ORACLE

Deploying OpenSearch

PeopleSoft Search Framework

OpenSearch is an open-source search engine that you can install and deploy as part of the PeopleSoft Search framework.
In this spotlight series we'll cover the following topics:

- In the first topic, PeopleSoft Search Framework and OpenSearch, we'll briefly discuss PeopleSoft Search Framework and the OpenSearch search engine.
- In the Understanding Clusters in OpenSearch topic, we'll briefly discuss the important concept of clusters in OpenSearch.
- PeopleSoft Deployment Packages, or DPKs, are used to deploy OpenSearch. In the Understanding OpenSearch DPK topic, we'll describe how to obtain the OpenSearch Kit (OSK) DPK and describe some of the software components it contains.
- In the Deploying OpenSearch topic we'll describe the OpenSearch deployment process, the individual deployment steps, and verifying the cluster setup.
- The Integrating OpenSearch with PeopleSoft topic takes you through the steps of configuring PeopleSoft Search Framework with OpenSearch using PIA or the ACM template.
- The Managing an OpenSearch Cluster topic discusses tasks that help you manage a cluster.
- And the final topic, Additional Resources, lists the resources for more information about the PeopleSoft Search Framework and lists the terminology used in the

OpenSearch product.

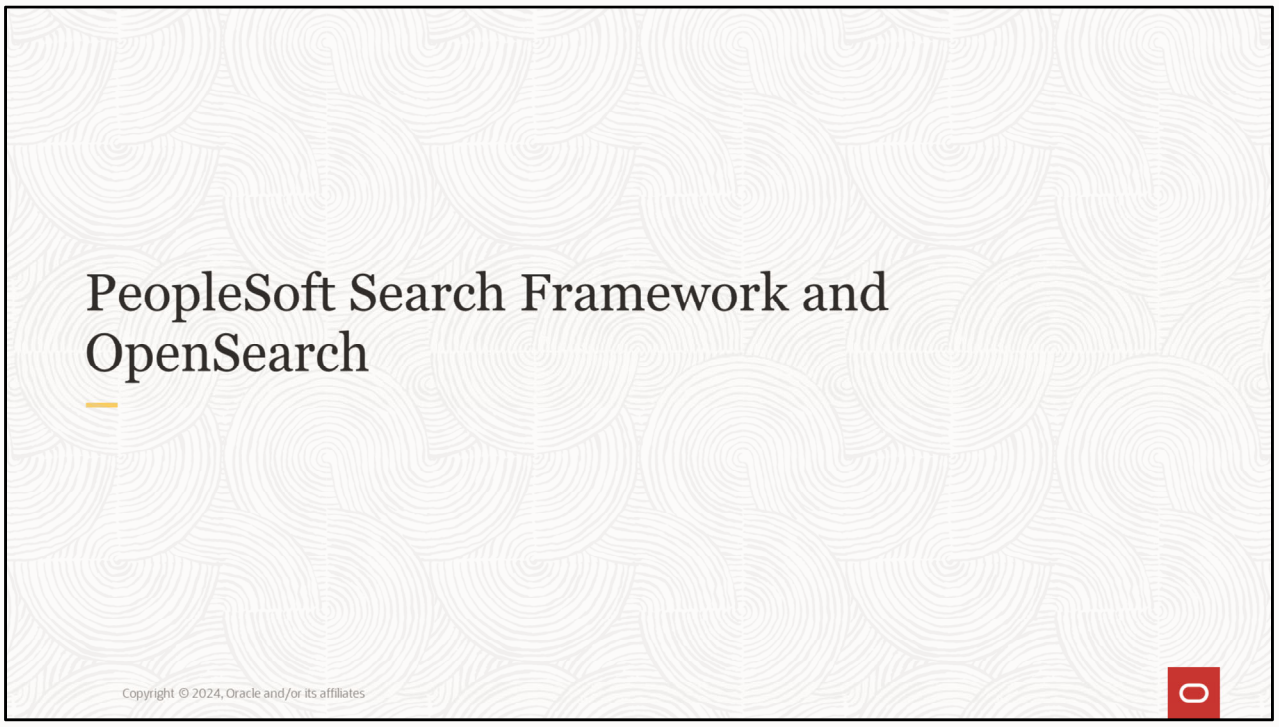

Let’s begin with a brief introduction to the PeopleSoft Search Framework and the new search engine, OpenSearch.

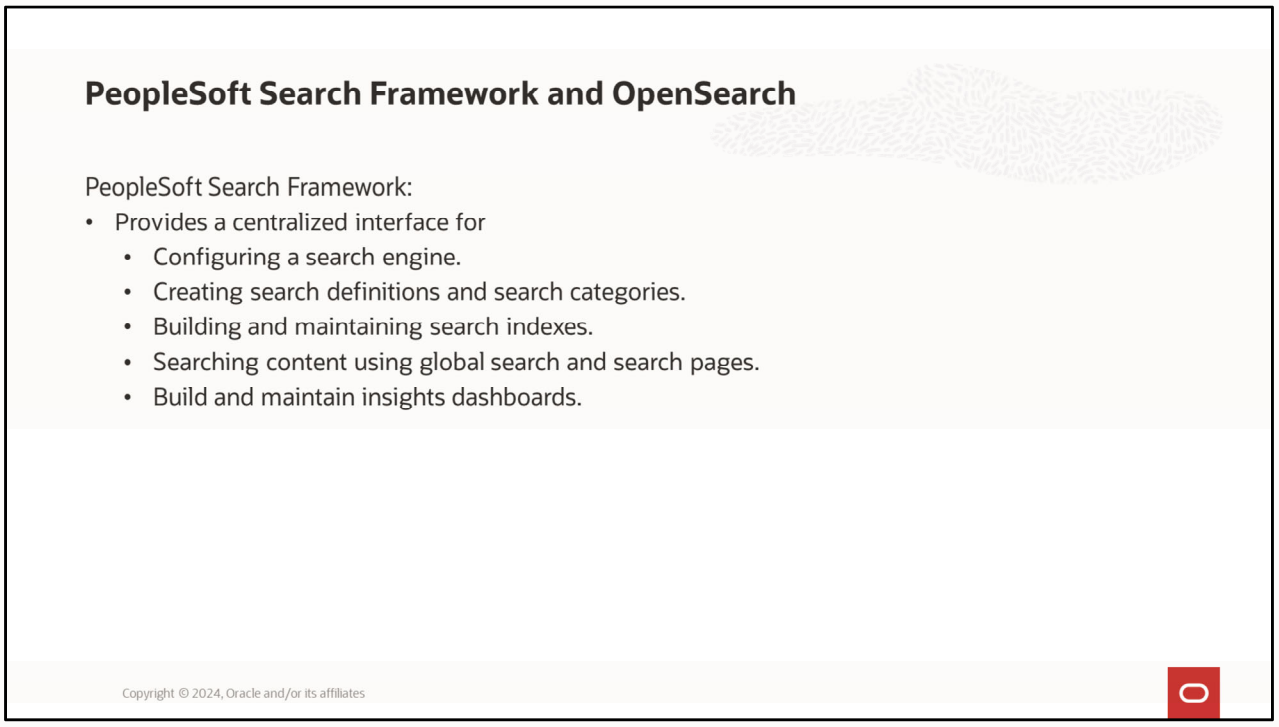

The PeopleSoft Search Framework is a centralized interface available within PeopleTools to:

- Configure a search engine.
- Create search definitions and search categories.
- Build and maintain search indexes.
- Search content using global search and search pages.
- Build and maintain Insights dashboards.

When you deploy the OpenSearch search engine, the processes of creating queries, creating search definitions and categories, deploying search definitions and categories, and viewing search results are not changed. The same search definitions and categories used in Elasticsearch can be used in OpenSearch as well.

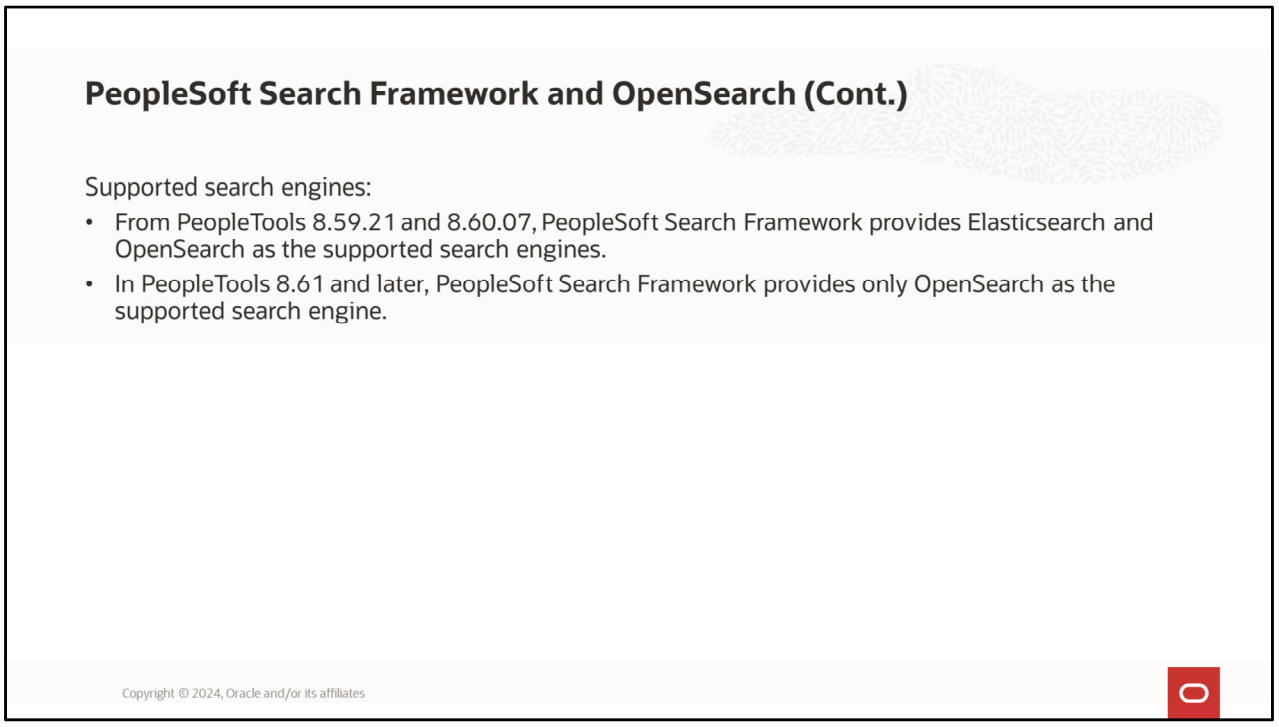

From PeopleTools 8.59.21 and 8.60.07, PeopleSoft Search Framework enables customers to use Elasticsearch or OpenSearch as their supported search engine. In PeopleTools 8.61 and later, PeopleSoft Search Framework will only support OpenSearch.
In this video, we will talk about how you can seamlessly migrate to OpenSearch in PeopleTools 8.59, 8.60, and 8.61.

Next, we'll discuss clusters and their importance in OpenSearch. In this section, we will discuss the basics of clusters.
We'll cover more details related to clusters in the subsequent sections.

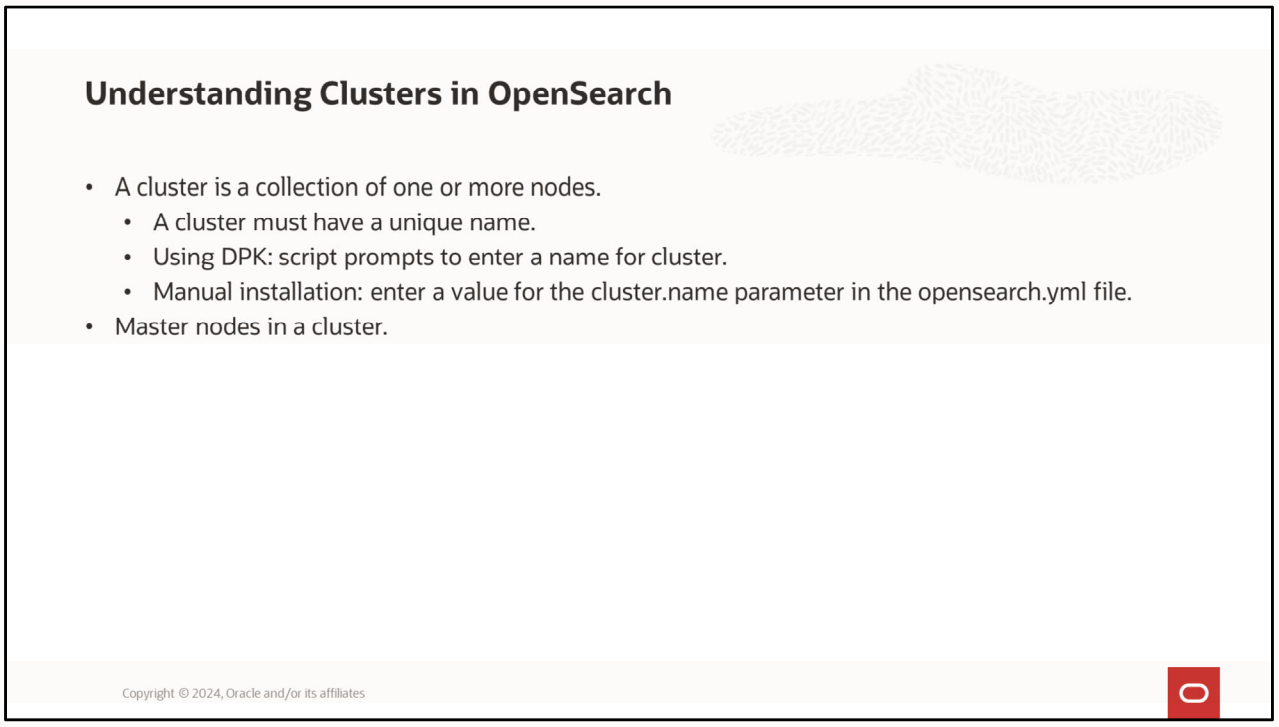

A cluster is a collection of nodes. Multiple nodes running on one or more hosts (physical or virtual) can be grouped into a cluster. Different types of nodes can be present in an OpenSearch cluster. In the PeopleSoft implementation of an OpenSearch cluster, only the master-data node type is supported, that is, the node can act both as a master node and as a data node. Multiple nodes in a cluster provide for redundancy and scaling.

It is important to note here that the following bullet points describe the parameters of clusters that should be defined during the installation of OpenSearch.

Before we discuss clusters, it is important to understand the opensearch.yml file. The opensearch.yml file is a configuration file that contains the most important settings that you need to configure a cluster. Note that the .yml file uses space indentation, so you should be careful when you edit it.

Cluster name must be unique. The default cluster name is OSCLUSTER. However, you can override the default cluster name. The cluster name is important because a node can only be part of a cluster if the node is set up to join the cluster by its name. Make sure that you don't reuse the same cluster name in different environments, otherwise you might end up with nodes joining the wrong cluster. For example, you could use logging-dev, logging-stage, and logging-prod for the development, staging, and production clusters. When you install OpenSearch using DPK, the script prompts for the name of the cluster. If you are installing OpenSearch manually, you

need to specify the cluster name in the cluster.name parameter in the opensearch.yml file.
Master nodes in a cluster. One node in the cluster is elected to be the master node, which is responsible for lightweight cluster-wide actions such as creating or deleting an index, tracking which nodes are part of the cluster, and deciding which shards to allocate to which nodes. Any master-eligible node (all nodes by default are master-eligible nodes) may be elected to become the master node by the master election process.
OpenSearch decides the initial set of master nodes through the cluster.initial_master_nodes parameter. The value for this parameter in the opensearch.yml configuration file indicates the number of master nodes in the cluster.
Oracle PeopleSoft recommends a cluster with three nodes for high availability. Then, depending on the hardware, memory availability, and search performance, you may add nodes to the existing cluster.

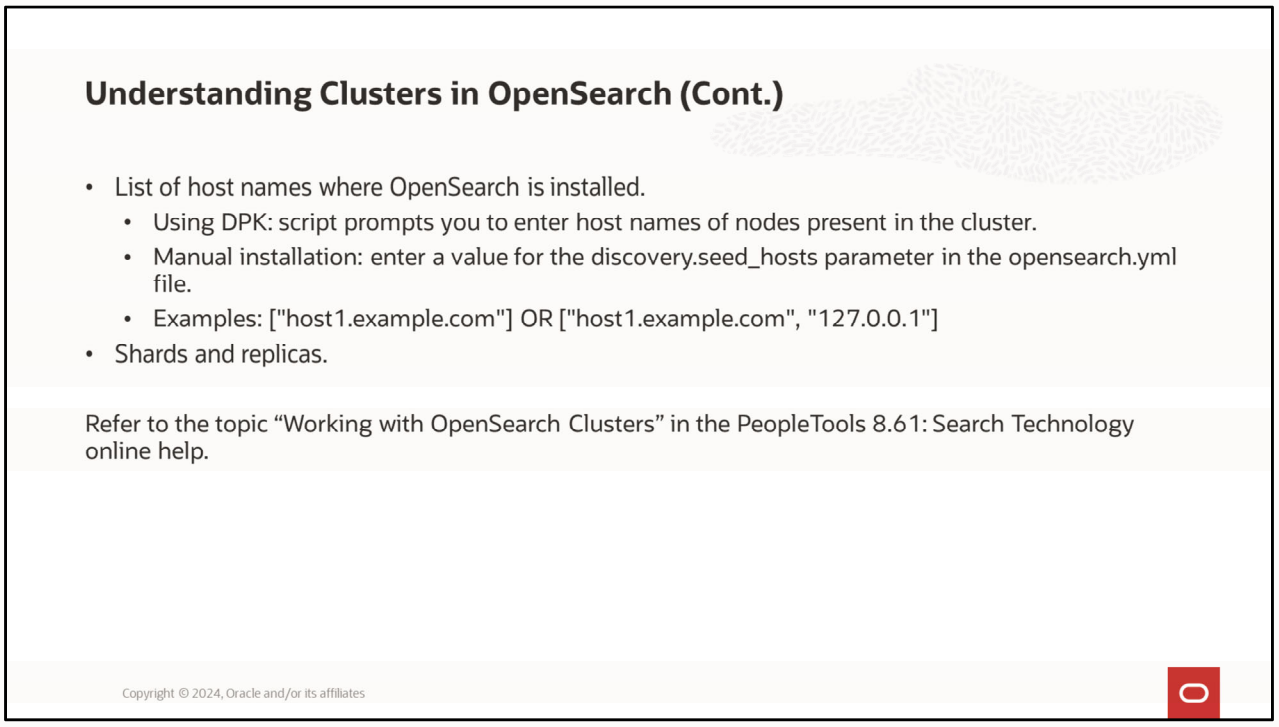

List of host names where OpenSearch is installed. A node name is a symbolic name for identifying the node. Host name is the IP or DNS of the machine where OpenSearch is installed. Host names (or IP/DNS) are required for letting each OpenSearch server be aware where it can ping and find other OpenSearch servers during booting up.
When you install OpenSearch using DPK, the script prompts you to enter host names of nodes present in the cluster. If you are installing OpenSearch manually, you need to specify the host names for the discovery.seed_hosts parameter in the opensearch.yml file. Enclose one or more host names in square brackets, with the host name or IP address in double quotes.
For example:
For one host, you would enter: ["host1.example.com"]
To list two or more hosts, use commas: ["host1.example.com", "127.0.0.1"]
In PeopleSoft Search Framework, each of these nodes needs to be added to the Search Instance Properties page (unless you are relying upon a load balancer), which can be used in a failover scenario by PeopleSoft so that search requests can be directed to an available node.
Shards and replicas are defined on the Search Options page in PIA after you deploy OpenSearch and connect it to the PeopleSoft application through a search instance. OpenSearch provides the ability to subdivide your index into multiple pieces called

shards. When you create an index, you can simply define the number of shards that you want. Each shard is a fully functional and independent 'index' that can be hosted on any node in the cluster. OpenSearch allows you to make one or more copies of your index's shards into what are called replica shards, or replicas for short. We'll discuss more on shards and replicas when we discuss the post deployment actions with respect to a cluster.
For more information, refer to the topic "Working with OpenSearch Clusters" in the PeopleTools 8.61: Search Technology online help.

