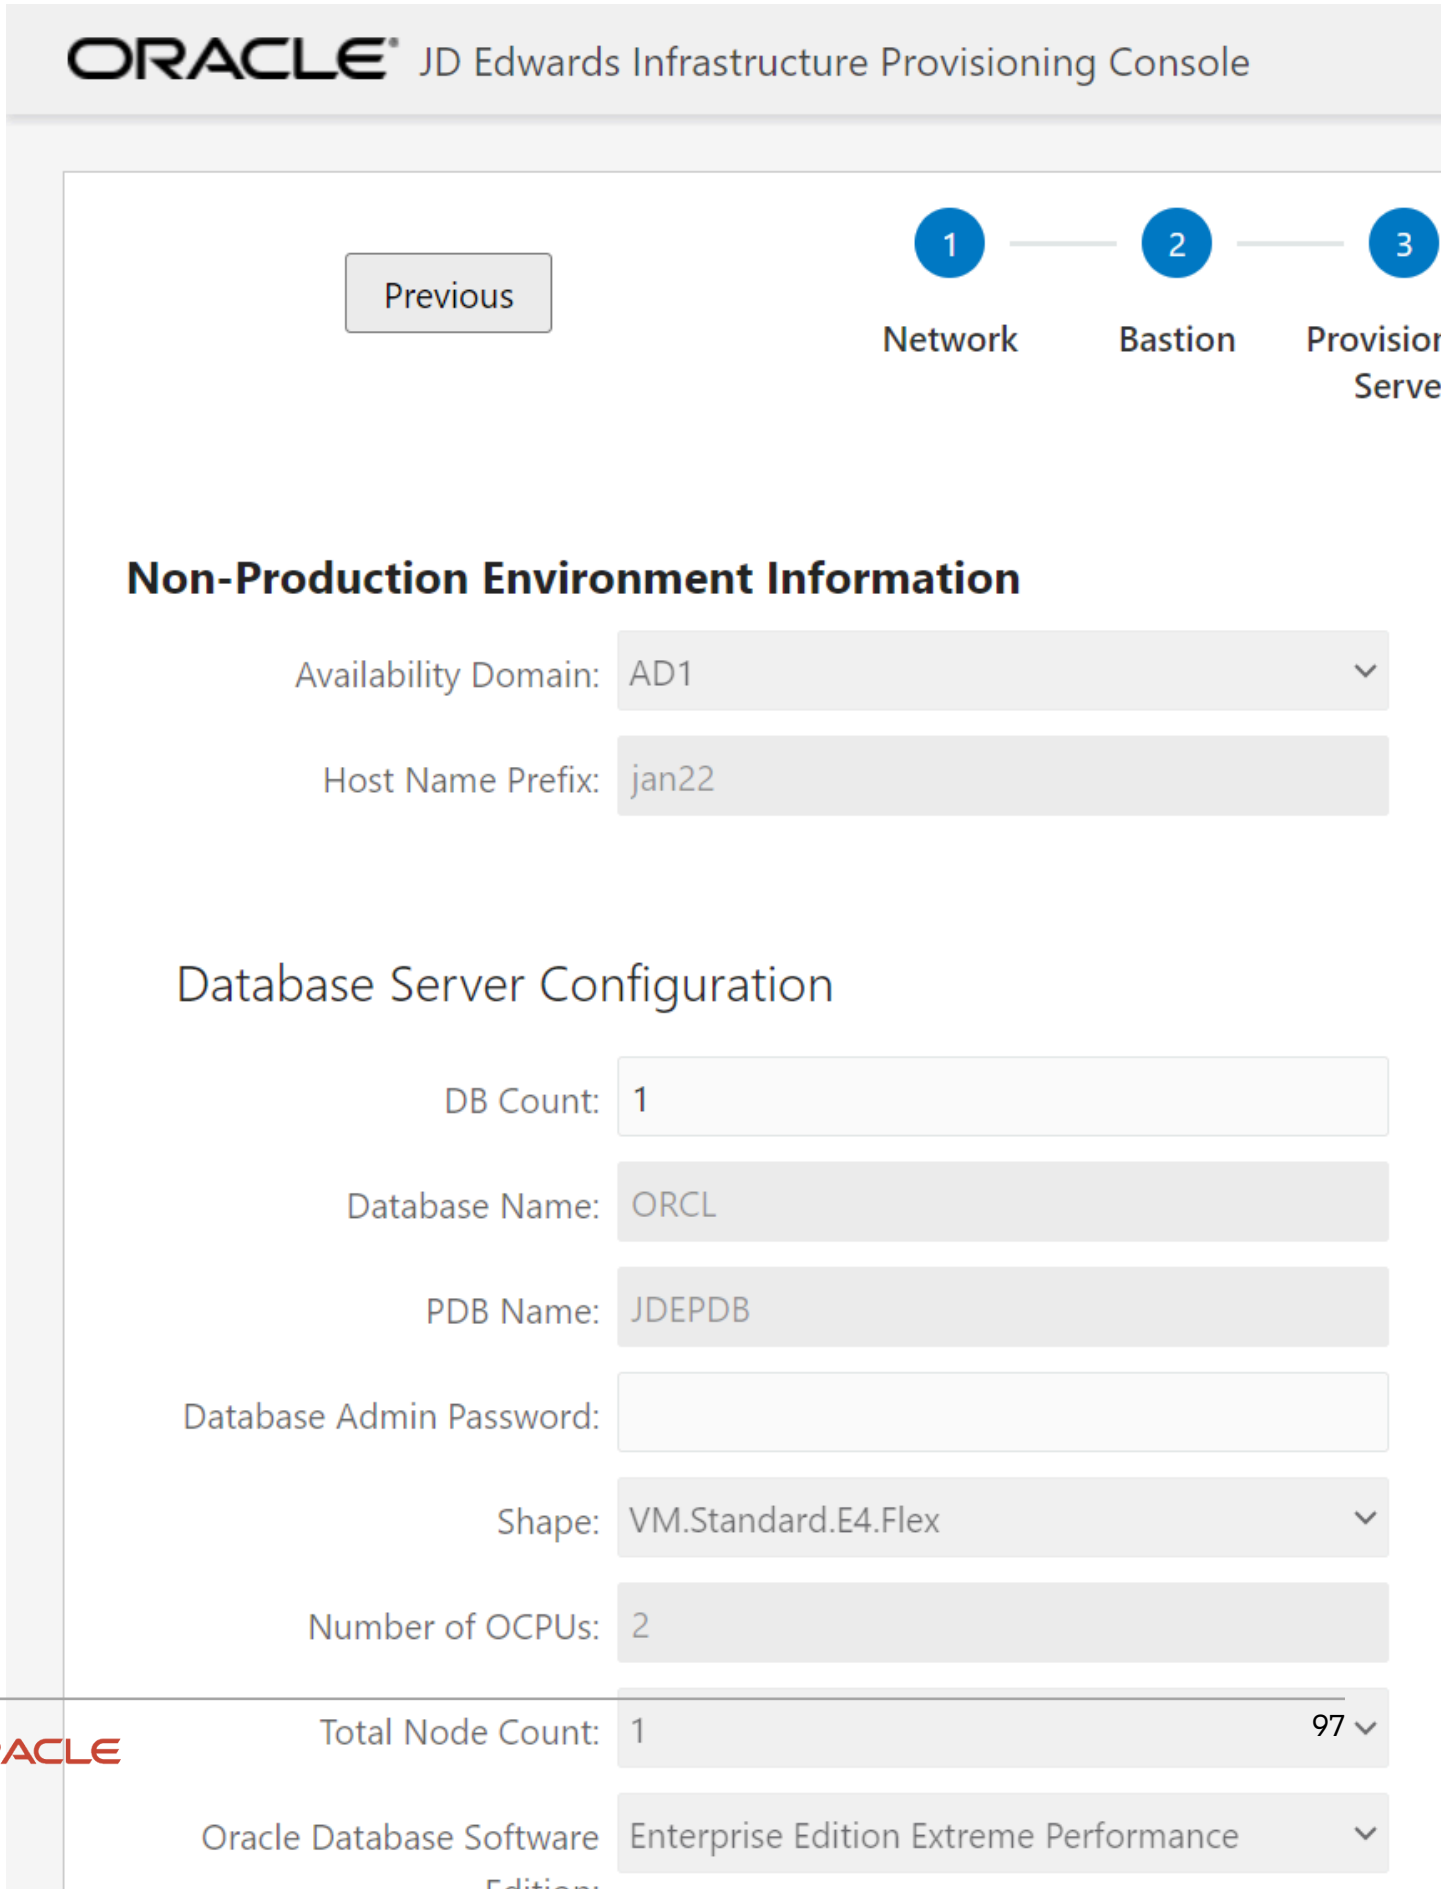

- *Availability Domain*

  Select the required availability Domain for the non- production servers. The available domains are: AD1, AD2, and AD3.

  > **Note:** Note: You can select only one availability domain for the non-production server.

- *Host Name Prefix*

  Enter the host name prefix for the non-production servers.

  > **Note:** All of the non-production instances are created using the prefix value entered in this field.

Database Server Configuration

- *Database Name*

  The value in this field is pre-populated.

Database Server Configuration

- *Database Name*

  The value in this field is pre-populated.
- *DB Count*

  Select the DB count from the drop-down list.
- *PDB Name*

  The value in this field is pre-populated.
- *Database Admin Password*

  Enter the password for the Database Administrator. The password must be 9 to 30 characters long, must contain at least two upper case letters, two lower case letters, two numbers, and two special characters.

  > **Note:** You must remember to enter this same password if you want to edit the information on this page in case of an error during the provisioning. You can enter a new password or use this same password if you want to edit the information or to increase the server count after the provisioning is successful. The new password will apply only to the newly added servers.

- *Shape*

  Select a supported shape from the pull-down list.

  > **Note:** Do not select a non-supported shape, such as Exadata, or else you will experience downstream run-time errors.

- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPUs recommended for the Database Server is 2.
- *Total Node Count*

  Select the required DB node count from the drop-down list.
- *Oracle Database Software Edition*

Select the Database software edition from the drop-down list.

- *Available Storage Size*

  Select the DB block storage size from the drop-down list.
- *License Type*

  This option is pre-selected. The available options are License Included and Bring your Own License (BYOL).
- *Database Version*

  Select the Database version from the drop-down list.
- *Database Workload*

  This option is pre-selected.
- *Character Set*

  The value in this field is pre-populated.
- *National Character Set*

  The value in this field is pre-populated.

Enterprise Server Configuration

- *Total Logic ES Count*

  Enter the total number of Enterprise Server count for the non-production environment as required.
- *Shape*

  Select the Logic Instance shape list.
- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPUs recommended for the Enterprise Server is 2.
- *Amount of Memory (GB)*

  For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the Enterprise Server is 32 GB.
- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
- *Burst Size*

  For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.
- *Block Volume (GB)*

  Enter the Logic Block Storage size.

  > **Note:** Additional storage is available at /u01.

WebLogic Server Configuration

- *WebLogic Server Count*

  Enter the WebLogic Server count for Non-Production.
- *WebLogic Server Version*

  The value in this field is pre-populated based on the WebLogic Server version that is supported by this release of Infrastructure Provisioning. This value is for informational purposes only; it cannot be changed.
- *WebLogic Server License Type*

  Use the pull-down menu to select your WebLogic Server license type.
- *WebLogic Admin Password*

  Enter the password for the WebLogic Administrator.

  The password must be 8 to 30 characters long, must contain at least one number or a special character.

  > **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.

  > **Note:** You must remember to enter this same password if you want to edit the information on this page in case of an error during the provisioning. You can enter a new password or use this same password if you want edit the information or to increase the server count after the provisioning is successful. The new password will apply only to the newly added servers.

- *Shape*

  Enter the WebLogic Server Instance shape list.
- *Number of OCPUs*

  Specify the number of OCPUs. The minimum number of OCPUs recommended for the WebLogic Server is 2.
- *Amount of Memory (GB)*

  For each OCPU you can select from 1 to 64 GB of memory with a maximum of 1024 GB total. The minimum amount of memory recommended for the WebLogic Server is 32 GB.
- *Burstable*

  For Flex shapes, you can enable the checkbox if you want to specify a burstable instance.

  If you enable a burstable instance, you can choose this option to provide a baseline level of CPU performance with the ability to burst to a higher level when required by your workload.
- *Burst Size*

  For Flex shapes, if you chose to create a burstable instance, you can also specify a Burst Size which you can adjust based on your requirements.

- *Block Volume (GB)*

  Enter the WebLogic Server Block storage size.

  > **Note:** Additional storage is available at /u01.

- *HTTPS Listen Port Range for Web Servers*

  Select the HTTPS listen port range required for all Non-Production HTML and AIS Servers.

  This port range is used to orchestrate the HTML Servers and the AIS Servers for the non-production pathcode.

  > **Note:** Both the http and the https ports will be assigned for this range. It is recommended to enter long ranges. For example, if you enter 8030 to 8050 (range of 20 ports), a maximum number of 15 HTML Servers or AIS Servers can be provisioned.

17. Click **Next.**

If you have selected **WebLogic Server Image** option for the deployment of WLS, the following dialog for License Agreement is displayed:

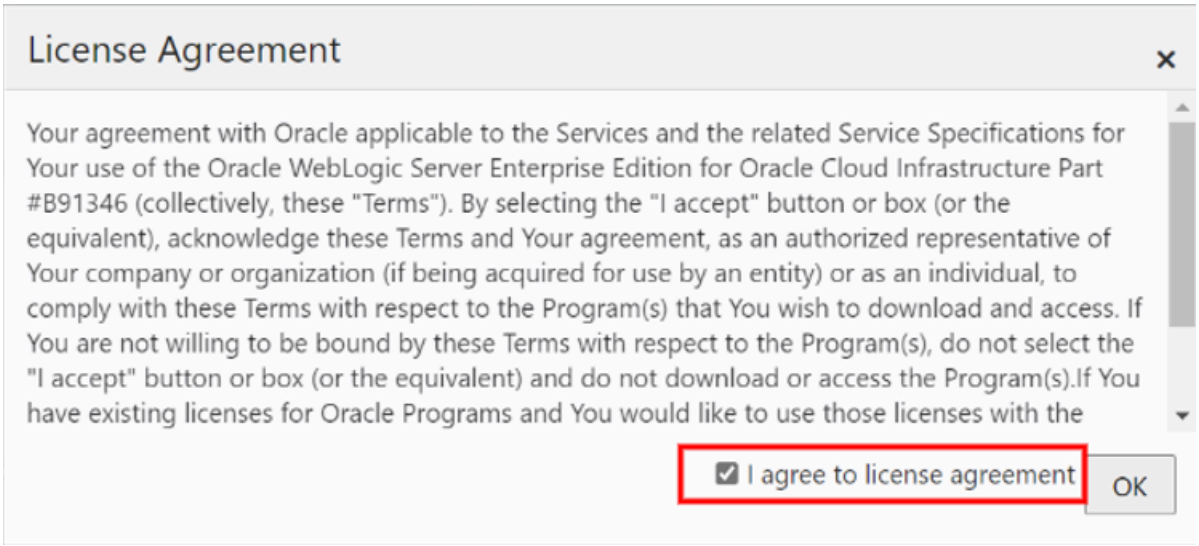

> **Note:** You must review and agree to the terms of the License Agreement by clicking the "**I agree to license agreement**" checkbox before you can proceed with Infrastructure Provisioning. Click the **OK** button to proceed.

> **Note:** The next step for Disaster Recovery is only applicable and supported for Production environments that are created with a single run (Production Environment + Disaster Recovery). This functionality is not supported in a separate run to only implement Disaster Recovery. Further, Disaster Recovery is not supported for non-Production environments. If none of these scenarios apply to your run of One-Click Infrastructure Provisioning, you can skip this step.

18. This step is only applicable if you want to implement Disaster Recovery (see **Note** above). On Disaster Recovery Environment Information, enter valid values as illustrated in the following example:

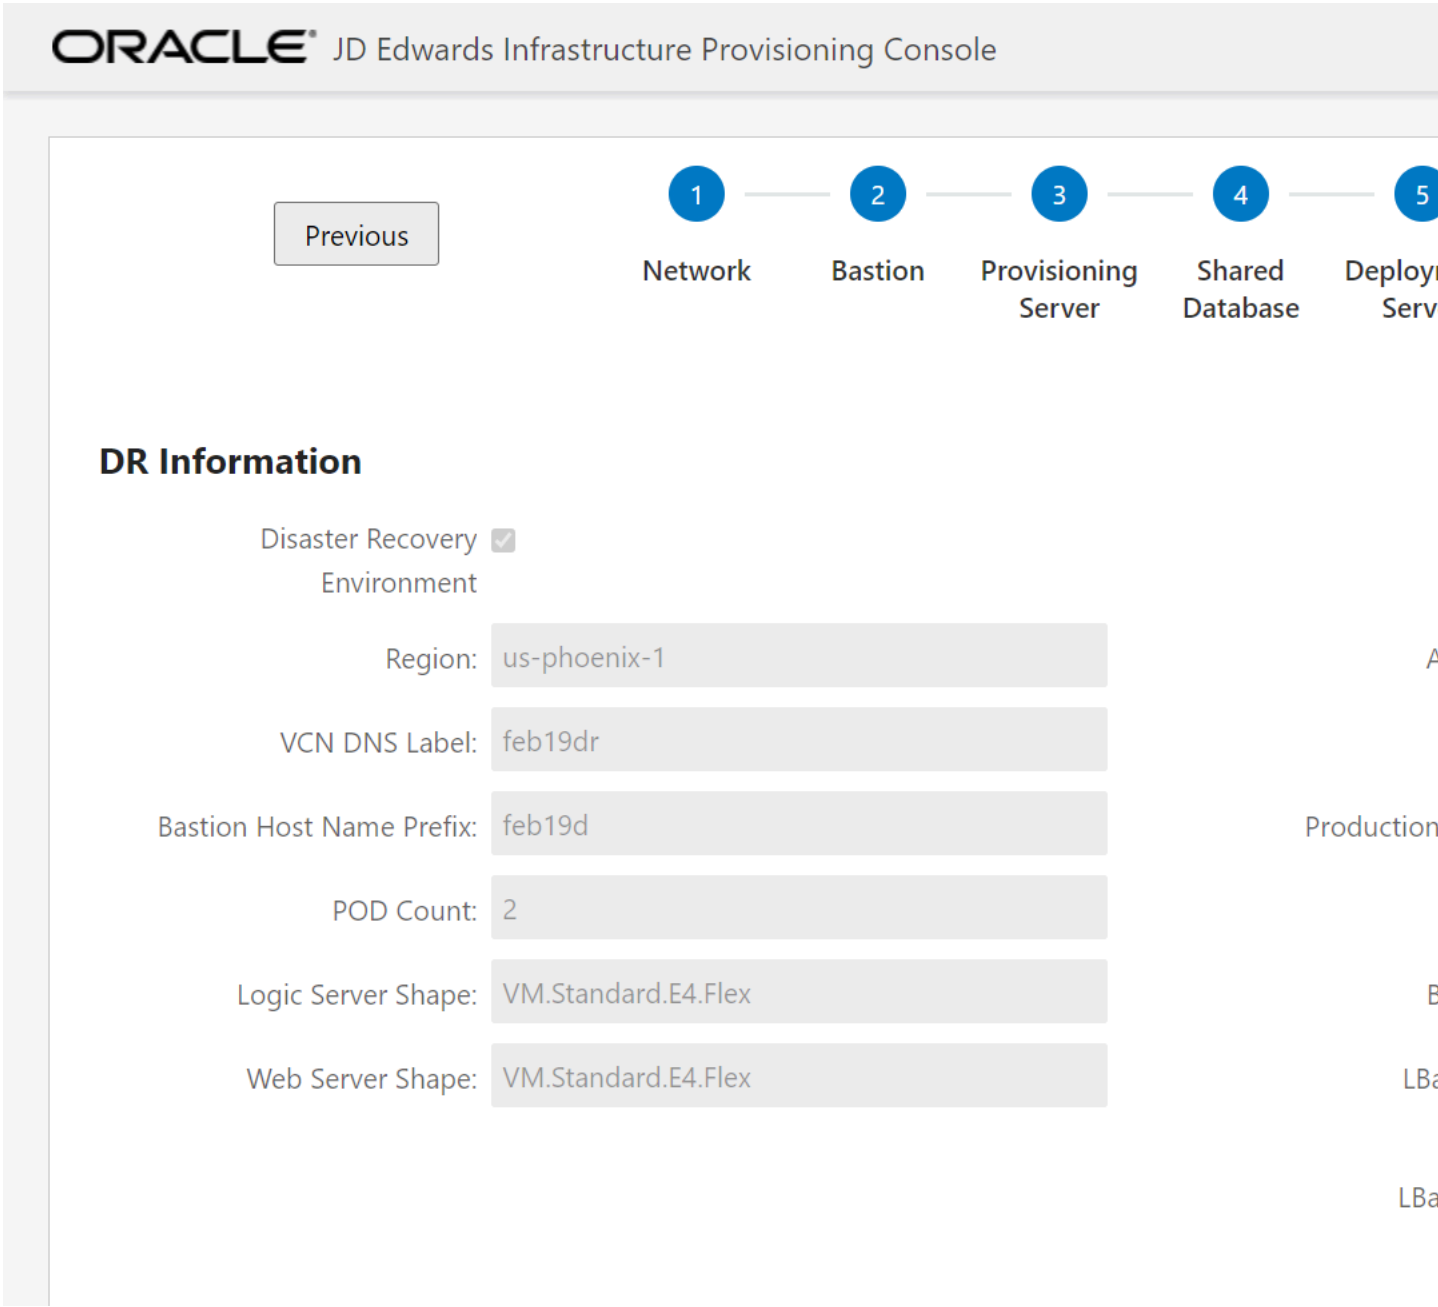

- *Region*

  Enter the **Region Identifier** of the region where disaster recovery resources need to be created.

  For example, for the Region Name US West Phoenix, the **Region Identifier** is **us-phoenix-1**.

Note: Oracle recommends the Disaster Recovery Region Identifier should not be same as the Region specified in the Network section (step 4 above). Therefore, in this example, if you specified the region as **us-phoenix-1** in the network section in step 4, you should specified a different region here for Disaster Recovery.

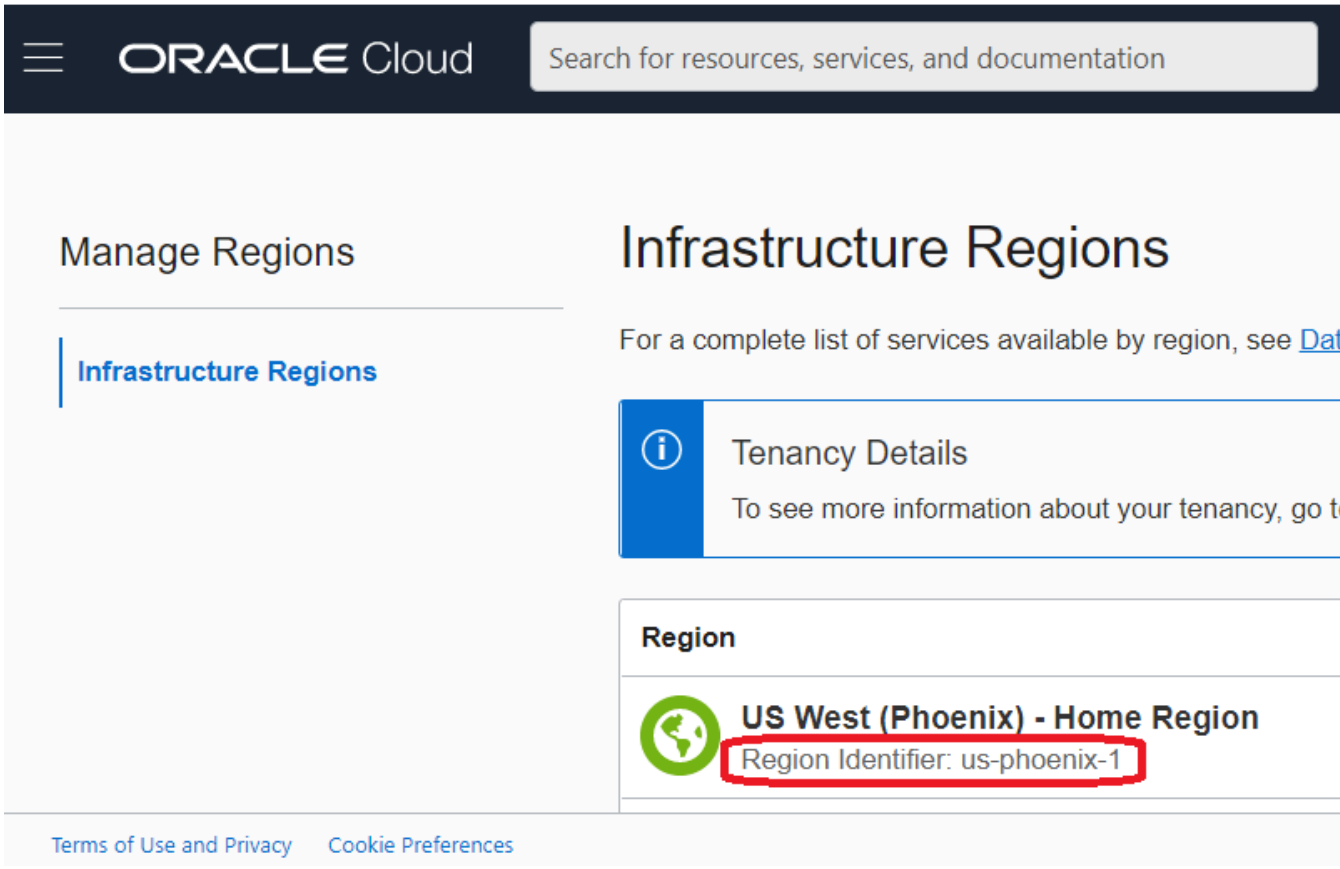

- *VCN DNS Label*

  Enter the value of the DNS Label for your VCN of Disaster Recovery environment.

  Note: The length of the field cannot exceed 15 characters. The name must start with a letter and must contain only alphanumeric characters. The value should only specify the globally unique domain name; do not include the prefix "www" or "local".

  Note: The value for this label **cannot** be the same as you have given on Network Details page.

- *Bastion Host Name Prefix*

  Enter the Bastion host name prefix for Disaster Recovery environment. The Bastion instance is created using the prefix value entered in this field. The length of the prefix must not exceed 6 characters. For example, if the prefix value is entered as **pdDR**, the server names will be created as follows:

  Bastion Host: pdDRbas[*xxx*][*N*]

  where ***[xxx]*** is the first three character of domain, and

where ***[N]*** is the number of servers, which can be 1-N.

> **Note:** The value for this label cannot be the same as you have given on Bastion Details page.

- *POD Count*

  The system populates this field with the value you have provided when you configured the Production Environment. You cannot update this field on this screen.

  > **Note:** A single POD count includes one Logic Server instance, one Batch Server instance, and one WebLogic Server instance. The minimum value for POD count is 2.

- *Logic Server Shape*

  The system populates this field with the shape of the Logic Server instance that you defined for the Disaster Recovery environment for the Production Environment. You cannot update this field on this screen.

- *Web Server Shape*

  The system populates this field with the shape of the Web Server instance that you defined for the Disaster Recovery environment for the Production Environment. You cannot update this field on this screen.

- *Availability Domain*

  Select the required availability domain for the Disaster Recovery environment. For most regions, the available domains are: AD1, AD2, and AD3.

- *VCN CIDR Block*

  The value for the VCN CIDR blocks indicates the network addresses that can be allocated to the resources of Disaster Recovery environment. The system calculates this value based on the VCN CIDR of the Production Environment. You cannot manually adjust this value.

- *Production Host Name Prefix*

  You should specify the host name prefix for the Disaster Recovery environment that you defined for the Production Environment.

  > **Note:** The value for this label cannot be the same as you have given for the production environment.

- *Database Server Shape*

  The system populates this field with the shape of the Database Server instance that you defined for the Disaster Recovery environment for the Production Environment. You cannot update this field on this screen.

- *Batch Server Shape*

  The system populates this field with the shape of the Batch Server instance that you defined for the Disaster Recovery environment for the Production Environment. You cannot update this field on this screen.

- *LBaaS Shape*

  The system populates this field with the shape of the LBaaS that you defined for the Disaster Recovery environment for the Production Environment. You cannot update this field on this screen.

19. Click **Finish**.

The OpenTofu scripts are automatically initiated. The following message is displayed if the provisioning is successful.

> **Note:** You can click **Edit** to go back and increase the server count.

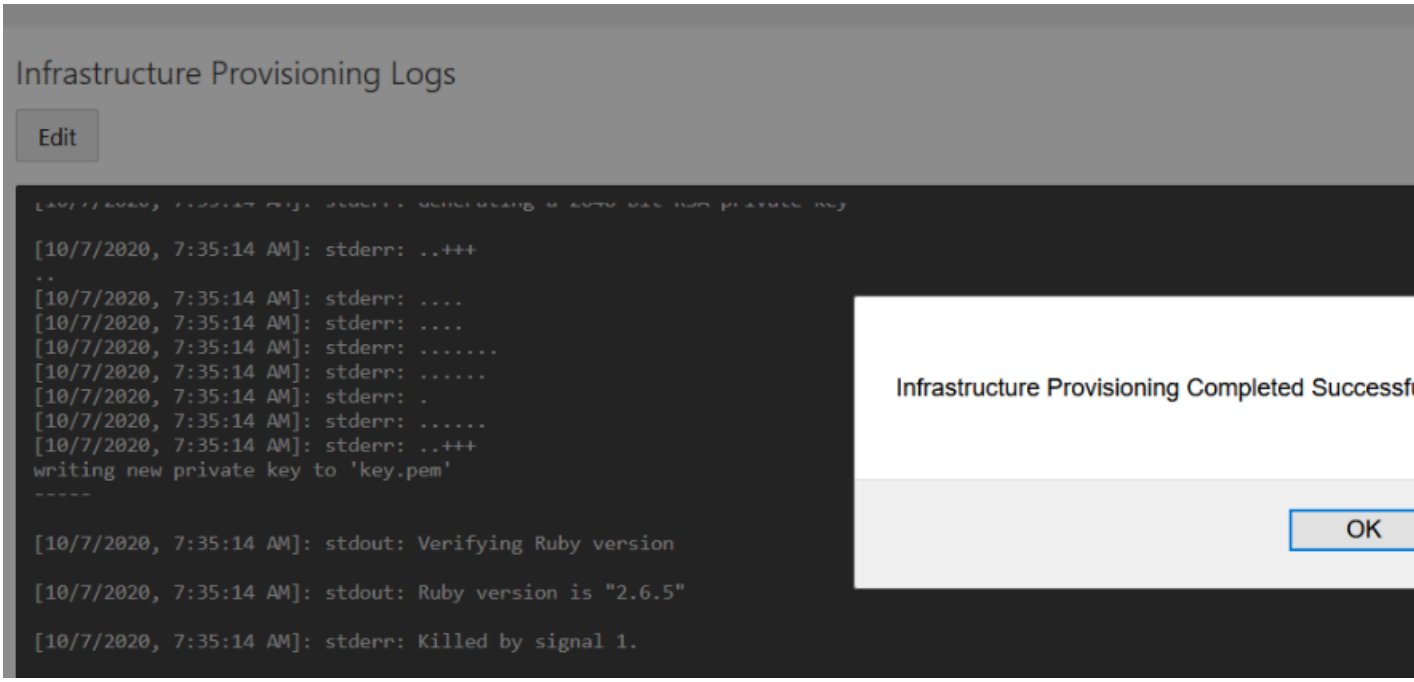

20. Verify that the files **pdOrch.json**, **nonpdOrch.json** and **infraOutput.json** are created in this path:

`JDERefArch_InfraProvisioning/E1InfraProvisionConsole/outputJson`

- `infraOutput.json` contains details about the production and non-production pathcodes, as well as the disaster recovery environment, if it was selected during Infrastructure Provisioning. This file also contains information such as the host name and Private IP of Logic Server Load Balancer, Batch Server Load Balancer, and WebLogic Server Load Balancer.
- `pdOrch.json` contains details of all the instances in the production environment, as well as the disaster recovery environment, if it was selected during Infrastructure Provisioning.
- `nonpdOrch.json` contains details of all the instances in the non-production environment. It is not possible to import this json file into the JD Edwards One-Click Provisioning Console. You can use the details in this file to orchestrate the components manually in the JD Edwards One-Click Provisioning Console.

> **Note:** You can review the log messages on the Infrastructure Provisioning Logs window. You can also find the log files in the following location. `JDERefArch_InfraProvisioning\\E1InfraProvisionConsole\logs`

21. Click the **Edit** button on the Infrastructure Provisioning Log window to correct the information provided in case of an error.

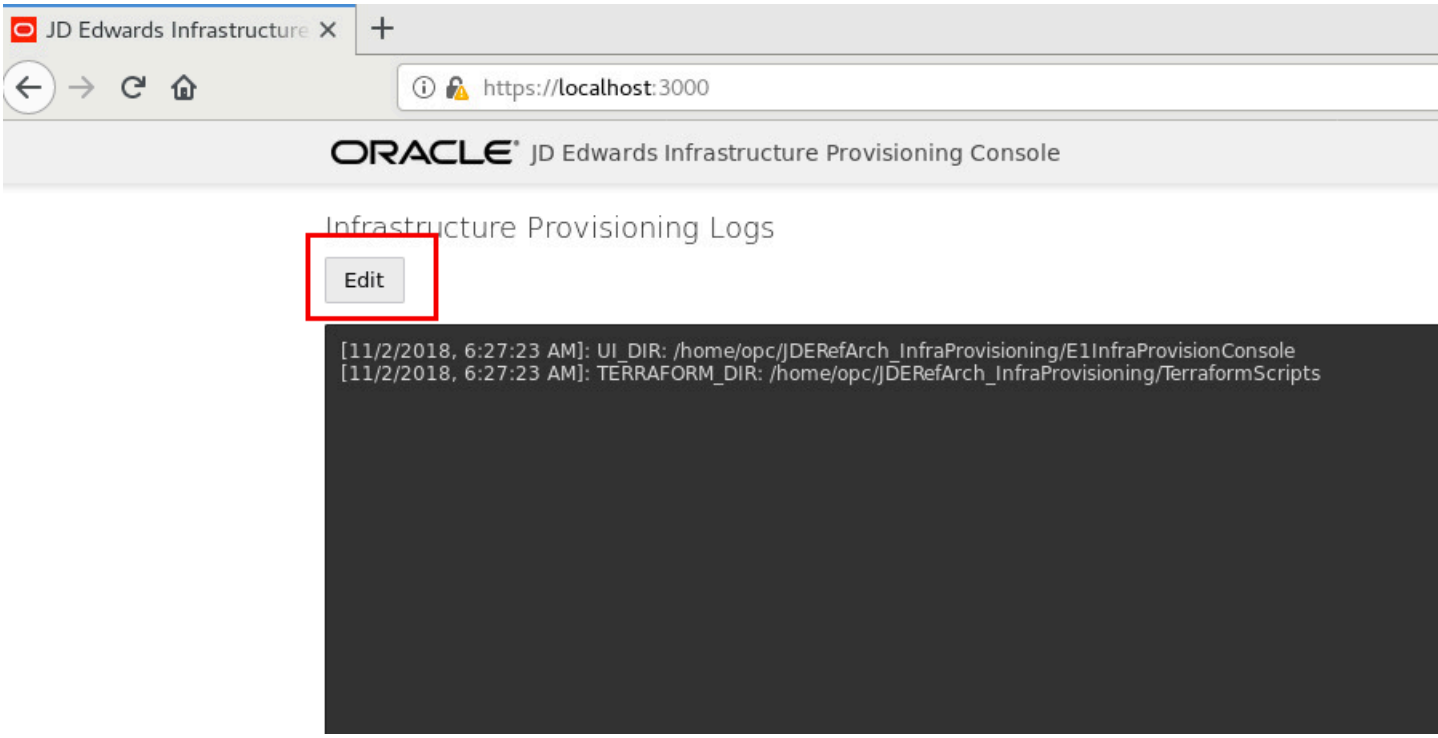

Applying Patches to the Oracle Database

Note: This purpose of this section is to emphasize the required sequence to apply patches to the Oracle database that you just deployed by JD Edwards EnterpriseOne Infrastructure Provisioning. You must not patch the Oracle database until after you have run JD Edwards EnterpriseOne One-Click Provisioning to populate the database, which is described in the next task in this tutorial. If you apply patches to the Oracle Database prior to running One-Click Provisioning, it may fail if the database is not as expected.

Copying Files from the OpenTofu Staging Server

This section shows you how to copy files from the OpenTofu Staging Sever.

After you have completed first-time configuration, you will need to copy files from the OpenTofu Staging Server to your local workstation by performing these steps:

1. Establishing a Secure FTP (SFTP) Connection
2. Copying Your Private SSH Keys from the OpenTofu Staging Server
3. Copying JSON Files from the OpenTofu Staging Server

Prerequisites

- Your OpenTofu Staging Server must be created and configured as described in the preceding OBEs in this Learning Path.
- You must have previously created key pairs in openssh format as described in the preceding OBE "**Generating Instance Key Pairs in openssh Format**" of this Learning Path.

  > **Tip:** The best practice is to use separate key pairs for accessing the Bastion host and all other servers.

- You must have an FTP program to copy your local private SSH keys to the OpenTofu Staging Server.

  > **Tip:** If you followed the recommendation in the referenced section above, you have given this file a significant name such as **jdeSSHKey.openssh**.

Establishing a Secure FTP (SFTP) Connection

Use this procedure to create a secure FTP (SFTP) connection from your Microsoft Windows workstation to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server.

> **Note:** This example in this procedure uses a program for Microsoft Windows called WinSCP. However, you can use any SFTP program to make the connection.

1. Start your SFTP program, such as WinSCP.

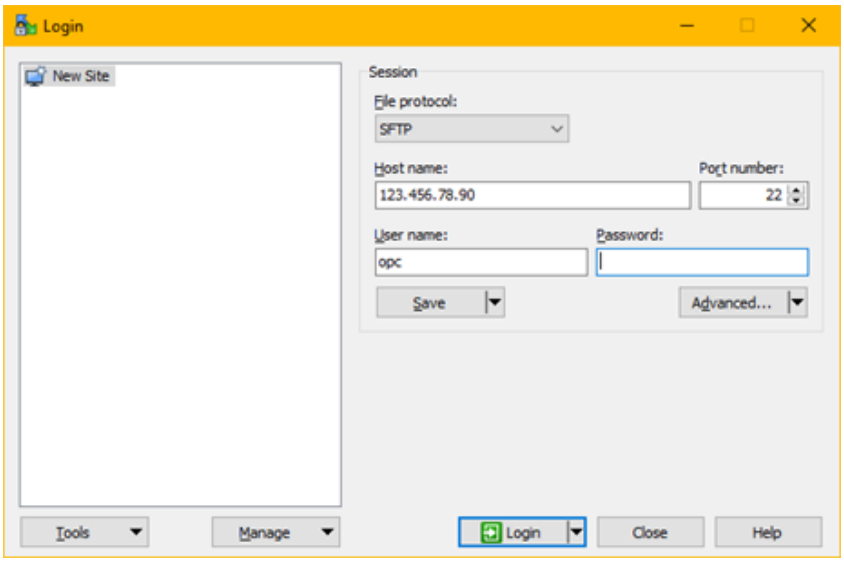

2. On the Login page, complete these fields:
   - *File protocol*FTP
   - *Host name*

     Enter the public IP address of the Linux instance you created in Oracle Cloud Infrastructure for the OpenTofu Staging Server.
   - *Port*Enter the value **22**.
   - *User name*Enter the value **opc**.
   - *Password*At this point, you cannot enter a password in the FTP program to complete the connection. Instead you must use the Advanced option and wait until prompted to enter a private key password by following the procedure described in Step 3 below.
   - *Advanced*

     Select **Advanced Site Settings** from the drop-down menu.

     In the tree structure in the left pane, navigate to SSH, Authentication.

     In the **Private key file** field, use the browse button to locate the private key file. If you followed the recommendation provided earlier, this file is named **jdeSSHKey.ppk**.

     Click the **OK** button to save the setting and return to the Login screen.
3. After completing all the above fields, click the **Login** option on the Login screen.

   After the SFTP program establishes a connection to the OpenTofu Staging Server, a warning is displayed.

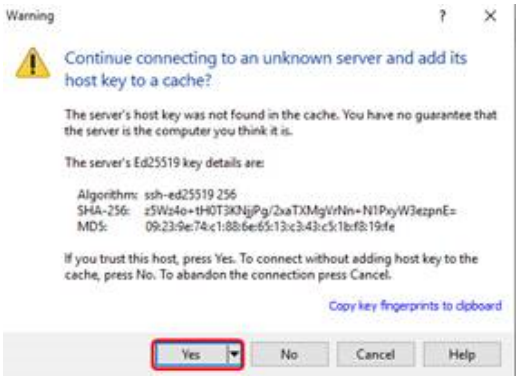

4. Unknown Server - Click Yes
5. On the Warning dialog box, click the **Yes** button to accept the connection to the “unknown server” and to add the host key of this server to the cache.

   After the connection is made, the SFTP program is ready to be used for file transfers from your workstation to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server.

Copying Your Private SSH Keys from the OpenTofu Staging Server

Use this procedure to copy your SSH keys from the OpenTofu Staging Server. These keys were created by the createStage.sh script run on the server and enable you to access all provisioned hosts, except the Bastion server. As a best practice, this key pair should be different than the one used to create and access the OpenTofu Staging Server.

1. Using an SFTP program (a secure connection to which can be created by following the preceding procedure), open a connection to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server.

2. Locate the **Bastion.openssh** file that was generated on the OpenTofu Staging Server when you ran the configuration. The file is located in this directory: `/home/opc/keys`
3. Copy the **Bastion.openssh** file to your local workstation as shown in the following example.

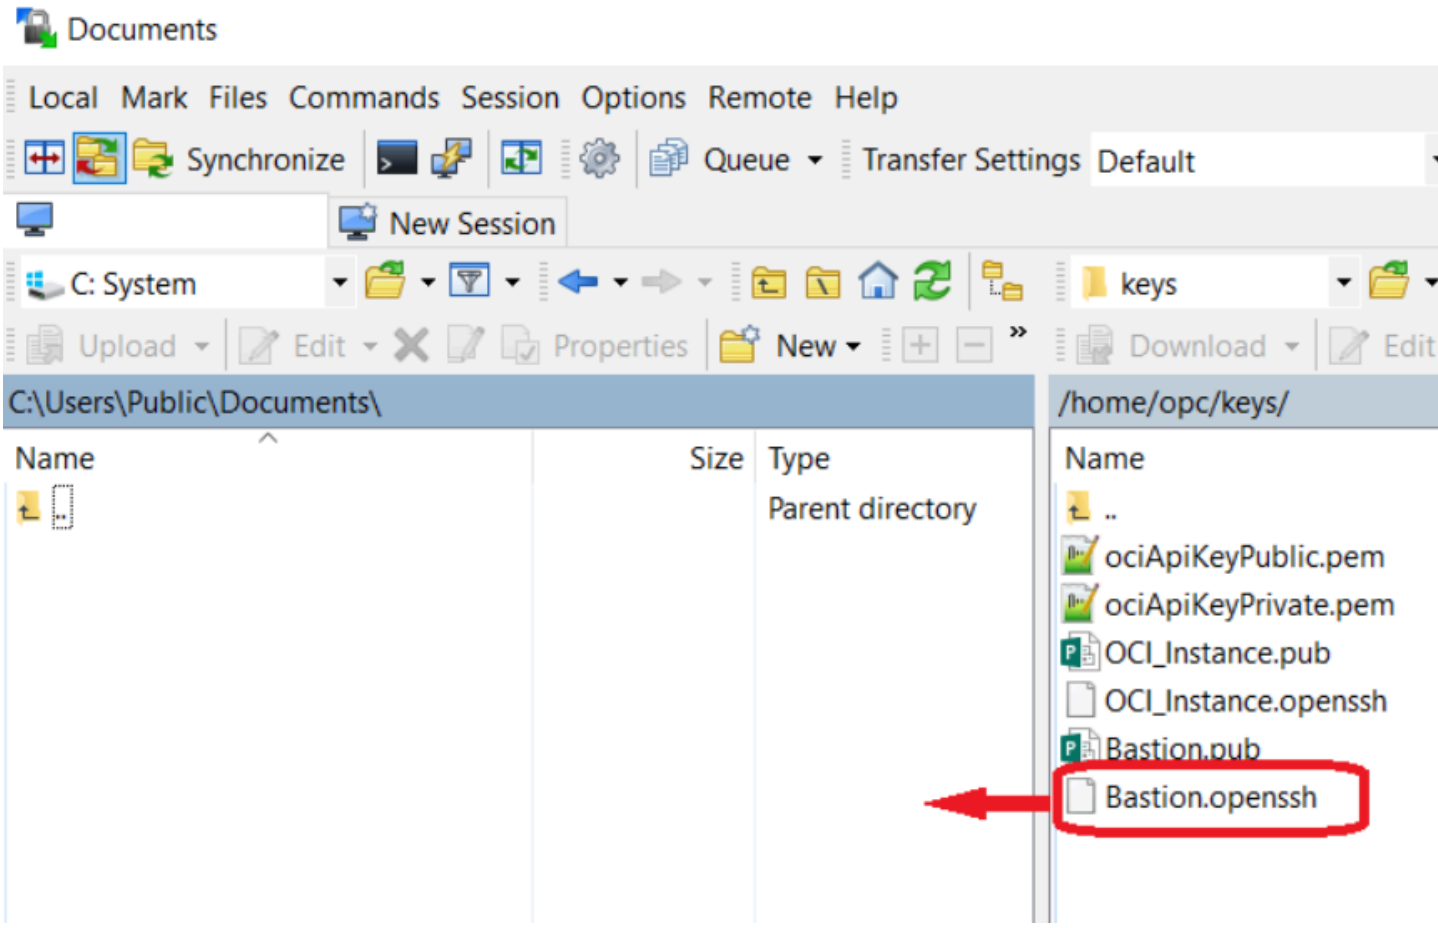

4. Locate the **OCI_Instance.openssh** file that was generated on the OpenTofu Staging Server when you ran the configuration. The file is also located in: `/home/opc/keys`

5. Copy the **OCI_Instance.openssh** file to your local workstation as shown in the following example.

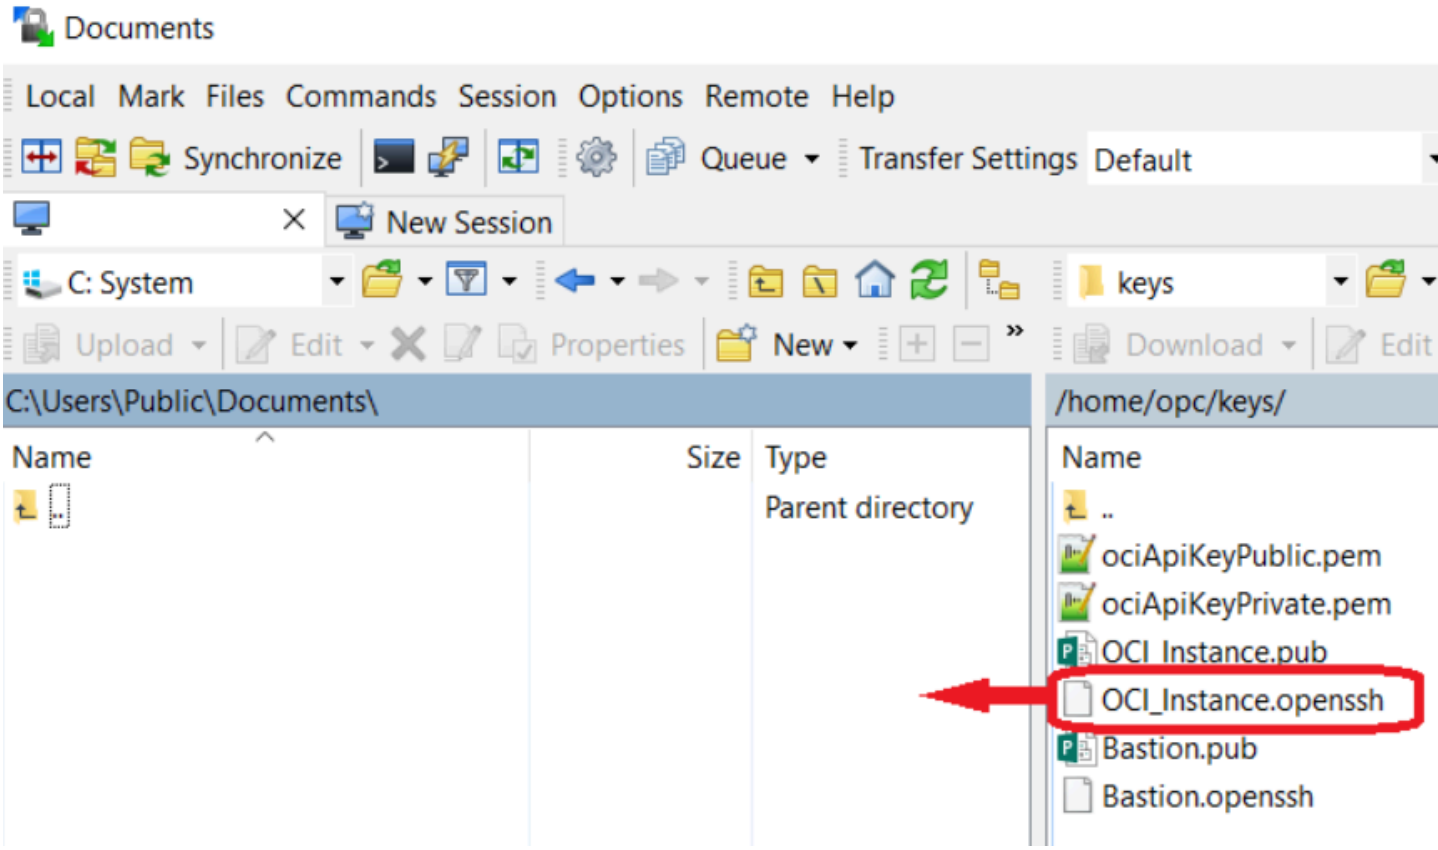

Copying JSON Files from the OpenTofu Staging Server

Use this procedure to copy the generated JSON files from the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server to your local workstation. These JSON files will be used to quickly configure the JD Edwards One-Click Provisioning Server to automatically define the servers on the infrastructure provisioned by OpenTofu.

