ORACLE®
PeopleSoft

PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.57)

October 2019

ORACLE®

PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.57)
Copyright © 2019, Oracle and/or its affiliates. All rights reserved.

Contents

Chapter 3

Chapter 4

Appendix A

About This Documentation

Understanding This Documentation

This documentation is designed to guide you through the deployment of the Oracle's PeopleSoft Deployment Packages. It is not a substitute for the documentation provided for PeopleSoft PeopleTools or PeopleSoft applications.

Typographical Conventions

To help you locate and understand information easily, the following conventions are used in this documentation:

Convention	Description
`Monospace`	Indicates a PeopleCode program or other code, such as scripts that you run during the install. Monospace is also used for messages that you may receive during the install process.
Italics	Indicates field values, emphasis, and book-length publication titles. Italics is also used to refer to words as words or letters as letters, as in the following example: Enter the letter *O*. Italics are also used to indicate user-supplied information. For example, the term *domain* is used as a placeholder for the actual domain name in the user's environment. When two such placeholders are used together, they may be set apart with angle brackets. For example, the path *<PS_CFG_HOME>/*appserv/*<domain>* includes two placeholders that require user-supplied information.
Initial Caps	Field names, commands, and processes are represented as they appear on the window, menu, or page.
lower case	File or directory names are represented in lower case, unless they appear otherwise on the interface.
Menu, Page	A comma (,) between menu and page references indicates that the page exists on the menu. For example, "Select Use, Process Definitions" indicates that you can select the Process Definitions page from the Use menu.

Convention	Description
Cross-references	Cross-references that begin with *See* refer you to additional documentation that will help you implement the task at hand. We highly recommend that you reference this documentation. Cross-references under the heading *See Also* refer you to additional documentation that has more information regarding the subject.
⇒ (line-continuation arrow)	A line-continuation arrow inserted at the end of a line of code indicates that the line of code has been wrapped at the page margin. The code should be viewed or entered as a continuous line of code, without the line-continuation arrow.
" " (quotation marks)	Indicate chapter titles in cross-references and words that are used differently from their intended meaning.
Note. Note text.	Text that begins with *Note.* indicates information that you should pay particular attention to as you work with your PeopleSoft system.
Important! Important note text.	A note that begins with *Important!* is crucial and includes information about what you need to do for the system to function properly.
Warning! Warning text.	A note that begins with *Warning!* contains critical configuration information or implementation considerations; for example, if there is a chance of losing or corrupting data. Pay close attention to warning messages.

Products

This documentation may refer to these products and product families:

- Oracle® BPEL Process Manager
- Oracle® Enterprise Manager
- Oracle® Tuxedo
- Oracle® WebLogic Server
- Oracle's PeopleSoft Application Designer
- Oracle's PeopleSoft Change Assistant
- Oracle's PeopleSoft Change Impact Analyzer
- Oracle's PeopleSoft Data Mover
- Oracle's PeopleSoft Process Scheduler
- Oracle's PeopleSoft Pure Internet Architecture
- Oracle's PeopleSoft Customer Relationship Management
- Oracle's PeopleSoft Enterprise Learning Management

- Oracle's PeopleSoft Enterprise Performance Management

- Oracle's PeopleSoft Financial Management

- Oracle's PeopleSoft Human Capital Management

- Oracle's PeopleSoft Interaction Hub

- Oracle's PeopleSoft Pay/Bill Management

- Oracle's PeopleSoft PeopleTools

- Oracle's PeopleSoft Staffing Front Office

- Oracle's PeopleSoft Supply Chain Management

See the Products area on the Oracle web site, http://www.oracle.com/us/products/product-list/products-a-z/index.html.

Related Information

Oracle provides reference information about PeopleSoft PeopleTools and your particular PeopleSoft Application. You can access documentation for recent releases of PeopleSoft PeopleTools and PeopleSoft Applications at the PeopleSoft Hosted Documentation site. You can also find documentation by searching for the product name on My Oracle Support.

- My Oracle Support. This support platform requires a user account to log in. Contact your PeopleSoft representative for information.

  To locate documentation on My Oracle Support, search for the title and select PeopleSoft Enterprise to refine the search results.

  See My Oracle Support, https://support.oracle.com.

- *PeopleTools: Getting Started with PeopleTools* for your release. This documentation provides a high-level introduction to PeopleTools technology and usage.

  See Oracle PeopleSoft Online Help, http://www.peoplesoftonlinehelp.com.

- PeopleSoft Application Fundamentals for your PeopleSoft Application and release. This documentation provides essential information about the setup, design, and implementation of your PeopleSoft Application.

To install additional component software products for use with PeopleSoft products, including those products that are packaged with your PeopleSoft products as well as products from other vendors, you should refer to the documentation provided with those products, as well as this documentation. For those additional components that are offered by Oracle, such as Oracle Middleware products, see the documentation on the Oracle Help Center.

See Also

Oracle Help Center, https://docs.oracle.com/en/

Comments and Suggestions

Your comments are important to us. We encourage you to tell us what you like, or what you would like changed about PeopleSoft documentation and other Oracle reference and training materials. Please send your suggestions to:

PSOFT-Infodev_US@oracle.com

While we cannot guarantee to answer every email message, we will pay careful attention to your comments and suggestions. We are always improving our product communications for you.

Chapter 1

Preparing to Deploy

This chapter discusses:

- Understanding Elasticsearch
- Prerequisites

Understanding Elasticsearch

This section discusses:

- Understanding Elasticsearch
- Understanding the Elasticsearch Deployment Package

Understanding Elasticsearch

Elasticsearch is an open-source search engine used for the PeopleSoft Search Framework for PeopleSoft 9.2 applications on the current PeopleTools releases. Elasticsearch is open-source software based on Apache Lucene™, a Java-based information retrieval library.

Oracle provides Elasticsearch as deployment packages (DPKs) for Microsoft Windows and Linux that deliver the required Elasticsearch software version, Java-based plug-ins needed for integration with PeopleSoft environments, and customized code where required. Be sure to obtain and use the Elasticsearch DPKs as described in this documentation.

The Elasticsearch DPK includes Kibana, an open-source analytics and visualization platform. The installation of Kibana is optional.

Elasticsearch is supported on Linux and Microsoft Windows operating systems. Oracle recommends that Elasticsearch be installed on servers that are separate from those used for PeopleSoft installations.

The deployment of Elasticsearch for PeopleSoft environments includes the following high-level steps:

1. Download the Linux or Microsoft Windows version of the Elasticsearch DPK from My Oracle Support or Oracle Software Delivery Cloud.

   See "Deploying the Elasticsearch Deployment Package," Obtaining the Elasticsearch DPK.

2. Extract the DPK zip file and run the script to deploy and set up Elasticsearch.

   See "Deploying the Elasticsearch Deployment Package."

3. Identify and apply any necessary bugs delivered with the latest PeopleSoft Update Image (PI) for your PeopleSoft application.

   See "Integrating Elasticsearch with the PeopleSoft Environment."

4. Set up the Elasticsearch instance in the PeopleSoft application.

   See "Integrating Elasticsearch with the PeopleSoft Environment."

Elasticsearch is the supported search engine for PeopleSoft 9.2 applications on PeopleSoft PeopleTools 8.57. For more information Elasticsearch, see these resources:

- For details on using Elasticsearch with the PeopleSoft Search Framework, see *PeopleTools: Search Technology.*

  See PeopleSoft Hosted Online Help, http://www.peoplesoftonlinehelp.com.

- You can find the most current version of this installation documentation, *PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.57),* as well as other information you need to implement Elasticsearch within the PeopleSoft Search Framework, on the Elasticsearch Home Page.

  See PeopleTools Elasticsearch Home Page, (select the tab Elasticsearch for PeopleTools 8.57), My Oracle Support, Doc ID 2205540.2, https://support.oracle.com/epmos/faces/DocumentDisplay?id=2205540.2.

- You can find links to the most current Elasticsearch DPK, which is available in the Patches & Updates area, on the Elasticsearch Home Page.

  See PeopleTools Elasticsearch Home Page, (select the tab Elasticsearch for PeopleTools 8.57), My Oracle Support, Doc ID 2205540.2, https://support.oracle.com/epmos/faces/DocumentDisplay?id=2205540.2.

- For more information about using PeopleSoft DPKs, see *PeopleSoft PeopleTools 8.57 Deployment Packages Installation.*

  See PeopleSoft PeopleTools Patches Home Page, My Oracle Support, Doc ID 2062712.2, https://support.oracle.com/epmos/faces/DocumentDisplay?id=2062712.2.

- The deployment of the VirtualBox version of the PeopleSoft DPKs for PeopleSoft Update Images includes the automatic installation and initialization of Elasticsearch.

  See *PeopleSoft Deployment Packages for Update Images Installation (PeopleSoft PeopleTools 8.57),* PeopleSoft Update Manager Home Page, My Oracle Support, Doc ID 1641843.2. https://support.oracle.com/epmos/faces/DocumentDisplay?id=1641843.2.

Understanding the Elasticsearch Deployment Package

To install Elasticsearch for the PeopleSoft Search Framework, you download and deploy the Elasticsearch DPK from Oracle Software Delivery Cloud or My Oracle Support. The Elasticsearch DPK includes:

- Open-source Elasticsearch software
- PeopleSoft-developed plug-ins for Elasticsearch
- PeopleSoft-delivered customized code for Elasticsearch
- Oracle Java JRE 1.8
- Open-source Python software
- PeopleSoft-developed Python scripts for Elasticsearch deployment automation
- Kibana 6.1.2

The Elasticsearch DPK installation enables automatic setup of Elasticsearch clusters, nodes, administrator user, and proxy user.

The following considerations were made when these DPKs were designed:

- System administration experience with PeopleTools is required. If you are new to the PeopleSoft system, it may be necessary to familiarize yourself with the PeopleSoft architecture before proceeding.

- System administration experience with Elasticsearch is required. If you are new to the Elasticsearch system, it may be necessary to familiarize yourself with the Elasticsearch architecture before proceeding.

Prerequisites

This section discusses:

- Reviewing Hardware Prerequisites
- Reviewing Software Prerequisites
- Reviewing Elasticsearch Recommendations

Reviewing Hardware Prerequisites

You can install the Elasticsearch DPK directly on a system running a Linux or Microsoft Windows operating system. The Elasticsearch DPK is certified to run on those Linux and Microsoft Windows operating systems that are certified for Elasticsearch for a PeopleSoft environment. The host can be a physical computer or a virtual machine.

- Host computer

  The Elasticsearch deployment packages can be installed on a Linux or Microsoft Windows host (bare metal or virtual). The recommendation is to install Elasticsearch on a server that is separate from those used for PeopleSoft installation.

- Host operating system

  The host operating system must be 64-bit Oracle certified platform.

  The integration of Elasticsearch and Kibana with PeopleSoft systems is supported for Microsoft Windows and Linux operating systems. For current support information for Elasticsearch and Kibana, see My Oracle Support.