Now, we'll discuss the contents of the OpenSearch DPK.

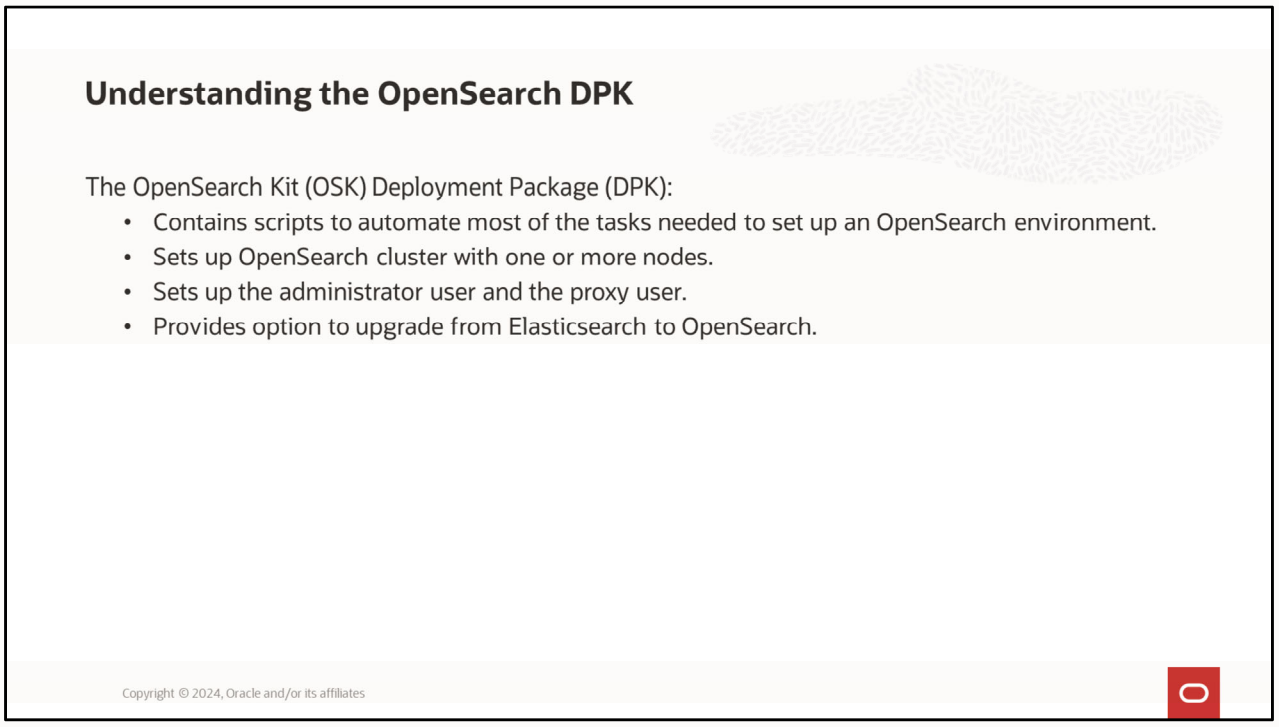

The OpenSearch Kit (OSK) Deployment Package (DPK) contains scripts to automate most of the tasks needed to set up an OpenSearch environment.
The OSK DPK sets up OpenSearch cluster with one or more nodes. You need to run the OSK DPK on each of the nodes to install OpenSearch and ensure that each node references the other nodes of the cluster. We've discussed these aspects in the preceding topic on OpenSearch clusters.
The OSK DPK also sets up the administrator user and the proxy user.
You can use the OSK DPK to switch from Elasticsearch to OpenSearch. Later in this video, we'll discuss the migration process from Elasticsearch to OpenSearch in detail.

Understanding the OpenSearch DPK (Cont.)

The OSK DPK includes:

- Open-source OpenSearch software.
- Open-source Logstash software.
- PeopleSoft-delivered plug-ins for OpenSearch.
  - Used for authentication, authorization, encryption and file web crawling.
- Software and scripts for OpenSearch deployment automation.
- Oracle Java JDK 11.

Refer to PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61) online help.

The OSK DPK includes:
The Oracle OpenSearch and OpenSearch Dashboards software.
The open-source Logstash software.
PeopleSoft-delivered plug-ins for OpenSearch. These include plug-ins for authentication, authorization, encryption, and file web crawling.
Python software and PeopleSoft-developed scripts that automate OpenSearch deployment.
and
Oracle Java JDK 11.

Now, we'll discuss the important process of deploying OpenSearch in a PeopleSoft implementation.

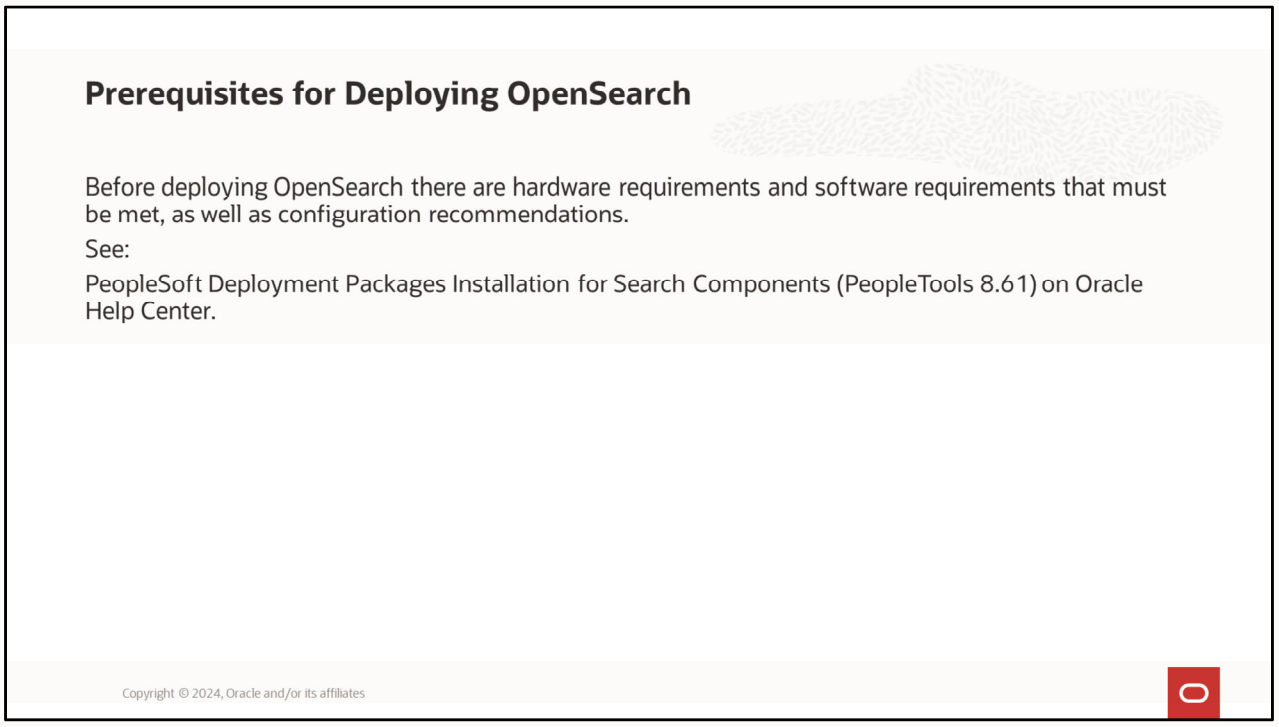

Before you begin deploying OpenSearch, review the hardware requirements, software requirements, and other configuration and performance recommendations for OpenSearch.
You can find this information in PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61) on Oracle Help Center.

Use Cases for Deploying OpenSearch

Consider the following use cases:

- Use case 1: Migrating from Elasticsearch to OpenSearch (in PT 8.59 and 8.60).
- Use case 1: Migrating from Elasticsearch to OpenSearch (in PT 8.61).
- Use case 2: New installation of OpenSearch.

Let's discuss three use cases.
The first use case discusses how existing users of Elasticsearch 7.0 or 7.10 can seamlessly migrate to OpenSearch without reindexing in PeopleTools 8.59 and 8.60.
The second use case discusses how an existing Elasticsearch user can migrate to OpenSearch without reindexing in PeopleTools 8.61.
And the third use case discusses a fresh installation of OpenSearch.

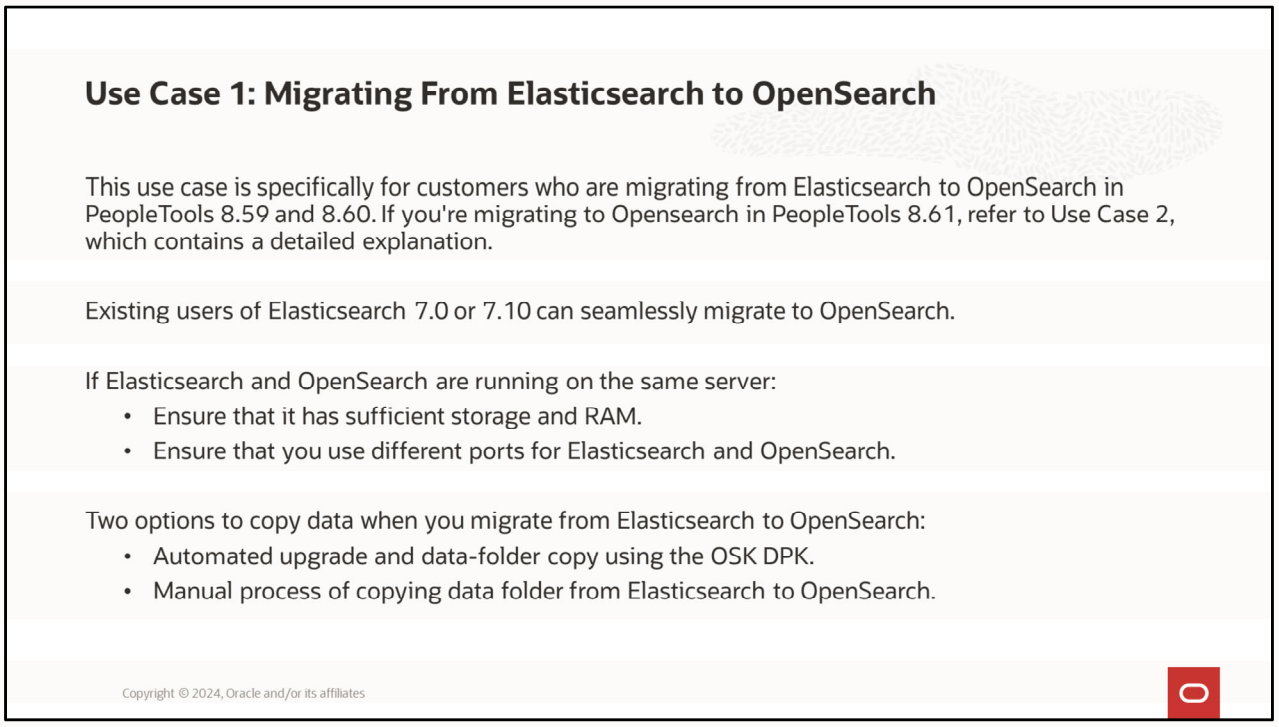

This use case is specifically for customers who are migrating from Elasticsearch to OpenSearch in PeopleTools 8.59 and 8.60. If you're migrating to Opensearch in PeopleTools 8.61, refer to Use Case 2, which contains a detailed explanation.
In this use case, the assumption is that you are already running Elasticsearch 7.0 or 7.10 in your PeopleSoft 9.2 application. Now, you want to migrate to OpenSearch.
You can install OpenSearch on a server that is different from the Elasticsearch server. However, if you’re planning to run Elasticsearch and OpenSearch on the same server, you should ensure that it has sufficient storage and RAM to sustain the additional volume of data. Also, you should ensure that you use different ports for Elasticsearch and OpenSearch.
You have two options to copy data when you switch to OpenSearch, that is, to use the automated process to install OpenSearch and copy the Elasticsearch data folder or manually copy the Elasticsearch data folder.

Use Case 1: Migrating From Elasticsearch to OpenSearch (Cont.)

Two options to migrate to OpenSearch:

1. Automated upgrade and data-folder copy using the DPK.
   - Run the upgrade script.
     ./psft-dpk-setup.sh --os_upgrade --install_base_dir BASE_DIR
   - Answer Yes to the upgrade option.
   - Answer Yes to the copy data option.
     If you answer No, you can manually copy the data folder.
   - Continue with the step to create an OpenSearch instance in the PeopleSoft application.

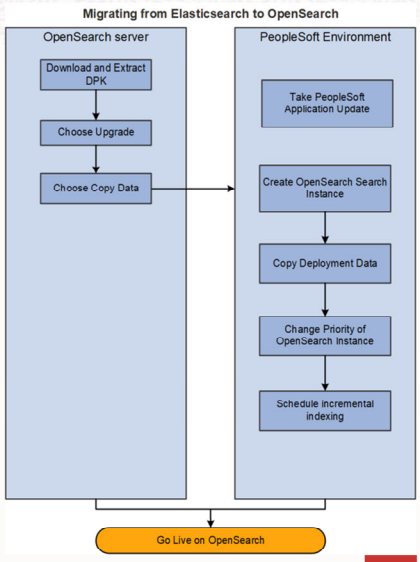

In this video, we're discussing the upgrade process on a Windows environment. For upgrade process on a Linux environment, refer to the PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61).
On Windows, due to directory permissions, you must stop the Elasticsearch instance or service to copy the data before the upgrade process.
The OSK DPK enables you to automate the process of installing OpenSearch and copying the indexed data in Elasticsearch to the OpenSearch data folder.
To run the upgrade process, you use the os_upgrade command as shown. Answer Yes to the prompt for upgrading. Answer Yes to the prompt for copying data from Elasticsearch to OpenSearch. After the data copy is completed, you may proceed with the steps listed on the PeopleSoft application.

Use Case 1: Migrating From Elasticsearch to OpenSearch (Cont.)

Steps to complete on the PeopleSoft application:

- Create an OpenSearch search instance.
- Copy deployment data.
- Change priority of the OpenSearch search instance.
- Schedule incremental indexing.

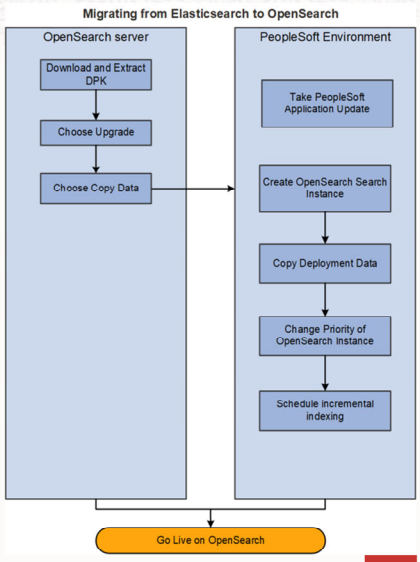

On the PeopleSoft application, you should complete the steps listed here. PeopleSoft applications provide updates to support terminology changes to OpenSearch and PeopleSoft Insights in the user interface. Please review the PeopleSoft Search and Insights Home Page on My Oracle Support, Doc ID 2205540.2, to identify and take any fixes that are relevant to your PeopleSoft application. This part is common to all the use cases discussed in this video.
Since you currently use Elasticsearch, you will have an Elasticsearch search instance, which is the primary search instance. At this point, it is important to remember to start the Elasticsearch service, which you had stopped. Create an OpenSearch instance and save it. Later in the video, we'll discuss about creating an instance.
Next, on the Search Instance page, click the Copy Deployment Data button. This action copies all the deployment related details from the primary search instance (Elasticsearch) to the new search instance (OpenSearch). These deployment details ensure that the incremental indexing from the last indexed time resumes in the new search instance without altering anything in the active Elasticsearch instance settings.
Now, you can set the OpenSearch instance as the primary search instance by changing its priority on the Search Instance Administration page.
Schedule incremental indexing to ensure that the new index is up to date.
You're now live on OpenSearch.

You can use your existing dashboards but you have to deploy OpenSearch Dashboards which is covered in another spotlight: Working With PeopleSoft Insights.

Use Case 1: Migrating From Elasticsearch to OpenSearch (Cont.)

Two options to migrate to OpenSearch:

2. Manually copy of data folder from Elasticsearch to OpenSearch.
   - Run the upgrade script or run the install script for a new installation.
     Answer No to the copy data option.
   - Ensure to stop the OpenSearch service.
   - Replace the data folder in OpenSearch with the data folder from Elasticsearch.
   - Start the OpenSearch service.
   - Continue with the step to create an OpenSearch instance in the PeopleSoft application.

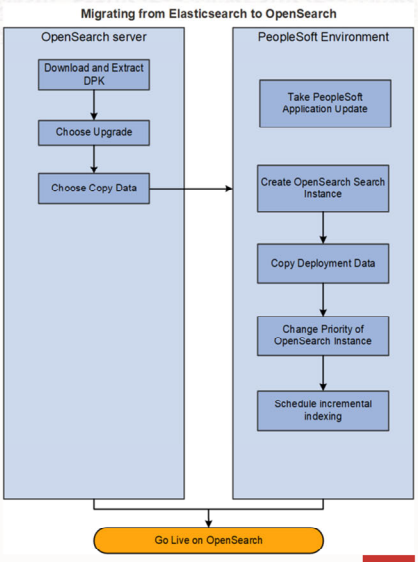

The second option is to manually copy the data folder from Elasticsearch to OpenSearch.
Run the upgrade script, which we discussed previously. Answer No to the prompt to copy data. Complete the upgrade process.
You should stop the OpenSearch service before you undertake the task of deleting the data folder from the OpenSearch home directory. Then, copy the data folder from the Elasticsearch home directory to the OpenSearch home directory.
Now, you can start the OpenSearch service.
You can now continue with the steps on the PeopleSoft application as discussed earlier in the video.

Use Case 2: Migrating From Elasticsearch to OpenSearch (in PT 8.61)

When you upgrade to OpenSearch in PeopleTools 8.61, consider and plan for the following tasks:

1. Automated upgrade using the DPK.
   - Run the upgrade script.
     ./psft-dpk-setup.sh --os_upgrade --install_base_dir BASE_DIR
   - Answer Yes to the upgrade option.
   - Continue with the step to use the existing Elasticsearch instance in the PeopleSoft application.
2. Steps to complete on the PeopleSoft application:
   - Use the existing Elasticsearch search instance.
   - Schedule incremental indexing.

