
PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.58)

October 2020

PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.58)
Copyright © 2020, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

The business names used in this documentation are fictitious, and are not intended to identify any real companies currently or previously in existence.

Documentation Accessibility

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

Access to Oracle Support

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

Contents

Preface

About This Documentation	7
Understanding This Documentation	7
Typographical Conventions	7
Products	8
Related Information	9
Comments and Suggestions	10

Chapter 1

Preparing to Deploy	11
Understanding Elasticsearch	11
Understanding Elasticsearch	11
Understanding the Elasticsearch Deployment Package	12
Prerequisites	13
Reviewing Hardware Prerequisites	13
Reviewing Software Prerequisites	14
Reviewing Elasticsearch Recommendations	14

Chapter 2

Deploying the Elasticsearch, Logstash, and Kibana Deployment Package	17
Understanding the Elasticsearch, Logstash, and Kibana Installation	17
Preparing for the Logstash Installation	18
Fulfilling Prerequisites	18
Obtaining the Integration Broker REST URL	19
Using Logstash with an SSL Setup	19
Generating JSON Files with an SSL Setup	20
Obtaining the Elasticsearch, Logstash, and Kibana DPK	22
Obtaining the ELK DPK from Oracle Software Delivery Cloud	22
Obtaining the ELK DPK from My Oracle Support	22
Using the PT-INFRA DPK for Additional Component Software	23
Installing Elasticsearch, Logstash, and Kibana on Linux	23
Installing Elasticsearch, Logstash, and Kibana on Linux Interactively	23
Installing Elasticsearch, Logstash, and Kibana on Linux in Silent Mode	28
Verifying the Elasticsearch DPK Installation on Linux	34
Removing the Elasticsearch Installation from Linux	35

- Removing the Kibana Installation from Linux 35
- Removing the Logstash Installation from Linux 35
- Performing Post-Installation Steps on Linux 35
- Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows 36
 - Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows Interactively 36
 - Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows in Silent Mode 40
 - Verifying the Elasticsearch Installation on Microsoft Windows 46
 - Removing the Elasticsearch Installation from Microsoft Windows 48
 - Removing the Kibana Installation from Microsoft Windows 48
 - Removing the Logstash Installation from Microsoft Windows 48
 - Performing Post-Installation Steps on Microsoft Windows 49

Chapter 3

- Upgrading Elasticsearch and Kibana 51**
 - Upgrading Elasticsearch and Kibana to a New Revision Interactively 51
 - Upgrading to a New Revision on Microsoft Windows 51
 - Upgrading to a New Revision on Linux 52
 - Upgrading Elasticsearch from 6.1.2 to 7.0.0 Interactively 53
 - Upgrading Interactively on Microsoft Windows 53
 - Upgrading Interactively on Linux 55
 - Upgrading Elasticsearch from 6.1.2 to 7.0.0 in Silent Mode 56

Chapter 4

- Integrating Elasticsearch with the PeopleSoft Environment 59**
 - Applying PeopleSoft Application Enhancements for Kibana 59
 - Setting Up the PeopleSoft Application for Elasticsearch 60
 - Understanding the PeopleSoft Application Setup 60
 - Verifying the Integration Broker Setup 60
 - Verifying PeopleSoft Roles for All Installations 61
 - Adding and Configuring an Elasticsearch Instance 62
 - Using the Automated Configuration Management SEARCH_TEMPLATE 62
 - Configuring the Search Instance on the Search Instance Properties Page 72

Chapter 5

- Performing Additional Tasks 73**
 - Modifying the Elasticsearch Configuration File (Optional) 73
 - Starting and Stopping an Elasticsearch Service 75
 - Adding Additional Elasticsearch Nodes 75
 - Bringing Up an Elasticsearch Node 76

Using the Elasticsearchuser Script	76
Adding Elasticsearch as a Service in Linux	77
Prerequisites	77
Adding an Elasticsearch Service	77
Verifying that the Elasticsearch Service Starts Automatically	78
Removing the Elasticsearch Service	79
Adding Kibana as a Service in Linux	80
Prerequisites	80
Adding a Kibana Service	80
Verifying that the Kibana Service Starts Automatically	81
Removing the Kibana Service	81
Reviewing the Logstash Configuration Files (Optional)	82

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.
<i>Italics</i>	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 <i>O</i> . Italics are also used to indicate user-supplied information. For example, the term <i>domain</i> 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 <code><PS_CFG_HOME>/appserv/<domain></code> 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	<p>Cross-references that begin with <i>See</i> refer you to additional documentation that will help you implement the task at hand. We highly recommend that you reference this documentation.</p> <p>Cross-references under the heading <i>See Also</i> refer you to additional documentation that has more information regarding the subject.</p>
⇒ (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 <i>Note.</i> 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 <i>Important!</i> 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 <i>Warning!</i> 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 page in the Oracle Help Center. You can also find documentation by searching for the product name on My Oracle Support.

- PeopleSoft on the Oracle Help Center

You can access PeopleSoft Online Help, or download the PeopleBooks PDFs, from the PeopleSoft page in the Oracle Help Center. Select PeopleTools or your PeopleSoft application from the navigation list on the left. On the page for the selected product application, select the PeopleTools release or image number at the top and go to the Online Help and PeopleBooks section.

See Oracle Help Center, <https://docs.oracle.com/en/applications/peoplesoft/index.html>.

- *PeopleTools: Getting Started with PeopleTools* for your release.

This documentation provides a high-level introduction to PeopleTools technology and usage.

See PeopleTools on the Oracle Help Center, <https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html>.

- PeopleSoft Application Fundamentals for your PeopleSoft Application and release

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

See Oracle Help Center, <https://docs.oracle.com/en/applications/peoplesoft/index.html>.

- Installation guides

You can find the installation guides for PeopleSoft PeopleTools and your PeopleSoft application on the appropriate Oracle Help Center page. Select your release or update image at the top and then go to the Install and Upgrade section.

- My Oracle Support

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

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

You can find several pages which compile documentation, links, and known issues for various PeopleSoft product areas. For a list of many of the PeopleSoft pages, select the PeopleSoft tab on the Oracle Information Center Catalog.

See Oracle Information Center Catalog, My Oracle Support, Doc ID 50.2.

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 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, Kibana, Logstash, 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.

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, Kibana and Logstash.
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 8.58. For more information Elasticsearch, see these resources:

- For details on using Elasticsearch with the PeopleSoft Search Framework, see *PeopleTools: Search Technology*.
See PeopleTools Online Help on the Oracle Help Center, <https://docs.oracle.com/en/applications/peoplesoft/index.html>.
- You can find the most current version of this installation documentation, *PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.58)*, 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 8.58), 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 8.58), 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 8.58 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 8.58)*, 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, Logstash, and Kibana DPK from Oracle Software Delivery Cloud or My Oracle Support. This documentation sometimes uses "ELK DPK" to refer to the Elasticsearch, Logstash, and Kibana DPK. The ELK DPK includes:

- **Elasticsearch**
Elasticsearch is an open-source search engine used for the PeopleSoft Search Framework.
- **Kibana**
Use Kibana to visualize data for Elasticsearch server monitoring and application index monitoring, and to create dashboards to visualize application search indexes. The installation of Kibana is optional.
- **Logstash**
Logstash collects JMX (Java Management Extensions) metrics from JMX servers and provides them to Elasticsearch. The PeopleSoft Health Center displays the charts and data using Kibana.
The ELK DPK includes an input plugin for the JMX agents used for the PeopleSoft Health Center. The installation of Logstash is optional. For the current release, Logstash is supported only for use with PeopleSoft Health Center. Due to security concerns, any other data indexed using this Logstash will not be supported.
- PeopleSoft-developed plug-ins for Elasticsearch, Kibana, and Logstash
- PeopleSoft-delivered customized code for Elasticsearch
- Oracle Java JRE 1.8
- Open-source Python software
- PeopleSoft-developed Python scripts for Elasticsearch deployment automation

The ELK 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, Logstash, and Kibana DPK directly on a system running a Linux or Microsoft Windows operating system. The ELK 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 ELK DPK can be installed on a Linux or Microsoft Windows host (bare metal or virtual).

Installing Kibana and Logstash are optional. Oracle recommends that you install the ELK DPK on a server that is separate from those used for PeopleSoft installations. Kibana and Logstash can be installed on one of the Elasticsearch nodes in the Elasticsearch cluster for improved performance, depending on the available memory and CPU. For PeopleSoft Health Center and Elasticsearch monitoring purposes, completely different Kibana and Elasticsearch instances can be used to isolate the Kibana used for application analytics.

If you want to use PeopleSoft Health Center, you must install Elasticsearch, Kibana, and Logstash.

See *PeopleTools: Search Technology*.

- Host operating system

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

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

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, Logstash, and Kibana integration is supported for PeopleSoft 9.2 applications on PeopleSoft PeopleTools 8.58.
 - The ELK DPKs are updated regularly to incorporate the latest JRE. ELK DPKs are released concurrently with PeopleTools patches, but are not dependent on the PeopleTools patch release. You can use the ELK 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 8.58.
 - Oracle recommends that you use the ELK 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 ELK 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/7.0/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, Logstash, and Kibana Deployment Package

This chapter discusses:

- Understanding the Elasticsearch, Logstash, and Kibana Installation
- Preparing for the Logstash Installation
- Obtaining the Elasticsearch, Logstash, and Kibana DPK
- Using the PT-INFRA DPK for Additional Component Software
- Installing Elasticsearch, Logstash, and Kibana on Linux
- Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows

Understanding the Elasticsearch, Logstash, and Kibana Installation

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

- Run the script interactively, or run it in silent mode with a prepared configuration file.
- Install Elasticsearch, Logstash, and Kibana together or separately.
- If you choose to install Kibana at the same time as you install Elasticsearch, the ELK DPK setup 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.
- If you choose to install Kibana at a different time, 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.
- If you are doing a new installation, you can use the PT-INFRA DPK to install the latest JRE version.
You can use the PT-INFRA DPK with ELK DPK version 02 and later (ELK-DPK-WIN-7.0.0_02.zip or ELK-DPK-LNX-7.0.0_02.zip).
See [Using the PT-INFRA DPK for Additional Component Software](#).
- Specify Elasticsearch clusters and nodes.
The Elasticsearch 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.

- See the requirements and recommendations for Logstash in the following section, Preparing for the Logstash Installation.
- When using the ELK DPK setup script:
 - The user who installs the ELK 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, Logstash, 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 *PeopleTools: Performance Monitor*, "Understanding PeopleSoft Health Center."

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

Task 2-1: Preparing for the Logstash Installation

This section discusses:

- Fulfilling Prerequisites
- Obtaining the Integration Broker REST URL
- Using Logstash with an SSL Setup
- Generating JSON Files with an SSL Setup

Task 2-1-1: Fulfilling Prerequisites

As mentioned, Logstash is used for PeopleSoft Health Center. Oracle recommends that you install Logstash on a different system than the systems on which the PeopleSoft system is installed.