1. Using an SFTP program (which can be created by following the preceding procedure), open a connection to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server.
2. Locate the JSON files generated on the OpenTofu Staging Server when you ran the configuration. The files are located in:

```
/home/opc/JDERefArch_InfraProvisioning/E1InfroProvisionConsole/outputJson
```

3. Copy all JSON files to your local workstation as shown in the following example.

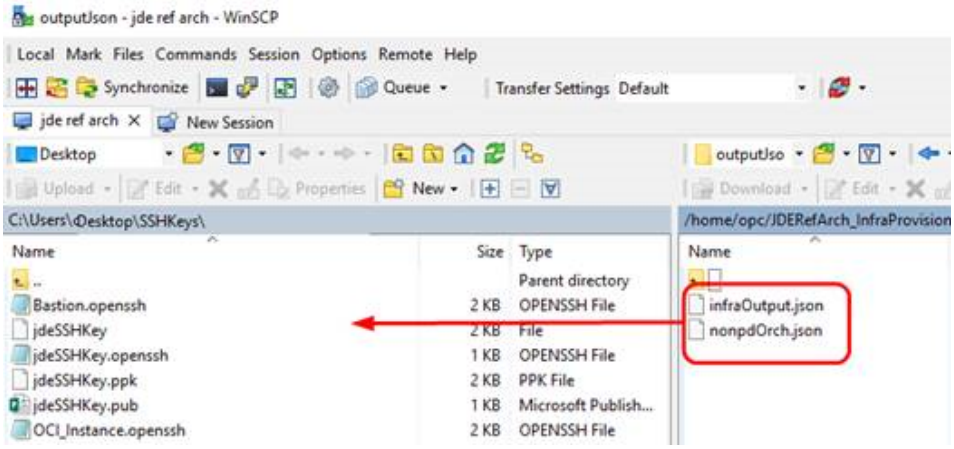

Converting Your Private SSH Keys to .ppk Format

This procedure describes how to convert the private key to .ppk format, which is the format required to connect to the Provisioning Server.

Prerequisite

You must have created a private SSH key by following the procedure described in the preceding OBE "Generating Instance Key Pairs in openssh Format" of this Learning Path.

Converting Your Private SSH Keys to .ppk Format

Use this procedure to convert your private SSH keys to .ppk format.

1. Run the puTTygen program.

2. From the Conversions menu item, select Import key.

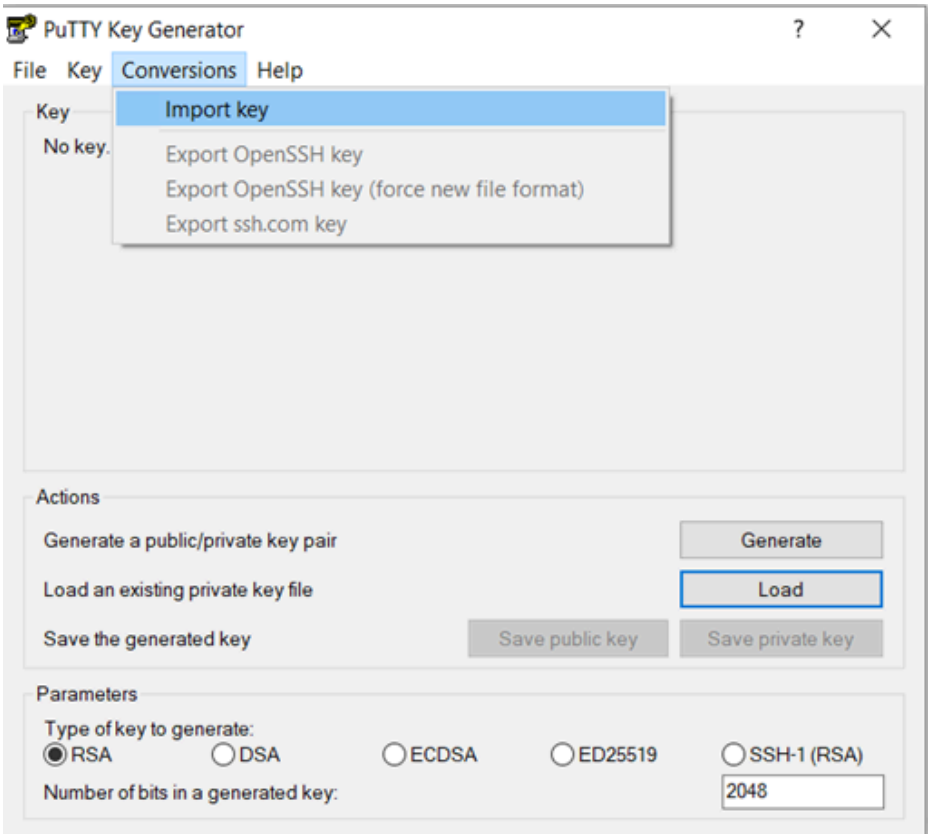

3. Browse to select the private key, which should be named: Bastion.openssh.

4. Click the Save private key button.

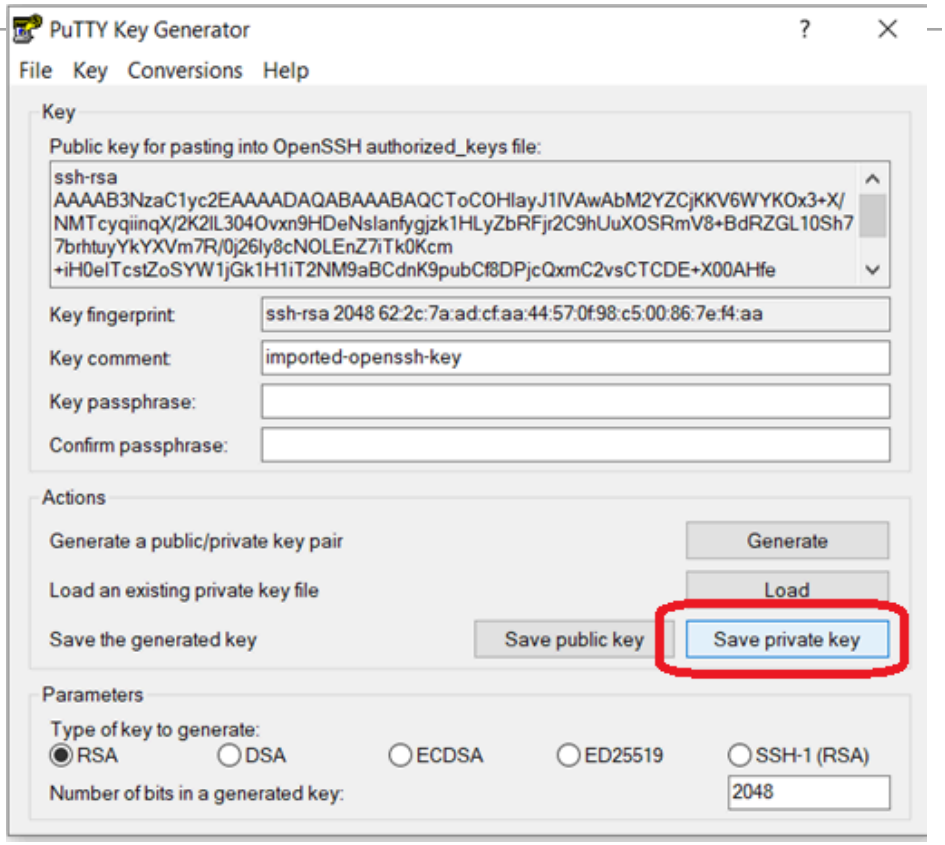

5. In the PuTTYgen Warning dialog box, click the Yes button.

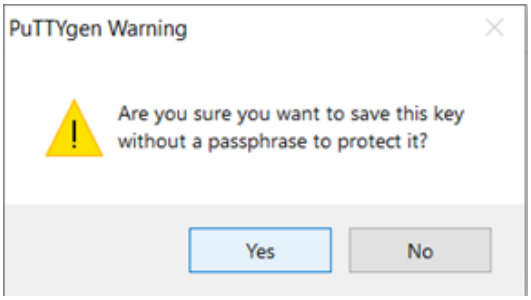

7 Using the One-Click Provisioning Console

Accessing the JD Edwards One-Click Provisioning Console

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

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

1. Configure – Provide the Server Manager details.
2. Orchestrate – Create the deployment plan.
3. Deploy – Initiate the scripts for the automated provisioning of the JD Edwards EnterpriseOne system.

Prerequisite

- 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 to Oracle Cloud Infrastructure and and an administrator account in the platform. For more information, refer to this site:
  Getting Started with Oracle Cloud Applications
- You must have installed PuTTY on your Microsoft Windows machine. By default, this installation includes the requisite software component called Pageant (Putty authentication agent). The program provides a Secure SHell (SSH) tunneling method for connecting to Unix or Linux machines through PuTTY.
- To access the JD Edwards One-Click Provisioning Console, you must use a supported browser. See Supported Browsers under *Browser Issues*.

Obtaining IP Addresses

You will need these IP addresses to complete this task:

- Public IP address of the bastion host
- Private IP address of the Provisioning Server (which also includes Server Manager)

These are the servers that were created by the Reference Architecture (OpenTofu) user interface using the Infrastructure Provisioning Console. The bastion host is the server to use as a gateway to the other servers. It is the only created server that has an external facing public IP address. The Provisioning Server (which includes Server Manager) is used to run the One-Click Provisioning Console, and it contains a pre-deployed Server Manager for JD Edwards EnterpriseOne.

Obtaining the Public IP Address of the Bastion Host

The name of the instance for the bastion host is formed by the value you enter in the Infrastructure Provisioning Console plus the system appended value `basxxxn`

where `bas` is a constant

where `xxx` is the 3 letter abbreviation for your region

where `n` is the suffix of that region name.

For example, the abbreviated value for the region `us-ashburn-1` is `ash1`.

Therefore, if you entered a value of `jdeprov` in the Infrastructure Provisioning Console as the name of the bastion server, the instance created by the OpenTofu scripts would be `jdeprovbasash1`.

You can determine and obtain the public IP address of the bastion host using these steps to navigate the user interface of Oracle Cloud Infrastructure.

1. Using the information above as criteria for the instance name, locate and click the instance for your bastion server for JD Edwards EnterpriseOne infrastructure provisioning.
2. Obtain and record the value from the Public IP Address field.

   You will need this value to access this instance.

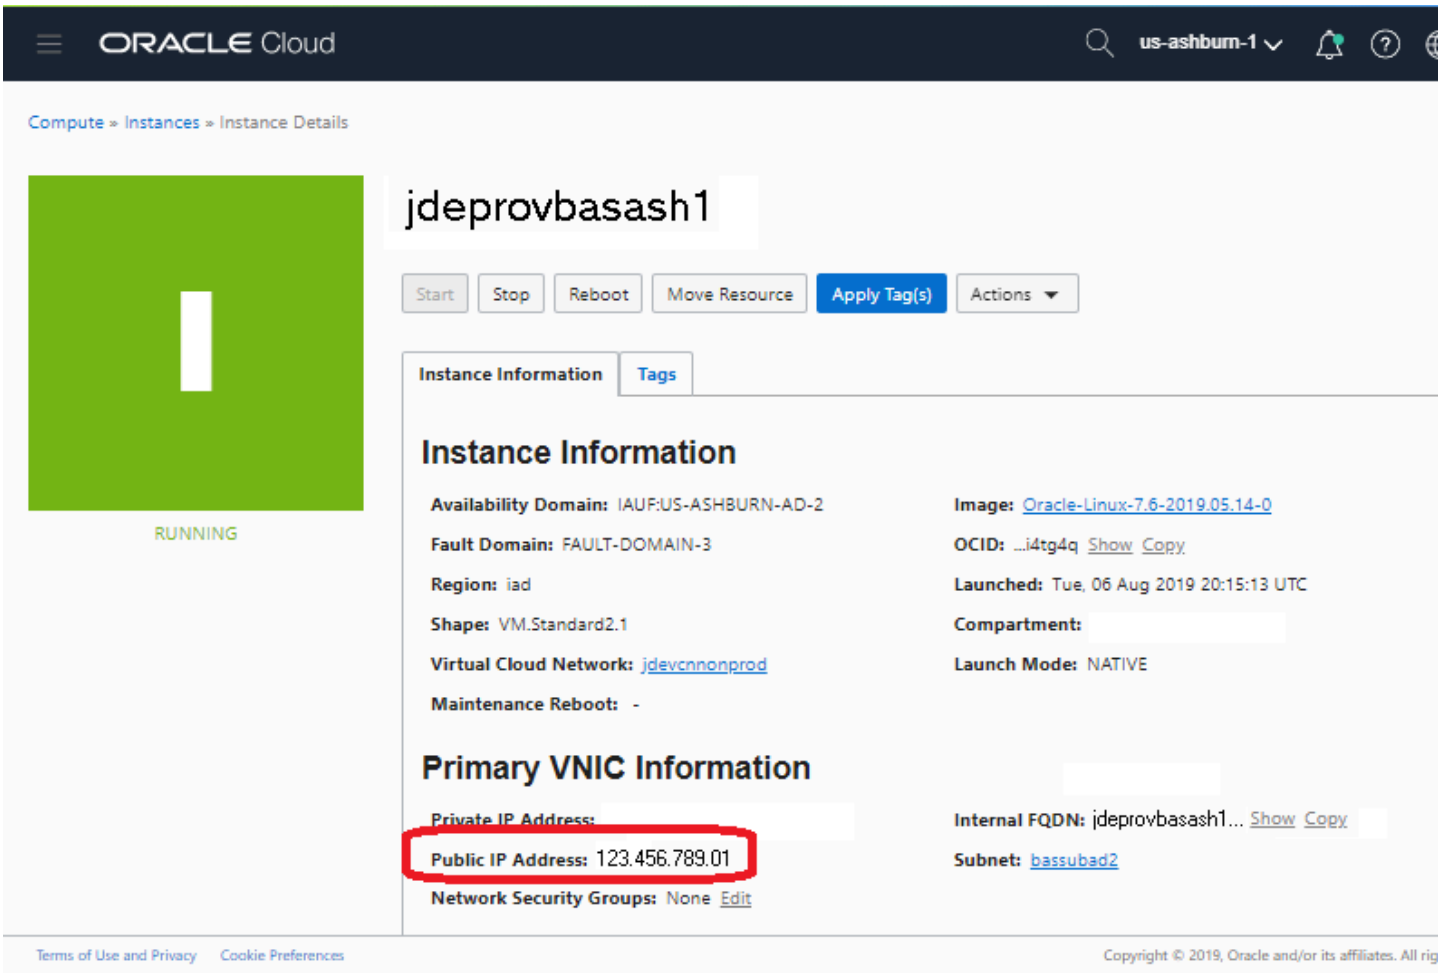

Obtaining the Private IP Address of the Provisioning Server (which also includes Server Manager)

The name of the instance for the Provisioning Server (also contains the Server Manager Console) is formed by the value you enter in the Infrastructure Provisioning Console plus the system appended value smc1.

For example, if you entered a value of jde in the Infrastructure Provisioning Console as the name of the Provisioning Server, the instance created by the OpenTofu scripts would be jdesmc1.

You can determine and obtain the private IP address of the Provisioning Server using the user interface of Oracle Cloud Infrastructure.

1. Using the information above as criteria for the instance name, in the region and compartment in which you ran the first-time setup scripts, locate and click the instance for your Provisioning Server for JD Edwards EnterpriseOne infrastructure provisioning.
2. Obtain and record the value from the Private IP Address field.

   You will need this value to access this instance.

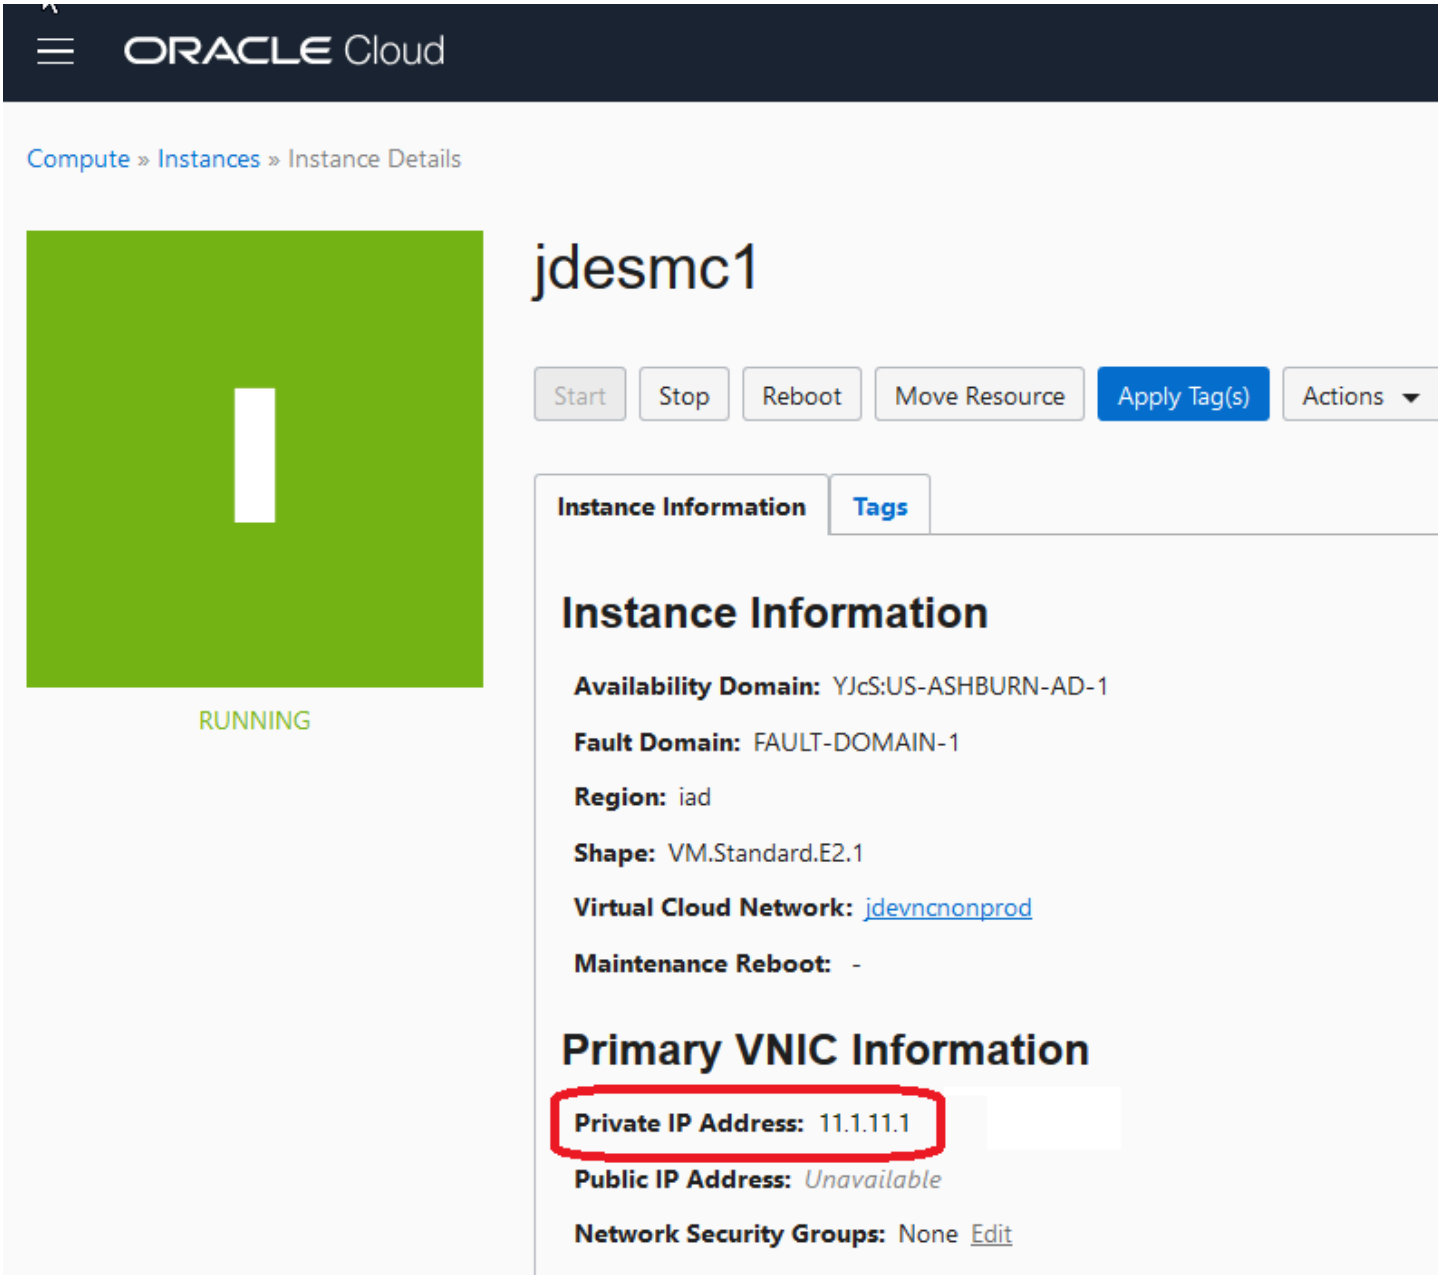

Connecting to a Provisioning Server in a Private Network Through the Bastion Host

To connect to a Provisioning Server in a private network, you must provide a private key. This key was converted from openssh format to ppk format in a preceding task in this Learning Path.

This procedure describes how to connect from a Microsoft Windows client to a Microsoft Windows host in a private network in Oracle Cloud Infrastructure through the bastion host that has been deployed during JD Edwards EnterpriseOne infrastructure provisioning.

1. On your Microsoft Windows client, search for the Pageant application (pageant.exe).

   > **Note:** As mentioned in the What Do You Need? section above, this program is a standard component of PuTTY for Microsoft Windows.

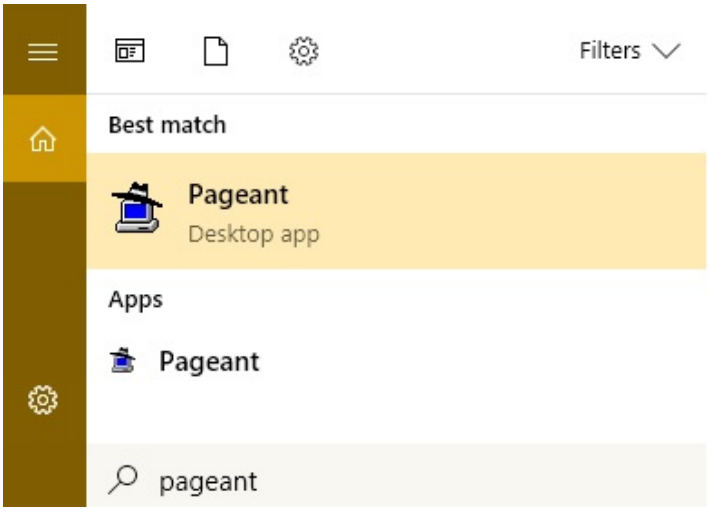

2. Start the pageant <Pageant> application to see the Pageant Key List window.

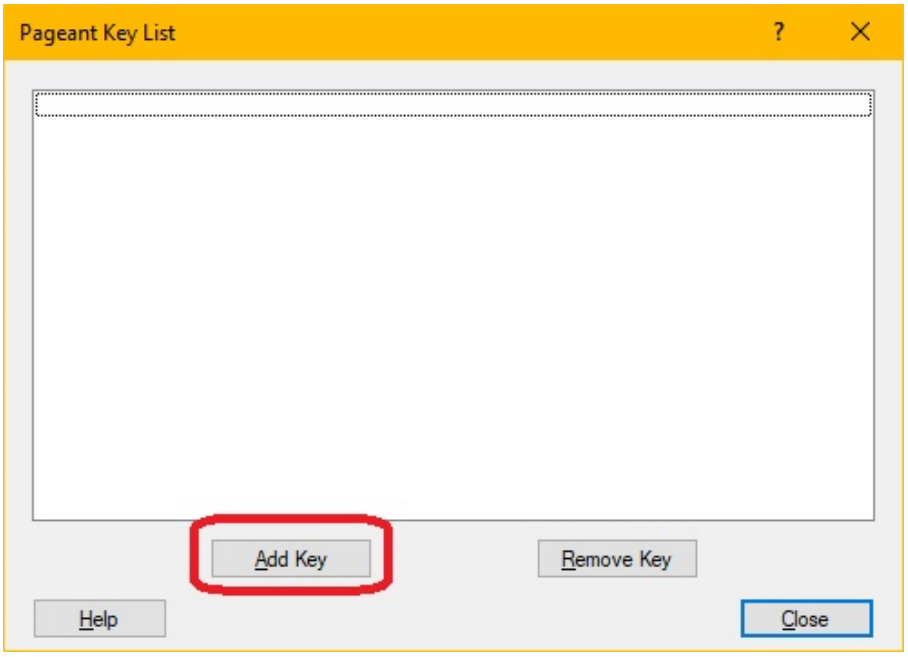

3. Click the Add Key button and browse and select the private key named OCI_Instance.ppk. This is the key that you converted to ppk format in your local workstation in the preceding task "Converting Your Private SSH Keys to .ppk Format" of this Learning Path.

   The following image shows an example of a private key in .ppk format.

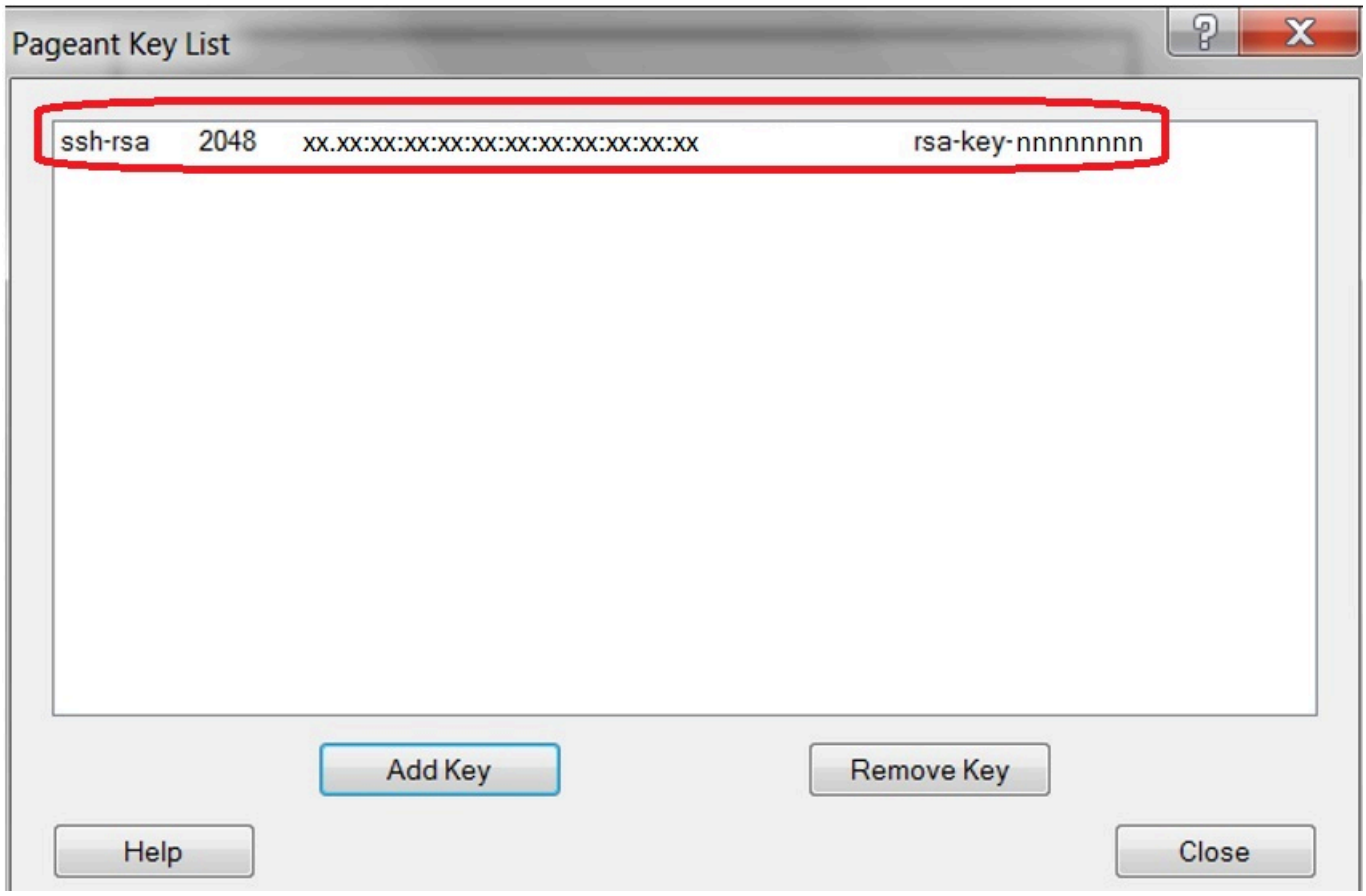

4. Open PuTTY and in the Host Name field, enter the public IP address of the bastion server.

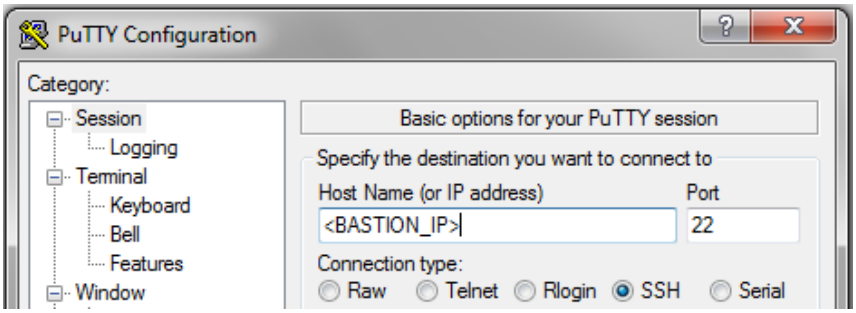

5. To log in to the connection, you need to log in as the opc user. Expand the Connection node and in the Data section, enter the value opc in the Auto-login username field.

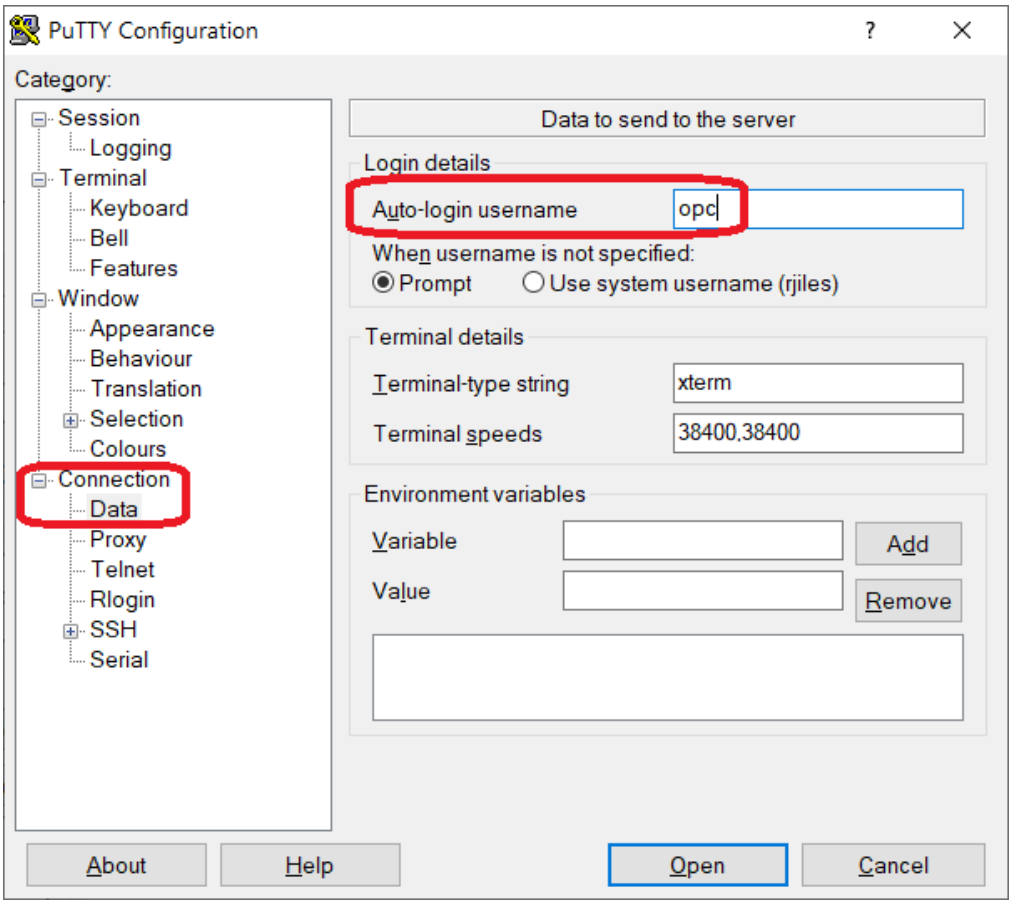

6. To create an SSH tunnel to the local host, in the Connection node, click SSH, and then click Auth.

7. In the PuTTY Configuration—Options controlling SSH authentication window, in the Authentication parameters section, ensure that the Allow agent forwarding option is selected.

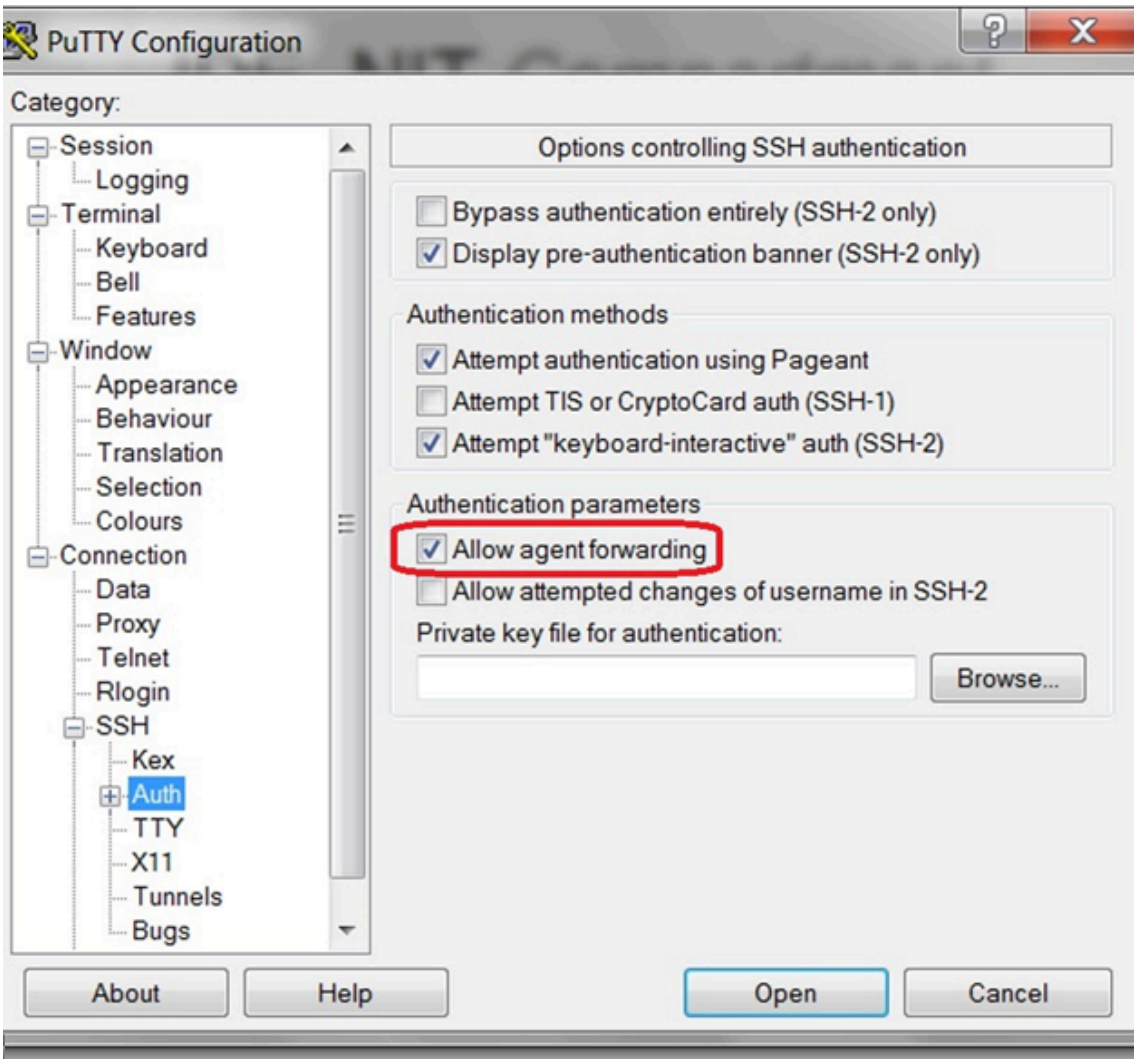

8. In the Private key file for authentication section, click the Browse button to select the Bastion.ppk key that you converted to ppk format in your local workstation by following the steps in the preceding OBE "Converting Your Private SSH Keys to .ppk Format" of this Learning Path.

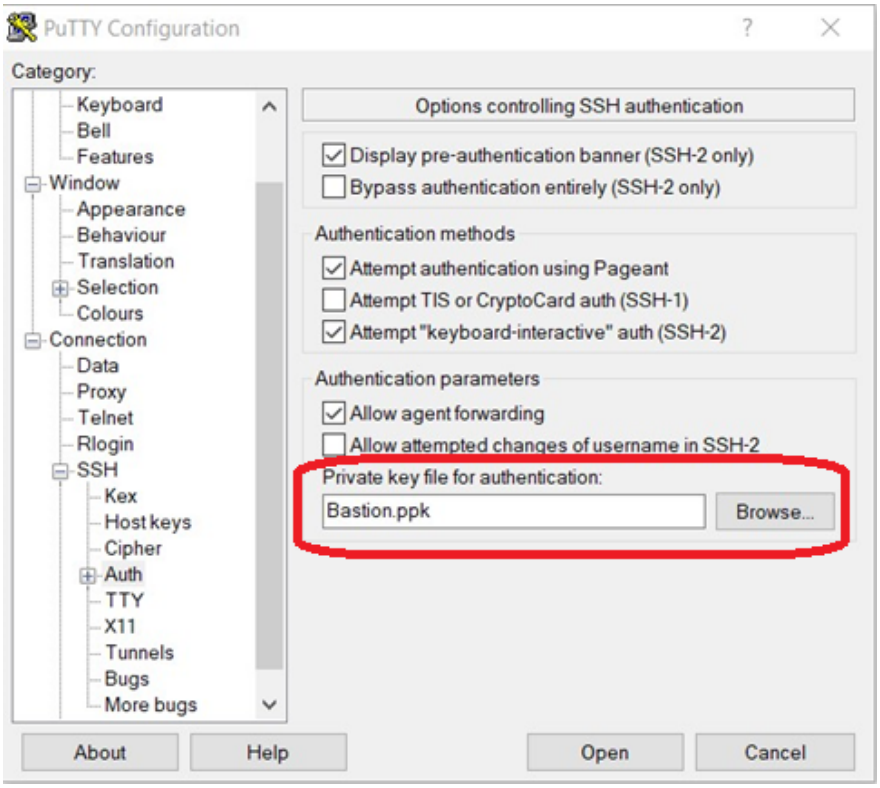

9. To create an SSH tunnel to the local host, in the Category section, in the Connections node, click SSH, and then click Tunnels.
10. In the Options controlling SSH port forwarding section, enter a port number in the Source Port field. This can be any port <You can enter the number of any port that is free on your local machine.

11. In the Destination field, enter the private IP address of the Provisioning Server followed by the port number 3000.

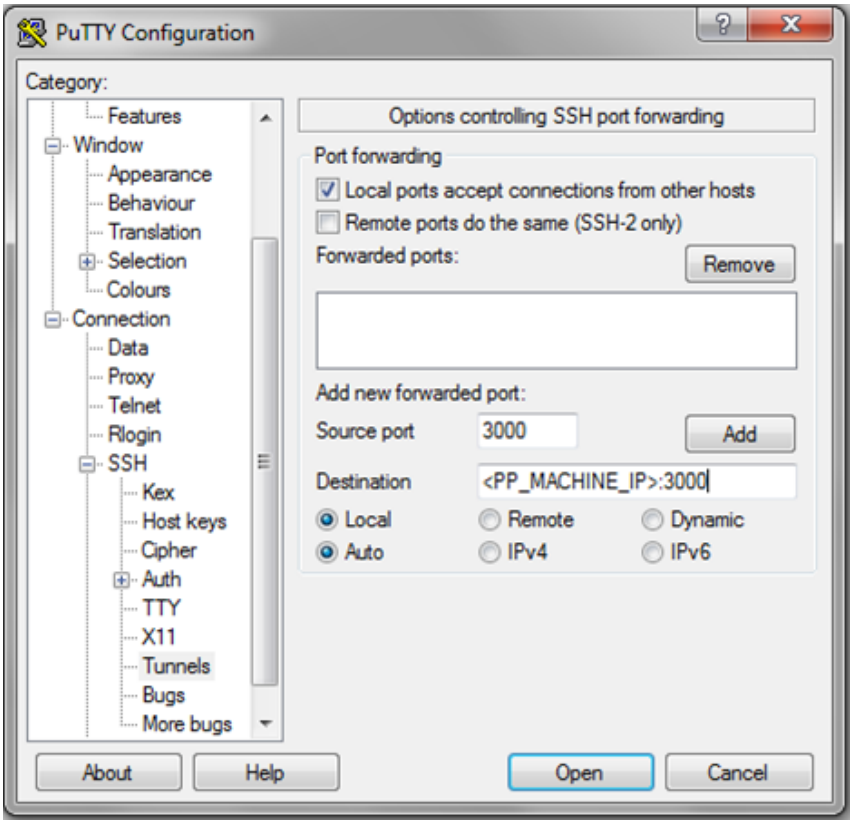

12. Click the Add button to add the port.

13. After you click the Add button, confirm that the IP address and port are added as shown below.

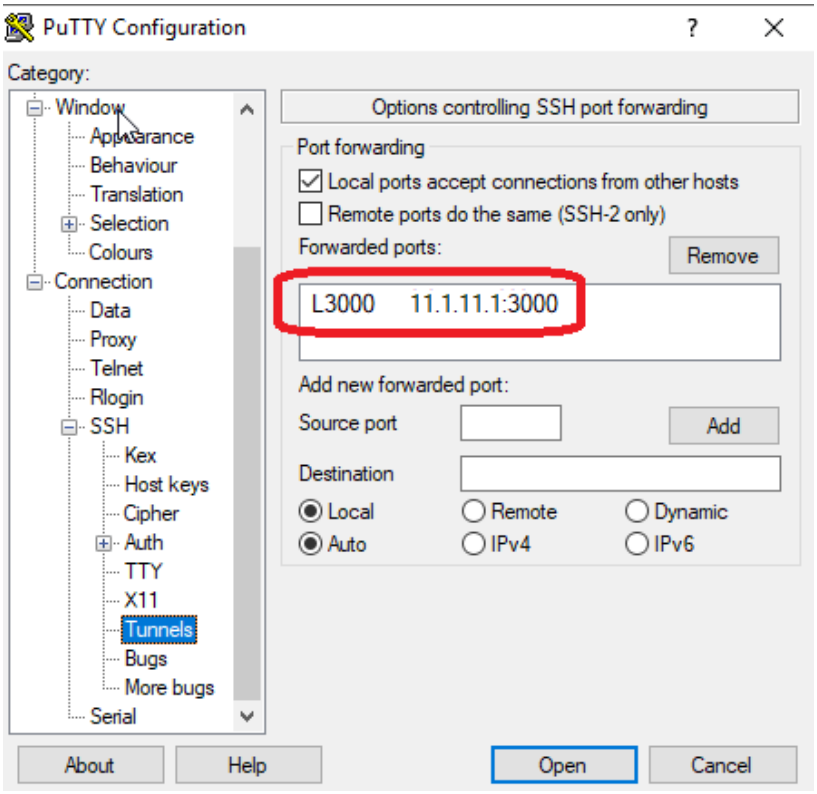

14. Ensure that these options are selected:

    Local ports accept connections from other hosts

    Local

    Auto

15. To access the Server Manager Console running on the Provisioning Server, you must repeat this procedure to add port 8998.

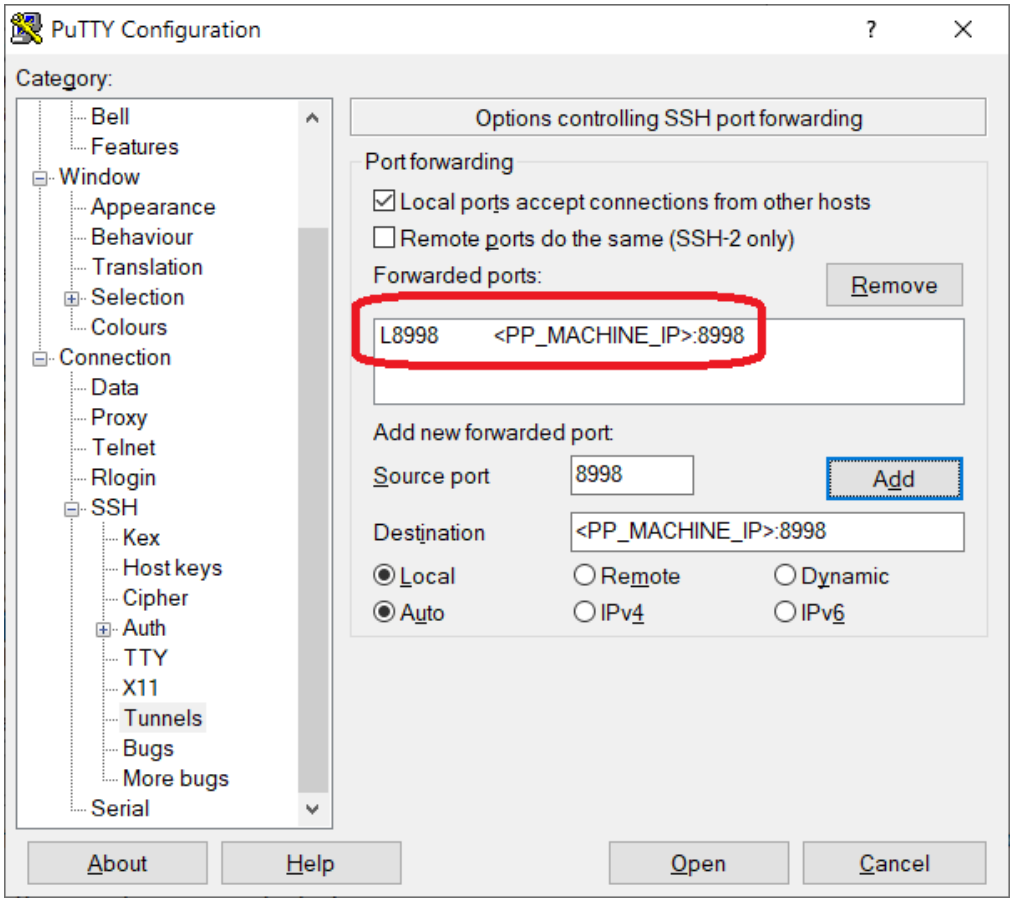

> **Note:** Before you proceed, as a best practice you should save this PuTTY session for future use when logging into machines in the private network known to this Bastion Host.