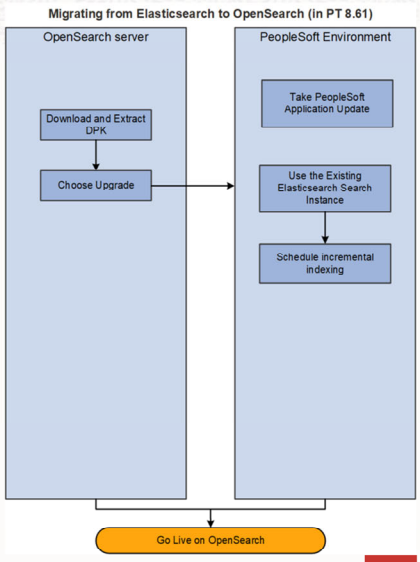

In this use case, we're discussing the upgrade process to OpenSearch in PeopleTools 8.61 on a Windows environment. You should note that OpenSearch is the only supported search engine in 8.61.

We assume that you are already running Elasticsearch 7.0 or 7.10 in your PeopleSoft 9.2 application and upgraded to PeopleTools 8.61. Now, you want to upgrade to OpenSearch.

You can install OpenSearch on the Elasticsearch server or on a server that is different from the Elasticsearch server. If you're planning to install OpenSearch on the Elasticsearch server, you should ensure that it has sufficient storage and RAM to sustain the additional volume of data. However, you must note that in PeopleTools 8.61, you can run only OpenSearch; you cannot run Elasticsearch.

On Windows, due to directory permissions, you must stop the Elasticsearch instance before the upgrade process.

The OSK DPK enables you to automate the process of installing OpenSearch.

To run the upgrade process, you use the os_upgrade command as shown. Answer Yes to the prompt for upgrading.

After the upgrade, you may proceed with the steps listed on the PeopleSoft application.

In order to get continuity in incremental indexing, it is essential to use the existing search instance, which you used for Elasticsearch, but you must ensure that the

configuration of the search instance points to the OpenSearch search engine.
Ensure that the following fields have the correct and valid values:
Search Provider - It should be OpenSearch.
Host name - Ensure that the URL points to the server where the OpenSearch search engine is running.
Port - Ensure that you use the port number where the OpenSearch search engine listens for requests.
Username and Password - Ensure that you use the same values that you entered when running the OSK DPK.
Proxy user and password - Ensure that you use the same values that you entered when running the OSK DPK.
Also ensure that other configuration fields on this page are entered correctly.
Once the configuration on the search instance is fine, schedule incremental indexing to ensure that the new index is up to date.
You're now live on OpenSearch.
You can use your existing dashboards but you have to deploy OpenSearch Dashboards which is covered in another spotlight: Working With PeopleSoft Insights.

Use Case 3: New Installation of OpenSearch

When you plan a new installation of OpenSearch, consider and plan for the following tasks:

- Planning for the installation.
- Deploying the DPK.
- Setting up PeopleSoft Search Framework.
- Taking PeopleSoft Application update.
- Adding an OpenSearch search instance.
- Completing other tasks:
  - Deploying search definitions and categories.
  - Building indexes.
  - Performing acceptance testing.

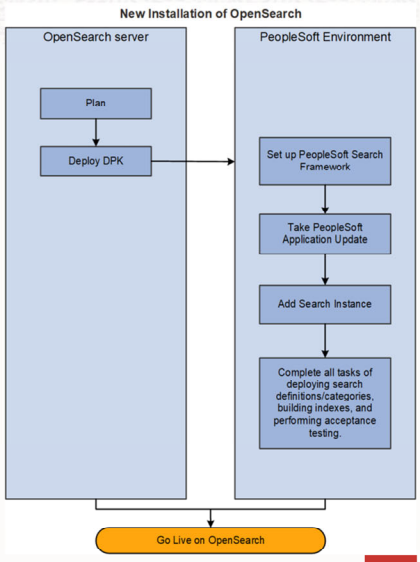

This use case presents a scenario where you are performing a fresh install of OpenSearch. This use case assumes the following:

• You are using PeopleSoft 9.2 application.

The PeopleSoft application must include the appropriate fixes for OpenSearch, as delivered with the latest PeopleSoft Update Image. You need to refer to PeopleSoft Search and Insights Home Page on My Oracle Support, Doc ID 2205540.2.

• The PeopleSoft application is on PeopleSoft PeopleTools 8.59.21 and higher, 8.60.07 and higher, and 8.61 patch level.

• You want to use OpenSearch as the search engine for the PeopleSoft Search Framework.

On the OpenSearch server, you need to:

Plan

That is, you estimate resource requirements and review settings needed for the host on which you install OpenSearch. Refer to the PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61).

Deploy DPK

Deploy the OpenSearch DPK and run the configuration script on a Microsoft Windows or Linux host.

On the PeopleSoft environment, you need to:

Set up PeopleSoft Search Framework

Refer to the “Administering PeopleSoft Search Framework” topic in the PeopleTools 8.61: Search Technology online help for instructions on setting up the search framework.

Take PeopleSoft Application Updates.

Make sure you have at least taken the bugs related to OpenSearch for your PeopleSoft Application as shown on the PeopleSoft Search and Insights Home Page.

Add Search Instance

Two options to create a search instance:

You use the Search Instance Properties page to create an instance.

Alternatively, you can use Automated Configuration Manager (ACM) to run a template to create a search instance.

Later in this video, we’ll view the description of the Search Instance Properties page.

Complete other tasks

To complete the setup of search functionality, you need to perform other tasks, such as, deploying search definitions and categories, building indexes, and testing the search functionality. On successfully completing all these tasks, you are set to go live on OpenSearch.

If you want to do more reading on any of the steps, refer to the 8.61 Search Technology online help and to the PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61).

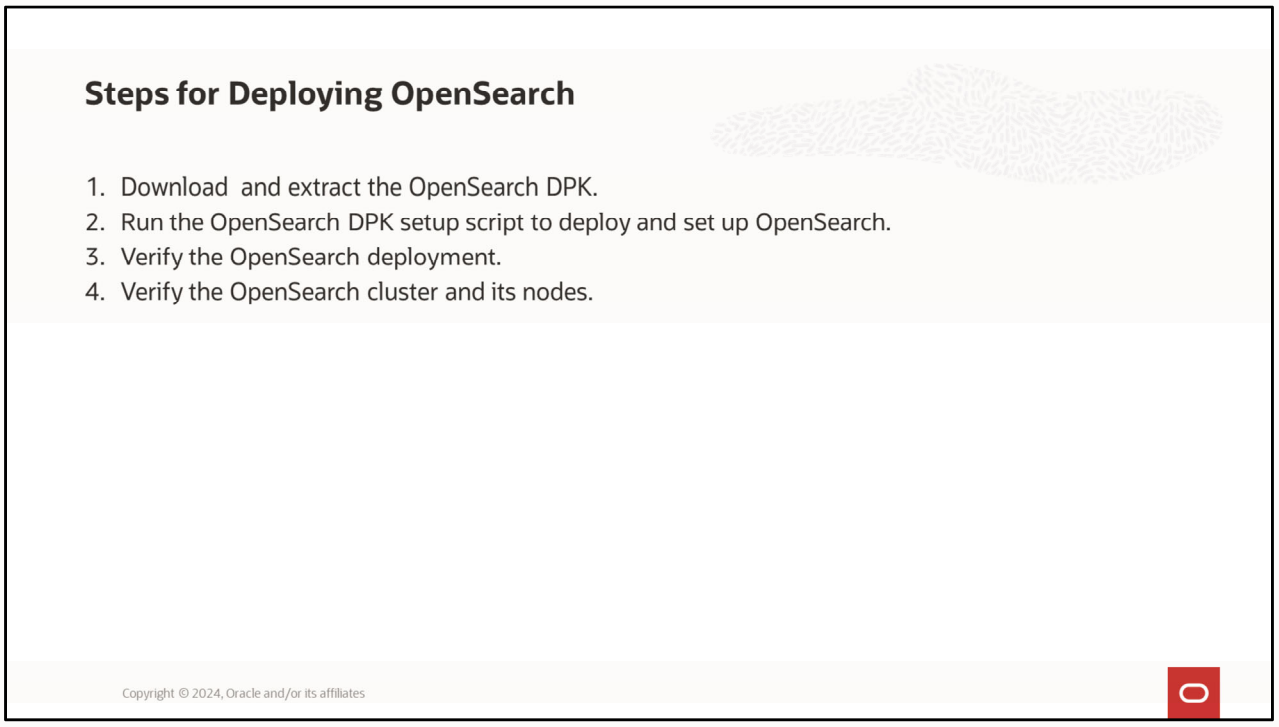

The high-level steps for deploying OpenSearch are as follows:
Step 1 - download and extract the OpenSearch DPK.
Step 2 - run the DPK setup script to deploy and set up OpenSearch.
Step 3 – verify the OpenSearch deployment.
Step 4 – verify the OpenSearch cluster and its nodes.
Note that in this video we describe deploying OpenSearch on a Windows environment.

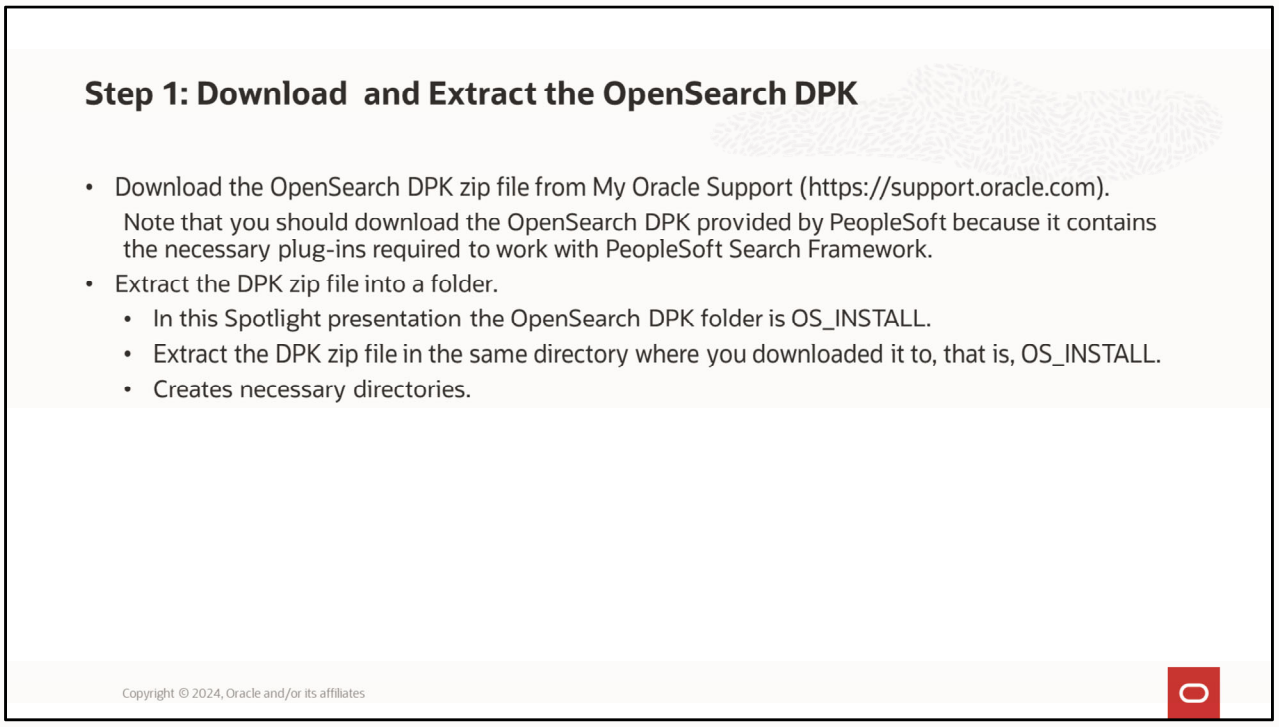

The first step in deploying OpenSearch is to download the OpenSearch DPK.
You can download the OpenSearch DPK from My Oracle Support website.
Note that you should download the OpenSearch DPK that is provided by PeopleSoft because it contains the necessary plug-ins required to work with the PeopleSoft Search Framework.
You must then extract the OpenSearch DPK zip file. Download into an empty directory and extract the entire contents of the zip file in the same download directory where you downloaded it to.
Note that in this video, we'll refer to the OS_INSTALL folder as the location where the OpenSearch DPK is downloaded and extracted.
The extraction process creates the following directories:
Archives directory that includes archives for deployment.
Setup directory that includes the setup scripts and a silent installation sample.
Opensearch-manifest file that includes the version information about OpenSearch and JDK.
And a Readme.txt file.

The next step is to run the OpenSearch DPK setup script.
We've discussed the use cases in this video. Here's a step-by-step demonstration of how to migrate to OpenSearch and copy the data folder, which is the first use case.

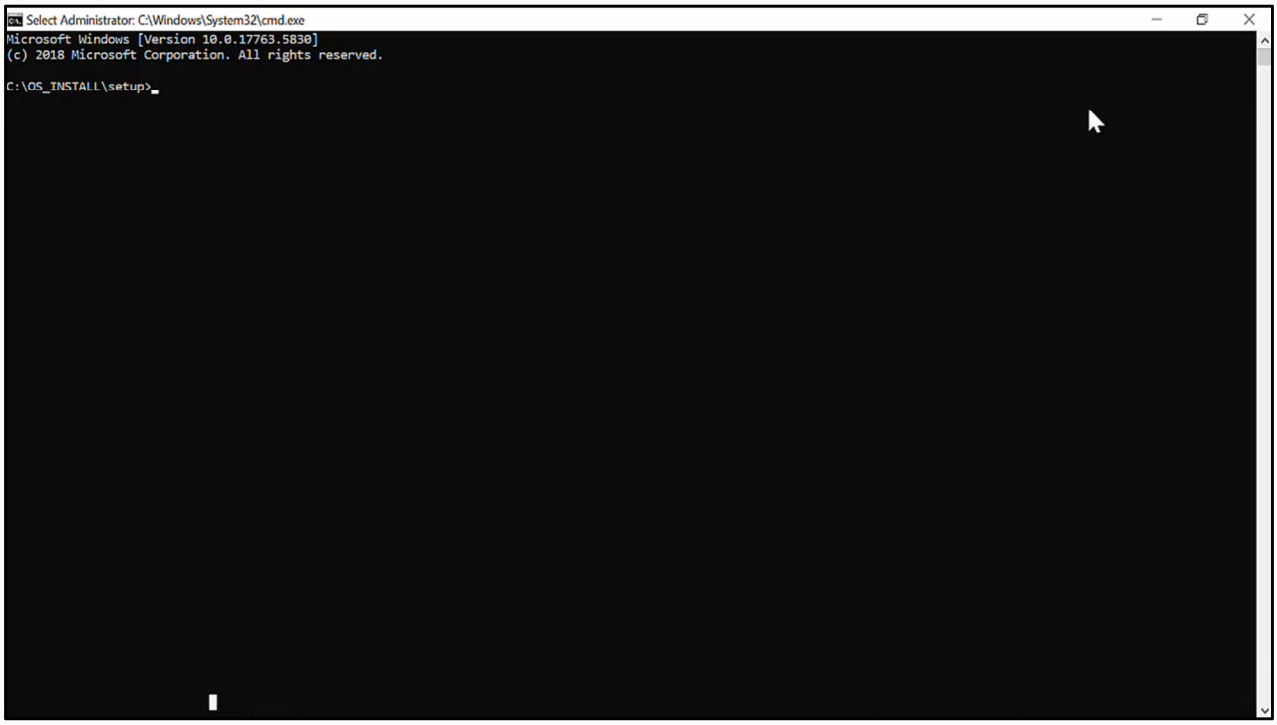

Run the OpenSearch DPK setup script and specify deployment options and settings.

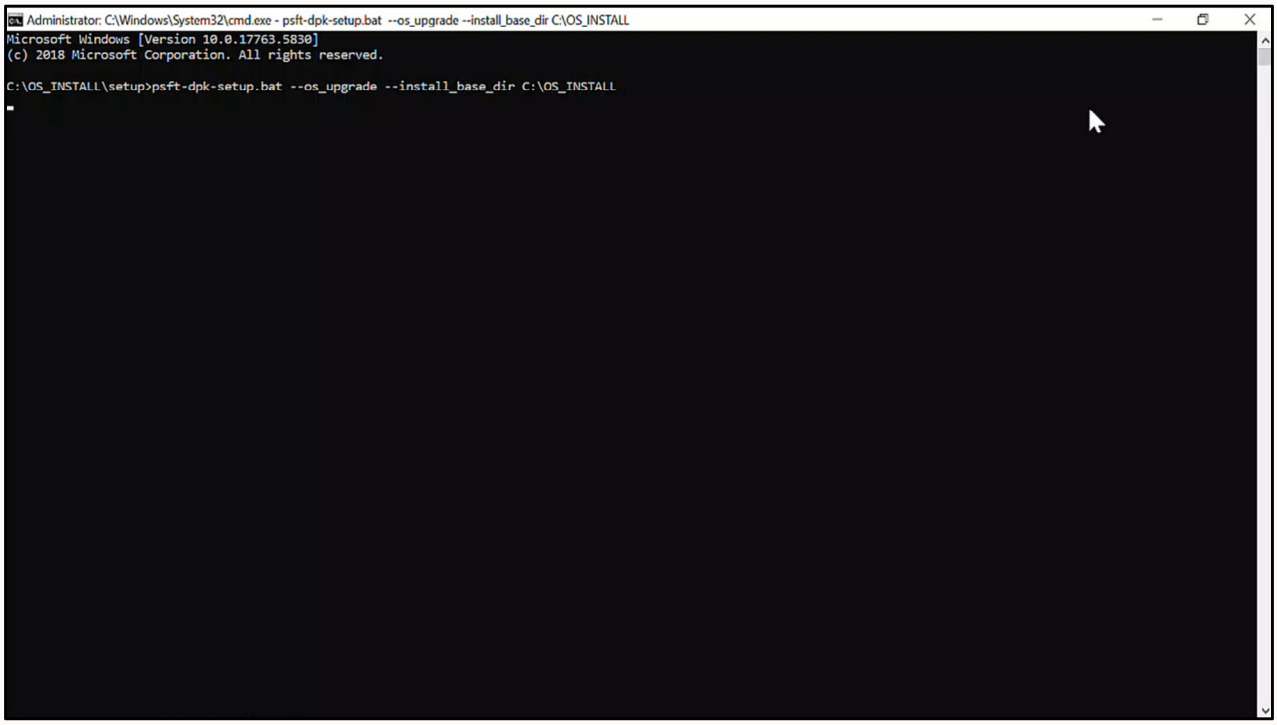

The deployment options and settings include the following:
Define the PeopleSoft base directory. This is the directory where the system installs OpenSearch components.