  See My Oracle Support, Certifications.

  See PeopleSoft PeopleTools Certifications, My Oracle Support, Doc ID 747587.1, for help searching PeopleSoft Certifications.

- RAM (Memory)

  Oracle recommends a minimum of 32 GB available RAM for running an Elasticsearch environment, and preferably 64 GB.

  Note. See the information on heap size in the section Reviewing Elasticsearch Recommendations.

- Disk space

  A minimum of 100 GB free disk space is required for the Elasticsearch deployed environment.

- CPU

  A minimum of 4 CPUs is recommended.

  Note. Choose a modern processor with multiple cores. If you need to choose between faster CPUs or more cores, choose more cores. The extra concurrency that multiple cores offers will far outweigh a slightly faster clock speed.

Reviewing Software Prerequisites

Here are the software requirements for using the Elasticsearch DPK:

- PeopleSoft environment

  - The Elasticsearch 6.1.2 integration is supported for PeopleSoft 9.2 applications on PeopleSoft PeopleTools 8.57.

  - The Elasticsearch DPKs are updated regularly to incorporate the latest JRE. Elasticsearch DPKs are released concurrently with PeopleTools patches, but are not dependent on the PeopleTools patch release. You can use the Elasticsearch DPKs with earlier PeopleTools patches, as long as the PeopleTools patch level is greater than the minimum patch level listed on My Oracle Support Certifications. See the notes for Elasticsearch in the certifications for PeopleTools 8.57.

  - Oracle recommends that you use the Elasticsearch 6.1.2 DPK for the latest PeopleTools patch release to take advantage of the latest fixes and features. See the Elasticsearch home page for up-to-date information on features and fixes that require specific PeopleTools patch releases.

    See PeopleTools Elasticsearch Home Page, My Oracle Support, Doc ID 2205540.2.

- Secure Shell (ssh) client

  You need an SSH client to connect to the host for any PeopleSoft administrative tasks after the environment setup.

- Zip utility

  You need a utility that can extract (unzip) the DPK zip file on your operating system.

Reviewing Elasticsearch Recommendations

These specifications apply to the computer where you install Elasticsearch.

- Elasticsearch prerequisites

  Before installing the Elasticsearch DPK, review the prerequisites on the Elasticsearch web site. Elasticsearch has strict bootstrap validations. If you do not fulfill the prerequisites, the Elasticsearch instance may not start.

  See Bootstrap Checks, Elastic Docs web site, https://www.elastic.co/guide/en/elasticsearch/reference/6.1/bootstrap-checks.html.

- Heap size

  To adjust memory usage after installation, you can adjust the memory settings in the jvm.options properties file. Locate the jvm.options file in the config directory under the installation directory, and modify the values for Xms and Xmx.

  Xms represents the initial size of the total heap space.

  Xmx represents the maximum size of the total heap space.

  The standard recommendations are to set Xms and Xmx to the same value, and to give 50% of the available memory to the Elasticsearch heap, while leaving the other 50% free. The memory is used by Lucene for caching in-memory data structures. As a standard practice never set the heap size greater than 30 GB, as setting a higher value would not use JAVA compressed pointers, wastes memory, reduces CPU performance, and makes the garbage collection (GC) struggle with large heaps.

  For example, if the available memory is 20 GB, set both Xms and Xmx to 10 GB:

  ```
  -Xms10g
  -Xmx10g
  ```

  After you modify the jvm.options file, start and stop Elasticsearch.

  See "Performing Additional Tasks," Starting and Stopping an Elasticsearch Service.

- Swapping

  Disable swapping. Swapping is expensive in terms of memory required, and thus affects performance.

You can disable swapping on Linux temporarily by running: `sudo swapoff -a`. To disable it permanently, you will need to edit the /etc/fstab file and comment out any lines that contain the word "swap."

You can disable swapping on Microsoft Window by disabling the paging file entirely. For example, select System Properties, Advanced. Click the Settings button in the Performance area. Select Advanced, Virtual memory, and change the value for the paging file. Alternatively, you can set the sysctl value vm.swappiness to 1.

If disabling swapping completely is not an option, you can decrease the "swappiness" value. This value controls how aggressively the operating system (OS) tries to swap memory. This prevents swapping under normal circumstances, but still allows the OS to swap under emergency memory situations. A swappiness of 1 is better than 0, since on some kernel versions a swappiness of 0 can invoke the out-of-memory (OOM) killer. If neither approach is possible, you should enable `mlockall file`. This allows the JVM to lock its memory and prevent it from being swapped by the OS. The recommendation is to set this parameter to true. To enable this parameter, set this value in the elasticsearch.yml configuration file:

```
bootstrap.mlockall: true
```

See "Performing Additional Tasks," Modifying the Elasticsearch Configuration File (Optional).

- Type of disk drive

  Solid-state drives (SSDs) are by far superior to any spinning media. SSD-backed nodes see boosts in both query and indexing performance. If you choose to use SSDs, the I/O scheduling should be set to Deadline/noop for optimal utilization of SSDs and increased performance.

  If you use spinning media, try to obtain the fastest disks possible (high-performance server disks, 15k RPM drives). Using RAID 0 is an effective way to increase disk speed, for both spinning disks and SSDs. There is no need to use mirroring or parity variants of RAID, since high availability is built into Elasticsearch via replicas. Avoid network-attached storage (NAS). NAS is often slower, displays larger latencies with a wider deviation in average latency, and is a single point of failure.

- File Descriptors and MMAP

  Lucene uses a very large number of files. You should increase your file descriptor count to something very large, such as 64,000. Elasticsearch uses a mix of new IO File system (NioFS) and memory-mapped file system (MMapFS) for the various files. Ensure that you configure the maximum map count so that there is ample virtual memory available for mmapped files. This can be set by modifying vm.max_map_count in /etc/sysctl.conf; for example, `vm.max_map_count =262144`.

Chapter 2

Deploying the Elasticsearch Deployment Package

This chapter discusses:

- Understanding the Elasticsearch and Kibana Installation
- Obtaining the Elasticsearch DPK
- Installing Elasticsearch and Kibana on Linux
- Installing Elasticsearch and Kibana on Microsoft Windows

Understanding the Elasticsearch and Kibana Installation

The Elasticsearch DPK zip file includes a setup script, psft-dpk-setup.bat for the a Microsoft Windows DPK and a psft-dpk-setup.sh for the Linux DPK. Include the following decisions in preparing to install with the Elasticsearch DPK setup script:

- Run the script interactively, or run it in silent mode with a prepared configuration file.
- Install Elasticsearch and Kibana together or separately.

  You can install Kibana as part of the Elasticsearch DPK deployment. The installation script prompts you to specify the Elasticsearch server that you want to connect to Kibana. You can connect to the server that you are in the process of installing, or enter the credentials for a different Elasticsearch server.

  You can also use the setup script to install Kibana only, but in this case you must have an existing Elasticsearch server to connect to.

  For example, you could install Elasticsearch first on server1, and then at a later date install Kibana on server2, and specify the Elasticsearch installation on server1.

- Specify Elasticsearch clusters and nodes.

  The Elasticsearch DPK deployment creates an Elasticsearch cluster with one or more Elasticsearch nodes. Briefly, an Elasticsearch node refers to the server where Elasticsearch is installed, and the cluster is comprised of nodes which all have the same cluster name. The DPK setup script asks for the cluster name and the names of existing nodes. If you provide an existing cluster name and list of existing nodes, the existing nodes will join the cluster. For information on adding additional nodes after installation, see the section Adding Additional Elasticsearch Nodes.

  See "Performing Additional Tasks," Adding Additional Elasticsearch Nodes.

When using the Elasticsearch DPK setup script:

- The Elasticsearch DPK setup script does not create the user "esuser," which was associated with the Puppet-based installation.

  The user who installs the Elasticsearch DPK owns all the Elasticsearch files, and only that user will be able to start the process.

- The installation does not require root access (on Linux) or administrator access (on Microsoft Windows).

For information on using Kibana, and on the Elasticsearch concepts mentioned in this task, such as clusters and nodes, see the PeopleTools product documentation and the Elasticsearch online help.

See *PeopleTools: Search Technology.*

See Elasticsearch online help, https://www.elastic.co/guide/en/elasticsearch/guide/current/index.html.

Task 2-1: Obtaining the Elasticsearch DPK

This section discusses:

- Obtaining the Elasticsearch DPK from Oracle Software Delivery Cloud
- Obtaining the Elasticsearch DPK from My Oracle Support

Task 2-1-1: Obtaining the Elasticsearch DPK from Oracle Software Delivery Cloud

To obtain the Elasticsearch DPK from Oracle Software Delivery Cloud:

1. Sign in to Oracle Software Delivery Cloud.

   See Oracle Software Delivery Cloud, https://edelivery.oracle.com.

2. Search for the current PeopleSoft PeopleTools release.

3. In the search results, locate the PeopleSoft PeopleTools download package, and click to add it to your cart.

4. Click Selected Software.

5. Locate Elasticsearch in the search results and select Microsoft Windows or Linux from the Platforms/Languages drop-down list.

6. Download the Elasticsearch file, for example V123456-01.zip.

   Be sure that the directory where you download the zip file has adequate available space. The directory should be a newly created directory with no other files present.

   This documentation refers to the downloaded zip file as *ES_FILENAME.zip,* and the directory where you download the zip file as *ES_INSTALL.*

Task 2-1-2: Obtaining the Elasticsearch DPK from My Oracle Support

To obtain the Elasticsearch DPK from My Oracle Support:

1. Sign in to My Oracle Support.

   See My Oracle Support, https://support.oracle.com.

2. Select the Patches & Updates tab.

3. Select Product or Family (Advanced), and search for PeopleSoft PeopleTools.

4. Select the current release from the Release drop-down list, and then click Search.

5. In the list of results, locate the Elasticsearch file for your operating system, ESK-DPK-*<Operating_System>-<Release.Patch>_<DPK_version>*.zip, where the file name includes the following:

   - *<Operating_System>* is LNX for Oracle Linux, or WIN for Microsoft Windows.
   - *<Release.Patch>* is the release and patch number for the product, such as 6.1.2.

- *<DPK_revision>* is a number specific to the DPK revision.

For example, ESK-DPK-WIN-6.1.2_03.zip or ESK-DPK-LNX-6.1.2_03.zip.

6. Download the file.

   Be sure that the directory where you download the zip file has adequate available space. The directory should be a newly created directory with no other files present.