16. You can click the Open button to verify your settings. This step should result a successful connection to the Provisioning Server in a private network through the bastion host.

Accessing the JD Edwards One-Click Provisioning Console Using a Browser

The recommended browsers for accessing the JD Edwards One-Click Provisioning Console and the Server Manager Console, both of which are running on the Provisioning Server, are:

- Google Chrome
- Mozilla Firefox

Open the browser and connect to the URL for the Provisioning Server using this syntax:

```
https://localhost:3000
```

Where:

- `https` is the only supported browser protocol.
- `localhost` is used to access the Provisioning Server with a private IP address using tunneling. The Provisioning Server will not have a public IP address.

- `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 JD Edwards One-Click Provisioning Console or experience a timeout, the next time you enter the JD Edwards One-Click 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 JD Edwards One-Click Provisioning Console.

Accessing the JD Edwards Server Manager Console Using a Browser

You can access the Server Manager Console that is running on the Provisioning Server using this URL:

```
https://localhost:8998/manage
```

Note: To access the Server Manager Console, you must have opened port 8998 in the PuTTY session as described in the preceding section "Connecting to a Provisioning Server in a Private Network Through the Bastion Host" of this OBE.

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.

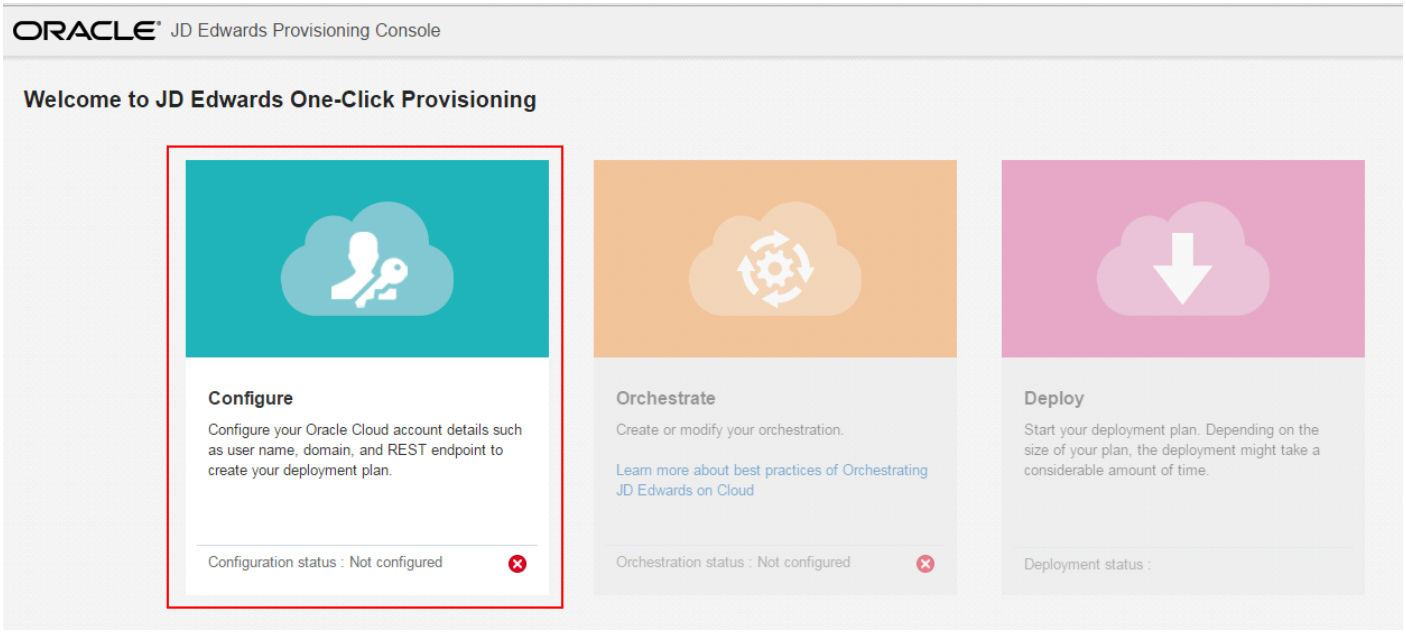

2. If you are a new user, as prompted by the JD Edwards Provisioning Console, you are also required to change the administrator passwords for WebLogic Server and Server Manager Console.

   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.

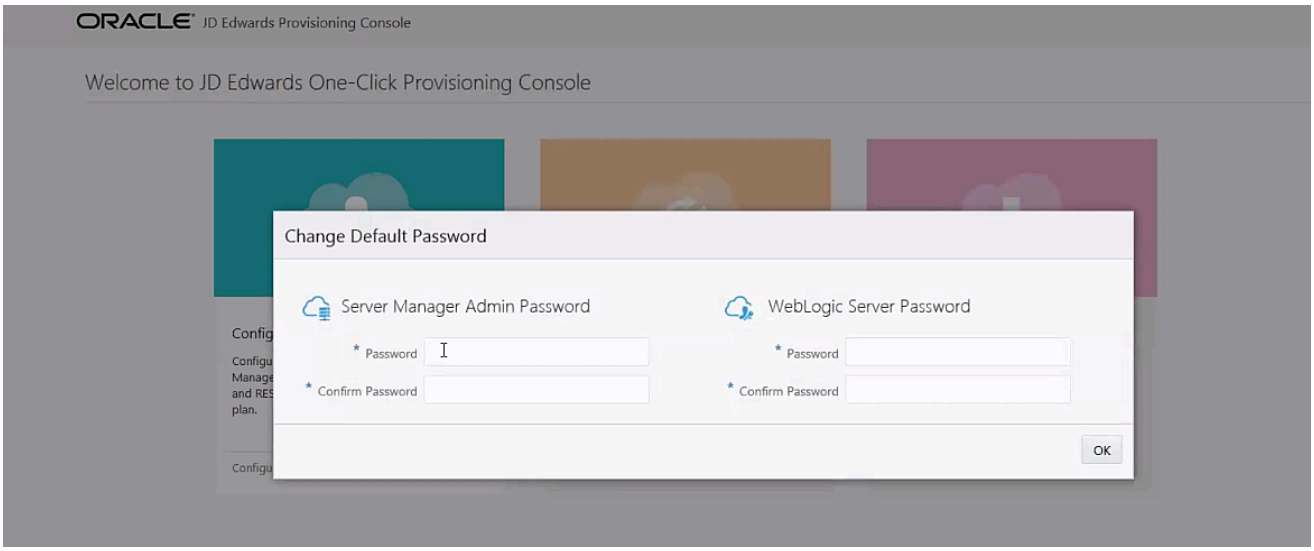

3. On the Success window, click the **OK** button.

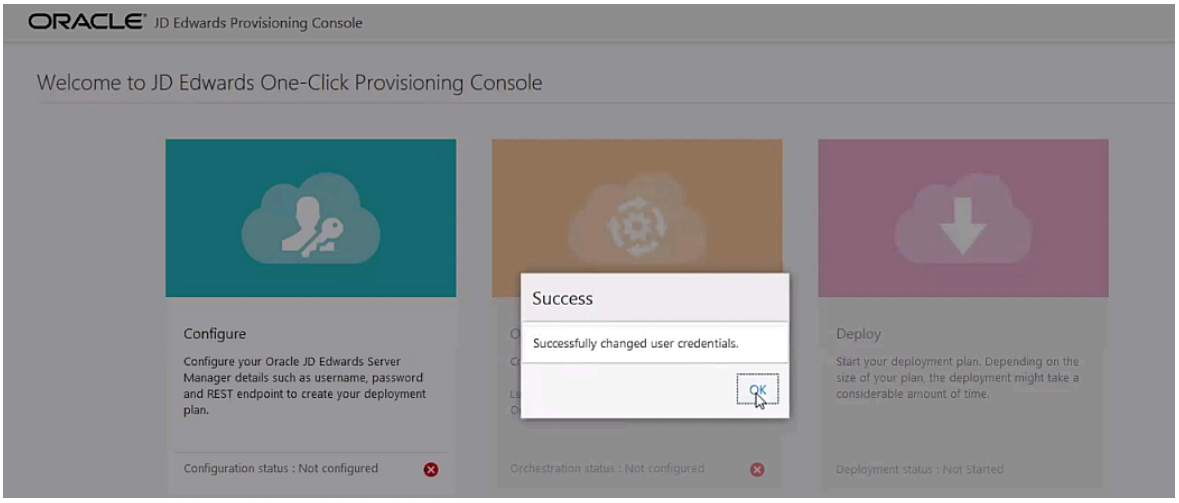

4. On Oracle JD Edwards Server Manager Details, reenter the Admin password for the Server Manager Console. The SMC Endpoint is pre-populated automatically.

> **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**.

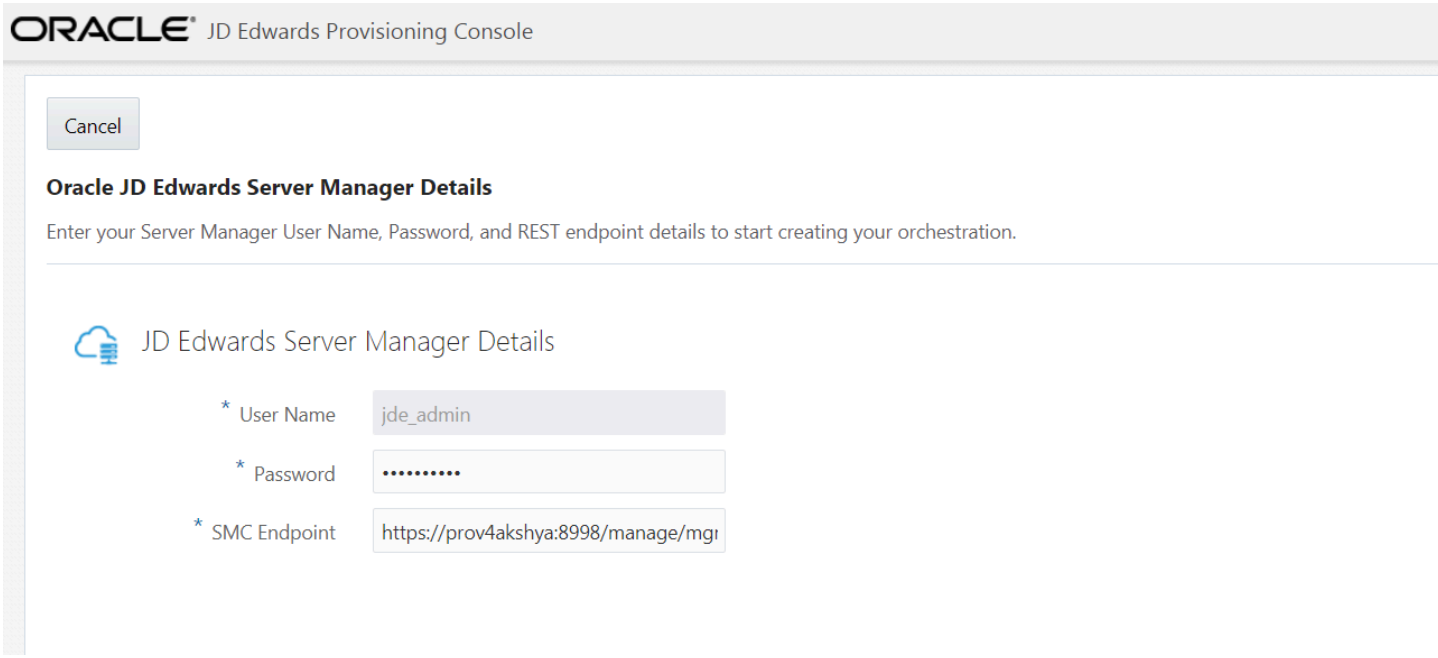

5. Click the **Save Configuration** button.

   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.

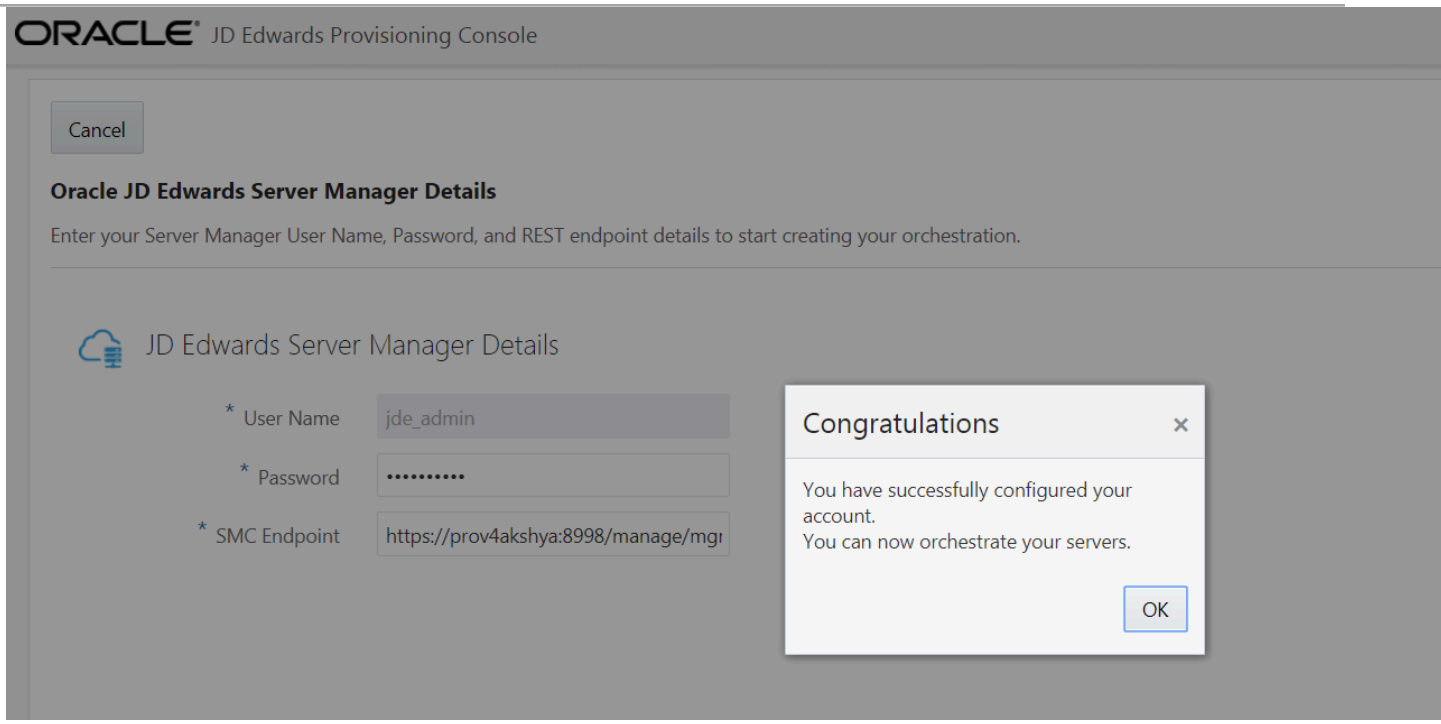

Automated Deployment of WebLogic Server as an Image

This quick tour explains how to use One-Click Infrastructure Provisioning to automate the deployment of WebLogic Server as an image.

Automated Deployment of WebLogic Server as an Image

Orchestrating Servers with the One-Click Provisioning Console

This section shows you how to orchestrate servers with the JD Edwards One-Click Provisioning Console.

After the infrastructure provisioning is completed, the JD Edwards Infrastructure Provisioning Console creates json files in this path:

`/home/opc/JDERefArch_InfraProvisioning/E1InfraProvisionConsole/outputJson`

- The **`infraOutput.json`** file contains details about the production and non-production pathcodes as well as disaster recovery environment, if it was selected during infrastructure provisioning. This file also contains

information about the public and private IP addresses and host name of the bastion and other provisioned instances.

- The `pdOrch.json` file contains details of all the instances in the production environment, as well as disaster recovery environment, if it was selected during infrastructure provisioning. This file can be imported into the One-Click Provisioning Console.
- The `nonpdOrch.json` file contains details of all the instances in the non-production environment. It is not possible to import this json file into the One-Click Provisioning Console. You can use the details in this file to orchestrate <import> the components manually in the One-Click Provisioning Console.

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 Servers with the One-Click Provisioning Console

> **Note:** Instances for non-production environment are not automatically orchestrated on the One-Click Provisioning Console. Instead, you must manually orchestrate instances for non-production environments that were created through the Infrastructure Provisioning Console. For instructions, refer to the OBE "Orchestrating Using Advanced Mode" in this Learning Path.You can get details of non-production environments that require manual orchestration at this location on the OpenTofu Staging Server: `/home/opc/JDERefArch_InfraProvisioning/E1InfraProvisionConsole/outputJson/nonpdOrch.json` Use this procedure to edit orchestrated servers for production environments with the One-Click Provisioning Console:

1. On the JD Edwards Provisioning Console, select the Orchestrate tile.

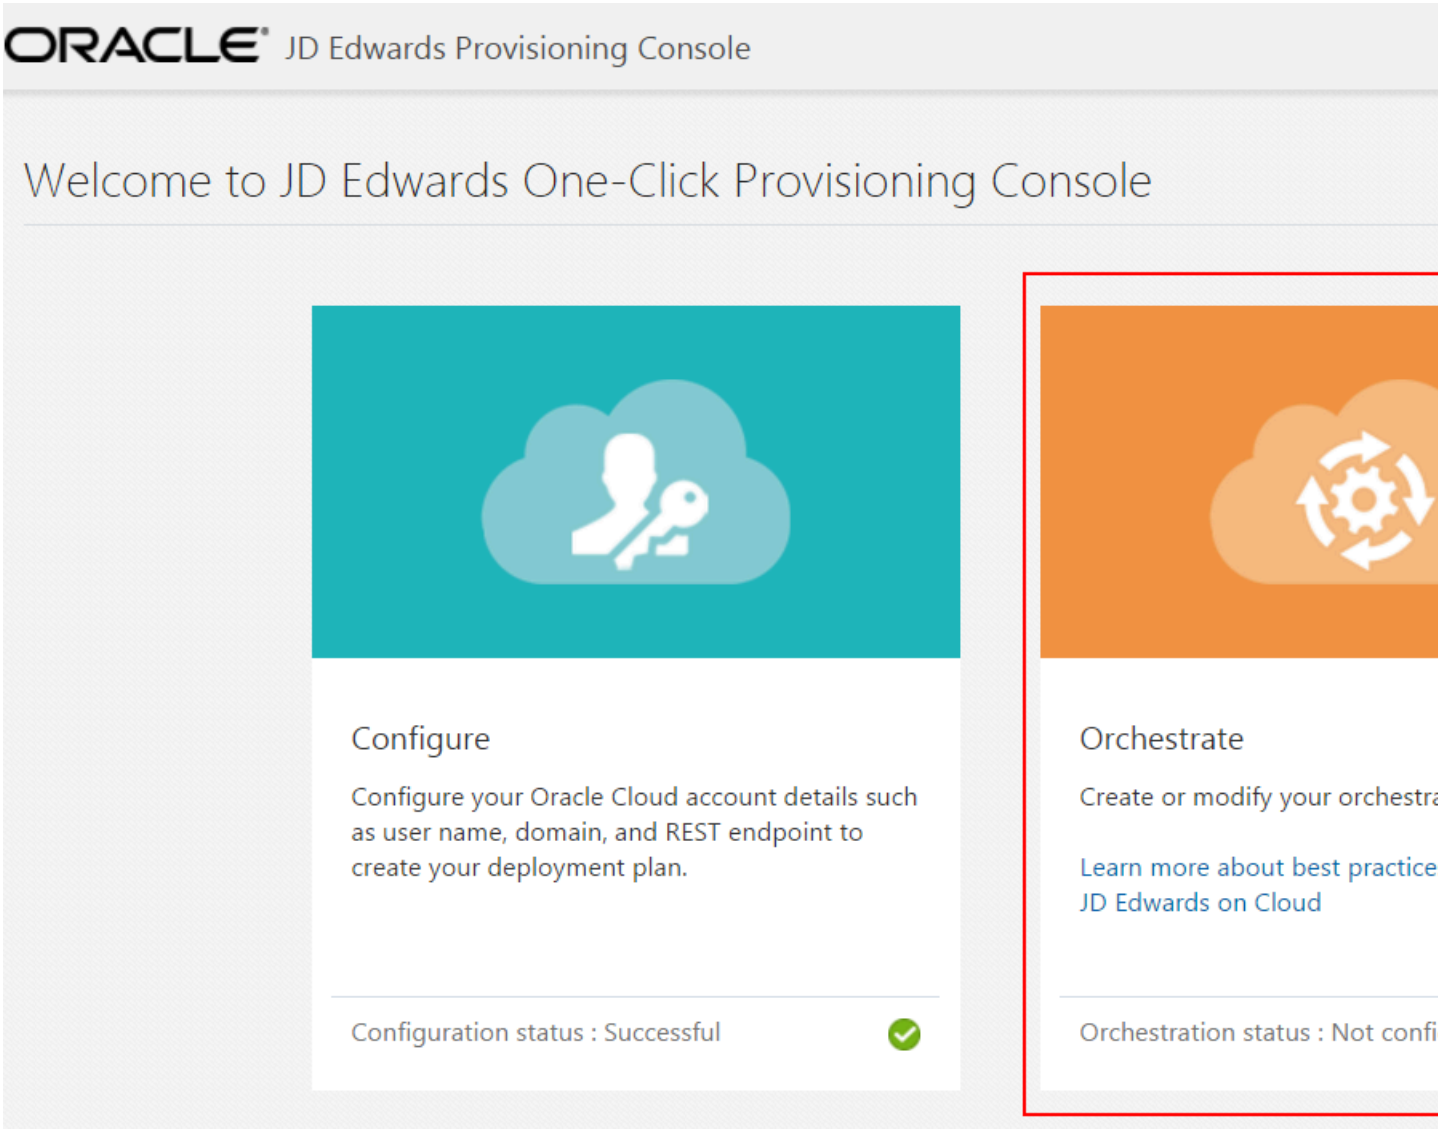

2. If this is the first pass through a JD Edwards One-Click Provisioning orchestration, the following Global Settings screen will appear first. In the **SSH Private Key** section, click the **View/Edit Private Key** button.

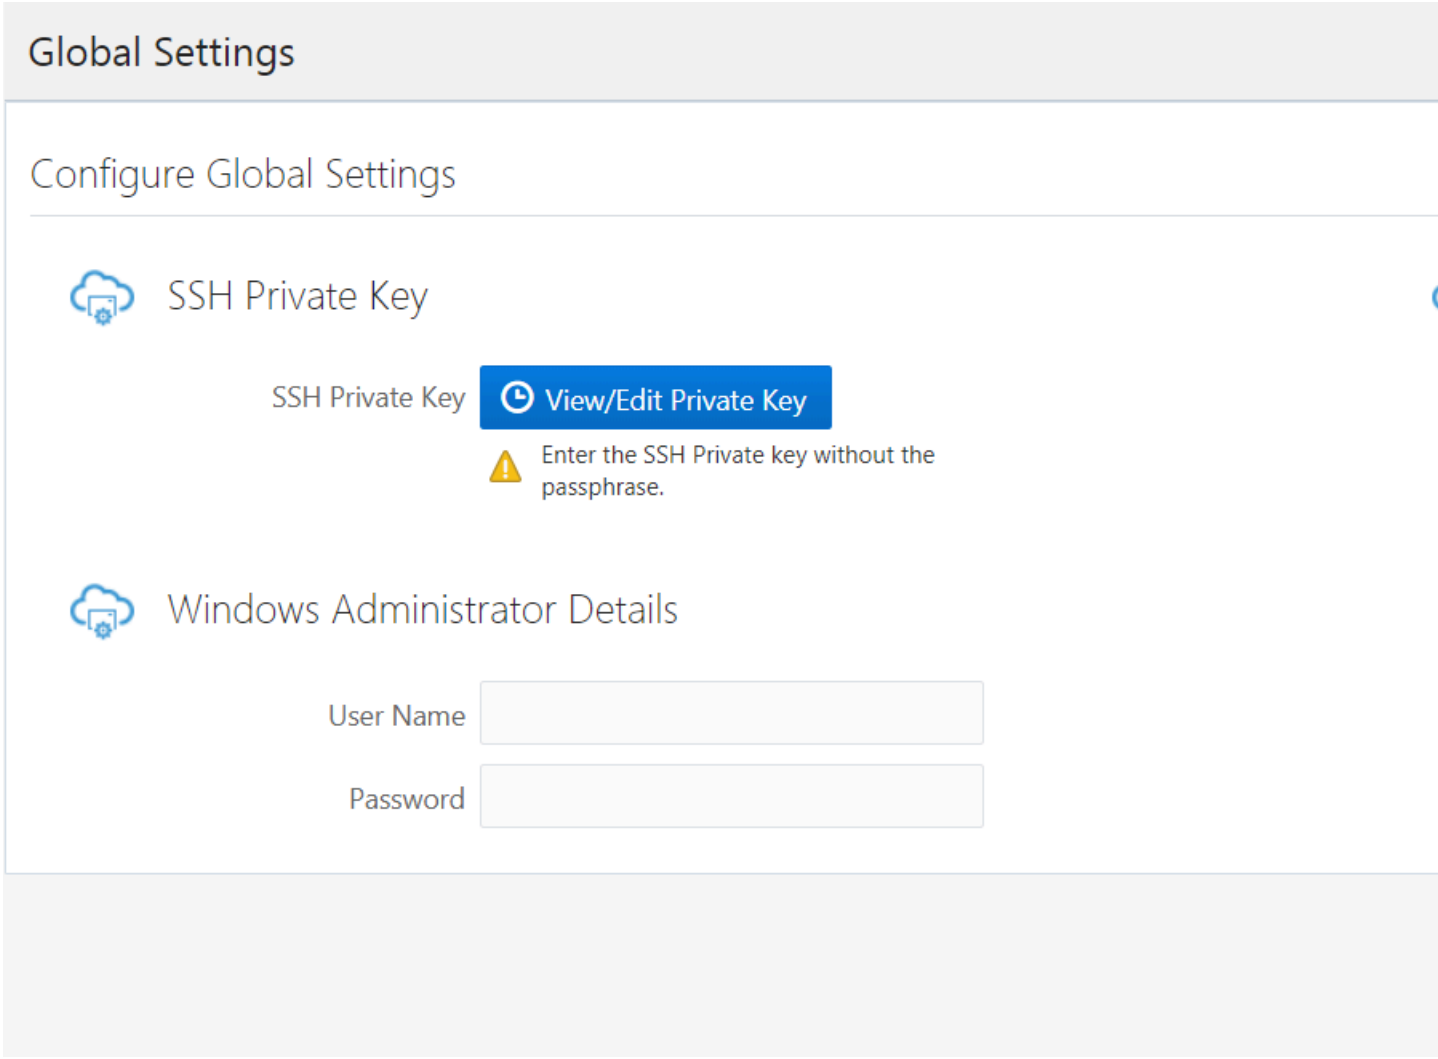

3. On the pop-up window **Private Key input for VM access**, you must specify the values for either the SSH private key text, or browse and select the file that contains the SSH private key contents for accessing all the instances that are provisioned using this tool.

   Select the **SSH Private Key File** option and then click the **Choose File** button. For more information regarding SSH keys, refer to the OBE "Generate Secure SHell (SSH) Key Pairs on Your Local System" of this Learning Path.

   > **Note:** Ensure that the SSH Private Key field is not empty. The One-Click Provisioning Console validates the private keys and it is not possible to save the Global Settings if this field is empty.

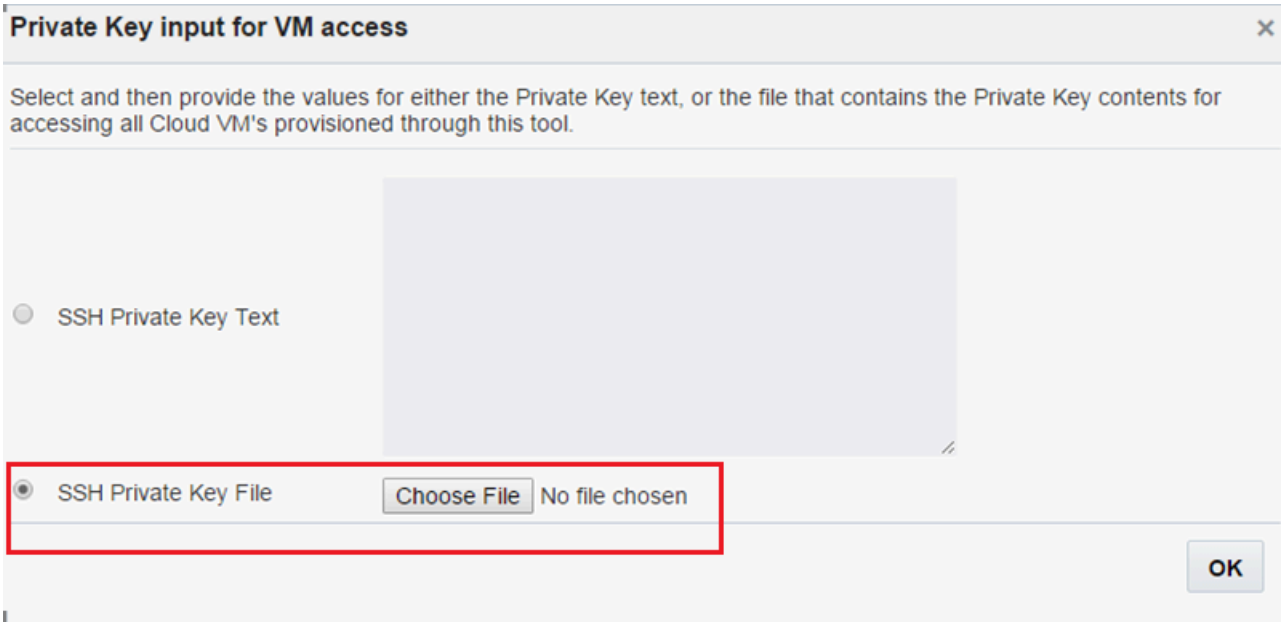

4. Browse and select the appropriate file, and then click OK. For example: OCI_Instance.openssh.

> **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.

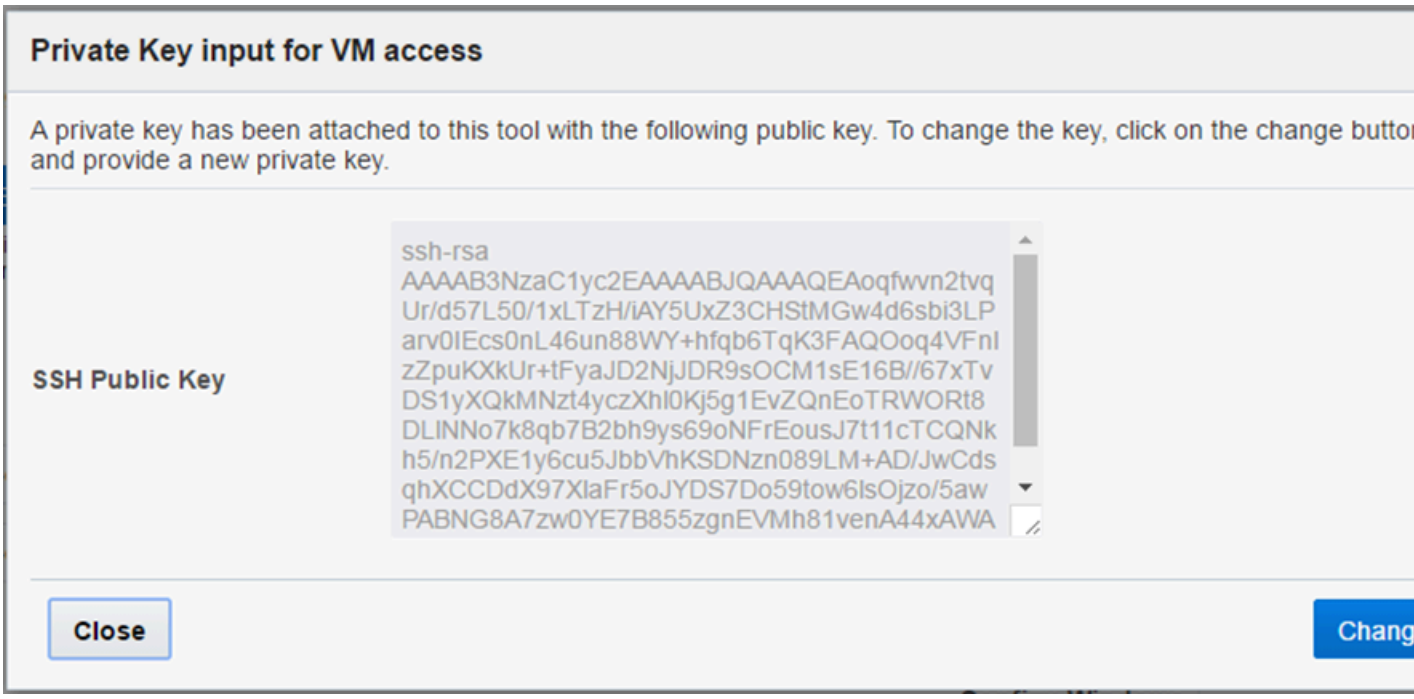

5. Click the **Close** button.
6. In the Windows Administrator Details section, enter the Windows user name and password. Ensure that the user name is entered as **opc** and that this user has the Administrator privileges.

> **Note:** You must enter the same password for this Windows Server that you previously specified for the Deployment Server in the section of this OBE entitled: "Completing Configuration in the JD Edwards Infrastructure Provisioning Console" of this Learning Path.

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 provided by JD Edwards Tools Release 9.2.4.5, the password must be between 12 and 30 characters. It can contain only alphanumeric characters, 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 an alphanumeric character, must be between 8 and 40 characters, and contain at least 2 uppercase letters, 2 lowercase letters, 2 numbers, and 2 underscore characters.

> **Tip:** The conditions to set the passwords appear in the tooltip window when you click the fields.

> **Note:** It is highly recommended that you keep a record of these critical passwords. If you have not already done so, you should ensure these values are recorded on the ***Pre-Install Worksheet*** .

8. Click the **Save** button to exit the Global Settings screen.

9. On Choose your Orchestration Template, click **Advanced**.

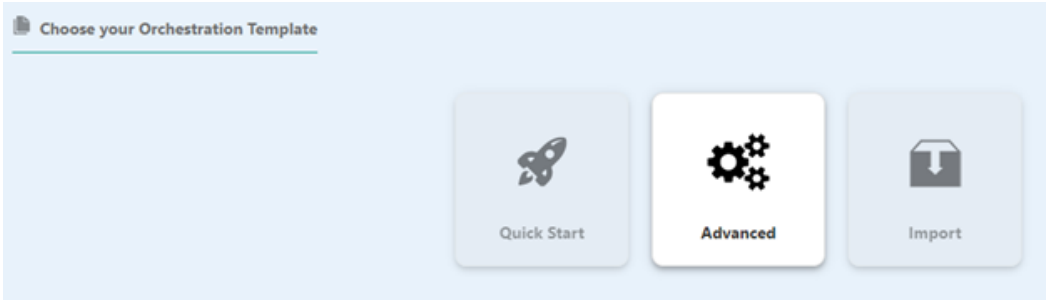

10. Click on Orchestrate to modify the JD Edwards Deployment Plan.

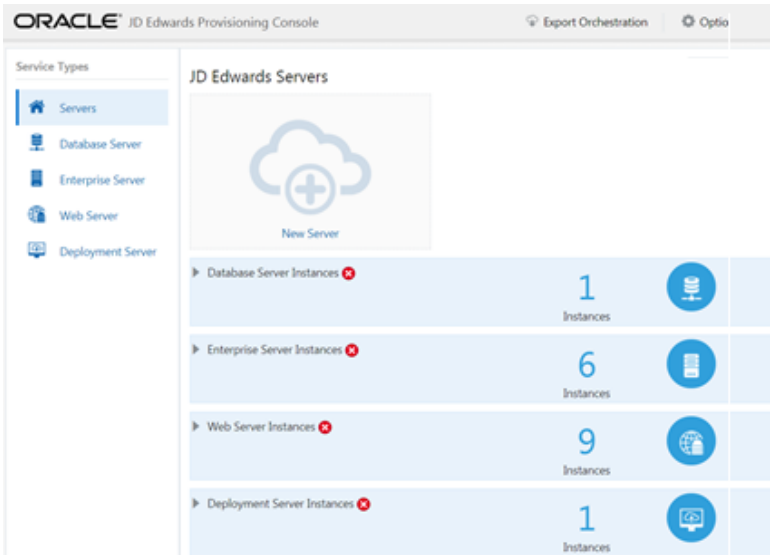

11. As shown above, the system displays a list of orchestrated servers with an error (red X icon) if any user input is required. To remedy, you must edit the instance.

12. Click on the orchestrated instance and click the **Edit** function.

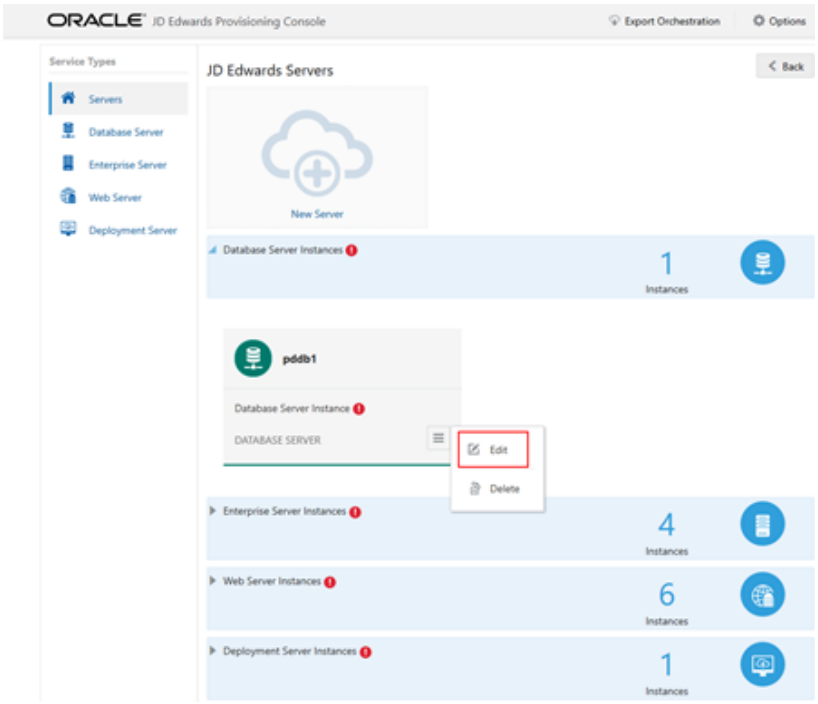

13. For each orchestrated server enter the required inputs:

Database Server

- DB Admin Password

  > **Note:** This password must be the same as the password you entered for the Database Server in the section of this Learning Path entitled: "Completing Configuration in the JD Edwards Infrastructure Provisioning Console".

- JDE DB Install Directory (for example, /u01/DataDB)
- JDE DB Table Directory (for example, /u01/ORATABLE)
- JDE DB Index Directory (for example, /u01/ORAINDEX)

  > **Note:** These are example paths for when Use ASM feature is disabled. If **Use ASM feature** is enabled see the following document: E1: OCI: What Significance Does The 'Use ASM Feature' Option Have In One-click Provisioning? (Doc ID 2880000.1)

Enterprise Server

To configure the Enterprise Server, complete the following sections:

Server Configuration

- *Platform*

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

  Create an instance name for the Enterprise Server. The conditions to set the instance name are 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:** 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.

  - *HA enabled*

    Turn on this toggle button if you want to provision multiple Enterprise Servers on same pathcode using a virtual hostname.

Web Server

- *WebLogic Password*

  > **Note:** This password must be the same as the password you entered for the WebLogic Server in the section of this Learning Path entitled: "Completing Configuration in the JD Edwards Infrastructure Provisioning Console".

- *HA enabled* Turn on this toggle button if you want to provision multiple Enterprise Servers on same pathcode using a virtual hostname.

  Add correct dependencies while orchestrating servers:

  - for htm4a1 instance, dependent server is logic1
  - for ais1 instance, dependent server is htm4a1
  - for htm1 instance, as it is standalone JAS server, you need to toggle button for standalone JAS and set dependent AIS server with ais1

Deployment Server

- Location

  > **Note:** 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).

14. For each orchestrated server that you edit, save the instance details.

15. After you have finished editing and saving the instance details, click on the **Back** button at the top right of the form. Then you can proceed to the following section entitled: "Deploying an Orchestration".

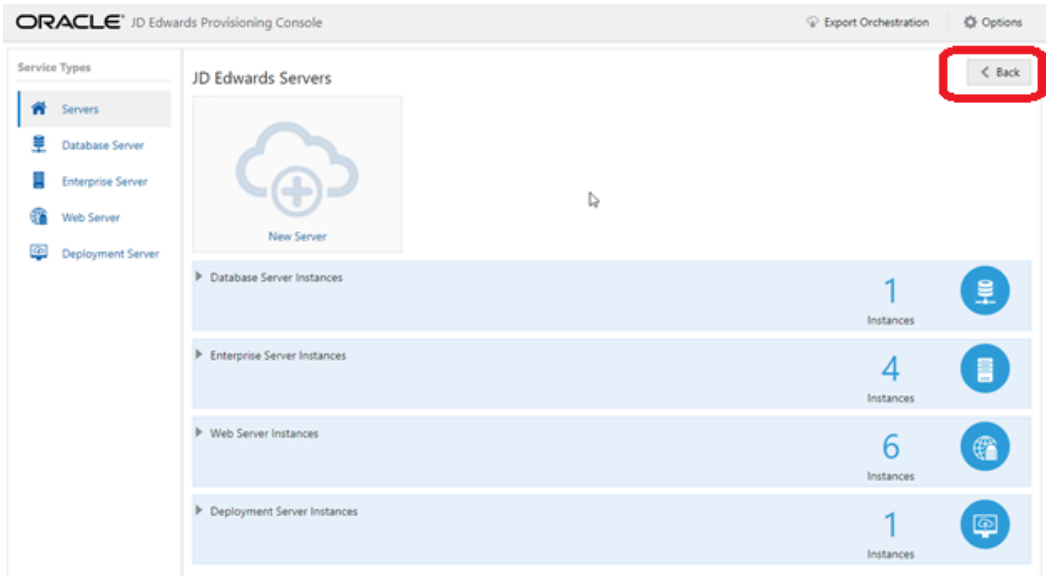

Orchestrating Using Advanced Mode

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.

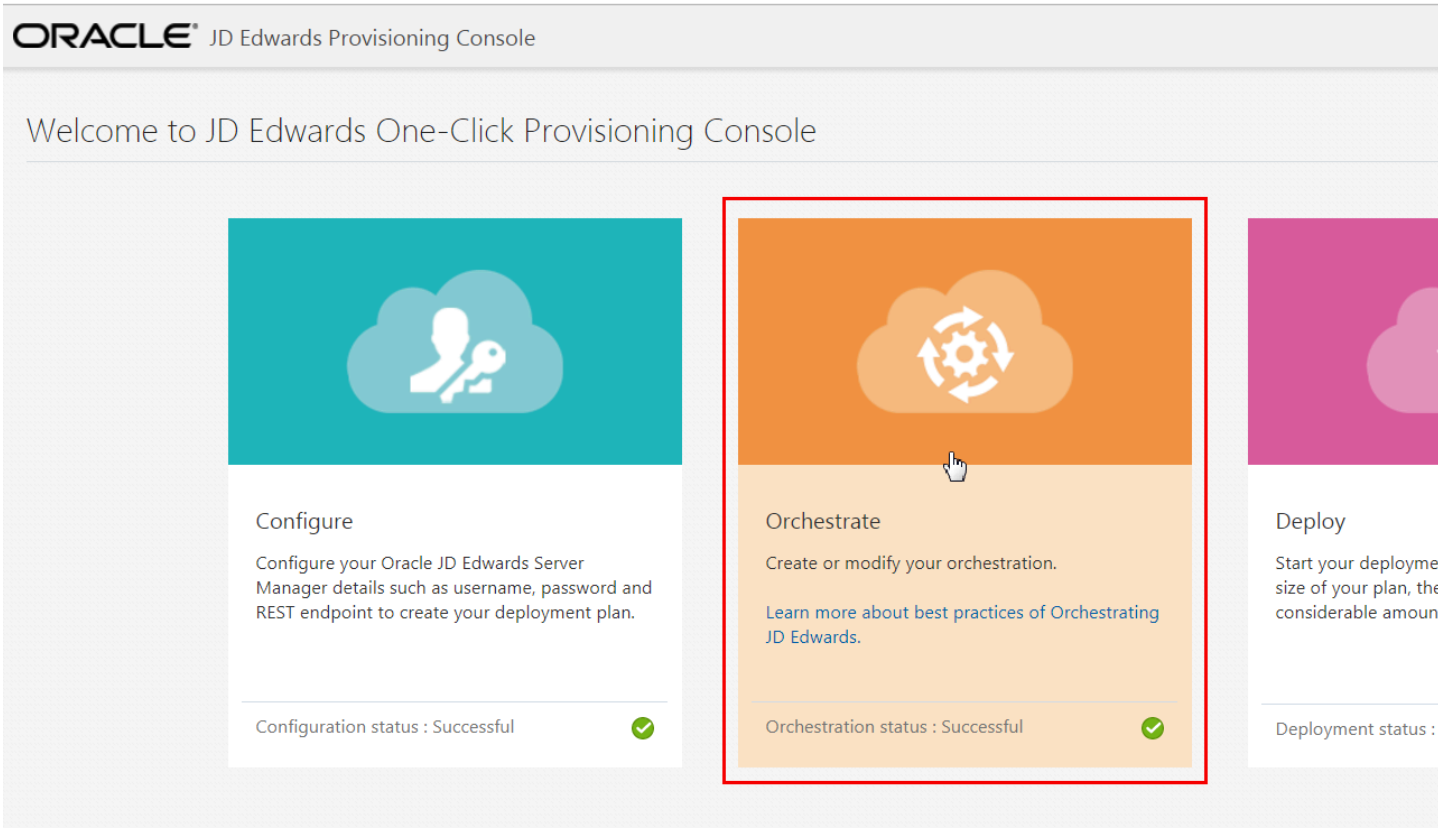

2. On JD Edwards Servers, click on the instance of an existing server, or click the **New Server** icon to add a new JD Edwards service. Alternately you can select the desired server from **Service Types** on the left tab, and then click **New Server** to add a New Server.