Before installing Logstash for the PeopleSoft Health Center, ensure that:

- PeopleSoft Pure Internet Architecture (PIA) is running.
- Integration Broker is configured.
- The Integration Broker user must be a PeopleSoft user who has access to the service operation PT_CREATEJSON_REST_GET.
- The role assigned to the Integration Broker user must include permission list PTPT4800 (ACM Administrator).

See *PeopleTools: Security Administration*.

- PPM agents are enabled.
- JMX agents are enabled.
- You have run the ACM plug-in PTSFMonitorConfiguration.

Make sure the PPM agents have equivalent JMX agents registered. The Agent IDs that the following SQL returns should have entries in the PS_PTPMJMXUSER table. If there are missing agents in the PS_PTPMJMXUSER table, it either means your JMX registration is not successful or your PSPMAGENT table is carrying agents that are no longer needed.

```
select * from PSPMAGENT where PM_DOMAIN_MONITOR='Y';
```

See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

Task 2-1-2: Obtaining the Integration Broker REST URL

Before beginning the Logstash installation, make a note of the Integration Broker REST service URL, user name and password. You can obtain the URL with these steps:

1. Sign in to the PeopleSoft installation (PIA) in a browser.
2. Select PeopleTools > Integration Broker > Integration Setup > Service Operations.
3. Locate the service operation PT_CREATEJSON_REST_GET.
4. On the General page for the service operation, in the URI grid, select the Validate link.
5. Select Generate URL and make a note of the REST URL.

See *PeopleTools: Integration Broker*, "Accessing and Viewing REST Service Operation Definitions."

If the setup is such that Performance Monitor is configured for remote monitoring, the Integration Broker REST URL should be the one that connects to the monitoring database. If the PPM and JMX agents registrations are successful, the registration details are available in a monitoring system database in the tables PSPMAGENT and PS_PTPMJMXUSER.

If you perform the ELK DPK installation before registering the PPM and JMX agents, carry out these steps to get the JSON configuration files for Logstash:

1. Register the PPM and JMX agents.
See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."
2. Enter the value for the Integration Broker REST URL in `LOGSTASH_HOME\pt\config\JsonLogstash.properties`.
See "Performing Additional Tasks," Reviewing the Logstash Configuration Files.
3. Go to `LOGSTASH_HOME\pt\bin` and run the script `CreateJSON.bat` (Microsoft Windows) or `CreateJSON.sh` (Linux) to get the JSON files.
See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

Note. If you want add a new application to be monitored, rerun `CreateJSON.bat` (Microsoft Windows) or `CreateJSON.sh` (Linux). It is not necessary to restart Logstash.

Task 2-1-3: Using Logstash with an SSL Setup

This section includes guidelines for using Logstash with an Elasticsearch installation that uses SSL signon.

1. Open `LOGSTASH_HOME\pt\config\LogstashPipeLine.CONF` for editing.
See "Performing Additional Tasks," Reviewing the Logstash Configuration Files.
2. Add a line specifying the Elasticsearch root certificate in the output section.

In this example, the cacert line is in bold font:

```
input {
  jmx {
    path => "C:\elk70\pt\Logstash7.0.0\pt\jmxmonitor"
    polling_frequency => 5
    type => "jmx"
    nb_thread => 15
  }
}
output {
  elasticsearch {
    hosts => [" https://<ES_host>:<ES_port>"]
    index => "psft_hc_metrics"
    user => "esadmin"
    password => "encrypted_password"
    cacert => 'C:\elk70\pt\elasticsearch7.0.0\plugins\orcl-security-plugin->
\config\properties\cacert.cer'
  }
}
```

Task 2-1-4: Generating JSON Files with an SSL Setup

The creation of the JSON configuration files requires connection to the Integration Broker REST URL. By default the connection to the Integration Broker REST URL is authenticated using Basic Authentication (Integration Broker User name and password). If your environment is set up to use SSL, you must create the JSON files manually after completing the ELK DPK installation.

1. Install Logstash, and answer *n* (no) to the prompt about JSON files.
2. Sign in to the PeopleSoft installation (PIA) in a browser and ensure that the Integration Broker user has permission list PTPT4800.
3. Select PeopleTools > Integration Broker > Integration Setup > Service Operations.
4. Locate the service operation PT_CREATEJSON_REST_GET.

5. Select *Basic Authentication and SSL* from the Req Validation drop-down list.

This example shows the top portion of the General page for the PT_CREATEJSON_REST_GET service operation.

The screenshot displays the 'General' tab of the Service Operation configuration page. Key fields include:

- Service Operation:** PT_CREATEJSON_REST_GET
- REST Method:** GET
- *Operation Description:** Create JSON
- Operation Comments:** (Empty text area)
- Owner ID:** PeopleTools
- Operation Alias:** PT_CREATEJSON_REST
- *Req Verification:** Basic Authentication and SSL (selected from a dropdown menu)
- REST Resource Definition:**
 - REST Base URL:** http://server.example.com:8000/PSIGWRESTListeningConnector/H92WS125/PT_CREATEJSON_REST.v1
 - URI Template Format Example:** weather/{state}/{city}?forecast={day}
 - URI Grid:**

Index	Template	Validate	Build
1	json={refValue}	Validate	Build
 - Document Template:** CREATE_JSON.v1

Service Operation General page

6. On the General page for the service operation, in the URI grid, select the Validate link.
7. Select the SSL check box and then select Generate URL.
8. Make a note of the REST URL on the Validate URI window.
9. Save the changes.
10. Obtain a certificate from a Certificate Authority (CA) and save it as certLS.jks.
11. Put the certLS.jks truststore file under `LOGSTASH_HOME\pt\config\`.
12. To create the JSON files after completing the ELK DPK installation, go to `LOGSTASH_HOME\pt\bin` and run `CreateJSON.bat` (Microsoft Windows) or `CreateJSON.sh` (Linux).

See Also

PeopleTools: Search Technology, "Configuring SSL between PeopleSoft and Elasticsearch"

PeopleTools: Security Administration, "Installing Web Server-Based Digital Certificates"

Task 2-2: Obtaining the Elasticsearch, Logstash, and Kibana DPK

This section discusses:

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

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

To obtain the ELK 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. Select Checkout.
5. On the Selected Software page, locate Elasticsearch in the search results and select Microsoft Windows or Linux from the Platforms/Languages drop-down list.

Clear the check boxes for the other items in the list.

6. Click Continue.
7. Review and accept the license download agreement.
8. 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 *ELK_FILENAME.zip*, and the directory where you download the zip file as *ELK_INSTALL*.

Task 2-2-2: Obtaining the ELK DPK from My Oracle Support

To obtain the ELK 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, ELK-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 7.0.0.
 - *<DPK_revision>* is a number specific to the DPK revision.

For example, ELK-DPK-WIN-7.0.0_03.zip or ELK-DPK-LNX-7.0.0_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 *ELK_FILENAME.zip*, and the directory where you download the zip file as *ELK_INSTALL*.

Task 2-3: Using the PT-INFRA DPK for Additional Component Software

The PT-INFRA DPK contains supporting (third-party) software that is required for a PeopleSoft installation. A separate PT-INFRA DPK is delivered as needed to provide security updates or other patches for one or more of the components. You can use this separate PT-INFRA DPK in a new installation to take advantage of up-to-date patches and security updates (CPUs).

The PT-INFRA DPK includes the following:

- Oracle Tuxedo
- Oracle WebLogic
- Oracle Database Client
- JDK and JRE

See *PT-INFRA Deployment Package Installation (PeopleSoft PeopleTools 8.58)*, Oracle's PeopleSoft PeopleTools 8.58 Home Page, Installation and Upgrade tab, My Oracle Support, Doc ID 2602329.2.

Task 2-4: Installing Elasticsearch, Logstash, and Kibana on Linux

This section discusses:

- Installing Elasticsearch, Logstash, and Kibana on Linux Interactively
- Installing Elasticsearch, Logstash, 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
- Removing the Logstash Installation from Linux
- Performing Post-Installation Steps on Linux

Task 2-4-1: Installing Elasticsearch, Logstash, and Kibana on Linux Interactively

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

- You have downloaded the required ELK DPK for Linux, referred to as *ELK_FILENAME.zip*, and saved it in a newly created directory accessible to the Linux host, referred to as *ELK_INSTALL*.
- There is enough space on the Linux host for the Elasticsearch installation and your estimated indexing requirements.
- Review the requirements described previously in Preparing for the Logstash Installation.

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 *ELK_INSTALL*.

```
cd ELK_INSTALL
```

3. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_INSTALL*.

```
unzip ELK_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 *ELK_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 *ELK_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/elk70
```

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

5. If you are using the PT-INFRA DPK with the ELK DPK, verify that you see the progress message:

```
Extracting PTINFRA DPK  
[OK]
```

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

```
You've chosen to do a fresh installation of Elasticsearch, Logstash  
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 ELK 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 full path location for the Elasticsearch data.

Oracle recommends that you do not use the default location, *BASE_DIR*/pt/elasticsearch7.0.0/data, with PeopleSoft environments. Instead, specify the full path for a data directory that is outside of *BASE_DIR*/pt/elasticsearch7.0.0.

Enter the path where you want the Elasticsearch data to reside [/home=>/elk70/pt/elasticsearch7.0.0/data] :

13. Enter the location for the Elasticsearch logs.

The default location is *BASE_DIR*/pt/elasticsearch7.0.0/logs.

Enter the path where you want the Elasticsearch Logs to be written to [=>/home/elk70/pt/elasticsearch7.0.0/logs] :

14. Enter the heap size in GB.

Enter a number 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]
```

```
SUCCESS: Specified value was saved. [OK]
Elasticsearch Installation Completed.
```

16. Answer *y* (yes) to install Logstash, or *n* (no) to continue.

```
Do you want to install Logstash (Needed for PeopleSoft Health Centre): =>
(y/n): y
```

17. Enter the URL for the Integration Broker REST service.

```
Enter the IB REST service URL:
```

18. Enter the Integration Broker user name.

Note. Be sure to fulfill the prerequisites for the Integration Broker user in the section Preparing for the Logstash Installation.

```
Enter the IB user:
```

19. Enter the Integration Broker password twice.

```
Enter the IB password:
Re-enter the IB password:
```

20. Enter the full path to the location to save the JSON files.

```
Enter the location where JSON files need to be created [/home/elk70/pt=>
/Logstash7.0.0/pt/jmxmonitor ]:
```

21. Enter the polling frequency, in seconds, for the JMX agents.

The polling frequency is mandatory input. There is no recommended or default value.

The value you enter sets the frequency with which the JMX metrics data are fetched from JMX servers. This means that every n seconds, data is pushed to Elasticsearch.

It is a good idea to regularly purge the `psft_hc_metrics` index. The interval you select depends upon your usage. See the information on the Monitoring Server page in the Search Technology product documentation.