   This documentation refers to the downloaded zip file as *ES_FILENAME.zip,* and the directory where you download the zip file as *ES_INSTALL.*

Task 2-2: Installing Elasticsearch and Kibana on Linux

This section discusses:

- Installing Elasticsearch and Kibana on Linux Interactively
- Installing Elasticsearch and Kibana on Linux in Silent Mode
- Verifying the Elasticsearch DPK Installation on Linux
- Removing the Elasticsearch Installation from Linux
- Removing the Kibana Installation from Linux
- Performing Post-Installation Steps on Linux

Task 2-2-1: Installing Elasticsearch and Kibana on Linux Interactively

Use this procedure on physical or virtual Linux hosts. This procedure assumes:

- You have downloaded the required Elasticsearch DPK for Linux, referred to as *ES_FILENAME.zip,* and saved it in a newly created directory accessible to the Linux host, referred to as *ES_INSTALL.*
- There is enough space on the Linux host for the Elasticsearch installation and your estimated indexing requirements.

Make a note of the values you supply for ports, passwords, and so on. When you configure the Elasticsearch instance for PeopleSoft, the values must match those specified here.

1. Open a terminal window.

2. Change directory to *ES_INSTALL.*

   ```
   cd ES_INSTALL
   ```

3. Extract the entire contents of *ES_FILENAME.zip* in the same directory, *ES_INSTALL.*

   ```
   unzip ES_FILENAME.zip
   ```

   Note. It is a good idea to extract into the same directory where you downloaded the zip files, and to extract into an empty directory for each new installation.

   The extraction creates the following directories and files in *ES_INSTALL*:

   - setup directory — includes the setup scripts and a silent installation sample
   - archives directory — includes archives for deployment
   - readme.txt file
   - elasticsearch–manifest — versions of Elasticsearch and JRE

4. Run the DPK setup script from *ES_INSTALL/*setup as follows:

    ```
    ./psft-dpk-setup.sh --install --install_base_dir BASE_DIR
    ```

    *   For the `install_base_dir` option, specify the full path where you want Elasticsearch installed. The installation directory is referred to in this documentation as *BASE_DIR*. For example:

        ```
        ./psft-dpk-setup.sh --install --install_base_dir /home/esk612
        ```

    *   Use double-dashes when specifying the script options; for example, `--install`.

5. Answer *y* (yes) to install Elasticsearch, or *n* (no) to exit.

    ```
    You've chosen to do a fresh installation of Elasticsearch and Kibana.
    Do you want to install Elasticsearch? (y/n): y
    ```

6. Enter the password two times for the Elasticsearch administrative user esadmin, at the following prompt.

    The esadmin user is used to authenticate requests on Elasticsearch.

    Note. The script does not display the password or any masking characters as you type.

    Note. The esadmin user is not the same as the user who installs the Elasticsearch DPK and owns the files.

    ```
    Enter the password for esadmin.
    Re-enter the password for esadmin:
    ```

7. Enter the password for the Elasticsearch proxy user, people.

    Note that this is not the same user as the PeopleSoft connect ID, which also has people as the default value.

    ```
    Enter the password for people.
    Re-enter the password for people:
    ```

8. Enter the name for the Elasticsearch cluster, or accept the default name, ESCLUSTER.

    ```
    Enter the ES cluster name [ ESCLUSTER ] :
    ```

9. Enter the Elasticsearch HTTP port.

    This is the port on which Elasticsearch listens for requests. The default is 9200.

    ```
    Enter the HTTP port for Elasticsearch [ 9200 ] :
    ```

10. Enter the host name for any nodes that are already members of a cluster.

    ```
    Enter the list of discovery hosts [ ["127.0.0.1", "[::1]"] ] :
    ```

    Be sure to use the following syntax:

    *   Enclose one or more host names in square brackets.
    *   Enclose the host name or IP address in *double quotes* (" ").
    *   Use commas to list two or more hosts.
    *   Use this as an example for one host: `["host1.example.com"]`
    *   Use this as an example for more than one host: `["host1.example.com", "127.0.0.1"]`

11. Enter the minimum number of nodes.

    ```
    Enter the mininum number of master nodes :
    ```

12. Enter the location for the Elasticsearch data.

Oracle recommends that you do not use the default location, *BASE_DIR/*pt/elasticsearch6.1.2/data, with PeopleSoft environments. Instead, specify the full path for a data directory that is outside of *BASE_DIR/*pt/elasticsearch6.1.2.

```
Enter the path where you want the Elasticsearch data to reside [ C:⇒
\esk612\pt\elasticsearch6.1.2\data ] :
```

13. Enter the location for the Elasticsearch logs.

    The default location is *BASE_DIR/*pt/elasticsearch6.1.2/logs.

```
Enter the path where you want the Elasticsearch Logs to be written to [⇒
 /home/esk612/pt/elasticsearch6.1.2/logs ] :
```

14. Enter the heap size as a number with the letter "G" as shown in this example:.

    See Prerequisites

```
Enter the Java Heap size for Elasticsearch in GB [ 2 ] : 7
```

15. Review the status messages as the script sets up the PeopleSoft environment.

```
Extracting the new ES Binary .... ... [OK]
Extracting the new JRE ............ . [OK]
Setting users/roles in ES .......    [OK]
Configuring Elasticsearch ......    [OK]
Starting Elasticsearch server ...  . [OK]

Elasticsearch Installation Completed.
```

16. Answer *y* (yes) to install Kibana, or *n* (no) to exit.

```
Do you want to install Kibana: (y/n): y
```

    The script displays information about the current server.

```
Elasticsearch Host: server1.example.com
Elasticsearch Port: 9200
```

17. Answer *y* (yes) to set up Kibana to connect to the Elasticsearch server you are currently installing, or *n* (no) to enter information about a different Elasticsearch server.

```
Do you want to use the same Elasticsearch(as above) for kibana: (y/n):
```

18. Enter the Kibana port, or accept the default, 5601.

```
Enter the server port for Kibana [ 5601 ] :
```

19. If you answered no to the previous prompt, enter the host name and port for the Elasticsearch server to connect to.

```
Enter the Elasticsearch host[http(s)://hostname]:

Enter the Elasticsearch port [ 9200 ] :
```

20. Wait until the installation is complete.

```
Elasticsearch service is running. Proceeding with the Kibana⇒
 installation.
Extracting the new Kibana Binary .......[OK]
Configuring Kibana ......
```

```
Kibana Keystore updated. [OK]
Kibana installation is completed.
```

A successful deployment includes the following:

- The Elasticsearch home directory is found in *BASE_DIR/*pt/elasticsearch6.1.2.
- The ES_HOME environment variable is set to *BASE_DIR/*pt/elasticsearch6.1.2 for the current terminal session.
- The Kibana home directory is found in *BASE_DIR/*pt/Kibana6.1.2.
- Java is installed to *BASE_DIR/*pt/es_jre1.8.0_*yy,* where yy is the JRE version.
- The JAVA_HOME environment variable is set to *BASE_DIR/*pt/es_jre1.8.0_*yy* for the current terminal session.

  See Performing Post-Installation Steps on Linux.

Task 2-2-2: Installing Elasticsearch and Kibana on Linux in Silent Mode

This section discusses:

- Extracting the DPK and Encrypting the Passwords on Linux
- Creating a Configuration File and Running the Silent Mode Installation on Linux

Extracting the DPK and Encrypting the Passwords on Linux

You can use the setup script to install Elasticsearch and Kibana in silent mode by preparing a text file that includes installation details.

The passwords that you include in the silent mode text file must be encrypted using the PSCipher utility. You must use the pscipher.jar and psvault files that are part of the Elasticsearch DPK. The psvault that is delivered with PeopleSoft PeopleTools will not work for Elasticsearch.

1. Open a terminal window.

2. Change directory to *ES_INSTALL.*

   ```
   cd ES_INSTALL
   ```

3. Extract the entire contents of *ES_FILENAME.zip* in the same directory, *ES_INSTALL.*

   ```
   unzip ES_FILENAME.zip
   ```

   Note. It is a good idea to extract into the same directory where you downloaded the zip files, and to extract into an empty directory for each new installation.

   The extraction creates the following directories in *ES_INSTALL*:

- setup directory — includes the setup script and a sample configuration file.
- archives directory — includes archives for deployment
- readme.txt file
- elasticsearch-manifest — lists the version information for Elasticsearch and JRE included in the DPK

4. Go to *ES_INSTALL*archives and extract pt-elasticsearch-6.1.2.tgz into a folder with the same name.

5. Copy pscipher.jar from pt-elasticsearch-6.1.2/plugins/orcl-security-plugin to a convenient location, such as /var/estemp.

6. Copy psvault from pt-elasticsearch-6.1.2/plugins/orcl-security-plugin/config/properties to /var/estemp.

7. In a command prompt, change directory to a temporary directory, such as /var/estemp.

8. Run the following command, specifying the passwords for esadmin and people.

    Copy the encrypted text without adding line feeds or spaces and paste them into the silentinstall.config file in the next section. The encrypted passwords are saved in es_output.txt.

    ```
    java -Dpscipher.path=<path_to_which_files_are_extracted> -cp <path_to_⇒
    which_files_are_extracted>/pscipher.jar com.peoplesoft.pt.elasticsearch⇒
    .pscipher.PSESEncrypt esadmin <password> people <password> <output_⇒
    path>/es_output.txt
    ```

Creating a Configuration File and Running the Silent Mode Installation on Linux

To install in silent mode:

1. Create a configuration file.

    The *ES_INSTALL/*setup directory include a sample configuration file, silentinstall.config. If you want to use this file, it is a good idea to make a backup copy of the original file before you continue.

    Edit the configuration file with the required environment parameters. Here is a sample:

    ```
    #Silent Install Configuration Values for installing Elasticsearch and⇒
     Kibana

    #Install Elasticsearch Y/N is mandatory field
    Install elasticsearch?[Y/N]= n

    #Elasticsearch cluster name (mandatory)
    cluster.name= ESCLUSTER

    #Elasticsearch host name (mandatory)
    network.host=

    #Elasticsearch port number (optional-default 9200)
    http.port= 9200

    #Elasticsearch data directory (optional-default ES_HOME/data)
    path.data=

    #Elasticsearch log directory (optional-default ES_HOME/logs)
    path.logs=

    #Elasticsearch discovery hosts (optional)
    discovery.hosts= ["127.0.0.1", "[::1]"]

    #Elasticsearch minimum master nodes (optional)
    minimum_master_nodes= 1

    #Elasticsearch heap size (optional-default 2)
    ES_HEAP_SIZE= 2

    #Elasticsearch encrypted password for esadmin (mandatory)
    ```

```
esadmin.password=

#Elasticsearch encrypted password for people (mandatory)
people.password=

#Install Kibana Y/N is mandatory field
Install kibana?[Y/N]= y

#Kibana host name (mandatory)
kibana.host=

#Kibana port number (optional-default 5601)
kibana.port= 5601

#Set the value to Y if Install Elasticsearch(Y) and Install Kibana(Y)
Use same ES?[Y/N]= n

#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=
#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=
```

2. Enter the encrypted passwords that you prepared for esadmin.password, people.password and ES.password.

3. If you want to install Kibana, specify values for the following:

```
Install Kibana Y/N is mandatory field
Install kibana?[Y/N]= y

#Kibana host name (mandatory)
kibana.host=

#Kibana port number (optional-default 5601)
kibana.port= 5601
```

4. If you want to install Kibana and connect to a previously-installed Elasticsearch server, specify the following:

```
#Set the value to Y if Install Elasticsearch(Y) and Install Kibana(Y)
Use same ES?[Y/N]= n

#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=
#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=
```

5. Run the DPK setup script from *ES_INSTALL/*setup as follows:

```
./psft-dpk-setup.sh --install_silent --install_base_dir BASE_DIR ⇒
--config_path full_path_configuration_file
```

Note. The command includes a line feed for readability. Do not include the line feed when you run.