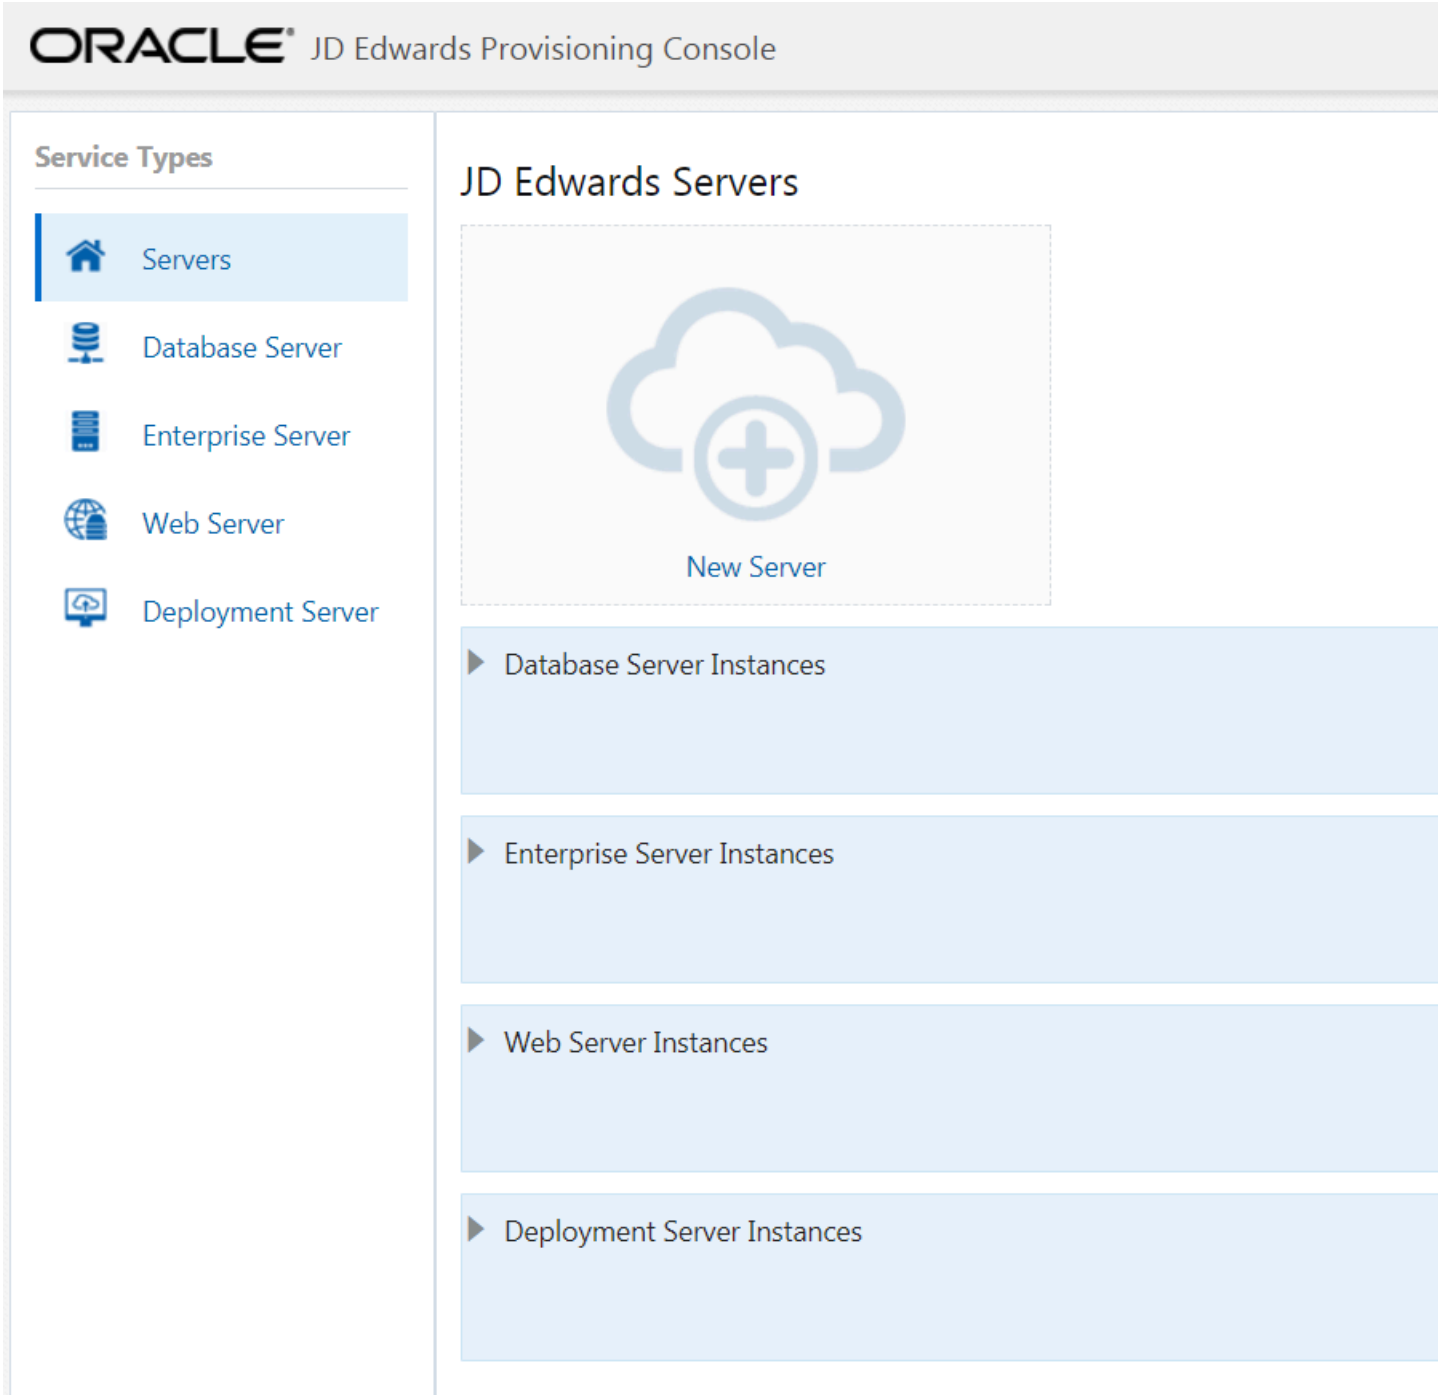

Note: The remainder of this procedure includes steps to either add (where allowed) or modify (existing) these instances:

- 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.

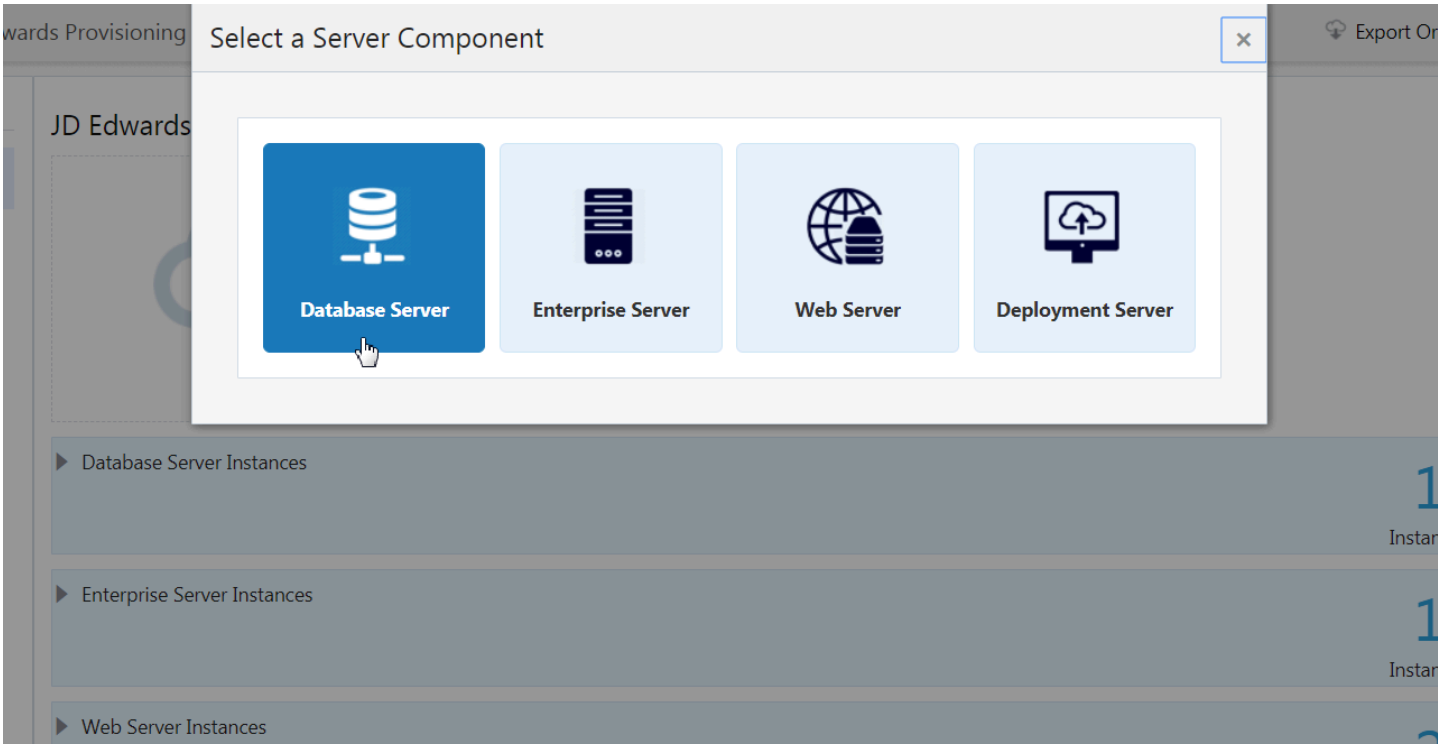

2. On Database Server Instance page, complete the following fields to create and configure the Database Server instance.

   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.

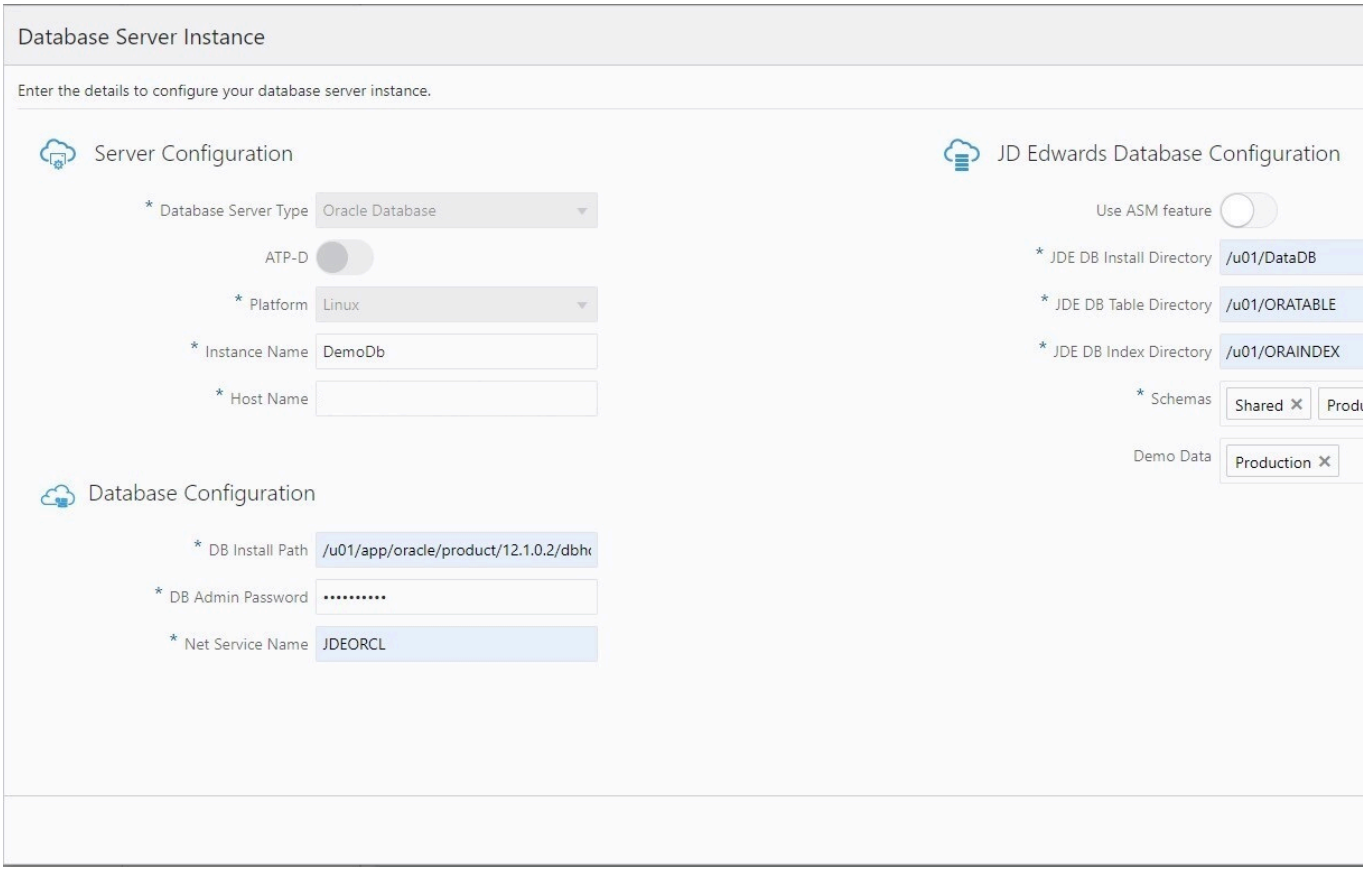

Oracle Autonomous Database

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.

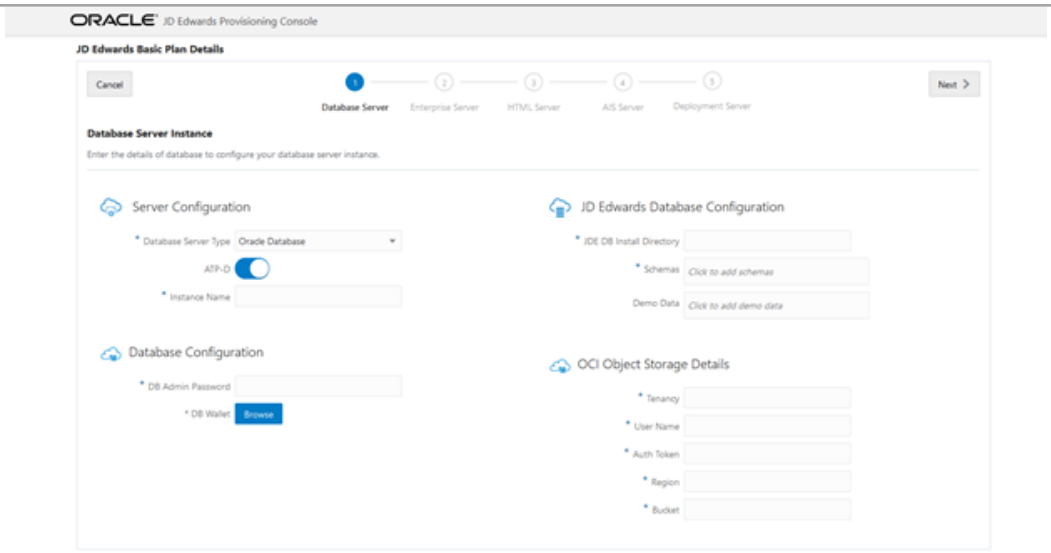

3. You should now be able to view multiple instances of the Database Server.

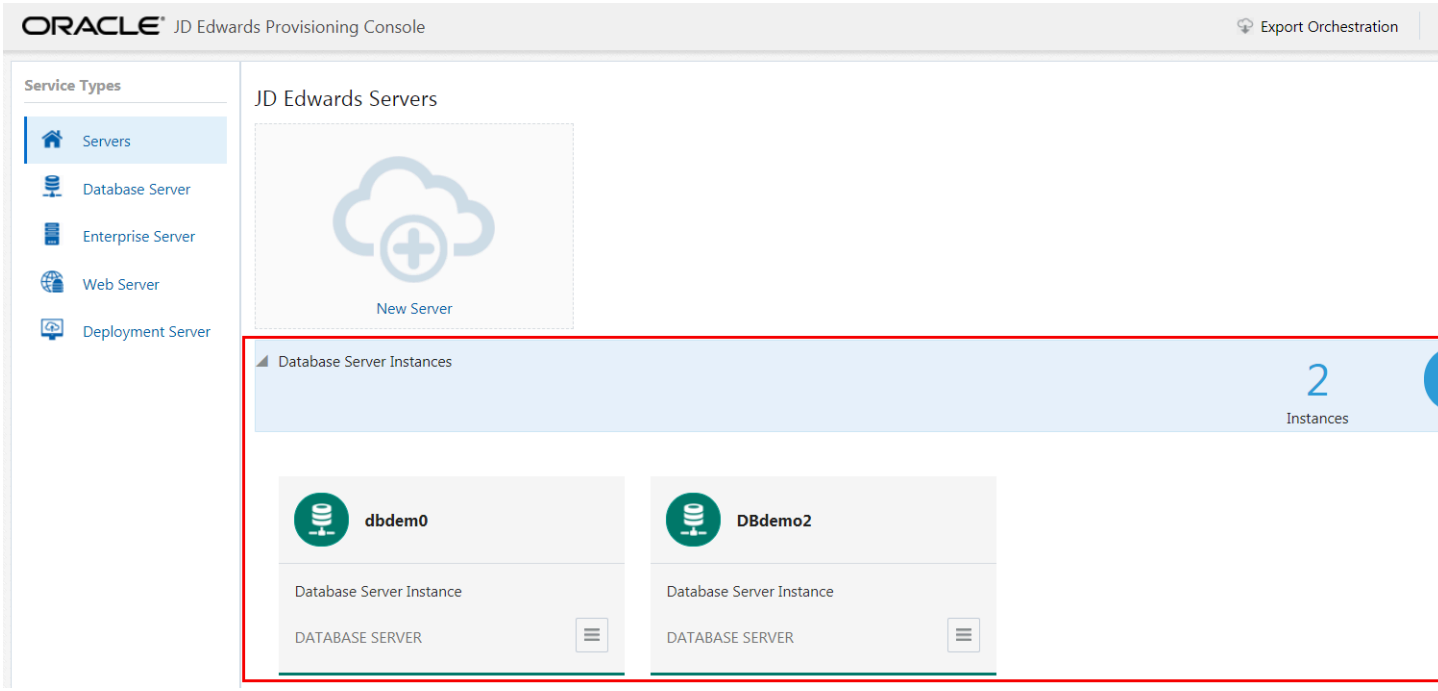

4. If a Database Server exists, click **View** from the Application Options tab to view the existing configuration for the Database Server. To modify the instance configuration use the Application Options tab and choose the **Edit** option.

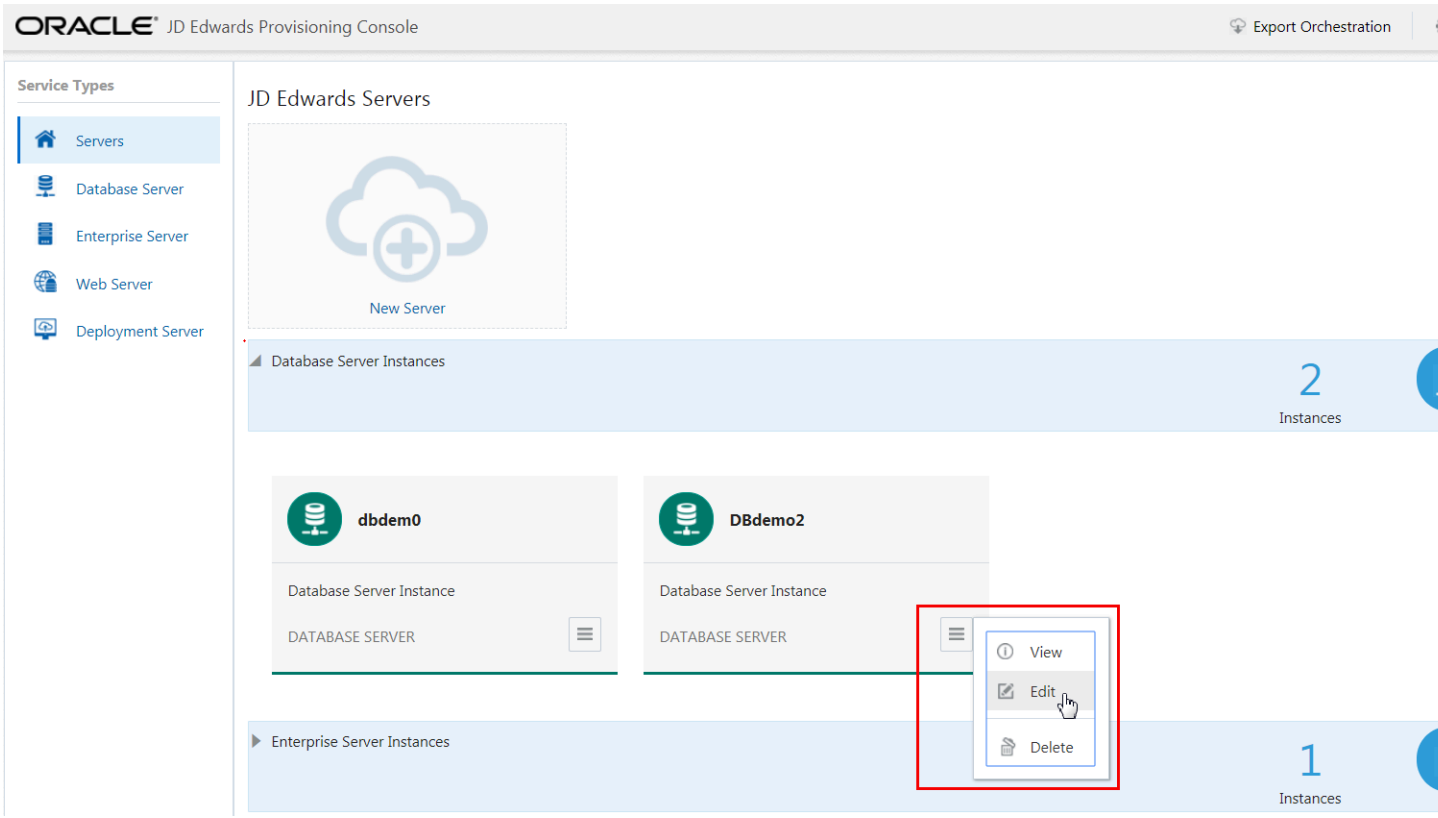

5. To delete the instance, use the Application Options tab to select **Delete.**

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.

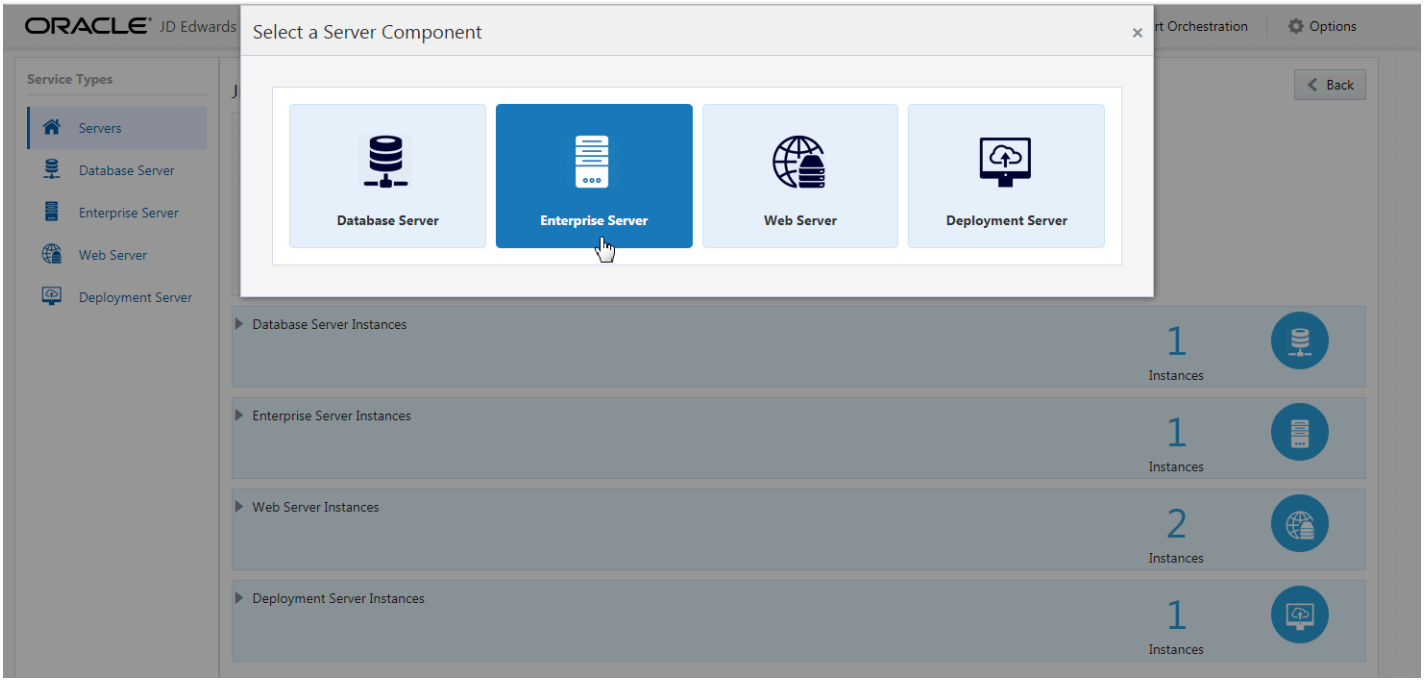

2. On Enterprise Server Instance page, complete these fields to create and configure the Enterprise Server instance.

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.

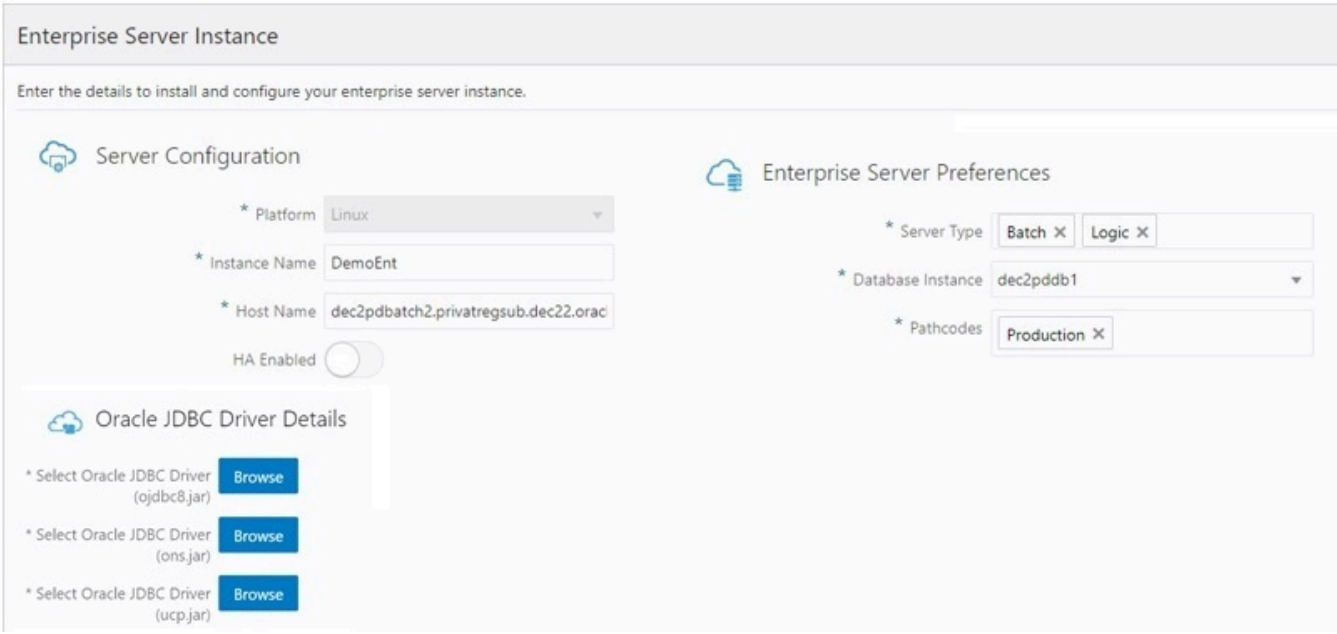

3. You should now be able to view multiple instances of the Enterprise Server.

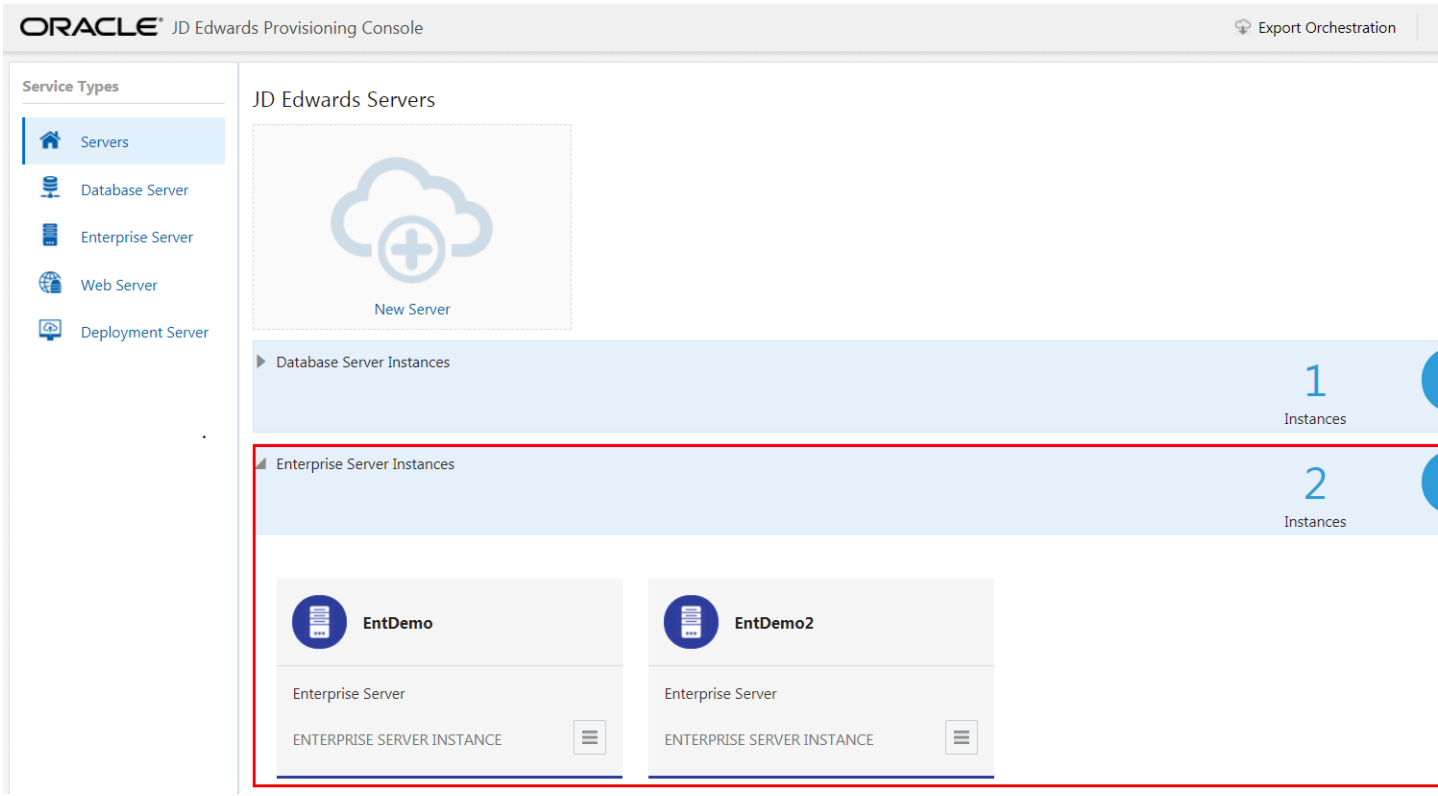

4. If an Enterprise Server exists, click **View** from the Application Options tab to view the existing configuration for the Enterprise Server. To modify the instance configuration use the Application Option tab and choose the **Edit** option. Use the Clone option to **clone** the Enterprise Server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

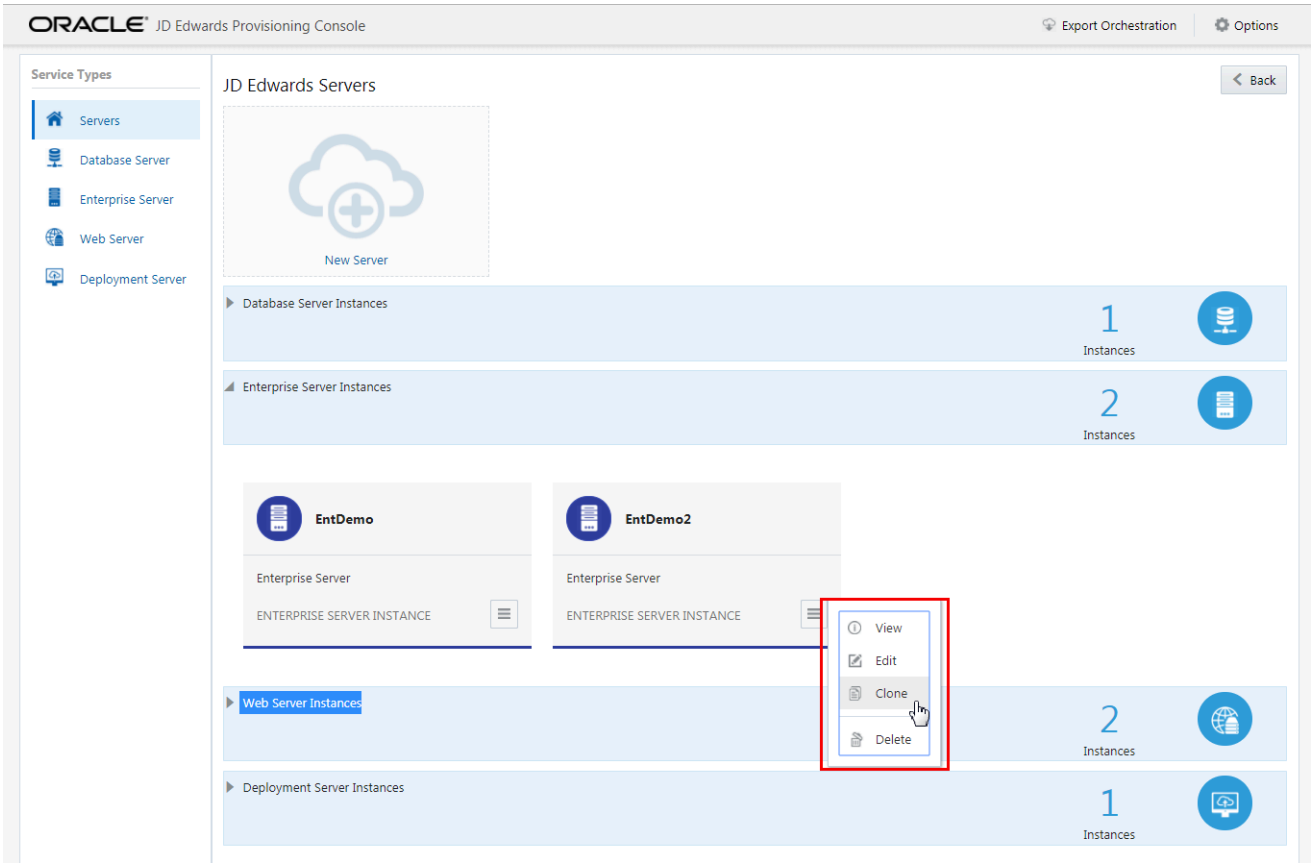

5. To delete the instance, use the Application option tab to select **Delete**.

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.

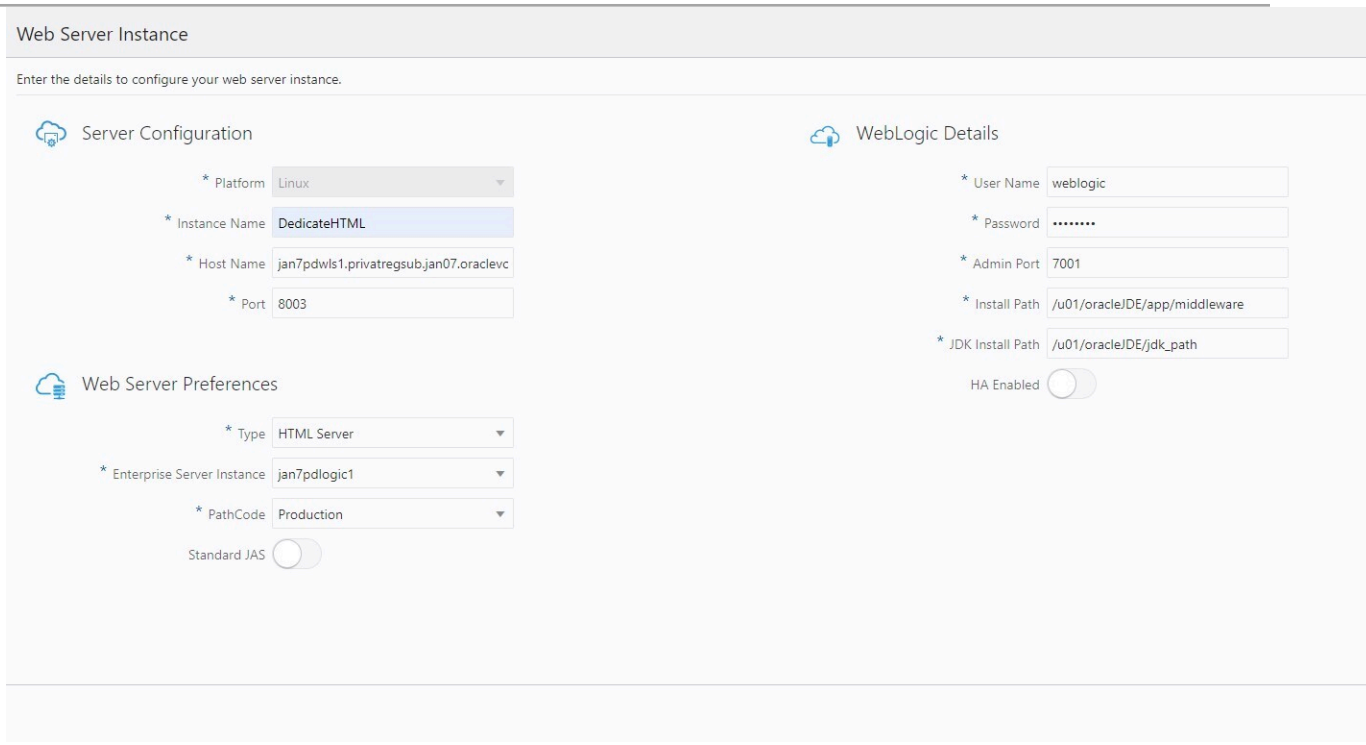
Web Server Instance
Enter the details to configure your web server instance.
Server Configuration
* Platform Linux
* Instance Name DedicateHTML
* Host Name jan7pdwls1.privatregsub.jan07.oraclevc
* Port 8003
WebLogic Details
* User Name weblogic
* Password ••••••••
* Admin Port 7001
* Install Path /u01/oracleJDE/app/middleware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

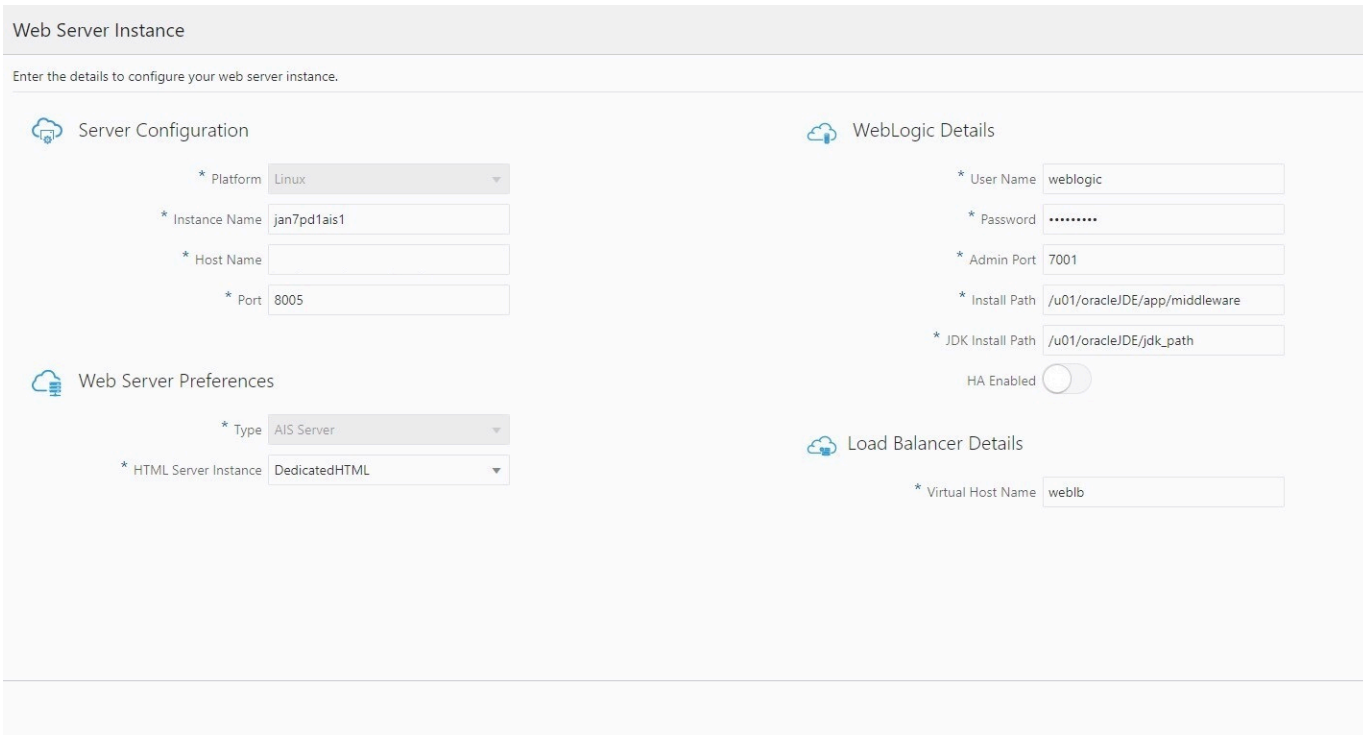
Web Server Instance
Enter the details to configure your web server instance.
Server Configuration
* Platform Linux
* Instance Name jan7pd1ais1
* Host Name
* Port 8005
WebLogic Details
* User Name weblogic
* Password ••••••••••
* Admin Port 7001
* Install Path /u01/oracleJDE/app/middleware
* JDK Install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

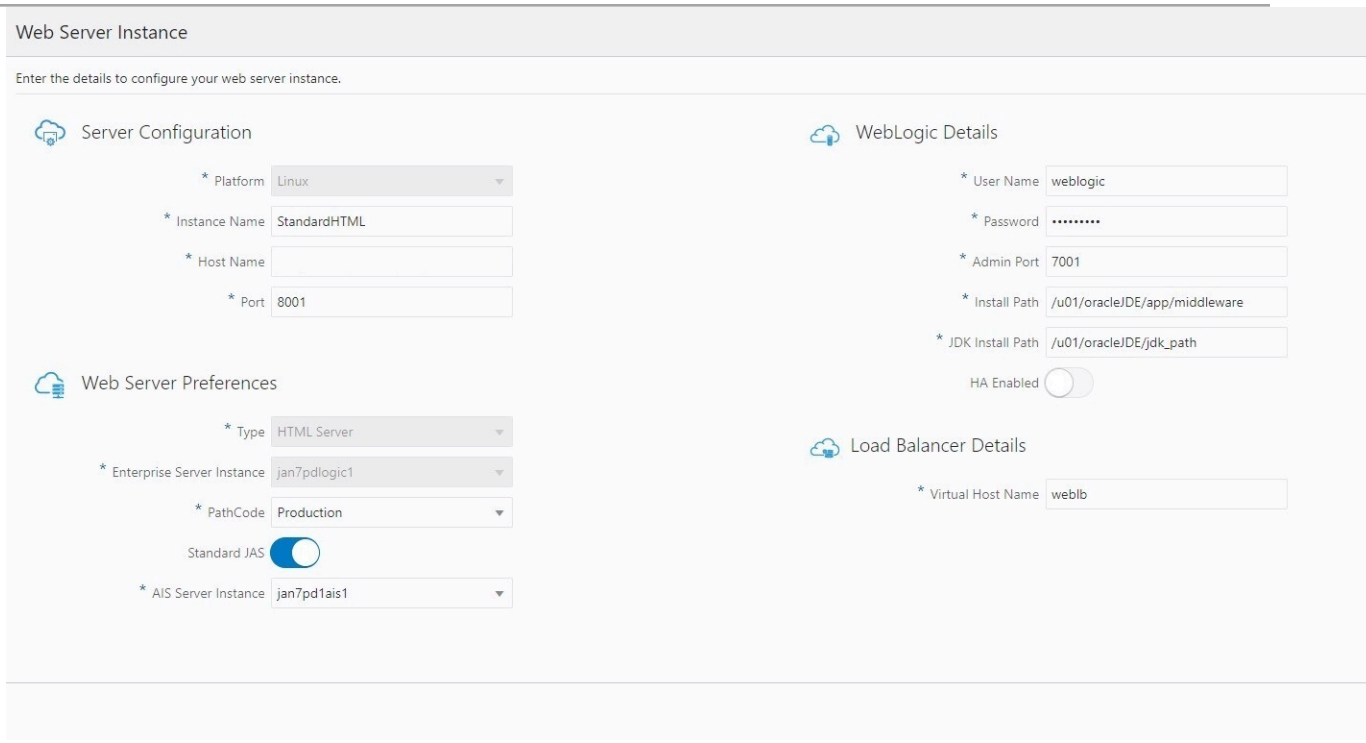

4. Click the **OK** button.

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

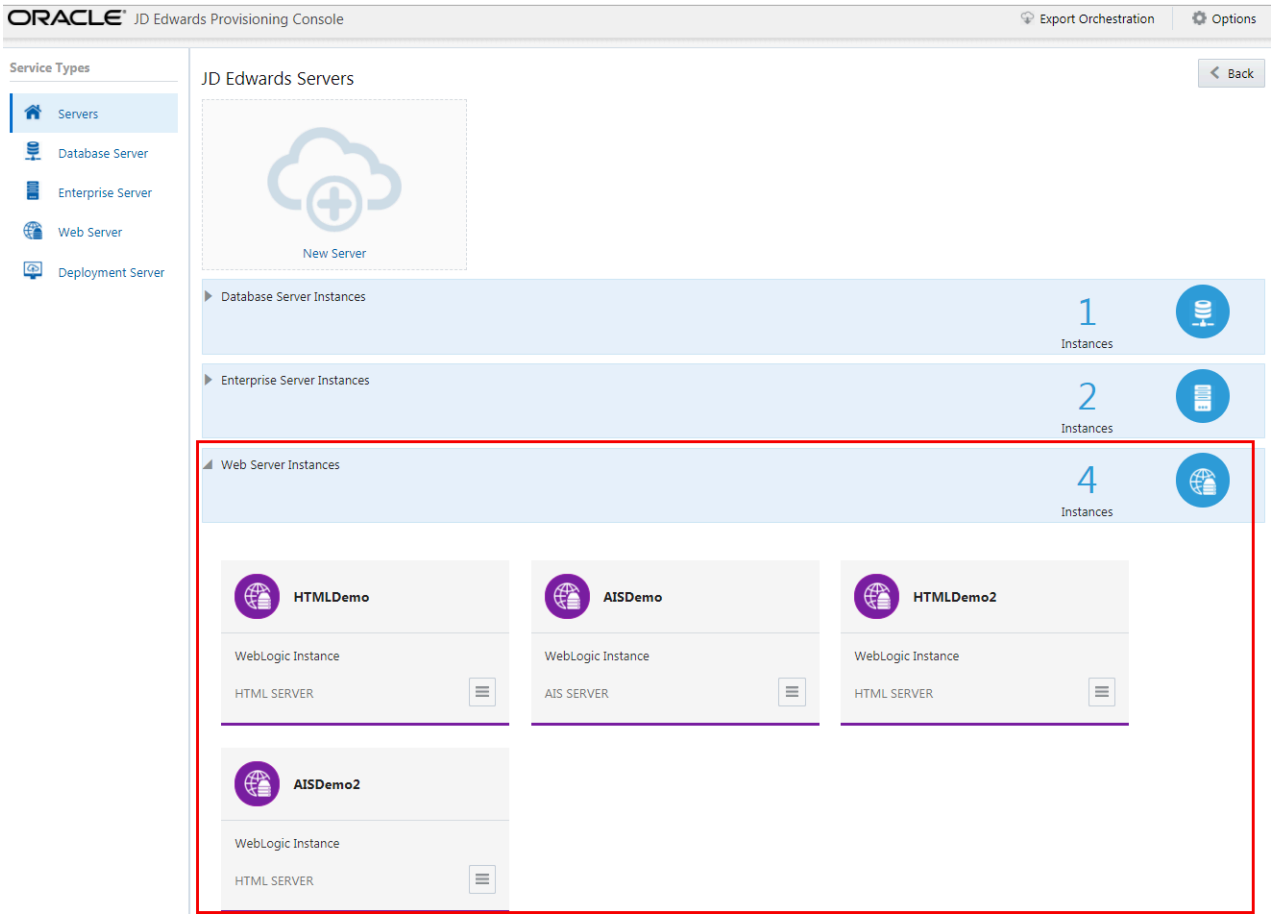

6. To delete any web instance, use the Application Options tab to select **Delete.**

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.