See *PeopleTools: Search Technology*, "Configuring the Monitoring Server."

```
Enter the polling frequency for JMX metrics:
```

22. Enter the number of threads that will be used to retrieve metrics and create events.

The number you select depends upon your usage. Increase or decrease it as needed.

```
Enter the number of threads:
```

23. Enter the information you used to set up the Elasticsearch host.

```
Enter the Elasticsearch host name:
```

```
Enter the Elasticsearch port:
```

```
Enter the Elasticsearch username:
```

```
Enter the Elasticsearch password:
Re-enter the Elasticsearch password:
```

24. Enter `y` (yes) if you want to create JSON files, or `n` (no) to continue.

If your environment is set up to use SSL, or if you have not completed the necessary prerequisite steps, enter `n` (no). See the section Preparing for the Logstash Installation for information on creating the JSON files manually.

```
Do you want to create JSON files? (y/n): y
```

25. Wait while Logstash is installed.

```
Extracting the Logstash Binary .....[OK]
Configuring Logstash.....[OK]
Verifying if Logstash config files are present in place....
Logstash config files are present.

Creating Json files. This may take some time...

Logstash installation completed
```

26. 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
```

27. 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):
```

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

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

29. 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 ] :
```

30. 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.
```

31. After you complete the Elasticsearch, Kibana, and Logstash installation, you must configure the integration with the PeopleSoft environment.

See "Integrating Elasticsearch with the PeopleSoft Environment."

32. To start and use Kibana, see the Search Technology product documentation.

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

33. To use Logstash, see the Performance Monitor product documentation.

Note. The ELK DPK installation creates JSON files and starts Logstash if the required conditions are met.

See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

A successful deployment includes the following:

- The Elasticsearch home directory is found in *BASE_DIR/pt/elasticsearch7.0.0*.
- The *ES_HOME* environment variable is set to *BASE_DIR/pt/elasticsearch7.0.0* for the current terminal session.
- The Kibana home directory is found in *BASE_DIR/pt/Kibana7.0.0*.
- The *LOGSTASH_HOME* environment variable is set to *BASE_DIR/pt/Logstash7.0.0* for the current terminal session.
- 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 2-4-2: Installing Elasticsearch, Logstash, and Kibana on Linux in Silent Mode

This section discusses:

- Encrypting the Elasticsearch Passwords on Linux
- Encrypting the Logstash Passwords on Linux
- Creating a Configuration File and Running the Silent Mode Installation on Linux

Encrypting the Elasticsearch 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 ELK DPK. The *psvault* that is delivered with PeopleSoft PeopleTools will not work for Elasticsearch.

1. Open a terminal window, and change directory to *ELK_INSTALL*.

```
cd ELK_INSTALL
```

2. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_INSTALL*.

```
unzip ELK_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 *ELK_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
3. Change directory to *ELK_INSTALL/archives*, and extract *pt-jre-1.8.0_x.tgz* into a directory with the same name.

It may be necessary to extract the file twice.

4. Change directory to *ELK_INSTALL/archives* and extract *pt-elasticsearch-7.0.0.tgz* into a directory with the same name.

It may be necessary to extract the file twice.

- bin
 - config
 - jdk
 - lib
 - logs
 - modules
 - plugins
5. Copy *psvault* from the directory *ELK_INSTALL/pt-elasticsearch-7.0.0/plugins/orcl-security-plugin/config/properties* to *ELK_INSTALL/pt-elasticsearch-7.0.0/plugins/orcl-security-plugin*.
 6. Run the following command, specifying the passwords for *esadmin* and *people*.

```
<path_to_java>/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>/<outputfile>
```

For example:

```
/home/ELK_INSTALL/archives/pt-jre1.8.0_x/bin/java -Dpscipher.path=/home⇒
/ELK_INSTALL/archives/pt-elasticsearch-7.0.0/plugins/orcl-security-⇒
plugin -cp /home/ELK_INSTALL/archives/pt-elasticsearch-7.0.0/plugins⇒
/orcl-security-plugin/pscipher.jar com.peoplesoft.pt.elasticsearch.pscip⇒
her.PSESEncrypt esadmin password1 people password2 /home/ELK_INSTALL/es_⇒
output.txt
```

7. Copy the encrypted text from *es_output.txt*, without adding line feeds or spaces. Paste the encrypted passwords into the *silentinstall.config* file for the *esadmin.password* and *people.password* parameters.

Encrypting the Logstash Passwords on Linux

The passwords that you include in the silent mode text file must be encrypted. For the Integration Broker password for the Logstash installation, encrypt the password using the *PSLSCipher* script that is part of the ELK DPK.

1. Open a terminal window, and change directory to *ELK_INSTALL*.
2. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_INSTALL*.

```
unzip ELK_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 *ELK_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. Change directory to *ELK_INSTALL/archives* and extract *pt-jre-1.8.0_x.tgz* into a directory with the same name.

It may be necessary to extract the file twice.

4. Change directory to *ELK_INSTALL/archives* and extract *pt-logstash-7.0.0.tgz* into a directory with the same name.

It may be necessary to extract the file twice. The extraction creates several directories and text files.

5. Set the following environment variables.

Add *JAVA_HOME* to the *PATH*. Use the location where you extracted *pt-jre-1.8.0_xx.tgz*. Set *LOGSTASH_HOME* to the location where you extracted *pt-logstash-7.0.0.tgz*.

```
export PATH=JAVA_HOME/Bin:$PATH
export LOGSTASH_HOME=ELK_INSTALL/archives/pt-logstash-7.0.0
```

6. Change directory to *ELK_INSTALL/pt-logstash-7.0.0/pt/bin*.

7. Run the script, supplying the password:

```
./PSLSCipher.sh password
```

8. Copy the encrypted password from the output on the screen, without adding line feeds or spaces, and paste it into the *silentinstall.config* file for the *IB_PWD* parameter.

The output has this format:

```
Encrypted text: {V2.1}encrypted password==
```

The ELK DPK installation includes a *psvault* key file, which is found in the directory *LOGSTASH_HOME/pt/properties*. The key in this *psvault* is not related to the keys available in Elasticsearch or the PeopleSoft Web server (PIA) *psvault*. It is recommended that appropriate access be used to protect the Logstash *psvault*, as with other similar files.

To generate and use a new version of *psvault*:

1. Change directory to *ELK_INSTALL/pt-logstash-7.0.0/pt/bin*.
2. Run the script as follows to build a new key:

```
./PSLSCipher.sh -buildkey
```

3. Regenerate the Integration Broker user ID and password, and the Elasticsearch password with one of these commands:

```
./PSLSCipher.sh password
```

Or

```
./PSLSCipher.sh user ID
```

4. Edit the Logstash configuration files in the directory *LOGSTASH_HOME\pt\config* with the new encrypted text.

- Update the values for the Integration Broker user ID and password in *JsonLogstash.properties*.
- Update the values for the Elasticsearch password in *LogstashPipeLine.CONF*.

See *Reviewing the Logstash Configuration Files (Optional)*.

See *PeopleTools: Security Administration*, "Securing the External Key File."

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

To install in silent mode:

1. Create a configuration file.

The `ELK_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 values. See the section `Installing Interactively` for guidance on specifying the values.

2. Enter `y` (yes) if you want to install Elasticsearch, or `n` (no) if you want to skip the Elasticsearch installation. You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory. If you enter no, you can leave them blank. Enter encrypted text for the passwords as described in the previous section.

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

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

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

```
#Elasticsearch host name (mandatory)
network.host=
```

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

```
#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=
```

```
#Elasticsearch minimum master nodes (optional)
minimum_master_nodes=
```

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

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

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

3. Specify `y` if you want to install Logstash, or `n` if you do not.

You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory. If you enter no, you can leave them blank. Enter encrypted text for the passwords as described in the previous section.

```
#Install Logstash Y/N (Needed for PeopleSoft Health Centre) is
  mandatory field
Install Logstash (for PHC)?[Y/N]=#IB REST service URL (mandatory)
IB_REST_URL=

#Enter the encrypted IB user [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
IB_USER=

#Enter the encrypted IB password [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
IB_PWD=

#location where JSON files need to be created(optional-default LS_HOME⇒
/pt/jmxmonitor)
JSON_LOC=

#The polling frequency for JMX metrics (mandatory)
polling_freq =

#The number of threads (mandatory)
no_of_threads =

#The Elasticsearch host name (mandatory)
ES_host =

#The Elasticsearch port (mandatory)
ES_port =

#The Elasticsearch username (mandatory)
ES_user =

#The encrypted Elasticsearch password [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
ES_pwd =

#Do you want to create JSON files(Y/N) (mandatory)
JSON_files?[Y/N] =
```

4. Specify *y* if you want to install Kibana, or *n* if you do not.

You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory in this section. If you enter no, you can leave them blank.

```
#Install Kibana Y/N is mandatory field
Install kibana?[Y/N]=
#Kibana host name (mandatory)
kibana.host=

#Kibana port number (optional-default 5601)
```



```
kibana.port=
```

5. If you want to install Kibana and connect to a previously-installed Elasticsearch server, specify *n* (no) for the following:

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

6. Specify the host name, port, and encrypted password for the Elasticsearch server.

```
#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=
```

7. If the Elasticsearch is SSL enabled, enter the path to the certificate.

```
#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=
```

8. Enter *n* to indicate you do not want to upgrade.

You must enter a value for this field.

```
#####Silent Upgrade Configuration Values for upgrading⇒
Elasticsearch
from 6.1.2 to 7.0.0 #####
```

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

9. In a terminal window, run the DPK setup script from *ELK_INSTALL/setup* as follows:

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 BASE_DIR ⇒
--config_file 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 a line feed for readability. Do not include the line feed when you run.

```
./psft-dpk-setup.sh --install_silent --install_base_dir /home/elk70 ⇒
--config_file /home/temp_install/silentinstall.config
```

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

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

11. After you complete the Elasticsearch, Kibana, and Logstash installation, you must configure the integration with the PeopleSoft environment.

See "Integrating Elasticsearch with the PeopleSoft Environment."

12. To start and use Kibana, see the Search Technology product documentation.

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

13. To use Logstash, see the Performance Monitor product documentation.

Note. The ELK DPK installation creates JSON files and starts Logstash if the required conditions are met.

See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

Task 2-4-3: 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" : "7.0.0",
    "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"
}
```

Task 2-4-4: 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 2-4-5: 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 2-4-6: Removing the Logstash Installation from Linux

If you installed Logstash, use these steps to remove the Logstash 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 Logstash service:

```
ps -ef |grep Logstash
```

2. Kill the process, substituting the Logstash process ID for `<pid>`:

```
kill -9 <pid>
```

3. Remove the Logstash installation directory.

Task 2-4-7: Performing Post-Installation Steps on Linux

After you complete the Elasticsearch 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_jrel.8.0_yy
```

Task 2-5: Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows

This section discusses:

- Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows Interactively
- Installing Elasticsearch, Logstash, 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
- Removing the Logstash Installation from Microsoft Windows
- Performing Post-Installation Steps on Microsoft Windows

Task 2-5-1: Installing Elasticsearch, Logstash, 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 ELK DPK for Microsoft Windows, *ELK_FILENAME.zip*, and saved it in a newly created directory accessible to the Microsoft Windows host, referred to as *ELK_INSTALL*.
- There is enough space on the host for the Elasticsearch installation and your estimated indexing requirements.
- Review the requirements described previously in Preparing for the Logstash Installation.