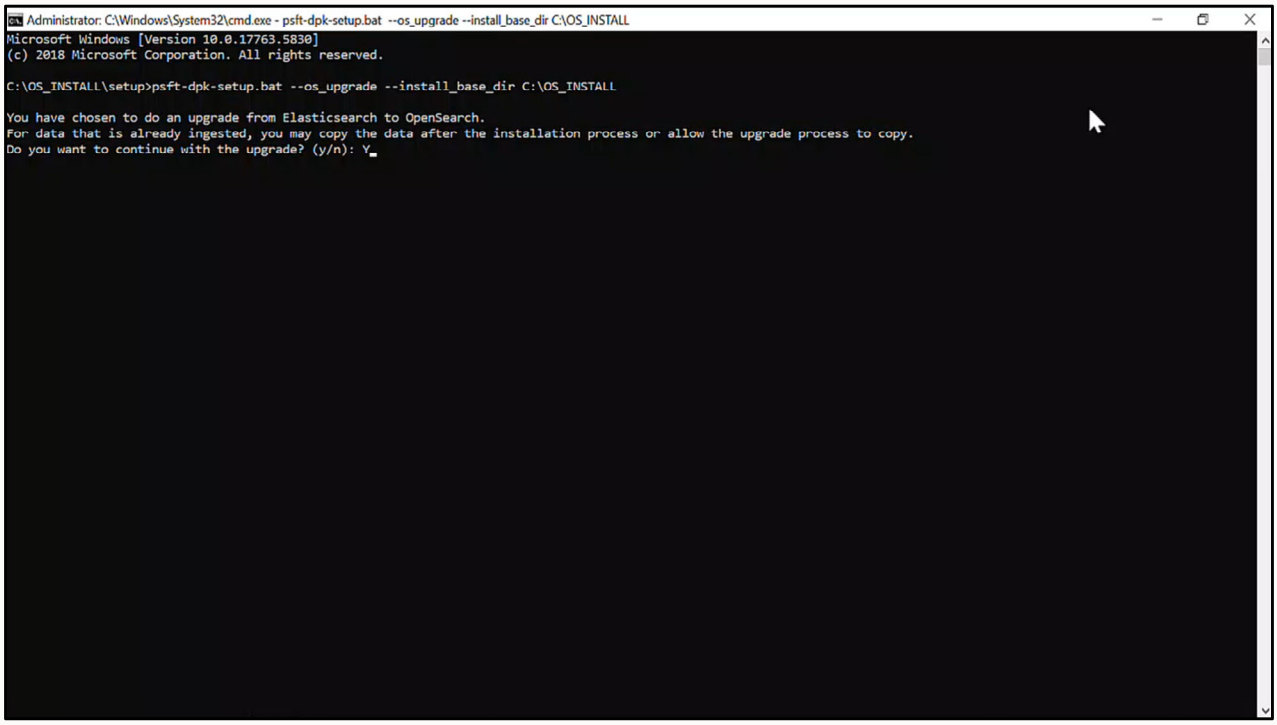

Answer Yes to the prompt for upgrading.

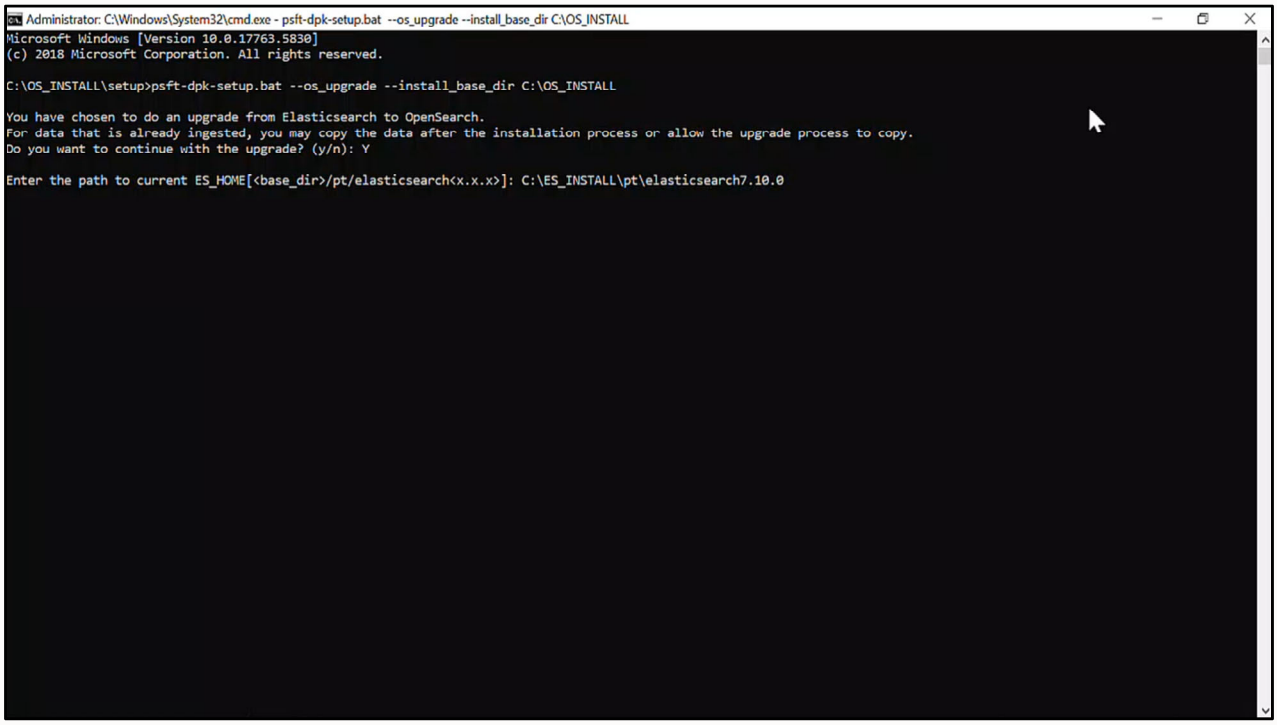

Specify the path to the current ES_HOME.

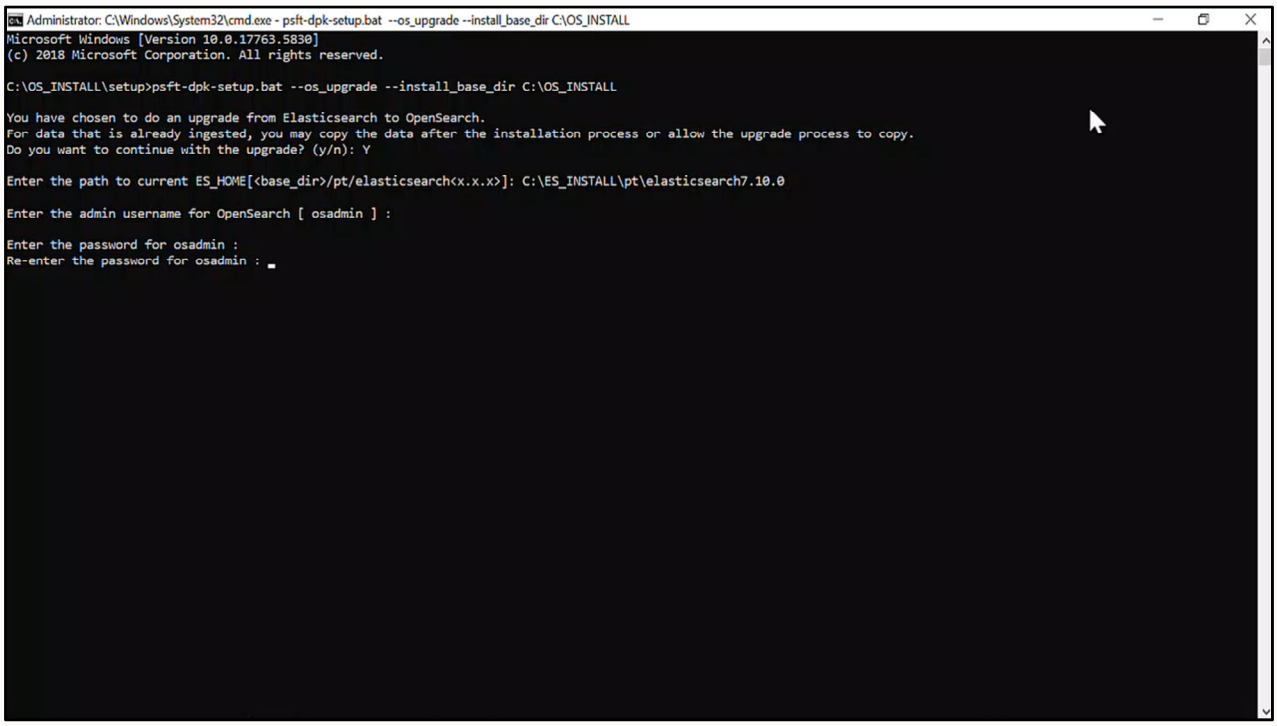

Define the Administrator user and password. You can use the default name or enter a name of your choice. The administrator user has permissions to authenticate requests on OpenSearch.

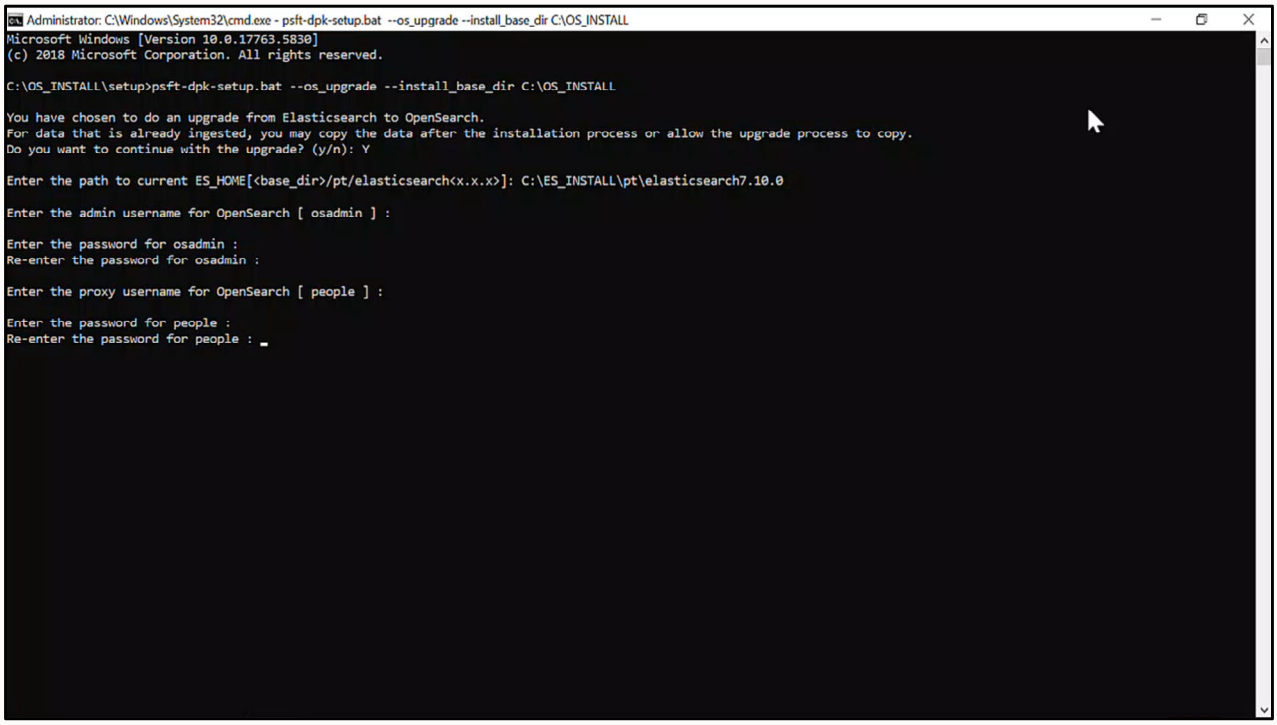

Define the Proxy user and password. You can use the default name or enter a name of your choice. The proxy user enables the PeopleSoft system to log into the search engine using a proxy identity to run a query.

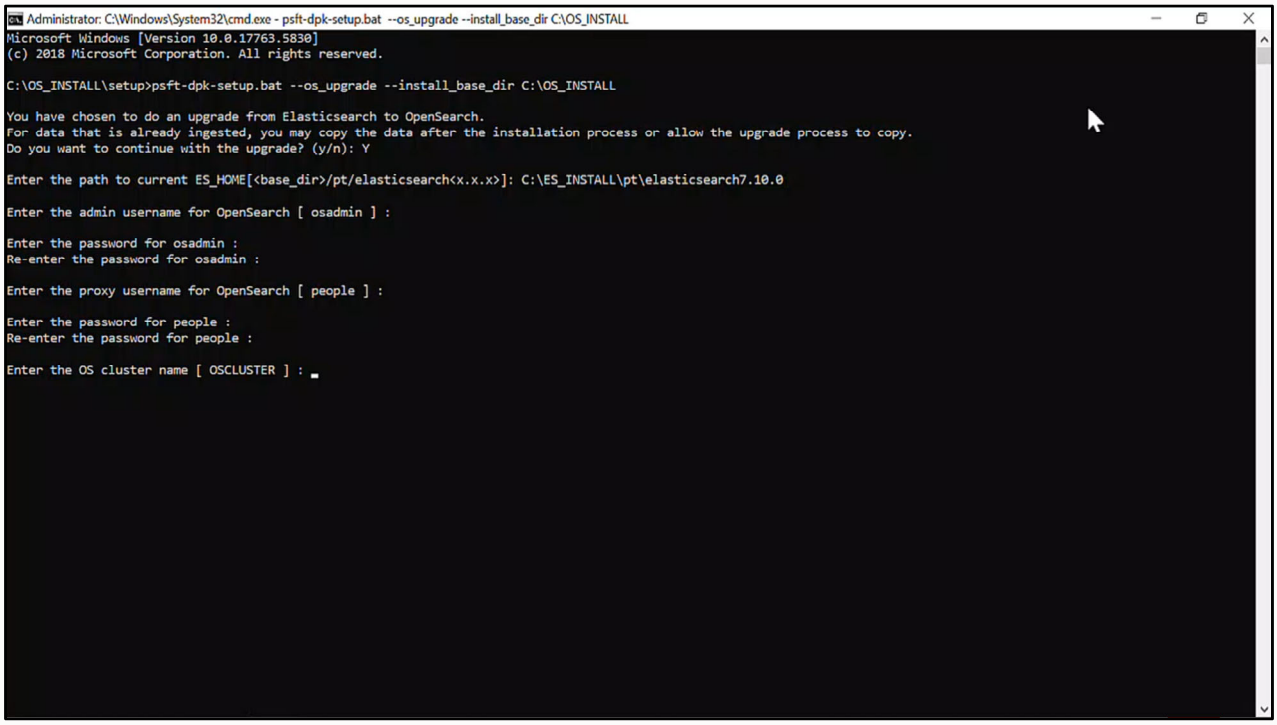

Next,
Specify a name for the OpenSearch cluster. Note that you need to consider the description for cluster name in the earlier part of this video.

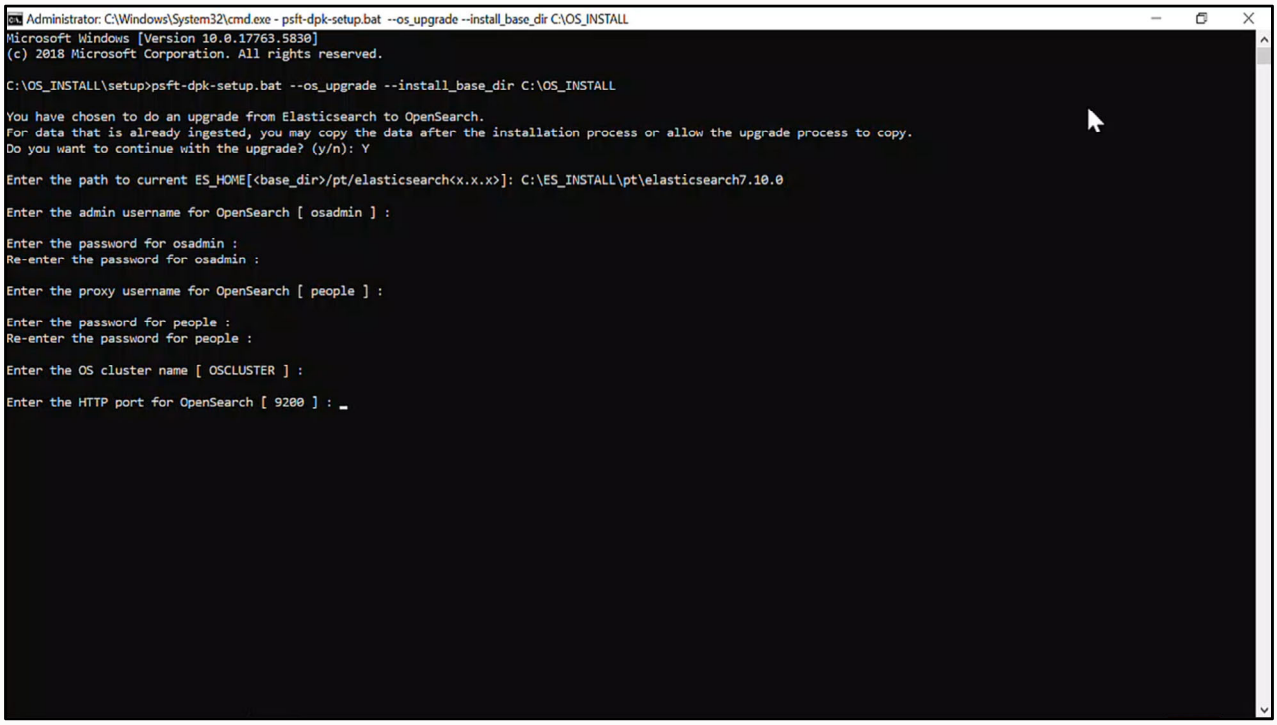

Specify an OpenSearch HTTP port.

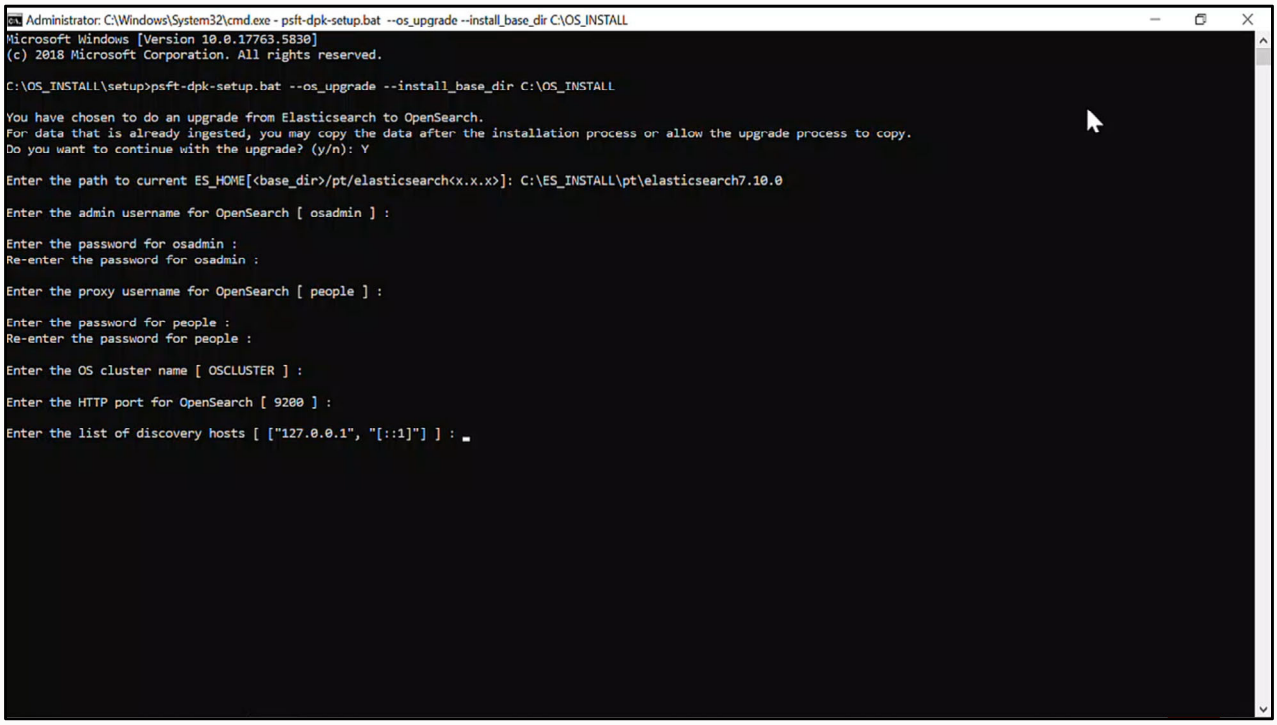

Enter the host name or names for any nodes that are already part of a cluster. Note that you need to consider the description for host name in the earlier part of this video.