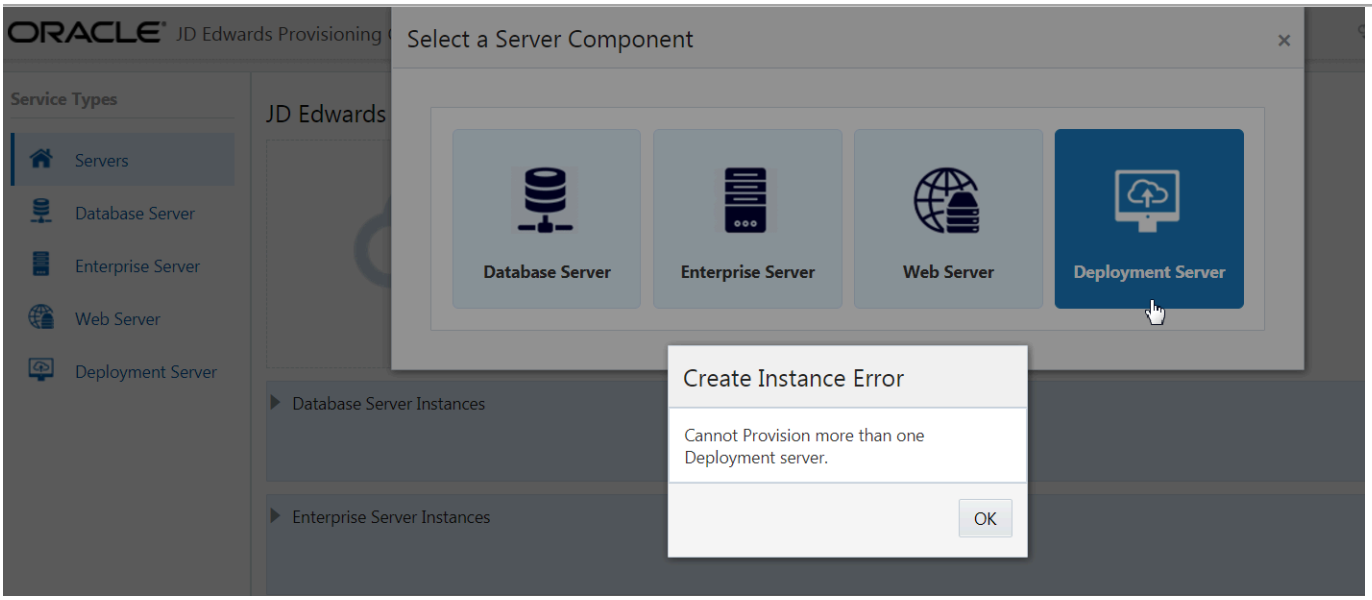
ORACLE JD Edwards Provisioning
Service Types
Servers
Database Server
Enterprise Server
Web Server
Deployment Server
JD Edwards
Database Server Instances
Enterprise Server Instances
Select a Server Component
Database Server
Enterprise Server
Web Server
Deployment Server
Create Instance Error
Cannot Provision 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**.

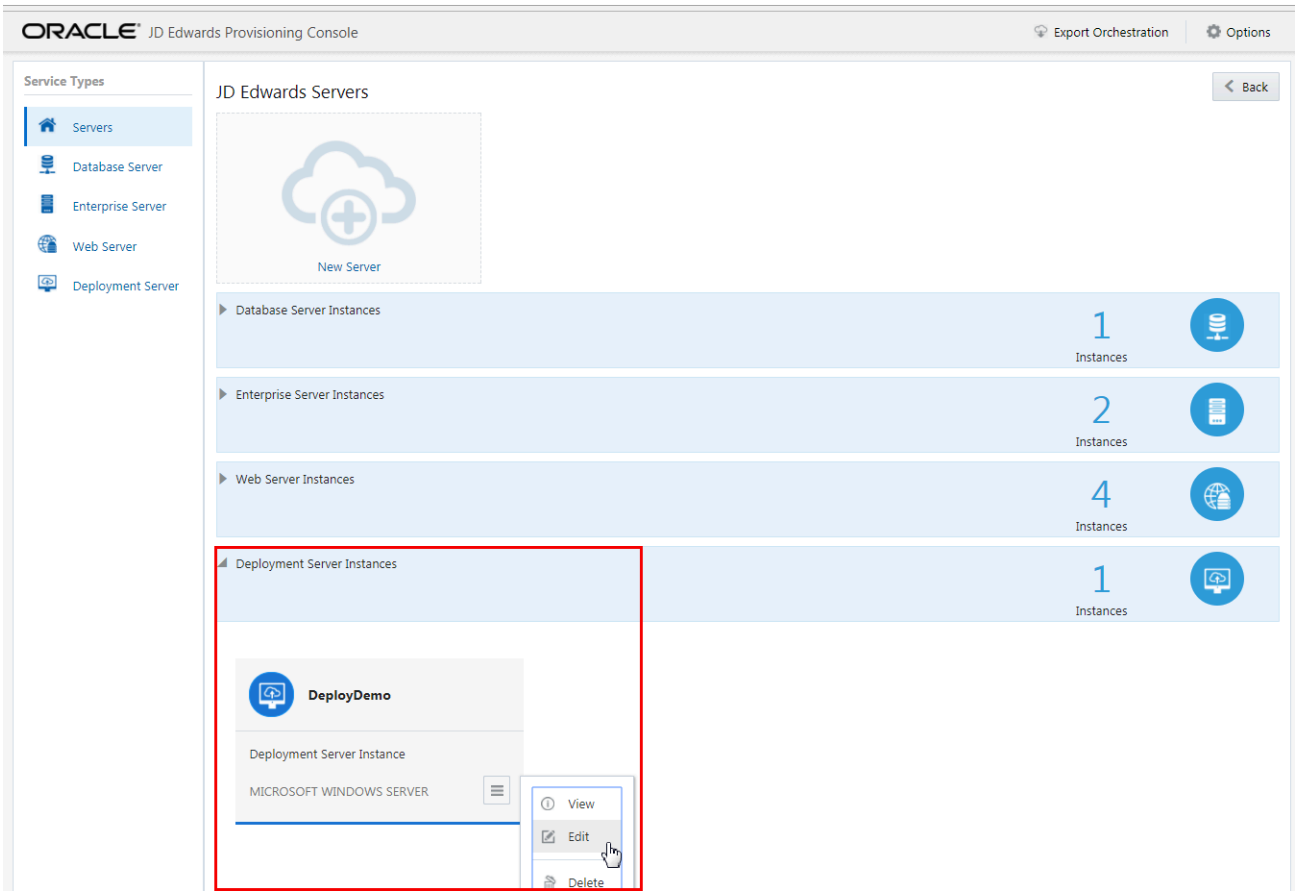

2. On JD Edwards Deployment Server page, complete these fields to create and configure your Deployment Server instance.

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.

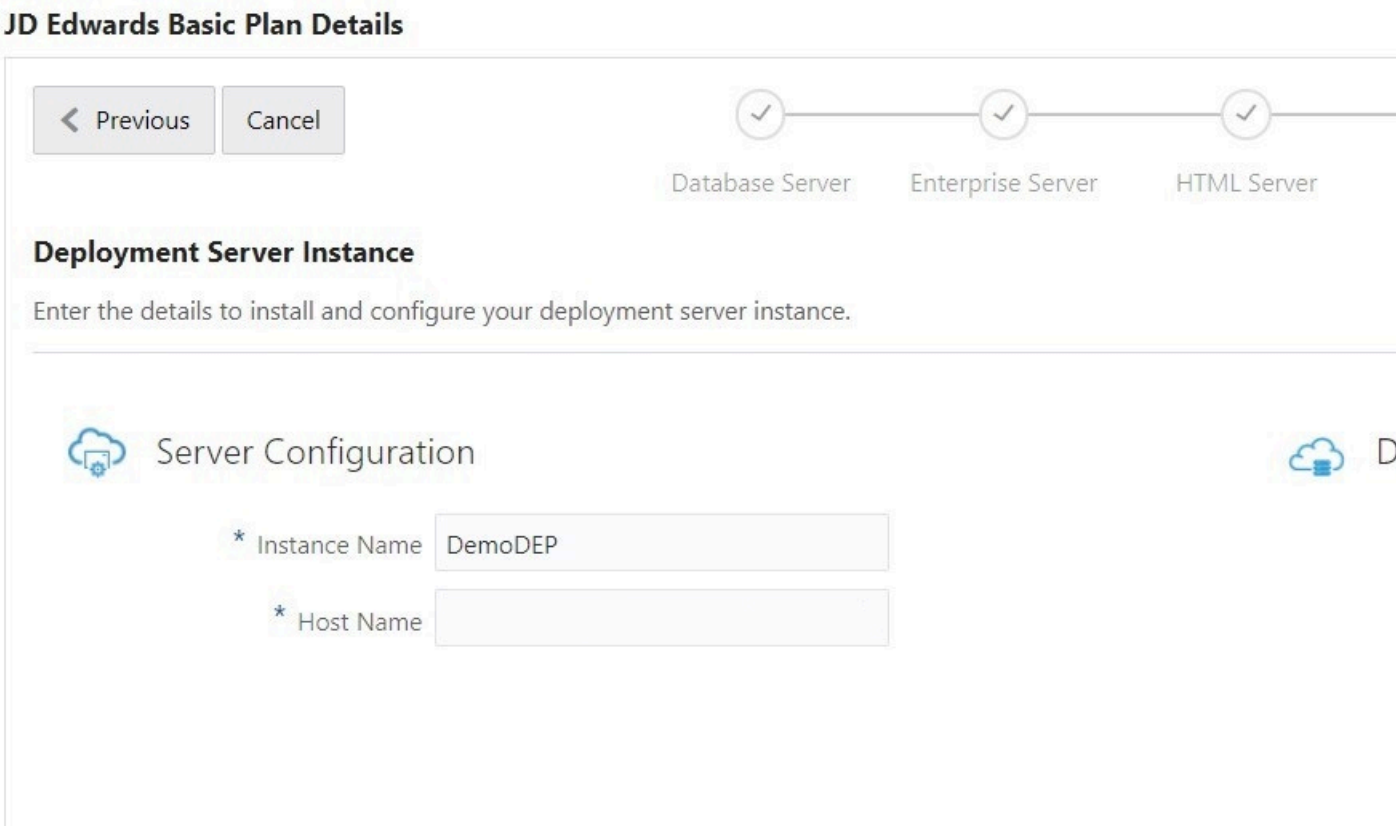

3. Verify that the Deployment Server instance you modified is displayed in the JD Edwards Servers window.
4. You can choose the **Delete** option in the action tab to delete the Deployment Server instance. After you delete the existing Deployment Server instance, you can click the **New Server** icon, 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.***

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.

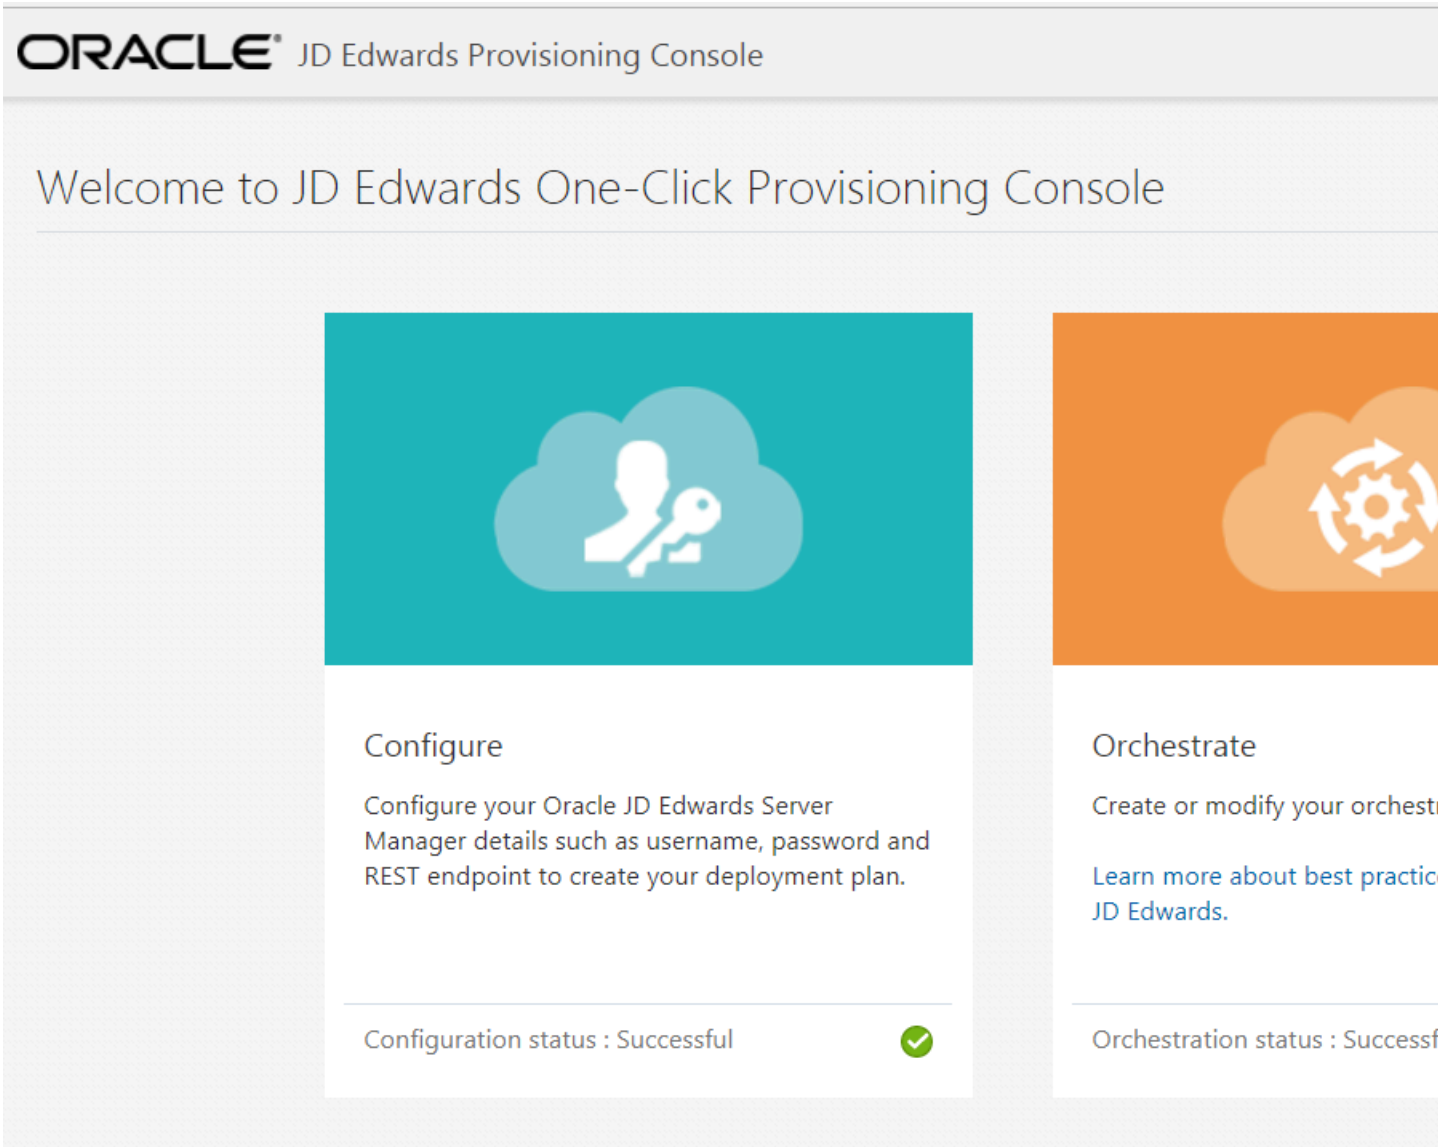

2. To view the Account details and Global Settings Summary, in the **Deployment Details** tab click the "**>**" icon for each server in the Deployment Details to see the details of the servers you provisioned.

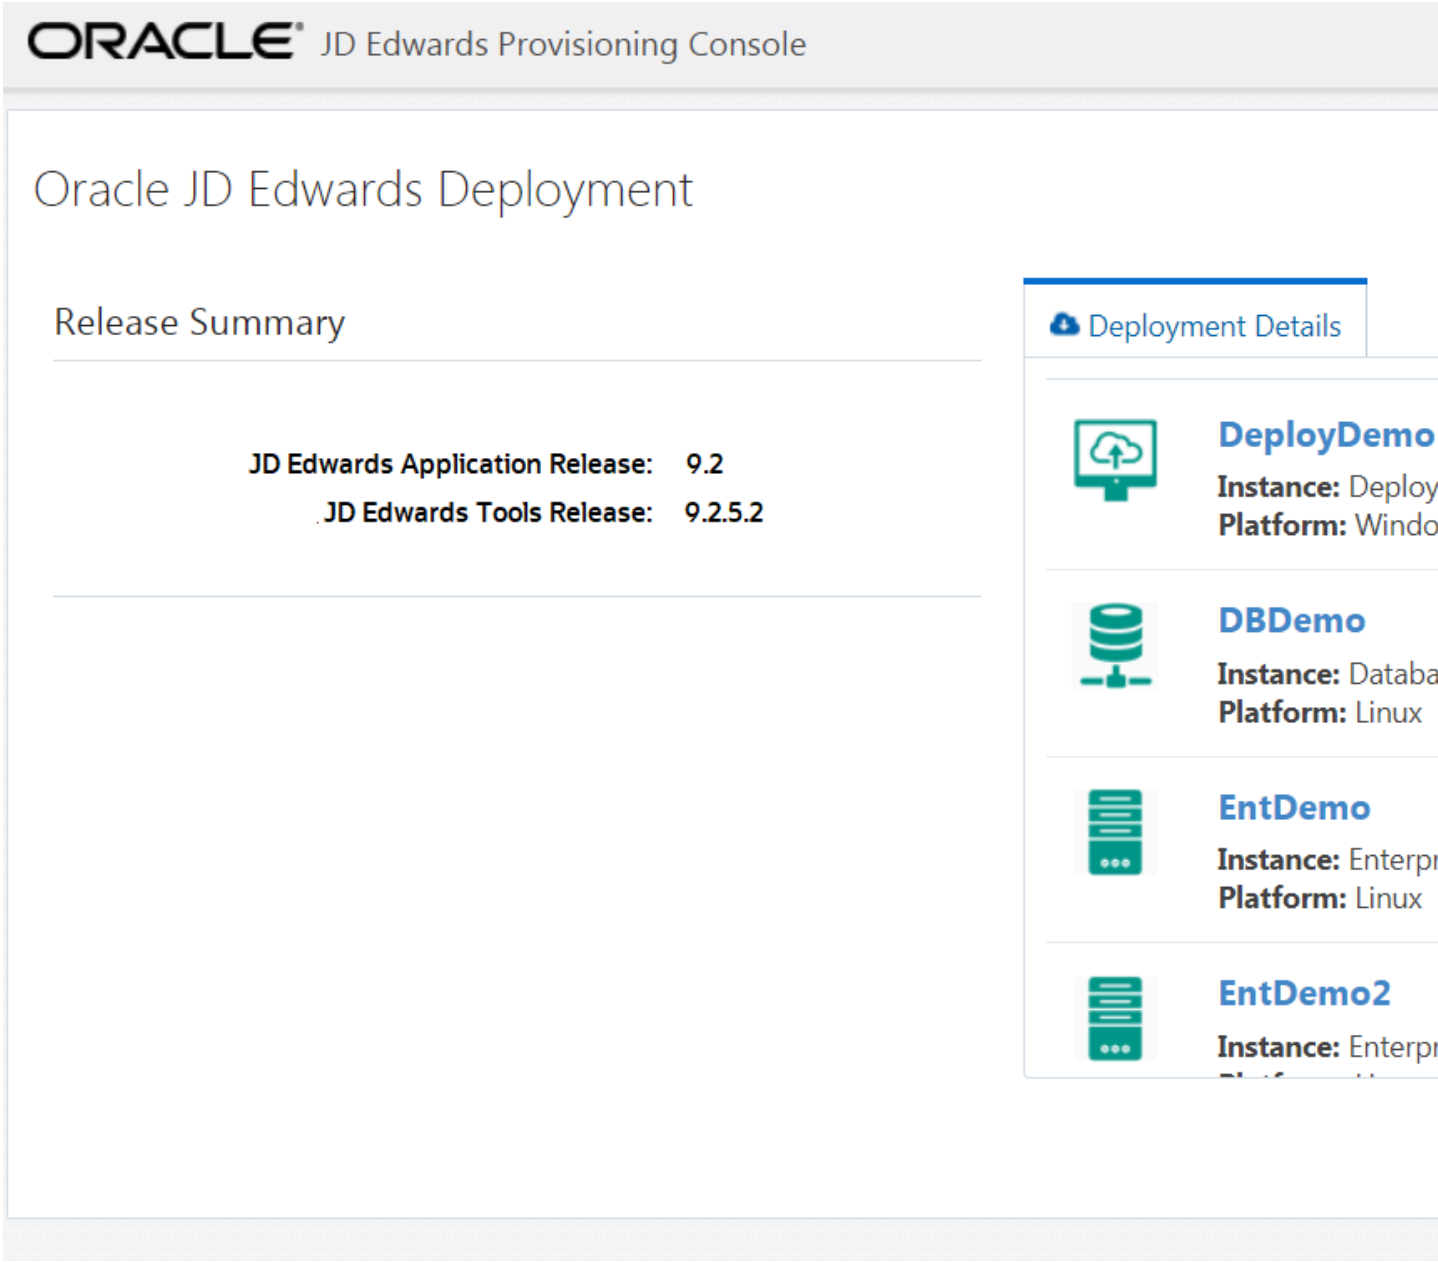

3. Click the **Back** button if required to make additional changes to the server instances.
4. To start the deployment, click the **Start Deployment** button.

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.

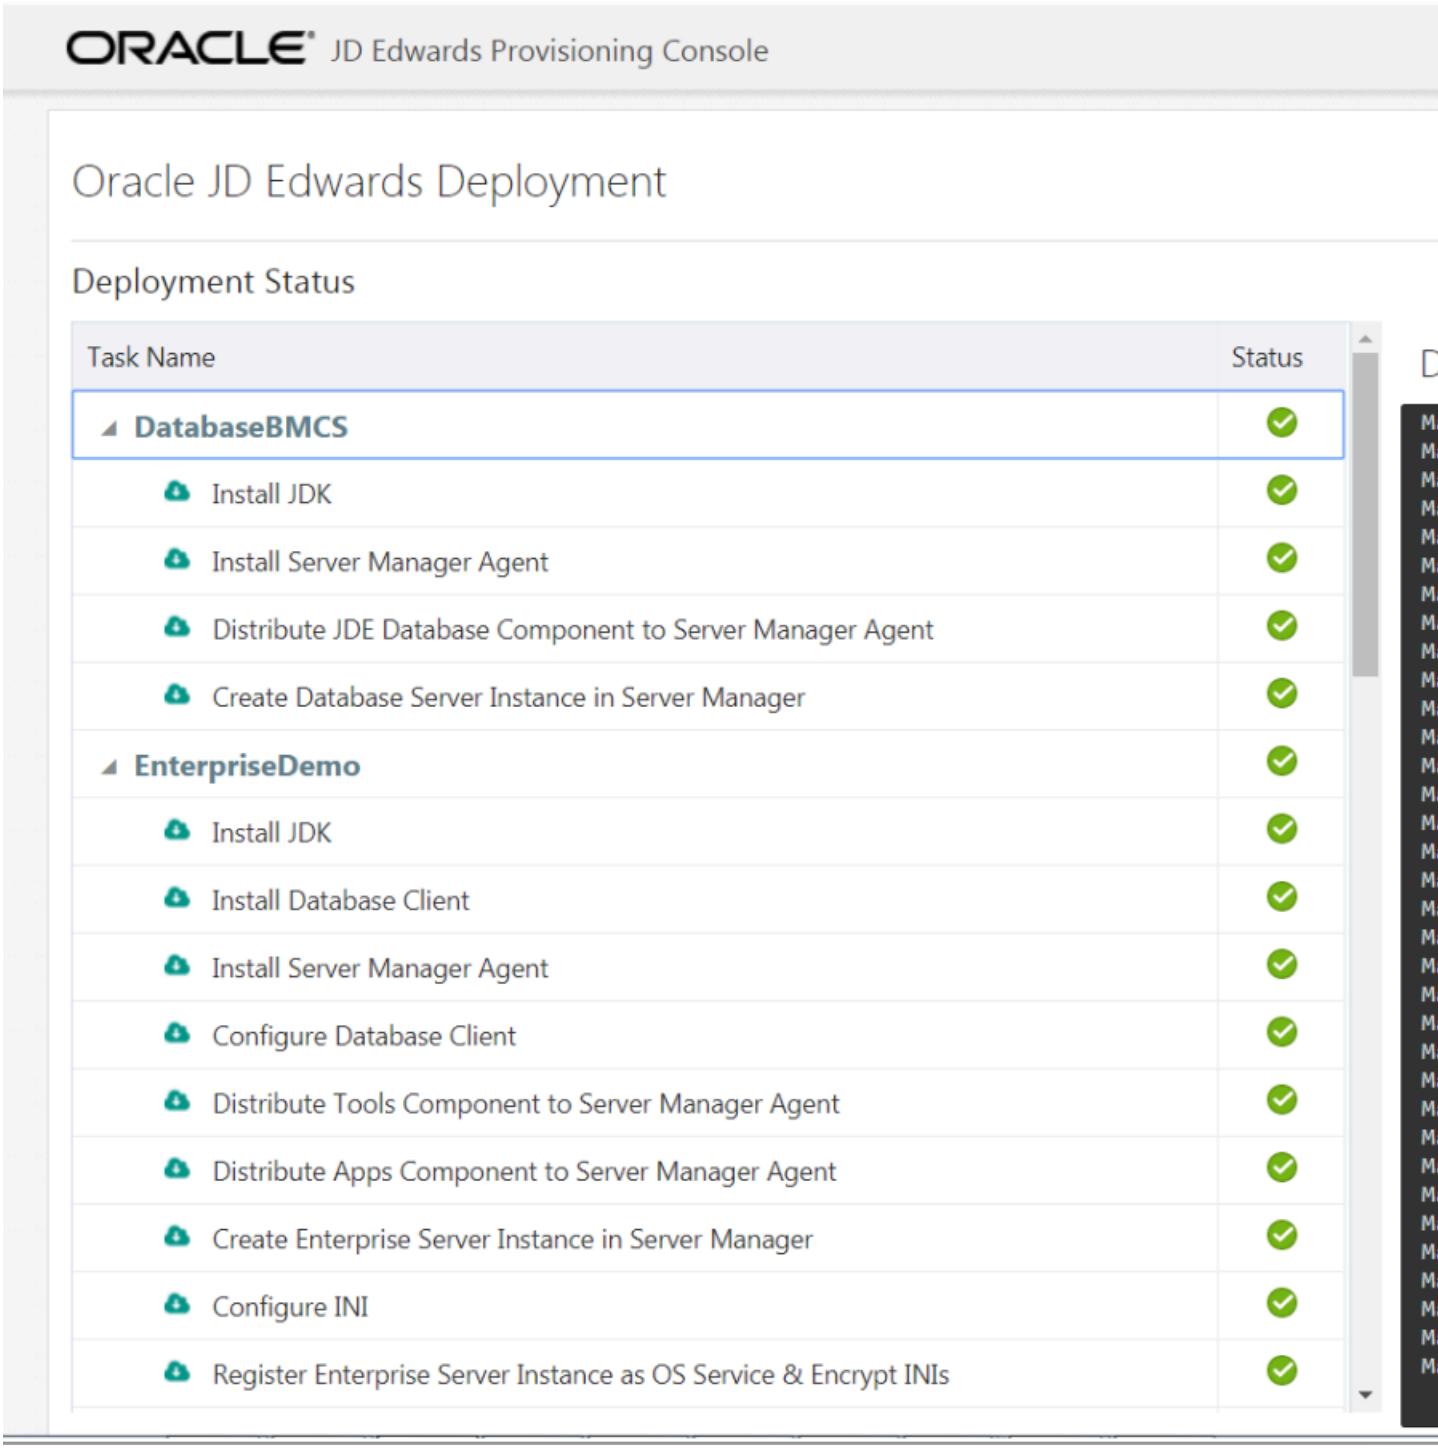

6. Click the **Back** button if you want to go back to the Deployment Summary window.

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.

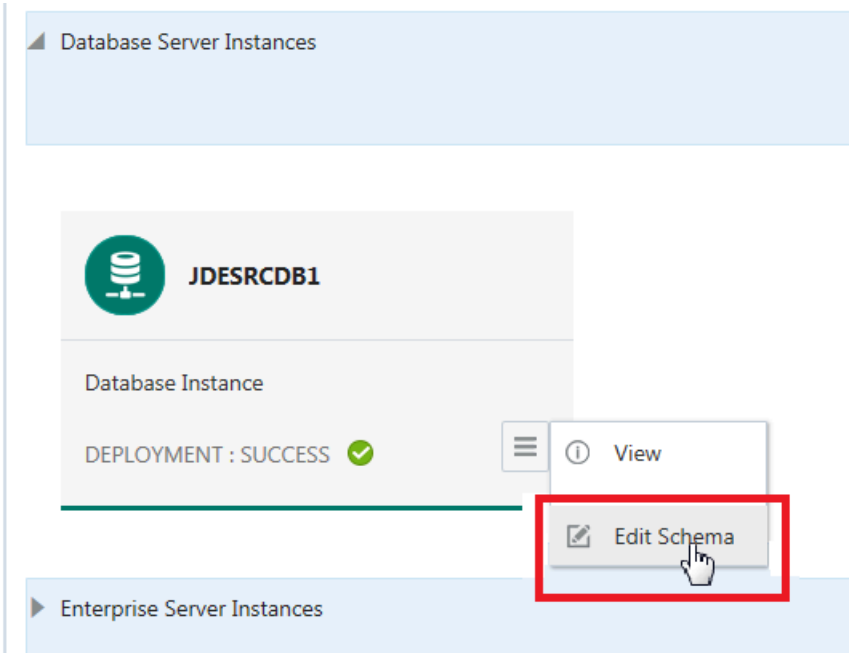

4. On Existing Database Instance Details, select the available pathcodes. An error message window appears if you attempt to delete the previously installed schema or the demo data.
5. Click the **OK** button.

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

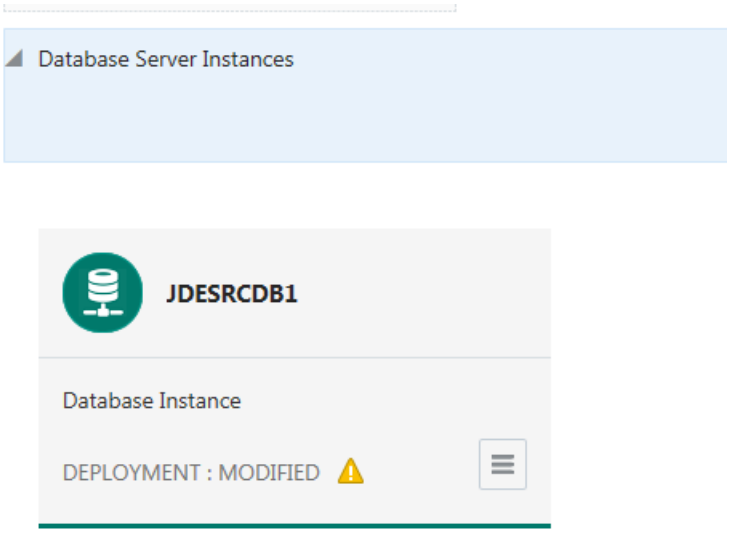

7. Click the **Back** button, and then click **Deploy** icon to start your modified deployment.

8 Configuring JD Edwards Components Post Deployment

Connecting to a Host in a Private Network Through the Bastion Host

This section shows you how to connect from a Microsoft Windows machine to a Linux host and Microsoft Windows host that have been deployed by JD Edwards EnterpriseOne Infrastructure Provisioning in Oracle Cloud Infrastructure. Such hosts are deployed in a private network and can only be accessed through the Bastion host as described in this section.

Prerequisite

- You must have installed PuTTY on your Microsoft Windows machine. By default, this installation includes the requisite software component called Pageant (Putty Authentication Agent). The program provides a Secure Shell (SSH) tunneling method for connecting to Unix or Linux machines through PuTTY.

Connecting to a Linux Host in a Private Network Through the Bastion Host

This procedure describes how to connect from a Microsoft Windows machine to a Linux host in a private network in Oracle Cloud Infrastructure through the Bastion host that has been deployed using JD Edwards EnterpriseOne infrastructure provisioning.

1. On your Microsoft Windows machine, search for the Pageant application (pageant.exe).

   Note: As mentioned in the "prerequisite" above, this program is a standard component of PuTTY for Microsoft Windows.

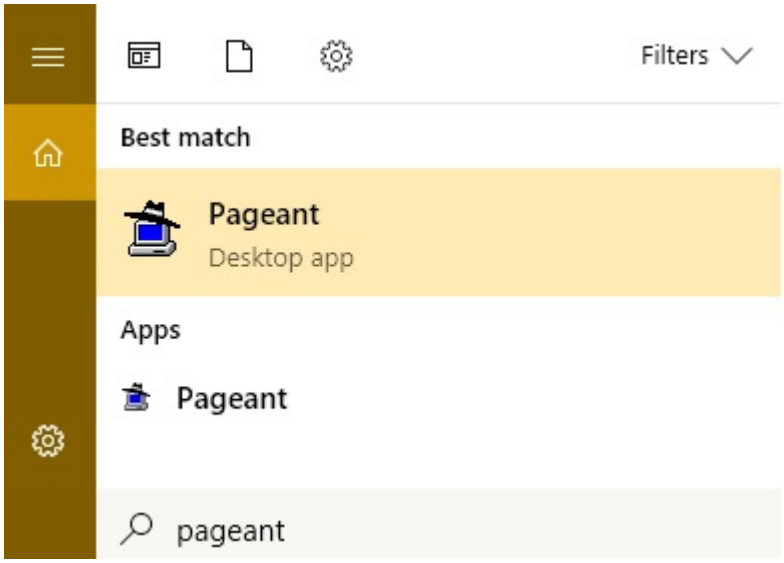

2. Start the Pageant application to access the Pageant Key List window.

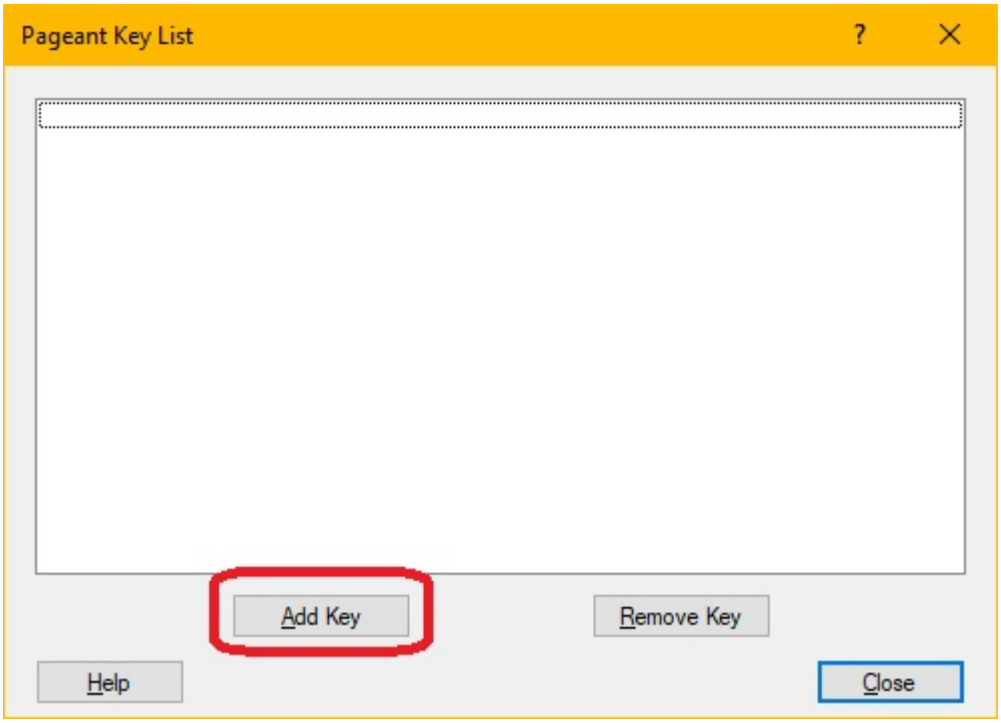

3. Click **Add Key** and browse to the private key you copied and converted to ppk format in your local workstation as described in the preceding task "Converting Your Private

SSH Keys to ppk Format" of this Learning Path. For example: **OCI_Instance.ppk**.

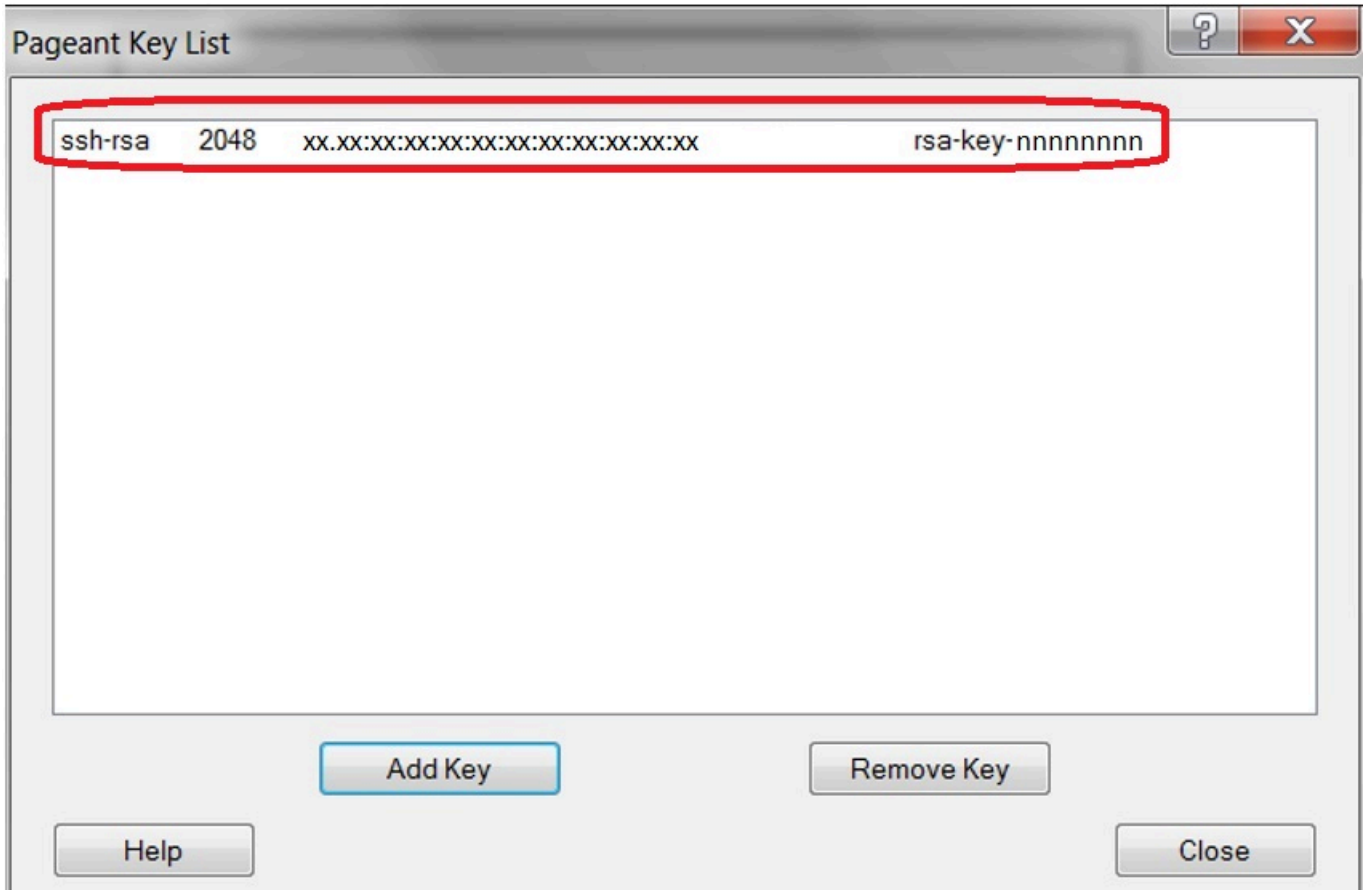

4. Open PuTTY and in the Host Name field, enter the public IP address of the Bastion server (also includes the NAT server and Server Manager Console).

Note: As a best practice, you can save this PuTTY session for future use when logging in to machines in the private network known to this Bastion host.

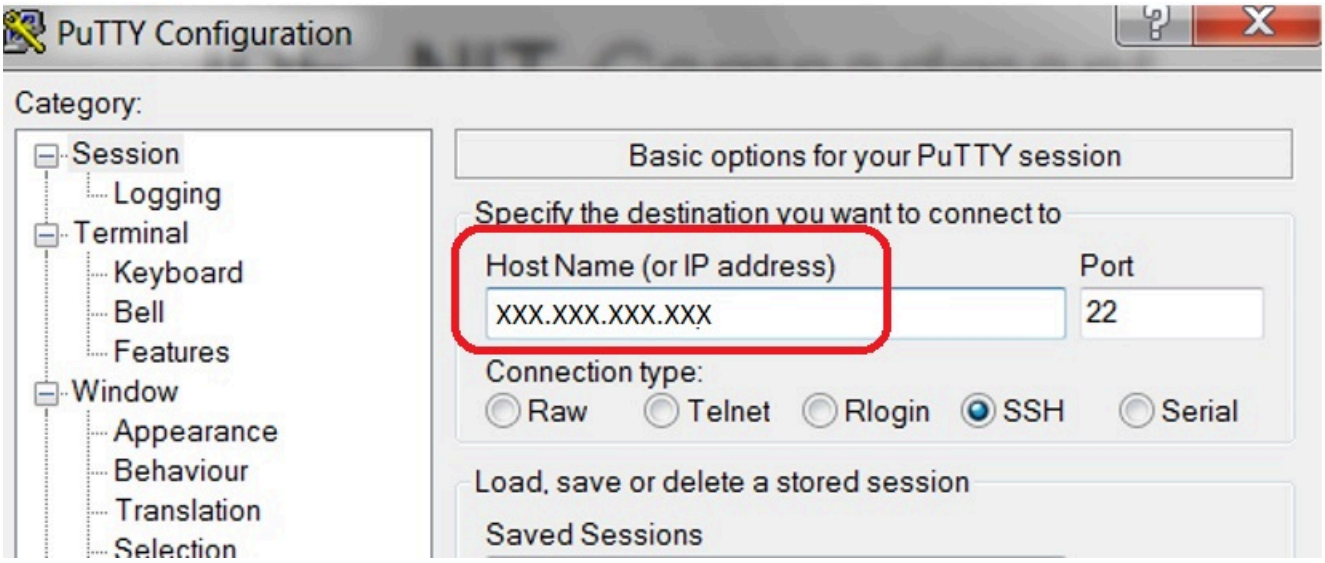

5. From the Connection type options, select **SSH,** and then click **Auth**.

6. To log in to the connection, you need to log in as the **opc** user. Expand the **Connection** node and in the **Data** section, enter the value **opc** in the **Auto-login username** field.

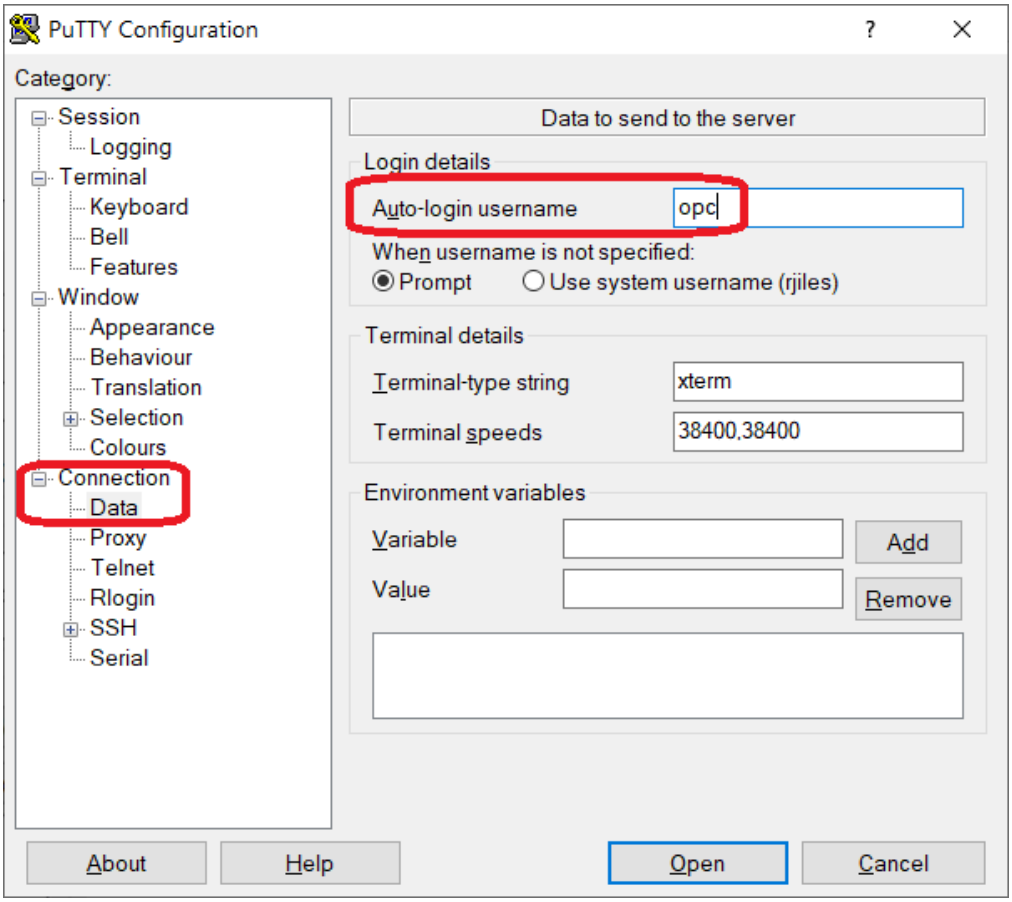

7. In the Options controlling SSH authentication section, in the **Authentication parameters** subsection, ensure that the **Allow agent forwarding** option is selected.