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 *ELK_INSTALL*.
2. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_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 *ELK_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 *ELK_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:\elk70
```

- 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, Logstash and⇒
Kibana.
```

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

7. If you are using the PT-INFRA DPK with the ELK DPK, verify that you see the progress message:

```
Extracting PTINFRA DPK
[OK]
```

8. 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 ELK DPK and owns the files.

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

9. 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:
```

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

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

11. 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 ] :
```

12. 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"]

13. Enter the full path location for the Elasticsearch data.

Oracle recommends that you do not use the default location, *BASE_DIR*/pt/elasticsearch7.0.0/data, with PeopleSoft environments. Instead, specify the full path for a data directory that is outside of *BASE_DIR*\pt\elasticsearch7.0.0.

Enter the path where you want the Elasticsearch data to reside [C:⇒ \elk70\pt\elasticsearch7.0.0\data] :

14. Enter the location for the Elasticsearch logs.

The default location is *BASE_DIR*\pt\elasticsearch7.0.0\logs.

Enter the path where you want the Elasticsearch Logs to be written to [⇒ C:\elk70\pt\elasticsearch7.0.0\logs] :

15. Enter the heap size in GB.

Enter a number 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 installs Elasticsearch.

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

```
SUCCESS: Specified value was saved. [OK]
Elasticsearch Installation Completed.
```

17. Answer *y* (yes) to install Logstash, or *n* (no) to continue.

Do you want to install Logstash (Needed for PeopleSoft Health Centre): ⇒ (y/n) : **y**

18. Enter the URL for the Integration Broker REST service.

Enter the IB REST service URL:

19. Enter the Integration Broker user name.

Note. Be sure to fulfill the prerequisites for the Integration Broker user in the section Preparing for the Logstash Installation.

Enter the IB user:

20. Enter the Integration Broker password twice.

Enter the IB password:
Re-enter the IB password:

21. Enter the full path to the location to save the JSON files.

Enter the location where JSON files need to be created [E:\elk70\pt⇒ \Logstash7.0.0\pt\jmxmonitor] :

22. Enter the polling frequency, in seconds, for the JMX agents.

The polling frequency is mandatory input. There is no recommended or default value.

The value you enter sets the frequency with which the JMX metrics data are fetched from JMX servers. This means that every n seconds data is pushed to Elasticsearch.

It is a good idea to regularly purge the `psft_hc_metrics` index. The interval you select depends upon your usage. See the information on the Monitoring Server page in the Search Technology product documentation.

See *PeopleTools: Search Technology*, "Configuring the Monitoring Server."

Enter the polling frequency for JMX metrics:

23. Enter the number of threads that will be used to retrieve metrics and create events.

The number you select depends upon your usage. Increase or decrease it as needed.

Enter the number of threads:

24. Enter the information for the Elasticsearch host.

Enter the Elasticsearch host name:

Enter the Elasticsearch port:

Enter the Elasticsearch username:

Enter the Elasticsearch password:

Re-enter the Elasticsearch password:

25. Enter *y* (yes) if you want to create JSON files, or *n* (no) to continue.

If your environment is set up to use SSL, or if you have not completed the necessary prerequisite steps, enter *n* (no). See the section Preparing for the Logstash Installation for information on creating the JSON files manually.

Do you want to create JSON files? (y/n): *y*

26. Wait while Logstash is installed.

```
Extracting the Logstash Binary .....[OK]
Configuring Logstash.....[OK]
Verifying if Logstash config files are present in place....
Logstash config files are present.
```

Creating Json files. This may take some time...

Logstash installation completed

27. 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

28. 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):

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

Enter the server port for Kibana [5601] :

30. 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] :

31. 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.
```

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

See "Integrating Elasticsearch with the PeopleSoft Environment."

33. To start and use Kibana, see the Search Technology product documentation.

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

34. To start and use Logstash, see the Performance Monitor product documentation.

Note. The ELK DPK installation creates JSON files if the required conditions are met. You must start Logstash manually.

See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

A successful deployment includes the following:

- The Elasticsearch home directory is found in *BASE_DIR*\pt\elasticsearch7.0.0.
- 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\elasticsearch7.0.0.
The environment variable persists until you close the command prompt window.
- The Kibana home directory is found in *BASE_DIR*\pt\Kibana7.0.0
- The LOGSTASH_HOME environment variable is set to *BASE_DIR*\pt\Logstash7.0.0.
This environment variable is set at the system level.
- 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.

Task 2-5-2: Installing Elasticsearch, Logstash, and Kibana on Microsoft Windows in Silent Mode

This section discusses:

- Encrypting the Elasticsearch Passwords on Microsoft Windows

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

Encrypting the Elasticsearch Passwords on Microsoft Windows

You can use the setup script to install Elasticsearch, Logstash, 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. For the Elasticsearch installation, encrypt the esadmin and people passwords using the PSCipher utility. You must use the pscipher.jar and psvault files that are part of the ELK DPK. The psvault that is delivered with PeopleSoft PeopleTools will not work for Elasticsearch.

To encrypt the passwords for the Logstash installation, see the next section.

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

```
unzip ELK_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 *ELK_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 *ELK_INSTALL* and extract pt-jre-1.8.0_x.tgz into a folder with the same name.

It may be necessary to extract the file twice.

4. In *ELK_INSTALL\archives* and extract pt-elasticsearch-7.0.0.tgz into a folder with the same name.

It may be necessary to extract the file twice. The extraction creates the following folders:

- bin
 - config
 - jdk
 - lib
 - logs
 - modules
 - plugins
5. Copy psvault from the folder *ELK_INSTALL\pt-elasticsearch-7.0.0\plugins\orcl-security-plugin\config\properties* to *ELK_INSTALL\pt-elasticsearch-7.0.0\plugins\orcl-security-plugin*.
 6. In a command prompt, run the following command, specifying the passwords for esadmin and people.

```
<path_to_java>/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>/<outputfile>
```

For example:

```
C:\ELK_INSTALL\archives\pt-jre1.8.0_x\bin\java -Dpscipher.path=C:\ELK_
INSTALL\archives\pt-elasticsearch-7.0.0\plugins\orcl-security-plugin ->
cp C:\ELK_INSTALL\archives\pt-elasticsearch-7.0.0\plugins\orcl->
security-plugin\pscipher.jar com.peoplesoft.pt.elasticsearch.pscipher.P=>
SESEncrypt esadmin password1 people password2 C:\ELK_INSTALL\es_>
output.txt
```

7. Copy the encrypted text from `es_output.txt`, without adding line feeds or spaces. Paste the encrypted passwords into the `silentinstall.config` file for the `esadmin.password` and `people.password` parameters.

The output has this format:

```
esadmin:{V2.1}encrypted password==
people:{V2.1}encrypted password==
```

Encrypting the Logstash Passwords on Microsoft Windows

The passwords that you include in the silent mode text file must be encrypted. For the Integration Broker password for the Logstash installation, encrypt the password using the PSLSCipher script that is part of the ELK DPK.

1. Go to `ELK_INSTALL`.
2. Extract the entire contents of `ELK_FILENAME.zip` in the same directory, `ELK_INSTALL`.

```
unzip ELK_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 `ELK_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 `ELK_INSTALL\archives` and extract `pt-jre-1.8.0_x.tgz` into a folder with the same name. It may be necessary to extract the file twice.
 4. Go to `ELK_INSTALL\archives` and extract `pt-logstash-7.0.0.tgz` into a folder with the same name. It may be necessary to extract the file twice. The extraction creates several folders and text files.
 5. In a command prompt, set the following environment variables.

Add `JAVA_HOME` to the `PATH`. Use the location where you extracted `pt-jre-1.8.0_x.tgz`. Set `LOGSTASH_HOME` to the location where you extracted `pt-logstash-7.0.0.tgz`.

```
PATH=%JAVA_HOME%/bin;%PATH%
LOGSTASH_HOME=ELK_INSTALL\archives\pt-logstash-7.0.0
```

6. Change directory to `ELK_INSTALL\pt-logstash-7.0.0\pt\bin`.

7. Run the script, supplying the password:

```
PSLSCipher.bat password
```

8. Copy the encrypted password from the output on the screen, without adding line feeds or spaces, and paste it into the silentinstall.config file for the IB_PWD parameter.

The output has this format:

```
Encrypted text: {V2.1}encrypted password==
```

The ELK DPK installation includes a psvault key file, which is found in the directory `LOGSTASH_HOME\pt\properties`. The key in this psvault is not related to the keys available in Elasticsearch or the PeopleSoft Web server (PIA) psvault. It is recommended that appropriate access be used to protect the Logstash psvault, as with other similar files.

To generate and use a new version of psvault:

1. Change directory to `ELK_INSTALL\pt-logstash-7.0.0\pt\bin`.
2. Run the script as follows to build a new key:


```
PSLSCipher.bat -buildkey
```
3. Regenerate the Integration Broker user ID and password, and the Elasticsearch password with this command:


```
PSLSCipher.bat password or user ID
```
4. Edit the Logstash configuration files in the directory `LOGSTASH_HOME\pt\config` with the new encrypted text.
 - Update the values for the Integration Broker user ID and password in `JsonLogstash.properties`.
 - Update the values for the Elasticsearch password in `LogstashPipeLine.CONF`.

See [Reviewing the Logstash Configuration Files \(Optional\)](#).

See *PeopleTools: Security Administration*, "Securing the External Key File."

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

To install in silent mode:

1. Create a configuration file.

The `ELK_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 values. See the section [Installing Interactively](#) for guidance on specifying the values.

2. Enter `y` (yes) if you want to install Elasticsearch, or `n` (no) if you want to skip the Elasticsearch installation.

You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory. If you enter no, you can leave them blank. Enter encrypted text for the passwords as described in the previous section.

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

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

```

#Elasticsearch cluster name (mandatory)
cluster.name=

#Elasticsearch host name (mandatory)
network.host=

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

#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=

#Elasticsearch minimum master nodes (optional)
minimum_master_nodes=

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

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

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

```

3. Specify *y* if you want to install Logstash, or *n* if you do not.

You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory. If you enter no, you can leave them blank. Enter encrypted text for the passwords as described in the previous section.

```

#Install Logstash Y/N (Needed for PeopleSoft Health Centre) is
mandatory field
Install Logstash (for PHC)?[Y/N]=#IB REST service URL (mandatory)
IB_REST_URL=

#Enter the encrypted IB user [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
IB_USER=

#Enter the encrypted IB password [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
IB_PWD=

#location where JSON files need to be created(optional-default LS_HOME⇒
/pt/jmxmonitor)
JSON_LOC=

#The polling frequency for JMX metrics (mandatory)

```

```

polling_freq =

#The number of threads (mandatory)
no_of_threads =

#The Elasticsearch host name (mandatory)
ES_host =

#The Elasticsearch port (mandatory)
ES_port =

#The Elasticsearch username (mandatory)
ES_user =

#The encrypted Elasticsearch password [encrypted using PSLSCipher.bat⇒
/PSLSCipher.sh] (mandatory)
ES_pwd =

#Do you want to create JSON files(Y/N) (mandatory)
JSON_files?[Y/N] =

```

4. Specify *y* if you want to install Kibana, or *n* if you do not.

You must enter a value for this field. If you enter yes, you must specify values for the items labelled mandatory in this section. If you enter no, you can leave them blank.