- Use double-dashes when specifying the script options; for example, `--install_silent`..
- For the `install_base_dir` option, specify the full path where you want Elasticsearch installed. The installation directory is referred to in this documentation as *BASE_DIR*.
- For the config_path option, specify the full path to the prepared configuration file. For example:

Note. The command includes a line feed for readability. Do not include the line feed when you run.

```
./psft-dpk-setup.sh --install_silent --install_base_dir /home/esk612 ⇒
--config_path /home/temp_install/silentinstall.config
```

Task 2-2-3: Verifying the Elasticsearch DPK Installation on Linux

To verify the Elasticsearch installation, in a terminal window, ensure that the following command gives an output:
```
ps —ef | grep elas
```

After verifying the process, use one of the following methods to verify the Elasticsearch installation.

- Run a REST call.

  This example uses the CURL utility for the REST call:
  ```
  curl --user <username>:<password> -XGET http://<host>:<port>
  ```

  For details about the correct usage for the CURL utility, see your operating system documentation.
- Open a browser and enter the following URL:

  http://<username>:<password>@<host>:<port>/

  For both the REST command and the browser URL, use these definitions:

  - username ― the Elasticsearch user administrator, esadmin

  - password ― the password you entered during the DPK setup script installation

  - host ― the Elasticsearch host name

  - port ― Elasticsearch (REST) port that you entered during the DPK setup script installation

In both cases, you should see a message similar to the following:

```
{
    "name" : "abc.abc.com",
    "cluster_name" : "ESCLUSTER",
    "cluster_uuid" : "2Lnh...",
    "version" : {
       "number" : "6.1.2",
       "build_hash" : "b9e4a",
       "build_timestamp" : "2018-04-21T16:03:47Z",
       "build_snapshot" : false,
       "lucene_version" : "7.1.0"
       "minimum_wire_compatibility_version" : "5.6.0"
       "minimum_index_compatibility_version" : "5.0.0"
    },
    "tagline" : "You Know, for Search"
}
```

After you complete the Elasticsearch installation, you must configure the integration with the PeopleSoft environment.

See "Integrating Elasticsearch with the PeopleSoft Environment."

To learn how to start and use Kibana, see the Search Technology product documentation.

See *PeopleTools: Search Technology,* "Setting Up Kibana in PeopleSoft Search Framework."

Task 2-2-4: Removing the Elasticsearch Installation from Linux

Use these steps to remove the Elasticsearch DPK installation from a Linux host:

Note. You must use the manual steps. There is no cleanup option for the psft-dpk-setup.sh script.

1. Use the following command to determine the Elasticsearch process ID (pid):

   ```
   ps -ef |grep elastic
   ```

2. Stop the process, substituting the Elasticsearch process ID for <pid>:

   ```
   stop <pid>
   ```

3. Remove the Elasticsearch installation directory.

Task 2-2-5: Removing the Kibana Installation from Linux

If you installed Kibana, use these steps to remove the Kibana installation from a Linux host:

Note. You must use the manual steps. There is no cleanup option for the psft-dpk-setup.sh script.

1. Use the following command to determine the process ID (pid) for the running Kibana service:

   ```
   ps -ef |grep node
   ```

2. Stop the process, substituting the Kibana process ID for <pid>:

   ```
   stop <pid>
   ```

3. Remove the Kibana installation directory.

Task 2-2-6: Performing Post-Installation Steps on Linux

After you complete the Elasticsearch DPK installation, if you want to perform operations described later in this chapter, such as starting and stopping Elasticsearch, or running the elasticsearchuser script, perform the manual steps described in this section.

- Ensure that the heap size is set to a value equal to or less than 50% of available memory, and not exceeding 30G.

  See "Preparing to Deploy," Reviewing Elasticsearch Recommendations.

- In a terminal for the current session, set the JAVA_HOME environment variable to the location installed by the DPK setup script; for example:

  ```
  export JAVA_HOME= BASE_DIR/pt/es_jre1.8.0_yy
  ```

Task 2-3: Installing Elasticsearch and Kibana on Microsoft Windows

This section discusses:

- Installing Elasticsearch and Kibana on Microsoft Windows Interactively
- Installing Elasticsearch and Kibana on Microsoft Windows in Silent Mode
- Verifying the Elasticsearch Installation on Microsoft Windows
- Removing the Elasticsearch Installation from Microsoft Windows
- Removing the Kibana Installation from Microsoft Windows
- Performing Post-Installation Steps on Microsoft Windows

Task 2-3-1: Installing Elasticsearch and Kibana on Microsoft Windows Interactively

Use this procedure on physical or virtual Microsoft Windows hosts. This procedure assumes that:

- You have downloaded the required Elasticsearch DPK for Microsoft Windows, *ES_FILENAME.zip,* and saved it in a newly created directory accessible to the Microsoft Windows host, referred to as *ES_INSTALL.*
- There is enough space on the host for the Elasticsearch installation and your estimated indexing requirements.

Make a note of the values you supply for ports, passwords, and so on. When you configure the Elasticsearch instance for PeopleSoft, the values must match those specified here.

1. Go to *ES_INSTALL.*
2. Extract the entire contents of *ES_FILENAME.zip* in the same directory, *ES_INSTALL.*

   Note. It is a good idea to extract into the same directory where you downloaded the zip files, and to extract into an empty directory for each new installation.

   The extraction creates the following directories and files in *ES_INSTALL*:

   - setup directory — includes the setup script and a silent installation sample
   - archives directory — includes archives for deployment
   - readme.txt file
   - elasticsearch–manifest — versions of Elasticsearch and JRE

3. Open a command prompt.
4. Change directory to *ES_INSTALL/*setup.
5. Run the DPK setup script with these options:

   ```
   psft-dpk-setup.bat --install --install_base_dir BASE_DIR
   ```

   - For the `install_base_dir` option, specify the full path where you want Elasticsearch installed. The installation directory is referred to in this documentation as *BASE_DIR*; for example:

     ```
     psft-dpk-setup.bat --install --install_base_dir C:\esk612
     ```

   - Use double-dashes when specifying the script options; for example, `--install`.

6. Answer *y* (yes) to install Elasticsearch, or *n* (no) to exit.

```
You've chosen to do a fresh installation of Elasticsearch and Kibana.
Do you want to install Elasticsearch? (y/n): y
```

7. Enter the password two times for the Elasticsearch administrative user esadmin, at the following prompt.

The esadmin user is used to authenticate requests on Elasticsearch.

Note. The script does not display the password or any masking characters as you type.

Note. The esadmin user is not the same as the user who installs the Elasticsearch DPK and owns the files.

```
Enter the password for esadmin.
Re-enter the password for esadmin:
```

8. Enter the password for the Elasticsearch proxy user, people.

Note that this is not the same user as the PeopleSoft connect ID, which also has people as the default value.

```
Enter the password for people.
Re-enter the password for people:
```

9. Enter the name for the Elasticsearch cluster, or accept the default name, ESCLUSTER.

```
Enter the ES cluster name [ ESCLUSTER ] :
```

10. Enter the Elasticsearch HTTP port.

This is the port on which Elasticsearch listens for requests. The default is 9200.

```
Enter the HTTP port for Elasticsearch [ 9200 ] :
```

11. Enter the host name for any nodes that are already members of a cluster.

```
Enter the list of discovery hosts [ ["127.0.0.1", "[::1]"] ] :
```

Be sure to use the following syntax:

- Enclose one or more host names in square brackets.
- Enclose the host name or IP address in *double quotes* (" ").
- Use commas to list two or more hosts.
- Use this as an example for one host: `["host1.example.com"]`
- Use this as an example for more than one host: `["host1.example.com", "127.0.0.1"]`

12. Enter the minimum number of nodes.

```
Enter the mininum number of master nodes :
```

13. Enter the location for the Elasticsearch data.

Oracle recommends that you do not use the default location, *BASE_DIR/*pt/elasticsearch6.1.2/data, with PeopleSoft environments. Instead, specify the full path for a data directory that is outside of *BASE_DIR*pt\elasticsearch6.1.2.

```
Enter the path where you want the Elasticsearch data to reside [ C:⇒
\esk612\pt\elasticsearch6.1.2\data ] :
```

14. Enter the location for the Elasticsearch logs.

The default location is *BASE_DIR*\pt\elasticsearch6.1.2\logs.

```
Enter the path where you want the Elasticsearch Logs to be written to [⇒
 C:\esk612\pt\elasticsearch6.1.2\logs ] :
```

15. Enter the heap size as a number with the letter "G" as shown in this example:.

    See Prerequisites

    ```
    Enter the Java Heap size for Elasticsearch in GB [ 2 ] : 7
    ```

16. Review the status messages as the script sets up the PeopleSoft environment.

    ```
    Extracting the new ES Binary .... ... [OK]
    Extracting the new JRE ............ . [OK]
    Setting users/roles in ES .......    [OK]
    Configuring Elasticsearch ......     [OK]
    Starting Elasticsearch server ...  . [OK]

    Elasticsearch Installation Completed.
    ```

17. Answer *y* (yes) to install Kibana, or *n* (no) to exit.

    ```
    Do you want to install Kibana: (y/n): y
    ```

    The script displays information about the current server.

    ```
    Elasticsearch Host: server1.example.com
    Elasticsearch Port: 9200
    ```

18. Answer *y* (yes) to set up Kibana to connect to the Elasticsearch server you are currently installing, or *n* (no) to enter information about a different Elasticsearch server.

    ```
    Do you want to use the same Elasticsearch(as above) for kibana: (y/n):
    ```

19. Enter the Kibana port, or accept the default, 5601.

    ```
    Enter the server port for Kibana [ 5601 ] :
    ```

20. If you answered no to the previous prompt, enter the host name and port for the Elasticsearch server to connect to.

    ```
    Enter the Elasticsearch host[http(s)://hostname]:

    Enter the Elasticsearch port [ 9200 ] :
    ```

21. Wait until the installation is complete.

    ```
    Elasticsearch service is running. Proceeding with the Kibana⇒
     installation.
    Extracting the new Kibana Binary .......[OK]
    Configuring Kibana ......
    Kibana Keystore updated. [OK]
    Kibana installation is completed.
    ```

A successful deployment includes the following:

- The Elasticsearch home directory is found in *BASE_DIR*\pt\elasticsearch6.1.2.