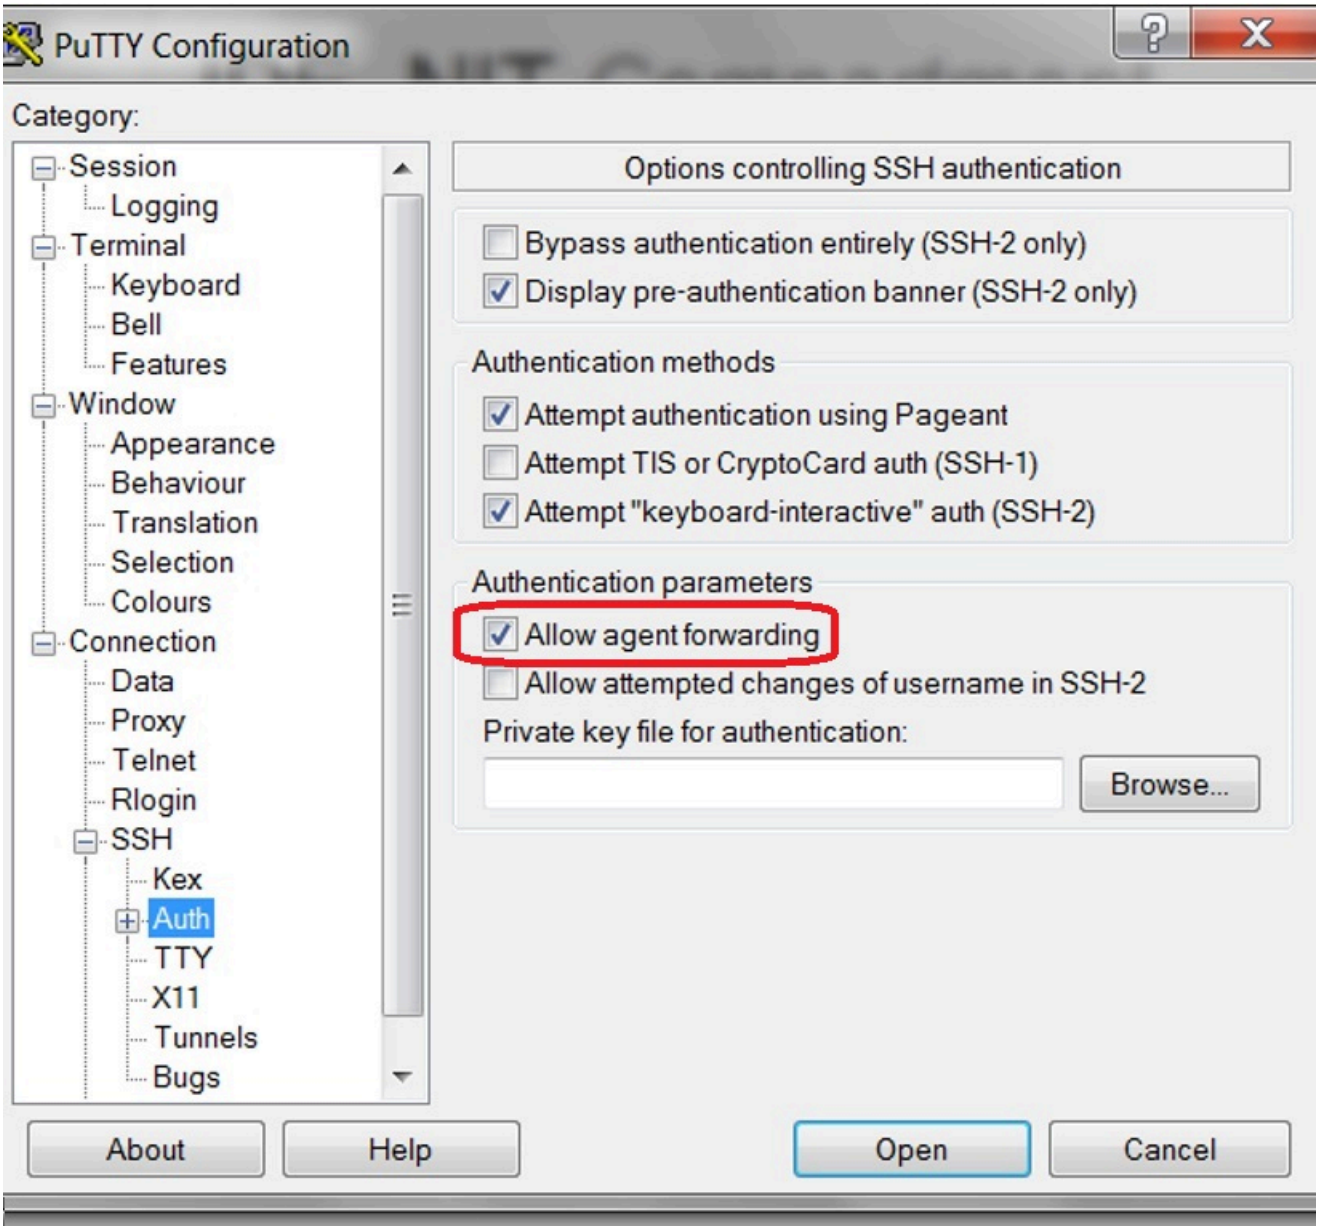

8. Under Private key file for authentication, click the **Browse** button to select the `Bastion.ppk` key that you converted to ppk format in your local workstation as described in the

preceding task "Converting Your Private SSH Keys to .ppk Format" of this Learning Path.

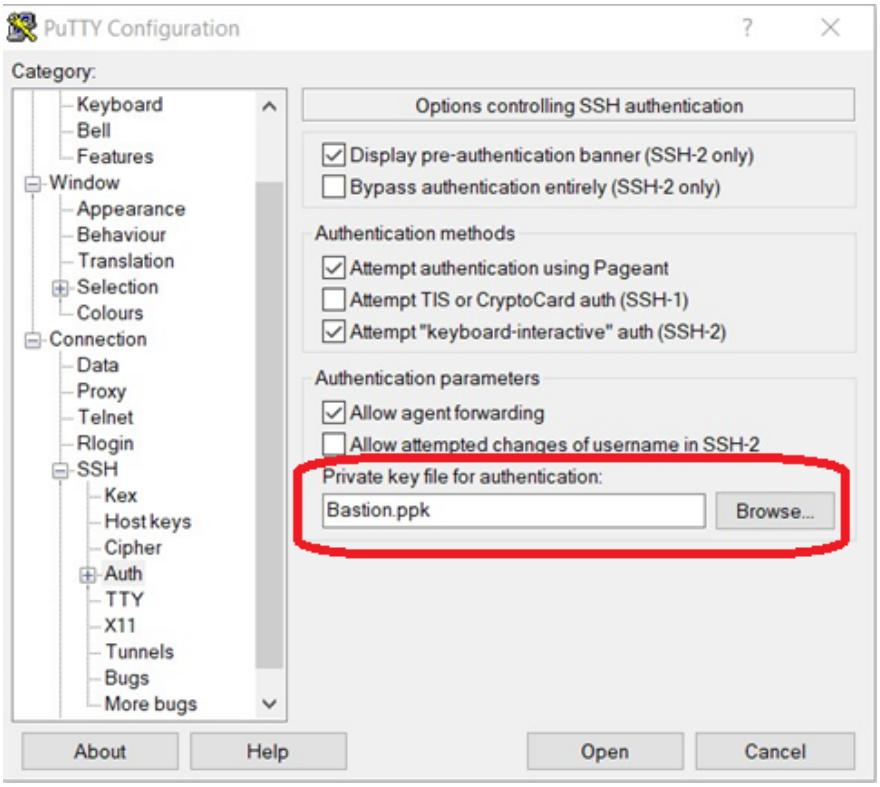

9. Go back to **Session**, click the **Save** button.
10. Click the **Open** button to open a console on the Bastion Server.

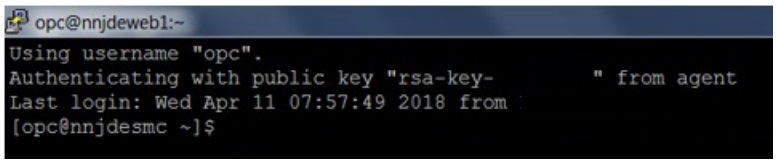

11. As this point you can use the below command to access the required JD Edwards EnterpriseOne Server using its private IP address.

```
ssh opc@<private_ip_of_any_instance>
```

For example: `ssh opc@10.0.0.0`

Connecting to a Windows Host in a Private Network Through the Bastion Host

This procedure describes how to connect from a Microsoft Windows machine to a Windows host in a private network in Oracle Cloud Infrastructure through the Bastion host that has been deployed using JD Edwards EnterpriseOne infrastructure provisioning.

1. On your Microsoft Windows machine, search for the Pageant application (pageant.exe).

   Note: As mentioned in the section "prerequisite" above this program is a standard component of PuTTY for Microsoft Windows.

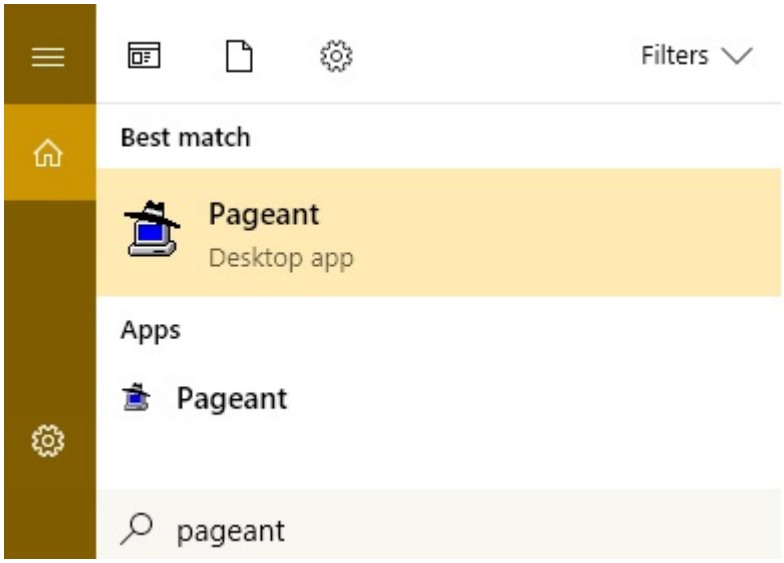

2. Start the pageant application to access the Pageant Key List window.

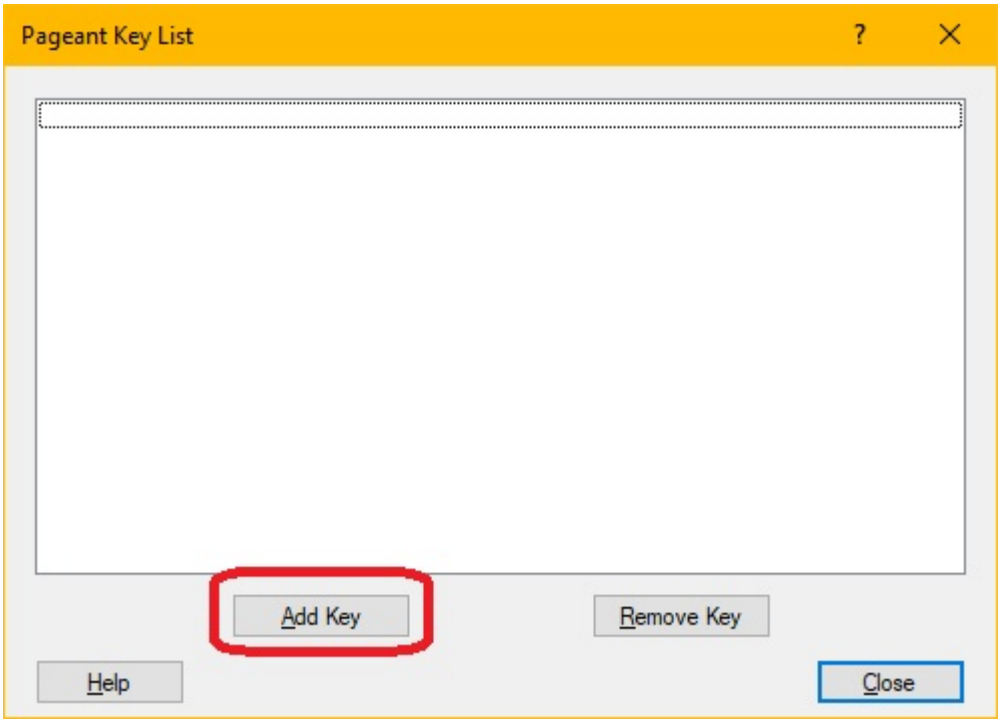

3. Click **Add Key** and browse to the private key you converted to ppk format in your local workstation as described in the preceding task "Converting Your Private

SSH Keys to ppk Format" of this Learning Path. For example: `OCI_Instance.ppk`.

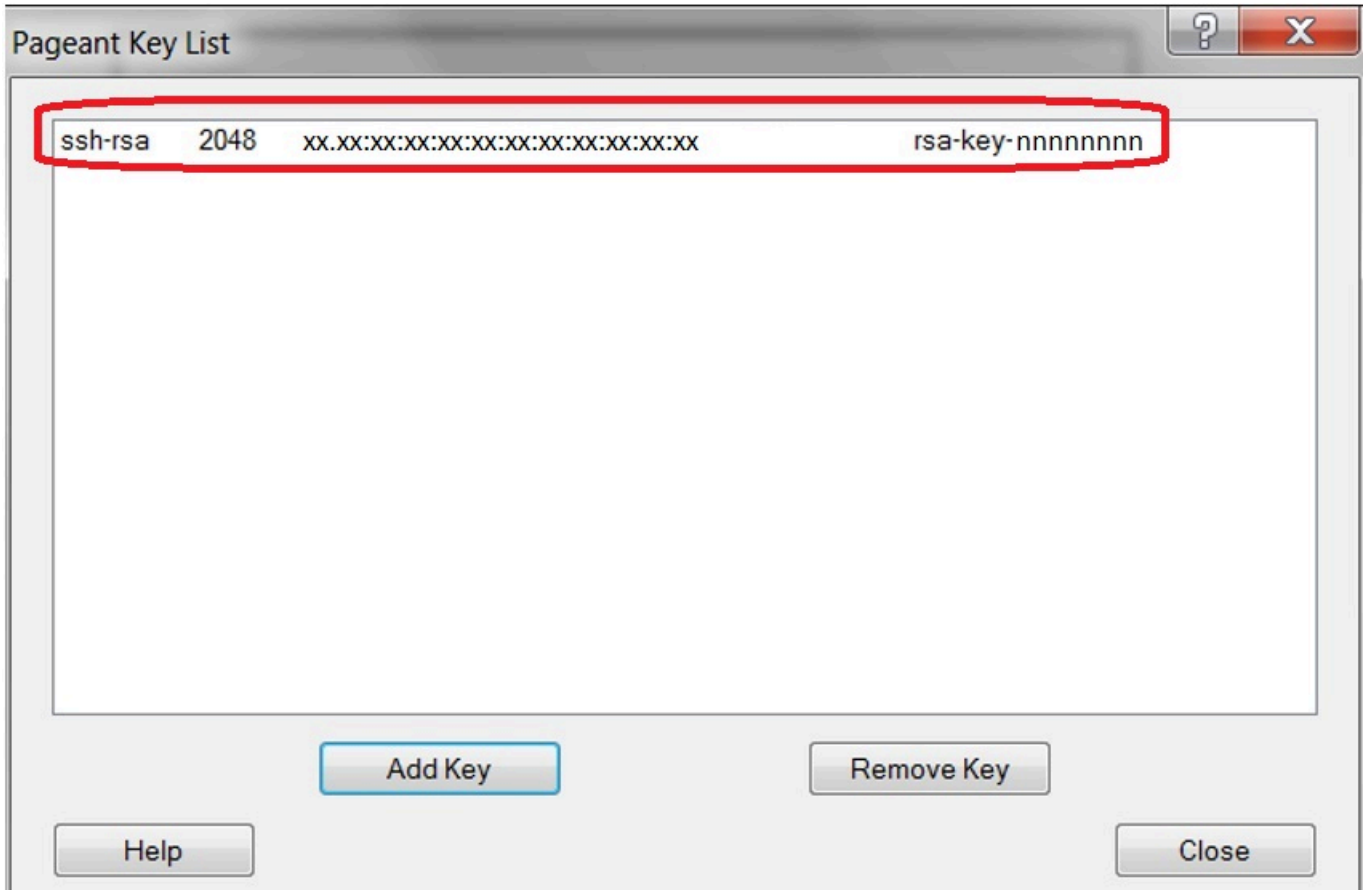

4. Open PuTTY and in the Host Name field, enter the public IP address of the Bastion server (also includes the NAT server and Server Manager Console).

Note: As a best practice, you can save this PuTTY session for future use when logging in to machines in the private network known to this Bastion Host.

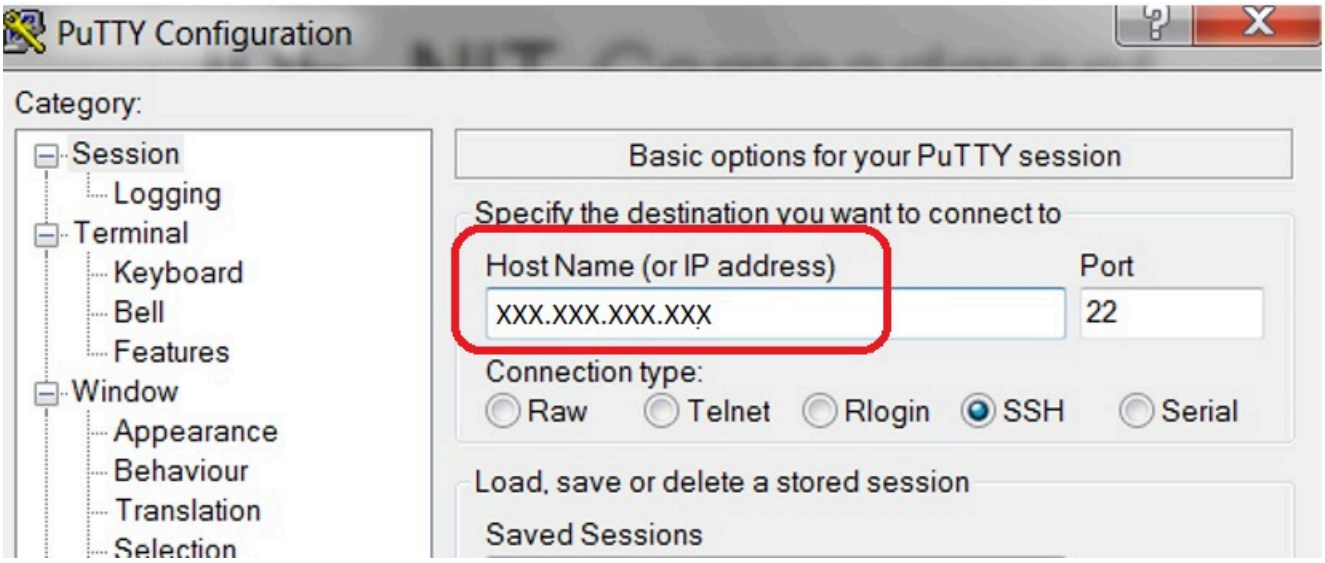

5. To create an SSH tunnel to the local host, in the Connection node, click **SSH**, and then click **Tunnels**.

6. In the Options controlling SSH authentication section, in the Authentication parameters subsection, ensure that the **Allow agent forwarding** option is selected.

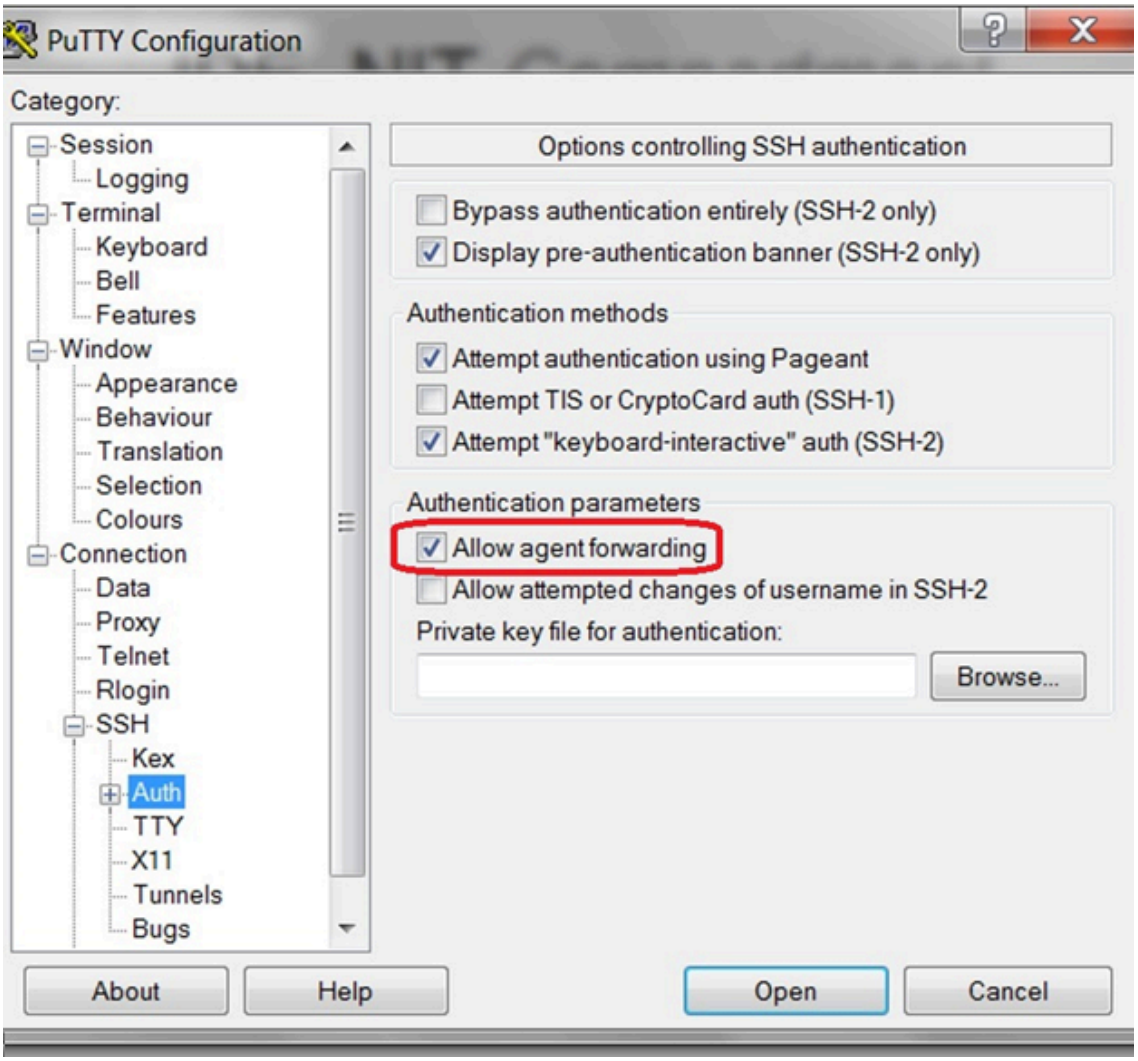

7. Under Private key file for authentication, click the **Browse** button to select the Bastion.ppk key that you converted to ppk format in your local workstation as described in the

preceding task "Converting Your Private SSH Keys to .ppk Format" of this Learning Path.

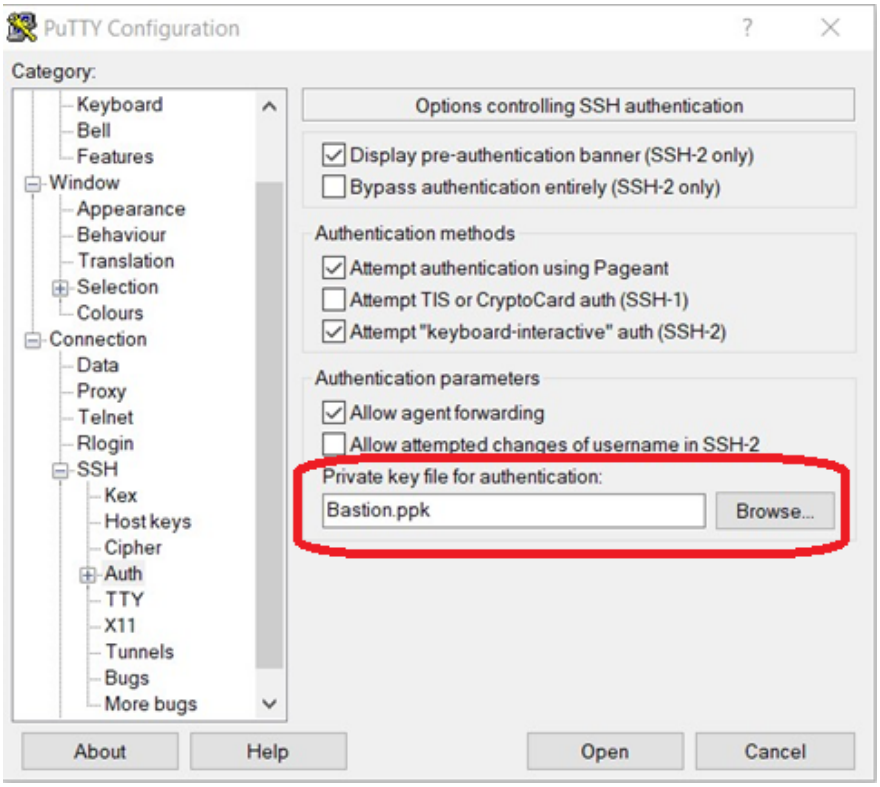

8. To create an SSH tunnel to the local host, in the **Category** section, expand the **Connection** node, expand the **SSH** node, and click **Tunnels**.
9. In the Options controlling SSH port forwarding section, enter the port number in the **Source Port** field. You can choose any port that is free on your local machine. For example, 33389.
10. In the **Destination** field, enter the private IP address of the server you want to connect to and port 3389.
11. Click the **Add** button to add the port.

12. After you click the **Add** button, confirm that the IP address and the port are added as shown below.

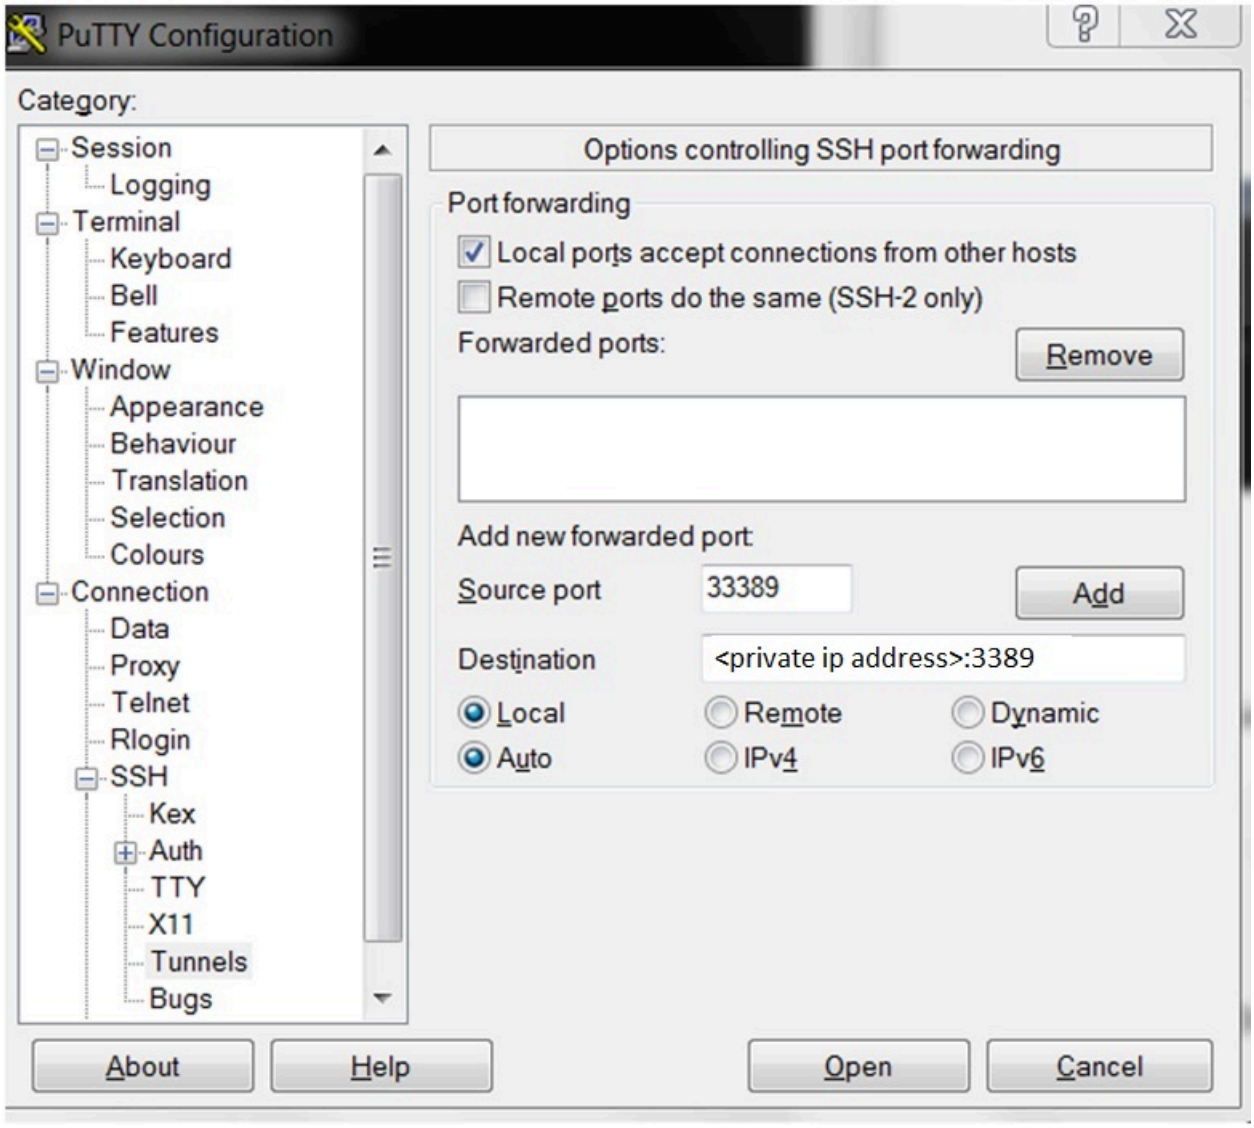

13. Ensure that these options are selected:
    - **Local ports accept connections from other hosts**
    - **Local**
    - **Auto**

> **Note:** Before you proceed, as a best practice you should save this PuTTY session for future use when logging in to machines in the private network known to this Bastion host.

14. Launch Remote Desktop Connection and connect to localhost: 3389.

Note: You will lose access to this local host if the PuTTY session becomes inactive.

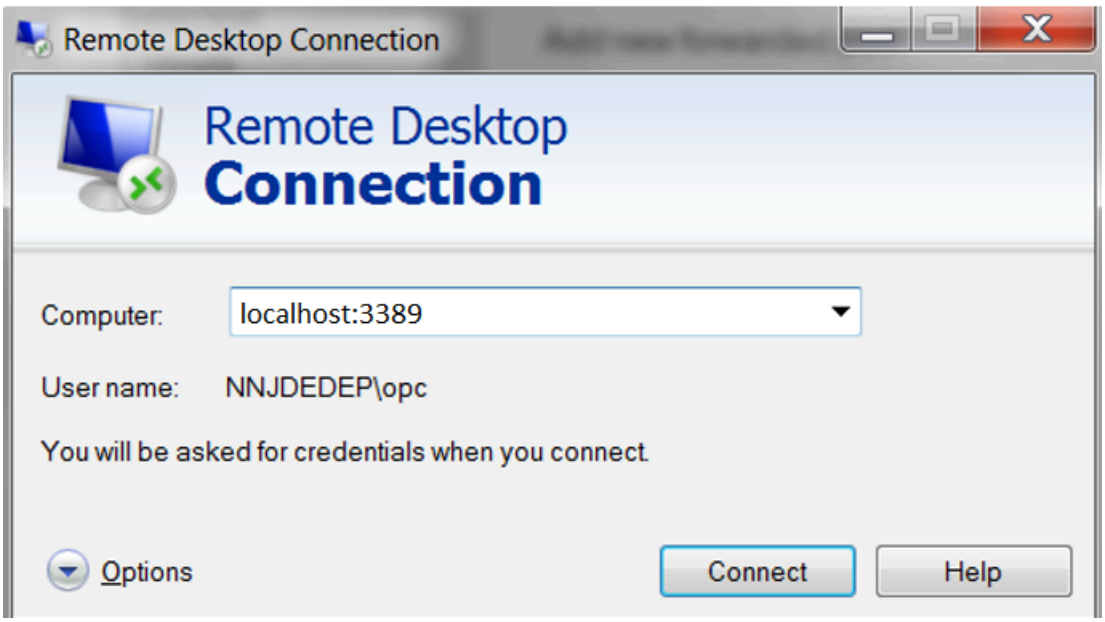

Performing Post Infrastructure Provisioning Tasks

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

This section describes these tasks:

- Re-enabling SELINUX
- Updating the Microsoft Windows Host File for JD Edwards Enterprise Servers with Load Balancer
- Starting the Database in the Secondary Node in RAC

Prerequisite

- You must have successfully completed the deployment of an orchestration using the JD Edwards One-Click Provisioning Console.

1. **Re-enabling SELINUX**
   Use this command to check the status of Security Enhanced Linux (SELINUX) on each instance that was creating by Infrastructure Provisioning in support of JD Edwards EnterpriseOne servers for logic, batch and Oracle WebLogic servers.
   sudo getenforce
   If the returned status is **Disabled**, you can re-enable the extra security restriction provided by SELINUX by modifying this file:
   /etc/selinux/config
   Edit the /etc/selinux/config file to change the SELINUX=setting to either of these values:

SELINUX=**Enforcing**

or

SELINUX=**Permissive**

You must reboot the machine for the change to take affect any time you change the /etc/hostname file or when you change security settings.

After the reboot is complete, run sudo getenforce to confirm that the status of SELINUX is **Enforcing** or **Permissive**, whichever you set.

2. **Updating the Microsoft Windows Host File for JD Edwards Enterprise Servers with Load Balancer**

   You must update the host file of the Microsoft Windows with the following entry if the JD Edwards Enterprise Servers are configured with the load balancer:

   <ip address of actual logic server> <virtual hostname of logic server>

   <ip address of actual batch server> <virtual hostname of batch server>

   For example, if the private IP address of the logic server is 10.0.0.10, the private IP address of the batch server is 10.0.0.11, the virtual host name of the logic server is logiclb and the virtual host name of the batch server is batchlb, then update the host file with the following entry:

   ```
   10.0.0.10 logiclb
   10.0.0.11 batchlb
   ```

   > **Note:** You must make these entries in both the Deployment Server and the Development Client.

   If there are multiple logic and batch servers, when you do package builds on the Deployment Server, you must further modify the /hosts file to include the private IP address of each logic and batch server. For example, if you want to deploy a package on a specific private IP address, you must modify the /hosts file to include that private IP address. For example, if you have a logic server with an IP address of 10.0.0.12 and a batch server with an IP address of 10.0.0.13, you would change the above entries (10.0.0.10 and 10.0.0.11) so that only the below entries are in the /hosts file:

   ```
   10.0.0.12 logiclb
   10.0.0.13 batchlb
   ```

   Similarly, you will need to add entries for each additional logic and batch server for which you want to build packages.

3. **Starting the Database in the Secondary Node in RAC**

   This section is only applicable for the Production environment, which is the only environment where RAC is enabled. By following 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: ***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 a node goes down while processing a transaction resulting in the failure of the database connection, then there is a 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 a 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.

> **Note: Disaster Recovery.** If you are pointing virtual hostname of enterprise servers in hosts file to the instance in Disaster Recovery (DR) region then you may have to make additional changes described below:

a. Edit this file:
```
C:\JDE\oracle19\product\19.0\client_1\network\admin\tnsnames.ora
```

b. Update the below **bolded** values for HOST= and SERVICE NAME = based on what you have in the DR region:
```
JDEORCL = (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST =
 jdepddb1dr.privatregsub.vcndr.oraclevcn.com)(PORT = 1521)) (CONNECT_DATA =
```

```
(SERVER = DEDICATED) (SERVICE_NAME = jdepdb.privatregsub.vcndr.oraclevcn.com) ) )
```

c. Add the host entry of the standby database host in the hosts file of your instance of Microsoft Windows. For example:

d. `jdepddb1dr.privatregsub.vcndr.oraclevcn.com`

To view details of the secondary node in RAC:

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

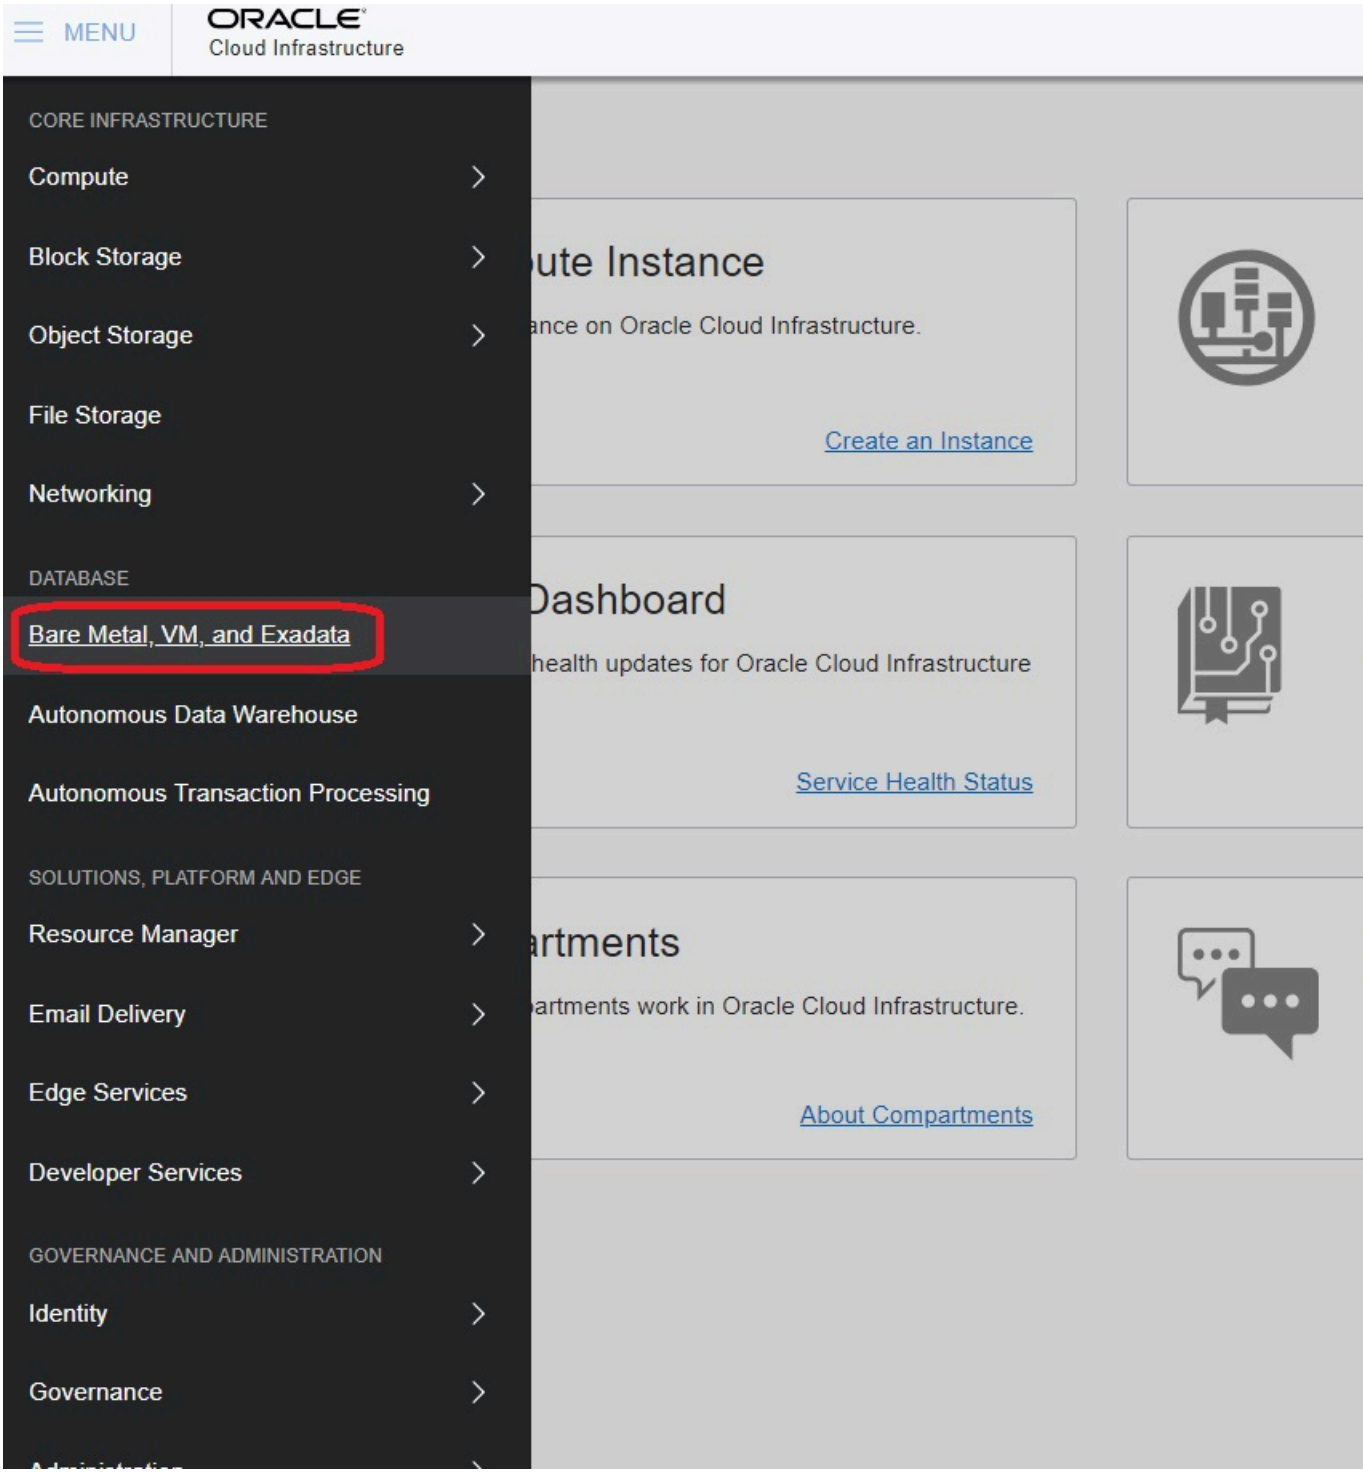

b. From the Navigation Menu, in the **Database** section, select **Bare Metal, VM, and Exadata**.
c. In the left panel ensure that DB Systems is selected, and also ensure that under **List Scope** the applicable **COMPARTMENT** is selected which contains the DB System with RAC that you created.

d. Click the link for your DB System with RAC.

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

The system displays the details of the two nodes as shown below:

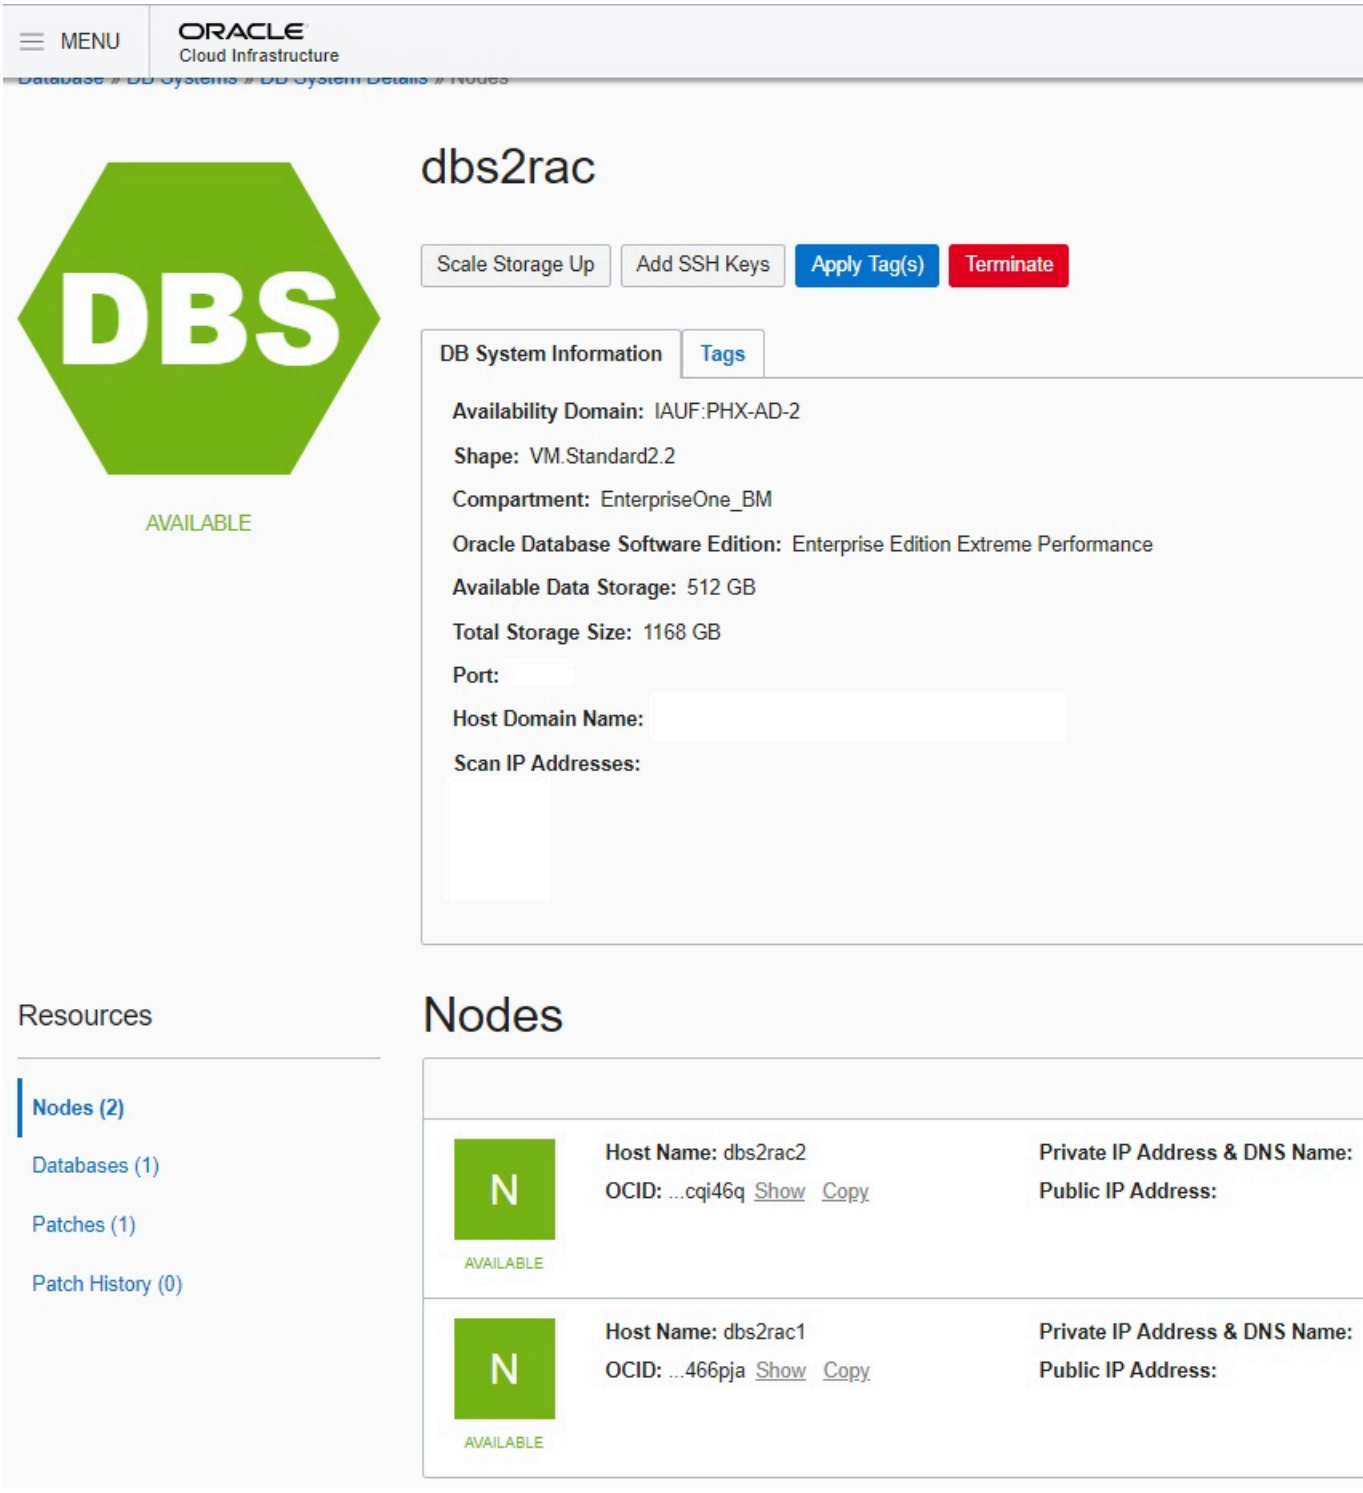

Setting up Additional Security Configurations

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.

Enabling Oracle Data Guard on Oracle Cloud Infrastructure

This section describes how to enable Oracle Data Guard on Oracle Cloud Infrastructure by performing these tasks:

- Enabling Oracle Data Guard in Oracle Cloud Infrastructure

- Postinstallation of JD Edwards EnterpriseOne Servers
- Switch Over to the Disaster Recovery (DR) Region
- Limitations of Switch Over

Note: It is your responsibility to synchronize and maintain the primary and DR region. For example, such activities would include such functions as applying security patches, and JD Edwards EnterpriseOne package build and deploy.

Refer to these Oracle web sites for additional details:

- *Plan Disaster Recovery for Databases*
- *Using Oracle Data Guard*

Enabling Oracle Data Guard in Oracle Cloud Infrastructure

Use this procedure to enable Oracle Data Guard in Oracle Cloud Infrastructure.

1. Go to the DB system for which you want to enable Oracle Data Guard.
2. Click Databases on the left panel.
3. Click on an existing database (for example, ORCL).
4. Click on Data Guard Associations on the left panel.
5. Click the Enable Data Guard button.
6. On the popup screen to enter details, complete these fields:
    - Display name

      Enter any name that is different than the primary instance.
    - Region

      Select the same Disaster Recovery region you have previously chosen.
    - Availability domain

      Select the same availability domain that is specified for the primary instance.
    - Select a shape

      Select shape is based on your requirement.