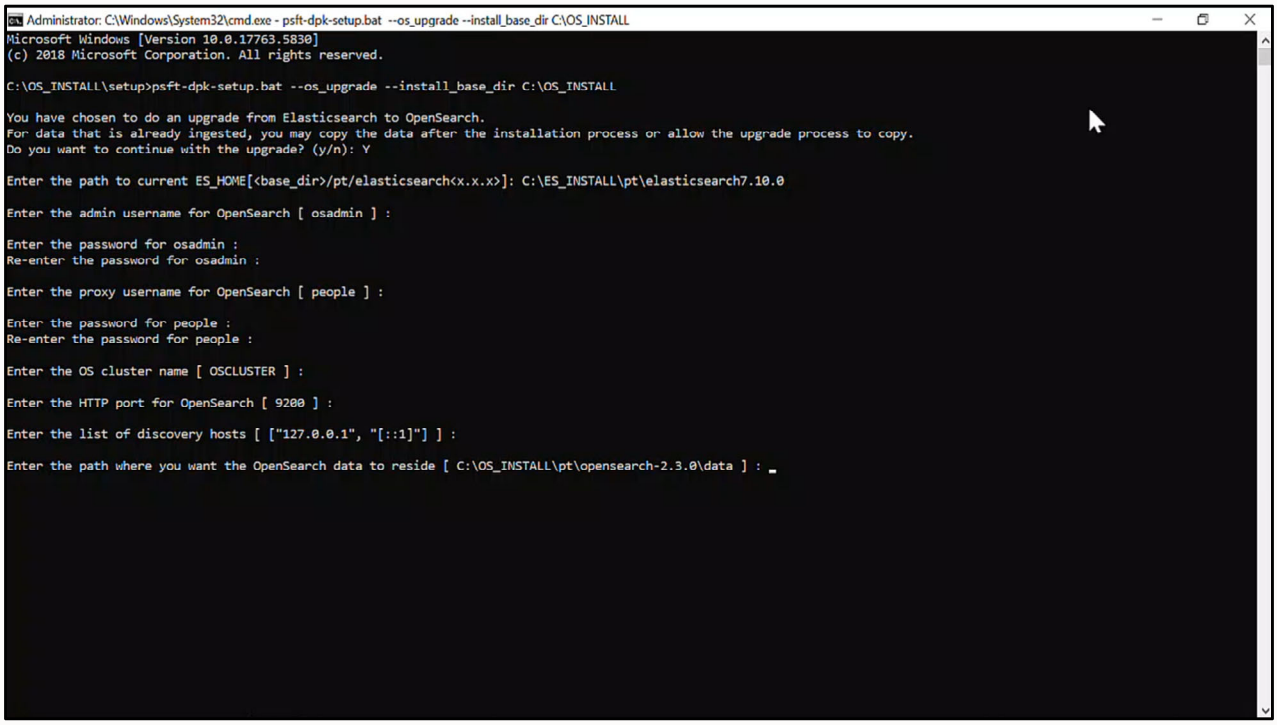

Specify the path where OpenSearch data should reside. A default path is listed, but it is recommended to change the path to a location outside OSK_HOME (BASE_DIR).

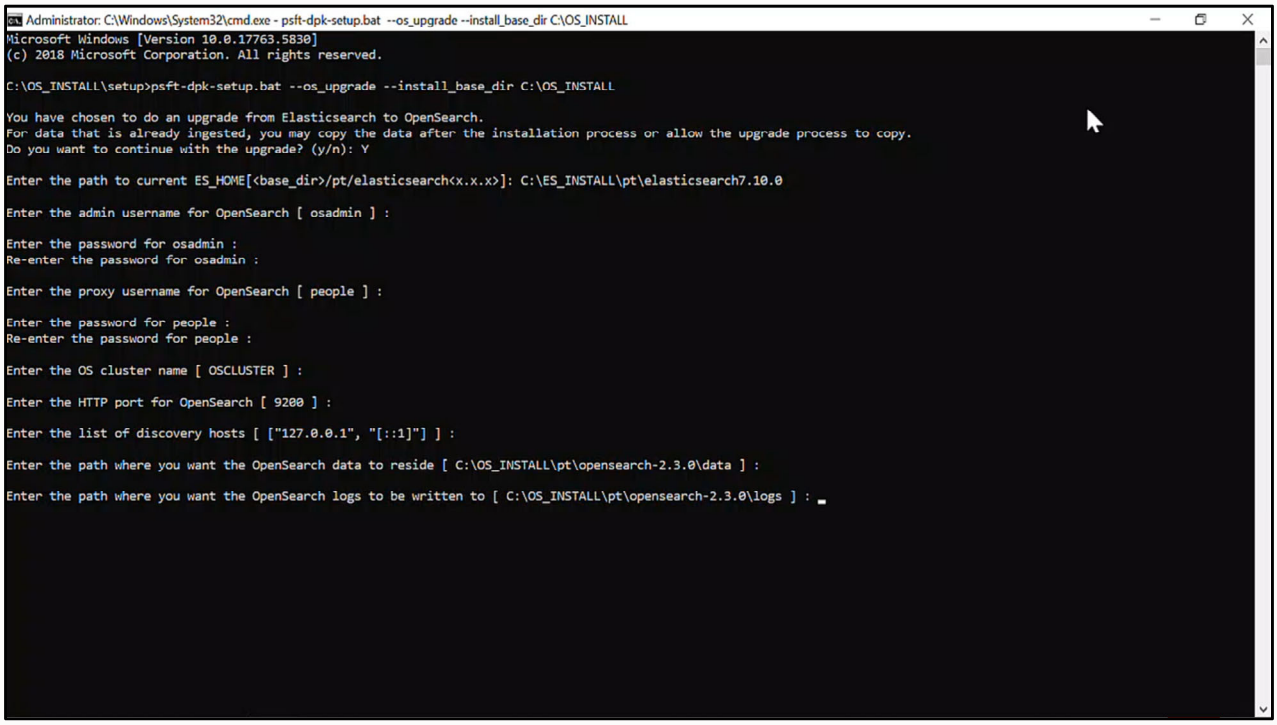

Specify the path where OpenSearch logs should reside. A default path is listed, but you may specify a path if required.

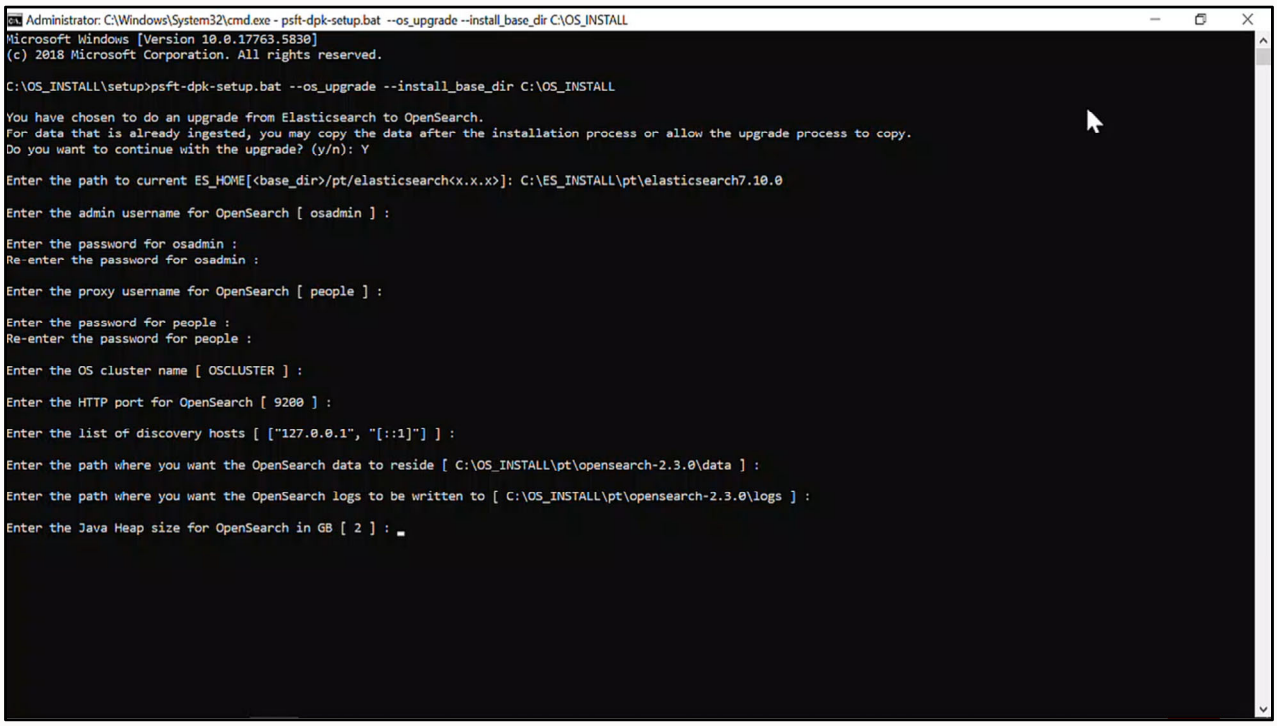

Define the Java heap size. The default value is 2 gigabytes. You could define the heap size as more than 2. To understand the allocation of heap size, refer to the PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61).

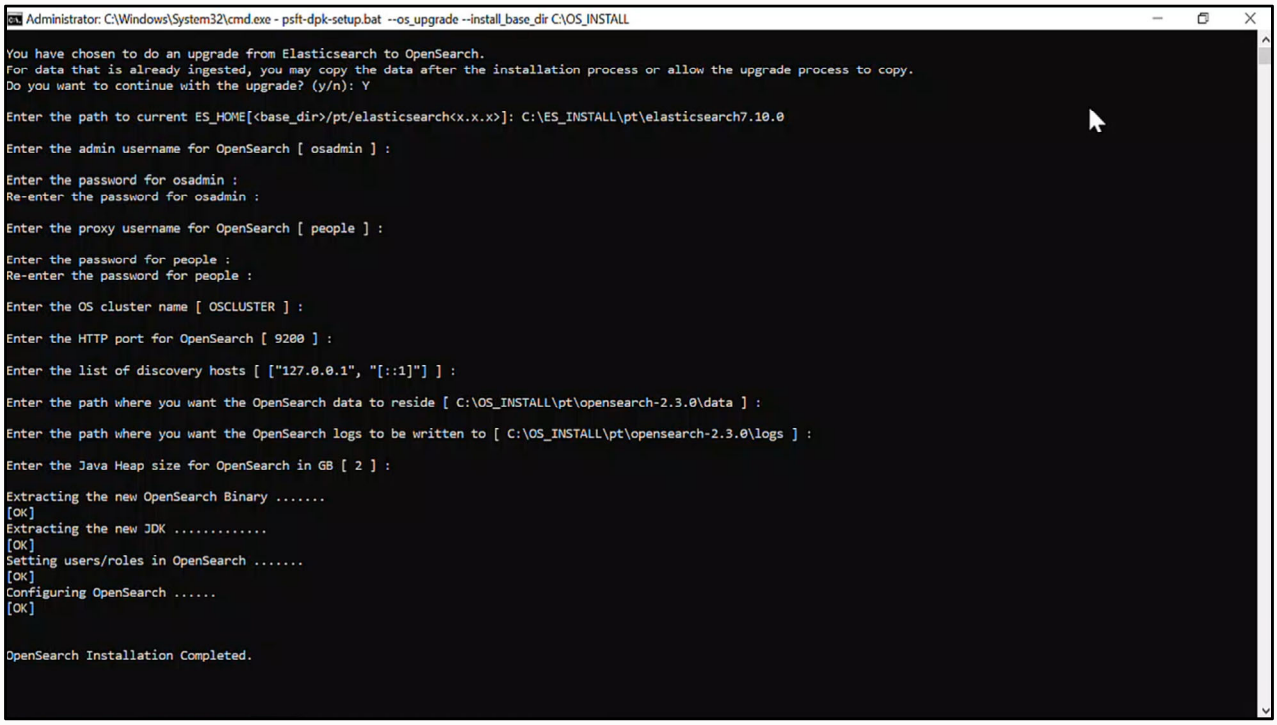

Review the status messages as the script runs and sets up the environment for OpenSearch.

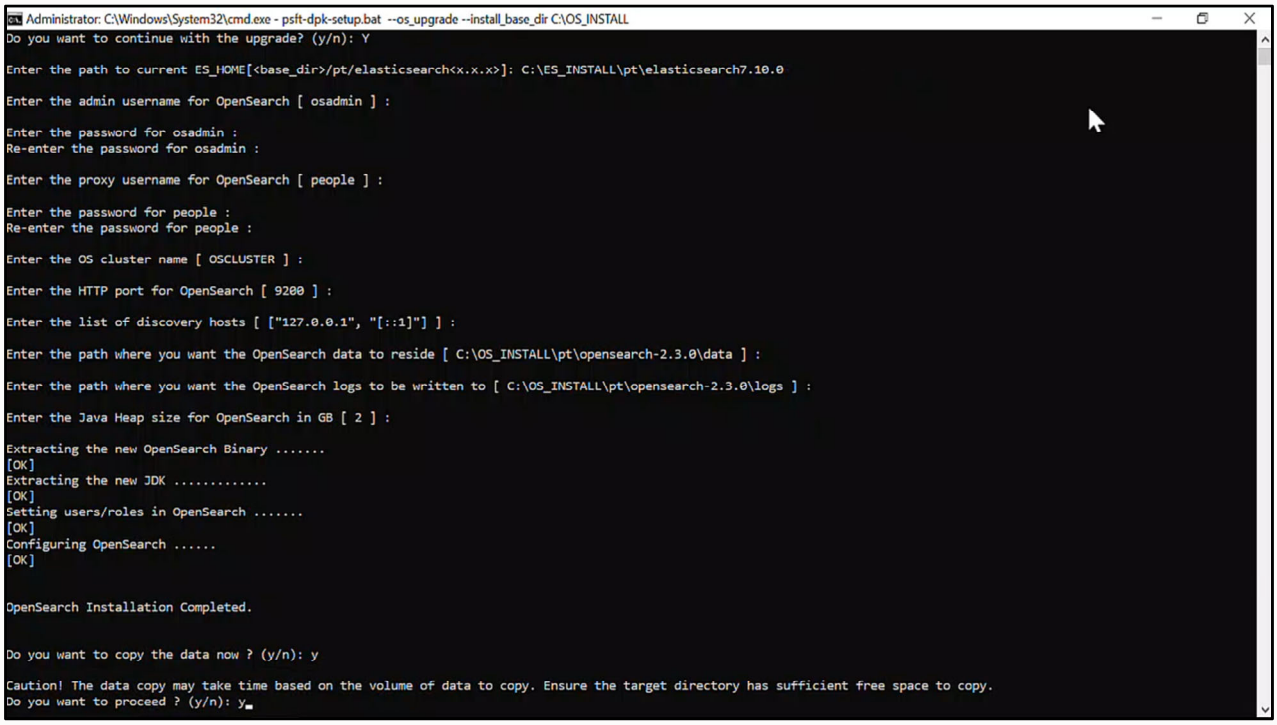

Answer Yes to the prompt for copying data from Elasticsearch to OpenSearch. The data copy may take time depending on the volume of data.

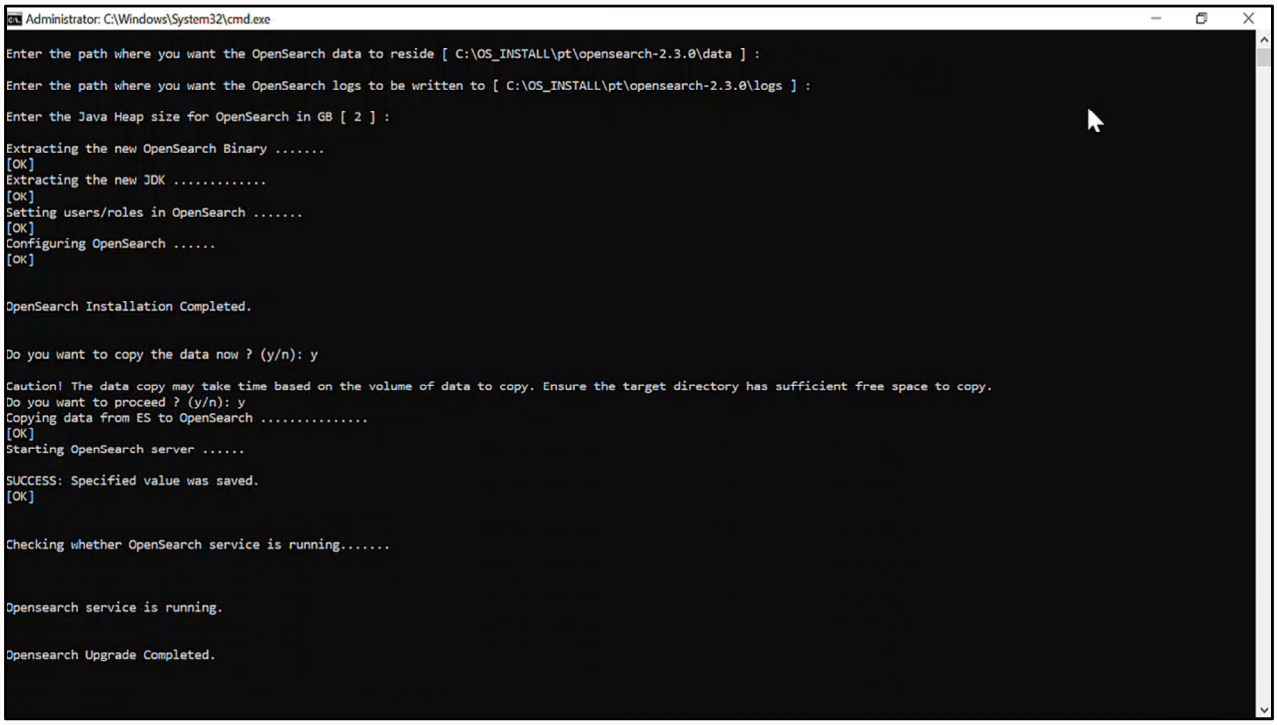

Once the data copy is complete, follow the instructions on the PeopleSoft application that we previously covered in this video.