- The Elasticsearch data and logs directories are installed to the locations you specified.

- The Elasticsearch service is installed and running.

- The ES_HOME environment variable is set to *BASE_DIR*\pt\elasticsearch6.1.2.

The environment variable persists until you close the command prompt window.

- The Kibana home directory is found in *BASE_DIR*\pt\Kibana6.1.2

- Java is installed to *BASE_DIR*\pt\es_jre1.8.0_*yy,* where yy is the JRE version.

- The JAVA_HOME environment variable is set to *BASE_DIR*\pt\es_jre1.8.0_*yy.*

  This environment variable is set at the system level.

- The ES_HEAP_SIZE environment variable is set to the value you enter when running the setup script.

  This environment variable is set at the system level.

Task 2-3-2: Installing Elasticsearch and Kibana on Microsoft Windows in Silent Mode

This section discusses:

- Extracting the DPK and Encrypting the Passwords on Microsoft Windows

- Creating a Configuration File and Running the Silent Mode Installation on Microsoft Windows

Extracting the DPK and Encrypting the Passwords on Microsoft Windows

You can use the setup script to install Elasticsearch and Kibana in silent mode by preparing a text file that includes installation details.

The passwords that you include in the silent mode text file must be encrypted using the PSCipher utility. You must use the pscipher.jar and psvault files that are part of the Elasticsearch DPK. The psvault that is delivered with PeopleSoft PeopleTools will not work for Elasticsearch.

1. Go to *ES_INSTALL.*

2. Extract the entire contents of *ES_FILENAME.zip* in the same directory, *ES_INSTALL.*

   ```
   unzip ES_FILENAME.zip
   ```

   Note. It is a good idea to extract into the same directory where you downloaded the zip files, and to extract into an empty directory for each new installation.

   The extraction creates the following directories in *ES_INSTALL*:

   - setup directory — includes the setup script and sample configuration file

   - archives directory — includes archives for deployment

   - readme.txt file and other files

   - elasticsearch-manifest — lists the version information for Elasticsearch and JRE included in the DPK

3. Go to *ES_INSTALL*\archives and extract pt-elasticsearch-6.1.2.tgz into a folder with the same name.

   It may be necessary to extract the file twice.

4. Copy pscipher.jar from pt-elasticsearch-6.1.2/plugins/orcl-security-plugin to a convenient location, such as C:\estemp.

5. Copy psvault from pt-elasticsearch-6.1.2/plugins/orcl-security-plugin/config/properties to C:\estemp.

6. In a command prompt, change directory to C:\estemp.

7. Run the following command, specifying the passwords for esadmin and people.

Copy the encrypted text without adding line feeds or spaces and paste them into the silentinstall.config file in the next section. The encrypted passwords are saved in es_output.txt.

```
java -Dpscipher.path=<path_to_which_files_are_extracted> -cp <path_to_⇒
which_files_are_extracted>/pscipher.jar com.peoplesoft.pt.elasticsearch⇒
.pscipher.PSESEncrypt esadmin <password> people <password> <output_⇒
path>/es_output.txt
```

Creating a Configuration File and Running the Silent Mode Installation on Microsoft Windows

To install in silent mode:

1. Create a configuration file.

   The *ES_INSTALL/*setup directory includes a sample configuration file, silentinstall.config. If you want to use this file, it is a good idea to make a backup copy of the original file before you continue.

   Edit the configuration file with the required environment parameters. Here is a sample:

```
#Silent Install Configuration Values for installing Elasticsearch and⇒
 Kibana

#Install Elasticsearch Y/N is mandatory field
Install elasticsearch?[Y/N]= n

#Elasticsearch cluster name (mandatory)
cluster.name= ESCLUSTER

#Elasticsearch host name (mandatory)
network.host=

#Elasticsearch port number (optional-default 9200)
http.port= 9200

#Elasticsearch data directory (optional-default ES_HOME/data)
path.data=

#Elasticsearch log directory (optional-default ES_HOME/logs)
path.logs=

#Elasticsearch discovery hosts (optional)
discovery.hosts= ["127.0.0.1", "[::1]"]

#Elasticsearch minimum master nodes (optional)
minimum_master_nodes= 1

#Elasticsearch heap size (optional-default 2)
ES_HEAP_SIZE= 2

#Elasticsearch encrypted password for esadmin (mandatory)
esadmin.password=

#Elasticsearch encrypted password for people (mandatory)
```

```
people.password=

#Install Kibana Y/N is mandatory field
Install kibana?[Y/N]= y

#Kibana host name (mandatory)
kibana.host=

#Kibana port number (optional-default 5601)
kibana.port= 5601

#Set the value to Y if Install Elasticsearch(Y) and Install Kibana(Y)
Use same ES?[Y/N]= n

#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=
#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=
```

2. If you want to install Kibana, specify values for the following:

```
Install Kibana Y/N is mandatory field
Install kibana?[Y/N]= y

#Kibana host name (mandatory)
kibana.host=

#Kibana port number (optional-default 5601)
kibana.port= 5601
```

3. Enter the encrypted passwords that you prepared for esadmin.password, people.password and ES.password.

4. If you want to install Kibana and connect to a previously-installed Elasticsearch server, specify the following:

```
#Set the value to Y if Install Elasticsearch(Y) and Install Kibana(Y)
Use same ES?[Y/N]= n

#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=
#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=
```

5. In a command prompt window, run the DPK setup script from *ES_INSTALL/*setup as follows:

Note. The command includes line feeds for readability. Do not include the line feed when you run.

```
psft-dpk-setup.bat --install_silent --install_base_dir BASE_DIR ⇒
--config_path full_path_configuration_file
```

- Use double-dashes when specifying the script options; for example, --install_silent.
- For the install_base_dir option, specify the full path where you want Elasticsearch installed. The

installation directory is referred to in this documentation as *BASE_DIR*.

- For the config_path option, specify the full path to the prepared configuration file. For example:

Note. The command includes line feeds for readability. Do not include the line feed when you run.

```
psft-dpk-setup.bat --install_silent --install_base_dir C:/esk612 ⇒
--config_path C:/tmp_install/silentinstall.config
```

6. When the script completes, you see a message such as:

```
Elasticsearch Installation Completed.
Silent mode installation of Kibana
```

Task 2-3-3: Verifying the Elasticsearch Installation on Microsoft Windows

The Elasticsearch deployment sets up a Windows service. To verify the installation:

1. Launch the Services dialog box, for example by opening Task Manager and selecting the Services tab.

2. Verify that the "elasticsearch-service-x64" service is present and has status "Running," as in this example:
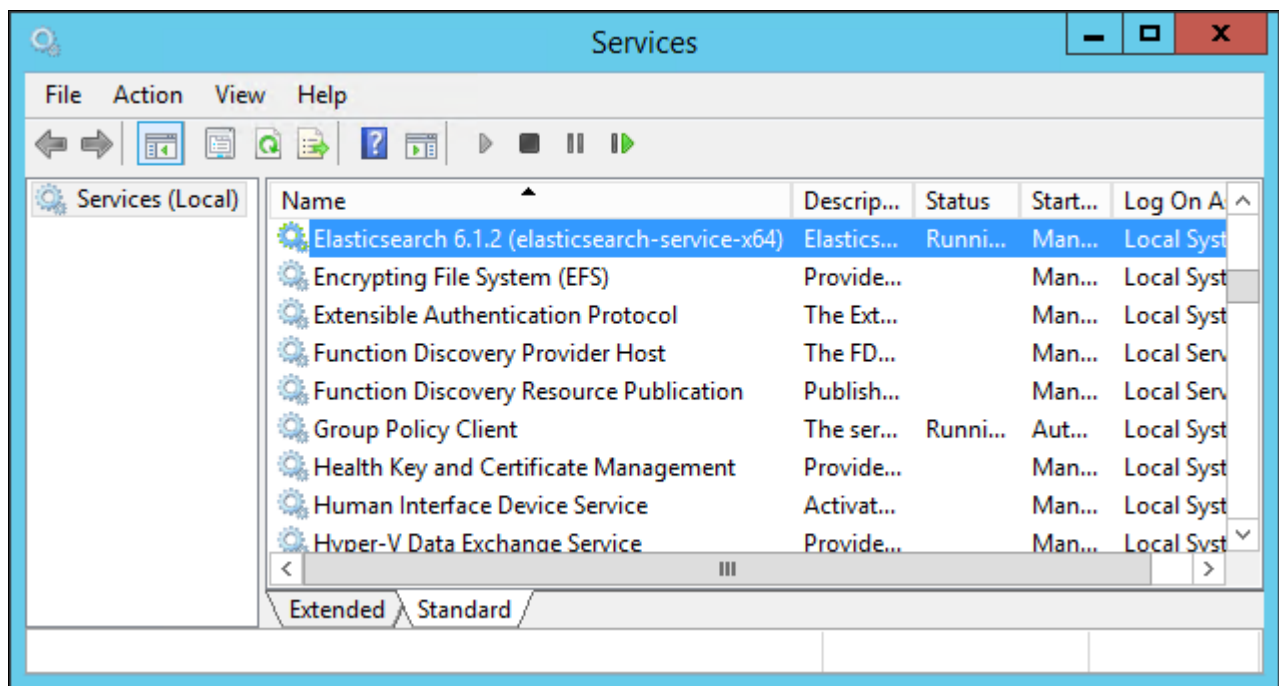
Services dialog box with the Elasticsearch service

3. If the Elasticsearch service is not listed, you can start it manually:

   a. Open a command prompt.

   b. Go to *ES_HOME*\bin and enter the following command:
   ```
   elasticsearch-service.bat install
   ```

   c. Start the service with the following command:
   ```
   elasticsearch-service.bat start
   ```

After verifying the service, use one of the following methods to verify the Elasticsearch installation.

- Run a REST call.

  This example uses the CURL utility to run the REST call:

  ```
  curl --user <username>:<password> -XGET http://<host>:<port>
  ```

  For details about the correct usage of the CURL utility, see your operating system documentation.

- Open a browser and enter the following URL:

  http://<username>:<password>@<host>:<port>/

  For both the REST command and the browser URL, use these definitions:

  - username — the Elasticsearch user administrator, esadmin

  - password — the password you entered during the DPK setup script installation

  - host — the Elasticsearch host name.

  - port — Elasticsearch (REST) port that you entered during the DPK setup script installation

In both cases, you should see a message similar to the following:

```
{
    "name" : "abc.abc.com",
    "cluster_name" : "ESCLUSTER",
    "cluster_uuid" : "2Lnh...",
    "version" : {
        "number" : "6.1.2",
        "build_hash" : "b9e4a",
        "build_timestamp" : "2018-04-21T16:03:47Z",
        "build_snapshot" : false,
        "lucene_version" : "7.1.0"
        "minimum_wire_compatibility_version" : "5.6.0"
        "minimum_index_compatibility_version" : "5.0.0"
    },
    "tagline" : "You Know, for Search"
}
```

After you complete the Elasticsearch installation, you must configure the integration with the PeopleSoft environment.