      Note: Initially only shapes with a 1 Node Count RAC are supported for use with JD Edwards EnterpriseOne using Oracle Data Guard on Oracle Cloud Infrastructure.
    - Virtual Cloud Network

      Select the VCN in the Disaster Recovery region in which you want to provision the standby database system.
    - Subnet

      Select the private subnet for use by this Oracle Data Guard.
    - Configure network security groups (NSGs)

      Leave this checkbox unchecked.
    - Hostname prefix

      Provide hostname prefix of your choice. Ensure this value is different than that specified for the primary instance.

- Configure standby database

Provide the database password. The standby database admin password must be the same as the primary database admin password.

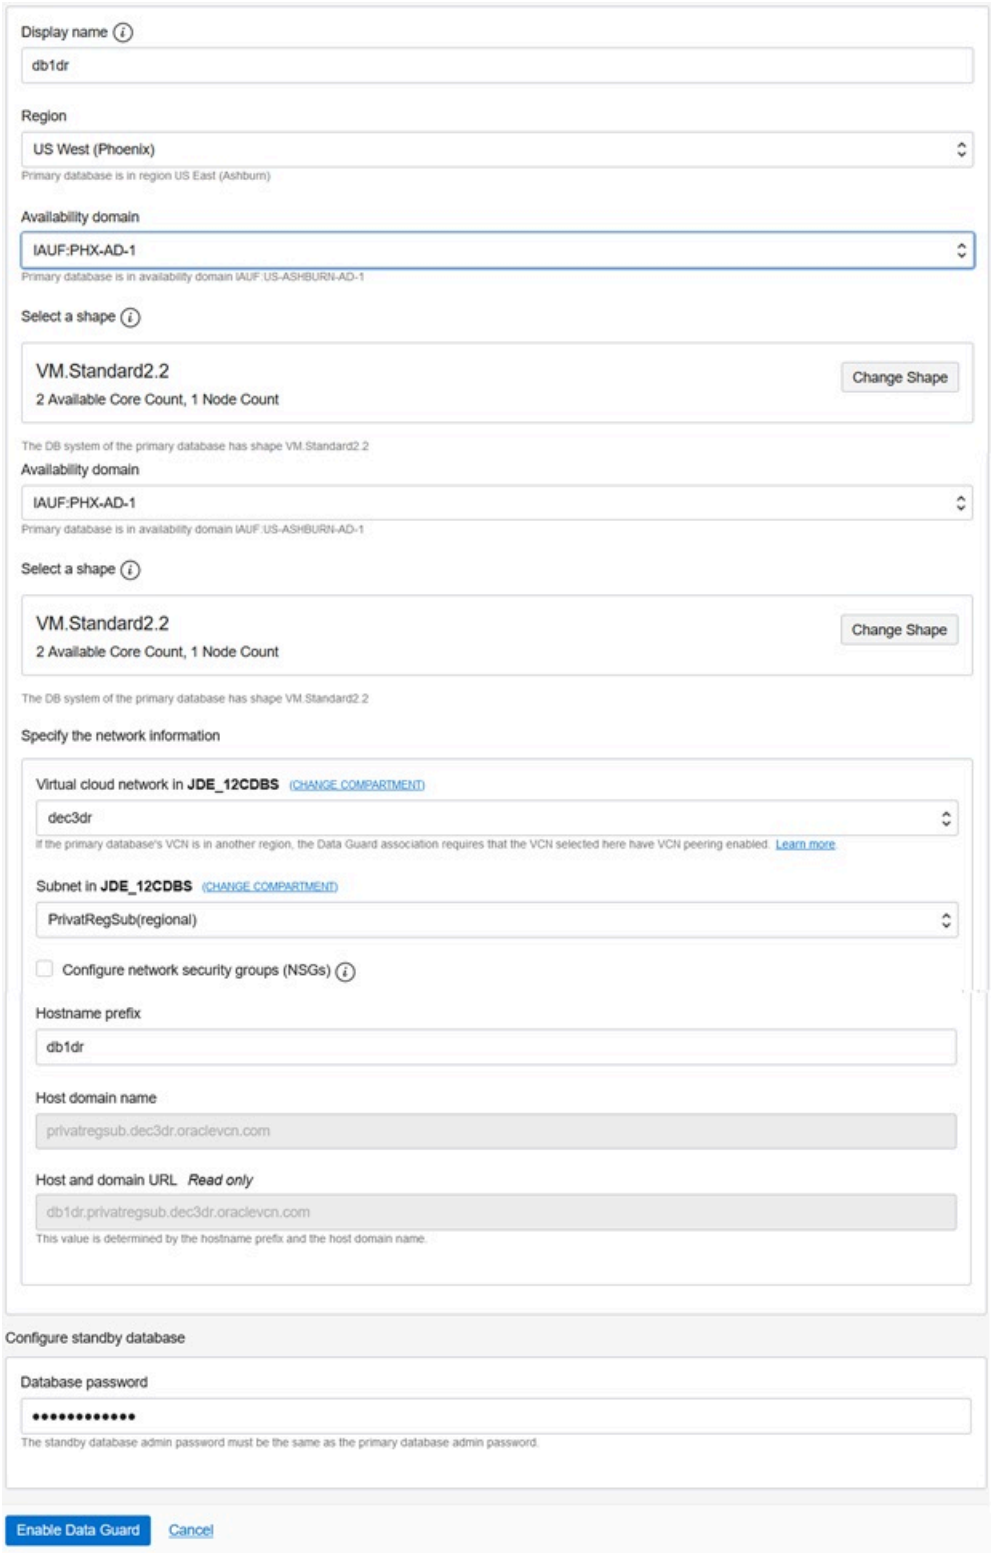

7. Click the Enable Data Guard button to complete the form.

Postinstallation of JD Edwards EnterpriseOne Servers

After you have enabled Oracle Data Guard as described in the preceding procedure, use this procedure to perform postinstallation for these JD Edwards EnterpriseOne servers:

- Enterprise Server
- HTML Server
- AIS Server

Note: For Deployment Server related additional configuration, refer to the section in this Learning Path entitled: Performing Post Infrastructure Provisioning Tasks.

Enterprise Server

1. Determine the DB system server name. For example, jdepddb1 or jdepddb1-scan.
2. Login to Enterprise server instance using Putty.
3. Edit the tnsnames.ora file at below location:
   /u01/oracle19/product/19.0/client_1/network/admin/tnsnames.ora
4. Update the below bolded values for HOST= and SERVICE= using valid values for your DR region:

```
JDEORCL =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP)
 (HOST = jdepddb1dr.privatregsub.vcndr.oraclevcn.com)
 (PORT = 1521))
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = jdepdb.privatregsub.vcndr.oraclevcn.com)
 )
 )
```

5. From the JD Edwards EnterpriseOne Server Manager Console, navigate to the Enterprise Server instance.
6. In the Enterprise Server instance, navigate to Advance Configuration > Database > Database Server Name.
7. For Database Server Name, enter the DB system name for the DR region.
8. If multiple Enterprise Server are in the same DR region, repeat this process for each.
9. Restart each Enterprise Server instance that you modified.

HTML Server

1. Determine the DB system server name. For example, jdepddb1 or jdepddb1-scan.
2. From the JD Edwards EnterpriseOne Server Manager Console, navigate to the HTML Server instance.
3. In the HTML Server instance, navigate to Advance Configuration > Database > Database Server Name.
4. For Database Server Name, enter the DB system name for the DR region.
5. On the same screen, update the TNSNAMES.ora file with the below bolded values for HOST= and SERVICE= using valid values the DR region for your DR region:

```
JDEORCL =
 (DESCRIPTION =
 (ADDRESS = (PROTOCOL = TCP)
 (HOST = jdepddb1dr.privatregsub.vcndr.oraclevcn.com)
 (PORT = 1521))
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = jdepdb.privatregsub.vcndr.oraclevcn.com)
 )
```

6. If multiple HTML Servers are in the same DR region, repeat this process for each.
7. Restart each HTML Server instance that you modified.

AIS Server

1. Determine the DB system server name. For example, jdepddb1 or jdepddb1-scan.
2. From the JD Edwards EnterpriseOne Server Manager Console, navigate to the AIS Server instance.
3. In the AIS Server instance, navigate to Advance Configuration > Miscellaneous > JDBC URL.
4. Update the below setting with values using valid values for your DR region:

```
jdbc:oracle:thin:@jdepddb1:1521/jdepdb1.privatregsub.vcndr.oraclevcn.com
```

5. If multiple AIS Servers are in the same DR region, repeat this process for each.
6. Restart each AIS Server instance that you modified.

Switch Over to the Disaster Recovery (DR) Region

Use this procedure to manually initiate a switch over (also called switchover or failover) to the DR region when the primary region is malfunctioning or is down.

1. Ensure all servers in DR region are up and running.
2. Go to the DB system instance in the active primary region.
3. In the left panel, click Databases to display the list of all of your databases in the region.
4. Click on the database you want to switch over to the DR region. For example, your database might be named ORCL.
5. In the left panel, click Data Guard Associations.
6. On the right side, click the ellipses and select Switchover.

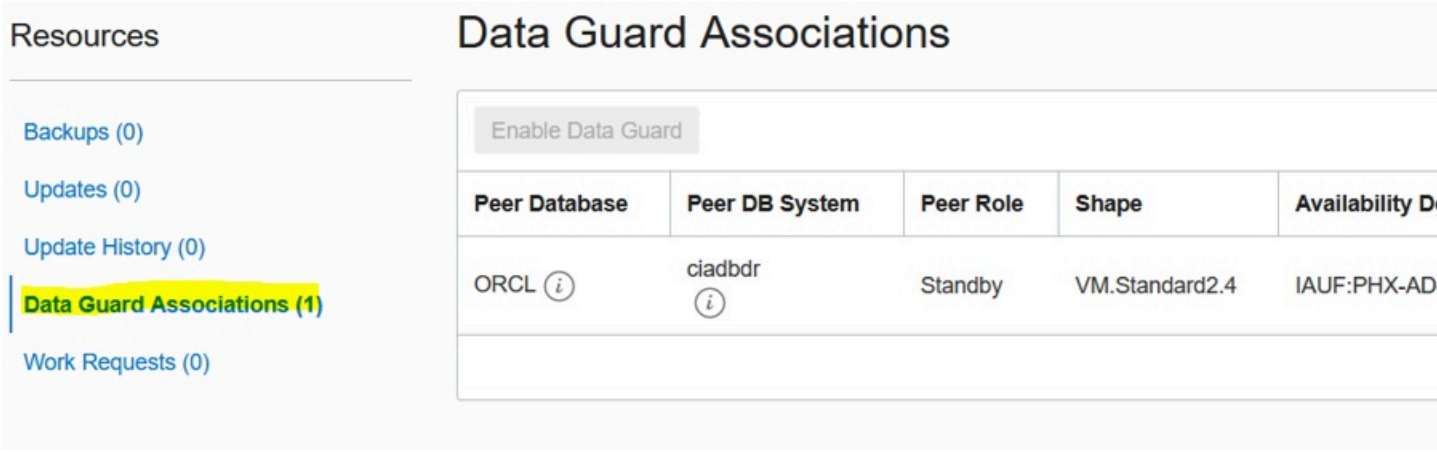

7. The system will take a few minutes to change the DB system that was in the Standby state in the DR region to the Active state.

At this point users can access the HTML and AIS server instances using the ip/port of the load balancer in the DR region.

Limitations of Switch Over

The only instance that was running in primary region and the instance running in DR region that is in sync is the DB System, which is using Data Guard. When the switch over occurs:

- Any active session of web components at the time of the switch over are lost. Users must login again to the switch over instances of in the DR region,
- Any currently running JD Edwards EnterpriseOne business function (BSFN) or batch processes (UBE) at the time of the switch over will fail. To recover, a failed BSFN or UBE must be resubmitted in an instance of DR region.

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.

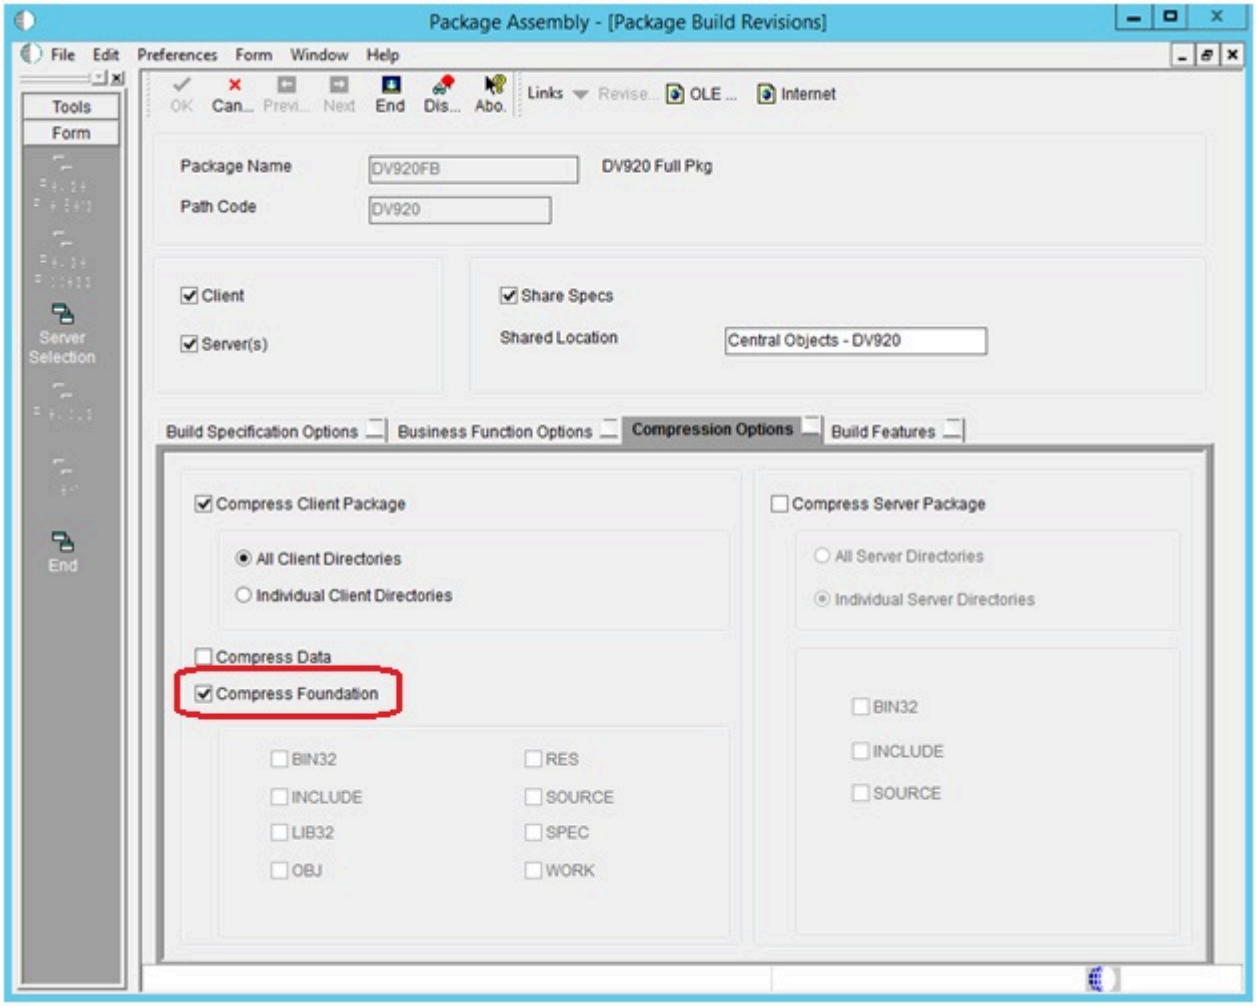

Note: If you encounter NER failures during your full package build (server side only), to resolve these failures refer to the My Oracle Support Doc Id 1950295.1 at this link: *https://support.oracle.com/epmos/faces/DocumentDisplay?id=1950295.1*

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**

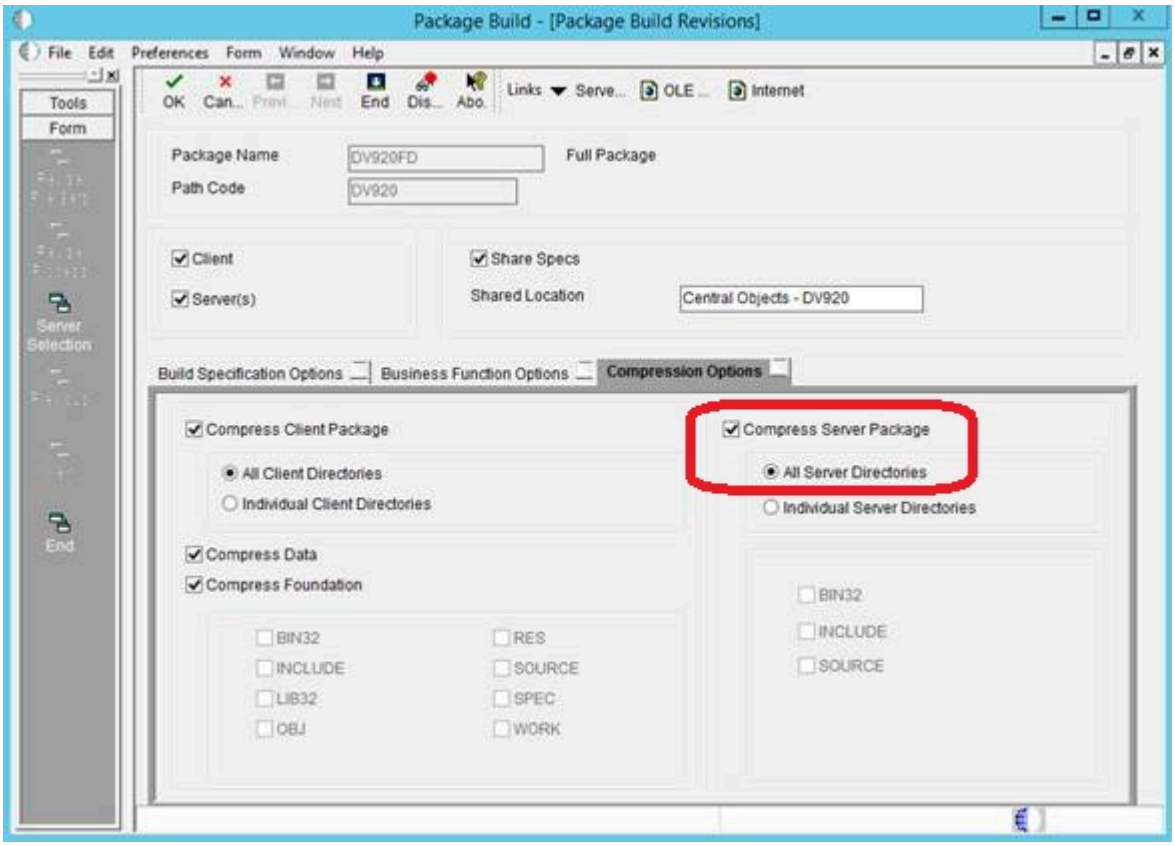

Considerations for Development Client

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:

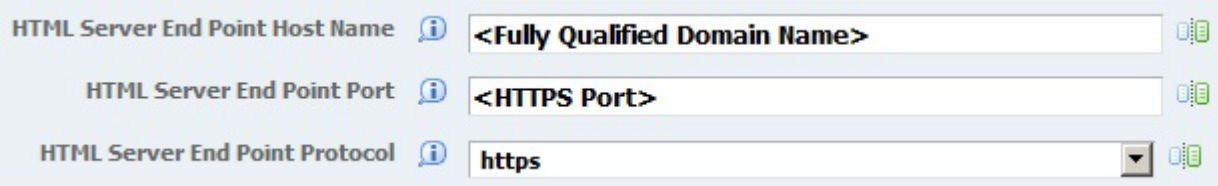

Tip: After you have SSL configured and tested, it is recommended that you disable all the non-SSL ports.

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

Accessing the HTML Server and AIS Server Through the Bastion Host

This section shows you how to access the HTML Server and the AIS Server through the Bastion host. This procedure includes two separate subtasks:

- Accessing the HTML Server Without Using the Load Balancer
- Accessing the HTML Server Using the Load Balancer

Prerequisite

- You must have installed PuTTY on your Microsoft Windows machine. By default, this installation includes the requisite software component called Pageant (Putty Authentication Agent). The program provides a Secure Shell (SSH) tunneling method for connecting to Unix or Linux machines through PuTTY.

Accessing the HTML and AIS Server

This procedure describes how to connect from a Microsoft Windows machine to the JD Edwards EnterpriseOne HTML or AIS Server in a private network in Oracle Cloud Infrastructure through the Bastion host that has been deployed using JD Edwards EnterpriseOne Infrastructure Provisioning.

1. On your Microsoft Windows machine, search for the Pageant application (pageant.exe).

   Note: As mentioned in the section "prerequisite" this program is a standard component of PuTTY for Microsoft Windows.

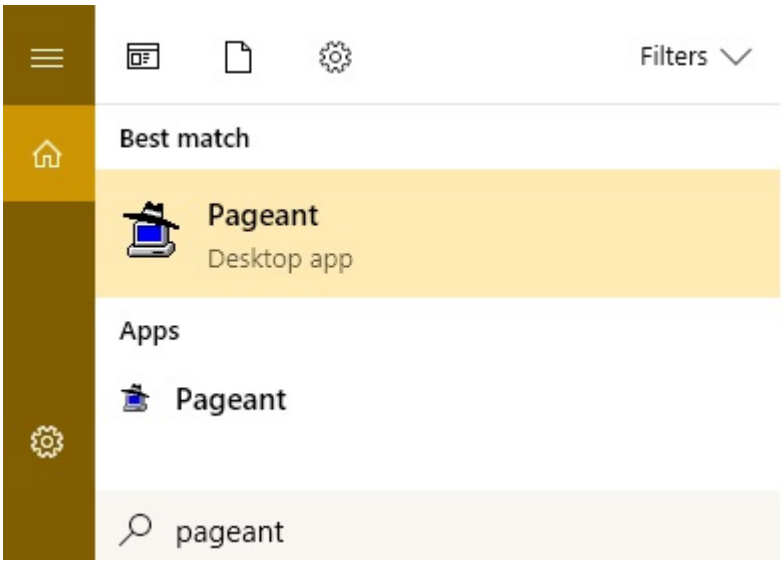

2. Start the Pageant application for the Pageant Key List window to be displayed.

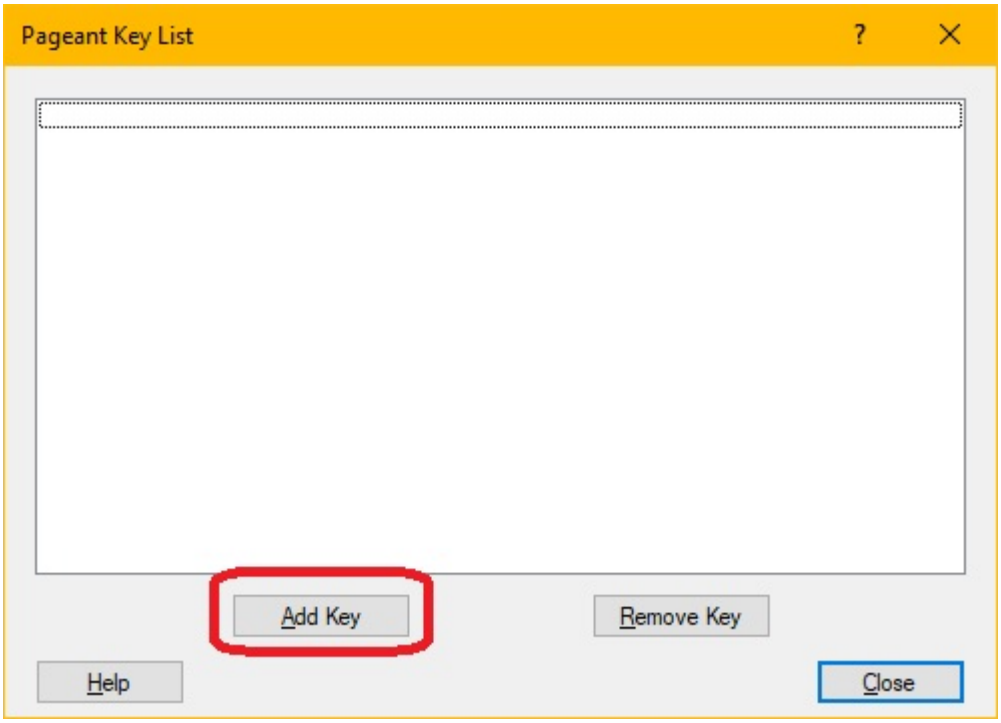

3. Click **Add Key** and browse to the private key you provided to the Infrastructure Provisioning Console, which is used to create Compute instances in Oracle Cloud Infrastructure. For example:

   `OCI_Instance.ppk`

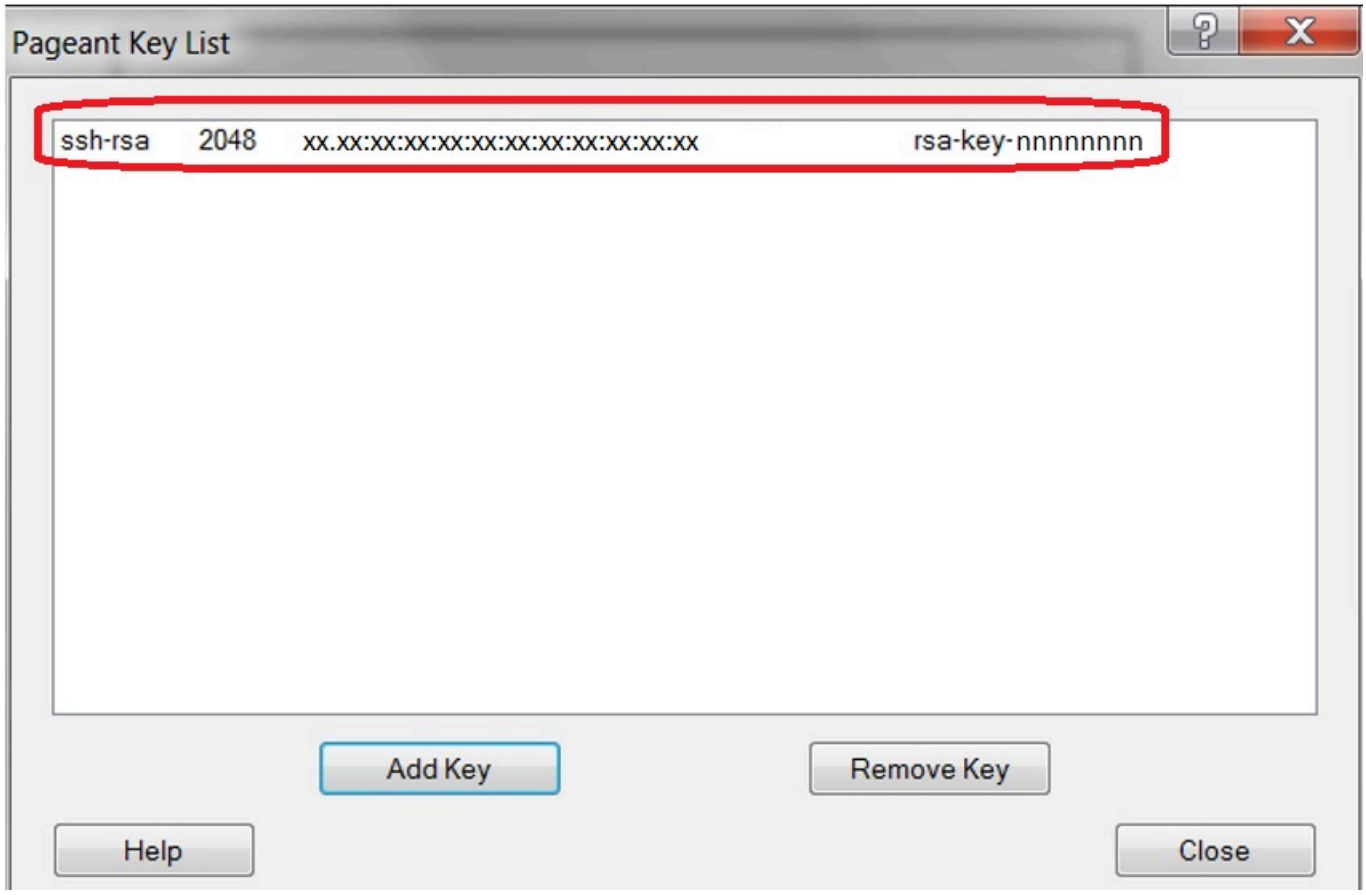

4. Open PuTTY and in the Host Name field, enter the public IP address of the Bastion server (also includes the NAT server and Server Manager Console).

> **Note:** As a best practice, you can save this PuTTY session for future use when logging in to machines in the private network known to this Bastion Host.

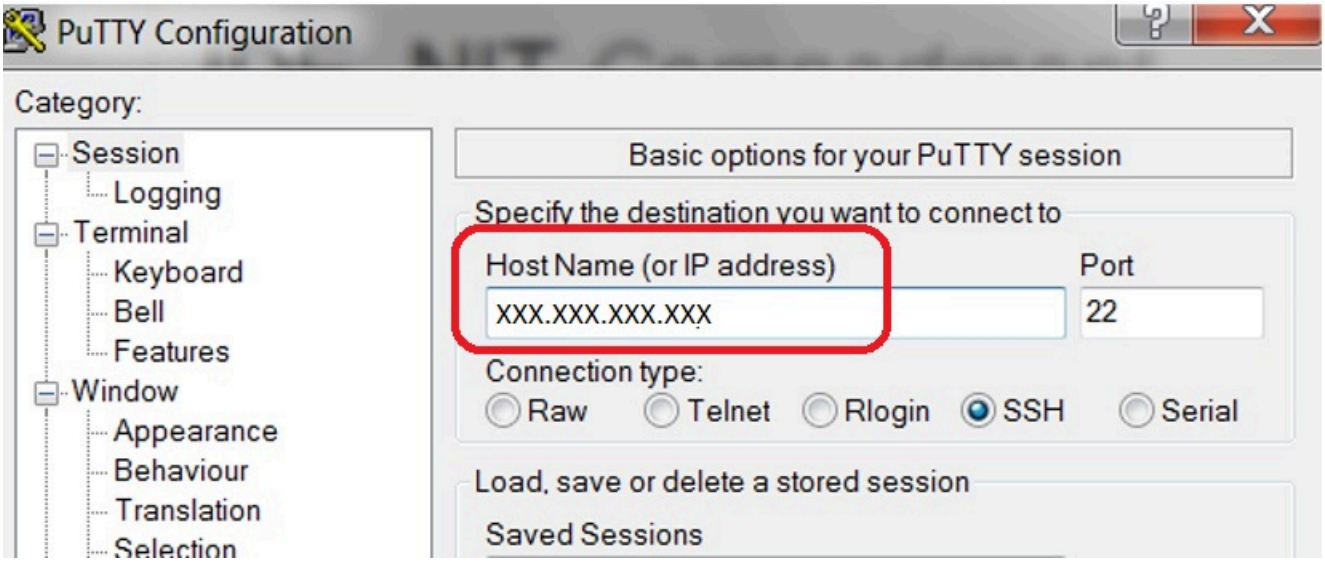

5. To create an SSH tunnel to the local host, in the **Category** section, expand the **Connections** node, expand the **SSH** node, and click **Tunnels**.
6. On the window for the Options controlling SSH port forwarding, enter the port number in the Source Port field. You can use any port that is free on your local machine.

7. The subsequent procedure differs based on whether you want to access these servers using or without using a load balancer.

Accessing an HTML or AIS Server Without Using a Load Balancer

a. On the PuTTY Configuration window, in the **Destination** field, enter the private IP address of the HTML or AIS Server and the port on which the JD Edwards EnterpriseOne HTML or AIS Server is listening.
b. Click the **Add** button to add the port.
c. Similarly, one by one you can add ports to as many HTML and AIS servers as you need to access.

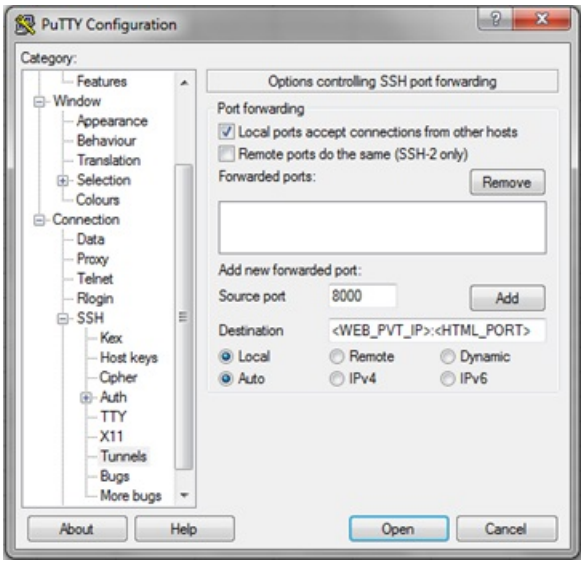

Accessing an HTML or AIS Server using Load Balancer

a. In the **Destination** field, enter the private IP address of the WebLogic load balancer and the port on which the load balancer is listening for the JD Edwards EnterpriseOne HTML or AIS Server.

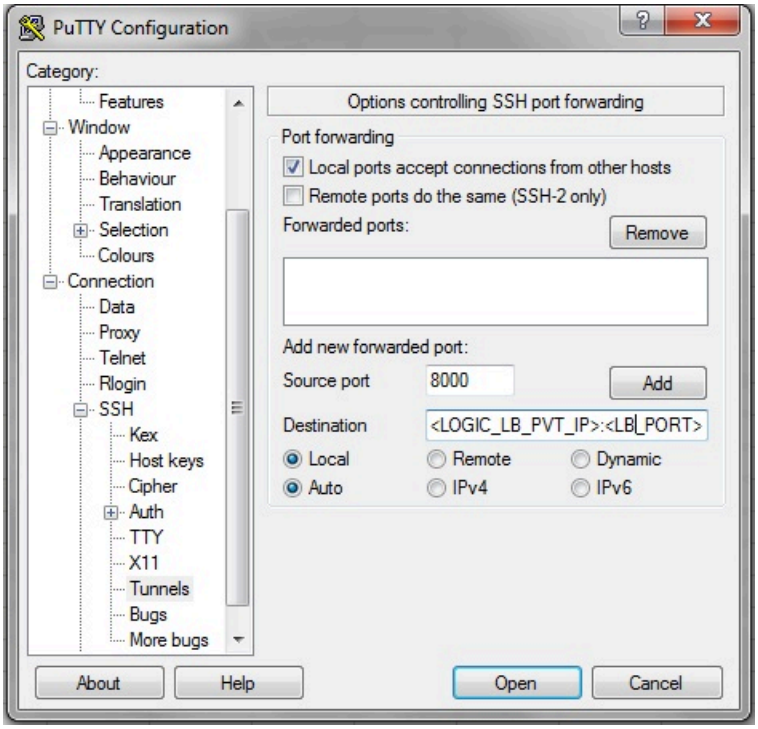

LOGIC_LB_PVT_IP

You can obtain this private IP address value from this file on the Infrastructure Staging machine:

```
<JDERefArch_InfraProvisioning>/E1InfraProvisionConsole/outputJson/infraOutput.json
```

LB_PORT

This value of this port was set when entered it in the Infrastructure Provisioning Console as the LBaaS Listen Port for HTML. You can obtain this port value from the Oracle Cloud Infrastructure user interface as shown below:

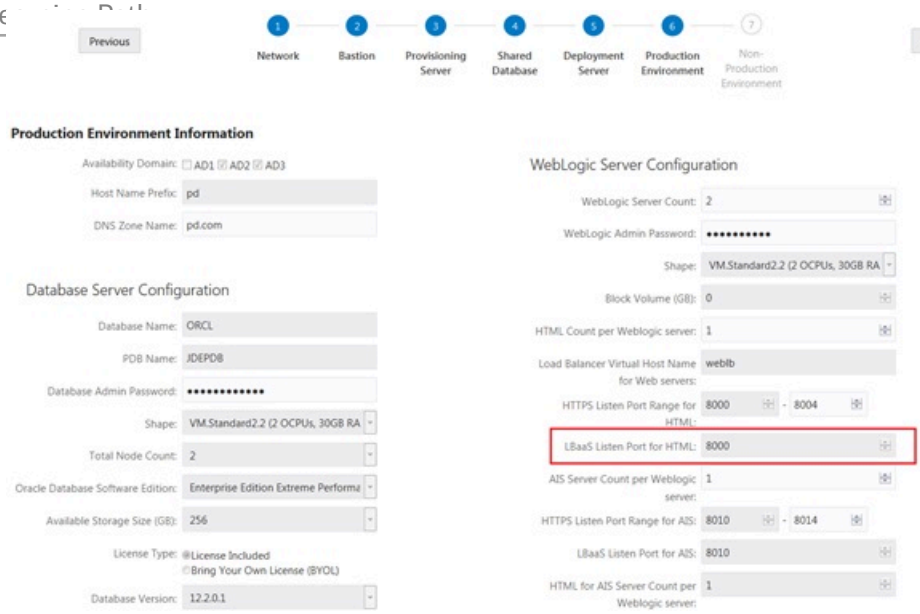

b. On the PuTTY Configuration window, click the **Add** button to add the port (if you are not using the load balancer) or the IP address and the port (if you are using the load balancer).
c. Similarly, one by one you can add ports to as many HTML and AIS Servers or load balancers as needed.

8. Ensure that these options are selected:
   - **Local ports accept connections from other hosts**
   - **Local**
   - **Auto**
9. After adding all the entries, save these changes and open a new PuTTY session, which will be using the previously saved changes from the preceding steps.
10. On your local machine, you can enter the following URL to open a Putty session.

    For example, the below URL is for an HTML or AIS Server:

    ```
    https://localhost:8000/jde
    ```

    In the above URL, 8000 is the source port that is configured in the PuTTY tunnel.

9 Performing Optional Tasks

Starting and Stopping the Infrastructure Provisioning Console

This section shows you how to start and stop the service for the Infrastructure Provisioning Console.

Prerequisite

You must have already created a Linux instance and set up the OpenTofu staging server on it.

Starting the Infrastructure Provisioning Console

On the OpenTofu staging machine on which you have completed all the setup, use this command to start the service for the Infrastructure Provisioning Console:

```
sudo systemctl start E1InfraProvisionConsole.service
```

Stopping the Infrastructure Provisioning Console

On the OpenTofu staging machine on which you have completed all the setup, use this command to stop the service for the Infrastructure Provisioning Console:

```
sudo systemctl stop E1InfraProvisionConsole.service
```

Using Reentrant Mode

This section shows you how to use the reentrant mode of the Infrastructure Provisioning Console for these use cases:

1. Infrastructure Provisioning Failed
2. Scaling Up (Adding) Servers

Prerequisite

- You must have already created a Linux instance and set up the OpenTofu staging server on it.

Using Re-entrant Mode

Two scenarios are supported for the reentrant mode of the Infrastructure Provisioning Console:

- Infrastructure Provisioning Failed
- Scaling Up (Adding) Servers

Infrastructure Provisioning Failed

In this case, the user has previously started the infrastructure provisioning process by clicking the **Finish** button in the Infrastructure Provisioning Console, and upon execution, the infrastructure provisioning process failed at

some point and did not complete as expected. An example of a failed process is shown in the screenshot below:

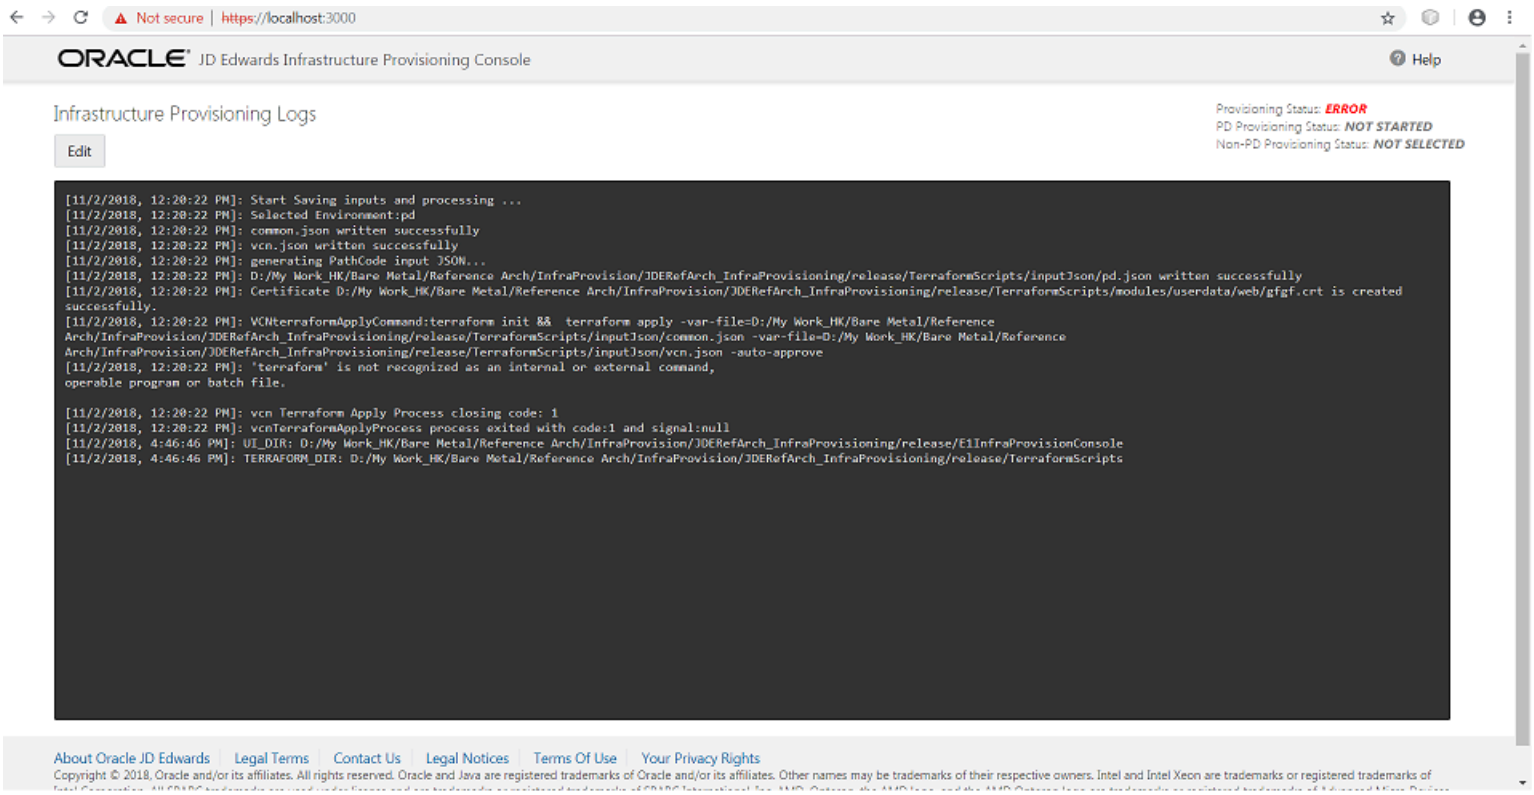

To invoke the reentrant mode, click the Edit button to return to the input screen of the Infrastructure Provisioning Console. At this point you can navigate to the final screen in the Provisioning Console and again click the Finish button. The infrastructure provisioning process will resume from the point where it previously stopped.

Scaling Up (Adding) Servers

If you have completed infrastructure provisioning successfully, you can reenter the Infrastructure Provisioning Console to scale up (add) servers.

1. Start the Infrastructure Provisioning Console on the OpenTofu staging server. For this tutorial the server is located here:
   ```
   https://localhost:3000
   ```
2. After the Infrastructure Provisioning Console is running, if required you can increase the server count in the input screen. When all the counts are set, click the Finish button to provision additional servers.

   > **Note:** When you are adding Enterprise Servers or HTML Web Servers, you should expect downtime for the existing servers that are running until the infrastructure provisioning is complete.

Creating an Oracle Web Application Firewall (WAF) Policy

This section shows you how to configure an Oracle Cloud Infrastructure service for a Web Application Firewall (WAF) policy for use with Oracle JD Edwards EnterpriseOne.

WAF is designed to protect applications from malicious and unwanted internet traffic with a cloud-based, PCI-compliant, global web application firewall service. By combining threat intelligence with consistent rule enforcement on Oracle Flexible Load Balancer, Oracle Cloud Infrastructure Web Application Firewall strengthens defenses and protects internet-facing application servers and internal applications. For additional information, refer to this Oracle Cloud Infrastructure documentation: *Web Application Firewall Policies*.

This section describes how to create a WAF policy with these characteristics:

1. Basic Information
2. Access Control
3. Rate Limiting
4. Protections
5. Select Enforcement Point
6. Review and Create the WAF Policy

Creating an Oracle Web Application Firewall (WAF) Policy

In Oracle Cloud Infrastructure, search for "Web Application Firewall".

1. Click "Create WAF Policy”.
2. On **Basic information** complete fields for name and compartment.
3. You also need to define the Actions as per your requirement.

4. Click the **Next** button.

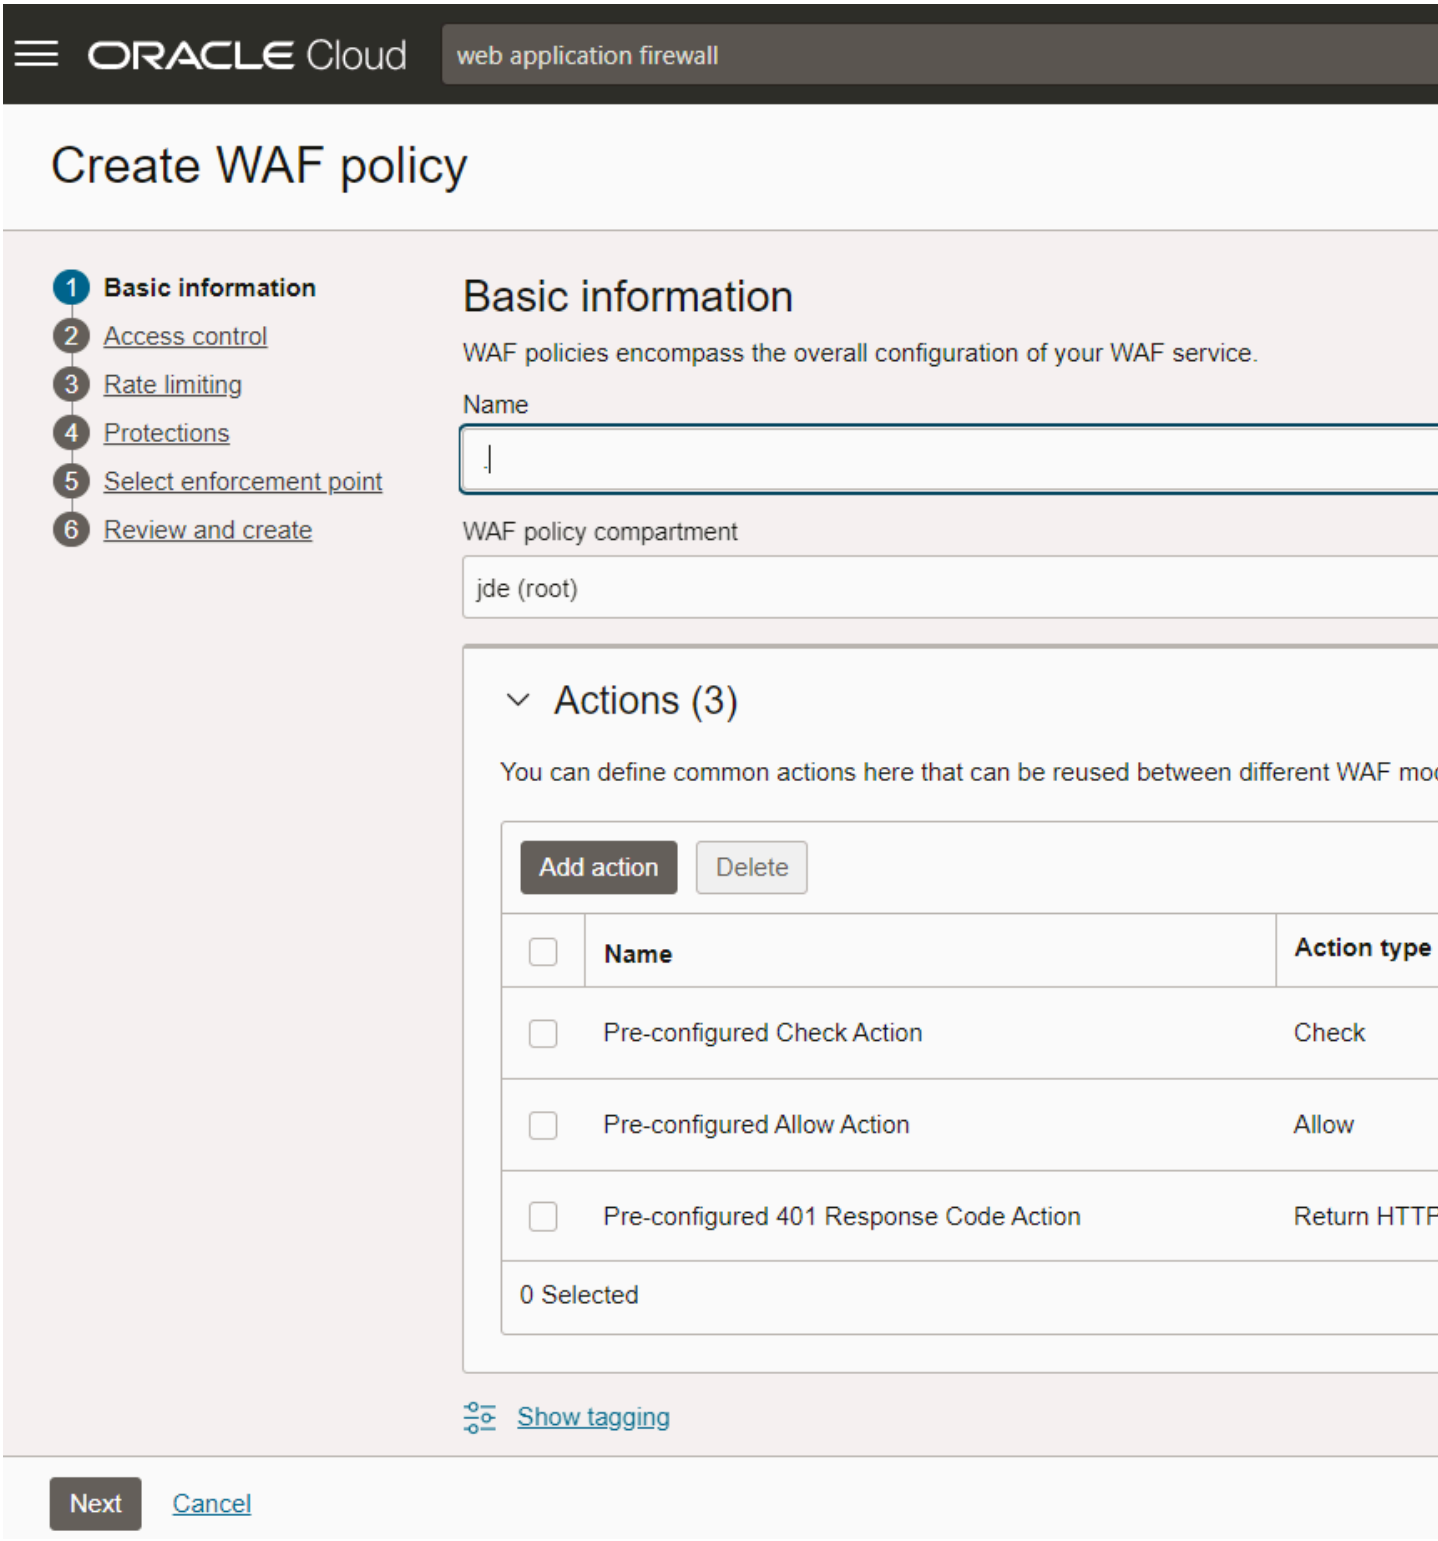

5. On Access Control (optional), if you want to add access rules, click the **Add access rule** button.

> **Note:** Since access control requirements vary based on organizational security policies, Oracle recommends that customers configure rules based on their specific requirements.