```

#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=

```

5. If you want to install Kibana and connect to a previously-installed Elasticsearch server, specify *n* (no) for the following:

```

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

```

6. Specify the host name, port, and encrypted password for the Elasticsearch server.

```

#Provide the Elasticsearch instance details mapped to Kibana
ES.host[http(s)://hostname]=
ES.port= 9200
ES.password=

```

7. If the Elasticsearch is SSL enabled, enter the path to the certificate.

```

#Provide SSL certificate path if Elasticsearch is SSL enabled.
ES.cacert.path=

```

8. Enter *n* to indicate you do not want to upgrade.

You must enter a value for this field.

```

#####Silent Upgrade Configuration Values for upgrading⇒
Elasticsearch

```

```

from 6.1.2 to 7.0.0 #####

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

```

9. In a command prompt window, run the DPK setup script from *ELK_INSTALL\setup* as follows:

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

```

psft-dpk-setup.bat --install_silent --install_base_dir BASE_DIR ⇒
--config_file 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:/elk70 ⇒
--config_file C:/tmp_install/silentinstall.config

```

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

```

Elasticsearch Installation Completed.
Silent mode installation of Kibana

```

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

See "Integrating Elasticsearch with the PeopleSoft Environment."

12. To start and use Kibana, see the Search Technology product documentation.

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

13. To use Logstash, see the Performance Monitor product documentation.

Note. The ELK DPK installation creates JSON files and starts Logstash if the required conditions are met.

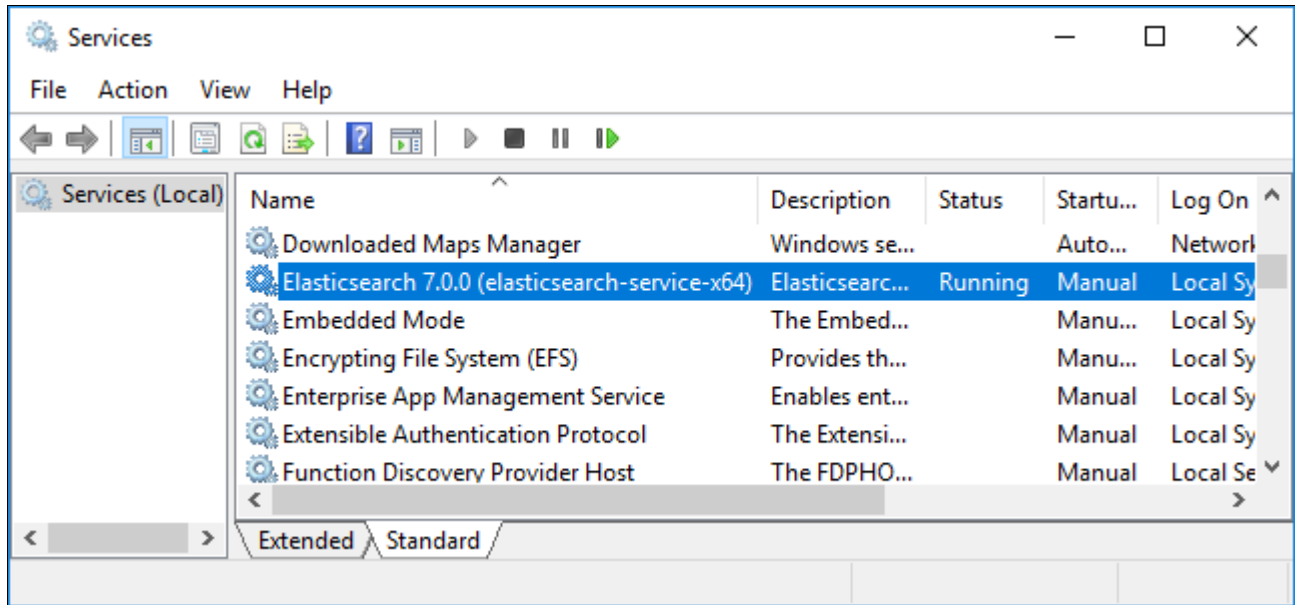
See *PeopleTools: Performance Monitor*, "Configuring PeopleSoft Health Center."

Task 2-5-3: 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 `ELK_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" : "7.0.0",
      "build_hash" : "b9e4a",
      "build_timestamp" : "2018-04-21T16:03:47Z",
      "build_snapshot" : false,
      "lucene_version" : "8.0.0"
      "minimum_wire_compatibility_version" : "6.7.0"
      "minimum_index_compatibility_version" : "6.0.0"
    },
    "tagline" : "You Know, for Search"
  }
}

```

Task 2-5-4: 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\elasticsearch7.0.0`, for `ELK_HOME`.

```

ELK_HOME\bin\elasticsearch-service.bat stop
ELK_HOME\bin\elasticsearch-service.bat remove

```

3. Remove the Elasticsearch installation directory.

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

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

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

1. Stop the Kibana script 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.
2. Remove the Kibana installation directory.

Task 2-5-6: Removing the Logstash Installation from Microsoft Windows

If you installed Logstash, use these steps to remove the Logstash installation from Microsoft Windows.

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

1. Stop the Logstash script if it is running.
If the command window that you used to start the Logstash script is open, either terminate the Logstash script by pressing `Ctrl+C`, or close the command window.
2. Remove the Logstash installation directory.

Task 2-5-7: 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

Upgrading Elasticsearch and Kibana

This chapter discusses:

- Upgrading Elasticsearch and Kibana to a New Revision Interactively
- Upgrading Elasticsearch from 6.1.2 to 7.0.0 Interactively
- Upgrading Elasticsearch from 6.1.2 to 7.0.0 in Silent Mode

Task 3-1: Upgrading Elasticsearch and Kibana to a New Revision Interactively

This section discusses:

- Upgrading to a New Revision on Microsoft Windows
- Upgrading to a New Revision on Linux

Task 3-1-1: Upgrading to a New Revision on Microsoft Windows

Use this process to upgrade an existing installation of Elasticsearch 7.0 and Kibana 7.0 to a later ELK DPK revision using the DPK setup script interactively. For example, if you installed using `ELK-DPK-WIN-7.0.0_01.zip`, you can upgrade to `ELK-DPK-WIN-7.0.0_02.zip` or later.

1. Download the current ELK DPK and save it in a newly created directory, referred to here as `ELK_INSTALL`.
2. Go to `ELK_INSTALL`.
3. Extract the entire contents of `ELK_FILENAME.zip` in the same directory, `ELK_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 `ELK_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
4. Open a command prompt.
 5. Change directory to `ELK_INSTALL/setup`.
 6. Run the DPK setup script with these options:

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

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

```
psft-dpk-setup.bat --upgrade --install_base_dir C:\elk70
```

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

7. Answer *y* (yes) to upgrade Elasticsearch, or *n* (no) to exit.

```
You've chosen to do an upgrade, it'll bring down the elasticsearch for⇒
a while, do you still want to continue with the upgrade? (y/n): y
```

8. Review the status messages as the script stops and upgrades the existing Elasticsearch.

```
Stopping Elasticsearch server.....[OK]
Extracting the new ES Binary.....[OK]
Upgrade in process.....[OK]
Cleaning up.....[OK]
Starting Elasticsearch server.....[OK]
```

```
Elasticsearch Upgrade Completed.
```

9. Answer *y* (yes) to upgrade Kibana, or *n* (no) to exit.

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

The script displays progress messages.

```
Elasticsearch service is running. Proceeding with the Kibana⇒
installation.
Extracting the new Kibana Binary .....[OK]
Upgrade in porcess..... [OK]
Kibana Upgrade Completed.
```

Task 3-1-2: Upgrading to a New Revision on Linux

Use this process to upgrade an existing installation of Elasticsearch 7.0 and Kibana 7.0 to a later ELK DPK revision using the DPK setup script interactively. For example, if you installed using ELK-DPK-LNX-7.0.0_01.zip, you can upgrade to ELK-DPK-LNX-7.0.0_02.zip or later.

1. Download the current ELK DPK and save it in a newly created directory, referred to here as *ELK_INSTALL*.
2. Open a terminal window.
3. Change directory to *ELK_INSTALL*.

```
cd ELK_INSTALL
```
4. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_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 *ELK_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
5. Open a command prompt.
 6. Change directory to *ELK_INSTALL/setup*.
 7. Run the DPK setup script with these options:


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

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


```
./psft-dpk-setup.sh --upgrade --install_base_dir /home/elk70
```
 - Use double-dashes when specifying the script options; for example, `--upgrade`.
 8. Answer *y* (yes) to upgrade Elasticsearch, or *n* (no) to exit.

```
You've chosen to do an upgrade, it'll bring down the elasticsearch for⇒
a while, do you still want to continue with the upgrade? (y/n): y
```

9. Review the status messages as the script stops and upgrades the existing Elasticsearch.

```
Stopping Elasticsearch server.....[OK]
Extracting the new ES Binary.....[OK]
Upgrade in process.....[OK]
Cleaning up.....[OK]
Starting Elasticsearch server.....[OK]
```

```
Elasticsearch Upgrade Completed.
```

10. Answer *y* (yes) to upgrade Kibana, or *n* (no) to exit.

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

The script displays progress messages.

```
Elasticsearch service is running. Proceeding with the Kibana⇒
installation.
Extracting the new Kibana Binary .....[OK]
Upgrade in porcess..... [OK]
Kibana Upgrade Completed.
```

Task 3-2: Upgrading Elasticsearch from 6.1.2 to 7.0.0 Interactively

This section discusses:

- Upgrading Interactively on Microsoft Windows
- Upgrading Interactively on Linux

Task 3-2-1: Upgrading Interactively on Microsoft Windows

Use this process to upgrade an existing installation of Elasticsearch 6.1.2 to Elasticsearch 7.0 using the ELK DPK setup script interactively.

Note. The upgrade does not include Kibana or Logstash. Instead, you must perform a new installation of Kibana and Logstash and start them manually.