This completes your demonstration.

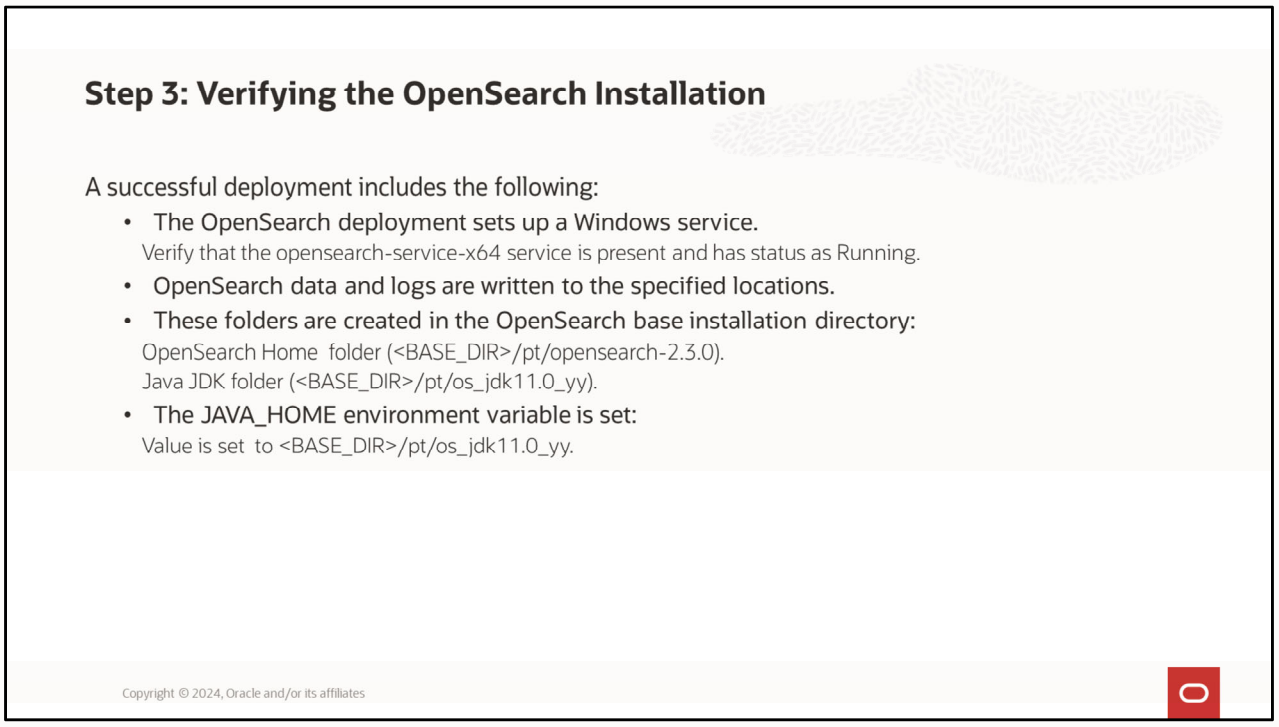

On your environment, you may perform checks on the file structure and environment variables to ensure that the deployment is successful, which includes:

- OpenSearch Windows service.
- Data and log files.
- Folders for OpenSearch home and Java JDK.
- Environment variable for Java_Home.

Refer to the PeopleSoft Deployment Packages Installation for Search Components (PeopleTools 8.61), which contains additional information that is not discussed in this video, such as:

- Additional deployment verification checks.
- Steps to remove an OpenSearch installation on Windows.
- Instructions on deploying OpenSearch on a Linux environment.

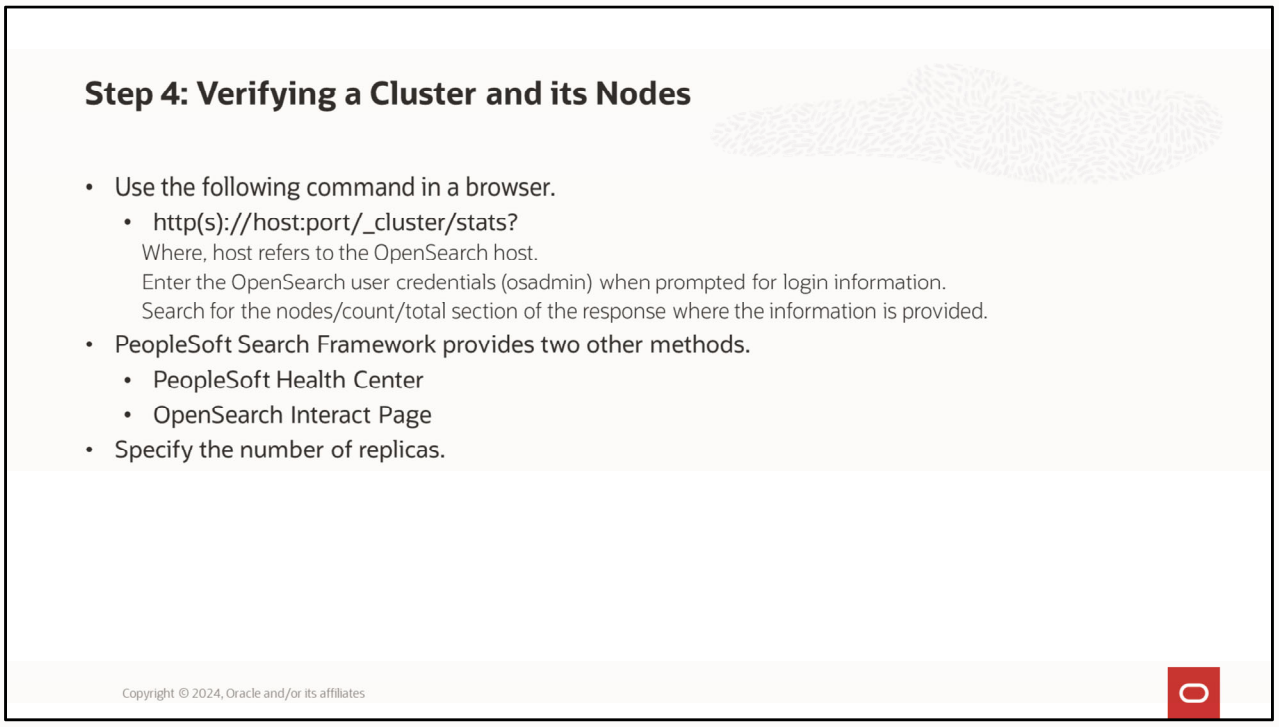

At this moment in the video, we've installed OpenSearch, but we've not integrated OpenSearch with the PeopleSoft environment. Before we integrate, we'll use the command as shown on your screen to verify a cluster and its nodes.
http(s)://host:port/_cluster/stats?
In the command, the host refers to the OpenSearch host.
You will be prompted to enter the OpenSearch user credentials. Enter osadmin.
In the response that is displayed, search for the nodes or count or total section. The information on cluster and nodes is provided in the section.
After we integrate the OpenSearch search engine with the PeopleSoft environment, we can take advantage of the options provided by PeopleSoft Search Framework to verify a cluster and its nodes. PeopleSoft Search Framework provides a couple of ways to verify a cluster and its nodes.

- PeopleSoft Health Center
- OpenSearch Interact page

These two options are not discussed in this video. Refer to the Working with OpenSearch Clusters topic in PeopleTools 8.61: Search Technology online help.
In a multi-node cluster, you must ensure that the replica values are set correctly. Replicas are set on the Search Options page, which we'll discuss later in this video.

After installing OpenSearch, the next task is to integrate OpenSearch with the PeopleSoft application.

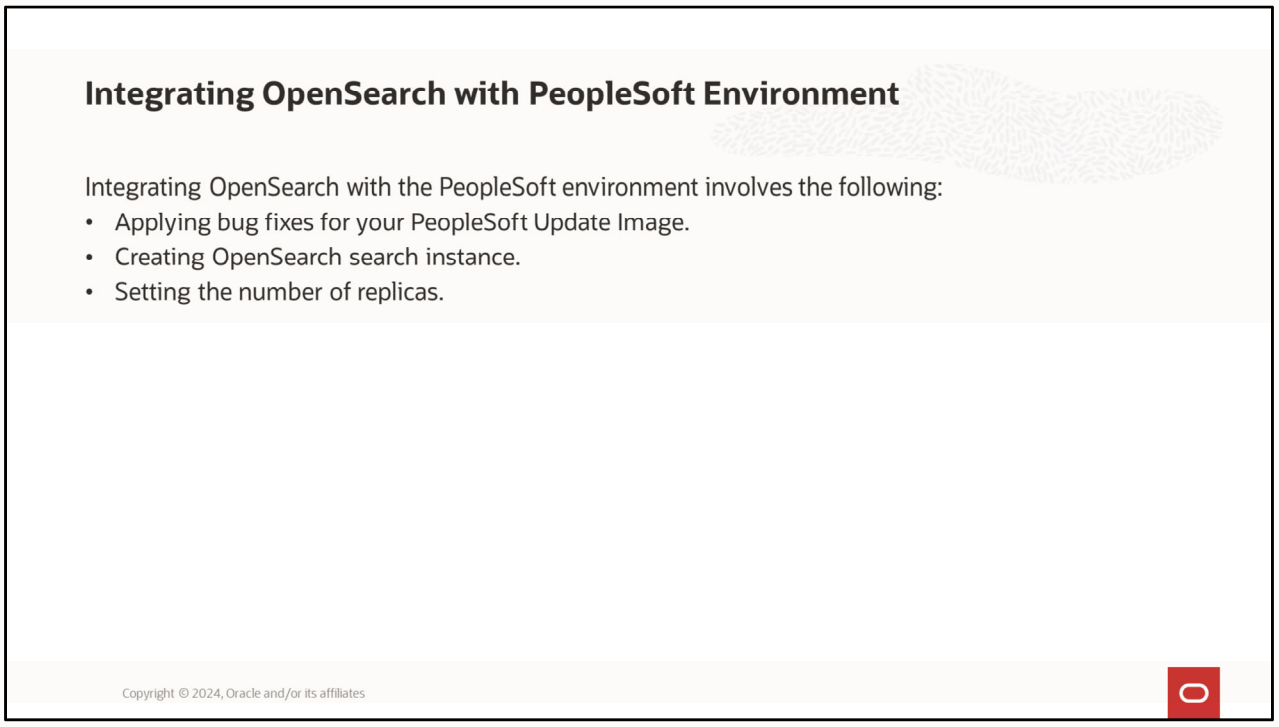

To complete the integration of the OpenSearch search engine with the PeopleSoft environment, the following tasks must be performed:

- You must identify and apply any necessary bugs delivered with the latest PeopleSoft Update Image (PI) for your PeopleSoft application.
- You must create a search instance using either the Search Instance Properties page or the configuration template (SEARCH_TEMPLATE) in Automated Configuration Manager (ACM).
- After creating a search instance, you need to specify the number of replicas on the Search Options page.

Let's look at each of these tasks separately.

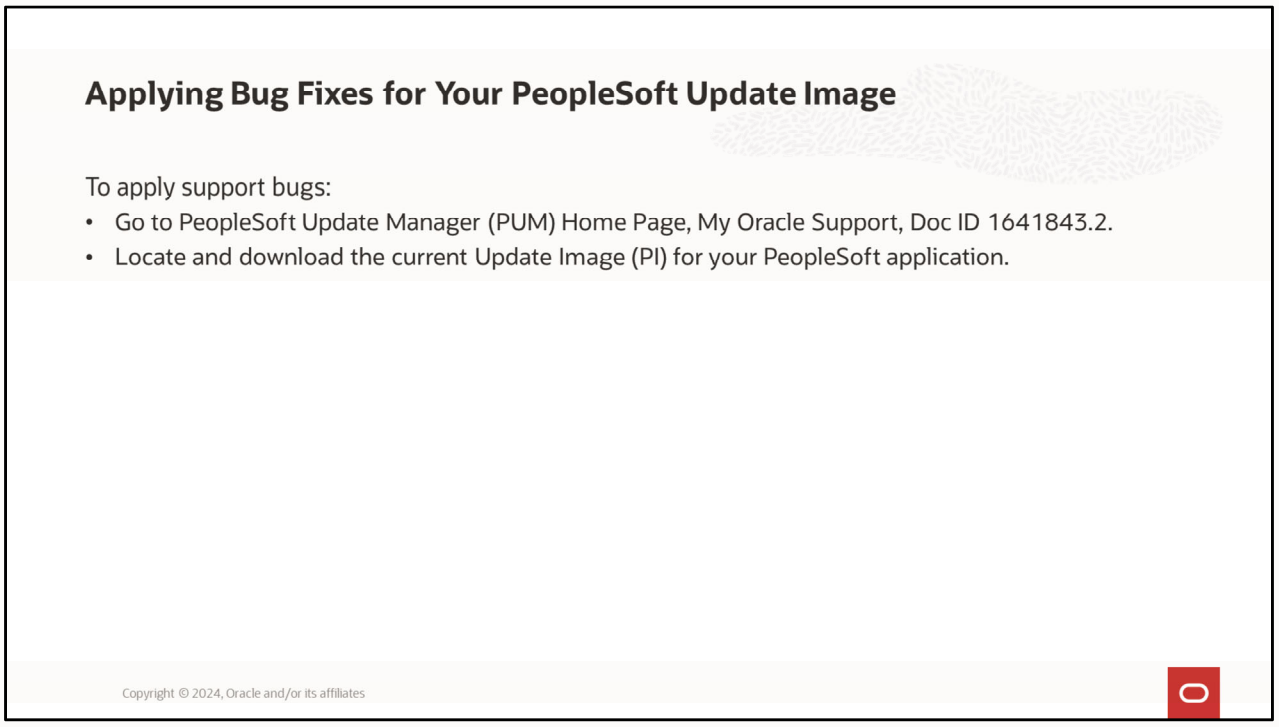

To identify and apply any necessary bugs delivered with the latest PeopleSoft Update Image, check the PeopleSoft Update Manager (PUM) Home Page on My Oracle Support (Doc ID 1641843.2) and locate and download the current Update Image (PI) for your PeopleSoft application.

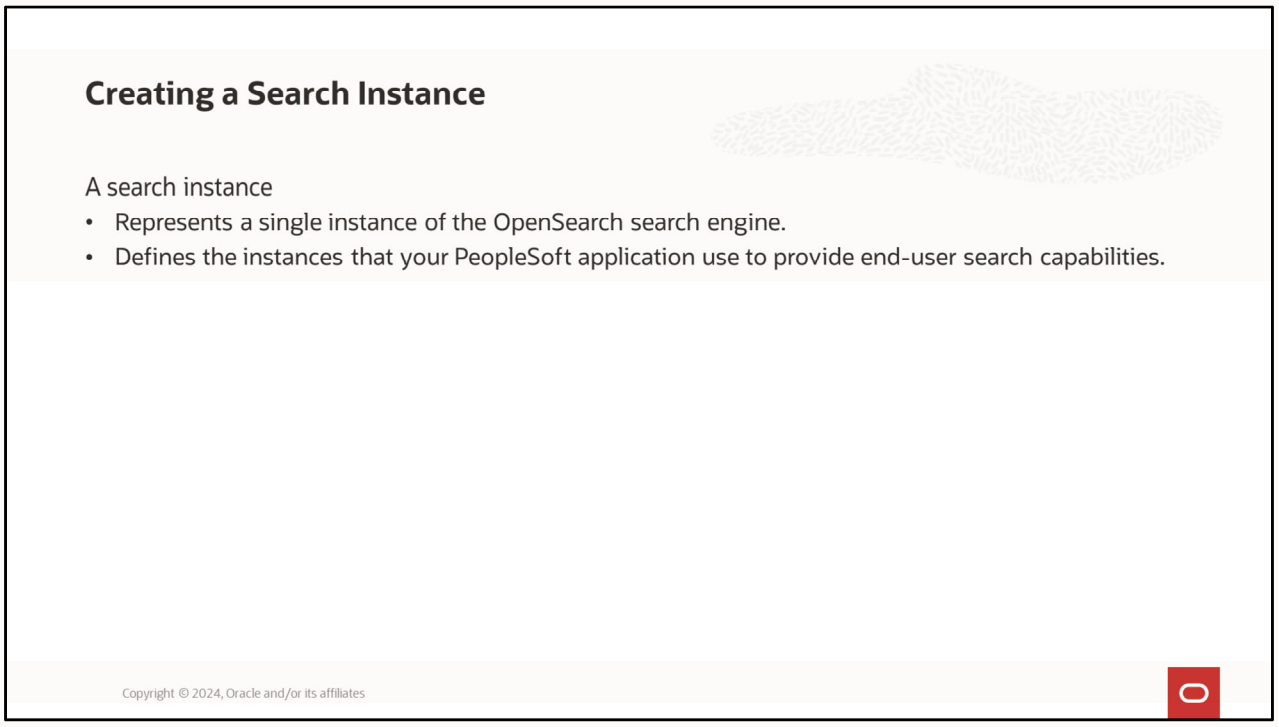

A search instance in the PeopleSoft Search Framework represents a single instance of the search engine.
The search instance(s) created using the PeopleSoft Search Framework define the instances that will be used by the PeopleSoft applications to provide search capabilities for an end user.

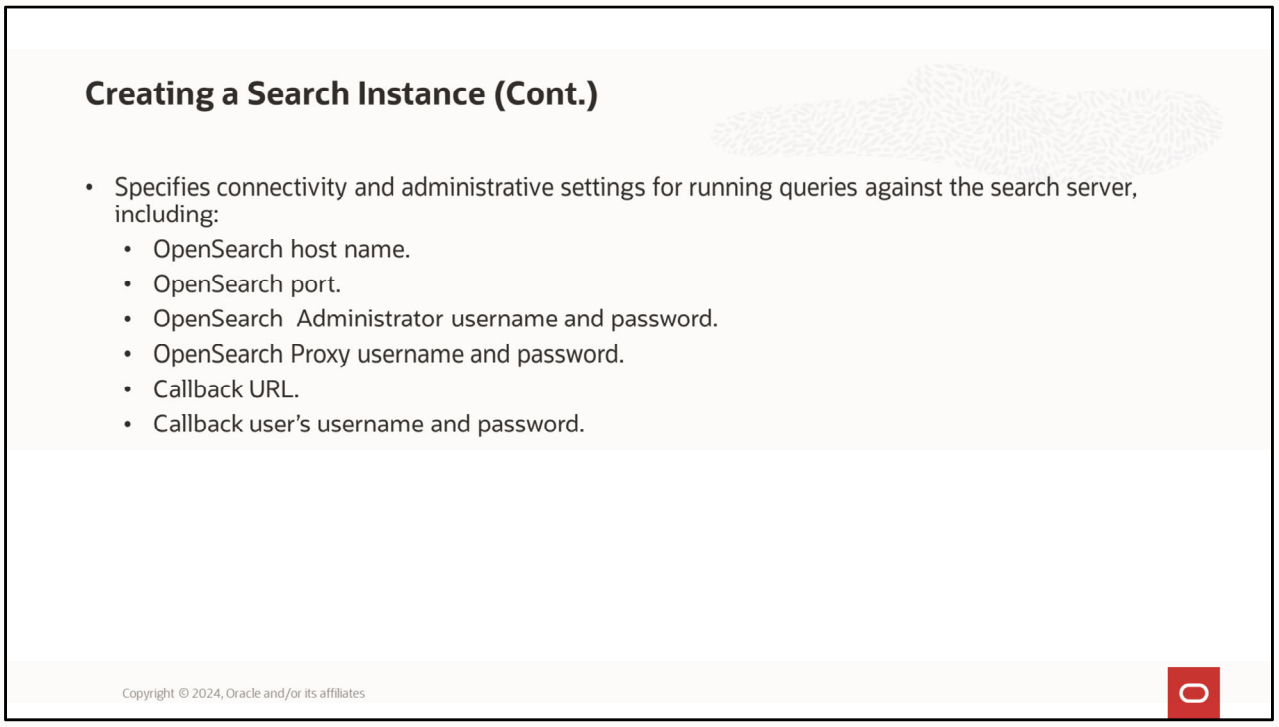

The search instance definition specifies connectivity and other administrative settings required for running queries against the search server and for administrative tasks. In a search instance, you provide values for:

- Host name of the search provider (OpenSearch)
- The port value.
- The administrator username and password that you entered while running the OpenSearch DPK.
- The proxy username and password that you entered while running the OpenSearch DPK.
- Callback URL that points to the gateway defined in Integration Broker.
- Callback user's username and password. This user must have the following roles - PeopleSoft User, Search Query Administrator, and Search Server.