6. Click the **Next** button.

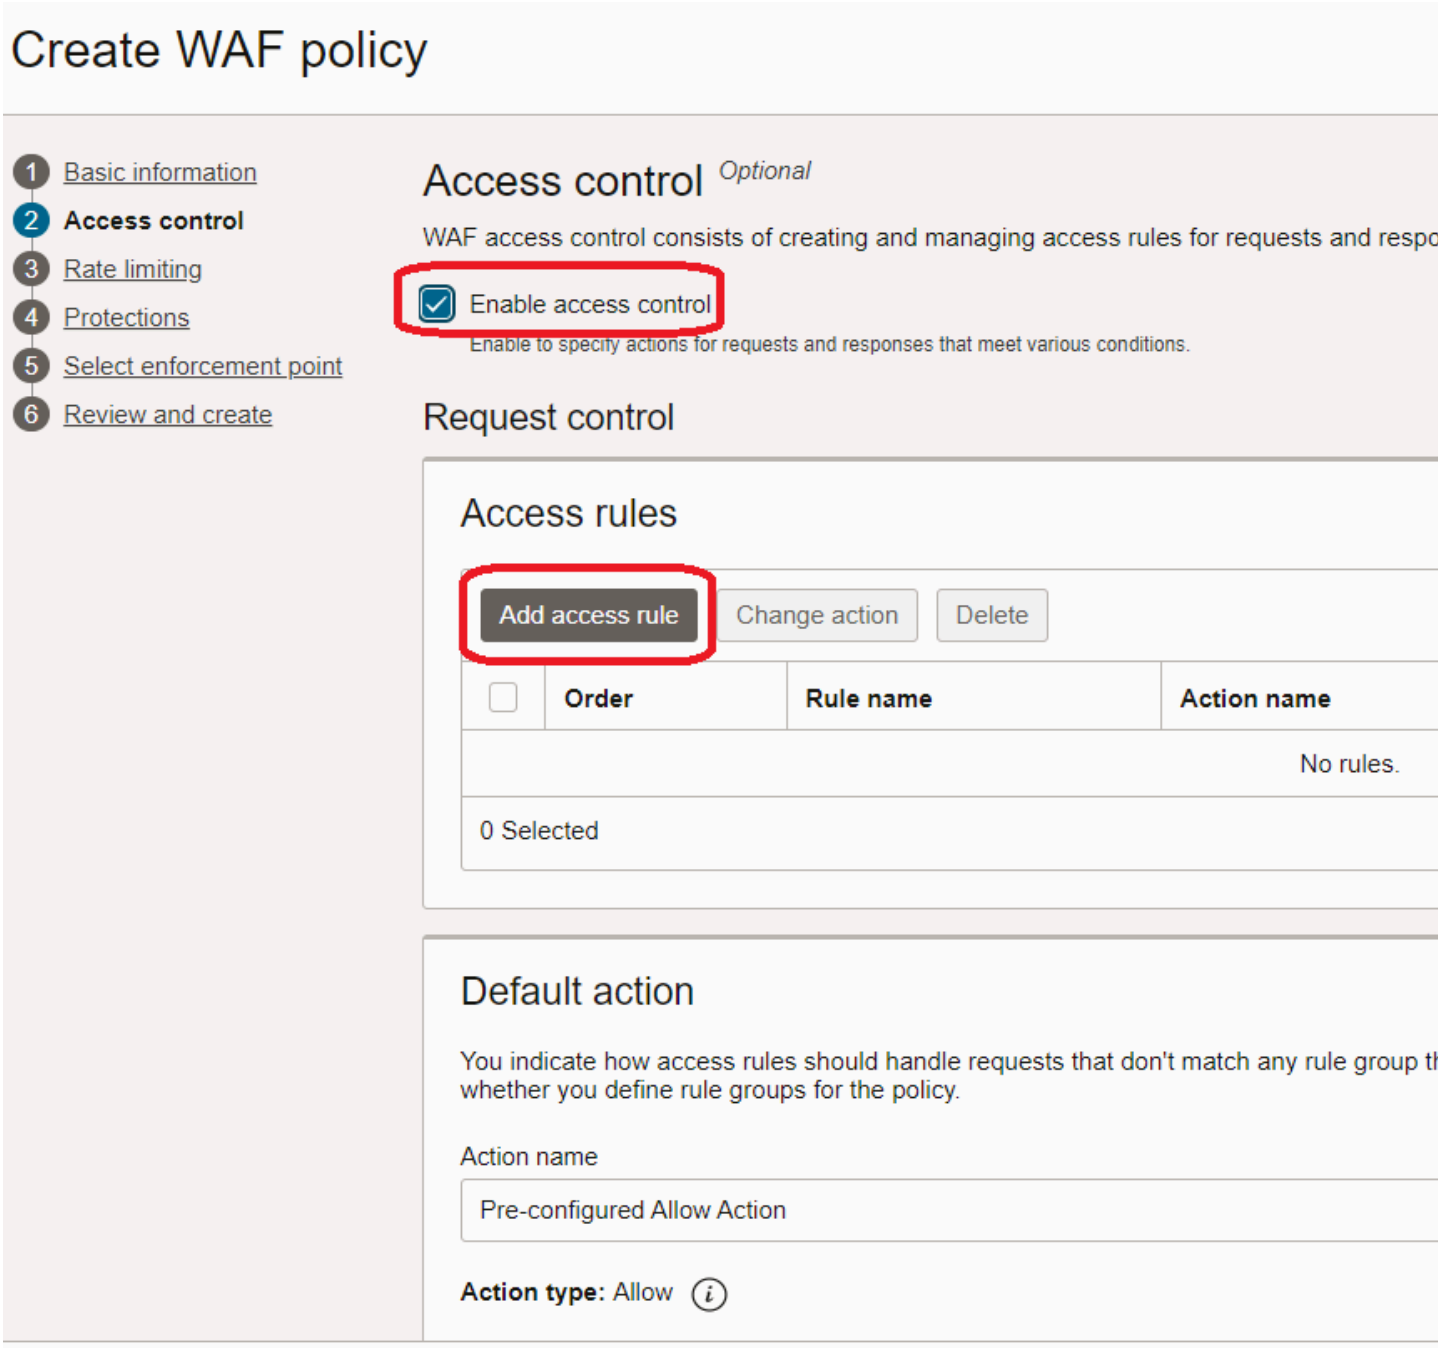

7. On Rate Limiting, if you want to allow inspection of HTTP connection properties and limit the frequency of requests for a given key., click the **Add rate limiting rule** button.

> **Note:** Since rate limiting requirements vary based on organizational security policies, Oracle recommends that customers configure rules based on their specific requirements.

8. Click the **Next** button.

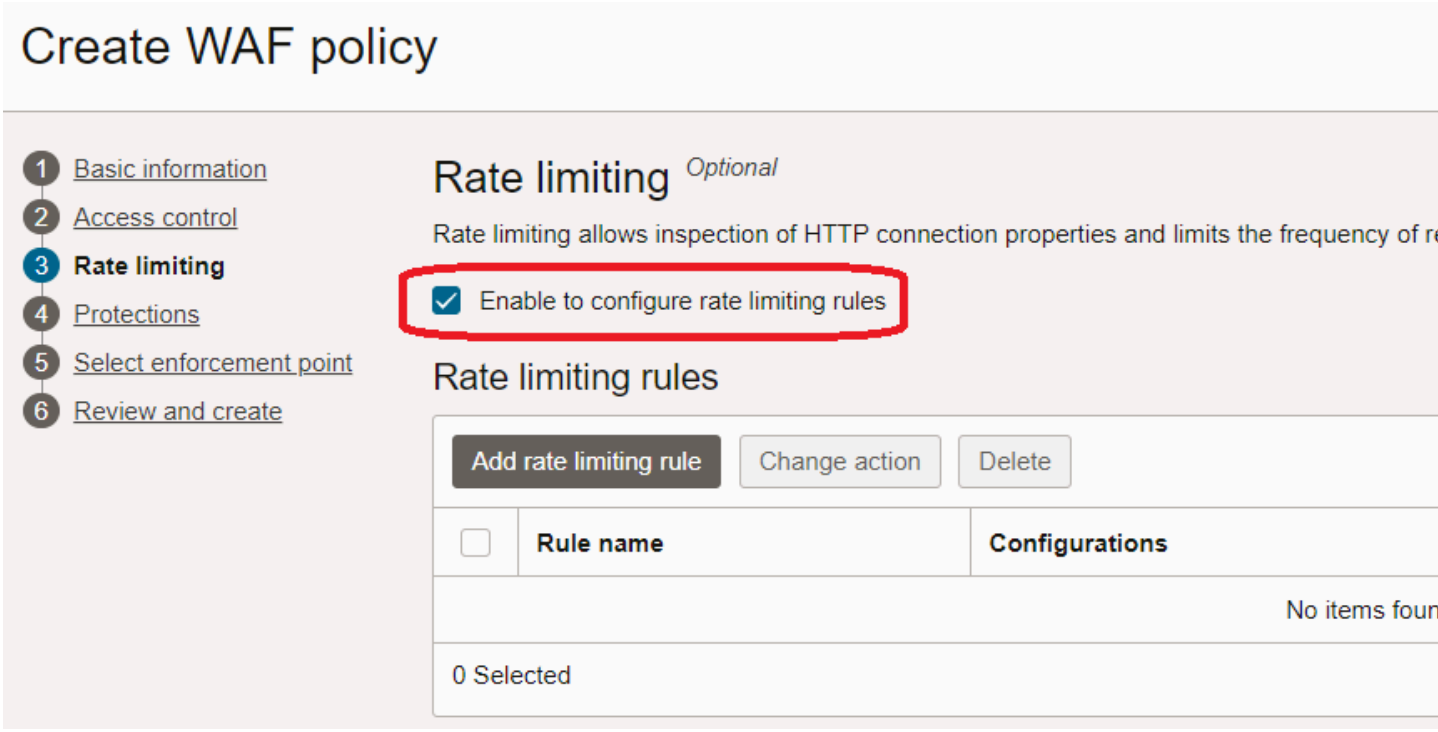

9. On Protections, you can create rules to determine if a network request is allowed but logged, or is blocked entirely.

   For use with JD Edwards EnterpriseOne, Oracle recommends that you configure Protection Capabilities as follows:

   - Click the **Enable to configure protection rules** button.
   - Select the Condition type for the protection rules you are adding.

Define all the rules you want to set for incoming requests by clicking the **Choose Protection Capabilities** button.

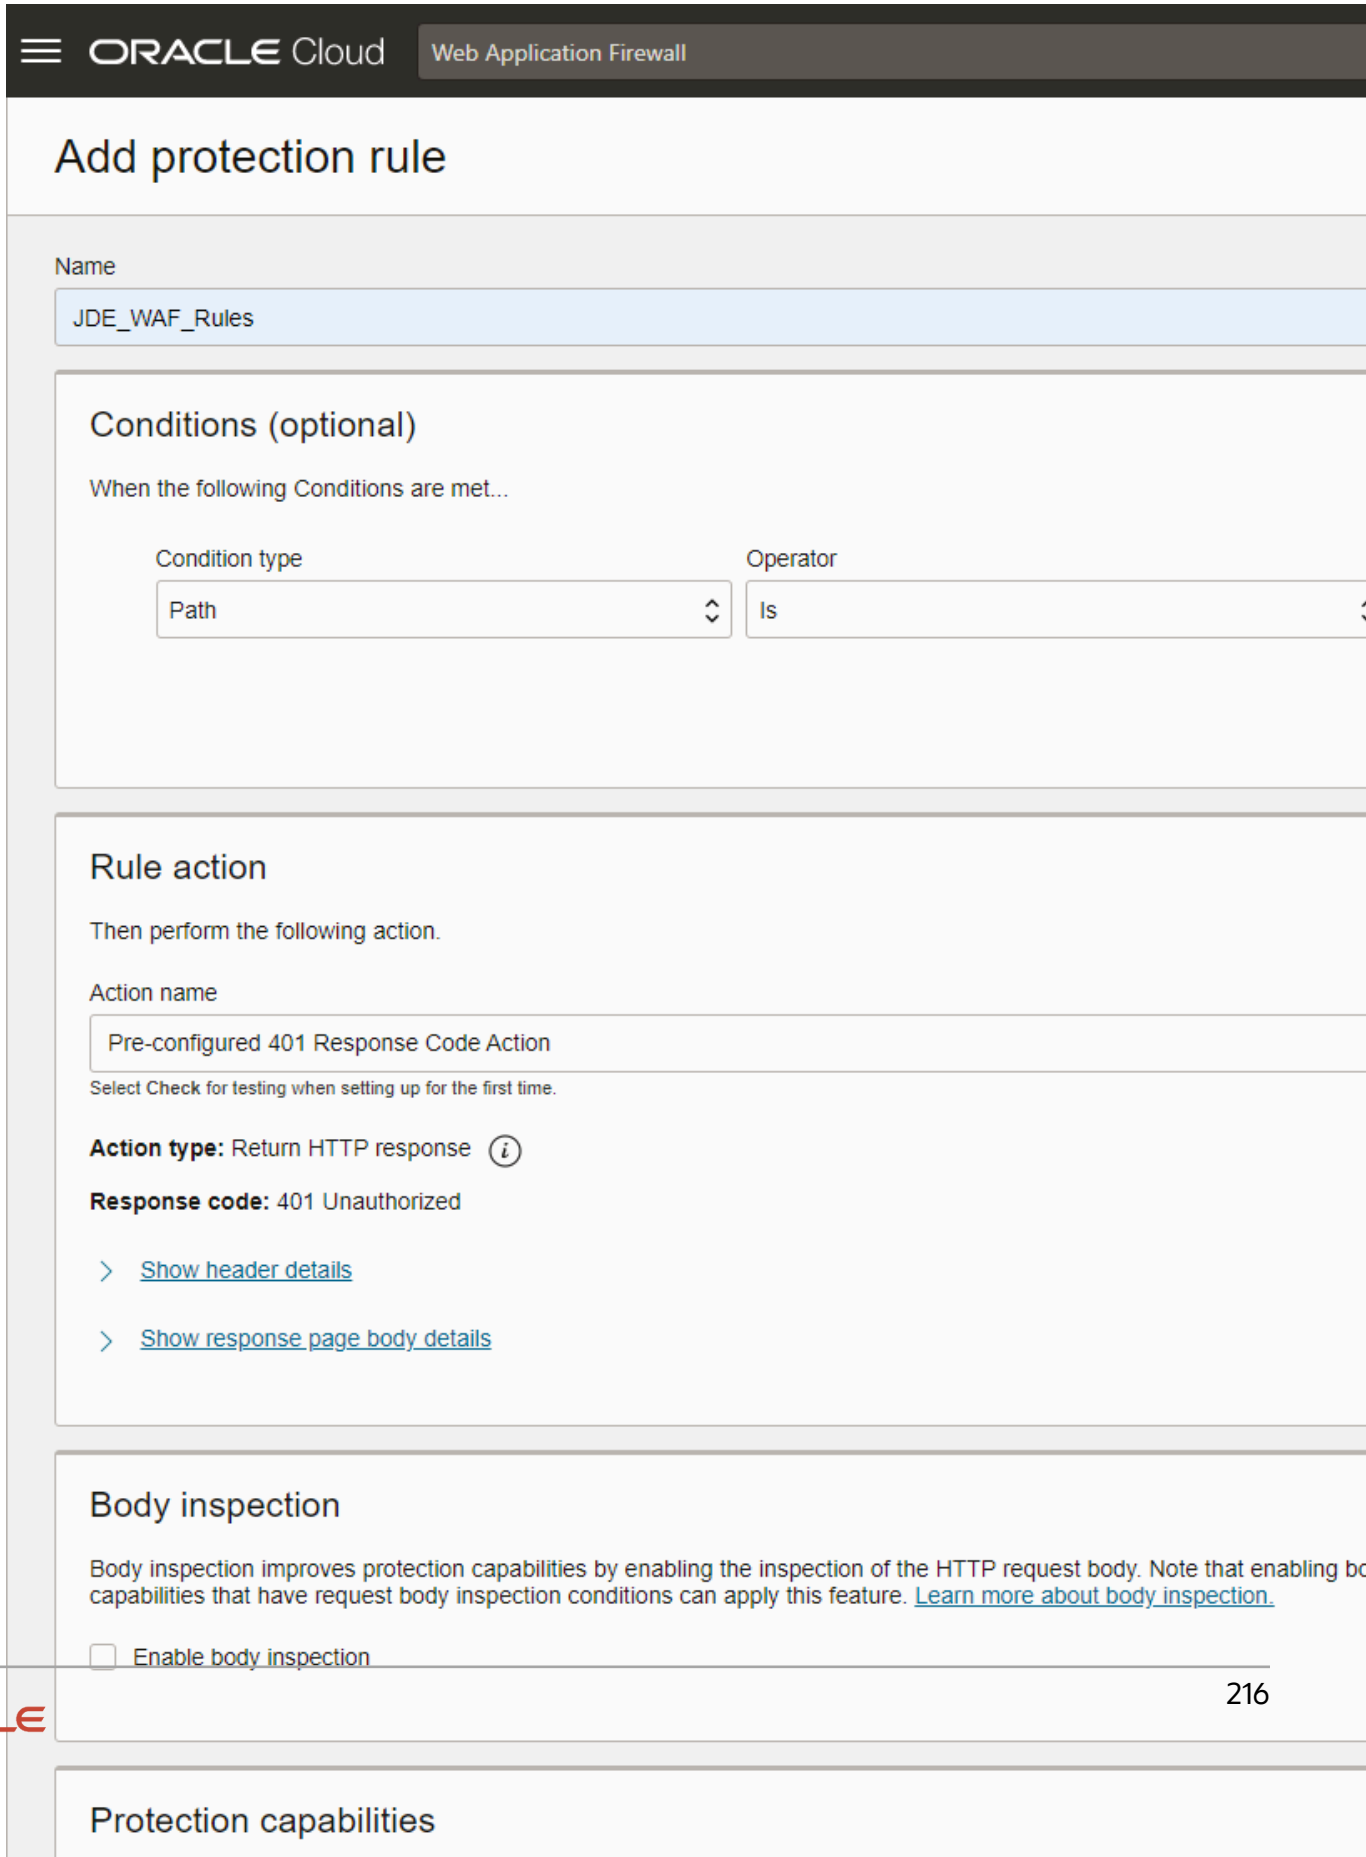

- As shown in the screen shot above, the recommended set of Protection Capability IDs for JD Edwards EnterpriseOne includes:
  - 9420000
  - 9410000
  - 9330000
  - 9320001
  - 9320000
  - 930120
  - 9300000
  - 920390
  - 920320
  - 911100

> **Note:** Since protection requirements vary based on organizational security policies, Oracle recommends that customers configure rules based on their specific requirements. For any rules that you create, Oracle recommends that you follow all the system-suggested actions for a triggered rule based on the condition you are setting.

10. Click the **Next** button.

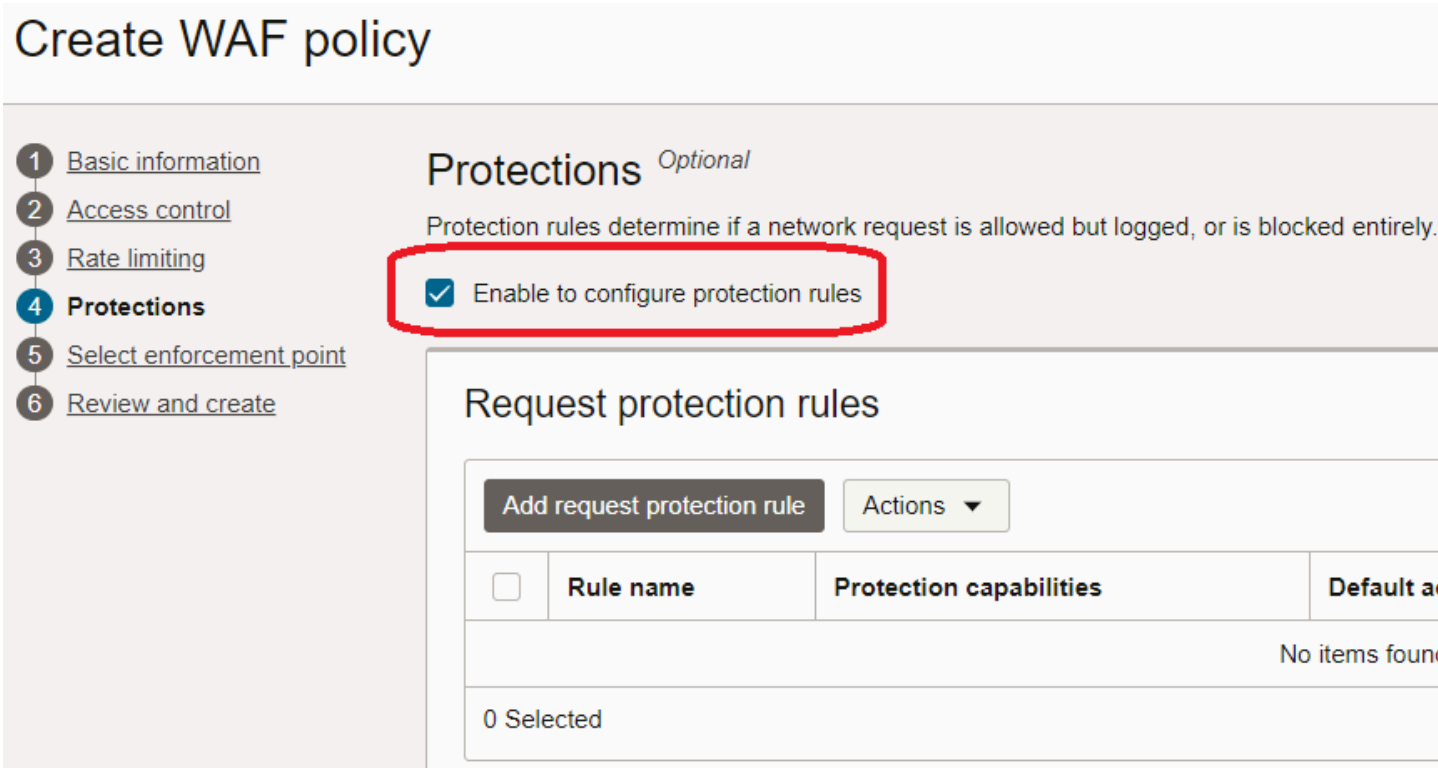

11. On Select Enforcement Point, you can configure your Load balancer with an HTTP listener. Additionally, you can add firewalls by selecting the specific in-region application delivery resources to secure.

> **Note:** Since enforcement requirements vary based on organizational security policies, Oracle recommends that customers configure this functionality based on their specific requirements.

12. Click the **Next** button.

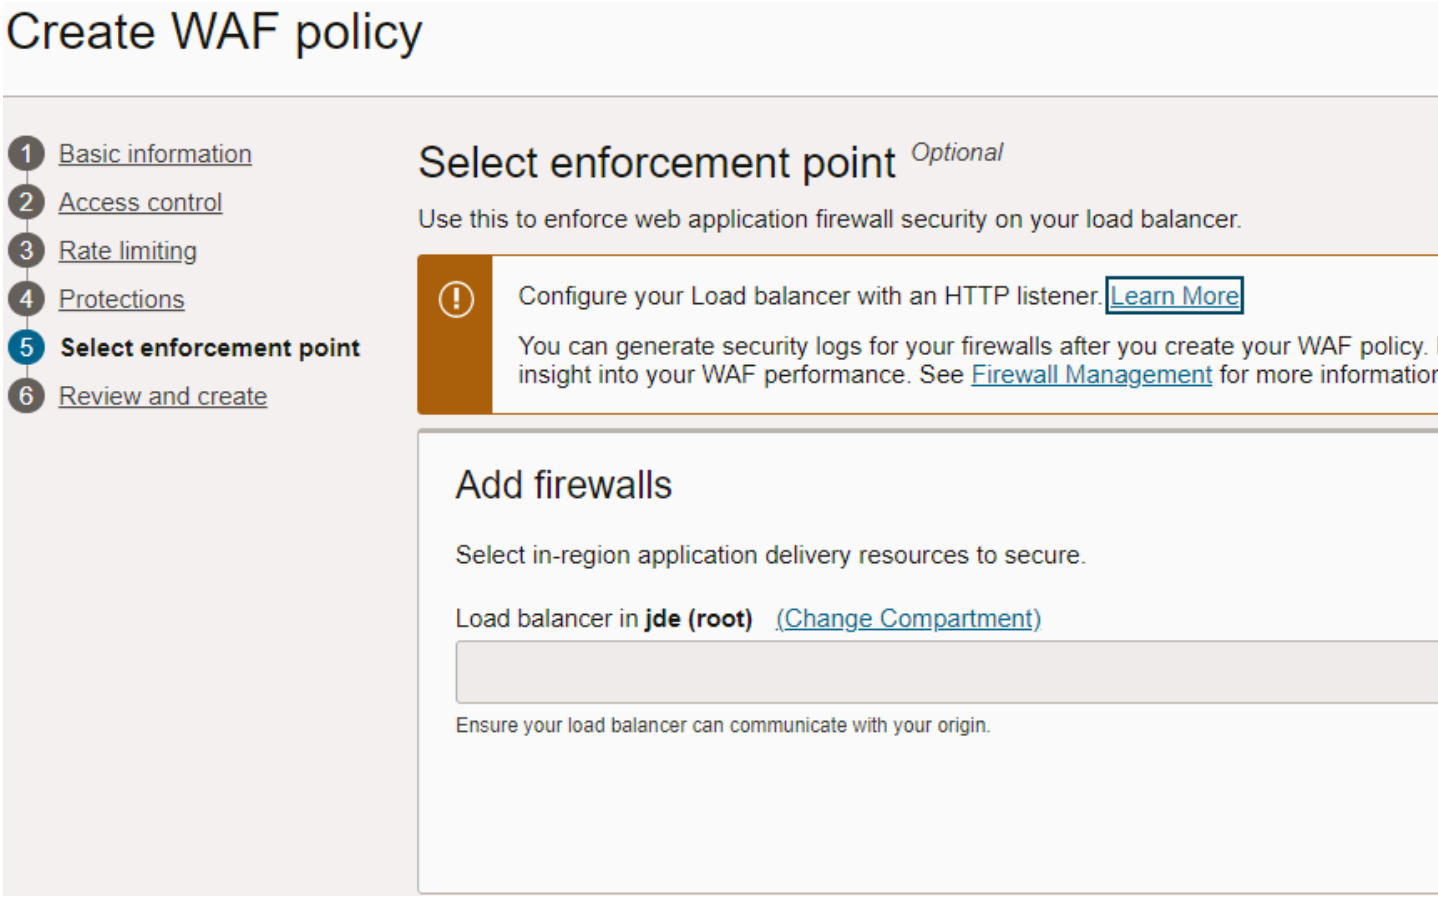

13. On Review and Create, review your settings and when approved, click the **Create WAF policy** button to complete the process.

10 Troubleshooting Your Infrastructure Provisioning Deployment

Troubleshooting Your Infrastructure Provisioning Deployment

This section lists the troubleshooting topics and resolutions.

Finding the Infrastructure Provisioning Logs

The infrastructure provisioning log files on the OpenTofu staging server are located in this directory:

`/home/opc/JDERefArch_InfraProvisioning/E1InfraProvisionConsole/logs`

Viewing the One-Click Provisioning Console Logs

To view the logs for the One-Click Provisioning Console:

1. You must use a private IP address to connect to the Provisioning Console running on the One-Click provisioning server.
2. To connect to the provisioning server, use the procedure explained in the section “Connecting to a Linux Host in a Private Network Through the Bastion Host” which is located in the OBE "Connecting to a Host in a Private Network Through the Bastion Host."
3. After a connection is established with the provisioning server, run the following command in the shell to view the logs for the One-Click Provisioning Console:

   $ sudo journalctl -u E1CloudConsole
4. Press Page-Down to scroll through the logs.
5. Enter Shift + G to scroll to the end of the log.

Note: The One-Click Provisioning Console logs are completely regenerated each time the console is run.

Finding the One-Click Provisioning Deployment Logs

Use this procedure to create a secure FTP (SFTP) connection from your Microsoft Windows workstation to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server.

This example in this procedure uses a program for Microsoft Windows called WinSCP. However, you can use any SFTP program to make the connection.

1. Start your SFTP program, such as WinSCP.

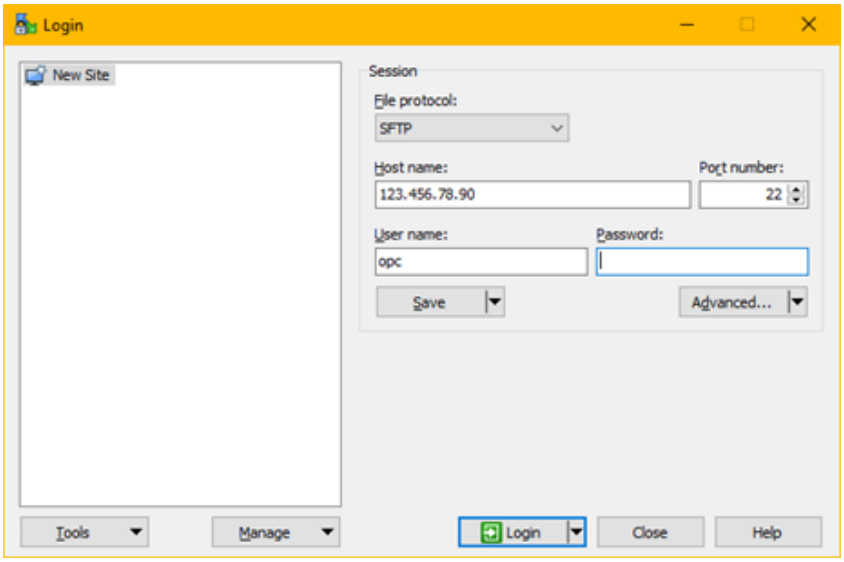

2. On the Login window, complete these fields:
   - **File protocol**

     SFTP
   - **Host name**

     Enter the private IP address of the provisioning server.
   - **Port**

     Enter the value 22.
   - **User name**

     Enter the value opc.
   - **Password**

     At this point you cannot enter a password in the FTP program to complete the connection. Instead you must use the Advanced option and wait until prompted to enter a private key password as described in Step 3 below.

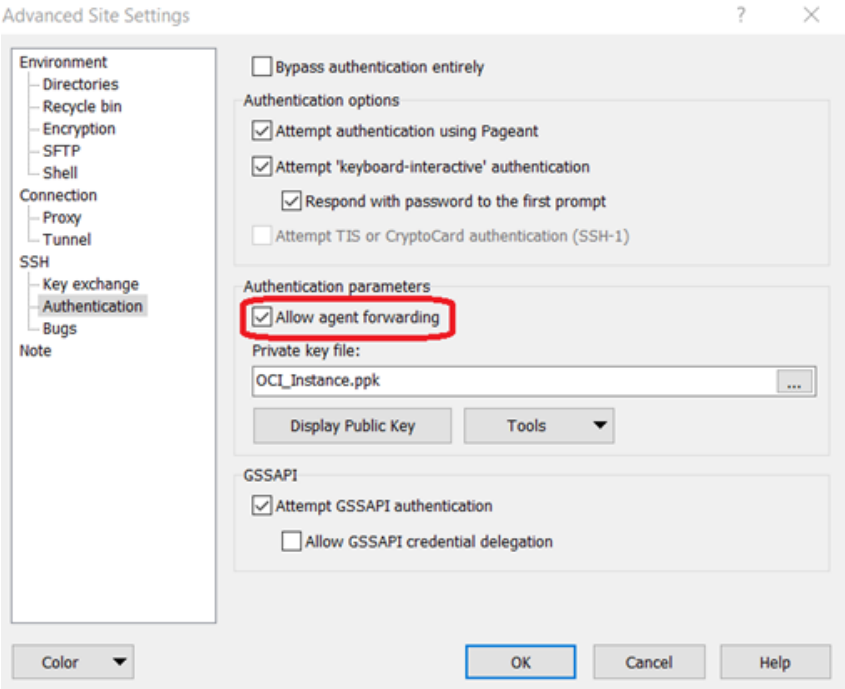

   - Advanced

     No changes are required in this section for this purpose.
3. In the tree structure in the left pane, navigate to SSH and then Authentication.
4. In the Private key file field, use the browse button to locate the private key file. If you followed the recommendation in the OBE "Connecting to a Host in a Private Network Through the Bastion Host", this file is named **OCI_Instance.ppk**.

   > **Note:** Ensure that you have checked the option Allow agent forwarding under Authentication parameters.

5. In the tree structure in the left pane, navigate to Connection and then Tunnel.

6. Complete the following fields:
   - **Host name**

     Enter the Public IP address of the Bastion Host.<Enter the public IP address of the Bastion host.>
   - **User name**

     Enter the value opc.
   - **Private key file**

     Use the Browse button to locate the Private <private> key file. If you followed the recommendation in the OBE "Connecting to a Host in a Private Network Through the Bastion Host", this file is named Bastion.ppk.

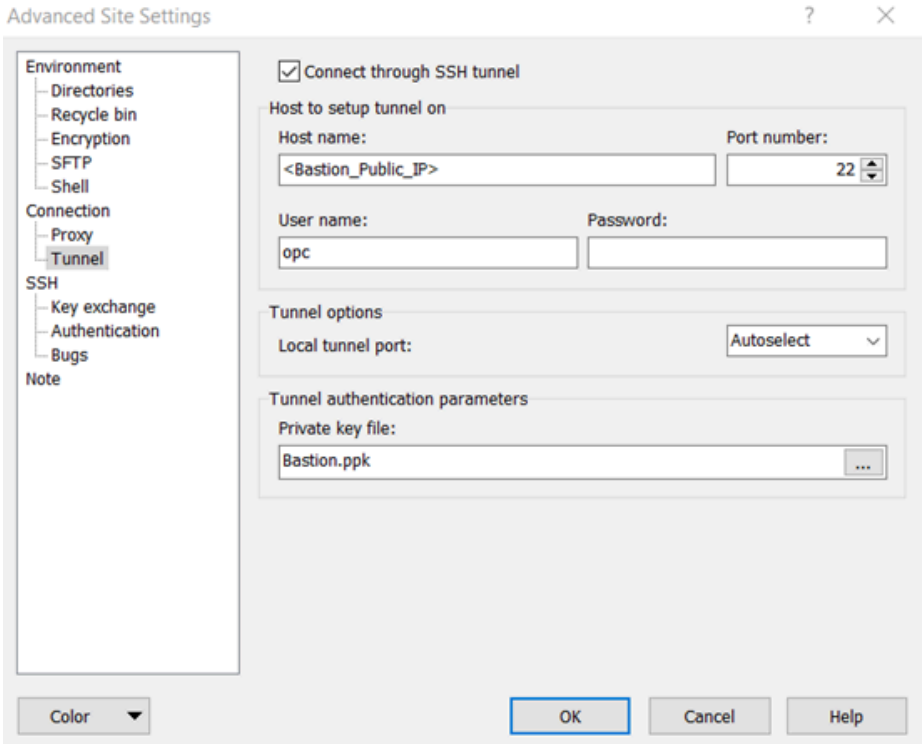

7. Click the **OK** button to save the setting and return to the Login screen.
8. On the Login screen with all the above fields completed, click the **Login** button. After the SFTP program establishes a connection to the OpenTofu staging server, the following warning is displayed:

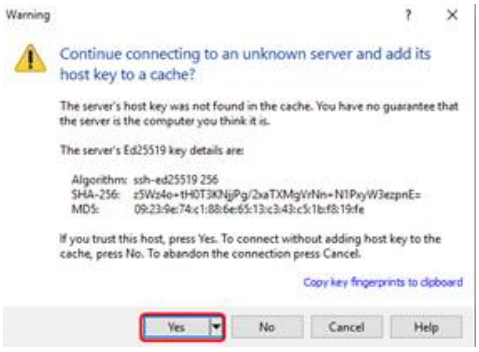

9. On the Warning dialog box, click the **Yes** button to accept the connection to the “unknown server” to add the host key of this server to the cache.

10. After the connection is made, the SFTP program is ready to be used for file transfers from your workstation to the JD Edwards EnterpriseOne Reference Architecture OpenTofu Staging Server. For example:

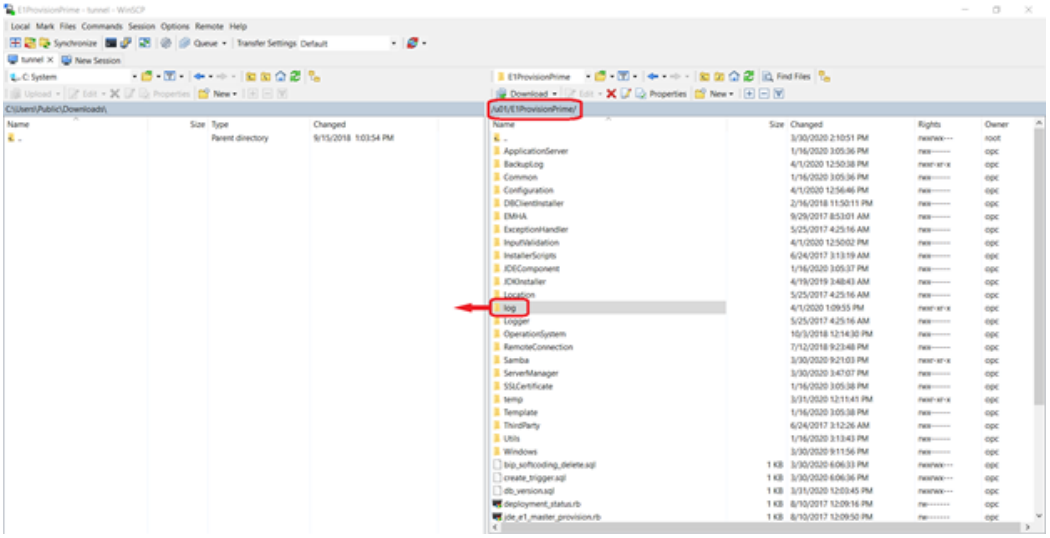

11. You can locate logs in this directory:

    `/u01/E1ProvisionPrime/log`

12. To inspect the logs, you can copy the directory to your local workstation.

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.

Cleaning Up a Deployment of Infrastructure Provisioning

This section shows you how to clean up your Oracle Cloud Infrastructure environment after a successful or failed deployment of JD Edwards EnterpriseOne infrastructure provisioning.

Removing All Instance Resources for a Production Environment

Use this procedure to remove all the instance resources in Oracle Cloud Infrastructure that were provisioned by the deployment of JD Edwards infrastructure provisioning for a production environment in a single availability domain (AD).

1. Navigate to this directory on your OpenTofu staging server:
   `JDERefArch_InfraProvisioning/TerraformScripts/pd`
2. Run this command as a single contiguous line command (line breaks below added for clarity):

```
tofu destroy
-var-file=../inputJson/common.json
-var-file=../inputJson/pdsubnet.json
-var-file=../inputJson/pd.json
-var db_admin_password='<dbpwd>'
-var dep_password='<depPwd>'
-var dep_subname='depSubName'
```

> **Note:** In the above command, the values `db_admin_password`, `dep_password`, and `dep_subname` do not have to be the literal values. That is, you use any syntactically valid string such as abc. You must provide a value here, otherwise the cleanup script will stop at this point and prompt for values. It is not necessary to provide the actual list of subnets to destroy, although you can provide a list of those subnets if you prefer. The list of subnets includes the admin subnet, db subnet , mid subnet, and pdDNSZoneName.

Removing All Instance Resources for a Non-Production Environment

Use this procedure to remove all the instance resources in Oracle Cloud Infrastructure that were provisioned by the deployment of JD Edwards EnterpriseOne infrastructure provisioning for a non-production environment in a single AD.

1. Navigate to this directory on your OpenTofu staging server:
   `JDERefArch_InfraProvisioning/TerraformScripts/nonpd`
2. Run this command as a single contiguous line command (line breaks below added for clarity):

```
tofu destroy
-var-file=../inputJson/common.json
```

```
-var-file=../inputJson/pdsubnet.json

-var-file=../inputJson/nonpd.json

-var db_admin_password='<dbPwd>'

-var dep_password='<depPwd>'

-var dep_subname='depSubName'
```

Note: In the above command, the values `db_admin_password`, `dep_password`, and `dep_subname` do not have to be the literal values. That is, you use any syntactically valid string such as abc. You must provide a value here, otherwise the cleanup script will stop at this point and prompt for values. It is not necessary to provide the actual list of subnets to destroy, although you can provide a list of those subnets if you prefer. The list of subnets includes the admin subnet, db subnet , mid subnet, and pdDNSZoneName.

Removing a WebLogic Subscription

After you have removed all resources for Production and Non-Production environments, you must run these commands to remove WebLogic subscription information.

Production Environment

1. Navigate to this directory on your OpenTofu staging server:

   ```
   cd <root>/JDERefArch_InfraProvisioning/TerraformScripts/wlsImageSubscription/pd Subscription
   ```
2. Run this command as a single contiguous line command:

   ```
   tofu destroy -var-file=/home/opc/JDERefArch_InfraProvisioning/TerraformScripts/inputJson/common.json
   ```

Non-Production Environment

1. Navigate to this directory on your OpenTofu staging server:

   ```
   cd <root>/JDERefArch_InfraProvisioning/TerraformScripts/wlsImageSubscription/nonpd Subscription
   ```
2. Run this command as a single contiguous line command:

   ```
   tofu destroy-var-file=/home/opc/JDERefArch_InfraProvisioning/TerraformScripts/inputJson/common.json
   ```

Removing All Instance Resources for a Disaster Recovery Environment

Use this procedure to remove all instance resources for a disaster recovery environment that was provisioned by the deployment of JD Edwards EnterpriseOne infrastructure provisioning.

1. Navigate to this directory on your OpenTofu staging server:

   ```
   JDERefArch_InfraProvisioning/TerraformScripts/pdDR
   ```
2. Run this command as a single contiguous line command:

   ```
   tofu destroy -var-file=../inputJson/drpd.json -var-file=../inputJson/drcommon.json -var- file=../
   inputJson/drpdsubnet.json
   ```

Removing Virtual Cloud Network (VCN) Resources for a Primary Site

Use this procedure to remove all the Virtual Cloud Network (VCN) resources in Oracle Cloud Infrastructure that were provisioned by the deployment of JD Edwards EnterpriseOne infrastructure provisioning.

1. Navigate to this directory on your OpenTofu staging server:

   ```
   JDERefArch_InfraProvisioning/TerraformScripts/global
   ```
2. Run this command as a single contiguous line command:

   ```
   tofu destroy -var-file=../inputJson/common.json -var-file=../inputJson/vcn.json
   ```

Removing Virtual Cloud Network (VCN) Resources for a Disaster Recovery Site

Use this procedure to remove all the Virtual Cloud Network (VCN) resources in Oracle Cloud Infrastructure that were provisioned by the deployment of JD Edwards EnterpriseOne infrastructure provisioning.

1. Navigate to this directory on your OpenTofu staging server:

   ```
   JDERefArch_InfraProvisioning/TerraformScripts/globalDR
   ```
2. Run this command as a single contiguous line command:

   ```
   tofu destroy -var-file=../inputJson/drpd.json -var-file=../inputJson/drcommon.json -var-file=../
   inputJson/drpdsubnet.json
   ```

Use Case Examples

You may need to clean up infrastructure provisioning in the following cases:

- You want to destroy an existing provisioned infrastructure that you no longer require and want to start a fresh deployment of infrastructure provisioning.
- You have tried multiple times to rerun a deployment of infrastructure provisioning but the deployment has failed every time. In this case you may want to start a fresh deployment of infrastructure provisioning.

To remove environments that were provisioned in a single AD, execute these commands in the given order. In these commands, each paragraph is a single contiguous line.

Non-Production

```
cd JDERefArch_InfraProvisioning/TerraformScripts/nonpd
```

```
tofu destroy -var-file=../inputJson/common.json -var-file=../inputJson/pdsubnet.json -var-file=../inputJson/
nonpd.json -var db_admin_password='<dbPwd>' -var dep_password='<depPwd>' -var dep_subname='depSubName'
```

Production

```
cd JDERefArch_InfraProvisioning/TerraformScripts/pd
```

```
tofu destroy -var-file=../inputJson/common.json -var-file=../inputJson/pdsubnet.json -var-file=../inputJson/
pd.json -var db_admin_password='<dbPwd>' -var dep_password='<depPwd>' -var dep_subname='depSubName'
```

Global (Virtual Cloud Network Resources – Primary site)

```
cd JDERefArch_InfraProvisioning/TerraformScripts/global
```

```
tofu destroy -var-file=../inputJson/common.json -var-file=../inputJson/vcn.json
```

Production Disaster Recovery

```
cd JDERefArch_InfraProvisioning/TerraformScripts/pdDR

tofu destroy -var-file=../inputJson/drcommon.json -var-file=../inputJson/drvcn.json
```

Global Disaster Recovery (Virtual Cloud Network Resources – Disaster Recovery site)

```
cd JDERefArch_InfraProvisioning/TerraformScripts/globalDR

tofu destroy -var-file=../inputJson/drcommon.json -var-file=../inputJson/drvcn.json
```

11 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>
```

12 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