See "Integrating Elasticsearch with the PeopleSoft Environment."

To learn how to start and use Kibana, see the Search Technology product documentation.

See *PeopleTools: Search Technology,* "Setting Up Kibana in PeopleSoft Search Framework."

Task 2-3-4: Removing the Elasticsearch Installation from Microsoft Windows

Use these steps to remove the Elasticsearch DPK installation from Microsoft Windows.

Note. You must use the manual steps. There is no cleanup option for the psft-dpk-setup.bat script.

1. Open a command prompt.

2. Enter the following commands, substituting the Elasticsearch installation directory, such as *BASE_DIR* pt\elasticsearch6.1.2, for *ES_HOME.*

   ```
   ES_HOME\bin\elasticsearch-service.bat stop
   ```

```
ES_HOME\bin\elasticsearch-service.bat remove
```

3. Remove the Elasticsearch installation directory.

Task 2-3-5: Removing the Kibana Installation from Microsoft Windows

If you installed Kibana, use these steps to remove the Elasticsearch DPK installation from Microsoft Windows.

Note. You must use the manual steps. There is no cleanup option for the psft-dpk-setup.bat script.

1. Open a command prompt.
2. Stop the Kibana service if it is running.

   If the command window that you used to start the Kibana script is open, either terminate the Kibana script by pressing Ctrl+C, or close the command window.
3. Remove the Kibana installation directory.

Task 2-3-6: Performing Post-Installation Steps on Microsoft Windows

After you complete the Elasticsearch DPK installation, if you want to perform operations described later in this chapter, such as starting and stopping Elasticsearch, or running the elasticsearchuser script, perform the manual steps described in this section.

- Ensure that the heap size is set to a value equal to or less than 50% of available memory, and not exceeding 30G.

  See "Preparing to Deploy," Reviewing Elasticsearch Recommendations.

- Verify that the JAVA_HOME environment variable was set to the location installed by the DPK setup script. If not, set it, for example:

```
set JAVA_HOME=BASE_DIR\pt\es_jre1.8.0_yy
```

Chapter 3

Integrating Elasticsearch with the PeopleSoft Environment

This chapter discusses:

- Applying PeopleSoft Application Support Bugs for Elasticsearch
- Setting Up the PeopleSoft Application for Elasticsearch
- Adding and Configuring an Elasticsearch Instance

Task 3-1: Applying PeopleSoft Application Support Bugs for Elasticsearch

In order to integrate your PeopleSoft environment with Elasticsearch, you must identify and apply any necessary bugs delivered with the latest PeopleSoft Update Image (PI) for your PeopleSoft application.

1. Go to PeopleSoft Update Manager (PUM) Home Page, My Oracle Support, Doc ID 1641843.2.

2. Locate and download the current Update Image (PI) for your PeopleSoft application.

   The features and fixes required for Elasticsearch are delivered beginning with the following images:

   - PeopleSoft Human Capital Management (HCM) 9.2.016
   - PeopleSoft Enterprise Learning Management (ELM) 9.2.013
   - PeopleSoft Financials and Supply Chain Management (FSCM) 9.2.015
   - PeopleSoft Customer Relationship Management (CRM) 9.2.011
   - PeopleSoft Interaction Hub (IH) 9.1.03

3. Include the relevant PeopleSoft application support bugs for Elasticsearch in your change package.

   Review the following article to find the minimum PeopleSoft Update Image and necessary bugs for your PeopleSoft application that are required to implement Elasticsearch.

   See PeopleSoft Applications Updates and Fixes for Elasticsearch, My Oracle Support, Doc ID 2181988.1.

Task 3-2: Setting Up the PeopleSoft Application for Elasticsearch

This section discusses:

- Understanding the PeopleSoft Application Setup
- Verifying the Integration Broker Setup
- Verifying PeopleSoft Roles for All Installations

Understanding the PeopleSoft Application Setup

Make sure your PeopleSoft environment meets these requirements:

- The Integration Broker and the integration gateway are up and running.

  When you use the PeopleSoft DPKs to install an environment, the Integration Broker configuration is performed as part of the installation.

  See Setting Up Integration Broker for additional information.

- The PeopleSoft roles required for Elasticsearch are set for both types of installation.

  See Verifying PeopleSoft Roles for All Installations.

After you satisfy these requirements, and complete the subsequent tasks to configure and deploy Elasticsearch, test the connection on the Search Instance Properties page.

See *PeopleTools: Search Technology,* "Working with Search Instances."

Task 3-2-1: Verifying the Integration Broker Setup

Use these instruction if you need to verify that Integration Broker is set up. The Integration Broker configuration includes the following setup tasks:

- Define the integration gateway.

  See *PeopleTools: Integration Broker Administration,* "Administering Integration Gateways."

- Define the integration gateway properties, including the keystore password setup.

  See *PeopleTools: Integration Broker Administration,* "Configuring Security and General Properties"

- Define service operations, including web services target and REST target locations.

  See *PeopleTools: Integration Broker Administration,* "Using the Target Locations Page to Set Target Locations for Services."

- Define nodes, including portal and content URIs.

  See *PeopleTools: Portal Technology,* "Defining Portal Nodes."

You can use the Integration Broker Activity Guide to carry out the necessary configuration. The activity guide provides centralized access to the PeopleSoft Pure Internet Architecture (PIA) pages used to configure PeopleSoft Integration Broker and the Integration Network.

See *PeopleTools: Integration Broker Administration,* "Understanding the Integration Broker Configuration Activity Guide."

You also have the option of configuring Integration Broker using the Automated Configuration Manager (ACM). See the product documentation for information on how to use the delivered plug-ins for Integration Broker configuration.

See *PeopleTools: Automated Configuration Manager.*

Task 3-2-2: Verifying PeopleSoft Roles for All Installations

The user who will set up the Elasticsearch integration must have the Search Administration, Search Developer, Search Server and ACM Administration roles.

If not, the Security Administrator should add the roles as follows:

1. Sign in to the PeopleSoft application in a browser.

2. Select PeopleTools, Security, User Profiles, User Profiles.

3. Select a User ID.

4. On the User Profiles page, select the Roles tab.

5. Verify that the roles are present, or add them if necessary.

    This example shows the Roles list with Search Administrator, Search Developer, and Search Server. The fourth required role, ACM Administration, is not shown here.

User Profiles page: Roles

Task 3-3: Adding and Configuring an Elasticsearch Instance

This section discusses:

•   Using the Automated Configuration Management SEARCH_TEMPLATE

•   Configuring the Search Instance on the Search Instance Properties Page

Task 3-3-1: Using the Automated Configuration Management SEARCH_TEMPLATE

The ACM framework enables you to store environment configuration settings in a template stored in the database or an external template file, which you can reapply when needed. You can use Automated Configuration Management (ACM) to configure the Elasticsearch instance. Using ACM allows you to automate the deployment and indexing.

See *PeopleTools: Automated Configuration Management.*

1. Verify that you have the ACM Administrator role, as mentioned in the previous section.

   See Setting Up the PeopleSoft Application for Elasticsearch.

2. Select PeopleTools, Automated Config Manager, ACM Templates, Template Definition.

3. Search for and open SEARCH_TEMPLATE.

4. On the Configuration Template Definition page, verify that Configure Search Instance is selected, as in this example, and then click the Properties icon.

   Note. Do not select Deploy Search Definition at this point.

Configuration Template Definition page for SEARCH_TEMPLATE

5. Specify the properties for the plug-in on the Configure Search Instance page.

   The values that begin and end with an AT sign ("@") are ACM template variables that can be used across multiple plug-ins in the template. As mentioned, PeopleSoft Integration Broker configuration is required for these tasks. These variables will be taken from Integration Broker and shared. You do not have to manually fill them in at this point.

   See *PeopleTools: Automated Configuration Management.*

   This table describes the properties for the Configure Search Instance plug-in (PTSFConfigureSrchInstance). Use the properties to configure the PeopleSoft system to talk to Elasticsearch and assign roles.

Property	Default Value	Description
env.ptsf_search_instance	PTSF_SEARCH	Search instance name
env.search_nodes	1	The number of nodes that have been set up for the Elasticsearch cluster. Elasticsearch provides High Availability by forming a cluster of multiple nodes. If you have set up a cluster with multiple nodes, specify the number here. The page expands with additional property fields for the multiple nodes. See *PeopleTools: Search Technology* for information on High Availability. See Adding Additional Elasticsearch Nodes for information on adding additional nodes.
env.node1_search_host	@searchhost@	The host name for the node. Enter the server name, including the domain, such as myhost.example.com, or an IP address.
env.node1_search_port	@searchport@	The port on which Elasticsearch listens for requests. Enter the value supplied when installing the Elasticsearch DPK. The default is 9200.

Property	Default Value	Description
env.node1_search_use_ssl	False	Flag to determine whether the configuration is secure or non-secure. Valid values are: • False — indicates non-secure (HTTP) configuration • True — indicates secure (HTTPS) configuration. Set this property to true only when the PeopleSoft environment is configured for SSL and that URL has been used for the Elasticsearch callback. See *PeopleTools: Search Technology,* "Configuring SSL Between PeopleSoft and Elasticsearch."
env.node1_search_admin_user	NA	The administrative user for Elasticsearch, esadmin. You cannot change this user during the installation. After the installation is complete, you can use the elasticsearchuser script to add users and assign them administrative roles. See *PeopleTools: Search Technology,* "Creating User and Assigning Roles in Elasticsearch."
env.node1_search_admin_password	NA	The password supplied for the esadmin user when installing the Elasticsearch DPK.
env.node1_search_read_user	NA	The Elasticsearch proxy user, people. The proxy user cannot be changed during the installation.
env.node1_search_read_password	NA	The password supplied for the proxy user when installing the Elasticsearch DPK.
env.search_call_back_user	@userid@	The user ID for Elasticsearch to access the PeopleSoft system for access control list (ACL) values. This must be a valid PeopleSoft user ID with Search Server role assigned. **Note.** It should not be necessary to change this value.
env.search_call_back_password	NA	Password for the call-back user ID.

Property	Default Value	Description
env.ps_search_administrator_user	@userid@	Search administrator user name See Setting Up the PeopleSoft Application for Elasticsearch.
env.ps_search_developer_user	@userid@	Search developer user name See Setting Up the PeopleSoft Application for Elasticsearch.
env.gateway_host	@host@.@domain@	Host where the Integration Broker gateway is installed.
env.gateway_port	@httpport@	Port number for the Integration Broker gateway.
env.gateway_ssl_port	@sslport@	Gateway SSL port.
env.use_ssl_gateway	False	Flag to determine whether the gateway is secure or non-secure for callback. Valid values are: • False ― indicates non-secure (HTTP) configuration • True ― indicates secure (HTTPS) configuration. See *PeopleTools: Search Technology,* "Configuring SSL Between PeopleSoft and Elasticsearch."
env.default_local_node	@nodename@	Default local node
env.enable_global_menu_search	All	Enable global search and menu search for All, Local, or a comma-separated list or portals.