1. Download the current ELK DPK and save it in a newly created directory, referred to here as *ELK_INSTALL*.
 2. Go to *ELK_INSTALL*.
 3. Extract the entire contents of *ELK_FILENAME.zip* in the same directory, *ELK_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 *ELK_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
4. Open a command prompt.
 5. Change directory to *ELK_INSTALL/setup*.
 6. Run the DPK setup script with these options:

```
psft-dpk-setup.bat --full_upgrade --install_base_dir ELK7_BASE_DIR
```

- For the `install_base_dir` option, specify the full path to where you want Elasticsearch installed. For example:
- ```
psft-dpk-setup.bat --full_upgrade --install_base_dir C:\elk70
```
- Use double-dashes when specifying the script options; for example, `--full_upgrade`.
7. Answer *y* (yes) to upgrade Elasticsearch, or *n* (no) to exit.

```
You've chosen to do an upgrade, it'll bring down the elasticsearch for⇒
a while, do you still want to continue with the upgrade? (y/n): y
```

8. Enter the path to the current *ES\_HOME*, such as `C:\esk612\pt\elasticsearch6.1.2`:

```
Enter the path to current ES_HOME[<base_dir>/pt/elasticsearch6.1.2]: C:⇒
\esk612\pt\elasticsearch6.1.2
```

9. Enter the current Elasticsearch user name, such as `esadmin`:

```
Enter the current Elasticsearch username: esadmin
```

10. Review the status messages as the script stops and upgrades the existing Elasticsearch.

```
Stopping Elasticsearch server.....[OK]
Extracting the new ES Binary.....[OK]
Extracting the new JRE.....[OK]
Upgrade in process.....[OK]
Configuring Elastcisearch.....[OK]
Starting Elasticsearch server.....[OK]

SUCCESS: Specified value was saved. [OK]
```

```

Checking if Elasticsearch service is running.....
Elasticsearch is running.....[OK]

Cleaning up security cache.....[OK]

Elasticsearch Upgrade Completed..

```

## Task 3-2-2: Upgrading Interactively on Linux

Use this process to upgrade an existing installation of Elasticsearch 6.1.2 to Elasticsearch 7.0 using the ELK DPK setup script interactively.

---

**Note.** The upgrade does not include Kibana. Instead, you must perform a new installation of Kibana 7.0.

---

1. Download the current ELK DPK and save it in a newly created directory, referred to here as *ELK\_INSTALL*.
2. Open a terminal window.
3. Change directory to *ELK\_INSTALL*.

```
cd ELK_INSTALL
```

4. Extract the entire contents of *ELK\_FILENAME.zip* in the same directory, *ELK\_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 *ELK\_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

5. Change directory to *ELK\_INSTALL/setup*.
6. Run the DPK setup script with these options:

```
./psft-dpk-setup.sh --full_upgrade --install_base_dir ELK7_BASE_DIR
```

- For the *install\_base\_dir* option, specify the full path where you want Elasticsearch installed; for example:

```
./psft-dpk-setup.sh --full_upgrade --install_base_dir /home/elk70
```

- Use double-dashes when specifying the script options; for example, *--full\_upgrade*.

7. Answer *y* (yes) to upgrade Elasticsearch, or *n* (no) to exit.

```
You've chosen to do an upgrade, it'll bring down the elasticsearch for⇒
a while, do you still want to continue with the upgrade? (y/n): y
```

8. Enter the path to the current *ES\_HOME*, such as */home/esk612/pt/elasticsearch6.1.2*:

```
Enter the path to current ES_HOME[<base_dir>/pt/elasticsearch6.1.2]: ⇒
/home/esk612/pt/elasticsearch6.1.2
```

9. Enter the current Elasticsearch user name, such as *esadmin*:

Enter the current Elasticsearch username: **esadmin**

10. Review the status messages as the script stops and upgrades the existing Elasticsearch.

```
Stopping Elasticsearch server.....[OK]
Extracting the new ES Binary.....[OK]
Extracting the new JRE.....[OK]
Upgrade in process.....[OK]
Configuring Elasticsearch.....[OK]
Starting Elasticsearch server.....[OK]

SUCCESS: Specified value was saved. [OK]

Checking if Elasticsearch service is running.....
Elasticsearch is running.....[OK]

Cleaning up security cache.....[OK]

Elasticsearch Upgrade Completed..
```

## Task 3-3: Upgrading Elasticsearch from 6.1.2 to 7.0.0 in Silent Mode

---

You can use the setup script to upgrade Elasticsearch in silent mode on either Microsoft Windows or Linux by preparing a text file that includes details about the current installation.

**Note.** The upgrade does not include Kibana or Logstash. Instead, you must perform a new installation of Kibana and Logstash and start them manually.

---

To upgrade in silent mode:

1. Create a configuration file.

The *ELK\_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.

2. To indicate that you do not want to install Elasticsearch, Kibana, or Logstash, enter *n* (no) to these questions:

---

**Note.** This sample shows only a few lines from the *silentinstall.config* file.

---

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

#####
#Install Logstash Y/N Needed for PeopleSoft Health Centre) is
mandatory field
Install Logstash (for PHC)?[Y/N]= n

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

3. Enter *y* (yes) to indicate you want to upgrade Elasticsearch.

You must enter a value for this field.



```
#####Silent Upgrade Configuration Values for upgrading
Elasticsearch from 6.1.2 to 7.0.0 #####
```

```
#Upgrade Elasticsearch Y/N is mandatory field
Upgrade elasticsearch?[Y/N]= y
```

4. Edit the configuration file with the location and host name for the current Elasticsearch 6.1.2 installation.

You must enter values for both of these fields.

```
#Current Elasticsearch home path in the format [<base_dir>/pt⇒
/elasticsearch6.1.2] (mandatory)
current.es.home=
```

```
#Current Elasticsearch host name (mandatory)
current.esuser=
```

5. On Microsoft Windows, in a command prompt window, run the DPK setup script from *ELK\_INSTALL/setup* as follows:

---

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

---

```
psft-dpk-setup.bat --full_upgrade_silent
--install_base_dir <ELK7_BASE_DIR>
--config_file <full_path_configuration_file>
```

- Use double-dashes when specifying the script options; for example, `--full_upgrade_silent`.
- For *ELK7\_BASE\_DIR*, supply the location for Elasticsearch 7.0.
- For the *config\_file* option, specify the full path to the prepared configuration file. For example:

---

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

---

```
psft-dpk-setup.bat --full_upgrade_silent --install_base_dir C:/elk70⇒
--config_file C:/tmp_install/silentinstall.config
```

6. On Linux, in a terminal window, run the DPK setup script from *ELK\_INSTALL/setup* as follows:

---

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

---

```
./psft-dpk-setup.sh --full_upgrade_silent
--install_base_dir <ELK7_BASE_DIR>
--config_file full_path_configuration_file
```

- Use double-dashes when specifying the script options; for example, `--full_upgrade_silent`.
- For *ELK7\_BASE\_DIR*, supply the location for Elasticsearch 7.0.
- For the *config\_file* option, specify the full path to the prepared configuration file. For example:

---

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

---

```
./psft-dpk-setup.sh --full_upgrade_silent /home/elk7⇒
```

```
--config_file /home/tmp_install/silentinstall.config
```

7. Wait until you see a message indicating the upgrade is complete.

```
Stopping Elasticsearch server[OK]
Extracting the new ES Binary ... [OK]
Extracting the new JRE [OK]
Upgrading ES [OK]
Configuring Elasticsearch [OK]
Starting Elasticsearch server [OK]

Checking if Elasticsearch service is running.
Elasticsearch is running.. ..[OK]

Cleaning up security cache..... [OK]

Elasticsearch Upgrade Completed.
```

## Chapter 4

# Integrating Elasticsearch with the PeopleSoft Environment

This chapter discusses:

- Applying PeopleSoft Application Enhancements for Kibana
- Setting Up the PeopleSoft Application for Elasticsearch
- Adding and Configuring an Elasticsearch Instance

## Task 4-1: Applying PeopleSoft Application Enhancements for Kibana

---

After installing the Elasticsearch, Logstash and Kibana DPK, you must apply updates and fixes provided by your PeopleSoft application to configure Kibana for your PeopleSoft environments. To find the updates and fixes required to implement Kibana, install the latest update image for your PeopleSoft application and use the following tracking group to pull all of the required bugs into a change package for easy deployment.

Tracking Group Name: PeopleSoft Features

Value: Kibana for PT 8.58 or above

To locate the tracking group and create the change package, see *PeopleTools: Change Assistant and Update Manager*, "Defining Change Packages," Selecting Other Criteria.

See PeopleTools 8.58 Online Help on the Oracle Help Center, <https://docs.oracle.com/en/applications/peoplesoft/peopletools/index.html>.

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

After you pull all of the required updates and fixes into a change package, and deploy the change package with PeopleSoft Update Manager, all specified components will be enabled for using Kibana to create dashboards to visualize search indexes.

This applies to the installation of Kibana from the ELK 7.0 DPK, which is required for PeopleTools 8.58. The updates and fixes for Kibana 7.0 are delivered beginning with the following update images:

- PeopleSoft Campus Solutions (CS) 9.2.018
- PeopleSoft Customer Relationship Management (CRM) 9.2.019
- PeopleSoft Enterprise Learning Management (ELM) 9.2.021
- PeopleSoft Financials and Supply Chain Management (FSCM) 9.2.036
- PeopleSoft Human Capital Management (HCM) 9.2.035
- PeopleSoft Interaction Hub (IH) 9.1.011

Note that upgrading to Elasticsearch 7.0 from version 6.1.2 requires no specific application fixes. To uptake new functionality on PeopleTools 8.58, follow the instructions in this section.

## **Task 4-2: 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 4-2-1: 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*.

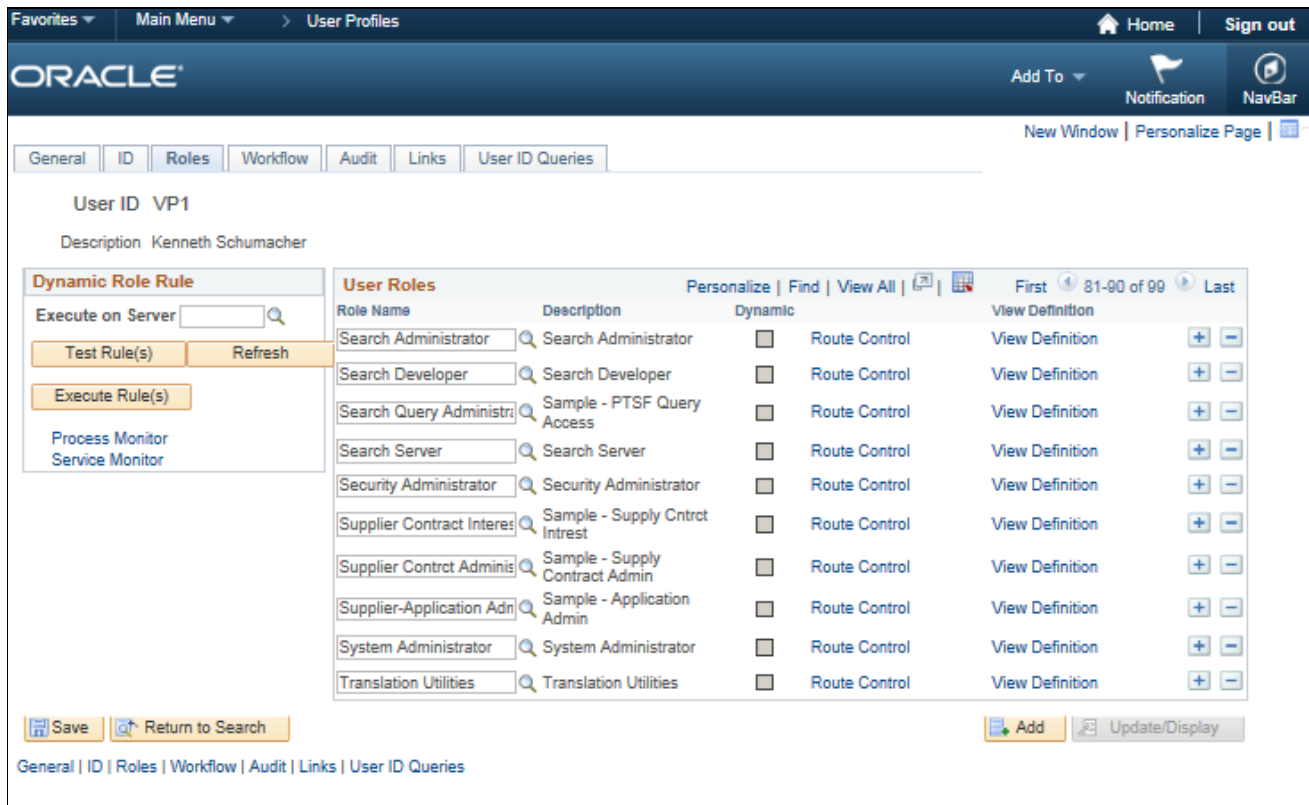
### Task 4-2-2: Verifying PeopleSoft Roles for All Installations