Creating a Search Instance (Cont.)

Prerequisites:

- Verify that you have the following security roles:
  - Search Administrator
  - Search Developer
  - Search Server
  - ACM Administrator
  - Insights Administrator

See the Search Technology online help and the PeopleSoft Deployment Packages Installation for Search Components document for more information.

Before creating the instance, verify that you have the necessary roles - Search Administrator, Search Developer, Search Server, Insights Administrator, and ACM Administrator – which are described in the installation documentation.

Creating a Search Instance (Cont.)

PeopleTools delivers two methods for creating search instances:

- Search Administration Activity Guide
  - Configuration activity / Search Instance Task.
  - Search Instance Properties page.
  - Creates only the search instance.

  Tasks of deploying search definition and search categories, and indexing are performed separately.
- Automated Configuration Manager
  - Use an ACM template.
  - Creates search instance, deploys search definitions and search categories, and builds indexes.

There are two ways to create an OpenSearch search instance.
The Search Administration Activity Guide Configuration activity (using PIA) features a Search Instance task. In this task you use the Search Instance Properties page to create the search instance.
Using the Automated Configuration Management (ACM) enables you to automate the configuration of an OpenSearch instance. You can also use Automated Configuration Manager to automate the deploying of search definitions and categories, as well as automate indexing.

Creating a Search Instance (Cont.)

Using the Search Administration Activity Guide/
Search Instance Properties Page
Navigation:
1. Select PeopleTools, Search Framework,
Search Admin Activity Guide.
2. In the Configuration activity, click the Search Instance task.

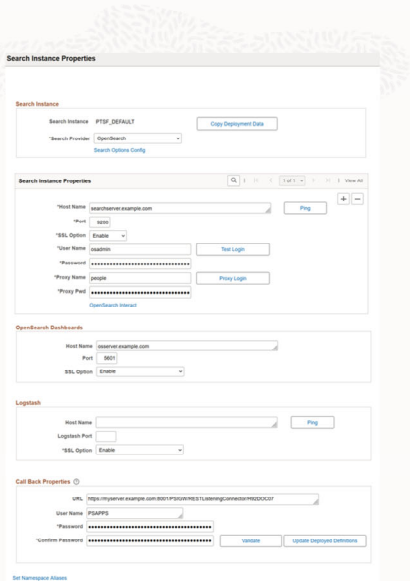

The first method of creating a search instance is by using the Search Instance Properties page.
This example shows the Search Instance Properties page located in the Search Administration Activity Guide.
To navigate to this page, select PeopleTools, Search Framework, Search Admin Activity Guide.
In the Configuration activity, click the Search Instance task.

Creating a Search Instance (Cont.)

Using Automated Configuration Manager – Configuration Template Definition Page Navigation:

1. Select PeopleTools, Automated Config Manager, ACM Templates, Template Definition.
2. Select the SEARCH_TEMPLATE.
3. Edit the template.
4. Save and run the template.

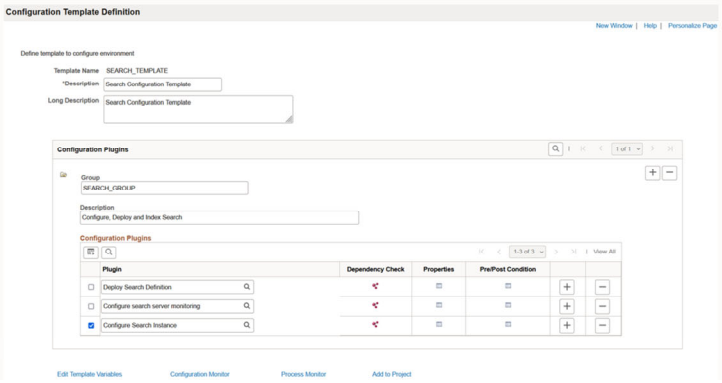

To use Automated Configuration Manager to create an OpenSearch instance, you use a delivered template (SEARCH_TEMPLATE) that automates the configuration process.

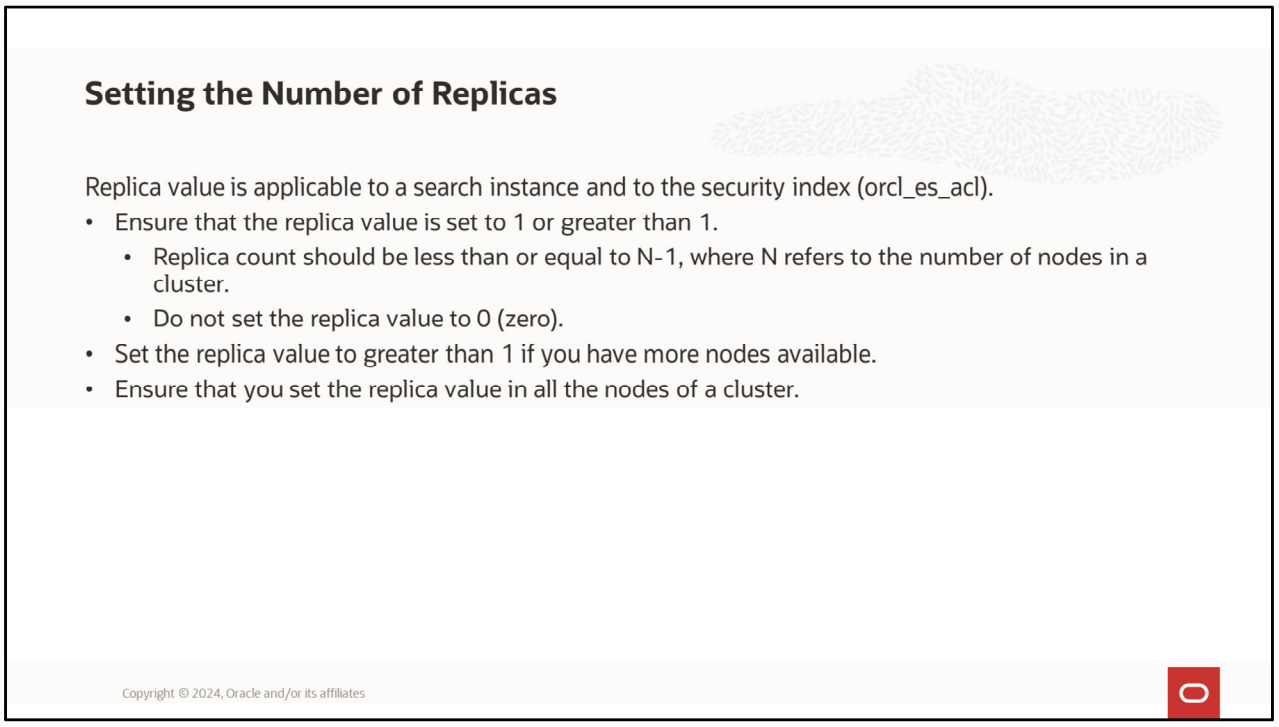

On the Search Options page, when you specify a replica value in PeopleTools 8.61, the replica value is applicable to a search instance and to the security index (orcl_es_acl).
Ensure that the replica value is set to 1 or greater than 1 so that in the event of a node failure data loss does not happen. The maximum number of nodes that can be down without loss of data will be equal to the number of replicas (replica count should be less than or equal to N-1), where N refers to the number of nodes in a cluster.
You may set the replica value to greater than 1 if you have more nodes available. For example, if you have 4 nodes in a cluster, and you want to provide high availability even if two nodes fail, then you need to set the replica value to 2.
If the replica value for an index is greater than N-1 (where N is the number of nodes in the cluster), you may find that the status of the cluster as yellow indicating a warning that at least one replica is missing.

Setting the Number of Replicas (Cont.)

Using the Search Options page
Navigation:
1. Select PeopleTools, Search Framework, Search Admin Activity Guide.
2. In the Configuration activity, click the Search Options task.
3. In the Number of Replicas field, enter a value.

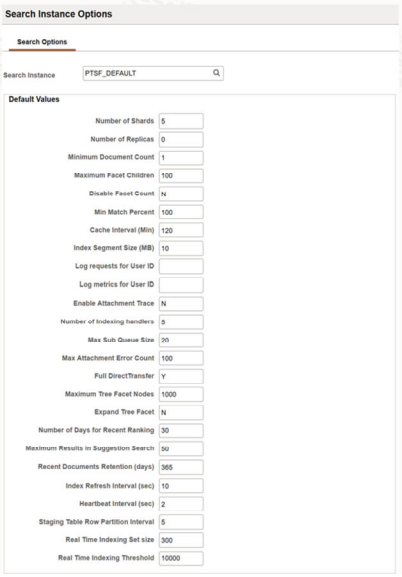

Specify the number of replicas on the Search Options page. This example shows the Search Options page located in the Search Administration Activity Guide.
To navigate to this page, select PeopleTools, Search Framework, Search Admin Activity Guide.
In the Configuration activity, click the Search Options task, and enter a value in the Number of Replicas field.
With these steps, we've completed the integration of OpenSearch with the PeopleSoft application.

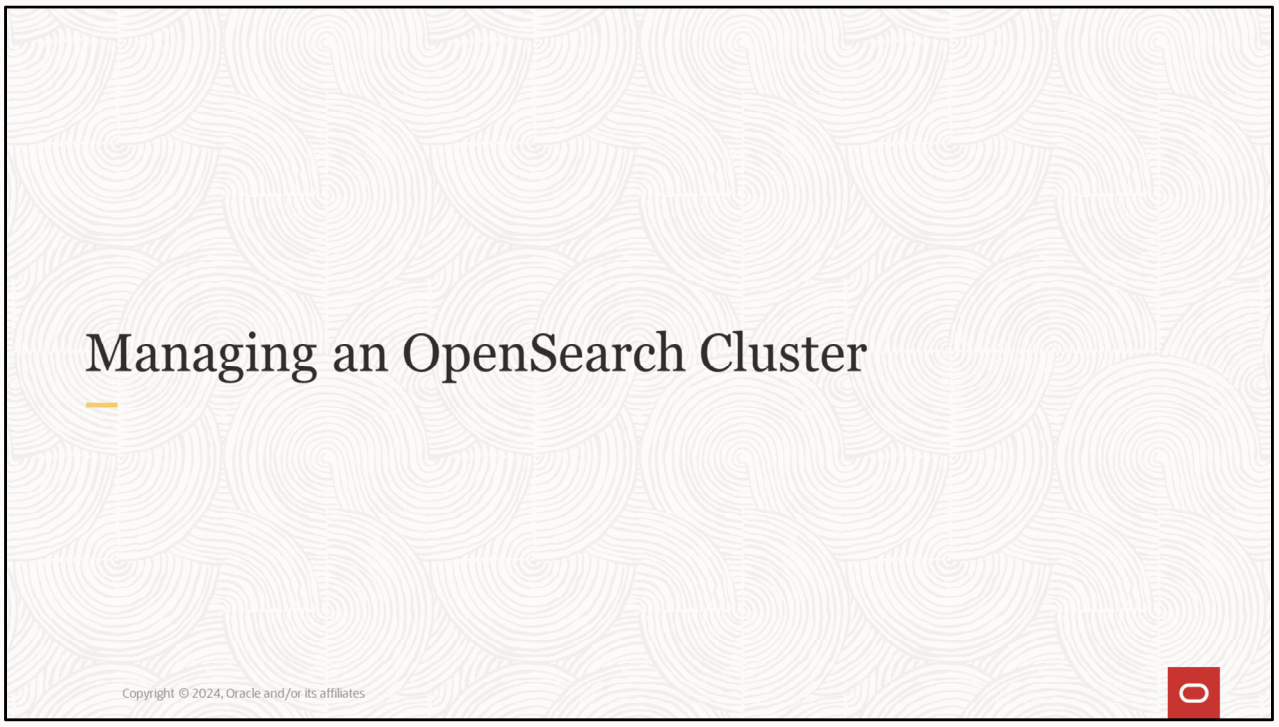

Let us now discuss the OpenSearch cluster that is created and some additional tasks to maintain the cluster.

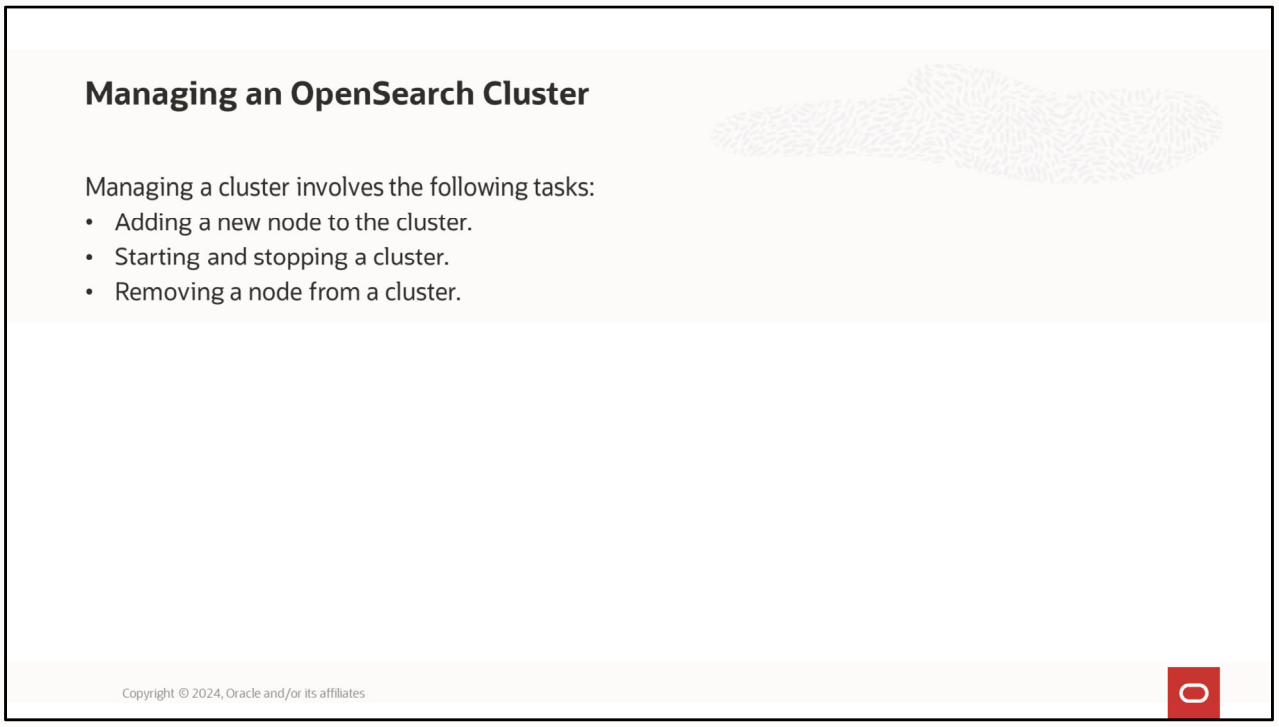

With the deployment of OpenSearch, a cluster has been created. Depending on your business requirements, you may have to perform these tasks to manage the cluster:

- Adding a new node to the cluster.
- Starting and stopping a cluster.
- Removing a node from a cluster.

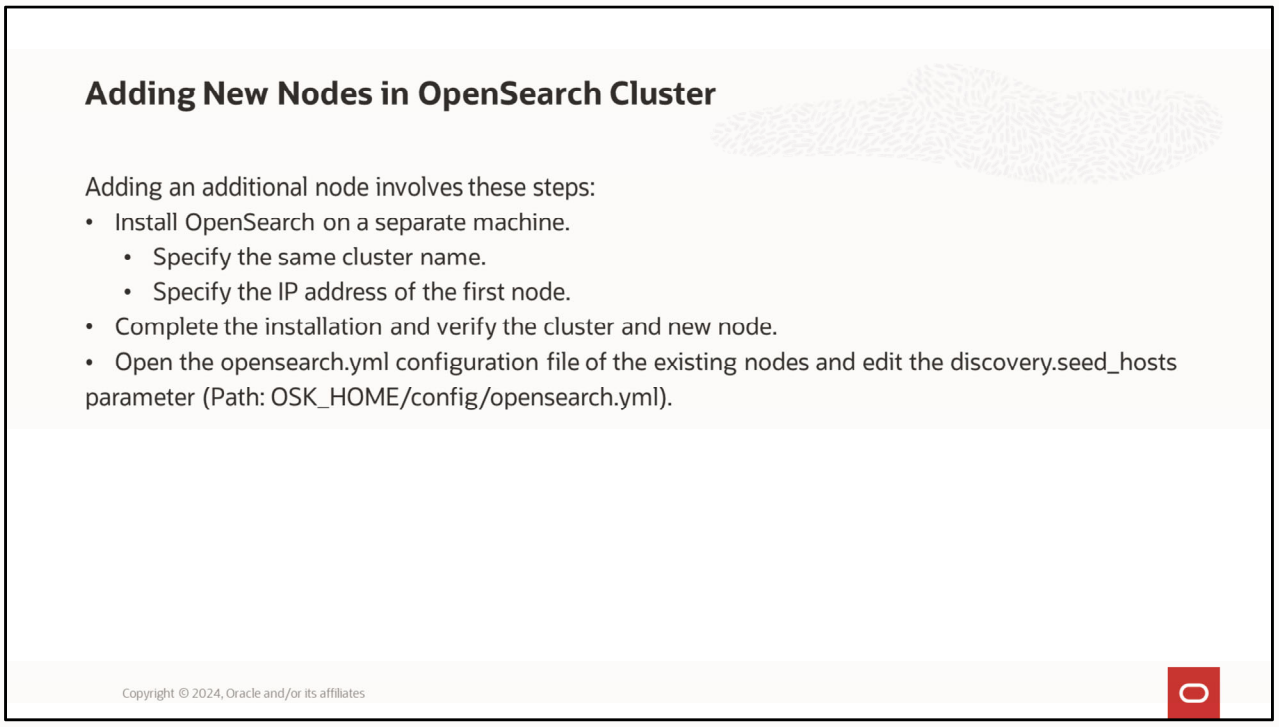

OpenSearch manages failover through multiple nodes within a cluster, so that when one node fails the other nodes in the cluster provide the indexed data when a search is initiated.
Use these steps to add an additional OpenSearch node after you have completed the OpenSearch DPK installation.
Note that the new node is added to the same cluster.
The first step is to install OpenSearch on a separate machine using the OpenSearch DPK.
During installation provide the following:

- For the cluster name, specify the name of the existing OpenSearch cluster. You need to enter this value so that the new node will join the cluster.
- Specify the IP address of the nodes present in the existing OpenSearch cluster. Follow the instructions on entering host names as described earlier in the video.