This example shows the left side of the page.

Configure Search Instance page

6. Save the template and click Execute to begin the program run.

7. To verify that the search instance page has been configured for Elasticsearch, select PeopleTools, Search Framework, Administration, Search Instance.

   On the Search Instance Properties page, make sure the Search Provider is Elasticsearch, and that the Callback URL includes RESTListeningConnector, as in this example:

Search Instance Properties page

8. Select PeopleTools, Automated Config Manager, ACM Templates, Template Definition.

9. Search for and open SEARCH_TEMPLATE.

10. On the Configuration Template Definition page, verify that Deploy Search Definition is selected and click the Properties icon.

   Note. Do not select Configure Search Instance.

11. Specify the properties for the plug-in on the Deploy Search Definition page and click OK.

This table describes the properties for the Deploy Search Definition plug-in (PTSFAdministerSearch) in the SEARCH_TEMPLATE. Use the plug-in to deploy, undeploy, and schedule index generation.

Property	Default Value	Description
env_ptsf_search instance	PTSF_SEARCH	Elasticsearch instance name

env.ptsf_selection_type	GLOBAL	Valid values:
		
		• ALL
		
		  Deploys all search definitions and categories excluding the ones listed in env.ptsf_exclude_definitions.
		
		• GLOBAL
		
		  Deploys search definitions and categories used for Global search.
		
		• LIST
		
		  Deploys the search definitions and their categories mentioned in env.ptsf_include_definitions.

env.ptsf_include_definitions	NA	Comma separated list of search category names to be included. You can use % as an operator in any part of the name. For example, EP_CS%, %CS_DOC%, %CS%DOC% and so on.
env.ptsf_exclude_definitions	NA	Comma separated list of search category names to be excluded. You can use % as an operator in any part of the name. For example, EP_CS%, %CS_DOC%, %CS%DOC% and so on.
env.ptsf_check_audit_errors	True	If true check for access to query/connected query, or invalid objects. For a search definition if audit errors are found, it will not proceed further with the action specified in the ptsf_admin_operations property for this particular search definition. It will continue with the next one.

Property	Default Value	Description
env.ptsf_admin_operations	DEPLOY,INDEX	Valid values: • DEPLOY — the search definitions will be deployed • INDEX — the search index will be scheduled. The INDEX option will create run control ids for both full indexing and incremental indexing, but schedules only full indexing for the first execution of the SEARCH_TEMPLATE. The full indexing run control ids will have the naming convention \<SearchDefinition\>_FULL and incremental indexing will have the naming convention \<SearchDefinition\>_INCR. When the SEARCH_TEMPLATE is executed the second and subsequent times, it will run incremental indexing only if the previously done FULL indexing is successful; otherwise it will again schedule full indexing. **Note.** To schedule recurring incremental indexing, you must set the recurrence manually using the incremental run control ids created by ACM framework from the Schedule Index page. ACM will not do this automatically. • UNDEPLOY — the search definition will be removed. • DEPLOY,INDEX
env.ptsf_index_all_lanaguages	False	If true the schedules are created to index all languages.
env.report_schedule_status_after_min utes	NA	Maximum minutes to wait before reporting scheduling status. • Enter 0 to wait till finish. • Leave blank to skip report. • Enter the time to wait to show status. For example, enter 5 to show the status once after 5 minutes.

Property	Default Value	Description
env.ptsf_schedule_on_server	NA	Specify the Process Scheduler to run the indexing on. Leave this blank to use master scheduler.

This example shows the left-hand side of the page.
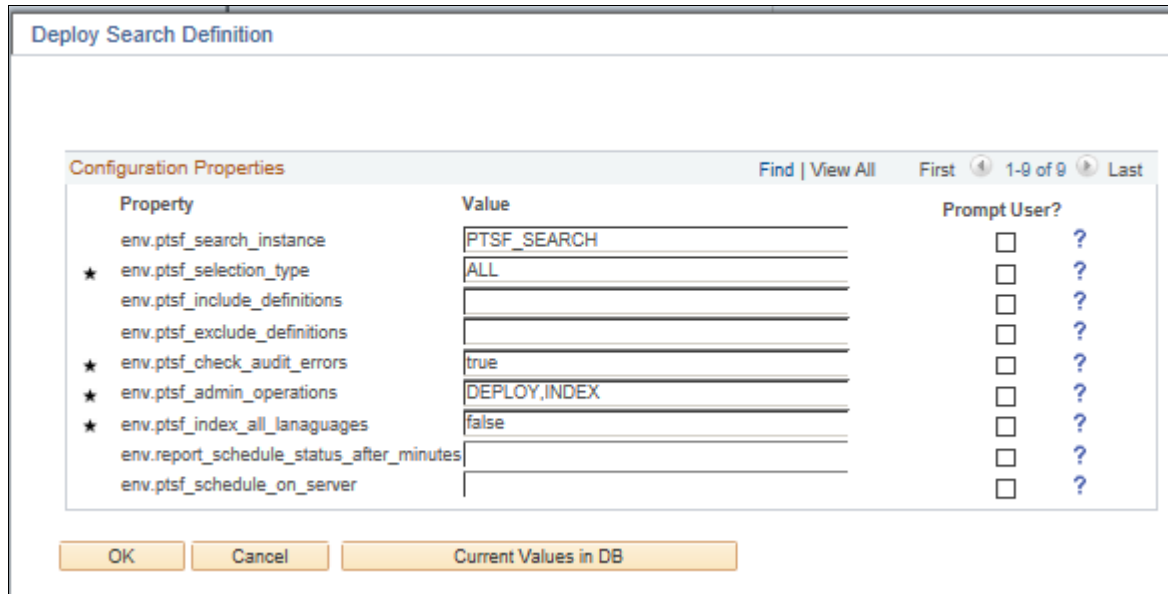
Deploy Search Definition page

12. Save the template and click Execute to begin the deployment.

Task 3-3-2: Configuring the Search Instance on the Search Instance Properties Page

If you need to add other search instances, you also have the option of adding the search instance using the Search Instance Properties page mentioned in the previous section. See the section on working with search instances in the product documentation.

Note. This initial configuration can be done instead of the ACM configuration described above, but one benefit of using the ACM Deploy option is that it will build run controls for the user running the ACM.

See *PeopleTools: Search Technology,* "Administering PeopleSoft Search Framework."

Chapter 4

Performing Additional Tasks

This chapter discusses:

- Modifying the Elasticsearch Configuration File (Optional)
- Starting and Stopping an Elasticsearch Service
- Adding Additional Elasticsearch Nodes
- Bringing Up an Elasticsearch Node
- Using the Elasticsearchuser Script
- Adding Elasticsearch as a Service in Linux

Task 4-1: Modifying the Elasticsearch Configuration File (Optional)

If you need to override the default values provided by the Elasticsearch software, you can use the elasticsearch.yml file. Go to the *ES_HOME/*config directory to locate the elasticsearch.yml file, and modify it in a text editor for your environment. It is probably a good idea to make a backup copy before modifying the file.

Note. Because the Elasticsearch DPK setup script automates the configuration, modifying elasticsearch.yml should not normally be necessary.

- cluster.name — a unique name for the cluster.

  This parameter identifies the cluster for auto-discovery. Make sure the name is unique. Do not reuse the same cluster names in different environments, because you might end up with nodes joining the wrong cluster.

- node.name — any meaningful name, such as *hostname,* which would make it easy to identify where the node is running.

- path.data — the path to the location where you want to store the Elasticsearch data.

  To include multiple paths, use commas to separate the paths.

- path.logs — the path to the location where you want to store the Elasticsearch logs.

- bootstrap.mlockall — if set to *True,* this parameter locks the memory when the Elasticsearch instance is started.

  The recommendation is to set this to True.

- network.host — the IP address or hostname of the machine.

- http.port — the port where Elasticsearch should listen for incoming requests.

  Set this property whether or not SSL is configured. Elasticsearch can work either in https or http mode, but not in both modes at the same time.

- discovery.zen.ping.unicast.hosts — the list of hosts that the node has to ping to discover other nodes to join the cluster.

  Elasticsearch nodes will find each other via unicast. Provide the address in the format: ["host1", "host2"].

- discovery.zen.minimum_master_nodes — configure the majority of nodes (total number of master nodes / 2 + 1) in order to avoid an issue referred to as "split brain."

  For information on setting up nodes and the definition of "split brain" see the Elasticsearch product documentation.

- gateway.expected_nodes — the number of data or master nodes that are expected to be in the cluster.

  Recovery of local shards will start as soon as the expected number of nodes have joined the cluster.

Note. This parameter may be specified in the elasticsearch.yml file as needed.

- gateway.recover_after_nodes — the number of data or master nodes required for recovery.

  Recovery will take place as long as this many data or master nodes have joined the cluster. Once the recover_after_time duration has passed, (the default is 5 minutes), recovery will start as long as the gateway.recover_after_nodes condition is met. Whenever you add a new user or role to Elasticsearch, make sure that it is added to each of the nodes.

- node.max_local_storage_nodes — the number of nodes on a single system

  Enter "1" to disable multiple nodes on a single system.

- action.destructive_requires_name — When set to *True,* this will require explicit names when deleting indexes.

- orclssl.http.ssl — enable or disable https

  Accepts values true or false.

- orclssl.transport.ssl — enable or disable transport layer encryption

  Accepts values true or false.

- orclssl.keystore — path to the keystore

  The keystore file must be placed under *ES_HOME/*config.

- orclssl.keystore_password — keystore password.

  Provide an encrypted password, which is obtained using the encryption mechanism in elasticsearchuser script.

  See Using the Elasticsearchuser Script.

- orclssl.truststore — path to the truststore

  The truststore file must be placed under *ES_HOME/*config.

- orclssl.truststore_password — truststore password

  Provide an encrypted password which is obtained using the encryption mechanism in elasticsearchuser script.

  See Using the Elasticsearchuser Script.

- acl.cache.delete.interval — Interval to run the delete process to clear the ACL security values cache. By default it is 24h.

- acl.cache.delete.bulk_size — Number of deletions to run in one bulk delete request.

  By default it is 10000.

Task 4-2: Starting and Stopping an Elasticsearch Service

If you modify the elasticsearch.yml or jvm.options file, use these instructions to stop and restart the Elasticsearch services.

On Microsoft Windows, the Elasticsearch service is installed by the DPK setup script.