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 4-3: 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 4-3-1: Using the Automated Configuration Management SEARCH\_TEMPLATE**

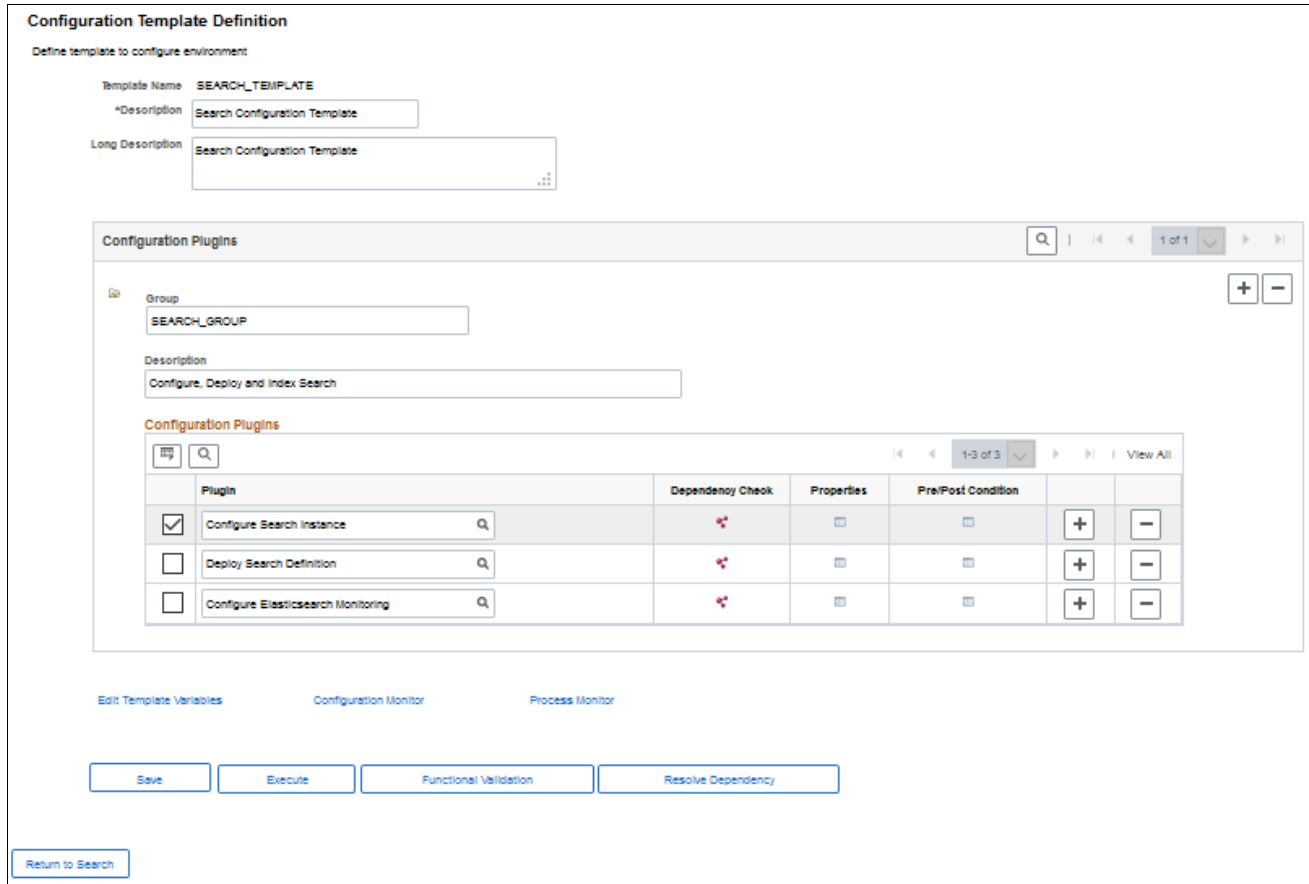
The ACM framework enables you to store environment configuration settings in a template stored in the database or an external template file, which you can reapply when needed. You can use Automated Configuration Management (ACM) to configure the Elasticsearch instance. Using ACM allows you to automate 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 or Configure Elasticsearch Monitoring 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.<br><br>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.<br><br>See <i>PeopleTools: Search Technology</i> for information on High Availability.<br><br>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         | <p>Flag to determine whether the configuration is secure or non-secure. Valid values are:</p> <ul style="list-style-type: none"> <li>False — indicates non-secure (HTTP) configuration</li> <li>True — indicates secure (HTTPS) configuration.</li> </ul> <p>Set this property to true only when the PeopleSoft environment is configured for SSL and that URL has been used for the Elasticsearch callback.</p> <p>See <i>PeopleTools: Search Technology</i>, "Configuring SSL Between PeopleSoft and Elasticsearch."</p> |
| env.node1_search_admin_user     | NA            | <p>The administrative user for Elasticsearch, esadmin.</p> <p>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.</p> <p>See <i>PeopleTools: Search Technology</i>, "Creating User and Assigning Roles in Elasticsearch."</p>                                                                                                                                                             |
| env.node1_search_admin_password | NA            | <p>The password supplied for the esadmin user when installing the Elasticsearch DPK.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| env.node1_search_read_user      | NA            | <p>The Elasticsearch proxy user, people.</p> <p>The proxy user cannot be changed during the installation.</p>                                                                                                                                                                                                                                                                                                                                                                                                              |
| env.node1_search_read_password  | NA            | <p>The password supplied for the proxy user when installing the Elasticsearch DPK.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| env.search_call_back_user       | @userid@      | <p>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.</p> <p><b>Note.</b> It should not be necessary to change this value.</p>                                                                                                                                                                                                                                                                  |
| env.search_call_back_password   | NA            | <p>Password for the call-back user ID.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

| Property                         | Default Value   | Description                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| env.ps_search_administrator_user | @userid@        | Search administrator user name<br>See Setting Up the PeopleSoft Application for Elasticsearch.                                                                                                                                                                                                                                                                |
| env.ps_search_developer_user     | @userid@        | Search developer user name<br>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: <ul style="list-style-type: none"> <li>False — indicates non-secure (HTTP) configuration</li> <li>True — indicates secure (HTTPS) configuration.</li> </ul> See <i>PeopleTools: Search Technology</i> , "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

Configuration Properties Find | View 15 First 1-19 of 19 Last

| Property                         | Value           | Prompt User?               |
|----------------------------------|-----------------|----------------------------|
| ★ env.ptsf_search_instance       | PTSF_SEARCH     | <input type="checkbox"/> ? |
| ★ env.search_nodes               | 1               | <input type="checkbox"/> ? |
| env.node1_search_host            | @searchhost@    | <input type="checkbox"/> ? |
| env.node1_search_port            | 9200            | <input type="checkbox"/> ? |
| env.node1_search_use_ssl         | false           | <input type="checkbox"/> ? |
| env.node1_search_admin_user      | esadmin         | <input type="checkbox"/> ? |
| env.node1_search_admin_password  | *****           | <input type="checkbox"/> ? |
| env.node1_search_read_user       | people          | <input type="checkbox"/> ? |
| env.node1_search_read_password   | *****           | <input type="checkbox"/> ? |
| ★ env.search_call_back_user      | @userid@        | <input type="checkbox"/> ? |
| ★ env.search_call_back_password  | *****           | <input type="checkbox"/> ? |
| env.ps_search_administrator_user | @userid@        | <input type="checkbox"/> ? |
| env.ps_search_developer_user     | @userid@        | <input type="checkbox"/> ? |
| ★ env.gateway_host               | @host@@@domain@ | <input type="checkbox"/> ? |
| ★ env.gateway_port               | 8000            | <input type="checkbox"/> ? |
| env.gateway_ssl_port             | 8001            | <input type="checkbox"/> ? |
| env.use_ssl_gateway              | false           | <input type="checkbox"/> ? |
| ★ env.default_local_node         | @nodename@      | <input type="checkbox"/> ? |
| env.enable_global_menu_search    | ALL             | <input type="checkbox"/> ? |

Configure Search Instance page

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

- 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**

**Search Instance**

|                 |               |                                       |
|-----------------|---------------|---------------------------------------|
| Search Instance | PTSF_DEFAULT  | <a href="#">Search Options Config</a> |
| Search Provider | Elasticsearch |                                       |

**Search Instance Properties** 1 of 1 | View All

\*Host Name  [Ping](#) + -

\*Port

\*SSL Option  [Test Login](#)

\*User Name

\*Password

\*Proxy Name  [Proxy Login](#)

\*Proxy Pwd

[Elasticsearch Interact](#)

**Kibana**

Host Name

Port

SSL Option

**Call Back Properties** ?

URL

User Name

\*Password

\*Confirm Password

[Validate](#)  
Update deployed definitions

[Set Namespace Aliases](#)

[Save](#) [Return to Search](#) [Notify](#) [Add](#)