Next, complete the installation and then verify the cluster and the new node. If you are adding the new node after connecting OpenSearch server with the PeopleSoft application, then you can use the options provided by Search Framework to verify a cluster and nodes.
Finally, open the opensearch.yml configuration file of the existing nodes in the cluster and edit the parameters. The configuration file is available in OSK_HOME/config folder.

- discovery.seed_hosts - Add the host of the new node to the existing list of hosts. OpenSearch nodes will find each other via unicast. Ensure that you enter the address in the correct format, which is described in an earlier part of this video.

Note that the .yml file uses space indentation, so ensure that you edit the file properly and save it.

Also note that when you add a new node to a cluster, you also need to update the search instance with the details of the new node. You need to complete this task on the Search Instance Properties page. Refer to the PeopleTools 8.61: Search Technology online help.

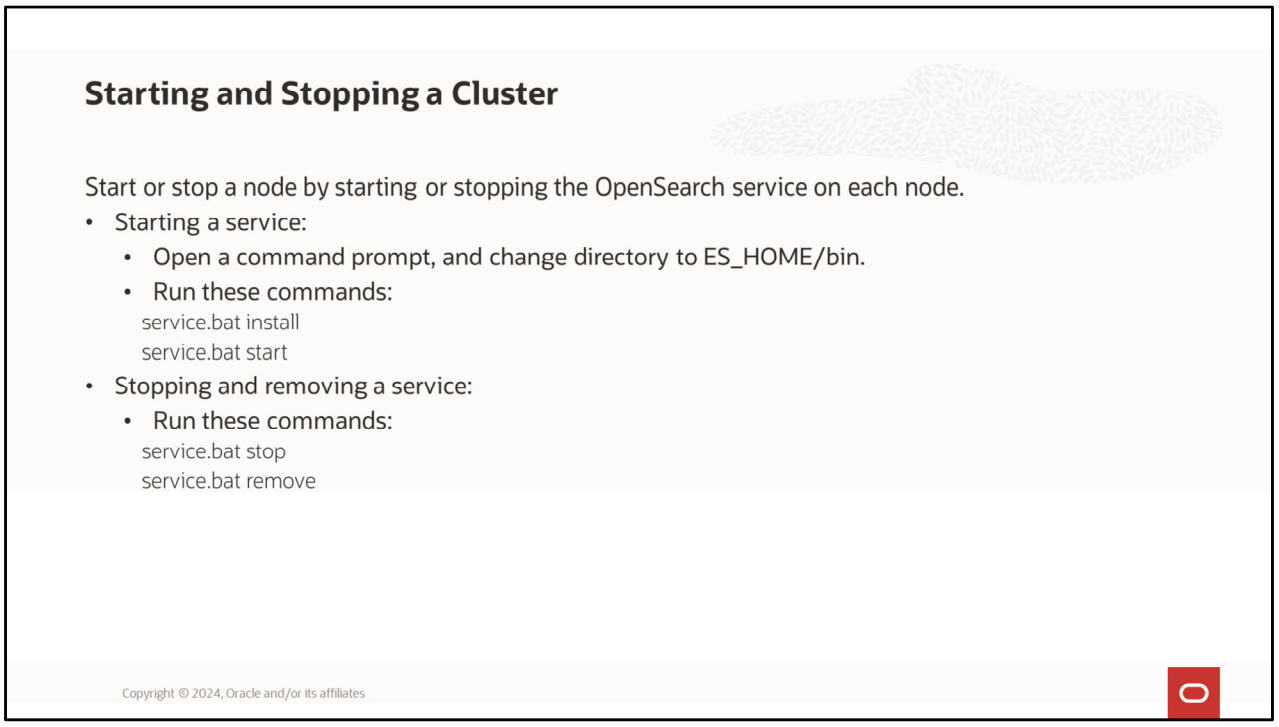

Starting a cluster would mean starting all the nodes of a cluster. Similarly, stopping a cluster would mean stopping all the nodes of a cluster. You start or stop a node by starting or stopping the OpenSearch service on each node.
Before you commence to stop a node, you must ensure that no indexing requests or administration-related tasks are being made or performed on the cluster because during indexing, the cluster meta data gets updated and master nodes broadcast the meta data to all other master nodes. If you stop a node during indexing, the cluster meta data might get corrupted, and the cluster could become non-operational (red color code). To ensure that no instances of PTSF_GENFEED are running, check the Process Monitor. Then you may stop all the nodes in a cluster and make the required modifications. After completing the modifications, you may start all the nodes of the cluster.
In this video, we're discussing the starting and stopping of an OpenSearch service on a Windows environment.
To start a service:
Open a command prompt window and change the directory to OS_HOME/bin and run these commands:
- Service.bat install
- Service.bat start

To stop and remove a service:

Open a command prompt window and change the directory to OS_HOME/bin and run these commands:

- Service.bat stop
- Service.bat remove

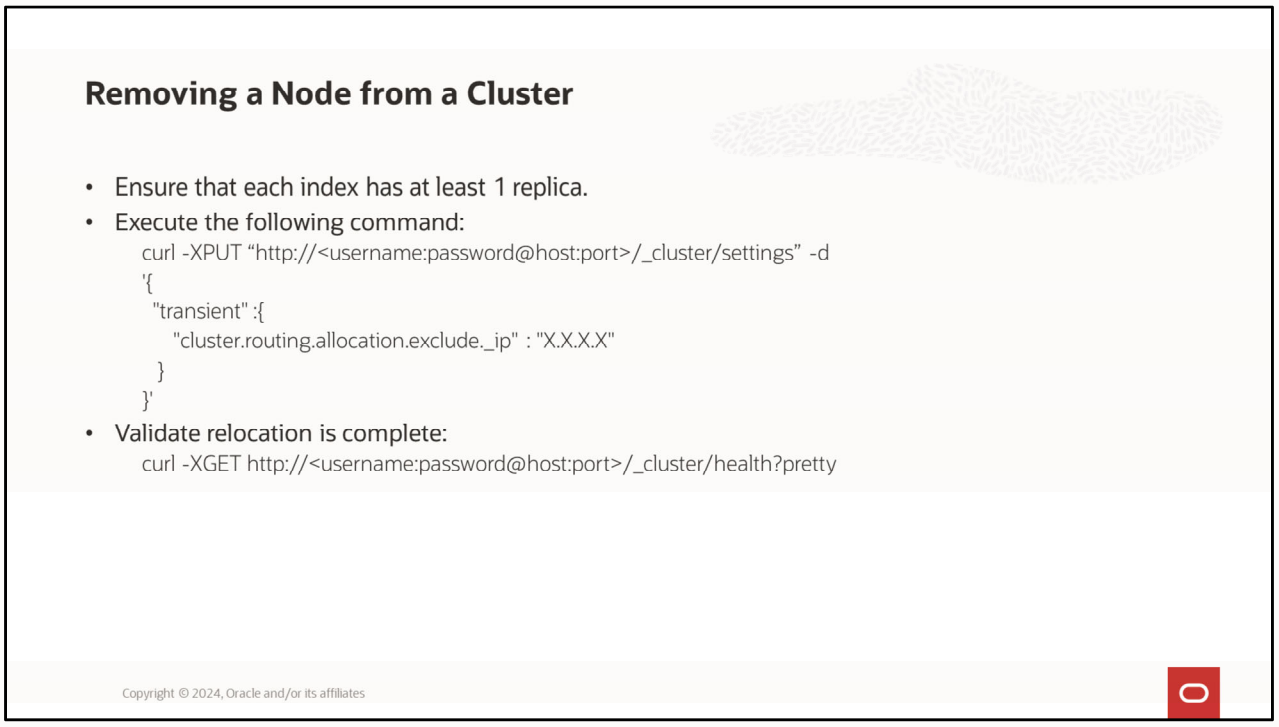

Before you commence the process of removing a node, ensure that each index has at least 1 replica. You can use the OpenSearch Interact page to verify the replica count by choosing the option highlighted on your screen from the Cluster API drop-down list.
Alternatively, you can use the OpenSearch Cluster Dashboard to view the replica information in PeopleSoft Health Center.
Identify the IP address of the OpenSearch node that needs to be removed from the cluster.
Open a command prompt window and execute the given command where X.X.X.X stands for the IP address of the node that needs to be removed and also replace the host:port with the IP or port of any of the nodes in the cluster.
When the command is executed, OpenSearch tries to move the existing shards out of the node that will be removed and moves it to other nodes in the cluster.
In a command prompt window, run the given command to ensure that the relocation is complete and to validate that the relocating_shards attribute shows value as 0 (zero). Alternatively, on the OpenSearch Interact page, you may use the option highlighted on your screen from the Cluster API drop-down list.

Removing a Node from a Cluster (Cont.)

- Shut down the node.
- Set the exclusion rules to empty:

```
curl -XPUT “http://username:password@host:port/_cluster/settings” -d
'{
  "transient" :{
     "cluster.routing.allocation.exclude._ip" : ""
  }
}'
```

- Update the Search Instance Properties page.
- Verify the status of the cluster.

Shutdown the node that was identified for removal.
Set the exclusion rules to empty so that, in the future, the removed node can be added back to the cluster and can be used to hold the shards.
On the Search Instance Properties page, delete the row corresponding to the removed node.
You can verify the status of the cluster by using the OpenSearch Cluster Dashboard. The cluster’s status should be either green (cluster is fully operational) or yellow (some replicas are not assigned, but the cluster is fully functional). Refer to the Monitoring OpenSearch Cluster Using PeopleSoft Health Center topic in PeopleTools 8.61: Search Technology online help.

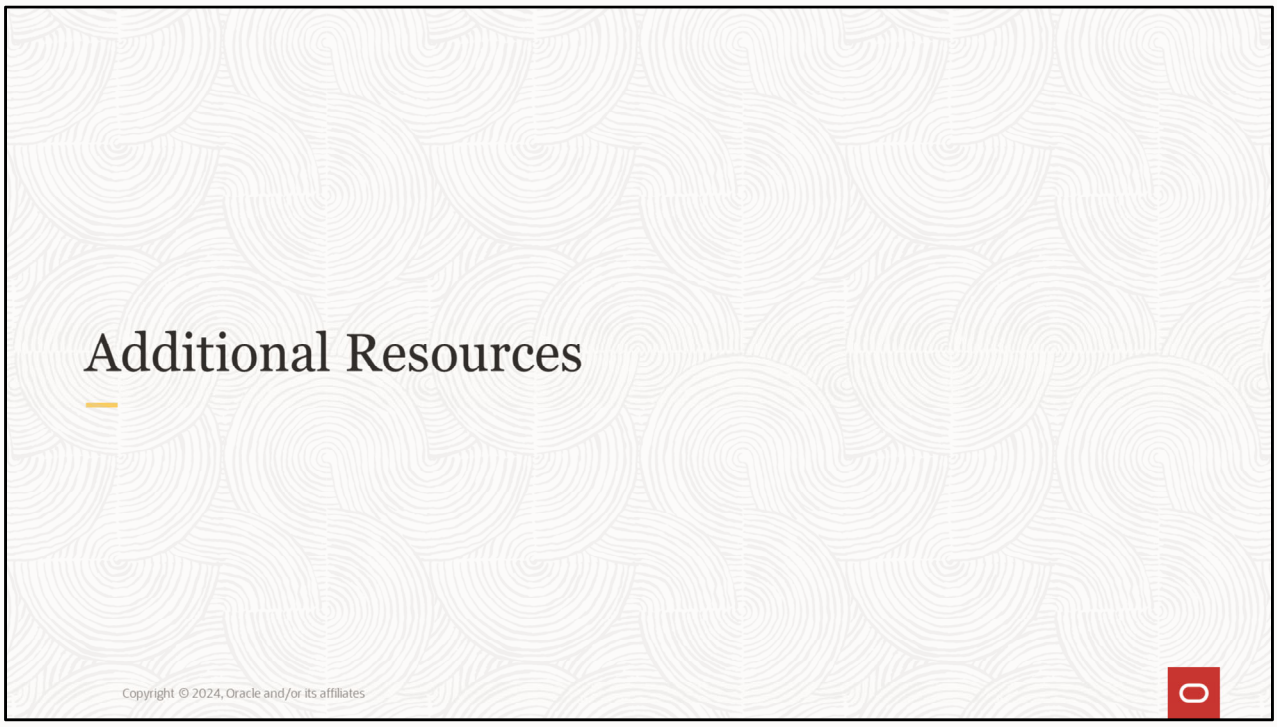

Before we complete this session, let's take a quick look at the various sources of information available on PeopleSoft Search Framework and OpenSearch. It will be beneficial to take a look at the table containing the terminology used in OpenSearch and Search Framework.

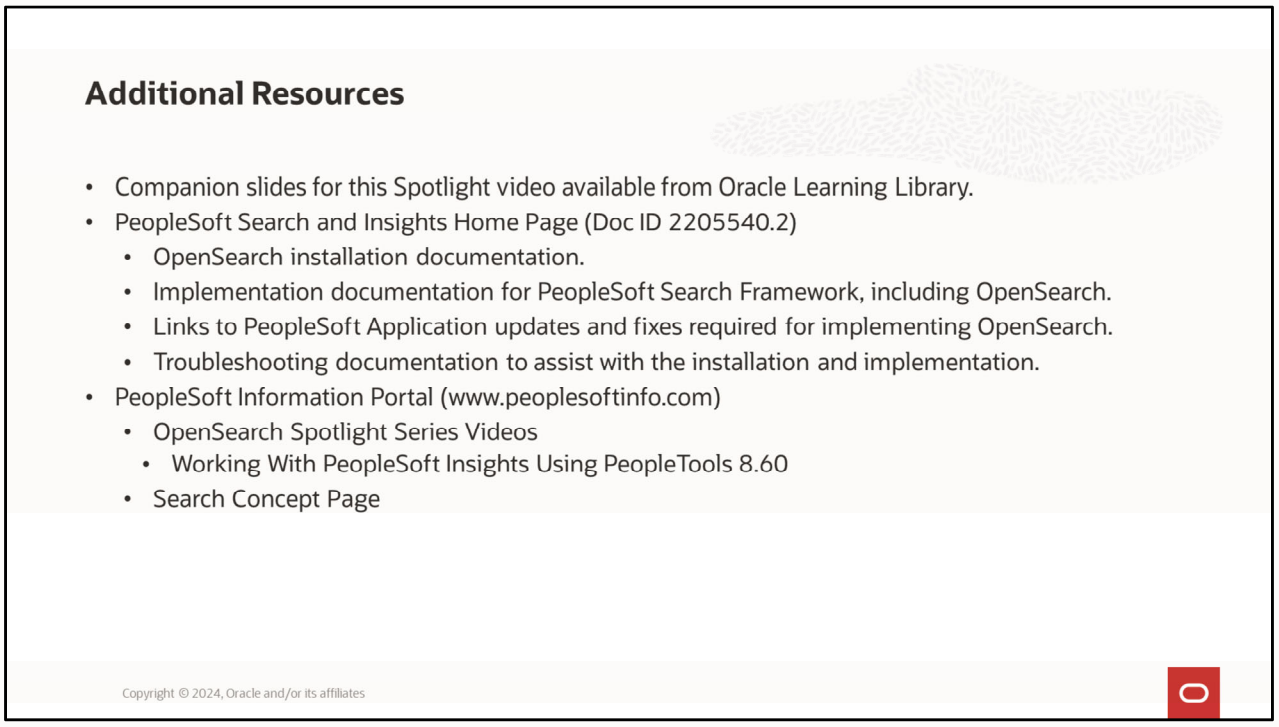

You can download the Companion Slides, with images and notes from this video, from the Oracle Learning Library.
The PeopleSoft Search and Insights Home Page includes installation documentation and implementation documentation for PeopleSoft Search Framework, including OpenSearch. Contains links to troubleshooting documentation to assist with the installation and implementation.
On the PeopleSoft Information Portal (peoplesoftinfo.com), you can find OpenSearch Spotlight Series videos and PeopleSoft Insights Concept page that contain information about how to use OpenSearch Dashboards in PeopleSoft.

Additional Resources

OpenSearch Terminology

Term	Use in OpenSearch	Use in Search Framework
Cluster	A collection of one or more nodes (servers) that together holds your entire data and provides federated indexing and search capabilities across all nodes.	A cluster is specified (named) during OpenSearch deployment.
Node	A node is a single server that is part of your cluster, stores your data, and participates in the cluster's indexing and search capabilities.	A node corresponds to a search instance in PeopleSoft Search Framework.
Index	An index is a collection of documents that have somewhat similar characteristics. In a single cluster, you can define as many indexes as you want.	In a PeopleSoft implementation, by default, all the search definitions/search categories are deployed as part of a single index. However, you can create multiple indexes.

Additional Resources (Cont.)

OpenSearch Terminology

Term	Use in OpenSearch	Use in Search Framework
Type	Within an index, you can define one or more types. A type is a logical category/partition of your index whose semantics is completely up to you.	In a PeopleSoft implementation, each search definition corresponds to a type.
Alias	Alias is a reference to an OpenSearch index. An alias can be mapped to more than one index. That is, an alias enables you to group indexes.	In a PeopleSoft implementation, a search category is mapped as an Alias. The mapping of a search category to Alias allows you to group the search definitions (types in OpenSearch) so that when search is performed on a search category, it returns documents from all the search definitions under it.

Additional Resources (Cont.)

OpenSearch Terminology

Term	Use in OpenSearch	Use in Search Framework
Document	A document is a basic unit of information that can be indexed. This document is expressed in JavaScript Object Notation (JSON) format.	Connected Query returns parent and child rows. In a PeopleSoft implementation, each row returned from the main query corresponds to a document and child information is attached to the main query and is sent as one document.
Shards	Elasticsearch provides the ability to subdivide your index into multiple pieces called shards. Thus, shards allow you to horizontally split/scale your content volume. When you create an index, you can define the number of shards that you want. Default value is 5.	Specify the number of shards on the Search Options page in PIA. After an index is created, you cannot change the number of shards.

Additional Resources (Cont.)

OpenSearch Terminology

Term	Use in OpenSearch	Use in Search Framework
Replica Shards	OpenSearch allows you to make one or more copies of your index's shards into what are called replica shards, or replicas for short. Replica shards provide high availability in case a shard or a node fails.	Specify the number of replica shards on the Search Options page in PIA. After an index is created, you may change the number of replicas dynamically anytime.


ORACLE