1. Open a command prompt, and change directory to *ES_HOME/*bin.

2. To see the usage for the service command:

    ```
    elasticsearch-service
    elasticsearch-service.bat install|remove|start|stop|manager [SERVICE_ID]
    ```

3. To stop and remove the Elasticsearch service:

    ```
    elasticsearch-service.bat remove
    ```

4. To start the service, run these commands from *ES_HOME/*bin:

    ```
    elasticsearch-service.bat install
    elasticsearch-service.bat start
    ```

On Linux to start and stop the process:

1. In a terminal window, change directory to *ES_HOME/*bin.

2. To start the Elasticsearch process:

    ```
    nohup ./elasticsearch &
    ```

3. To stop the process:

    a. Use this command to find the Elasticsearch process ID:

    ```
    ps -ef | grep elas
    ```

    b. Use this command to stop the process, substituting the process ID for pid:

    ```
    kill <pid>
    ```

Task 4-3: Adding Additional Elasticsearch Nodes

Use these steps to add an additional Elasticsearch node after you have completed the Elasticsearch DPK installation. The additional node will be added to the same cluster.

1. Install Elasticsearch on a second server and provide the following information during the DPK setup script:

    *   For the cluster name, specify the same name as that of the first Elasticsearch node.

    *   At the prompt "Enter the host name of nodes which are already member of a cluster", specify the IP address for the first Elasticsearch node.

2. On the second Elasticsearch node (the second server), open *ES_HOME/*config/elasticsearch.yml for editing.

3. In elasticsearch.yml, set the value for discovery.zen.minimum_master_nodes to *(N/2) + 1* where N is the number of nodes in the Elasticsearch cluster, including the new one being added.

    To avoid "split brain" problem, it is always recommended to have an odd number of nodes (N) in the cluster if N is less than 4. .For more information, search for details about configuring minimum master nodes in the Elasticsearch online help.

See Elastic Web site, https://www.elastic.co/.

4. Stop and restart the second Elasticsearch node.

See Starting and Stopping an Elasticsearch Service.

Task 4-4: Bringing Up an Elasticsearch Node

If one of the nodes in an existing cluster is down, use these steps to bring it up:

1. Open *ES_HOME/*config/elasticsearch.yml for editing.

See Modifying the Elasticsearch Configuration File.

2. Set the value for discovery.zen.minimum_master_nodes to *(N/2) + 1* where N is the number of nodes in the Elasticsearch cluster including the one that is being brought up.

3. If you are running on Linux, set these environment variables:

   • Verify that the heap size is set to a value equal to or less than 50% of available memory, and not exceeding 30G.

     See "Preparing to Deploy," Reviewing Elasticsearch Recommendations.

   • Set the JAVA_HOME environment variable, where yy is the JRE version.

     ```
     export JAVA_HOME= BASE_DIR/pt/es_jre1.8.0_yy
     ```

4. Start the Elasticsearch node.

See Starting and Stopping an Elasticsearch Service.

Task 4-5: Using the Elasticsearchuser Script

To add users or roles after installation, use the elasticsearchuser script, found in *ES_HOME/*bin. The elasticsearchuser script has the following uses:

• To add a new user or change password for a user:

  ```
  elasticsearchuser adduser [user]
  ```

• To add roles for an existing user:

  ```
  elasticsearchuser addrole [user]
  ```

• To view existing users:

  ```
  elasticsearchuser listusers
  ```

• To view roles of a user:

  ```
  elasticsearchuser listrole [user]
  ```

• To remove a user:

  ```
  elasticsearchuser removeuser [user]
  ```

• To encrypt the given text:

  ```
  elasticsearchuser encrypt [text]
  ```

  This is used for encrypting the keystore password while configuring SSL. The password needs to be encrypted

in the elasticsearch.yml file.

Task 4-6: Adding Elasticsearch as a Service in Linux

This section discusses:

- Prerequisites
- Adding an Elasticsearch Service
- Verifying Automatic Startup
- Removing the Elasticsearch Service

Prerequisites

Use the instructions in this section to run Elasticsearch as a service on a Linux host, and to start automatically upon rebooting. Ensure that you fulfill these requirements:

- The Elasticsearch process should not be running.

  If Elasticsearch is running, ensure that it is not being used, and then kill the process.

  See Starting and Stopping an Elasticsearch Service.

- The script to install the service must be run by the root user.
- Elasticsearch is installed on the Linux server where you run this procedure.
- The Elasticsearch and JRE installation folders are located under the DPK base folder; that is:

  - The DPK base folder is referred to in this documentation as *BASE_DIR*, such as /home/esk612.
  - The Elasticsearch installation directory, *ES_HOME*, is found in *BASE_DIR/*pt/elasticsearch6.1.2.
  - The Java installation directory, *JAVA_HOME*, is found in *BASE_DIR/*pt/es_jre1.8.0_*yy*, where yy is the JRE version.

- Elasticsearch must be run by the user who owns *ES_HOME*.

  This is due to the fact that the script does not input the user name.

- You have downloaded and extracted the required Elasticsearch DPK for Linux, in a directory referred to as *ES_INSTALL*.

Task 4-6-1: Adding an Elasticsearch Service

The script to add or delete the service uses the following arguments:

- `-h or --help`

  Show the help message and exit

- `--add`

  An Elasticsearch service will be added.

- `--delete`

  The Elasticsearch service will be removed.

- `--install_base_dir INSTALL_BASE_DIR`

  Enter the base directory where Elasticsearch is installed.

To add Elasticsearch as a service:

1. Open a terminal window, running as root.

2. Change directory to *ES_INSTALL/*setup:

   ```
   cd ES_INSTALL/setup
   ```

3. Run this command to add the service:

   ```
   # ./psft-es-service.sh --add --install_base_dir BASE_DIR
   ```

4. Use one of these methods to verify that the service was added:

   - The output of the following ps command must show a running Elasticsearch process:

     ```
     # ps -ef | grep elastic
     ```

   - Use this chkconfig command:

     ```
     # chkconfig --list | grep elastic
     ```

     This should give an output such as:

     ```
     elasticsearch 0:off 1:off 2:on 3:on 4:on 5:on 6:off
     ```

   - Use this service command:

     ```
     # service elasticsearch status
     ```

     This should give an output such as:

     ```
     elasticsearch (pid 21292) is running...
     ```

Task 4-6-2: Verifying Automatic Startup

After you install the Elasticsearch service and verify the installation, test to make sure the service starts automatically after you reboot the Linux server.

1. Reboot the Linux server.

2. Use this command to verify that the service has come up automatically:

   ```
   # service elasticsearch status
   ```

   This should give an output such as:

   ```
   elasticsearch (pid 5028) is running...
   ```

Task 4-6-3: Removing the Elasticsearch Service

To remove the Elasticsearch service:

1. Open a terminal window, running as root.

2. Change directory to *ES_INSTALL/*setup:

   ```
   cd ES_INSTALL/setup
   ```

3. Run this command to remove the service:

   ```
   # ./psft-es-service.sh --delete
   ```

4. Use one of these methods to verify that the service was deleted:

- The output of the following ps command should not include any Elasticsearch process:

  ```
  # ps -ef | grep elastic
  ```

- This chkconfig command should give an empty output:

  ```
  # chkconfig --list | grep elastic
  ```

- Use this service command:

  ```
  # service elasticsearch status
  ```

  This should give an output such as:

  ```
  elasticsearch: unrecognized service
  ```

Appendix A

Moving from Elasticsearch 2.3.2 to Elasticsearch 6.1.2 with a PeopleTools Upgrade

This appendix gives an overview of the steps needed to transition from an environment with PeopleSoft PeopleTools 8.55 or 8.56 and Elasticsearch 2.3.2 to one with PeopleSoft PeopleTools 8.57 and Elasticsearch 6.1.2.

This overview assumes the following:

- You have a PeopleSoft 9.2 application.

  The PeopleSoft application must include the appropriate fixes for Elasticsearch, as delivered with the latest PeopleSoft Update Image.

  See "Integrating Elasticsearch with the PeopleSoft Environment," Applying PeopleSoft Application Support Bugs for Elasticsearch.

- The PeopleSoft application is on PeopleSoft PeopleTools 8.55 or 8.56.

- The environment is using Elasticsearch 2.3.2 for the PeopleSoft Search Framework.

- You want to deploy, configure, and test Elasticsearch 6.1.2 with the same environment and at the same time that you continue using Elasticsearch 2.3.2.

- You want to upgrade to PeopleSoft PeopleTools 8.57.04.

This table outlines the steps and resources:

Task	Description	Reference
1. Apply PeopleTools patch.	Apply a PeopleTools patch to update from PeopleTools 8.55.*xx* to PeopleTools 8.55.26, or from PeopleTools 8.56.*yy* to PeopleTools 8.56.12.	PeopleSoft PeopleTools Deployment Packages Installation (PeopleTools 8.55 or 8.56), "Applying PeopleTools Patches Using DPKs." See PeopleTools Patches Home Page, select tab Best Practices, My Oracle Support, Doc ID 2062712.2.

Task	Description	Reference
2. Install the Elasticsearch 6.1.2 DPK.	Obtain and install the PeopleTools 8.55 or PeopleTools 8.56 version of the Elasticsearch 6.1.2 DPK. You have the option to install on the same server that you use for Elasticsearch 2.3.2.	See *PeopleTools Deployment Packages for Elasticsearch Installation (PeopleTools 8.55)* PeopleTools Elasticsearch Home Page, select the tab for Elasticsearch on PeopleTools 8.55, My Oracle Support, Doc ID 2205540.2. Or See *PeopleTools Deployment Packages for Elasticsearch Installation (PeopleTools 8.56)* PeopleTools Elasticsearch Home Page, select the tab for Elasticsearch on PeopleTools 8.56, My Oracle Support, Doc ID 2205540.2.
3. Move from Elasticsearch 2.3.2 to 6.1.2 with live-cutover.	Configure Elasticsearch 6.1.2 while continuing to run Elasticsearch 2.3.2. This includes creating the search instance, building indexes, and setting the priority of the Elasticsearch 6.1.2 search instance.	"PeopleSoft Applications 9.2: Moving from Elasticsearch 2.3.2 to Elasticsearch 6.1.2" See PeopleTools Elasticsearch Home Page, My Oracle Support, Doc ID 2205540.2.
4. Upgrade to PeopleTools 8.57.	Perform a PeopleTools-only upgrade to PeopleSoft PeopleTools 8.57.04.	See PeopleSoft PeopleTools 8.57 Upgrade Home Page, My Oracle Support, Doc ID 2442414.1.
5. Reuse the Elasticsearch 6.1.2. indexes.	Use the search feature in the upgraded environment.	See *PeopleTools 8.57: Search Technology.*
6. Install Kibana (Optional).	If you want Kibana, install the Elasticsearch and Kibana 6.1.2 DPK for PeopleTools 8.57. Run the setup script and install Kibana only. When running the installation script, specify the Elasticsearch host and port from the installation in step 2 of this table.	See *PeopleTools Deployment Packages for Elasticsearch Installation (PeopleTools 8.57),* PeopleTools Elasticsearch Home Page, select tab Elasticsearch on PeopleTools 8.57, My Oracle Support, Doc ID 2205540.2. See *PeopleTools: Search Technology,* "Setting Up Kibana in PeopleSoft Search Framework."