Search Instance Properties Page

- Select PeopleTools, Automated Config Manager, ACM Templates, Template Definition.
- Search for and open SEARCH\_TEMPLATE.
- 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: <ul style="list-style-type: none"> <li>• ALL<br/>Deploys all search definitions and categories excluding the ones listed in env_ptsf_exclude_definitions.</li> <li>• GLOBAL<br/>Deploys search definitions and categories used for Global search.</li> <li>• LIST<br/>Deploys the search definitions and their categories mentioned in env_ptsf_include_definitions.</li> </ul> |
| 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.<br><br>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.                                                                                       |
| env_ptsf_admin_operations    | DEPLOY,INDEX  | Valid values: <ul style="list-style-type: none"> <li>• DEPLOY — the search definitions will be deployed</li> <li>• INDEX — the search index will be scheduled.</li> </ul>                                                                                                                                                                                                                     |

| Property                                 | Default Value | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                          |               | <p>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 &lt;SearchDefinition&gt;_FULL and incremental indexing will have the naming convention &lt;SearchDefinition&gt;_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.</p> <p><b>Note.</b> 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.</p> <ul style="list-style-type: none"> <li>• UNDEPLOY — the search definition will be removed.</li> <li>• DEPLOY,INDEX</li> </ul> |
| env.ptsf_index_all_languages             | False         | If true the schedules are created to index all languages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| env.report_schedule_status_after_minutes | NA            | <p>Maximum minutes to wait before reporting scheduling status.</p> <ul style="list-style-type: none"> <li>• Enter 0 to wait till finish.</li> <li>• Leave blank to skip report.</li> <li>• Enter the time to wait to show status. For example, enter 5 to show the status once after 5 minutes.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 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.

| Property                                 | Value        | Prompt User?               |
|------------------------------------------|--------------|----------------------------|
| env.ptsf_search_instance                 | PTSF_SEARCH  | <input type="checkbox"/> ? |
| ★ env.ptsf_selection_type                | ALL          | <input type="checkbox"/> ? |
| env.ptsf_include_definitions             |              | <input type="checkbox"/> ? |
| env.ptsf_exclude_definitions             |              | <input type="checkbox"/> ? |
| ★ env.ptsf_check_audit_errors            | true         | <input type="checkbox"/> ? |
| ★ env.ptsf_admin_operations              | DEPLOY,INDEX | <input type="checkbox"/> ? |
| ★ env.ptsf_index_all_languages           | false        | <input type="checkbox"/> ? |
| env.report_schedule_status_after_minutes |              | <input type="checkbox"/> ? |
| env.ptsf_schedule_on_server              |              | <input type="checkbox"/> ? |

Buttons: OK, Cancel, Current Values in DB

Deploy Search Definition page

12. Save the template and click Execute to begin the deployment.
13. Use the plug-in Configure Elasticsearch Monitoring to configure the Elasticsearch monitoring server, deploy the PeopleSoft Health Center dashboards to Kibana, and deploy the system monitoring and Elasticsearch index metrics dashboards to Kibana.

See the information on the delivered plug-in PTSFMonitorConfiguration in the Automated Configuration Management product documentation.

See *PeopleTools: Automated Configuration Management*, "Delivered Configuration Plug-ins."

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

## Task 4-3-2: 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 5

# 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
- Adding Kibana as a Service in Linux
- Reviewing the Logstash Configuration Files (Optional)

## Task 5-1: 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 5-2: 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 5-3: 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 5-4: 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 5-5: 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 5-6: Adding Elasticsearch as a Service in Linux

---

This section discusses:

- Prerequisites
- Adding an Elasticsearch Service
- Verifying that the Elasticsearch Service Starts Automatically
- 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 stop 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 directories are located under the DPK base directory; that is:
  - The DPK base directory is referred to in this documentation as *BASE\_DIR*, such as `/home/elk70`.
  - The Elasticsearch installation directory, *ELK\_HOME*, is found in `BASE_DIR/pt/elasticsearch7.0.0`.
  - 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 *ELK\_HOME*.  
This is due to the fact that the script does not input the user name.
- You have downloaded and extracted the required ELK DPK for Linux, in a directory referred to as *ELK\_INSTALL*.

### Task 5-6-1: 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` *BASE\_DIR*  
Enter the base directory (*BASE\_DIR*) where you installed the ELK DPK.

To add Elasticsearch as a service:

1. Open a terminal window, running as root.
2. Change directory to `ELK_INSTALL/setup`:

```
cd ELK_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
```

- Beginning with ELK DPK version 04, use this `systemctl` command:

```
systemctl status elasticsearch
```

This should give an output with the status "active (running)," as shown in bold font in this sample:

```
elasticsearch.service - Elasticsearch
 Loaded: loaded (/etc/systemd/system/elasticsearch.service;⇒
 enabled; vendor preset: disabled)
 Active: active (running) since Tue 2020-09-01 22:24:38 PDT; 12s ago
 Docs: http://www.elastic.co
 Main PID: 59416 (java)
 Tasks: 45
 Memory: 2.2G
 CGroup: /system.slice/elasticsearch.service
```

- For the ELK DPK version 03 and earlier versions, 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
```

- For the ELK DPK version 03 and earlier versions, use this `service` command:

```
service elasticsearch status
```

This should give an output such as:

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

## Task 5-6-2: Verifying that the Elasticsearch Service Starts Automatically

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. Verify that the service has come up automatically.
  - Beginning with ELK DPK version 04, use this command to verify that the service has come up automatically:

```
systemctl status elasticsearch
```

This should give an output with the status "active (running)," as shown in bold font in this sample:

```
elasticsearch.service - Elasticsearch
 Loaded: loaded (/etc/systemd/system/elasticsearch.service;⇒
 enabled; vendor preset: disabled)
 Active: active (running) since Tue 2020-09-01 22:24:38 PDT; 12s ago
 Docs: http://www.elastic.co
 Main PID: 59416 (java)
 Tasks: 45
 Memory: 2.2G
 CGroup: /system.slice/elasticsearch.service
```

- For the ELK DPK version 03 and earlier versions, 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 5-6-3: Removing the Elasticsearch Service

To remove the Elasticsearch service:

1. Open a terminal window, running as root.
2. Change directory to *ELK\_INSTALL/setup*:

```
cd ELK_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
```

- Beginning with ELK DPK version 04, use this command:

```
systemctl status elasticsearch
```

- For the ELK DPK version 03 and earlier versions, use this `chkconfig` command. The command should give an empty output:

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

- For the ELK DPK version 03 and earlier versions, use this `service` command:

```
service elasticsearch status
```

This should give an output such as:

```
elasticsearch: unrecognized service
```

## Task 5-7: Adding Kibana as a Service in Linux

---

This section discusses:

- Prerequisites
- Adding a Kibana Service
- Verifying that the Kibana Service Starts Automatically
- Removing the Kibana Service

### Prerequisites

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

- The Kibana process should not be running.  
If Kibana is running, ensure that it is not being used, and then stop the process.  
See Starting and Stopping an Elasticsearch Service.
- The script to install the service must be run by the root user.
- Kibana is installed on the Linux server where you run this procedure.
- The Kibana and JRE installation folders are located under the DPK base directory; that is:
  - The DPK base directory is referred to in this documentation as *BASE\_DIR*, such as */home/elk70*.
  - The Kibana installation directory, *KIBANA\_HOME*, is found in *BASE\_DIR/pt/Kibana7.0.0*.
  - The Java installation directory, *JAVA\_HOME*, is found in *BASE\_DIR/pt/es\_jre1.8.0\_yy*, where *yy* is the JRE version.
- Kibana must be run by the user who owns *KIBANA\_HOME*.  
This is due to the fact that the script does not input the user name.
- You have downloaded and extracted the required ELK DPK for Linux, in a directory referred to as *ELK\_INSTALL*.

### Task 5-7-1: Adding a Kibana Service

You can add a Kibana service beginning with ELK DPK version 04.

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

- `-h` or `--help`  
Show the help message and exit
- `--add`  
A Kibana service will be added.
- `--delete`  
The Kibana service will be removed.
- `--install_base_dir` *BASE\_DIR*  
Enter the base directory (*BASE\_DIR*) where you installed the ELK DPK.



To add Kibana as a service:

1. Open a terminal window, running as root.
2. Change directory to `ELK_INSTALL/setup`:

```
cd ELK_INSTALL/setup
```

3. Run this command to add the service:

```
./psft-kibana-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 Kibana process:

```
ps -ef | grep kibana
```

- Use this `systemctl` command:

```
systemctl status kibana
```

This should give an output with the status "active (running)," as shown in bold font in this sample:

```
kibana.service - Kibana
 Loaded: loaded (/etc/systemd/system/kibana.service; enabled;⇒
 vendor preset: disabled)
 Active: active (running) since Tue 2020-09-01 23:05:38 PDT; 9s ago
 Main PID: 67149 (node)
 Tasks: 11
 Memory: 111.7M
 CGroup: /system.slice/kibana.service
```

## Task 5-7-2: Verifying that the Kibana Service Starts Automatically

After you install the Kibana 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. Verify that the service has come up automatically.

```
systemctl status kibana
```

This should give an output with the status "active (running)," as shown in bold font in this sample:

```
kibana.service - Kibana
 Loaded: loaded (/etc/systemd/system/kibana.service; enabled; vendor⇒
 preset: disabled)
 Active: active (running) since Tue 2020-09-01 23:05:38 PDT; 9s ago
 Main PID: 67149 (node)
 Tasks: 11
 Memory: 111.7M
 CGroup: /system.slice/kibana.service
```

## Task 5-7-3: Removing the Kibana Service

To remove the Kibana service:

1. Open a terminal window, running as root.

2. Change directory to *ELK\_INSTALL/setup*:

```
cd ELK_INSTALL/setup
```

3. Run this command to remove the service:

```
./psft-kibana-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 Kibana process:

```
ps -ef | grep kibana
```

- Use this `systemctl` command:

```
systemctl status kibana
```

## **Task 5-8: Reviewing the Logstash Configuration Files (Optional)**

When you install Logstash, the DPK installation creates configuration files in *LOGSTASH\_HOME/pt/config*. The values in the files are supplied by the installation process, so you should not need to make any changes.

- `JsonLogstash.properties` includes the following parameters:
  - `IBServiceURL` — the Integration Broker REST service URL
  - `JSONLocation` — the location that you specified for the JSON files
  - `IBusername` — encrypted name for the Integration Broker user
  - `IBpassword` — encrypted password for the Integration Broker user
- `LogstashPipeline.CONF` includes input and output values for the Logstash event processing pipeline. See the Logstash information on the Elastic web site for a description of the Logstash pipeline process.

The file has the following format:

```
input {
 jmx {
 path => "C:\elk70\pt\Logstash7.0.0\pt\jmxmonitor"
 polling_frequency => 5
 type => "jmx"
 nb_thread => 15
 }
}
output {
 elasticsearch {
 hosts => "server.example.com:5601"
 index => "psft_hc_metrics"
 user => "esadmin"
 password => "encrypted_password"
 }
}
```

The password in the output section is the `esadmin` password, which is encrypted using the key available in the `psvault`.

Use these guidelines if you must change the configuration files:

- Enter the same value for the path in the input section of `LogstashPipeline.CONF` and the `JSONLocation` in the `JsonLogstash.properties` file.
- The `polling_frequency` parameter in the input section of `LogstashPipeline.CONF` is the frequency with which the JMX data is pulled from JMX agents and pushed to Elasticsearch.  
This value is mandatory. There is no default or recommended value. Set the value based on the requirements of your environment.
- The `nb_thread` parameter in the input section of `LogstashPipeline.CONF`, which refers to the number of threads used for retrieving data, can be increased or decreased based on your environment if necessary.

